

The Furthest-Site Geodesic Voronoi Diagram*

Boris Aronov,¹ Steven Fortune,² and Gordon Wilfong²

¹ Courant Institute of Mathematical Sciences, New York University,
251 Mercer Street, New York, NY 10012, USA

² AT&T Bell Laboratories, 600 Mountain Avenue,
Murray Hill, NJ 07974, USA

Abstract. We present an $O((n+k)\log(n+k))$ -time, $O(n+k)$ -space algorithm for computing the furthest-site Voronoi diagram of k point sites with respect to the geodesic metric within a simple n -sided polygon.

1. Introduction

A common goal of much recent research in computational geometry is to extend algorithms that have been developed for the Euclidean metric to the more complicated geodesic metric inside a simple polygon. The geodesic distance between two points in a simple polygon is the length of the shortest path connecting the points that remains inside the polygon. For example, Toussaint [T] gives an algorithm for the “relative convex hull” of a set of points inside a simple polygon; Aronov [A] gives an algorithm for the nearest-neighbor geodesic Voronoi diagram; and Pollack *et al.* [PSR] give an algorithm for the “geodesic center” of a simple polygon.

A classic structure in the Euclidean metric is the “furthest-site Voronoi diagram.” Given a finite collection of point sites in the plane, the furthest-site Voronoi diagram partitions the plane into Voronoi cells, one cell per site. For each point in a cell, the owner site is the site furthest from the point, among all sites. Using

* A preliminary version of this paper appeared in the *Proceedings of the Fourth ACM Symposium on Computational Geometry*, 1988. The work of Boris Aronov was supported by an AT&T Bell Laboratories Ph.D. Scholarship. Part of the work was performed while he was at AT&T Bell Laboratories, Murray Hill, NJ, USA. His current address is Computer Science Department, Polytechnic University, Brooklyn, NY 11201, USA.

well-know algorithms, the Euclidean furthest-site Voronoi diagram of k sites can be computed in time $O(k \log k)$ and space $O(k)$ [PS].

The content of this paper is an efficient algorithm for computing the furthest-site Voronoi diagram defined by the geodesic metric inside a simple polygon. The algorithm uses $O((n + k) \log(n + k))$ time and $O(n + k)$ space, where n is the number of bounding edges of the polygon and k is the number of sites. The best previous algorithm for this problem had running time $O(n^3 \log \log n)$ [AT], and just computed (a superset of) the vertices of the furthest-site Voronoi diagram of the n corners of the polygon. We also remark that our furthest-site geodesic Voronoi diagram algorithm is a factor of $O(\log n)$ faster than the best known *nearest-site* geodesic Voronoi diagram algorithm [A].

The problem of computing the furthest-site Voronoi diagram is an extension of the “furthest-neighbor problem” which is: “Given a finite collection of points, for each point identify the element in the collection that is maximally distant from it.” Suri [S] shows how to solve a special case of the furthest-neighbor problem in the geodesic metric inside a simple polygon. Specifically, he gives an algorithm that for each corner of the polygon computes the corner that is furthest from it in the geodesic metric. His algorithm runs in time $O(n \log n)$ and space $O(n)$, where n is the number of bounding edges of the polygon.

The geodesic furthest-site Voronoi diagram generalizes the geodesic furthest-neighbor mapping of Suri [S] in two ways. First, the Voronoi diagram provides a planar partition of the polygon together with its interior into furthest-site Voronoi cells. Consequently, arbitrary furthest-site queries can be answered using a planar point-location algorithm. Second, the set of sites is not restricted to the corners of the polygon. Rather, the sites can be arbitrarily situated in the polygon. Both of these generalizations have substantial technical impact on the algorithm for computing furthest-site Voronoi diagrams.

There are many analogies between the Euclidean furthest-site Voronoi diagram and the geodesic furthest-site Voronoi diagram. In the Euclidean case, if a site has nonempty Voronoi cell, then it is extreme, i.e., it lies on the convex hull of the set of sites and it does not appear on the line segment between two sites. The counterclockwise sequence of Voronoi cells (at infinity) is the same as the counterclockwise sequence of sites on the convex hull. In the geodesic case, a site with nonempty Voronoi cell is also extreme, i.e., it lies on the relative convex hull of the set of sites and does not appear on a geodesic between two sites. The counterclockwise order of Voronoi cells along the boundary of the polygon is a subsequence of the counterclockwise order of sites on the relative convex hull. In the Euclidean case every extreme site has nonempty Voronoi cell; in the geodesic case the cell of an extreme site may be empty, roughly because the polygon is not large enough for the cell to appear.

A further analogy between the two cases is the structure of the Voronoi diagram itself. In the Euclidean case the Voronoi diagram forms a tree with root at the Euclidean center of the set of sites. (The center of a set of point sites is the point that minimizes the maximum distance to any site.) If edges are directed toward the root, then this orientation is consistent with geometric direction toward the

center. In the geodesic case exactly the same properties hold, substituting “geodesic center” for “Euclidean center” and “geodesic direction” for “direction.”

The algorithm for computing the furthest-site Voronoi diagram consists of two steps. First, we compute the restriction of the Voronoi diagram to the boundary of the polygon. Intuitively, the boundary of the polygon in the geodesic case corresponds to points “at infinity” in the Euclidean case. Second, we extend the diagram to the interior of the polygon. Because the Voronoi diagram forms a tree with root at the geodesic center, the second step is easy. It can be performed by a “reverse geodesic sweep” toward the geodesic center.

The first step, the computation of the Voronoi diagram on the boundary of the polygon, is much more involved. We use a technique developed by Suri [S] for determining furthest neighbors: we reduce the problem to three instances of the “two-fragment problem”; an instance of the two-fragment problem consists of a fragment of the boundary of the polygon and a fragment of the relative convex hull of the set of sites. The relative convex hull fragment contains the furthest sites of all points on the polygon boundary fragment. We solve an instance of the two-fragment problem using divide and conquer in the following manner. The polygon boundary fragment is split at its midpoint; this implies a corresponding split of the convex hull fragment. Thus one instance of the two-fragment problem results in two simpler instances. Eventually instances become small enough to be solved directly.

We refine Suri’s two-fragment technique in two ways. First, Suri’s algorithm always splits the polygon boundary fragment at a corner of the polygon. This is sufficient for the furthest-neighbor problem, because furthest-neighbor information for points along a wall is not of interest. For the furthest-site Voronoi diagram, simply splitting at corners is insufficient, since potentially many Voronoi cells meet a single wall of the polygon. If necessary, we further split each wall into subsegments so that the shortest-path tree from a point in a subsegment to the sites is combinatorially invariant over the entire subsegment. The combinatorial invariance of the shortest-path tree implies that the Voronoi partition of the subsegment can be easily computed.

The second refinement of Suri’s technique concerns the complexity analysis of the recursion. Suri’s original algorithm required a step called “trimming”; trimming a two-fragment instance introduces a different subproblem that could be solved directly. This operation is necessary in Suri’s analysis in order to maintain the linearity of the total size of all subproblems at a particular level of recursion. We show that even without trimming, the total size of all subproblems at a particular level is linear. This observation simplifies the recursive structure of the “two-fragment” algorithm so that it actually matches the description given above. The analysis has been incorporated into the final version of Suri’s furthest-neighbor algorithm [S].

The best lower bound that we know for computing the furthest-site geodesic Voronoi diagram is $\Omega(n + k \log k)$. It follows from known lower bounds for diameter computation in the Euclidean case. Conceivably, the current algorithm could be improved to match this lower bound.

2. The Furthest-Point Geodesic Voronoi Diagram

This section contains the definition of the furthest-point geodesic Voronoi diagram and some of its basic properties. We begin by discussing geodesics in Section 2.2. Section 2.3 contains a fairly extensive treatment of relative convex hulls. The notion of a “far side” of a relative convex hull is developed in Section 2.4; this is a technical idea used to prove the Ordering Lemma in Section 2.7. Sites are assumed to be in “general position”; this assumption and its consequences are discussed in Section 2.5. In Section 2.6 we actually give the definition of the Voronoi diagram. Section 2.7 contains the Ordering Lemma, which states that the order of Voronoi cells around the boundary of the containing polygon (i.e., the sequence of sites, as determined by the Voronoi cells meeting consecutive sections of the boundary) is the same as the order of sites around the relative convex hull of the set of sites. In Section 2.8 we define a refined form of the Voronoi diagram and use it to show a linear bound on the descriptive complexity of the (unrefined) Voronoi diagram. Finally, in Section 2.9 we show that the Voronoi diagram forms a tree directed toward the geodesic center of the set of sites. The algorithm for computing the diagram appears in Section 3.

2.1. Preliminaries

The universe U is a compact region in the plane whose boundary ∂U is a simple n -sided polygon. The set of *sites* S is a collection of k points of U . A vertex of ∂U is called a *corner* and a segment of ∂U is a *wall*. A corner is *reflex* if the measure of its interior angle is more than π and *convex* if it is less than π . If x and y are distinct points of ∂U , then *boundary fragment* $\partial U[x, y]$ is the portion of ∂U counterclockwise from x to y inclusive. The symbol ∂ denotes the boundary of a set relative to the whole plane, rather than to any proper subset of the plane. The terms “relative boundary” and “relative interior” used without any other qualification mean “relative to U .”

A *polygonal path* is a simple path comprised of a sequence of line segments. If p is a polygonal path, then the *size* of p , $|p|$, is the number of maximal segments contained in p and not containing a corner of ∂U in their interiors, and the *length* of p is the sum of their (Euclidean) lengths. A *polygonal region* is a compact set whose boundary is the union of a finite number of line segments. We allow points as (degenerate) line segments, so for example a finite set of points is also a polygonal region.

2.2. Geodesics

For points $x, y \in U$, the *geodesic path* $g(x, y)$ is the shortest path in U connecting x and y . Such a shortest path is called simply a *geodesic*. This path is unique. In fact, it is a polygonal path with interior vertices only at reflex corners of ∂U [LP]. We often consider $g(x, y)$ directed from x to y . A *link* of $g(x, y)$ is a maximal segment

of $g(x, y)$ not containing any corners in its interior; clearly an endpoint of a link is either x , y , or a reflex corner. The first endpoint of the last link of $g(x, y)$ is the anchor of y with respect to x ; it is either a reflex corner of ∂U or x itself.

We make heavy use of the fact that the intersection of any two geodesics is connected and is itself a geodesic. This fact follows immediately from uniqueness of geodesics. Two geodesics *overlap* if they intersect in more than a single point.

The *geodesic distance* $d(x, y)$ between points x and y is the length of $g(x, y)$. The geodesic distance is a metric; in particular, it is continuous as a function of both x and y and satisfies the triangle inequality $d(x, z) \leq d(x, y) + d(y, z)$. Furthermore, by uniqueness of geodesics, $d(x, y) + d(y, z) = d(x, z)$ if and only if y lies on $g(x, z)$. We often write d_u for the function defined by $d_u(x) = d(u, x)$.

Lemma 2.2.1 [PSR, Lemma 1]. *For any $u, v, w \in U$, d_u is a convex function on $g(v, w)$ with unique local minimum (possibly at v or w). In particular, for any $z \in g(v, w)$, $z \neq v, w$, $d_u(z) < \max\{d_u(v), d_u(w)\}$.*

The *geodesic direction* $\theta(x, y)$ from x to $y \neq x$ is the direction (i.e., unit vector) from x toward the anchor of x with respect to y , that is the direction given by the first link of $g(x, y)$. For fixed y , $\theta(x, y) = -\nabla d_y(x)$, where $\nabla d_y(x)$ is the gradient of d_y with respect to x , evaluated at x [A, 3.12]. At any point x not an anchor with respect to y , $d_y(x)$ is differentiable as a function of x and $\theta(x, y)$ is continuous as a function of both x and y . For H a closed subset of U , let $\theta(x, H)$ be the set of directions from x to the points of H , i.e., $\{\theta(x, h) : h \in H, h \neq x\}$.

The *geodesic angle* $\angle xyz$ is the angle counterclockwise from $\theta(y, x)$ to $\theta(y, z)$. The measure of $\angle xyz$ is written $m\angle xyz$. The *angle between* $\theta(y, x)$ and $\theta(y, z)$ is the smaller of the two angles $\angle xyz$ and $\angle zyx$.

Lemma 2.2.2 [PSR, Corollary 2]. *If $x, y, z \in U$, $y \neq x, z$, and the angle between $\theta(y, x)$ and $\theta(y, z)$ is at least $\pi/2$, then $d(x, z) > \max\{d(x, y), d(y, z)\}$.*

The *shortest-path tree* from s , $T(s)$, is the union of the sets of links of $g(s, y)$ taken over corners y of ∂U . It has $n - 1$ or n links, depending on whether or not s itself is a corner of U [GHL⁺].

Let $P_a(s)$ be the set of points in U that have anchor a with respect to s . The *shortest-path partition* of U from s is the collection $\{P_a(s) : P_a(s) \neq \emptyset\}$. It is a planar polygonal subdivision of U ; it can be computed and in fact triangulated in linear time given a triangulation of U [GHL⁺]. We can describe the bounding edges of the shortest-path partition as follows. Suppose $P_a(s)$ is not empty, where $a \neq s$. Let ab be the first link of the geodesic $g(a, s)$ (clearly, ab is the second link of all geodesics $g(x, s)$ for $x \in P_a(s)$). Since $P_a(s)$ is not empty and a is a reflex corner of ∂U , we can extend link ab past a into U . First suppose that neither wall of ∂U incident to a overlaps this extension. Let y be the first point past a of the extension so that $y \in \partial U$. Then segment ay is the *shortest-path partition edge* (from s with anchor a), denoted $p_a(s)$. Now suppose some wall of ∂U overlaps the extension of segment ab ; then we simply define $p_a(s)$ to be this wall. It can be checked that the




Fig. 1. Geodesic $g(x, y)$, its shadow \bar{y} and foreshadow x^* (they are unique in this case), the corresponding boundary geodesic $\bar{g}(x, y) = g(x^*, \bar{y})$, and geodesic direction $\theta(x, y)$. $U[y, x]$ is shaded.

boundary of a cell of the shortest-path partition consists of an alternating sequence of shortest-path partition edges and sections of ∂U .

A set $A \subseteq U$ is *relatively convex* with respect to U if $g(x, y) \subseteq A$ whenever $x, y \in A$; the *relative convex hull* of set F , denoted $R(F)$, is the smallest relatively convex set containing F , i.e., the intersection of all relatively convex sets containing F [T]. Relatively convex sets are discussed in detail in the next section. A set is *degenerate* if it is contained in a single geodesic.

If x and y are distinct points of U , then a *shadow* of $g(x, y)$ is a point $\bar{y} \in \partial U$ so that segment $y\bar{y}$ extends the last link of $g(x, y)$ while staying in U . Clearly, y is a shadow of $g(x, y)$ only if $y \in \partial U$. Shadows are not unique, indeed it is possible that every point on a subsegment of a wall is a shadow of $g(x, y)$. Similarly, a *foreshadow* of $g(x, y)$ is a point $x^* \in \partial U$ lying on a segment contained in U extending the first link of $g(x, y)$ backward. Equivalently, a foreshadow of $g(x, y)$ is a shadow of $g(y, x)$. For an illustration of these definitions, see Fig. 1.

