

Geodesic-Preserving Polygon Simplification^{*}

Oswin Aichholzer¹, Thomas Hackl¹, Matias Korman², Alexander Pilz¹, and
Birgit Vogtenhuber¹

¹ Institute for Software Technology, Graz University of Technology, Austria,
[oaich|thackl|apilz|bvogt]@ist.tugraz.at.

² Dept. Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Spain,
matias.korman@upc.edu.

Abstract. Polygons are a paramount data structure in computational geometry. While the complexity of many algorithms on simple polygons or polygons with holes depends on the size of the input polygon, the intrinsic complexity of the problems these algorithms solve is often related to the reflex vertices of the polygon. In this paper, we give an easy-to-describe linear-time method to replace an input polygon \mathcal{P} by a polygon \mathcal{P}' such that (1) \mathcal{P}' contains \mathcal{P} , (2) \mathcal{P}' has its reflex vertices at the same positions as \mathcal{P} , and (3) the number of vertices of \mathcal{P}' is linear in the number of reflex vertices. Since the solutions of numerous problems on polygons (including shortest paths, geodesic hulls, separating point sets, and Voronoi diagrams) are equivalent for both \mathcal{P} and \mathcal{P}' , our algorithm can be used as a preprocessing step for several algorithms and makes their running time dependent on the number of reflex vertices rather than on the size of \mathcal{P} .

1 Introduction

A *simple polygon* is a closed connected domain in the plane that is bounded by a sequence of straight line segments (edges) such that any two edges may intersect only in their endpoints (vertices), and such that in every vertex exactly two edges intersect. Let \mathcal{P} be a simple polygon and let $\mathcal{H}_1, \dots, \mathcal{H}_k$ be a set of pairwise-disjoint simple polygons such that \mathcal{H}_i is contained in the interior of \mathcal{P} for $1 \leq i \leq k$. Then the closure \mathcal{Q} of $\mathcal{P} \setminus \bigcup_{1 \leq i \leq k} \mathcal{H}_i$ is called a *polygon with holes*. The polygons \mathcal{H}_i are called the *holes* of \mathcal{Q} , and the vertices and edges of \mathcal{Q} are the vertices and edges of \mathcal{P} and all \mathcal{H}_i , respectively. We regard \mathcal{P} and \mathcal{Q} as closed sets, i.e., they include their boundary. All polygons we consider are

^{*} Research supported by the ESF EUROCORES programme EuroGIGA - ComPoSe, Austrian Science Fund (FWF): I 648-N18 and grant EUI-EURC-2011-4306. T.H. supported by the Austrian Science Fund (FWF): P23629-N18 ‘Combinatorial Problems on Geometric Graphs’. M.K. received support of the Secretary for Universities and Research of the Ministry of Economy and Knowledge of the Government of Catalonia and the European Union. A.P. is recipient of a DOC-fellowship of the Austrian Academy of Sciences at the Institute for Software Technology, Graz University of Technology, Austria.

either simple polygons or polygons with holes. A vertex of a polygon is reflex if its inner angle is larger than 180° . Given a polygon \mathcal{P} with n vertices, the *geodesic path* $\pi_{\mathcal{P}}(p, q)$ between two points p and q of \mathcal{P} is defined as the shortest path that connects p and q among all the paths that stay within \mathcal{P} . The length of that path is called the *geodesic distance*. For any pair of points in \mathcal{P} , such a path always exists and, when \mathcal{P} is simple, is unique. Moreover, such a path is a polygonal chain whose vertices (other than p and q) are reflex vertices of \mathcal{P} . When the path $\pi_{\mathcal{P}}(p, q)$ is a straight line segment, we say that p *sees* q (and vice versa).

We say that two polygons \mathcal{P} and \mathcal{P}' have the same reflex vertices if for any reflex vertex $v \in \mathcal{P}$ there is a reflex vertex $v' \in \mathcal{P}'$ at the same point in the plane and vice versa. We say that \mathcal{P}' *subsumes* \mathcal{P} if \mathcal{P} and \mathcal{P}' have the same reflex vertices and $\mathcal{P} \subseteq \mathcal{P}'$. See Fig. 1 for examples.




