

On the Structure of Straight Skeletons

Kira Vyatkina

Saint Petersburg State University, Department of Mathematics and Mechanics,
28 Universitetsky pr., Stary Peterhof, Saint Petersburg 198504, Russia

kira@meta.math.spbu.ru

<http://meta.math.spbu.ru/~kira>

Abstract. For a planar straight line graph G , its straight skeleton $S(G)$ can be partitioned into two subgraphs $S^c(G)$ and $S^r(G)$ traced out by the convex and by the reflex vertices of the linear wavefront, respectively. By further splitting $S^c(G)$ at the nodes, at which the reflex wavefront vertices vanish, we obtain a set of connected subgraphs M_1, \dots, M_k of $S^c(G)$. We show that each M_i is a pruned medial axis for a certain convex polygon Q_i closely related to G , and give an optimal algorithm for computation of all those polygons, for $1 \leq i \leq k$. Here “pruned” means that M_i can be obtained from the medial axis $M(Q_i)$ for Q_i by appropriately trimming some (if any) edges of $M(Q_i)$ incident to the leaves of the latter.

1 Introduction

The straight skeleton was first introduced for simple polygons by Aichholzer et al. [2], and soon generalized to the case of planar straight line graphs by Aichholzer and Aurenhammer [1]. It has promptly found a number of applications in such areas as surface reconstruction [3], computational origami [6], and many others; besides, it has served as a basis for another type of skeleton called linear axis [9,10,12]. Its advantage over the only previously known skeleton – the medial axis – resides in the fact that all its edges are straight line segments, while the medial axis for a non-convex polygonal domain necessarily contains parabolic edges as well.

Both the straight skeleton and the medial axis can be defined through wavefront propagation; we shall illustrate this for the case of polygons. Initially, the wavefront coincides with the given polygon. To obtain the straight skeleton, we let the wavefront edges move inside the polygon at equal speed, thereby remaining parallel to themselves, and keep track of the movement of its vertices. The underlying process is referred to as a *linear wavefront propagation*. To obtain the medial axis, we apply a *uniform wavefront propagation*, during which all the wavefront points move inside at constant speed. Thus, at time $t > 0$, the uniform wavefront consists of the interior points of the polygon at the distance t from its boundary. In the process, the wavefront vertices trace out the edges of the medial axis.

However, the straight skeleton is computationally more expensive than the medial axis: for an n -gon with r reflex vertices, the fastest known deterministic

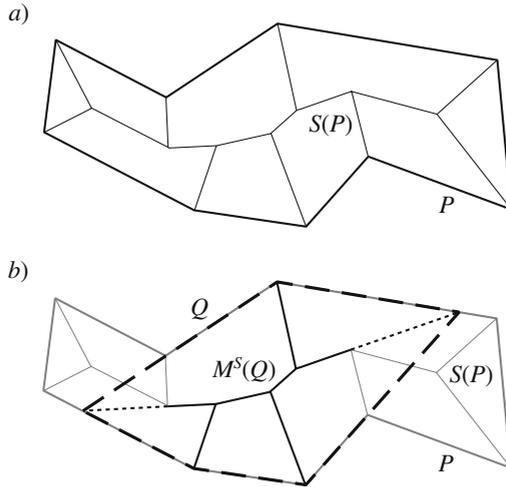


Fig. 1. a) A simple polygon P and its straight skeleton $S(P)$. b) The subtree $M^S(Q)$ of $S(P)$ is a pruned medial axis for the convex polygon Q .