A *boundary geodesic* is a geodesic connecting two distinct points of ∂U . Let x^* and \bar{y} be the foreshadow closest to x and shadow closest to y of $g(x, y)$, respectively. Then $\bar{g}(x, y)$ denotes the boundary geodesic $g(x^*, \bar{y})$. Geodesic $\bar{g}(x, y)$ splits U into two simply connected polygonal regions $U[x, y]$ and $U[y, x]$ with disjoint interiors; $\partial(U[x, y])$ is $\partial U[x^*, \bar{y}] \cup g(\bar{y}, x^*)$ and $\partial(U[y, x])$ is $g(x^*, \bar{y}) \cup \partial U[\bar{y}, x^*]$. (Note that $\partial(U[x, y])$ is distinct from $\partial U[x, y]$.) $U[y, x]$ is shaded in Fig. 1. Intuitively, $U[x, y]$ contains points lying on or to the right of $\bar{g}(x, y)$, while points on the geodesic or to the left of it constitute $U[y, x]$. Notice that $\bar{g}(x, y)$ is exactly the common boundary of $U[x, y]$ and $U[y, x]$; hence any geodesic from a point in $U[x, y]$ to a point in $U[y, x]$ must intersect $\bar{g}(x, y)$. Both $U[x, y]$ and $U[y, x]$ are relatively convex, since a geodesic connecting two points of, say, $U[x, y]$ must have connected intersection with $\bar{g}(x, y)$.

Boundary geodesics are intended to model infinite Euclidean lines, just as geodesics correspond to line segments. Unfortunately, there may be more than one boundary geodesic containing $g(x, y)$; $\bar{g}(x, y)$ is meant to be the canonical choice. It is the smallest boundary geodesic containing $g(x, y)$. Similarly, $U[x, y]$ models the Euclidean half-plane lying to the right of the line directed from x to y . (A different notion of “half-plane” based on geodesic distance from a point is introduced in Section 2.6).

The following two lemmas capture some basic properties of boundary geodesics. Suppose point w lies on boundary geodesic $g(u, v)$ and point x is in $U[u, v]$. The first lemma concerns shadows of $g(w, x)$; the second shadows of $g(x, w)$. As an exercise, the reader is encouraged to compare them with corresponding statements in Euclidean geometry or, equivalently, with the case when U is convex. The number of conditions in the lemmas is due to the large variety of degeneracies that may occur in the geodesic universe.

Lemma 2.2.3. *Suppose $u, v \in \partial U$, $u \neq v$, $w \in g(u, v)$, $x \in U[u, v]$, and $w \neq x$.*

- (1) *If $x \notin g(u, v)$, then any shadow \bar{x} of $g(w, x)$ lies in $\partial U[u, v]$.*
- (2) *If $x \in g(u, v)$, then some shadow \bar{x} of $g(w, x)$ lies in $\partial U[u, v]$.*

Proof. (1) Suppose $x \notin g(u, v)$. Then $g(w, \bar{x})$ cannot intersect $\bar{g}(u, v)$ again after x , so $\bar{x} \in \partial U[u, v]$.

(2) Suppose $x \in g(u, v)$; without loss of generality assume w, x, v are in that order along $g(u, v)$. We can choose $\bar{x} = v$ unless $g(u, v)$ bends at or after x . If $g(u, v)$ bends right at some point c at or after x , then since $U[u, v]$ lies locally to the right of $g(u, v)$, c must be a reflex corner of $\partial U[u, v]$, and we can choose $\bar{x} = c$. If $\bar{g}(u, v)$ bends left at some point c , then the straight-line continuation of $g(w, c)$ at c enters the interior of $U[u, v]$ and thus will not intersect $g(u, v)$ again. Hence we can choose \bar{x} to be any shadow of $g(w, c)$ distinct from c . □

Lemma 2.2.4. *Suppose $u, v \in \partial U$, $u \neq v$, $w \in g(u, v)$, $x \in U[u, v]$, $w \neq x$, and \bar{w} is the closest shadow of $g(x, w)$.*

- (1) *If $w \notin \partial U$ and $g(x, w)$ does not overlap $g(u, v)$, then $\bar{w} \in \partial U[v, u]$.*
- (2) *If v is the closest shadow of $g(u, w)$ and $g(u, w)$ is not an initial portion of $g(u, x)$ (i.e., $w \notin g(u, x)$), then $\bar{w} \in \partial U[v, u]$.*

Proof. (1) If $g(x, w)$ does not overlap $g(u, v)$, then $\theta(w, x)$ cannot be $\theta(w, u)$ or $\theta(w, v)$. Since $w \notin \partial U$, $\bar{w} \neq w$ and $\theta(w, \bar{w}) = -\theta(w, x)$ must enter the interior of $U[v, u]$ at w . As $g(x, \bar{w})$ cannot intersect $g(u, v)$ again, $\bar{w} \in \partial U[v, u]$.

(2) The statement is trivial if $w \in \partial U$. If not, then $w \neq \bar{w}$. We might have $\theta(w, x) = \theta(w, u)$, in which case $\bar{w} = v$. We cannot have $\theta(w, x) = -\theta(w, u)$, else $w \in g(u, x)$. Otherwise $\theta(w, \bar{w})$ must enter the interior of $U[v, u]$ and as before $\bar{w} \in \partial U[v, u]$. □

Boundary geodesic $g(x, y)$ separates points a and b if $a \in U[x, y]$ and $b \in U[y, x]$, or vice versa. Similarly, boundary geodesic $g(x, y)$ separates sets A and B if $A \subseteq U[x, y]$ and $B \subseteq U[y, x]$, or vice versa. Note that separation does not imply disjointness, indeed if (degenerate) set A is contained in boundary geodesic $g(x, y)$, then $g(x, y)$ separates A from itself.

For a compact set $F \subseteq U$ and $z \in U$, let $rad(z, F) = \max_{x \in F} d_z(x)$. The center of F is the point z that minimizes $rad(z, F)$. Pollack *et al.* [PSR] show that the center of the set of vertices of U is unique. In fact, with minor modifications their proof shows that the center of any compact set $F \subseteq U$ is unique.




Fig. 2. Geodesic triangle Δ_{xyz} .

For points $x, y, z \in U$ we define the *geodesic triangle* Δ_{xyz} as follows [PSR]. If $\{x, y, z\}$ is degenerate, then Δ_{xyz} is the smallest geodesic containing x, y , and z . Otherwise we can choose points x', y' , and z' so that x' is the point at which $g(x, y)$ and $g(x, z)$ diverge, and similarly for y' and z' . Refer to Fig. 2. Then the circuit $g(x', y'), g(y', z'), g(z', x')$ is a simple polygon $\Delta_{x'y'z'}$. We define Δ_{xyz} to be the union of $g(x, x'), g(y, y'), g(z, z')$, and $\Delta_{x'y'z'}$ together with its interior. All of the interior angles of $\Delta_{x'y'z'}$ are reflex except at x', y', z' [PSR]. The geodesic triangle is a special case of the relative convex hull of a finite set of points discussed in Section 2.3. A seemingly more natural definition of Δ_{xyz} as $U[x, y] \cap U[y, z] \cap U[z, x]$ could lead to unexpected results, such as inclusion of common shadows of $g(y, x)$ and $g(z, x)$ into Δ_{xyz} . The proof of the following lemma is in the same spirit as the proofs of Lemmas 2.2.3 and 2.2.4.

Lemma 2.2.5 (Triangle Lemma). *Suppose $\{x, y, z\}$ is not degenerate. Let \bar{y} and \bar{z} be shadows of $g(x, y)$ and $g(x, z)$ respectively. Assume $z \in U[y, x]$. If $u \in \Delta_{xyz}$ and $u \notin g(x, y) \cap g(x, z)$, then there is a shadow \bar{u} of $g(x, u)$ so that $\bar{u} \in \partial U[\bar{y}, \bar{z}]$ and $g(x, \bar{u})$ intersects $g(y, z)$. If u is in the interior of Δ_{xyz} , then, for any shadow \bar{u} of $g(x, u)$, $\bar{u} \in \partial U[\bar{y}, \bar{z}]$ and $g(x, \bar{u})$ intersects $g(y, z)$. If $u \in U[y, x] \cap U[x, z] \cap U[y, z]$, then $g(x, u)$ intersects $g(y, z)$ and there is a shadow \bar{u} of $g(x, u)$ in $\partial U[\bar{y}, \bar{z}]$.*

2.3. Relatively Convex Sets

This section develops properties of relatively convex sets. The main result is Lemma 2.3.4. It states that the “extreme” points of a set F can be ordered so that the relative convex hull of F is the intersection of all “cones” defined by consecutive triples of extreme points. An immediate consequence of Lemma 2.3.4 is a decomposition of a relatively convex set into a collection of simple polygons and connecting geodesics. Also, the order of extreme points extends to a natural notion of a traversal of the boundary of a relatively convex set. This ordering has a number of useful consequences, given in Lemmas 2.3.6–2.3.9.

Lemma 2.3.1. *Any relatively convex set R is simply connected.*

Proof. We show that the region enclosed by a simple cycle γ in R lies completely in R . Suppose $x \in \gamma$, y is in the interior of γ , and \bar{y} is a shadow of $g(x, y)$. Since $y\bar{y}$

connects a point inside γ to a point on or outside γ , there is a point w in the intersection of $y\bar{y}$ and γ . Since R is relatively convex, $g(x, w) \subseteq R$, so $y \in R$. \square

We now give a long sequence of definitions related to relatively convex sets. To motivate the terms, consider the simpler case of a finite point set F and its (ordinary) convex hull H in the plane. For a point x in the plane, the set $\theta(x, F)$ is a finite set of directions; define its “convex closure” $span(x, F)$ by including any direction in the angle between any α, α' in $\theta(x, F)$. Then clearly $span(x, F) = \theta(x, H)$. We can classify a point x in the plane as interior to H (meaning $span(x, F) = \theta(x, H)$ has measure 2π), exterior to H (meaning $span(x, F)$ has measure at most π), or as a thin point of H (meaning H is a line segment and x is an interior point of the line segment). Notice that an exterior point of H could well lie on the boundary of H . A special case of a boundary point is an extreme point, when $span(x, F)$ has measure strictly less than π ; of course all extreme points of H must be points of F . For x an exterior point of H , we can define the clockwise extreme point $r(x)$ to be the point of F so that $\theta(x, F)$ is as clockwise as possible (in $span(x, F)$); possibly there are two such extreme points, in which case $r(x)$ is chosen to be the furthest from x . Similarly, we can define the counterclockwise extreme point $l(x)$; then, for extreme point x , $l(r(x)) = r(l(x)) = x$. For an extreme point x of F , the sequence $x_0 = x, x_{i+1} = r(x_i)$ eventually repeats and forms a counterclockwise traversal of the boundary of H . Furthermore, H is the intersection of half-planes with bounding lines through $x_i x_{i+1}$.

Much of the remainder of this section gives analogues of these definitions in the relatively convex case and establishes basic properties of the definitions. The technical details are quite intricate because of the complexity of the geodesic setting; however, the development follows the outline just given. For the following, let F be a nonempty polygonal region contained in U .

For $x \in U$, we define $span(x, F)$ to be the smallest set of directions containing $\theta(x, F)$ so that whenever $\alpha, \alpha' \in \theta(x, F), \alpha \neq -\alpha'$, and the angle between α and α' is contained in U near x , then every direction in this angle is in $span(x, F)$. If $x \in U$ is not a reflex corner of ∂U , then it is easy to see that either $span(x, F)$ has a single component or $span(x, F)$ consists of two opposite directions. If x is a reflex corner of ∂U , then $span(x, F)$ may have two components that are not opposite directions; in fact one or both components may have positive measure. Since F is a polygonal region, it is easy to check that the endpoint of any component of $span(x, F)$ is $\theta(x, y)$ for some $y \in F$. We show below (Lemma 2.3.3) that $span(x, F) = \theta(x, R(F))$.

If $span(x, F)$ consists of a single component of measure less than 2π , then x is an *exterior point* of F . Notice $span(x, F)$ can have measure exceeding π but less than 2π only if x is a reflex corner of ∂U . If $span(x, F)$ consists of a single component of measure less than π and $x \in F$, then x is an *extreme point* of F . If $span(x, F)$ consists of two connected components, then x is a *thin point* of F . Set F is *extreme* if every element of it is extreme.

In Fig. 3, for $F = \{1, 2, 3, 4, 5\}$, point 5 and all points of ∂U except a are exterior (but none are extreme), points 1, 2, and 3 are extreme points, and every point on $g(1, a)$ except 1 is a thin point. In Fig. 3, $\{1, 2, 3\}$ is extreme.

For the remainder of this section, $F \subseteq U$ is a finite set of points containing at least two elements.




Fig. 3. Relative convex hull of {1, 2, 3, 4, 5}.

Let x be an exterior point of F or a point of ∂U . We wish to define clockwise and counterclockwise extreme points of F from x . To do this we first define a point $opp_F(x) \in \partial U$ so that, instead of the cyclic counterclockwise order on ∂U , we may speak of linear counterclockwise order on $\partial U - \{opp_F(x)\}$. If $x \in \partial U$ define $opp_F(x)$ to be x , otherwise define $opp_F(x)$ to be the first point of ∂U intersected by the ray with endpoint x directed opposite the bisector of $span(x, F)$. The *clockwise extreme point* of F from x , denoted $r(x)$, is the point $f \in F$ so that there is no geodesic extending $g(x, f)$ to a point $f' \in F$, and among all such points the closest shadow of $g(x, f)$ is as clockwise as possible in $\partial U - \{opp_F(x)\}$. Similarly define the *counterclockwise extreme point* of F from x , denoted $l(x)$. For example, in Fig. 3, $l(2) = 1$, $r(2) = 3$, $l(a) = 1$, and $r(a) = 2$. There are two subtleties to these definitions. First, r and l depend upon the set F , but, except in the proof of Lemma 2.4.3, we leave this dependence implicit. Second, there are two distinct (though overlapping) cases in the definition: either x is exterior, so that $span(x, F)$ consists of a single component, or $x \in \partial U$ and, even if $span(x, F)$ consists of two components, there is still a natural way to define $r(x)$ and $l(x)$. Notice that there is no natural definition of $r(x)$ and $l(x)$ if $span(x, F)$ has measure 2π or if $span(x, F)$ has two components and $x \notin \partial U$. The following lemma describes properties of functions r and l ; its proof is highly technical, detailed, and sometimes unintuitive.

Lemma 2.3.2. *Let x be an exterior point of F or a point of ∂U , $r = r(x)$, $l = l(x)$, let \bar{r} be the closest shadow of $g(x, r)$, and let \bar{l} be the closest shadow of $g(x, l)$.*

- (1) *If x is exterior, then $span(x, F) = \angle rxl$. If $x \in \partial U$ or $m\angle rxl \neq \pi$, then $F \subseteq U[r, x] \cap U[x, l]$; if $m\angle rxl = \pi$, then $F \subseteq U[r, l]$.*
- (2) *If $r \neq l$, then $r\bar{r}$ and $l\bar{l}$ are disjoint; if also $x \in \partial U$, then x, \bar{r}, \bar{l} are in that counterclockwise order around ∂U .*
- (3) *Both r and l are extreme points of F .*
- (4) *If x is an extreme point of F , then $r(l(x)) = l(r(x)) = x$.*
- (5) *Let $x \in \partial U$ or $m\angle rxl < \pi$. Then $f = r(x)$ if and only if $f \in F$, $F \subseteq U[f, x]$, and $g(x, f)$ cannot be extended to $g(x, f')$ for any other $f' \in F$.*

Proof. (1) We must always have $span(x, F) \subseteq \angle rxl$ by definition of $span(x, F)$ and of points r and l ; certainly $\theta(x, r), \theta(x, l) \in span(x, F)$. If x is exterior, then $span(x, F)$ is connected, so we must have $\angle rxl = span(x, F)$. For the second statement, first

suppose $x \in \partial U$ or $m \angle rxl \neq \pi$. Let x^* be the closest foreshadow of $g(x, r)$. We claim $\bar{l} \in \partial U[\bar{r}, x^*]$: this follows immediately from the definition of r and l if $m \angle rxl < \pi$, $x \in \partial U$, or $m \angle rxl > \pi$ (since in the last case necessarily $x \in \partial U$). For any $f \in F$, either $f \in g(x, r)$, $f \in g(x, l)$, or the closest shadow of $g(x, f)$ lies in $\partial U[\bar{r}, \bar{l}] \subseteq \partial U[\bar{r}, x^*]$, so $f \in U[r, x]$, and $F \subseteq U[r, x]$. Similarly, $F \subseteq U[x, l]$. The argument that $m \angle rxl = \pi$ implies $F \subseteq U[r, l]$ is similar. (For an illustration of this somewhat unintuitive case, refer to Fig. 3, placing x at any point of segment ab other than a or b . For such a point x , $r(x) = 1$, $l(x) = 3$, and indeed $F = \{1, 2, 3, 4, 5\} \subseteq U[1, 3]$. On the other hand, $F \not\subseteq U[x, 3] = U[a, 3]$.)