Fig. 1. Simple polygons (dashed) drawn on top of subsuming polygons (solid).

Observation 1. Let $\mathcal{P}, \mathcal{P}'$ be two simple polygons such that \mathcal{P}' subsumes \mathcal{P} . Then, for any $p, q \in \mathcal{P}$ we have $\pi_{\mathcal{P}}(p, q) = \pi_{\mathcal{P}'}(p, q)$.

For algorithms that solely rely on geodesic paths inside \mathcal{P} , the output will remain equivalent if we replace \mathcal{P} by \mathcal{P}' in the input. It is therefore desirable to construct a subsuming polygon \mathcal{P}' such that (1) \mathcal{P}' has few vertices and (2) the construction can be done in time linear in the size of \mathcal{P} . For the creation of a subsuming polygon of minimal size, one has to connect the reflex vertices by pairwise non-intersecting paths in the union of the exterior and the boundary of the original polygon. Guibas, Hershberger, Mitchell, and Snoeyink [14] address various aspects of the problem of approximating polygonal paths and polygons by simpler ones. They show that the problem of finding a minimum link simple polygon (of a given homotopy class) having its boundary inside a given region R is NP-hard. In this paper, we show the following result.

Theorem 1. For any polygon \mathcal{P} (possibly with holes) of n vertices out of which $r > 0$ are reflex, there exists a polygon \mathcal{P}' with $O(r)$ vertices that subsumes \mathcal{P} . Moreover, \mathcal{P}' can be computed in $O(n)$ time and will have the same number of holes as \mathcal{P} .

Before giving the proof of Theorem 1 in Section 3, we review some basic properties of pointed pseudo-triangulations (also known as *geodesic triangulations*)

in Section 2. Section 4 gives an overview of the various applications of our result. In Section 5 we summarize our approach and give a short account on the open problem of computing an optimal subsuming polygon. Due to space constraints, the proofs of some claims are deferred to the full version of the paper [1].

Related Work

We follow the common aim of relating the time and space complexity of algorithms on polygons not only to the input size, but also to the number of reflex vertices. This is often a more significant parameter for the actual difficulty of the problem instance. Hertel and Mehlhorn [18] give an algorithm for triangulating a simple polygon of n vertices, r of which are reflex, in $O(n \log r)$ time. Bose et al. [8] give a $\Theta(m + n \log r)$ algorithm for computing a geodesic ham-sandwich cut of two given sets of m points in a simple polygon with r reflex vertices.

A related way of giving a more fine-grained analysis of algorithms on polygons is by expressing its complexity in the number of edges in the visibility graph (i.e., the number of point pairs that see each other), see for example [11, p. 68]. Note that, e.g., for computing the minimum weight triangulation of a simple polygon, the currently known worst case appears when there are no reflex vertices [7].

The term “polygon simplification” is also used in connection with operations that are used to compress polygons and to reduce noise in the representation. A well-known algorithm to smoothen polygonal chains is the Douglas-Peucker algorithm [10]. Guibas et al. [14] address several variations of the approach to fatten existing polygonal chains and approximate the chain inside the fattened region. There also exists work on constructing simple polygons that contain given ones and fulfill certain properties, as a generalization of the convex hull of simple polygons [4]; the main objective there is to approximate the shape.

2 Pseudo-Triangulations

In this section, we recall several properties of pseudo-triangles and pointed pseudo-triangulations of simple polygons. For details on pseudo-triangulations in various contexts see the survey by Rote, Santos, and Streinu [24].

A *pseudo-triangle* is a simple polygon with exactly three convex vertices. The three convex vertices are the *corners* of the pseudo-triangle, and the three polygonal chains between the corners are called the *side chains*. Note that a side chain might consist only of one edge.

