Straight Line Skeleton in Linear Time, Topologically Equivalent to the Medial Axis

Mirela Tănase Remco C. Veltkamp
Institute of Information and Computing Sciences,
Utrecht University, The Netherlands
{mirela, Remco.Veltkamp}@cs.uu.nl

1. Introduction

Skeletons have long been recognized as an important tool in computer graphics, computer vision and medical imaging. The most known and widely used skeleton is the *medial axis* [1]. However, the presence of parabolic arcs may be a disadvantage for the medial axis. The straight line skeleton [2], on the other hand, is composed only of line segments. In [3] a polygon decomposition based on the straight line skeleton was presented. The results obtained from the implementation indicated that this technique provides plausible decompositions for a variety of shapes. However, sharp reflex angles have a big impact on the form of the straight line skeleton, which in turn has a large effect on the decomposition (see figure 1).

This paper presents the *linear axis*, a new skeleton for polygons. It is related to the medial axis and the straight line skeleton, being the result of a wavefront propagation process. The wavefront is linear and propagates by translating edges at constant speed. The initial wavefront is an altered version of the original polygon: zero-length edges are added at reflex vertices. The linear axis is a subset of the straight line skeleton of the altered polygon. In this way, the counter-intuitive effects in the straight line skeleton caused by sharp reflex vertices are alleviated. We introduce the notion of ε-equivalence between two skeletons, and give an algorithm that computes the number of hidden edges for each reflex vertex which yield a linear axis ε-equivalent to the medial axis. This linear axis and thus the straight line skeleton can then be computed from the medial axis in linear time for non-degenerate polygons, i.e. polygons with a constant number of "nearly co-circular" sites.

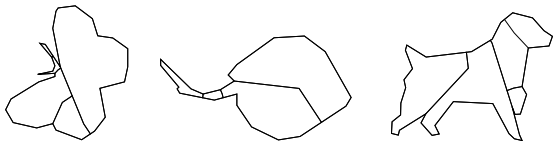

Fig. 1. A decomposition [3] based on split events of the straight line skeleton gives counter-intuitive results if the polygon contains sharp reflex vertices.

2. Linear axis

Let P be a simple, closed polygon. The set of sites defining the Voronoi diagram $VD(P)$ of P is the set of line segments and reflex vertices of P. The set of points $\mathcal{U}(d)$ inside P, having some fixed distance d to the polygon is called a *uniform wavefront*. As the distance d increases, the wavefront points move at equal, constant velocity along the normal direction. This process is called *uniform wavefront propagation*. During the propagation, the breakpoints between the line segments and circular arcs in $\mathcal{U}(d)$ trace the Voronoi diagram $VD(P)$. The medial axis $M(P)$ is a subset of $VD(P)$; the Voronoi edges incident to the reflex vertices are not part of the medial axis (see figure 2(a)).

The straight line skeleton is also defined as the trace of adjacent components of a propagating wavefront. The wavefront is linear and is obtained by translating the edges of the polygon in a self-parallel manner, keeping sharp corners at reflex vertices (edges incident to a reflex vertex will grow in length, see figure 2(b)). This also means that points in the wavefront move at different speeds: wavefront points originating from reflex vertices move faster than points originating from the edges incident to those vertices. In fact, the speed of the wavefront points originating from a reflex vertex gets arbitrary large when the internal angle of the

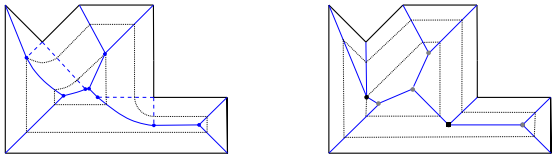

Fig. 2. Medial Axis (a) vs. Straight Line Skeleton (b). Instances of the propagating wavefront generating the skeletons are drawn with dotted line style in both cases. In (a) the Voronoi edges that are not part of the medial axis are drawn with dashed line style.