(2) Suppose $l \neq r$. If \bar{l} and \bar{r} met, either l would lie on $g(x, r)$ or vice versa, contrary to the definition of r, l . The ordering of x, \bar{r} , and \bar{l} follows immediately by definition.

(3) We assume $F \subseteq U[r, x]$, otherwise a similar argument works using $U[r, l]$. To show r extreme, we need to show $span(r, F)$ is connected and has measure less than π . Whether or not $r \in \partial U$, $\theta(r, U[r, x])$ is connected. As $F \subseteq U[r, x]$, $span(r, F) \subseteq \theta(r, U[r, x])$. Hence it suffices to show that, for $f \in F$, $\theta(r, f)$ lies in the angle from $\theta(r, x)$ clockwise to but not including $-\theta(r, x)$. Now $\theta(r, f)$ cannot be $-\theta(r, x)$ else $g(x, f)$ would extend $g(x, r)$. Also $\theta(r, f)$ can be clockwise of $-\theta(r, x)$ only in the case that r is a reflex corner of ∂U lying to the right of $g(x, r)$; again this is impossible because $g(x, f)$ would extend $g(x, r)$. The possibility that $span(r, F)$ has measure less than π but is not connected is excluded by a similar argument.

(4) By (1), $F \subseteq U[r, x]$. For any $f \in F$ not appearing on $g(x, r)$, f is in the relative interior of $U[r, x]$. Thus, by Lemma 2.2.3(1), the closest shadow of $g(r, f)$ lies on $\partial U[\bar{r}, x^*]$, where x^* is the closest foreshadow of $g(x, r)$. This implies $x = l(r) = l(r(x))$ by definition of counterclockwise extreme point.

(5) Clearly, there can be at most one point f in F satisfying “ $F \subseteq U[f, x]$ and $g(x, f)$ cannot be extended to $g(x, f')$ for any $f' \in F$.” Since $f = r(x)$ is one such point, the claim follows. □

Lemma 2.3.3. For any $x \in U$, $span(x, F) = \theta(x, R(F))$.

Proof. Suppose $y, z \in F$ are such that $0 < m \angle yxz < \pi$ and $\angle yxz$ is (locally around x) contained in U . By examining $\triangle xyz$ and using $g(y, z) \subseteq R(F)$ we see that $\angle yxz \subseteq \theta(x, R(F))$. Hence $span(x, F) \subseteq \theta(x, R(F))$.

If $span(x, F)$ has measure 2π , then it is immediate that $span(x, F) = \theta(x, R(F))$. Suppose x is an exterior point of F ; let $r = r(x)$ and $l = l(x)$. If $m \angle rxl \neq \pi$, then $R(F) \subseteq U[r, x] \cap U[x, l]$ since $F \subseteq U[r, x] \cap U[x, l]$ by Lemma 2.3.2(1), and $U[r, x]$ and $U[x, l]$ are relatively convex. Hence

$$\theta(x, R(F)) \subseteq \theta(x, U[r, x] \cap U[x, l]) = \angle rxl = span(x, F),$$

where the first equality follows by definition and the second by Lemma 2.3.2(1). If $m \angle rxl = \pi$, then $R(F) \subseteq U[r, l]$, so

$$\theta(x, R(F)) \subseteq \theta(x, U[r, l]) = \angle rxl = span(x, F).$$

Finally, if x is a thin point of F , we can find a segment through x splitting U into two simple polygons, each containing one of the components of $span(x, F)$. Let F_1 and F_2 be the intersections of these two polygons with F , and let H_1, H_2 be their respective convex hulls (relative to U). It is easy to verify that x is an exterior point of both F_1 and F_2 and thus, by the first part of the proof,

$$span(x, F_1) \cup span(x, F_2) = \theta(x, H_1) \cup \theta(x, H_2) \subseteq \theta(x, R(F)).$$

Since trivially $span(x, F_1) \cup span(x, F_2) \subseteq span(x, F)$, it is sufficient to show that the above inclusion is an equality. Suppose there is a point $y \in R(F) - \{x\}$ such that $\theta(x, y) \in \theta(x, R(F)) - \theta(x, H_1) \cup \theta(x, H_2)$. As $y \in R(F) = R(H_1 \cup H_2)$, there are two points, $z_1 \in H_1$ and $z_2 \in H_2$ so that $y \in g(z_1, z_2)$. Notice $x \notin g(z_1, z_2)$, hence $\Delta xz_1z_2 \subseteq H(F)$ is not degenerate and $\theta(x, g(z_1, z_2)) \subseteq span(x, F)$ is the angle between $\theta(x, z_1)$ and $\theta(x, z_2)$, contradicting the assumption that x is a thin point and $\theta(x, z_1)$ and $\theta(x, z_2)$ lie in different components of $span(x, F)$. \square

It is immediate that any two sets with the same relative convex hull have the same extreme, exterior, and thin points.

The next lemma is the main result of this section. It gives a characterization of the relative convex hull of F . In the Euclidean plane a convex hull is bounded by a chain of segments connecting extreme points of F . The hull can be constructed by intersecting all half-planes defined by pairs of consecutive extreme points. The analogue of the first statement holds in geodesic metric inside a simple polygon. The second statement, however, is false, if $U[x, y]$ is taken to play the role of a half-plane defined by x and y . Instead we represent $R(F)$ as the intersection of "geodesic cones" originating from each extreme point and containing F . More specifically, suppose $m \angle xyz < \pi$ and $z \in U[x, y]$. Let s be the segment from y to the closest foreshadow y^* of $g(y, x)$, open at y and closed at y^* . The geodesic cone $U[x, y, z]$ is $U[x, y] \cap U[y, z] - \{s\}$. (This is an analogue of a Euclidean cone; see Fig. 4.) Notice that s intersects $U[x, y] \cap U[y, z]$ only if s is also a foreshadow segment of $g(y, z)$. It is easily checked that $U[x, y, z]$ is relatively convex and that the boundary of $U[x, y, z]$ consists of $g(\bar{z}, y)$, $g(y, \bar{x})$, and $\partial U[\bar{x}, \bar{z}]$, where \bar{x} and \bar{z} are the closest shadows of $g(y, x)$ and $g(y, z)$, respectively. Furthermore, if f is an extreme point of F , then $F \subseteq U[r(f), f, l(f)]$.

Lemma 2.3.4. $H = R(F)$ is a simply connected polygonal region. The extreme




Fig. 4. Geodesic cones $U[2, 1, 3]$ (solid line) and $U[3, 2, 1]$ (dashed line).




Fig. 5. Illustration to the proof of Lemma 2.3.4.

elements of F can be labeled $f_1, f_2 = r(f_1), \dots, f_0 = f_m = r(f_{m-1}), f_1 = f_{m+1} = r(f_m)$ so that $H = \bigcap_{i=1}^m U[f_{i+1}, f_i, f_{i-1}]$ and $\partial H = \bigcup_{i=1}^m g(f_i, f_{i+1})$.

Proof. F must contain some extreme element, since $r(u)$ is extreme for any $u \in \partial U$. Let f_1 be an extreme element of F , and consider the sequence $f_2 = r(f_1), f_3 = r(f_2), \dots$. Since F is finite and $l \circ r$ is the identity (Lemma 2.3.2(4)), it must be that $f_{m+1} = f_1$ for some m with f_1, \dots, f_m distinct. If F is degenerate, then $m = 2$ and the lemma follows easily. We assume F is nondegenerate, hence $m > 2$ and points in each triple f_{i-1}, f_i, f_{i+1} are distinct.

Let $I = \bigcap_{i=1}^m U[f_{i+1}, f_i, f_{i-1}]$. We first show that $\partial I = \bigcup_{i=1}^m g(f_i, f_{i+1})$. Let $g(f_i, f_{i+1})$ have closest foreshadow s_i and closest shadow t_i . A schematic view of the situation is given in Fig. 5 which also indicates the relative positions of these points, to be justified below. Then

$$\partial(U[f_{i+1}, f_i, f_{i-1}]) = g(s_{i-1}, f_i) \cup g(f_i, t_i) \cup \partial U[t_i, s_{i-1}].$$

By Lemma 2.3.2(1) we have, for $j = 1, \dots, m, f_{j-1} \in U[f_{j+1}, f_j] = U[t_j, s_j]$, so using Lemmas 2.2.3 and 2.2.4, we have $s_{j-1}, s_j, t_{j-1}, t_j$ in that counterclockwise order on ∂U (possibly $s_j = t_{j-1}$). Consequently, $\bigcup_{j=1}^m \partial U[s_j, t_j] = \partial U$. Hence

$$\bigcap_{j=1}^m \partial U[t_j, s_{j-1}] \subseteq \bigcap_{j=1}^m \partial U[t_j, s_j] \subseteq \{s_1, \dots, s_m, t_1, \dots, t_m\}.$$

We now claim that $\partial U[t_i, s_{i-1}] \cap I \subseteq \bigcup_{j=1}^m g(s_j, t_j)$. Suppose $x \in \partial U[t_i, s_{i-1}] \cap I$; either $x \in \bigcap_{j=1}^m \partial U[t_j, s_{j-1}]$, in which case the claim is immediate, or

$$x \notin \bigcap_{j=1}^m \partial U[t_j, s_{j-1}].$$

In the second case, $x \notin \partial U[t_j, s_{j-1}]$ for some j , but $x \in U[f_{j+1}, f_j, f_{j-1}]$ since $x \in I$, and since $x \in \partial U$, it must be that $x \in \partial(U[f_{j+1}, f_j, f_{j-1}])$, so $x \in g(f_j, t_j) \cup g(s_{j-1}, f_j)$, establishing the claim. We now show $\partial I = \bigcup_{i=1}^m g(f_i, f_{i+1})$. By elementary topology,

$$\begin{aligned} \partial I &= \bigcup_{i=1}^m (I \cap \partial(U[f_{i+1}, f_i, f_{i-1}])) \\ &= \bigcup_{i=1}^m ((g(f_i, t_i) \cup g(f_i, s_{i-1})) \cap I) \cup \bigcup_{i=1}^m (\partial U[t_i, s_{i-1}] \cap I) \\ &= \bigcup_{i=1}^m g(s_i, t_i) \cap I, \end{aligned}$$

Since segments $s_i f_i$ and $t_{i-1} f_i$ are in $U[f_{i+1}, f_i, f_{i-1}]$ only if $s_i = t_{i-1} = f_i \in \partial U$, we have $\partial I = \bigcup_{i=1}^m g(f_i, f_{i+1})$.

For each i , we have $H \subseteq U[f_{i+1}, f_i, f_{i-1}]$ since $U[f_{i+1}, f_i, f_{i-1}]$ is relatively convex and contains F . Hence $H \subseteq I$. Also $\partial I \subseteq H$ since each geodesic is contained in H ; since H is simply connected, $H = I$. Every extreme element f of F must be in H and in fact on ∂H , since f is on the boundary of $U[r(f), f, l(f)]$. Since no extreme element can lie on a geodesic between two other extreme elements, each extreme element is equal to f_i for some i . □

Since $R(F)$ is a connected polygonal region, it can be decomposed into *plateaus* and *bridges* as follows: a *nondegenerate plateau* is a maximal compact two-dimensional subset of $R(F)$ whose boundary is a simple polygon and a *bridge* is either a maximal polygonal path of positive length not containing any plateau points in its interior or a single point shared between two plateaus. We require that both endpoints of a bridge lie in a plateau, thus introducing a *degenerate plateau* at the endpoint of any positive-length bridge that does not end in a point of a nondegenerate plateau. In Fig. 3, point 1 and geodesic triangle $\triangle 23a$ are plateaus (degenerate and nondegenerate, respectively); segment $1a$ is a bridge. Each plateau is relatively convex and each bridge is a geodesic. Since $R(F)$ is simply connected, this decomposition forms a tree, with plateaus taken as nodes and bridges as edges. An extreme point x of $R(F)$ is a convex vertex of a plateau if $\text{span}(x, F)$ has positive measure, or is a plateau itself if $\text{span}(x, F)$ consists of a single direction. A thin point x of $R(F)$ is an interior point of a bridge if $\text{span}(x, F)$ consists of exactly two directions, an endpoint of a bridge of positive length if $\text{span}(x, F)$ consists of two components only one of which has positive measure, or a bridge by itself if $\text{span}(x, F)$ consists of two components both of positive measure.

For F labeled as in Lemma 2.3.4 and $H = R(F)$, we define the *counterclockwise traversal* of ∂H to be the circuit $g(f_1, f_2), \dots, g(f_m, f_1)$. The counterclockwise traversal visits every thin point twice and other points once. Hence for x, y not thin points we can unambiguously define $\partial H[x, y]$ to be the section of the counterclockwise traversal of ∂H from x to y . Similarly, if the points in $H' \subseteq \partial H$ are not thin points, they have an unambiguous counterclockwise ordering on ∂H . If F is extreme, then the counterclockwise ordering covers every point of F . Even

if the sequence of points x_1, \dots, x_k on ∂H includes thin points, it is still meaningful to say that they appear in counterclockwise order, if they are visited in that order in a single counterclockwise traversal of ∂H . Of course, another order differing in the position of thin points may be consistent with a counterclockwise traversal as well.

Corollary 2.3.5. *The center of F is the center of $R(F)$.*

Proof. It is sufficient to show that, for any $x \in U$, $rad(x, F) = rad(x, R(F))$. A point of $H = R(F)$ maximally distant from x must lie on ∂H . By the previous lemma any point on ∂H lies on a geodesic connecting two points of F ; hence by Lemma 2.2.1, the maximally distant point must be an endpoint of such a geodesic, that is, a point of F . Thus a point of H furthest from x is necessarily a point of F and $rad(x, H) = rad(x, F)$. \square

Lemma 2.3.6. *If f, f' are extreme points of F , then every extreme point counterclockwise from f to f' is in $U[f, f']$, and all but f and f' lie in the relative interior of $U[f, f']$.*

Proof. If $r(f) = f'$, the claim is vacuously true. Otherwise we prove it by induction on the extreme points of F lying in counterclockwise order between f and f' . If $r(f) \neq f'$, then since f' must lie in the relative interior of $U[r(f), f']$, $r(f)$ must lie in the relative interior of $U[f, f']$. Now suppose f'' is counterclockwise of f before f' and $r(f'') \neq f'$. Inductively, assume that f'' is in the relative interior of $U[f, f']$. Again we have f' in the relative interior of $U[r(f''), f']$, so $r(f'')$ must be in the relative interior of $U[f'', f']$. Now $U[f'', f']$ is contained in $U[f, f']$ except possibly for a region bounded by segments $f's'$ and $f's''$ and $\partial U[s', s'']$, where s' and s'' are the closest shadows of $g(f, f')$ and $g(f'', f')$, respectively. However, $g(f'', r(f''))$ cannot intersect segment $f's'$, since f' is extreme. Hence $r(f'') \in U[f, f']$. Since $r(f'')$ cannot lie on $\bar{g}(f, f')$ as $f, f', r(f'')$ are extreme and distinct, in fact $r(f'')$ is in the relative interior of $U[f, f']$. \square

An immediate consequence of Lemma 2.3.6 is that counterclockwise order of extreme points is an absolute order, not depending on other extreme points: if a, b, c are extreme points of both sets A and B , then their order around A is the same as it is around B .