A *pseudo-triangulation* of a simple polygon P is a partition of P into pseudo-triangles, such that the union of the vertices of the pseudo-triangles is exactly the vertex set of P . A vertex v of a pseudo-triangulation is called *pointed* if it is incident to a face in which the angle at v is larger than 180° . In a *pointed pseudo-triangulation* every vertex is pointed. Throughout this work, we are only concerned with pointed pseudo-triangulations. It can be shown that a pointed pseudo-triangulation of a simple polygon with c convex vertices has $c - 2$ pseudo-triangles and adds $c - 3$ diagonals to the polygon (see, e.g., [24]). Our main result

heavily relies on that fact. Guibas et al. [13] showed that, given a vertex v of a triangulated simple polygon P , the set of all shortest paths between v and the vertices of P (i.e., the shortest path tree of v) can be constructed in linear time (this algorithm was later simplified by Hershberger and Snoeyink [17]). The union of all these shortest paths gives a pointed pseudo-triangulation of P . Hence, given a triangulation of a simple polygon P , a pointed pseudo-triangulation of P can be constructed in linear time.³

Let ∇ be a pseudo-triangle. The line ℓ bisecting the angle at any corner of ∇ separates the two adjacent side chains C_1 and C_2 , and will leave ∇ through the third chain C_3 . The bisecting line ℓ' of the angle of a second corner, say, the one joining C_1 and C_3 , separates C_1 and C_3 . Therefore, ℓ and ℓ' intersect inside ∇ and hence separate C_1 from C_2 and C_3 . This allows us to make the following basic observation.

Observation 2. *For any side chain C of a pseudo-triangle ∇ , the bisectors of the angles at the corners of C define a wedge that separates C from the remaining boundary of ∇ .*

We call this wedge the *separating wedge* of the side chain. Observation 2 is the crucial property of pseudo-triangles that we will use in the next section.

3 Proof of the Main Theorem

We call a simple polygonal chain C *hull-honest* if it is completely contained in the boundary of its convex hull $\mathcal{CH}(C)$. We call a simple chain $C = \langle v_1, v_2, \dots, v_k \rangle$ *simplifiable*, if it is hull-honest and if the ray $r(v_1, v_2)$ (i.e., the ray from v_1 through v_2) and the ray $r(v_k, v_{k-1})$ intersect. See Fig. 2. Note that if a chain is simplifiable, then any of its subchains is simplifiable as well. For any simplifiable chain C of vertices v_1, v_2, \dots, v_k we introduce an operation called the *simplification* of C as follows: If $k \leq 3$ then C remains unchanged. Otherwise, consider the rays $r(v_1, v_2)$ and $r(v_k, v_{k-1})$. As $k > 3$, these two rays intersect at a point m not on the chain. In that case, we replace C by the chain $C' = \langle v_1, m, v_k \rangle$.

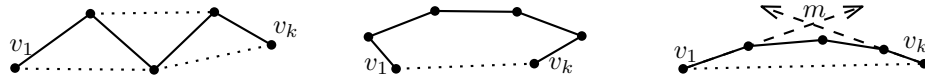


Fig. 2. A chain that is not hull-honest (left), a hull-honest chain (center), and a simplifiable chain (right).

The basic idea of our construction is to simplify long convex chains along the boundary of \mathcal{P} . The main challenge is to avoid that the edges introduced by the simplification intersect with the remaining boundary of the resulting polygon.

³ The connection between pointed pseudo-triangulations and the shortest path tree was mentioned by Speckmann and Tóth [25]; the concept of pseudo-triangulations has been developed after the writing of [13].

Lemma 1. *All the side chains of any pseudo-triangle are simplifiable. When simplifying an arbitrary number of pairwise interior-disjoint subchains, the resulting polygon is again a pseudo-triangle, and, in particular, its boundary is not self-intersecting.*

Using this result we can now prove our main theorem.

Proof (Theorem 1). We first consider the case where \mathcal{P} is a simple polygon with at least one reflex vertex. Let $\mathcal{CH}(\mathcal{P})$ be its convex hull (which can be computed in linear time using, e.g., Melkman's algorithm [19]). The set $\mathcal{CH}(\mathcal{P}) \setminus \mathcal{P}$ is the union of (the interiors of) simple polygons whose interiors are pairwise disjoint. We call these polygons the *pockets*. Each pocket P_i is defined by exactly one convex hull edge, which is not part of \mathcal{P} and is called a *lid edge*, and a subchain C of the boundary of \mathcal{P} . The c_i convex vertices of a pocket therefore consist of the reflex vertices of \mathcal{P} along C and the two vertices of the lid edge. Thus, a pointed pseudo-triangulation of P_i has $c_i - 3$ diagonals. We call the lid edges and the pseudo-triangulation diagonals the *support edges*. For p pockets, the number of support edges is $p + \sum_{i=1}^p (c_i - 3) = p + r + 2p - 3p = r$. Since the only vertices possibly shared by two pockets are the convex hull vertices, we can construct a pointed pseudo-triangulation of each pocket in accumulated $O(n)$ time for all pockets. See Fig. 3 for an example of a pseudo-triangulated pocket.