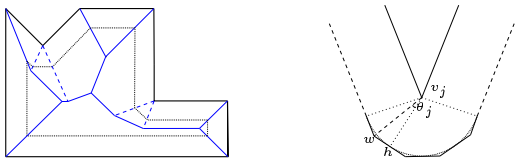

Fig. 3. (a) The linear axis in the case when one hidden edge is inserted at each reflex vertex. A linear wavefront is drawn in dotted line style; the dashed lines are the bisectors that are not part of the linear axis. (b) The linear offset (solid line) of a reflex vertex with 3 associated edges.

vertex gets arbitrary close to 2π. And that is why sharp reflex vertices have such a big impact on the form of the straight line skeleton. The straight line skeleton $SLS(P)$ of P is the trace in the propagation of the vertices of the wavefront.

Let $\{v_1, v_2, \ldots, v_n\}$ denote the vertices of a simple polygon P and let $\kappa = (k_1, k_2, \ldots, k_n)$ be a sequence of natural numbers. If v_i is a convex vertex of P, $k_i = 0$, and if it is reflex vertex, $k_i \geq 0$. Let $\mathcal{P}^\kappa(0)$ be the polygon obtained from P by replacing each reflex vertex v_i with $k_i + 1$ identical vertices. We assume that these vertices are the endpoints of k_i zero-length edges, which will be referred to as the *hidden edges* associated with v_i. The directions of the hidden edges are chosen such that the reflex vertex v_i of P is replaced in $\mathcal{P}^\kappa(0)$ by $k_i + 1$ "reflex vertices" of equal internal angle. The polygon $\mathcal{P}^\kappa(t)$ represents the linear wavefront, corresponding to a sequence κ of hidden edges, at moment t. The propagation consists of translating edges at constant unit speed, in a self-parallel manner.

Definition 1 *The **linear axis** $L^\kappa(P)$ of P, corresponding to a sequence κ of hidden edges, is the trace of the convex vertices of the linear wavefront $\mathcal{P}^\kappa$ in the above propagation process.*

Obviously, $L^\kappa(P)$ is a subset of $SLS(\mathcal{P}^\kappa(0))$; we only have to remove the bisectors traced by the reflex vertices of the wavefront (see figure 2 (a)).

Lemma 1 *If any reflex vertex v_j of internal angle $\alpha_j \geq 3\pi/2$ has at least one associated hidden edge, then no new reflex vertices are introduced to the wavefront during the propagation of $\mathcal{P}^\kappa$.*

For the rest of this paper, we will assume that each selection κ of hidden edges that is considered, satisfies the condition in lemma 1. This ensures that $L^\kappa(P)$ is a connected graph.

A *site* of P is a line segment or a reflex vertex. If S is an arbitrary site of P, we denote by $\mathcal{P}^\kappa_S(t)$ the points in $\mathcal{P}^\kappa(t)$ originating from S. We call $\mathcal{P}^\kappa_S(t)$ the *(linear) offset of S at moment t.*

Individual points in $\mathcal{P}^\kappa$ move at different speeds. The fastest moving points in $\mathcal{P}^\kappa$ are its reflex vertices. The slowest moving points have unit speed. Let v_j be a reflex vertex of P, of internal angle α_j, and with k_j associated hidden edges. Let $\mathcal{P}^\kappa_{v_j}(t)$ be the offset of v_j at some moment t.

Lemma 2 *Points in $\mathcal{P}^\kappa_{v_j}(t)$ move at speed at most*

$$s_j = 1/\cos\left((\alpha_j - \pi)/(2 * (k_j + 1))\right).$$

Each convex vertex of $\mathcal{P}^\kappa(t)$ is a breakpoint between two linear offsets. Lemmas 3 and 4 describe the trace of the intersection of two adjacent offsets in the linear wavefront propagation. If v is a vertex and e is an edge non-incident with v, we denote by $P(v, e)$ the parabola with focus v and directrix the line supporting e. For any real value $r > 1$, we denote by $H^r(v, e)$ the locus of points whose distances to v and the line supporting e are in constant ratio r. This locus is a hyperbola branch. If u and v are reflex vertices, we denote by $C^r(u, v)$ the Apollonius circle associated with u, v and ratio $r \neq 1$. $C^r(u, v)$ is the locus of points whose distances to u and v, respectively, are in constant ratio r.