Lemma 2.3.7. *If points x, y, z, w (not necessarily all distinct) occur in that order in a counterclockwise traversal of the boundary of a relatively convex polygonal region R , then $g(x, z)$ intersects $g(y, w)$.*

Proof. By creating dummy plateaus if necessary, we can assume that x, y, z , and w each lie in a plateau. We proceed by induction on the number of plateaus in the decomposition of R . If there is only a single plateau, then R is a simple polygon and the claim is immediate. Otherwise there is some bridge $g(a, b)$ whose removal splits R into relatively convex polygonal regions R_1 and R_2 . If all of x, y, z, w are contained in one of R_1 or R_2 , the claim follows by induction. Otherwise, since x, y, z, w appear in that order in a counterclockwise traversal of the boundary of R ,

there are up to symmetry two cases: x, y, a in R_1 and z, w, b in R_2 , or x, y, z, a in R_1 and b, w in R_2 . In the first case $g(x, z)$ and $g(y, w)$ both contain $g(a, b)$. In the second case x, y, z, a appear in that order in a counterclockwise traversal of R_1 , hence by induction hypothesis, $g(x, z)$ intersects $g(y, a)$. Since $g(y, a) \subseteq g(y, w)$, $g(x, z)$ must intersect $g(y, w)$. \square

We say segment $fx \subseteq U$ connects a relatively convex set R to ∂U if f is an extreme point of R , $x \in \partial U$, and fx intersects R only at f (possibly $f = x$ or some point of fx besides x lies on ∂U).

Lemma 2.3.8 (Connection Lemma). *Suppose f_1, \dots, f_m are extreme points of a relatively convex set R and segments $f_i x_i$ are pairwise disjoint and connect R to ∂U . Then the order of the f_i 's around ∂R is the same as the order of x_i 's around ∂U .*

Proof. Suppose f_i, f_j, f_k are in counterclockwise order around ∂R ; we show $x_j \in \partial U[x_i, x_k]$, i.e., x_i, x_j, x_k are in counterclockwise order around ∂U . By Lemma 2.3.6, f_j is in the relative interior of $U[f_i, f_k]$. Since $g(f_i, f_k) \subseteq R$ and $f_j x_j$ meets R only at f_j , $f_j x_j$ cannot intersect $g(f_i, f_k)$. Clearly, f_j is in the region bounded by $\partial U[x_i, x_k]$, segment $f_k x_k$, $g(f_k, f_i)$, and segment $f_i x_i$. Since $x_j \in \partial U$ and $f_j x_j$ does not intersect $f_k x_k$, $g(f_k, f_i)$, or $f_i x_i$, $x_j \in \partial U[x_i, x_k]$ and $x_j \neq x_i$ and $x_j \neq x_k$. \square

Lemma 2.3.9. *Mappings r and l restricted to ∂U preserve order.*

Proof. Suppose $u_1, u_2, u_3 \in \partial U$ are in that counterclockwise order and $r_1 = r(u_1)$, $r_2 = r(u_2)$, and $r_3 = r(u_3)$ are distinct. Let \bar{r}_i be the closest shadow of $g(u_i, r_i)$.

We show that $\bar{r}_1, \bar{r}_2, \bar{r}_3$ lie in that counterclockwise order on ∂U ; this follows from the claim that $\partial U[u_j, \bar{r}_j] \not\subseteq \partial U[u_i, \bar{r}_i]$ for all $i \neq j$. To establish this claim, suppose for the sake of contradiction that $\partial U[u_j, \bar{r}_j] \subseteq \partial U[u_i, \bar{r}_i]$ for some $i \neq j$. Then $r_j \in U[u_i, r_i]$ since $u_j, \bar{r}_j \in U[u_i, r_i]$, but $r_j \in F \subseteq U[r_i, u_i]$ by Lemma 2.3.2(1), so $r_j \in g(u_i, \bar{r}_i)$. In fact, $r_j \in g(u_i, r_i)$ since r_j cannot lie past r_i on $g(u_i, \bar{r}_i)$. If $r_j \in \partial U$, then $r_j \in g(u_i, r_i)$ and $r_i \neq r_j$ imply that r_j is a reflex corner of ∂U ; using $u_j \in U[u_i, r_i]$, we have that $g(u_j, r_j)$ could be extended to $g(u_j, r_i)$, a contradiction. If $r_j \notin \partial U$, then neither $g(u_i, \bar{r}_i)$ nor $g(u_j, \bar{r}_j)$ can bend at r_j . Since $u_j, r_j \in U[u_i, r_i]$, the two geodesics must overlap in some link containing r_j . Furthermore, the link must be traversed in the same direction in both geodesics, since $u_j, u_i, \bar{r}_i, \bar{r}_j$ appear in that counterclockwise order on ∂U . However, then $g(u_j, r_j)$ could be extended to $g(u_j, r_i)$, a contradiction.

We claim that $r_i \bar{r}_i$ and $r_j \bar{r}_j$ are disjoint whenever $i \neq j$; the lemma then follows immediately by the Connection Lemma. Suppose $r_i \bar{r}_i$ intersects $r_j \bar{r}_j$. The intersection must be a single point distinct from both r_i and r_j . Hence, either $r_j \notin U[r_i, u_i]$ or $r_i \notin U[r_j, u_j]$, contradicting Lemma 2.3.2(1). \square

To summarize, we have described the general form of a polygonal relatively convex region and its boundary. This yields a natural notion of ordering on the boundary of the set, consistent with the ordering around ∂U .

2.4. Far Sides

We consider the notion of the “far side” of a set F from a point. Informally, an extreme point $f \in F$ is on the far side of F from x if the geodesic $\bar{g}(x, f)$ leaves $R(F)$ after f . There are two main reasons for studying far sides. First, we use this notion to relate the order of extreme points of F with the order of points on ∂U (Corollary 2.4.2). Second, we are able to compute relative convex hulls efficiently using far sides (Lemma 2.4.3).

Let F be a finite set of points in U and $x \in \partial U$. The *far side* of F from x is the list of all extreme points of F counterclockwise from $r(x)$ to $l(x)$, inclusive. Here $r(x)$ and $l(x)$ are the clockwise and counterclockwise extreme points of F from x , as before.

Lemma 2.4.1. *Suppose f is an extreme point of F , $x \in \partial U$, and $f \neq x$. Then the following are equivalent:*

- (1) f is on the far side of F from x .
- (2) f is an extreme point of $F \cup \{x\}$.
- (3) $\bar{f}\bar{f}$ connects $R(F)$ to ∂U , for any shadow \bar{f} of $g(x, f)$, and there is no geodesic extending $g(x, f)$ to an extreme point f' of F distinct from f .

Proof. It is possible that $x \in R(F)$ if $x \in \partial U$ is a boundary point of $R(F)$; then all three conditions are satisfied. So assume $x \notin R(F)$. Then x is an exterior point of F . Set $r = r(x)$ and $l = l(x)$. By Lemma 2.3.2(1) $f \in U[r, x] \cap U[x, l]$.

(1 \rightarrow 2) If f is on the far side of F from x , then $f \in U[r, l]$ by Lemma 2.3.6. If $f \neq r$ and $f \neq l$, then, by the Triangle Lemma, $g(x, f)$ intersects $g(l, r)$ at some point x' , where $x' \neq f$ since f is extreme. Hence

$$\theta(f, x) = \theta(f, x') \in \theta(f, g(r, l)) \subseteq \text{span}(f, F),$$

and $\text{span}(f, F \cup \{x\}) = \text{span}(f, F)$, so f is extreme in $F \cup \{x\}$. If $f = r$, then, by the definition of clockwise extreme point, $r(x)$ is unchanged if F is replaced with $F \cup \{x\}$, so r is extreme by Lemma 2.3.2(3). The case $f = l$ is similar.

(2 \rightarrow 3) Trivial.

(3 \rightarrow 1) We argue the contrapositive. If f is not on the far side of F , then, by Lemma 2.3.6, $f \in U[l, r]$. Since f cannot be on $\bar{g}(l, r)$, $f \in \Delta rxl$. If $f \in g(x, r)$ or $f \in g(x, l)$, then $g(x, f)$ can be extended to an extreme point of F distinct from f . Otherwise by the Triangle Lemma, for some shadow \bar{f} of $g(x, f)$, $g(x, \bar{f})$ intersects $g(r, l)$ at a point distinct from f , thus $\bar{f}\bar{f}$ cannot connect $R(F)$ to ∂U . □

Corollary 2.4.2. *Suppose $\{f_i\}$ is a set of distinct points on the far side of F from x and, for each i , \bar{f}_i is a shadow of $g(x, f_i)$. Then the order of \bar{f}_i on ∂U is the same as the order of f_i on $\partial R(F)$.*

Proof. If $i \neq j$, then $g(x, f_i)$ and $g(x, f_j)$ must diverge before reaching f_i or f_j . Hence $f_i\bar{f}_i$ does not intersect $f_j\bar{f}_j$. The result then follows from the Connection Lemma. □




Fig. 6. Putting convex hulls together: Dashes outline $R(F \cup \partial U[u, v])$.

Lemma 2.4.3. *Suppose $u, v \in \partial U$, $u \neq v$, $F \subseteq U[v, u]$, $F \not\subseteq g(u, v)$, and $H = R(F)$. Then a counterclockwise traversal of the boundary of $R(F \cup \partial U[u, v])$ is $\partial U[u, v]$, $g(v, r(v))$, $\partial H[r(v), l(u)]$, $g(l(u), u)$.*

Proof. In this proof we write $r_M(x)$ and $l_M(x)$ for the clockwise and counterclockwise extreme points of M from x . This indicates the dependence upon M explicitly. Thus in the statement of the lemma, $r(x)$ really refers to $r_F(x)$. We also assume $\partial U[u, v] \not\subseteq R(F)$, hence some point of $\partial U[u, v]$ is extreme in $R(F \cup \partial U[u, v])$. If $\partial U[u, v] \subseteq R(F)$, a similar and easier argument suffices.

Refer to Fig. 6. Let G be $\{u, v\}$ together with the set of convex corners of $\partial U[u, v]$. Clearly, $R(F \cup \partial U[u, v]) = R(F \cup G)$. We first show that $r_{F \cup G}(v) = r_F(v)$; we use Lemma 2.3.2(5). Some point of F is in the relative interior of $U[v, u]$ by assumption; it is easy to check that in fact $r_F(v)$ must be in the relative interior of $U[v, u]$. Hence $F \cup \partial U[u, v] \subseteq U[r_F(v), v]$. Any geodesic extending $g(v, r_F(v))$ must stay in the relative interior of $U[v, u]$, hence must avoid $\partial U[u, v]$, and must also avoid any $f' \in F$ distinct from $r_F(v)$ since $r_F(v)$ is extreme in F . Hence, by Lemma 2.3.2(5), $r_{F \cup G}(v) = r_F(v)$. By Lemma 2.3.2(3), $r_F(v)$ is extreme in $F \cup G$. By a similar argument, $l_{F \cup G}(u) = l_F(u)$ is extreme in $F \cup G$.

We claim $h = l_{F \cup G}(r_F(v))$ is the most counterclockwise point of G (ordered along $\partial U[u, v]$) that is extreme in $F \cup G$: h is either v , the convex corner immediately clockwise of v , or possibly u if $\partial U[u, v]$ contains no convex corners. We have $h \in U[r_F(v), v]$ and $v \in U[r_F(v), h]$. This is only possible if v, h , and $r_F(v)$ lie on a common geodesic. Since h and $r_F(v)$ are extreme in $F \cup G$, they must be the endpoints of the geodesic. Hence either $h = v$, as desired, or $v \in g(r_F(v), h)$. In the latter case, by definition of counterclockwise extreme point, the shadow \bar{h} of $g(r_F(v), h)$ has to be as counterclockwise as possible; using $h \in U[v, r_F(v)]$ and the assumption that some point in $\partial U[u, v]$ is extreme in $F \cup G$, it follows that $h = \bar{h}$ and h is either u or the convex corner of G immediately clockwise of v .

Set $f_1 = r_F(v)$, $f_2 = r_F(f_1)$, ..., $f_h = l_F(u) = r_F(f_{h-1})$. Let $G' = \{g_1, \dots, g_t\}$ be the subset of $G - F$ extreme in $F \cup G$, where the index is given by the ordering along $\partial U[u, v]$. We want to show that $r_{F \cup G}$ maps $f_1 \rightarrow f_2 \rightarrow \dots \rightarrow f_h \rightarrow g_1 \rightarrow \dots \rightarrow g_t \rightarrow f_1$.

We have already shown that $l_{F \cup G}$ maps f_1 to g_l , hence $r_{F \cup G}(g_l) = f_1$. By a similar argument $r_{F \cup G}(f_h) = g_1$. It is easy to see that $r_{F \cup G}(g_i)$ can only be g_{i+1} , for $1 < i < l$.

It remains to show that $r_{F \cup G}(f_i) = f_{i+1}$ for $1 < i < h$. This follows if we establish $\text{span}(f_i, F \cup G) = \text{span}(f_i, F)$ (which also establishes f_i extreme in $F \cup G$). Since $f_h = l_F(u)$ is extreme in $F \cup G$, f_h is extreme in $F \cup \{v\}$, hence f_h is on the far side of F from v , and hence f_1, \dots, f_h are on the far side of F from v . Similarly, f_1, \dots, f_h are on the far side of F from u . Since r and l preserve order on ∂U , the far side of F from any $w \in \partial U[u, v]$ includes f_1, \dots, f_h . In particular, for any $w \in \partial U[u, v]$ and any i , $1 < i < h$, we have $f_i \notin g(l(w), r(w))$ since f_i is extreme in F and $f_i \neq l(w), r(w)$. Moreover, as in the proof of Lemma 2.4.1,

$$\theta(f_i, w) \in \theta(f_i, g(l(w), r(w))) \subseteq \theta(f_i, R(F)) = \text{span}(f_i, F).$$

Hence $\text{span}(f_i, F \cup G) = \text{span}(f_i, F)$.

Now if $g_1 = u$ and $g_l = v$, we are done. If, say, $g_1 \neq u$, then u must appear on $g(f_k, g_1)$. To obtain the lemma split $g(f_k, g_1) = g(l_F(u), g_1)$ into $g(l_F(u), u)$ and $g(u, g_1)$, then merge $g(u, g_1)$ with $\partial U[g_1, g_k]$. A similar split applies if $g_l \neq v$. \square

2.5. The General Position Assumption

For the remainder of the paper we make the following general position assumption: no corner of ∂U is equidistant from two sites. This condition can always be satisfied by applying a slight perturbation to the positions of the sites or corners. If this assumption is not made, then it is possible for the set of points equidistant from two sites to include a two-dimensional region (see Fig. 7), introducing considerable complexity to the definition of bisectors and Voronoi cells, which we wish to avoid. This section contains some consequences of the general position assumption that are critical to the rest of the paper. We emphasize that none of Lemmas 2.5.1, 2.5.2, and 2.5.3 holds if the assumption is removed.

Lemma 2.5.1. *If s and t are distinct sites, $x \in U$, and $d_s(x) = d_t(x)$, then $\theta(x, s) \neq \theta(x, t)$.*

Proof. Suppose $\theta(x, s) = \theta(x, t)$. Then the geodesics $g(s, x)$ and $g(t, x)$ share their final link yx . Point y must be the anchor of x with respect to both s and t . Now




Fig. 7. A configuration not in general position. The entire shaded region is contained in $b(s, t)$ (see Section 2.6).