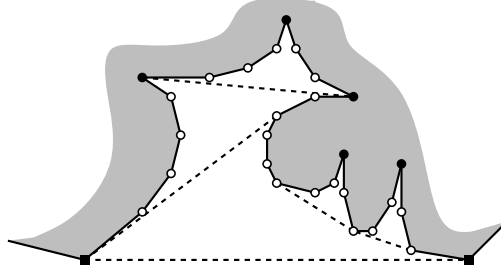


Fig. 3. A pocket and its pointed pseudo-triangulation. Support edges are drawn dashed.

Consider a pocket P_i and a pointed pseudo-triangulation $T(P_i)$ of P_i . A side chain of a pseudo-triangle of $T(P_i)$ consists of convex chains of \mathcal{P} , possibly separated by support edges. Note that any vertex of \mathcal{P} that is not on $\mathcal{CH}(\mathcal{P})$ is clearly part of a side chain of at least one pseudo-triangle of $T(P_i)$ in some pocket P_i of \mathcal{P} . For each pseudo-triangle in every pocket of \mathcal{P} we simplify all maximal subchains of its side chains that do not contain support edges. Due to Lemma 1, the resulting polygon is again simple and subsumes \mathcal{P} .

Observe that $\mathcal{CH}(\mathcal{P})$ consists of convex chains of \mathcal{P} , possibly separated by lid edges. Unlike the maximal convex chains inside pseudo-triangles, there might exist a maximal convex chain (not containing lid edges) $C = \langle v_i, \dots, v_j \rangle$ on $\mathcal{CH}(\mathcal{P})$ that is not simplifiable. The reason for this being, that the turn of the

chain is at least 180° , i.e., the rays $r(v_i, v_{i+1})$ and $r(v_j, v_{j-1})$ do not intersect. However, such a chain C appears at most once on \mathcal{P} and can be split into two simplifiable parts with a common vertex v^* .

It remains to count the number of convex vertices of the resulting polygon \mathcal{P}' . To this end, we charge the convex vertices of \mathcal{P}' either to reflex vertices or to one of the r support edges. Let $C = \langle v_i, \dots, v_j \rangle$ be a maximal convex chain not containing support edges of \mathcal{P}' . Each of the two end points of C is either a reflex vertex of \mathcal{P} (or v^*), or end point of a support edge, or both. If v_i is the end point of a support edge e , then we charge C to e . Otherwise, we charge C to the reflex point v_i (or to the special point v^*). Observe that each end point of a support edge can be at most once a starting point (v_i) of such a maximal convex chain C of \mathcal{P}' . The same is true for each reflex vertex of \mathcal{P} (and v^*). Thus, \mathcal{P}' consists of at most $2r + r + 1$ maximal convex chains.

Each such maximal convex chain of \mathcal{P}' consists of at most three vertices that might all be convex in \mathcal{P}' . Observe though, that the last vertex v_j of a chain C is either a reflex vertex or also the first vertex of another maximal convex chain $C' = \langle v_j, \dots, v_k \rangle$ of \mathcal{P}' . Hence, we need to count only at most two convex vertices per chain. Therefore, \mathcal{P}' has at most $6r + 2$ convex vertices.