Let e be an edge and v_j a reflex vertex of P. The points in the offset $\mathcal{P}^\kappa_e$ move at unit speed, the points in the offset $\mathcal{P}^\kappa_{v_j}$ move at speeds varying in the interval $[1, s_j]$, where s_j is given by lemma 2.

Lemma 3 *If the linear offsets of v_j and e become adjacent, the trace of their intersection is a polygonal chain that satisfies:*
1) it lies in the region between $P(v_j, e)$ and $H^{s_j}(v_j, e)$;
2) the lines supporting its segments are tangent to $P(v_j, e)$; the tangent points lie on the traces of the unit speed points in $\mathcal{P}^\kappa_{v_j}$;
3) $H^{s_j}(v_j, e)$ passes through those vertices which lie on the trace of a reflex vertex of $\mathcal{P}^\kappa_{v_j}$.

Let v_i and v_j be reflex vertices of P. The points

in the offset $\mathcal{P}^{\kappa}_{v_i}$ move at speeds varying in the interval $[1, s_i]$ while the points in the offset $\mathcal{P}^{\kappa}_{v_j}$ move at speeds varying in the interval $[1, s_j]$.

Lemma 4 *If the linear offsets of v_i and v_j become adjacent, the trace of their intersection is a polygonal chain that lies outside the Apollonius circles $C^{s_i}(v_i, v_j)$ and $C^{s_j}(v_j, v_i)$.*

3. Linear Axis vs. Medial Axis

Obviously, the larger the number of hidden edges associated with the reflex vertices, the closer the corresponding linear axis approximates the medial axis. For many applications of the medial axis, an approximation that preserves the main visual cues of the shape, though non-isomorphic with the medial axis, is perfectly adequate. The way we now define the ε-equivalence between the medial axis and the linear axis, will allow us to compute a linear axis that closely approximates the medial axis using only a small number of hidden edges.

Let $\varepsilon \geq 0$; an ε-edge is a Voronoi edge generated by four almost co-circular sites. Let $b_i b_j$ be a Voronoi edge, with b_i equally distanced to sites S_k, S_i, S_l, and b_j equally distanced to S_k, S_j, S_l.

Definition 2 *The edge $b_i b_j$ is an **ε-edge** if $d(b_i, S_j) < (1+\varepsilon) d(b_i, S_i)$ or $d(b_j, S_i) < (1+\varepsilon) d(b_j, S_j)$.*

A path between two nodes of $M(P)$ is an *ε-path* if it is made only of ε-edges. For any node b of M, a node b' such that the path between b and b' is an ε-path, is called an *ε-neighbour* of b. Let $N_\varepsilon(b)$ be the set of ε-neighbours of b. The set $\{b\} \cup N_\varepsilon(b)$ is called an *ε-cluster*.

Let (V_M, E_M) be the graph induced by $M(P)$ on the set of vertices V_M composed of the convex vertices of P and the nodes of degree 3 of $M(P)$. Let $(V_{L^\kappa}, E_{L^\kappa})$ be the graph induced by $L^\kappa(P)$ on the set of vertices V_{L^κ} composed of the convex vertices of P and the nodes of degree 3 of $L^\kappa(P)$.

Definition 3 *$M(P)$ and $L^\kappa(P)$ are **ε-equivalent** if there exists a surjection $f : V_M \to V_{L^\kappa}$ so that:*
i) $f(p) = p$, for all convex p of P;
ii) $\forall\, b_i, b_j \in V_M$ with $b_j \notin N_\varepsilon(b_i)$, $\exists$ an arc in E_M connecting b_i and $b_j \Leftrightarrow \exists$ an arc in E_{L^κ} connecting $f(b'_i)$ and $f(b'_j)$ where $b'_i \in \{b_i\} \cup N_\varepsilon(b_i)$ and $b'_j \in \{b_j\} \cup N_\varepsilon(b_j)$.