$y \neq s, t$ since otherwise $d(s, y) = d(t, y)$ would imply $s = t$. Hence y must be a reflex corner of ∂U equidistant from s and t , contradicting the general position assumption. \square

Lemma 2.5.2. *Suppose $u, v \in U, u \neq v$, and each of u and v is equidistant from distinct sites s and t . Then $g(s, u)$ does not intersect $g(t, v)$.*

Proof. Suppose x lies on both $g(s, u)$ and $g(t, v)$. Without loss of generality, assume $d(x, s) \leq d(x, t)$. Observe that

$$d_v(t) = d_v(s) \leq d_v(x) + d(x, s) \leq d_v(x) + d(x, t) = d_v(t).$$

Hence $d(x, s) = d(x, t)$, path $g(s, x) \cup g(x, v)$ is in fact geodesic $g(s, v)$, and similarly $g(t, x) \cup g(x, u)$ is $g(t, u)$. Since x lies on both geodesics $g(s, u)$ and $g(s, v)$, x must be a reflex corner of ∂U or an interior point of a common link of the two geodesics connecting two reflex corners of ∂U . In the second case the link must be common to $g(t, u)$ and $g(t, v)$ as well, hence in either case we can find a reflex corner of ∂U equidistant from s and t , violating the general position assumption. \square

Lemma 2.5.3. *There is at most one point equidistant from three distinct sites.*

Proof. Suppose to the contrary that points u and v are both equidistant from sites r, s, t . First note that r, s, t cannot lie on a common geodesic, for if say $r \in g(s, t)$, then, by Lemma 2.2.1, $d_u(r) < \max(d_u(s), d_u(t))$. Hence r, s, t are extreme points of $\{r, s, t\}$.

We claim that r, s, t are extreme points of $\{r, s, t, u\}$ (and also of $\{r, s, t, v\}$). In order to demonstrate this, we show $r \notin R(\{s, t, u\})$. Now r does not lie on $g(u, s)$ or $g(u, t)$, else s or t would be further from u than r . As argued before, r does not lie on $g(s, t)$. If r is in the interior of $R(\{s, t, u\})$, then, by the Triangle Lemma, $g(u, \bar{r})$ intersects $g(s, t)$, where \bar{r} is a shadow of $g(u, r)$. However, then using Lemma 2.2.1 again, we would have $d_u(r) < \max(d_u(s), d_u(t))$.

Now suppose one of u and v , say u , is extreme in $R(\{r, s, t, u, v\})$. Hence r, s, t, u are extreme in $R(\{r, s, t, u\})$; assume that they appear in that counterclockwise order. By Lemma 2.3.6, $r \in U[u, s], t \in U[s, u]$, and $r, t \notin \bar{g}(u, s)$. It must be that $\bar{g}(u, s) = g(u^*, \bar{s})$ intersects either $g(v, r)$ or $g(v, t)$; assume it is $g(v, t)$. To obtain a contradiction of Lemma 2.5.2, we show $g(u, s)$ in fact intersects $g(v, t)$. Now uu^* intersects $R(\{r, s, t, u, v\})$ only at u since u is extreme in $R(\{r, s, t, u, v\})$. Hence uu^* is disjoint from $R(\{r, s, t, v\})$. Also $s\bar{s}$ intersects $R(\{r, s, t, v\})$ only at s , since s is extreme in $R(\{r, s, t, v\})$ and some portion of $g(u, s)$ must lie in $R(\{r, s, t\})$, hence in $R(\{r, s, t, v\})$. Since $g(v, t) \subseteq R(\{r, s, t, v\})$, $g(v, t)$ must intersect $g(u, s)$.

If neither u nor v is extremal in $R(\{r, s, t, u, v\})$, then both $u, v \in R(\{r, s, t\})$ and the proof is similar. Geodesics $g(u, r), g(u, s)$, and $g(u, t)$ split $R(\{r, s, t\})$ into three geodesic triangles. Hence v lies in one of the triangles which again implies a contradiction of Lemma 2.5.2. \square




Fig. 8. Bisectors; dots indicate breakpoints.

2.6. Voronoi Cells

The bisector $b(s, t)$ of distinct sites s and t is $\{x \in U : d_s(x) = d_t(x)\}$ and the half-space closer to s , $H(s, t)$, is $\{x \in U : d_s(x) < d_t(x)\}$. Clearly, $H(s, t)$, $b(s, t)$, and $H(t, s)$ form a partition of U . A breakpoint of $b(s, t)$ is the intersection of $b(s, t)$ with a shortest-path partition edge from s or t . Figure 8 indicates the bisectors of three points, with breakpoints marked.

The following two lemmas depend upon the general position assumption.

Lemma 2.6.1 [A, 3.22]. *Bisector $b(s, t)$ is a smooth curve connecting two points on ∂U and having no other points in common with ∂U . It is the concatenation of $O(n)$ straight and hyperbolic arcs; the points where the arcs meet are precisely the breakpoints of $b(s, t)$. The tangent to $b(s, t)$ at x bisects the angle between $\theta(x, s)$ and $\theta(x, t)$.*

In particular, together with Lemma 2.5.1 this implies that, given $x \in b(s, t)$, $\theta(x, s)$ and $-\theta(x, t)$ enter $H(s, t)$ at x , while $-\theta(x, s)$ and $\theta(x, t)$ enter $H(t, s)$ (if they stay within U , that is).

Corollary 2.6.2. *The relative boundary of $H(s, t)$ and $H(t, s)$ is $b(s, t)$.*

Lemma 2.6.3 [A, 3.17]. *$H(s, t)$ is connected.*

Recall that S is the set of sites. The (geodesic furthest-site) Voronoi cell of site s is $V(s) = \bigcap_{t \neq s} H(t, s)$. The (geodesic furthest-site) Voronoi diagram V is

$$\{x \in b(s, t) : s, t \in S \text{ and } d_s(x) = \max_{r \in S} d_r(x)\}.$$

Figure 9 indicates the Voronoi diagram and Voronoi cells of the three points depicted in Fig. 8. A Voronoi edge $e(s, t)$ is $V \cap b(s, t)$ if the intersection consists of more than one point; else we say that $e(s, t)$ does not exist. A (Voronoi) vertex is a point $x \in V$ which has three or more sites furthest from it. By Lemma 2.5.3 above, there is at most one such point x for each triple of sites. A hitpoint is the intersection of a Voronoi edge with ∂U . Intuitively, a hitpoint corresponds to the “point at infinity” of an infinite Voronoi edge in a Euclidean furthest-site Voronoi diagram.




Fig. 9. Voronoi cells of the sites of Fig. 8.

Lemma 2.6.4. *Voronoi edge $e(s, t)$ is connected and has vertices or hitpoints as endpoints.*

Proof. Suppose r, s, t are distinct sites. Since d_r is continuous, a connected component of $b(s, t) \cap (b(r, s) \cup H(r, s))$ must have for each of its endpoints either an endpoint of $b(s, t)$, i.e., a hitpoint, or a point equidistant from s, t, r . However, by Lemma 2.5.3, there is at most one point equidistant from s, t, r , so

$$b(s, t) \cap (b(r, s) \cup H(r, s))$$

consists of a single connected component. Hence also

$$e(s, t) = b(s, t) \cap \bigcap_{r \neq s, t} (b(r, s) \cup H(r, s))$$

is connected and has hitpoints or vertices for endpoints. □

Lemma 2.6.5. *Suppose $s_1, \dots, s_k, s_{k+1} = s_1$ are the sites furthest (and thus equidistant) from vertex v , and directions $\theta(v, s_1), \dots, \theta(v, s_k)$ are in counterclockwise order. Then, for each i , edge $e(s_i, s_{i+1})$ is incident to v and extends away from v in direction bisecting $\angle s_{i+1}vs_i$, as long as that direction (locally) stays inside U .*

Proof. Elementary analysis, using $\nabla d_s(v) = -\theta(v, s_i)$. □

If vertex v appears on ∂U , then there is only one edge of V incident to v : as v cannot be a corner of ∂U by the general position assumption, it must be an interior point of a wall. (Of course, this means that v is not a vertex, but a hitpoint.) Hence only the edge bisecting $\angle svt$ remains within U , where directions $\theta(v, s)$ and $\theta(v, t)$ are the most clockwise and most counterclockwise directions toward sites furthest from v , respectively.

Lemma 2.6.6 (Extension Lemma). *If x lies on $g(s, y)$ and $x \in V(s)$ or $x \in e(s, t)$, for some other site t , then all of $g(s, y)$ past x lies in $V(s)$.*

Proof. Suppose $x \in V(s)$. For any site $r \neq s$ and for any $z \in g(x, y)$,

$$d_s(z) = d_s(x) + d(x, z) > d_r(x) + d(x, z) \geq d_r(z),$$

so $z \in V(s)$. Suppose $x \in e(s, t)$ and $y \neq x$. By the general position assumption x is not a corner so $\theta(x, y) = -\theta(x, s)$ stays in U at x . We show $\theta(x, y)$ enters $V(s) = \bigcap_{r \neq s} H(r, s)$ at x ; the result follows as before. Since $x \in e(s, t)$, for any $r \neq s$, $d_r(x) \leq d_s(x)$. If $d_r(x) = d_s(x)$, $-\theta(x, s) = \theta(x, y)$ enters $H(r, s)$; if $d_r(x) < d_s(x)$, any direction locally stays in $H(r, s)$. \square

An immediate consequence of this lemma is that every point in a Voronoi cell is connected to ∂U : if $x \in V(s)$, then segment $x\bar{x} \subseteq V(s)$, where \bar{x} is a shadow of $g(s, x)$.

Lemma 2.6.7. *Both $V(s) \cap \partial U$ and $V(s)$ are connected.*

Proof. Since every point of $V(s)$ is connected to a point of $V(s) \cap \partial U$, it suffices to show $V(s) \cap \partial U$ is connected. To prove it, we show that $H(r, s) \cap H(t, s) \cap \partial U$ is connected for every $r, t \neq s$. Label the hitpoints of $b(r, s)$ and $b(t, s)$ as x_r, y_r and x_t, y_t , respectively, so that

$$(\partial U \cap H(r, s)) \cup \{x_r, y_r\} = \partial U[x_r, y_r]$$

and

$$(\partial U \cap H(t, s)) \cup \{x_t, y_t\} = \partial U[x_t, y_t].$$

The only counterclockwise ordering of these points that disconnects

$$\partial U \cap H(r, s) \cap H(t, s)$$

is x_r, y_t, x_t, y_r . In particular, this implies that $H(r, s) \cup H(t, s)$ cover ∂U . We show this is impossible; suppose this were the ordering. Since $s \notin H(r, s) \cup H(t, s)$, $H(r, s) \cup H(t, s) \neq U$, so $b(r, s)$ must intersect $b(t, s)$ in at least two distinct points. However, this contradicts Lemma 2.5.3, since each of the two points would be equidistant from sites r, s , and t . \square

2.7. The Ordering Lemma

Let C be the relative convex hull of S , the set of sites. The first lemma of this section shows that if $V(s)$ is not empty, then s must be an extreme point of S and, in fact, lie on the far side of S from any point in the closure of $V(s) \cap \partial U$. Hence we can assume that all sites are extreme and that they are ordered by the counterclockwise traversal of ∂C . The main result of this section is the Ordering Lemma: the order of Voronoi cells around ∂U is exactly the order of sites around ∂C . In addition we show that there is a collection of at most three geodesics separating every point of ∂U from its furthest site(s). This is used to prove the rather remarkable fact that there are at most $O(n + k)$ distinct links in all geodesics that connect corners of ∂U to their respective furthest sites. The separating

geodesics are used in Section 3 to partition the problem of computing $V \cup \partial U$ into three recursively decomposable subproblems. The ordering that allows us to use recursion is based on the Ordering Lemma. We remark that Lemmas 2.7.2, 2.7.5, and 2.7.6 are generalizations of very similar lemmas proved by Suri [S].

Lemma 2.7.1. *If $V(s)$ is not empty, then s is an extreme point of S on the far side of C from any x in the closure of $V(s) \cap \partial U$.*

Proof. Suppose s is a site furthest from x . If $s \in g(r, t)$, for r, t sites distinct from s , then, by Lemma 2.2.1, $d_s(x) < \max(d_r(x), d_t(x))$, a contradiction. Suppose the last segment of $g(x, s)$ can be extended beyond s staying in C ; let s' be a point lying on such an extension and on the boundary of C . Then $d_{s'}(x) > d_s(x)$, s' must lie on $g(t, u)$ for some sites t, u , and $d_{s'}(x) \leq \max(d_t(x), d_u(x))$, contradicting the choice of s . Hence s is not an interior point of C , so s must be an extreme point. By Lemma 2.4.1, s is on the far side of C from x . □

Lemma 2.7.2. *Suppose s and t are furthest sites from $u, v \in U$, respectively, with $u \neq v, s \neq t$. Then $g(u, t)$ does not meet $g(v, s)$.*

Proof. Suppose to the contrary that $x \in g(u, t) \cap g(v, s)$. By the triangle inequality we have $d(s, u) \leq d(s, x) + d(x, u)$ and $d(t, v) \leq d(t, x) + d(x, v)$. Adding, we obtain $d(s, u) + d(t, v) \leq d(t, u) + d(s, v)$. However, since s and t are furthest sites from u and v , respectively, $d(s, u) \geq d(t, u)$ and $d(t, v) \geq d(s, v)$. Thus we must have $d(t, v) = d(s, v)$ and $d(s, u) = d(t, u)$. However, in this situation $g(u, t)$ cannot intersect $g(v, s)$ by Lemma 2.5.2, a contradiction. □

Lemma 2.7.3. *Suppose $V(t) \cap \partial U$ immediately follows $V(s) \cap \partial U$ in counterclockwise order around ∂U . If $u \neq s, t$ is another site on $\partial C[s, t]$, then $V(u)$ is empty.*

Proof. Let x be the hitpoint of $e(s, t)$ so that (near x) $V(s), x, V(t)$ are in counterclockwise order along ∂U . Let \bar{s} and \bar{t} be closest shadows of $g(x, s)$ and $g(x, t)$, respectively, so $g(x, \bar{s}) = \bar{g}(x, s)$ and $g(x, \bar{t}) = \bar{g}(x, t)$. (For a schematic diagram, refer to Fig. 10.) We show that x, \bar{s}, \bar{t} are distinct and appear in that order counterclockwise around ∂U . Note that, by Lemma 2.5.1, $\bar{g}(x, s)$ and $\bar{g}(x, t)$ are geodesics emanating from x with distinct initial directions; hence $\bar{s} \neq \bar{t}$. Near x , $V(t) \cap \partial U = H(s, t) \cap \partial U$ and $V(s) \cap \partial U = H(t, s) \cap \partial U$; also $\theta(x, s)$ enters $H(s, t)$ and $\theta(x, t)$ enters $H(t, s)$. Since $\bar{g}(x, s)$ does not intersect $\bar{g}(x, t)$ again, the ordering of x, \bar{s}, \bar{t} must be counterclockwise around ∂U .