Finally, observe that we can apply the same strategy to a polygon \mathcal{P} with holes \mathcal{H}_i by considering each hole as a simple polygon itself. Recall that our approach simplifies a polygon, preserving geodesics inside the polygon. We independently pseudo-triangulate the interior of each hole \mathcal{H}_i (in time linear in the number of vertices of \mathcal{H}_i) and apply the simplification strategy. For each \mathcal{H}_i we obtain a polygon whose complexity is proportional to the number of convex vertices of (the polygon) \mathcal{H}_i and preserves geodesics outside \mathcal{H}_i (i.e., paths that consider \mathcal{H}_i as an obstacle). Thus, the resulting simplification preserves geodesics inside \mathcal{P} . As each convex vertex of \mathcal{H}_i is a reflex vertex of \mathcal{P} , the complexity of this simplification will be proportional to the number of vertices that are reflex in \mathcal{P} . \square

Remark. Note that with this process we can explicitly give the coordinates of \mathcal{P}' . If the coordinates of the vertices of \mathcal{P} are rational, then the coordinates used for \mathcal{P}' are rational as well. Moreover, explicitly representing these vertices needs at most a constant times the number of bits used by input vertices. Alternatively, we can store the simplification in an implicit form: we store the first and last vertices of each simplified chain, allowing an easy identification between each connected component of $\mathcal{P}' \setminus \mathcal{P}$ and the simplified chains. In either case, at most $O(r)$ space is needed (for constant size vertex representation).

While our result is asymptotically optimal, we do not construct a subsuming polygon with the minimal number of vertices. Given a polygon \mathcal{Q} with h holes, Guibas et al. [14, Theorem 5] show how to construct a simple polygon \mathcal{P} with its boundary inside \mathcal{Q} such that \mathcal{P} contains all holes and has only $O(h)$ more vertices than the (probably non-simple) optimum. We note that, after knowing that the number of vertices of a minimal subsuming polygon is in $O(r)$, also the method by Guibas et al., after careful modification, could be used for constructing a subsuming polygon within the given bounds when applying it to each pocket.

However, our constructive proof of the bound gives a straight-forward algorithm that can be implemented using standard tools used for visibility problems in simple polygons.

4 Applications

Let \mathcal{A} be an algorithm that receives a polygon \mathcal{P} (and potentially other input I). We say that \mathcal{A} is *subsuming* if the result of executing \mathcal{A} with input \mathcal{P} and I is equal to the result obtained when the input is \mathcal{P}' and I , where \mathcal{P}' is a polygon that subsumes \mathcal{P} . Note that, if \mathcal{A} is subsuming, so will any other algorithm that solves the same problem (thus we say that the problem is *simplifiable*).

Theorem 2. *Let \mathcal{P} be a polygon of size n with r reflex vertices and let I be additional input of size m . Let $\mathcal{A}(\mathcal{P}, I)$ be an algorithm that solves a simplifiable problem and runs in $T(n, m)$ time using $S(n, m)$ space. Then, $\mathcal{A}(\mathcal{P}, I)$ can be modified to run in $O(n + T(r, m))$ time and $O(n + S(r, m))$ space.*

In the following we present several applications for Theorem 2 which show the versatility of our approach. Most of these problems can be stated in terms of shortest paths, which immediately implies that the problems are simplifiable. For each task, we briefly state the problem, mention the (to the best of our knowledge) fastest existing algorithm, and explain how our approach helps to reduce the running time.

4.1 Shortest Paths

Computing the shortest path that avoids a series of obstacles and connects two given points belongs to the most fundamental problems in computational geometry. When looking for the shortest path between two given points inside a simple polygon \mathcal{P} of size n , the currently best known algorithm is due to Guibas and Hershberger [12]: they provide a method that, after an $O(n)$ time preprocessing, can report the geodesic distance between any two points in \mathcal{P} in $O(\log n)$ time. If the geodesic is to be reported, their method needs $O(k + \log n)$ time instead, where k is the number of vertices of the path. By applying our polygon simplification strategy, we reduce the query time to $O(\log r)$.

Corollary 1. *We can preprocess a simple polygon \mathcal{P} of n vertices out of which $r > 0$ are reflex in $O(n)$ time and space, such that for any two points $p, q \in \mathcal{P}$ we can determine their geodesic distance $|\pi_{\mathcal{P}}(p, q)|$ in $O(\log r)$ time. Moreover, the geodesic path can be reported in $O(k + \log r)$ time, where k is the number of vertices of $\pi_{\mathcal{P}}(p, q)$.*

Many variations of the above problem have been studied in the literature (see [21] for a survey). Among several results, we highlight the algorithm of Hershberger and Suri [16] for computing shortest paths in the case where holes are also present. The running time of their preprocessing algorithm is bounded by $O(n \log n)$, which can again be improved by our approach.