algorithm for its construction, proposed by Eppstein and Erickson [7], requires $O(n^{1+\epsilon} + n^{8/11+\epsilon}r^{9/11+\epsilon})$ time, where ϵ is an arbitrarily small positive constant; the same bounds hold for planar straight line graphs. The best existing randomized algorithm by Cheng and Vigneron [4] computes the straight skeleton for a non-degenerate simple polygon in $O(n \log^2 n + r\sqrt{r} \log r)$ expected time; for a degenerate one, the expected time bound amounts to $O(n \log^2 n + r^{17/11+\epsilon})$. For a non-degenerate polygon with h holes, their algorithm takes $O(n\sqrt{h+1} \log^2 n + r\sqrt{r} \log r)$ expected time. But the medial axis for a simple polygon can be obtained in linear time [5], and for a planar straight line graph – in $O(n \log n)$ time [13]. It is a common belief that the straight skeleton can be computed in a more efficient way than it is possible nowadays. Yet development of such methods is likely to require investigation of additional properties of the straight skeleton. In this work, we take one step further in that direction.

Our main observation is that during the linear wavefront propagation for a planar straight line graph G , the pieces of the wavefront locally interact exactly in the same way as if they originated from the boundary of a (bounded or unbounded) convex polygon, the sides of which lie on the lines through certain edges of G . Since for a convex polygon, the two kinds of propagation proceed identically, this implies that some pieces of the medial axes for such polygons are embedded in the straight skeleton $S(G)$ for G . (Figure 1 illustrates this observation for a simple polygon P and its straight skeleton.) We formalize our ideas by indicating those pieces in $S(G)$, providing an efficient algorithm for computation of the corresponding convex polygons, and pointing out that for each such polygon Q , the piece $M^S(Q)$ of its medial axis $M(Q)$ present in $S(G)$ can be obtained by appropriately trimming the edges of $M(Q)$ incident to the

vertices of Q not being those of G , and possibly, the unbounded ones, if any exist. Consequently, we say that $M^S(Q)$ is a *pruned* medial axis for Q .

A preliminary version of this paper, in which the case of simple polygons only is addressed, has appeared as [11].

In the next section, we specify classification of events that occur during the linear wavefront propagation. Section 3 analyzes the structure of the straight skeleton for a simple polygon. In Section 4, the obtained results are extended to polygons of holes, and then to planar straight line graphs. We conclude by indicating a potential direction for future research.

2 Event Classification

For clarity of exposition, we shall first develop our reasoning for the case of simple polygons, and in Section 4, a passage to general cases will be performed.

Let P be a simple polygon. Consider the process of constructing the straight skeleton $S(P)$ for P through the linear wavefront propagation. During the propagation, the wavefront structure changes as certain events occur. The very first event taxonomy [2,1] distinguishes between *edge events* and *split events*. More elaborated variations [7,4,12] additionally recognize *vertex events*. The most refined version can be found in [8].

For our purposes, we find it convenient to follow the classification we propose below.

1. *Edge event*: an edge incident to two convex vertices shrinks to zero.
2. *Sticking event*: a reflex vertex runs into an edge, thereby giving rise to precisely one (convex) vertex in the wavefront.
3. *Split event*: a reflex vertex collides into an edge, thereby splitting a wavefront component into two, and giving rise to a convex vertex in either part.
4. *Vertex event*: two reflex vertices collide together, thereby giving rise to a new reflex vertex in the wavefront.

A sticking event is attended either with annihilation of a wavefront edge incident to the reflex vertex and adjacent to the edge involved in the event, or with an edge collision. In the latter case, two parallel wavefront edges meet, thereby producing an edge of the straight skeleton, and the nodes incident to the former are brought forth either by two sticking events, or by a sticking event and an edge event. If the second event is a sticking one as well, we shall be left with two components of the wavefront instead of one (the found skeleton edge is no longer part of the wavefront).

At a vertex event, either two reflex vertices and nothing else meet at the same point, or the vertex collision is attended with simultaneous annihilation of two wavefront edges, which form a chain connecting the two vertices, and are adjacent at a convex vertex of the wavefront. Vertex events often either require special handling [7], or are ruled out by non-degeneracy assumptions [4]. We do not have to exclude vertex events from our consideration. All types of events are illustrated in Fig. 2.