We claim $\bar{s}, \bar{u},$ and \bar{t} are in that counterclockwise order on ∂U , where \bar{u} is the closest shadow of $g(x, u)$. Let r and l be the clockwise and counterclockwise extreme points of S from x , with \bar{r} and \bar{l} the closest shadows of $g(x, r)$ and $g(x, l)$, respectively. By Lemma 2.3.2(2), $\partial U[\bar{r}, \bar{l}]$ does not contain x . By Lemma 2.7.1, s, t are on the far side of S from x . Hence by Corollary 2.4.2, $\bar{s}, \bar{t} \in \partial U[\bar{r}, \bar{l}]$, and in fact the counterclockwise order must be $\bar{r}, \bar{s}, \bar{t}, \bar{l}$, since $\partial U[\bar{s}, \bar{t}]$ does not contain x . By Corollary 2.4.2, r, s, t, l appear in that counterclockwise order on C ; since u is extreme and appears on ∂C between s and t , u is on the far side of S from x . Again by Corollary 2.4.2, $\bar{s}, \bar{u}, \bar{t}$ appear in that counterclockwise order on ∂U .




Fig. 10. Illustration to the proof of Lemma 2.7.3.

Now suppose contrary to the lemma $V(u)$ is not empty; we obtain a contradiction. Since Voronoi cells are connected to ∂U , there is $y \in V(u) \cap \partial U$. Now $y \neq x$ since Voronoi cells are relatively open, and x lies on the boundary of $V(s)$ (and $V(t)$).

As s, u , and t are distinct extreme elements of S , the set $\{s, u, t\}$ is extreme. As $x \in \partial U$, it is impossible for x to lie in the interior of $\triangle sut$. Moreover, s, u , and t are on the far side of $\{s, u, t\}$ from x (in that counterclockwise order). Thus either $x \in g(s, t)$ or $\{s, u, t, x\}$ is extreme (with s, u, t, x , occurring in that counterclockwise order). In either case $g(x, u)$ intersects $g(s, t)$ and enters $U[s, t]$ at some point u' , as u is an extreme point of C lying counterclockwise between s and t . Assume $y \in U[x, u]$, the case $y \in U[u, x]$ is similar. Lemma 2.7.2 implies that $g(y, t)$ does not intersect $g(x, u)$. As t lies in the relative interior of $U[u, x]$ and y lies in $U[x, u]$, $g(y, t)$ must intersect $u\bar{u}$ at some point $u'' \neq u$. This implies that the portion of $\bar{g}(x, u)$ from u' to u'' inclusive is contained in the triangle $\triangle yst$; in particular $u \in \triangle yst$. By Lemma 2.2.1, this contradicts the choice of $u \neq s, t$ as a site furthest from y . □

Corollary 2.7.4 (Ordering Lemma). *The ordering of sites with nonempty Voronoi cells around ∂C is the same as the ordering of Voronoi cells around ∂U .*

Lemma 2.7.5. *Suppose $\{u_1, \dots, u_m\}$ and $\{v_1, \dots, v_m\}$ are separated by boundary geodesic $g(x, y)$. Then there is a total of at most $O(m + n)$ distinct links in paths $g(u_i, v_i)$, $i = 1, \dots, m$.*

Proof. By a proof essentially identical to Lemma 4 of Suri [S], for each i there are at most three links of $g(u_i, v_i)$ not in $T(x) \cup T(y)$. (Three links are needed rather than one because u, v need not be corners of ∂U .) □

Lemma 2.7.6. *There are three boundary geodesics g_1, g_2, g_3 so that, for any point $u \in \partial U$ and any site s furthest from u , one of the geodesics separates u from s .*




Fig. 11. Illustration to the proof of Lemma 2.7.6.

Proof. Refer to Fig. 11. Fix $x \in S$ arbitrarily. Let y be a site furthest from x and let z be a site furthest from y . We argue the case that x, y, z are distinct and appear in that counterclockwise order around ∂C ; the case that $x = z$ is easier and the case that the order is x, z, y is similar.

We claim that we can choose a site w furthest from z in $\partial C[x, y]$. We first show that w can be chosen to lie in $\partial C[x, z]$: if all sites w furthest from z lie in $\partial C[z, x] - \{x\}$, then in particular $d(z, w) > d(z, y)$. We also have $d(x, y) \geq d(x, z)$ and $d(y, z) \geq d(y, w)$, so adding all three we get $d(x, y) + d(z, w) > d(x, z) + d(y, w)$. This contradicts the triangle inequality as can be seen by considering a point in the intersection of $g(y, w)$ and $g(x, z)$. Now w cannot be in $\partial C[y, z] - \{y\}$, else $g(x, w)$ intersects $g(y, z)$, a contradiction of Lemma 2.7.2, taking $u = x, t = w, v = z$, and $s = y$. Hence we can choose w in $\partial C[x, y]$.

Let x^*, y^*, z^* and $\bar{y}, \bar{z}, \bar{w}$ be the closest foreshadows and closest shadows of $g(x, y), g(y, z)$, and $g(z, w)$, respectively. We claim $y^* \in \partial U[x^*, \bar{y}]$: this is immediate if $y \in \partial U$ or $g(x, y)$ and $g(z, y)$ share a common final link; otherwise it follows from Lemma 2.2.4(1), since y^* is the closest shadow of $g(z, y)$. Similarly, $z^* \in \partial U[y^*, \bar{z}]$ and $x^* \in \partial U[z^*, \bar{w}]$.

Now, by the Extension Lemma, x^*, y^*, z^* have y, z, w as their respective furthest sites. By the Ordering Lemma, all points in $\partial U[x^*, y^*]$ have furthest sites in $\partial C[y, z]$; clearly, $\bar{g}(x, y)$ separates $\partial U[x^*, y^*]$ from $\partial C[y, z]$. Similarly, the sites furthest from $\partial U[y^*, z^*]$ lie in $\partial C[z, w]$ and these two sets are separated by $\bar{g}(y, z)$. Finally, the sites furthest from $\partial U[z^*, x^*]$ lie in $\partial C[w, y]$ and these two sets are separated by $\bar{g}(z, x)$. Note that in fact x^* can have furthest sites both in $\partial C[y, z]$ and $\partial C[w, y]$, and similarly for y^* and z^* , while each of the remaining points of ∂U has its furthest sites in exactly one of $\partial C[y, z], \partial C[z, w]$, or $\partial C[w, y]$. \square

2.8. The Refined Voronoi Diagram

The Voronoi diagram V clearly has $O(k)$ edges, since together with ∂U it forms a planar graph with at most k bounded regions and all Voronoi vertices have degree




Fig. 12. Refined Voronoi diagram of the sites of Fig. 8.

three or more. However, this is not an accurate description of the size complexity of the Voronoi diagram, since each Voronoi edge may consist of sections of several different hyperbolic arcs. This section discusses a refinement of the Voronoi partition, obtained by further subdividing each Voronoi cell $V(s)$ by the shortest-path partition from s . Each bounding edge of a refined Voronoi cell is a line segment or a section of a single hyperbolic arc. The main theorem is a linear bound on the size complexity of the refined Voronoi diagram. This implies an $O(n + k)$ bound on the size complexity of the Voronoi diagram itself.

The *refined Voronoi cell* of site s with anchor a , $V_a^*(s)$, is $V(s) \cap P_a(s)$. The *refined bisector edge* $e_{ab}^*(s, t)$ is $e(s, t) \cap P_a(s) \cap P_b(t)$. The *refined partition edge* (from s with anchor a), $p_a^*(s)$, is $V(s) \cap p_a(s)$. A *refined Voronoi edge* is a refined bisector edge or a refined partition edge. (Refer to Fig. 12.) Observe that distinct refined Voronoi edges are disjoint (except possibly at their endpoints). Empty refined Voronoi cells and edges and refined Voronoi edges consisting of a single point are disregarded.

Suppose $e_{ab}^*(s, t)$ is not empty. It is easy to see that each endpoint of $e_{ab}^*(s, t)$ is either a vertex of V , a hitpoint, or a breakpoint. Moreover, $e_{ab}^*(s, t)$ does not contain breakpoints (except possibly as endpoints). Consequently $e_{ab}^*(s, t)$ is a hyperbolic arc or a line segment.

Lemma 2.8.1. *Either $p_a^*(s)$ is empty, or it is all of $p_a(s)$, or it has an open endpoint at a breakpoint of $e(s, t)$ for some site t and closed endpoint on ∂U . (The latter two cases are illustrated in Fig. 12.)*

Proof. Suppose $p_a^*(s)$ is not empty and is not all of $p_a(s) = ay$ (where $y \in \partial U$). Then $p_a(s)$ must intersect some edge $e(s, t)$ before first entering $V(s)$. By the Extension Lemma, the intersection is a single point x , and $p_a^*(s) = xy - \{x\}$ is contained in $V(s)$. □

Suppose $p_a^*(s)$ is a refined partition edge of $V(s)$. Let $t(p_a^*(s))$ be the first link of the geodesic $g(a, s)$, directed toward s . Then by definition of the refined partition edge $p_a^*(s)$, $p_a(s)$ and $t(p_a^*(s))$ are collinear and meet at a . We claim that if $p_a^*(s)$ and $p_{a'}^*(s')$ are distinct refined partition edges, then $t(p_a^*(s))$ and $t(p_{a'}^*(s'))$ are distinct, at least as directed links. Clearly, $p_a^*(s) \neq p_{a'}^*(s')$ implies either $a \neq a'$ or $s \neq s'$. The claim is immediate if $a \neq a'$; otherwise suppose $s \neq s'$. Since $p_a^*(s) \subseteq V(s)$ and $p_{a'}^*(s') \subseteq V(s')$, we must have $\theta(a, s) \neq \theta(a, s')$. Since $t(p_a^*(s))$ and $t(p_{a'}^*(s'))$ have these two directions, they must be distinct.

Lemma 2.8.2. *Link $t(p_a^*(s))$ is a link of $g(v, s)$, where v can be chosen to be a corner lying in $\partial U \cap V(s)$ or a hitpoint of $V(s)$ (i.e., an endpoint of $\partial U \cap V(s)$).*

Proof. Let $p_a(s) = ay$, $y \in \partial U \cap V(s)$. Segment ay partitions U into two polygonal regions U_1 and U_2 so that U_1 contains s and every geodesic to s from a point in U_2 contains $t(p_a^*(s))$. If we traverse $\partial U \cap U_2$ starting at y , we must encounter either the endpoint of $\partial U \cap V(s)$ or a corner of ∂U lying in $V(s)$. \square

The *refined Voronoi diagram*, V^* , is the union of all refined Voronoi edges. A *vertex* of V^* is a vertex of V or a breakpoint. A *hitpoint* of V^* is a hitpoint of V or the point of intersection of a refined partition edge with ∂U (in the case when the refined partition edge consists of an entire shortest-path partition edge, only its nonanchor endpoint is considered a hitpoint).

Lemma 2.8.3. *There are at most $O(n + k)$ refined Voronoi edges and vertices.*

Proof. Clearly, V has at most k cells, hence k hitpoints and $O(k)$ edges. We show that there are only $O(n + k)$ refined partition edges. Since each refined partition edge contributes a single breakpoint, it follows that there are only $O(n + k)$ refined bisector edges. By planarity, there are only $O(n + k)$ vertices in V^* as well.

By Lemma 2.7.6, there are three boundary geodesics g_1, g_2, g_3 and a partition of ∂U into three fragments U_1, U_2, U_3 so that, for $i = 1, 2, 3$, g_i separates every point in U_i from its furthest site(s). (This is not strictly true for the endpoints of U_i , but the argument is similar.) Let W_i be the union of the sets of links of geodesics $g(v, w)$, where v is a corner of ∂U lying in U_i and w is the unique site furthest from v or where v is a hitpoint of V lying in U_i and w is one of the two sites whose Voronoi cell boundary contains v . Since there are at most k hitpoints and n corners in all of ∂U , there are certainly at most as many in each U_i . By Lemma 2.7.5, there are $O(n + k)$ links in W_i . Now modify W_i so that it contains for each link two oppositely directed links; this doubles its size. By Lemma 2.8.2, if p_a^* is a refined partition edge, then $t(p_a^*)$ is in W_i , for some $i = 1, 2, 3$. Since each $t(p_a^*)$ corresponds to a *unique* refined partition edge, there are at most $O(n + k)$ such edges. \square

2.9. Directing Edges of the Refined Voronoi Diagram

We let c be the center of C . Necessarily c lies on some Voronoi edge, since there must be (at least) two sites attaining the maximum distance from c .

Lemma 2.9.1. *Suppose s is a site furthest from $x \in U$. Then the angle between $\theta(x, s)$ and $\theta(x, c)$ is less than $\pi/2$.*

Proof. Suppose the angle between $\theta(x, s)$ and $\theta(x, c)$ is at least $\pi/2$. By Lemma 2.2.2, $d(c, s) > d(x, s)$. However, since s is a furthest site from x , $rad(x, C) = d(x, s)$. This contradicts the choice of c as the center of C , i.e., the point minimizing $rad(c, C)$. \square

We use this lemma to direct each refined Voronoi edge toward c ; if there is a Voronoi edge containing c in its interior, we split the edge at c . We show how to direct an edge $e(s, t)$ of the (unrefined) Voronoi diagram; this direction extends to each refined bisector edge. A similar argument directs each refined partition edge. Suppose $x \in e(s, t)$, $x \neq c$. Since $e(s, t)$ bisects $\angle sxt$ and both the angle between $\theta(x, c)$ and $\theta(x, s)$ and the angle between $\theta(x, c)$ and $\theta(x, t)$ are less than $\pi/2$, it is not possible that $\theta(x, c)$ is perpendicular to $e(s, t)$ at x . Hence we can direct $e(s, t)$ toward c locally at x . If $c \notin e(s, t)$, this direction extends globally to $e(s, t)$, since direction toward c is a continuous function away from the corners and can never be perpendicular to $e(s, t)$. If $c \in e(s, t)$, then $e(s, t)$ is split at c into two pieces, each of which is consistently directed toward c .

Lemma 2.9.2. *Suppose $v \notin \partial U$ is a vertex of V^* different from c . Then there are at least two edges of V^* entering v and exactly one edge of V^* leaving v .*

Proof. We consider the case that v is in fact a vertex of V ; the case that v is a breakpoint is similar. For simplicity assume there are exactly three sites equidistant from v . Label the sites r, s, t so that $\theta(v, r), \theta(v, s)$, and $\theta(v, t)$ are in counterclockwise order leaving v so that $m\angle rvt < \pi$. (This is possible by the previous lemma as each of the angles formed between $\theta(v, c)$ and $\theta(v, r), \theta(v, s)$, or $\theta(v, t)$ has measure less than $\pi/2$.) By Lemma 2.6.5, Voronoi edges $e(r, s), e(s, t)$, and $e(t, r)$ are incident to v and extend in directions bisecting angles $\angle svr, \angle tvs$, and $\angle rvt$, respectively. Hence edges $e(r, s)$ and $e(s, t)$ enter v and $e(t, r)$ leaves v . The general case of a vertex of arbitrarily high degree is handled analogously. Vertices $v \notin \partial U$ of degree two or less do not occur by definition of V^* . □

Suppose Voronoi edge $e(s, t)$ intersects ∂U at hitpoint x ; recall that, by the general position assumption, x is not a corner of ∂U . Then $e(s, t)$ is directed into the interior of U , because $\theta(x, c)$ makes an angle strictly less than $\pi/2$ with both $\theta(x, s)$ and $\theta(x, t)$, $e(s, t)$ bisects the angle between $\theta(x, s)$ and $\theta(x, t)$, and none of $\theta(x, s), \theta(x, t)$, and $\theta(x, c)$ can leave U at x . A similar argument shows that refined partition edges are directed away from ∂U at hitpoints.

Corollary 2.9.3. *The unrefined Voronoi diagram V forms a tree with root c and edges directed toward c .*

Proof. No cycles are possible because otherwise some Voronoi cell would be separated from ∂U . Only c has out-degree zero; every other vertex or hitpoint of V has out-degree 1, so V is a root-directed tree. □

The refined Voronoi diagram V^* consists of V together with refined partition edges. Each such edge must lie entirely within a single Voronoi cell, having one endpoint on ∂U and the other endpoint at a reflex corner or on V (by Lemma 2.8.1). However, V^* will not in general be connected, see for example Fig. 12.