Corollary 2. *We can preprocess a polygon \mathcal{P} with holes of n vertices out of which $r > 0$ are reflex in $O(n + r \log r)$ time and space, such that for any point $p \in \mathcal{P}$ we can determine the geodesic distance $|\pi_{\mathcal{P}}(p, q)|$ with respect to a fixed point $q \in \mathcal{P}$ in $O(\log r)$ time. Moreover, a geodesic path can be reported in $O(k + \log r)$ time, where k is the number of vertices of $\pi_{\mathcal{P}}(p, q)$.*

Shortest path computation has also been studied in other metrics, like, for example, the L_1 metric [20]. Although the proposed algorithm claims a running time of $O(n \log n)$, it is easy to see that convex vertices only contribute a linear fraction to the running time. Thus, the running time of the algorithm in [20] can be bounded by $O(n + r \log r)$ using a simple counting argument.

4.2 Geodesic Hull

A set $S \subseteq \mathcal{P}$ is called *geodesically* (or relative) convex if and only if for any $p, q \in S$, it holds that their geodesic $\pi_{\mathcal{P}}(p, q)$ is in S . The *geodesic hull* of a set $S \subseteq \mathcal{P}$ is defined as the (inclusion-wise) smallest geodesically convex set that contains S .

Given a set S of m points and a simple polygon \mathcal{P} of n vertices, Toussaint [27] studied the problem of determining whether S is geodesically convex, and — if not — computing its geodesic hull. The proposed algorithm runs in $O((n + m) \log(n + m))$ time, which we can reduce with our approach.

Corollary 3. *Given a simple polygon \mathcal{P} of n vertices out of which $r > 0$ are reflex, we can compute the geodesic hull of a given set S of m points in the interior of \mathcal{P} in $O(n + (m + r) \log(m + r))$ time using $O(n + m)$ space.*

4.3 Separating Point Sets

Given two sets R and B of m points in \mathbb{R}^2 and a geometric object ζ (usually a line) that partitions the plane into two components, we say that ζ *separates* R and B if each component of the plane only contains elements of one of the two sets. The extension of separability to polygons with holes was studied by Demaine et al. [9]. They showed that finding the minimum number of chords that separate the two sets inside a polygon with holes is NP-hard.

However, if \mathcal{P} is a simple polygon, the problem becomes easier: in [9] the authors give necessary and sufficient conditions for the existence of a separating geodesic in simple polygons, which results in an algorithm that runs in $O((n + m) \log(n + m))$ time to determine the separability of R and B . The combination of their result with our simplification technique reduces the dependency in n .

Corollary 4. *Given a simple polygon \mathcal{P} of n vertices out of which $r > 0$ are reflex, and two sets R and B of m points each, we can determine whether or not R and B are separable by a geodesic (and find a separating geodesic, if any exists) in $O(n + (m + r) \log(m + r))$ time using $O(n + m)$ space.*

4.4 Triple Orientation

Many algorithms for point sets in the plane only use the orientations of point triples, i.e., a ternary predicate $p(a, b, c)$. In analogy to unconstrained point sets, the orientation of a point triple inside a simple polygon is defined via the order in which the points appear on the geodesic hull of the triple [2].

Corollary 5. *We can preprocess a simple polygon \mathcal{P} of n vertices out of which $r > 0$ are reflex in $O(n)$ time and space so that, for any three points $a, b, c \in \mathcal{P}$, we can determine their orientation in $O(\log r)$ time.*

In [2], it was shown that each point set in a simple polygon corresponds to an *abstract order type*, a generalization of point set order types in the plane (roughly speaking, an order type is an equivalence class of point sets w.r.t. the predicate p). Algorithms that operate on abstract order types can be applied in this setting, using $O(\log r)$ time per orientation test. For example, using the results of [3] one can compute a halving geodesic through a given point of a set S of m points, and also the geodesic hull edges stabbed by (the extension of) a geodesic through two given points in $O(n + m \log r)$ time.