The following lemma gives a sufficient condition for the ε-equivalence of the two skeletons. The path between two disjoint Voronoi cells $VC(S_i)$ and $VC(S_j)$ is the shortest path in $M(P)$ between a point of $VC(S_i) \cap M(P)$ and a point of $VC(S_j) \cap M(P)$.

Lemma 5 *If the only sites whose linear offsets become adjacent in the propagation, are sites with adjacent Voronoi cells or sites whose path between their Voronoi cells is an ε-path, then the linear axis and the medial axis are ε-equivalent.*

We now present an algorithm for computing a sequence κ of hidden edges such that the resulting linear axis is ε-equivalent to the medial axis. The algorithm handles pairs of sites whose linear offsets must be at any moment disjoint in order to ensure the ε-equivalence of the two skeletons. These are sites with disjoint Voronoi cells and whose path between these Voronoi cells is not an ε-path (see lemma 5). However, we do not have to consider each such pair. The algorithm actually handles the pairs of *conflicting sites*, where two sites S_i and S_j (at least one being a reflex vertex) are said to be conflicting if the path between their Voronoi cells contains exactly one non-ε edge. When handling the pair S_i, S_j we check and, if necessary, adjust the maximal speeds s_i and s_j of the offsets of S_i and S_j, respectively, so that these offsets remain disjoint in the propagation. This is done in the subroutine HandleConflictingPair by looking locally at the configuration of the uniform wavefront and using the localization constraints for the edges of the linear axis given by lemmas 3 and 4.

The algorithm for computing a sequence κ of hidden edges can be summarized as follows.

Algorithm ComputeHiddenEdges
Input A simple polygon P and its medial axis.
Output A number of hidden edges for each reflex vertex.

1. For each reflex vertex S_j of P
   if $\alpha_j \geq 3\pi/2$ **then** $s_j \leftarrow 1/\cos((\alpha_j - \pi)/4)$
   else $s_j \leftarrow 1/\cos((\alpha_j - \pi)/2)$
2. Compute all pairs of conflicting sites
3. For each pair of conflicting sites S_i, S_j
   HandleConflictingPair(S_i,S_j)
4. For each reflex vertex S_j of P
   $k_j \leftarrow \lceil (\alpha_j - \pi)/(2\cos^{-1}(1/s_j)) \rceil$.

Theorem 6 *Algorithm ComputeHiddenEdges computes a sequence of hidden edges that leads to a linear axis ε-equivalent to the medial axis.*

The performance of this algorithm depends on the number of conflicting pairs. This in turn depends on the number of nodes in the ε-clusters of $M(P)$. If any ε-cluster of M has a constant number of nodes, there are only a linear number of con-

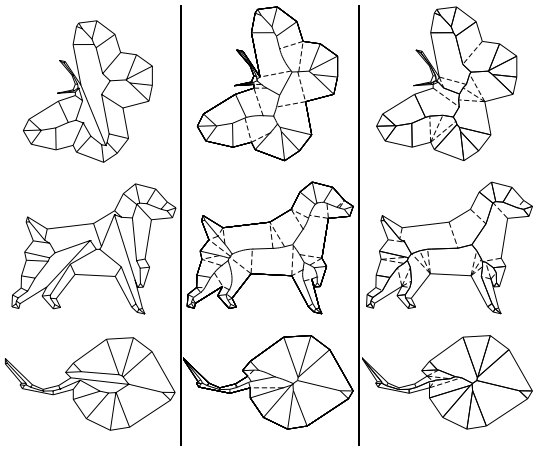

Fig. 4. A comparison of the linear axis (right) with the medial axis (middle) and the straight line skeleton (left). The skeletons are drawn in solid line style.

flicting pairs. Each conflicting pair is handled in constant time [4], thus in this case, ComputeHiddenEdges computes the sequence κ in linear time. There is only a limited class of shapes with a constant number of clusters, such that each has a linear number of nodes.