Procedure *gfv*

Input: A polygon U with n sides and a set $S \subseteq U$ of k sites.

Output: The refined furthest-site geodesic Voronoi diagram V^* .

1. Triangulate U .
2. Compute C , the relative convex hull of S , and discard all nonextreme sites of S .
3. Determine two or three two-fragment instances so that the union of the source fragments is ∂U .
4. Compute $V^* \cap \partial U$ by calling *rgfs*(u, v, s, t) for each two-fragment instance (u, v, s, t).
5. Call *sweep* to extend V^* to the interior of U .

Fig. 13. The algorithm.

3. The Algorithm

This section describes the algorithm for computing V^* . An outline of the algorithm is given in Fig. 13. It suffices to triangulate the polygon in time $O(n \log n)$ [GJPT]. The relative convex hull computation of the second step can be accomplished in time $O((n+k) \log(n+k))$ [T]. The third step also takes time $O((n+k) \log(n+k))$ and is described in Lemma 3.1.1 below. The fourth step, the most difficult of the algorithm, is the computation of $V^* \cap \partial U$, i.e., V^* restricted to ∂U . It is discussed in Sections 3.1–3.4. The last step is the extension of V^* to the interior of U . This is done using a “reverse geodesic sweeping” algorithm, discussed in Section 3.5 below. Both the fourth and fifth steps take time $O((n+k) \log(n+k))$.

The computation of $V^* \cap \partial U$ is quite similar in outline to Suri’s algorithm for furthest geodesic neighbors inside a simple polygon [S]. We first reduce the computation of $V^* \cap \partial U$ to at most three instances of the “two-fragment problem.” Roughly, an instance of the two-fragment problem consists of a fragment of ∂U and a fragment of ∂C so that all furthest sites of points in the fragment of ∂U are contained in the fragment of ∂C . Such a pair must also satisfy a technical condition given below; this reduction appears in Section 3.1. The algorithm to solve the two-fragment problem is based on a divide-and-conquer schema that splits an instance into two smaller instances. The basic properties of the divide-and-conquer schema appear in Section 3.2. Section 3.3 contains the exact splitting method and the procedures for handling the base cases of the recursion. The complexity analysis appears in Section 3.4. We show that the sum of all instance sizes at each level of recursion is linear in $n+k$. This implies a total running time of $O((n+k) \log(n+k))$.

We work with polygonal relatively convex sets in addition to simple polygons. Any such relatively convex set Q can be decomposed into a collection of plateaus and bridges. Clearly, a triangulation of Q can be obtained just by triangulating each plateau in the decomposition. This is easily done in time $O(m \log m)$ [GJPT] if m is the number of segments in the boundary of Q . Similarly, a shortest-path partition of Q from an arbitrary point in it can be obtained by using a shortest-path-partition algorithm in each plateau of the decomposition. If Q has been triangulated, this takes time $O(m)$ [GHL⁺].

3.1. The Two-Fragment Problem

A two-fragment instance is a quadruple (u, v, s, t) where $u, v \in \partial U$, $u \neq v$, s is a site furthest from u , t is a site furthest from v (possibly $s = t$), and $g(u, v)$ separates $\partial U[u, v]$ from $\partial C[s, t]$. The two-fragment problem is “Given a two-fragment instance (u, v, s, t) , compute $V^* \cap \partial U[u, v]$.” Observe that, by the Ordering Lemma, only the Voronoi cells of sites in $\partial C[s, t]$ can intersect $\partial U[u, v]$. The source fragment of the two-fragment instance (u, v, s, t) is $\partial U[u, v]$; the target fragment is $\partial C[s, t]$.

Lemma 3.1.1. *There exists a set of at most three instances of the two-fragment problem so that the union of the source fragments is ∂U . The instances each have size $O(n + k)$ and can be computed in time $O((n + k) \log(n + k))$ given a triangulation of U .*

Proof. Choose $x, y, z, w, x^*, y^*, z^*$ as in the proof of Lemma 2.7.6. It is clear that (x^*, y^*, y, z) , (y^*, z^*, z, w) , and (z^*, x^*, w, y) are two-fragment instances each of size at most $O(n + k)$; their source fragments cover ∂U . As for computing them, the choice of x was arbitrary. Site y can be determined in time $O((n + k) \log(n + k))$ by computing the shortest-path tree from x , then determining the distance from every site to x using a planar point-location algorithm in the resulting shortest-path partition. Sites z and w can be determined similarly. The points x^*, y^*, z^* can certainly be computed in additional time $O(n)$. The case when $x = z$ is handled similarly. □

In the following pages we perform a rather detailed analysis of the “anatomy” of an instance of the two-fragment problem. The next few definitions set the ground for this analysis. Let D be the relative convex hull of the sites on $\partial C[s, t]$. Clearly, the ordering of sites on D is the same as on C , with the addition that s immediately follows t in counterclockwise order. Let $R(u, v, s, t)$ be the relative convex hull of $\partial U[u, v]$ and D . See Fig. 14.




Fig. 14. Convex hull of a two-fragment instance. Dashes outline $R(u, v, s, t)$.

We say (u, v, s, t) is *degenerate* if D is contained in $g(u, v)$. If (u, v, s, t) is degenerate, there can be no sites in D besides s and t and, by Lemma 2.7.2, the order along $g(u, v)$ must be u, t, s, v . In this case we define $l = t$ and $r = s$.

Suppose (u, v, s, t) is not degenerate. Let l be the counterclockwise extreme point of D from u and let r be the clockwise extreme point of D from v . By Lemma 2.4.3, $\partial U[u, v]$, $g(v, r)$, $\partial D[r, l]$, and $g(l, u)$ constitute a counterclockwise traversal of the boundary of $R(u, v, s, t)$. It is an immediate consequence of the following lemma that $\partial D[r, l]$ is a subpath of $\partial C[s, t]$.

Lemma 3.1.2. *Sites s, r, l, t are in that counterclockwise order on D , not necessarily all distinct.*

Proof. The lemma is trivial if $r = l, s = t$, or if (u, v, s, t) is degenerate, so assume $s \neq t$ and (u, v, s, t) is not degenerate. We show that if s appears on $\partial D[r, l]$, then $r = s$ and if t appears on $\partial D[r, l]$, then $t = l$. Since s is the extreme point of D immediately counterclockwise from t , this implies the lemma.

Suppose t appears on $\partial D[r, l]$ and $t \neq l$. Since s is immediately counterclockwise from t , s is also an extreme point of D lying on $\partial D[r, l]$ and thus an extreme point of $R(u, v, s, t)$. Then t, s, u , and v appear in that order on a counterclockwise traversal of $\partial R(u, v, s, t)$, which by Lemma 2.3.7 implies that $g(t, u)$ meets $g(v, s)$, contrary to Lemma 2.7.2. \square

We wish to give a definition of “left side connector” and “right side connector” to capture the bounding edges of $R(u, v, s, t)$ not in $\partial C[r, l]$ and $\partial U[u, v]$. The obvious definitions are $g(u, l)$ and $g(v, r)$, respectively. Unfortunately, these definitions are not adequate. In Section 3.4 we analyze the size of side connectors; one crucial property used in our argument is that if $s \neq t$, then the left and right side connectors are disjoint except possibly at their endpoints (Lemma 3.1.3 below). Unfortunately, this is not true for side connectors defined as $g(u, l)$ and $g(v, r)$. See Fig. 14(b).

It is clear that geodesic $g(v, r)$ has a connected intersection with ∂D ; furthermore, if the intersection is more than a point it must be some final portion of geodesic $g(r', r)$, r' the site of D immediately clockwise of r . Let \hat{r} be r if $r = s$, otherwise let \hat{r} be the first point of $g(r', r)$ intersected by $g(v, r)$. The *right side connector* of (u, v, s, t) is $g(v, \hat{r})$. Similarly, we define the *left side connector* of (u, v, s, t) to be $g(u, \hat{l})$ where \hat{l} is l if $l = t$, otherwise \hat{l} is the first point of $g(l', l)$ intersected by $g(u, l)$, where l' is the site of D immediately counterclockwise from l . See Fig. 14. A *connector edge* is a link in either the left or right connector.

The *size* of (u, v, s, t) , denoted $|(u, v, s, t)|$, is $|\partial C[s, t]| + |\partial U[u, v]|$ plus the sizes of the side connectors. Since $\partial D[\hat{r}, \hat{l}]$ is a subpath of $\partial C[s, t]$, it is clear that $|\partial R(u, v, s, t)| \leq |(u, v, s, t)|$.

Lemma 3.1.3. *If the side connectors of a two-fragment instance (u, v, s, t) meet at a point other than one of their endpoints, then $s = t$.*

Proof. If (u, v, s, t) is degenerate, then $u, l = t, r = s, v$ appear in that order along

$g(u, v)$, and the side connectors are disjoint unless $l = r$, which in turn forces $s = t$. So suppose (u, v, s, t) is not degenerate and point $a \neq u, v, \hat{r}, \hat{l}$ is common to both side connectors. Without loss of generality we may assume that a is a reflex corner of ∂U . Deleting a splits $R(u, v, s, t)$ into two components, whose closures R_1 and R_2 are relatively convex polygonal regions with the property that a geodesic connecting two points of $R(u, v, s, t) - \{a\}$ passes through a if and only if one of the points lies in $R_1 - \{a\}$ and the other in $R_2 - \{a\}$. Since a lies on the left connector and thus on $g(u, l)$, either $u \in R_1$ and $l \in R_2$ or vice versa (and similarly for v and r). Since $\partial U[u, v]$, $g(v, r)$, $\partial C[r, l]$, $g(l, u)$ constitute a traversal of $\partial R(u, v, s, t)$ and such a traversal cannot visit any point more than twice, $\partial U[u, v]$ and $\partial C[r, l]$ do not meet a . In particular, u and v lie in the same component. Similarly, r and l must lie in the same component.

Without loss of generality, assume that R_1 contains u and v and R_2 contains r and l . Now also $s \notin R_1$, else a would appear on $\partial D[s, r]$ and thus could not lie on a side connector. Hence $s \in R_2$, similarly $t \in R_2$. Thus s is furthest from u , t is furthest from v , but both $g(u, t)$ and $g(v, s)$ pass through a , implying $s = t$ by Lemma 2.7.2. □

3.2. The Recursion Scheme

Figure 15 contains a recursive procedure *rgfs* for solving the two-fragment problem. Section 3.3 discusses the base cases and the choice of the splitting point w , while the complexity analysis is contained in Section 3.4. We now discuss some basic data structures needed for the recursion.

At each level of recursion, we need to have available the boundary of $R(u, v, s, t)$ and a triangulation of its interior. For the topmost level, the boundary of $R(u, v, s, t)$ can be constructed using the relative convex hull algorithm of Toussaint [T]; this takes time $O((n + k) \log(n + k))$. Then it can be triangulated in additional time $O((n + k) \log(n + k))$ [GJPT].

For the recursive step, we need to compute $f(w)$ and the boundaries and triangulations of $R(u, w, s, f(w))$ and $R(w, v, f(w), t)$ in total time $O(|(u, v, s, t)|)$. To compute $f(w)$ it suffices to know the geodesic distance $d_w(r)$ for every site r in $\partial C[s, t]$; $d_w(r)$ can be determined in constant time if the cell of the shortest-path partition of $R(u, v, s, t)$ from w containing r is known. The cell containing r for all sites r in $\partial C[s, t]$ can be determined in total time $O(|(u, v, s, t)|)$ as follows. We assume the shortest-path partition of $R(u, v, s, t)$ from w has been computed and

```

Procedure rgfs( $u, v, s, t$ )
if  $\partial U[u, v]$  or  $\partial C[s, t]$  is a base case
then compute  $\partial U[u, v] \cap V^*$  directly
else choose  $w \in \partial U[u, v]$ 
    locate a site  $f(w)$  furthest from  $w$ 
    call rgfs( $u, w, s, f(w)$ ) and rgfs( $w, v, f(w), t$ )
end
    
```

Fig. 15. Recursive procedure *rgfs*.

refined to a triangulation (this takes only linear additional time). First locate the triangle containing s ; this clearly can be done in the allowed time bound. Then traverse $\partial C[s, t]$, in one step moving to the next vertex of $\partial C[s, t]$ or to the next intersection of the current edge of $\partial C[s, t]$ with the boundary $\partial\Delta$ of the current triangle Δ of the shortest-path partition from w . Notice that the intersection of ∂C with Δ has at most three connected components, since C is relatively convex. Hence the traversal of $\partial C[s, t]$ takes total time $O(|(u, v, s, t)|)$, since the charge for a step to a vertex of $\partial C[s, t]$ can be allotted to the vertex and the step to an intersection with Δ can be allotted to one of the at most three connected components of $\Delta \cap \partial C[s, t]$.

We compute the boundary and triangulation of $R(u, w, s, f(w))$ as follows; handling $R(w, v, f(w), t)$ is similar. If $(u, w, s, f(w))$ is degenerate, the boundary of $R(u, w, s, f(w))$ can be easily obtained from the shortest-path partition of $R(u, v, s, t)$ from w . Otherwise compute r' , the clockwise extreme point of $\partial C[s, f(w)]$ from w using the shortest-path partition from w . Similarly, l' , the counterclockwise extreme point of $\partial C[s, f(w)]$ from u , can be determined by computing the shortest-path partition of $R(u, v, s, t)$ from u . Now since l' is extreme in $R(u, w, s, f(w))$, $g(w, l')$ splits $R(u, w, s, f(w))$ into two pieces, one piece lying to the left and one to the right. (Possibly one or the other is just $g(w, l')$.) A triangulation of the piece lying to the right can be obtained by refining the shortest-path partition of $R(u, v, s, t)$ from w . Similarly a triangulation of the piece lying to the left can be obtained by refining the shortest-path partition of $R(u, v, s, t)$ from l' . Notice the links in $g(u, l')$, $g(w, l')$, and $g(w, r')$ are used as triangle edges in this triangulation. The left and right side connectors of $R(u, w, s, f(w))$ are easily determined from $g(u, l')$ and $g(w, r')$. This computation can clearly be done in time $O(|(u, v, s, t)|)$.

3.3. Choosing Splitting Points and the Base Cases

We now discuss the recursion for an instance (u, v, s, t) . There are two base cases, one easy and one hard, and two splitting cases, one easy and one hard. Except for the time required to sort “partition points,” described below, the processing time for instance (u, v, s, t) is $|(u, v, s, t)|$, not counting the time for recursive calls.

The easy base case corresponds to $s = t$; then all of $\partial U[u, v]$ lies in $V(s)$. We need to find the refined partition edges of V^* intersecting $\partial U[u, v]$; it is sufficient to compute the shortest-path partition of $R(u, v, s, t)$ from s , which can be done in time $O(|(u, v, s, t)|)$ given the triangulation of $R(u, v, s, t)$ [GHL⁺].

The easy splitting case is if $s \neq t$ and u and v do not lie on the same wall of ∂U ; then the splitting point w is chosen simply as a corner of $\partial U[u, v]$ so that $|\partial U[u, w]|$ is within one of $|\partial U[w, v]|$.

The harder splitting case is if $s \neq t$ but u and v are the endpoints of the same wall of ∂U . In this case there is no obvious splitting point w . We perform a “partition” step: segment uv is split into subsegments so that within each subsegment the shortest-path tree from any point on the subsegment to the sites on $\partial C[s, t]$ is combinatorially invariant. This partitioning is described below. We introduce the partition points as dummy vertices, and use them as splitting points

in the divide and conquer, as before. Notice that the partitioning is performed at most once on a path from the topmost instance to a leaf instance in the recursion tree. We eventually show that we introduce only $O(n + k)$ such points altogether.