4.5 Voronoi Diagram

Voronoi diagrams are another fundamental data structure in computational geometry, hence, it is no surprise that the geodesic variant was studied in the late 80s [5]. Since then the algorithms have been improved, and the currently best algorithm is due to Papadopolou and Lee [23]. The furthest-site Voronoi diagram has also been studied in geodesic environments [6]. The fastest algorithms for computing either diagram run in $O((n+m) \log(n+m))$ time, and use $O(n+m)$ space [6, 23]. It is easy to see that, for any point $q \in \mathcal{P}$, its nearest or furthest site (w.r.t. geodesic distance) will be the same in \mathcal{P} and in any polygon that subsumes \mathcal{P} . Thus, in principle our approach can be used. Since the boundary is part of the Voronoi diagram, some post-processing will be necessary for this problem.

Corollary 6. *Given a simple polygon \mathcal{P} of n vertices out of which $r > 0$ are reflex, and a set S of m sites, we can compute the nearest- and furthest-site geodesic Voronoi diagram of S with respect to \mathcal{P} in $O(n + (m+r) \log(m+r))$ time using $O(n+m)$ space.*

4.6 Geodesic Diameter, Median, and Maximum Distance

Given a set S of n sites in a simple polygon \mathcal{P} , the *geodesic median* is the site of S that minimizes the maximum geodesic distance to points of S . The *geodesic diameter* of S is the maximal distance between any two sites of S . If we are given two sets of sites S_1 and S_2 instead, their *maximum geodesic distance* is defined as the largest distance between points of different sets. Toussaint [27] introduced these concepts and gave algorithms on how to compute them. The running times

of the proposed algorithms range between $O((n+m) \log(n+m))$ and $O(n^2+m^2)$. As pointed out by Toussaint, these problems can also be solved by computing the furthest-site geodesic Voronoi diagram and making $O(m)$ queries. If we simplify the polygon before making the queries, we obtain the following result.

Corollary 7. *Given a simple polygon \mathcal{P} of n vertices out of which $r > 0$ are reflex, and two sets S_1, S_2 of m sites, we can compute the geodesic median, the geodesic diameter of S_i (for $i \leq 2$), and the maximum geodesic distance between S_1 and S_2 in $O(n + (m+r) \log(m+r))$ time using $O(n+m)$ space.*

5 Conclusion

As already mentioned, while asymptotically tight, our method does not construct optimal subsuming polygons, i.e., subsuming polygons with the minimum number of vertices. In particular, every edge of the subsuming polygon \mathcal{P}' completely contains an edge of the initial polygon \mathcal{P} . One can easily construct examples where this is not the case for optimal subsuming polygons.

An optimal subsuming polygon consists of several disjoint paths with a minimum overall number of vertices. In general, the minimum link path problem is 3SUM-hard [22] even for a single path. However, the reduction used in [22] is not applicable to our restricted setting where the paths are in a domain without holes. Inside simple polygons, the minimum link path problem is solvable in linear time [26] for a single path.

Guibas et al. [14] show that finding a minimum link simple polygon having its boundary inside a given region R is NP-hard, but they point out that their reduction requires holes in R on both sides of the polygon boundary.

The pairwise-disjoint link paths problem inside a simple polygon is discussed by Gupta and Wenger [15]. They give a constant-factor approximation algorithm, and ask whether there exists a polynomial-time algorithm for computing the optimum. We state the same question for our restricted setting: Given a simple polygon \mathcal{P} , what is the complexity of constructing a subsuming polygon with the least number of (convex) vertices?

References

1. Aichholzer, O., Hackl, T., Korman, M., Pilz, A., Vogtenhuber, B.: Geodesic-preserving polygon simplification. ArXiv e-prints (2013), arXiv:1309.3858
2. Aichholzer, O., Korman, M., Pilz, A., Vogtenhuber, B.: Geodesic order types. Algorithmica (2013), to appear
3. Aichholzer, O., Miltzow, T., Pilz, A.: Extreme point and halving edge search in abstract order types. Computational Geometry: Theory and Applications 46(8), 970–978 (2013)
4. Arkin, E.M., Chiang, Y.J., Held, M., Mitchell, J.S.B., Sacristan, V., Skiena, S., Yang, T.H.: On minimum-area hulls. Algorithmica 21(1), 119–136 (1998)
5. Aronov, B.: On the geodesic Voronoi diagram of point sites in a simple polygon. In: SoCG. pp. 39–49 (1987)