Once we have a sequence κ that ensures the ε-equivalence between the corresponding linear axis and the medial axis, we can construct this linear axis. The medial axis can be computed in linear time [5]. Despite its similarity to the medial axis, the fastest known algorithms [6] [7] for the straight line skeleton are slower. Any of these algorithms can be used to compute the straight line skeleton of $\mathcal{P}^\kappa(0)$. The linear axis $L^\kappa(P)$ corresponding to the sequence κ is then obtained from $SLS(\mathcal{P}^\kappa(0))$ by removing the bisectors incident to the reflex vertices of P.

However, if M has only ε-clusters of constant size, $L^\kappa(P)$ can be computed from the medial axis in linear time by adjusting the medial axis. In computing the linear axis, we adjust each non-ε-edge of the medial axis to its counterpart in the linear axis. When adjusting an edge b_ib_j we first adjust the location of its endpoints to the location of the endpoints of its counterpart. If node b_i is part of an ε-cluster, we compute first the counterparts of the nodes in this cluster based on a local reconstruction of the linear wavefront. The adjustment of a node's location is done in constant time, if its ε-cluster has constant size, i.e. when the polygon is non-degenerate. Finally, we use lemmas 3-4 to replace the parabolic arc or the perpendicular bisector with the corresponding chain of segments.

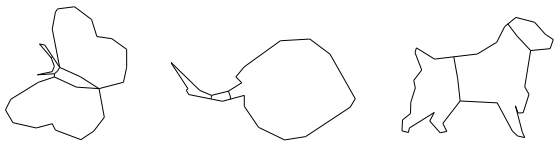

Fig. 5. A decomposition based on split events of the linear axis gives natural results even if the polygon contains sharp reflex vertices.

Theorem 7 *For a non-degenerate polygon P, the straight line skeleton of $\mathcal{P}^\kappa(0)$ can be computed in linear time.*

We have implemented the algorithm ComputeHiddenEdges and the algorithm that constructs the linear axis from the medial axis. Figure 4 illustrates the straight line skeleton (left column), medial axis (middle column) and the linear axis (right column) of three contours. We also used the linear axis instead of the straight line skeleton to decompose the contours in figure 1 based on wavefront split events [3]. The results are presented in figure 5; we see that the unwanted effects of the sharp reflex vertices are eliminated and the results of this first step in the decomposition look more natural.

References

[1] H. Blum, A Transformation for Extracting New Descriptors of Shape, *Symposium Models for Speech and Visual Form.* (ed: W. Wathen-Dunn. MIT Press, 1967) 362–381.

[2] O. Aichholzer, F. Aurenhammer, Straight Skeletons for General Polygonal Figures in the Plane, in: *Proc. 2nd International Computing and Combinatorics Conference COCOON '96.* 117–126.

[3] M. Tănase, R.C. Veltkamp, Polygon Decomposition based on the Straight Line Skeleton, in: *Proc. 19th ACM Symposium on Computational Geometry.* (2003) 58–67.

[4] M. Tănase, R.C. Veltkamp, Straight Line Skeleton in Linear Time, Topologically Equivalent to the Medial Axis, TR (2004).

[5] F. Chin, J. Snoeyink, C.-A. Wang, Finding the Medial Axis of a Simple Polygon in Linear Time. *Discrete Computational Geometry* **21(3)** (1999) 405–420.

[6] D. Eppstein, J. Erickson, Raising Roofs, Crashing Cycles, and Playing Pool: Applications of a Data Structure for Finding Pairwise Interactions, *Discrete and Computational Geometry* **22(4)** (1999) 569–592.

[7] S.-W. Cheng, A. Vigneron, Motorcycle Graphs and Straight Skeletons, in: *Proc. 13th ACM-SIAM Symp. Discrete Algorithms* (2002) 156–165.