The remaining problem is to handle a base case instance (u, v, s, t) , where $s \neq t$ and the shortest-path tree is combinatorially invariant on segment $uv = \partial U[u, v]$. Because of this invariance no refined partition edge of V^* intersects segment uv , in other words $uv \cap V^* = uv \cap V$ is the set of bisector hitpoints on uv . Notice that the partition induced on segment uv by V^* is exactly the partition induced by the upper envelope of the functions d_r , where r is a site in $\partial C[s, t]$. Again because of the combinatorial invariance of the shortest-path tree, each function d_r is “simple” on segment uv ; specifically $d_r(x)$ is of the form $c_1 + \sqrt{c_2(x)}$ where c_1 is constant and $c_2(x)$ depends quadratically upon the position of x on segment uv . (Observe that the purpose of partitioning the original wall was to ensure that c_1 and c_2 are fixed over the length of uv . Their values for each site $r \in \partial C[s, t]$ are defined by the identity of the anchor of $x \in uv$ with respect to r and the distance from r to this anchor; all of this information can be determined in linear time from the shortest-path tree of $R(u, v, s, t)$ from, say, u .) Thus in constant time it is possible to determine, for a fixed pair of sites r and r' , the unique point $x \in uv$, if any, for which $d_r(x) = d_{r'}(x)$. Now by the Ordering Lemma, Voronoi cells appear along segment uv in the same order as the corresponding sites appear along $\partial C[s, t]$. This implies that the upper envelope of the functions d_r on segment uv can be computed in time proportional to the number of sites (which is certainly $O(|(u, v, s, t)|)$). For example, an incremental algorithm is sufficient. Suppose that the partition of segment uv induced by an initial subsequence of the sites on $\partial C[s, t]$ has been computed. Then the partition induced by adding the next site in order can be determined in constant time plus time proportional to the number of cells deleted from the partition of segment uv computed so far.

We now describe the partition step, which uses a technique similar to that of [GHL⁺]. We determine partition points, which are points of segment uv intersected by some shortest-path partition edge from some z in $\partial C[s, t]$. Notice that for each z it suffices to consider the shortest-path partition from z within geodesic triangle Δuvz . We can do this as follows. Compute $T'(u)$, the shortest-path tree to all sites in $\partial C[s, t]$ from u . $T'(u)$ can be obtained from the shortest-path tree from u to $\partial R(u, v, s, t)$ by adding links of geodesics from u to sites of $\partial C[s, t]$ not appearing on $\partial R(u, v, s, t)$ (there can be at most one such link per site not already in $T'(u)$) and by deleting links in the resulting tree that appear only on geodesics to nonsite vertices of $\partial R(u, v, s, t)$. Tree $T'(u)$ can be computed in total time $O(|(u, v, s, t)|)$ using the technique described for computation of $f(w)$ in Section 3.2. Similarly compute $T'(v)$, the augmented shortest-path tree from v . For a site z in $\partial C[s, t]$, the geodesics $g(z, u)$ and $g(z, v)$ bounding geodesic triangle Δuvz are easily obtained from the shortest-path trees $T'(u)$ and $T'(v)$. The partition points resulting from Δuvz can be obtained by traversing $g(z, u)$ and $g(z, v)$ from z ; eventually the two geodesics diverge and the required partition points are obtained by extending links of the two geodesics until they intersect segment uv . For a single site z , the time required is proportional to the size of Δuvz ; by appropriately pruning the construction on a repeat visit to a vertex, all partition points for all

z can be found in time $O(|T'(u)| + |T'(v)|)$, which is clearly $O(|(u, v, s, t)|)$. The number of partition points is bounded by $|T'(u)| + |T'(v)|$. We give below a sharper bound on $|T'(u)| + |T'(v)|$, when we also account for the time required to sort partition points along segment uv .

3.4. Complexity Analysis

Lemma 3.4.1. *Let (u_i, v_i, s_i, t_i) be a topmost two-fragment instance. The total number of partition points created from subinstances of (u_i, v_i, s_i, t_i) is $O(n + k)$.*

Proof. Let a *partition instance* be an instance for which partitioning is performed. Each site of s can appear in at most two partition instances, since instances arising from the divide-and-conquer are disjoint except at endpoints and since partitioning is done at most once on a path from root to leaf in the recursion tree and is always applied to an instance that involves two or more sites. The number of partition points arising from a particular partition instance (u, v, s, t) is bounded by $|T'(u) + T'(v)|$, which is the number of distinct links in geodesics $g(u, z)$, $g(v, z)$ for sites z in $\partial C[s, t]$. Hence the total number of partition points arising from all partition instances is bounded by the number of distinct links in a set of geodesics connecting vertices in $\partial U[u_i, v_i]$ to sites in $\partial C[s_i, t_i]$, with at most four geodesics per site in the set. Now $g(u_i, v_i)$ separates $\partial U[u_i, v_i]$ from $\partial C[s_i, t_i]$, so by Lemma 2.7.5 the total number of links is $O(n + k)$. Hence the number of partition points is $O(n + k)$. \square

Lemma 3.4.2. *Let (u_i, v_i, s_i, t_i) be a topmost two-fragment instance. There are at most $O(n + k)$ distinct connector edges among all subinstances of (u_i, v_i, s_i, t_i) .*

Proof. The number of distinct subinstances is at most $O(n + k)$, since $O(n + k)$ bounds the number of vertices of U , sites in S , and partition points. Hence there are only $O(n + k)$ connectors. Since $g(u_i, v_i)$ separates $\partial U[u_i, v_i]$ from $\partial C[s_i, t_i]$, it separates the endpoints of each connector as well. Hence by Lemma 2.7.5, there are at most $O(n + k)$ distinct connector edges. \square

Lemma 3.4.3. *Let (u_i, v_i, s_i, t_i) be a topmost two-fragment instance. At each level of recursion, each connector edge appears in only a constant number of subinstances of (u_i, v_i, s_i, t_i) .*

Proof. We count the number of times edge ab can appear as a left connector edge directed from a to b . Notice a and b may be assumed to be reflex corners of ∂U , for otherwise ab is necessarily a first or last link on a side connector, and hence can appear in only one left connector (three if ab is the last link and the target fragment consists of the single site b in two of the instances). Let Q_a be the set of points x of U for which $g(x, b)$ passes through a . Similarly, define Q_b as the set of points that can reach a only through b . Clearly, Q_a and Q_b are nonempty and disjoint. Let A be $\partial U[u_i, v_i] \cap Q_a$ and let B be $\partial C[s_i, t_i] \cap Q_b$. It can be checked

that A is a single fragment of ∂U and that B is a single fragment of ∂C . Then, for $u \in \partial U[u_i, v_i]$ and $t \in \partial C[s_i, t_i]$, link ab appears in $g(u, t)$ exactly if $u \in A$ and $t \in B$.

We first claim that at each level of recursion there are at most two nonleaf instances of the two-fragment problem for which both the source fragment intersects A and the target fragment intersects B . To see this suppose (u_j, v_j, s_j, t_j) , $j = 1, 2, 3$, are instances at the same level, $\partial U[u_j, v_j] \cap A \neq \emptyset$, $\partial C[s_j, t_j] \cap B \neq \emptyset$, and u_1, u_2, u_3 (and hence s_1, s_2, s_3) are in that counterclockwise order. Since these instances are all at the same level of recursion, v_2 appears between u_2 and u_3 (possibly $v_2 = u_3$) and s_2 appears between t_1 and t_2 (possibly $s_2 = t_1$). However, then $v_2 \in A$, $s_2 \in B$, so the geodesic $g(v_2, s_2)$ contains link ab . Recalling that the relative interior of ab does not lie in Q_b and hence is disjoint from $\partial C[s_2, t_2] \subseteq B$, we deduce that the right connector of (u_2, v_2, s_2, t_2) contains the relative interior of ab . Since $g(u_2, t_2)$ also contains ab , the relative interior of ab lies in the left connector of this instance as well. Hence $s_2 = t_2$ by Lemma 3.1.3 and (u_2, v_2, s_2, t_2) is a leaf instance.

We now claim that at each level of recursion there are at most four instances of the two-fragment problem (leaf and nonleaf) with ab appearing as a left connector edge. Such an instance (u, v, s, t) must have $u \in A$ and $t \in B$. This is only possible if, for its parent instance, the source fragment intersects A and the target fragment intersects B . As just argued there are only two such parent instances. □

Theorem 3.4.4. $V^* \cap \partial U$ can be computed in time $O((n + k) \log(n + k))$.

Proof. Clearly, U can be triangulated in time $O(n \log n)$ [GJPT]. By Lemma 3.1.1, there are at most three two-fragment instances with the union of source fragments equal to ∂U that can be identified in time $O(n + k)$. We show that *rgfs* solves the two-fragment problem in total time $O((n + k) \log(n + k))$ for each (top-level) instance, proving the theorem.

Consider the work performed by *rgfs* for all instances at a particular level of recursion, ignoring recursive calls and the time required to partition walls as discussed in Section 3.3. It is linear in instance size which is the sum of the sizes of source and target fragments and the sizes of the side connectors. The total size of source and target fragments at a particular level of recursion is $O(n + k)$, because source and target fragments are partitioned disjointly except for endpoints, and there are only $O(n + k)$ possible endpoints. By Lemmas 3.4.2 and 3.4.3, the total size of all connectors at a particular level of recursion is also $O(n + k)$. Hence the total work at a particular level of recursion, summed over all instances at the level, is $O(n + k)$.

The total depth of recursion is $O(\log(n + k))$: at each step, except for partition steps, the size of a source fragment is split in half. At a partition step, the size of the source fragment increases to at most $O(n + k)$, and a partition step happens at most once on a path in the recursion tree from topmost instance to leaf instance.

The total work required for sorting partition points is $O((n + k) \log(n + k))$, since there are only $O(n + k)$ partition points. Hence the total work to solve a single top-level two-fragment instance is $O((n + k) \log(n + k))$. □

Input: Triangulated polygon U , refined Voronoi diagram V^* restricted to ∂U .
Output: V^* .
Data structures:
 L : a doubly linked circular list of refined Voronoi edges.
 Q : a priority queue of points of U , ordered by decreasing geodesic distance from c .

procedure *sweep*
 Compute c , the center of C .
 Compute the shortest-path partition of U from c .
 Initialize L to be $\partial U \cap V^*$.
 Initialize Q to contain all pairwise intersections of refined Voronoi edges immediately adjacent in L and all anchors of the refined partition edges appearing in L .
 while $Q \neq \emptyset$
 Extract from Q the point v of maximum geodesic distance from c .
 Delete from L all refined Voronoi edges with head v .
 Delete from Q any intersections and/or anchors involving just-deleted edges.
 if v is not an anchor **then**
 Insert into L the refined bisector edge e with tail v .
 Insert into Q any new intersections of e with adjacent refined Voronoi edges.
 end if
 end while
end *sweep*

Fig. 16. Procedure *sweep*.

3.5. Computing V^*

V^* is computed by the procedure *sweep* (Fig. 16), which is a “reverse geodesic sweeping algorithm”; it progresses from ∂U toward c , the center of C .

Theorem 3.5.1. *Procedure sweep computes V^* . It can be implemented to run in time $O((n + k) \log(n + k))$ and space $O(n + k)$.*

Proof. For positive real r , let $D_c(r)$ be the geodesic disk of radius r centered at c , i.e., the set of all points of U at geodesic distance at most r from c . We claim that the **while** loop maintains the invariant that L is exactly the refined Voronoi edges intersected by $\partial D_c(r)$, in order around $\partial D_c(r)$. This follows from Lemma 2.9.2, using standard sweepline arguments [BO]. Hence procedure *sweep* computes V^* .

List L can be implemented simply as a circular doubly linked list, so each list operation takes constant time. Q can be implemented as a heap, so that each operation takes time $O(\log(n + k))$. The geodesic center c of C can be computed in time $O((n + k) \log(n + k))$ as follows. Pollack *et al.* [PSR] show how to compute the center of (the set of vertices of) a simple polygon; their algorithm extends easily to the case of a polygonal simply connected set. Since geodesics restricted to lie inside a relatively convex set with respect to U are identical to geodesics inside U , it suffices to compute the center of $C = R(S)$ using this extended version of the algorithm of [PSR].

The shortest-path partition of U from c can be computed in time $O(n)$ since U is triangulated [GHL⁺]. Given the shortest-path partition from c , the geodesic distance from a point $x \in U$ to c can be computed in time $O(\log n)$, using a planar subdivision search algorithm to locate the shortest-path partition cell containing x such as that of [ST]. By Theorem 3.4.4, $\partial U \cap V^*$ can be computed in time $O((n+k) \log(n+k))$. Hence L and Q can be initialized in total time $O((n+k) \log(n+k))$.

Each iteration of the **while** loop uses a number of operations on Q and L proportional to the degree d_v of the current vertex v ; each iteration also uses one geodesic-distance computation for each item inserted in Q . Hence each iteration takes time $O(d_v \log(n+k))$. Using Lemma 2.8.3, the total running time of the entire algorithm is $O((n+k) \log(n+k))$. Clearly, the space usage is $O(n+k)$. \square

Acknowledgments

We wish to thank Micha Sharir for suggesting the idea of the proof for Lemma 2.5.3 and the referee for his careful reading of the manuscript.

References

- [A] B. Aronov, On the geodesic Voronoi diagram of point sites in a simple polygon, *Algorithmica* **4**(1), 1989, 109–140.
- [AT] T. Asano, G. T. Toussaint, Computing the geodesic center of a simple polygon, in *Perspectives in Computing: Discrete Algorithms and Complexity, Proceedings of Japan-US Joint Seminar*, D. S. Johnson, A. Nozaki, T. Nishizeki, H. Willis, eds, Academic Press, Boston, June 1986, pp. 65–79.
- [BO] J. L. Bentley, T. A. Ottmann, Algorithms for reporting and counting geometric intersections, *IEEE Transactions on Computers* **28** (1979), 643–647.
- [CG] B. Chazelle, L. Guibas, Visibility and intersection problems in plane geometry, *Discrete and Computational Geometry* **4** (1989), 551–581.
- [GHL⁺] L. Guibas, J. Hershberger, D. Leven, M. Sharir, R. E. Tarjan, Linear-time algorithms for visibility and shortest-path problems inside triangulated simple polygons, *Algorithmica* **2** (1987), 209–233.
- [GJPT] M. R. Garey, D. S. Johnson, F. P. Preparata, R. E. Tarjan, Triangulating a simple polygon, *Information Processing Letters* **7**(4) (1978), 8–21.
- [LP] D. T. Lee, F. P. Preparata, Euclidean shortest paths in the presence of rectilinear barriers, *Networks* **14**(3) (1984), 393–410.
- [PS] F. P. Preparata, M. I. Shamos, *Computational Geometry: an Introduction*, Springer-Verlag, New York, 1985.
- [PSR] R. Pollack, M. Sharir, G. Rote, Computing the geodesic center of a simple polygon, *Discrete and Computational Geometry* **4** (1989), 611–626.
- [S] S. Suri, Computing geodesic furthest neighbors in simple polygons, *Journal of Computer and Systems Sciences* **39** (1989), 220–235.
- [ST] N. Sarnak, R. E. Tarjan, Planar point location using persistent search trees, *Communications of the ACM* **29**(7) (1986), 660–679.
- [T] G. Toussaint, An optimal algorithm for computing the relative convex hull of a set of points in a polygon, in *Signal Processing III: Theories and Applications, Proceedings of EUSIPCO 86*, Part 2, North-Holland, Amsterdam 1986, pp. 853–856.

Received November 11, 1988, and in revised form October 15, 1992.