6. Aronov, B., Fortune, S., Wilfong, G.T.: The furthest-site geodesic Voronoi diagram. *Discrete and Computational Geometry* 9, 217–255 (1993)
7. Bern, M.W., Eppstein, D.: Mesh generation and optimal triangulation. In: *Computing in Euclidean Geometry*, pp. 47–123. No. 4 in *Lecture Notes Series on Computing*, World Scientific (1995)
8. Bose, P., Demaine, E.D., Hurtado, F., Iacono, J., Langerman, S., Morin, P.: Geodesic ham-sandwich cuts. In: *SoCG*. pp. 1–9 (2004)
9. Demaine, E.D., Erickson, J., Hurtado, F., Iacono, J., Langerman, S., Meijer, H., Overmars, M.H., Whitesides, S.: Separating point sets in polygonal environments. *International Journal of Computational Geometry and Applications* 15(4), 403–420 (2005)
10. Douglas, D.H., Peucker, T.K.: Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *The Canadian Cartographer* 10(2), 112–122 (1973)
11. Ghosh, S.: *Visibility Algorithms in the Plane*. Cambridge University Press, New York, NY, USA (2007)
12. Guibas, L.J., Hershberger, J.: Optimal shortest path queries in a simple polygon. *Journal of Computer and System Sciences* 39(2), 126–152 (1989)
13. Guibas, L.J., Hershberger, J., Leven, D., Sharir, M., Tarjan, R.E.: Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica* 2, 209–233 (1987)
14. Guibas, L.J., Hershberger, J., Mitchell, J.S.B., Snoeyink, J.: Approximating polygons and subdivisions with minimum link paths. *Journal of Computer and System Sciences* 3(4), 383–415 (1993)
15. Gupta, H., Wenger, R.: Constructing pairwise disjoint paths with few links. *ACM Transactions on Algorithms* 3(3) (2007)
16. Hershberger, J., Suri, S.: An optimal algorithm for euclidean shortest paths in the plane. *SIAM Journal of Computing* 28(6), 2215–2256 (1999)
17. Hershberger, J., Snoeyink, J.: Computing minimum length paths of a given homotopy class. *Computational Geometry: Theory and Applications* 4, 63–97 (1994)
18. Hertel, S., Mehlhorn, K.: Fast triangulation of simple polygons. In: *FOCS*. vol. 158, pp. 207–218 (1983)
19. Melkman, A.A.: On-line construction of the convex hull of a simple polyline. *Information Processing Letters* 25(1), 11–12 (1987)
20. Mitchell, J.S.B.: L_1 shortest paths among polygonal obstacles in the plane. *Algorithmica* 8, 55–88 (1992)
21. Mitchell, J.S.B.: Shortest paths and networks. In: *Handbook of Discrete and Computational Geometry*, pp. 607–642. Chapman & Hall/CRC, 2nd edn. (2004)
22. Mitchell, J.S.B., Polishchuk, V., Sysikaski, M.: Minimum-link paths revisited. *ArXiv e-prints* (2013), arXiv:1302.3091
23. Papadopoulou, E., Lee, D.T.: A new approach for the geodesic Voronoi diagram of points in a simple polygon and other restricted polygonal domains. *Algorithmica* 20(4), 319–352 (1998)
24. Rote, G., Santos, F., Streinu, I.: Pseudo-triangulations—a survey. In: *Surveys on Discrete and Computational Geometry—Twenty Years Later*, pp. 343–410. *Contemporary Mathematics* (2008)
25. Speckmann, B., Tóth, Cs.D.: Allocating vertex π -guards in simple polygons via pseudo-triangulations. *Discrete and Computational Geometry* 33(2), 345–364 (2005)
26. Suri, S.: A linear time algorithm for minimum link paths inside a simple polygon. *Computer Vision, Graphics, and Image Processing* 35(1), 99–110 (1986)

27. Toussaint, G.T.: Computing geodesic properties inside a simple polygon. *Revue D'Intelligence Artificielle* 3(2), 9–42 (1989)