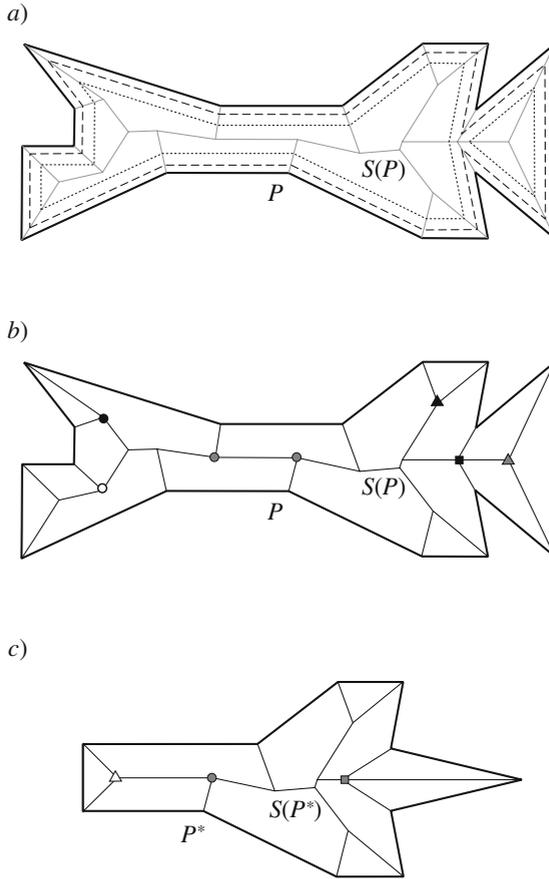


Fig. 2. a) A polygon P and its straight skeleton $S(P)$. The linear wavefront at two times is shown: soon after the propagation starts (dashed), and soon after the vertex event occurs (dotted). b) For some inner nodes of $S(P)$, the type of the corresponding event is indicated. Black triangle: an edge event. Gray triangle: three simultaneous edge events, which lead to annihilation of a wavefront component. Black circle: a sticking event. Two gray circles: two sticking events that occur simultaneously; at the same moment, the edge of $S(P)$ between the two nodes is generated entirely. White circle: a split event. Black box: a vertex event, at which two reflex vertices and nothing else meet at the same place. c) A polygon P^* and its straight skeleton $S(P^*)$. For three nodes of $S(P^*)$, the type of the underlying event is indicated. The gray circle and the white triangle denote a sticking event and an edge event, respectively; those two events occur simultaneously, and at the same moment, the edge of $S(P^*)$ between the corresponding nodes is generated entirely. The gray box denotes a vertex event attended with annihilation of the two wavefront edges forming a chain between the two colliding vertices.

A vertex event may lead to appearance of a degenerate vertex in the wavefront, with an internal angle of π . Such vertices should be handled as reflex ones.

Any inner node u of $S(P)$, which emerged not from a vertex event, but has degree $d \geq 4$, is produced by $(d-2)$ edge and/or split and/or sticking events that simultaneously occur at the same location. Those events can be handled one at a time, with any two consecutive ones being separated by a zero time interval. Therefore, any such node of $S(P)$ can be interpreted as $(d-2)$ coinciding nodes of degree three connected by $(d-3)$ edges of zero length in such a way that the subgraph induced by those nodes is a tree.

On the other hand, at any point, at most one vertex event can occur. Thus, if a node u of degree $d \geq 5$ has an associated vertex event, then it is produced by that vertex event and $(d-4)$ edge and/or split and/or sticking events that simultaneously occur at the same location. All those events can be handled as in the above case, and u admits an analogous interpretation.

Consequently, we may further suppose that any node of $S(P)$ resulting from a vertex event has degree four, and any other its inner node has degree three. A wavefront component annihilation then corresponds to three simultaneous edge events at the same point.

3 Separation of the Pruned Medial Axes

In this section, we outline step by step both extraction of the pruned medial axes from the straight skeleton and reconstruction of the respective convex polygons.

3.1 Partition of the Straight Skeleton

Let P be a simple polygon with n vertices, r of those being reflex; assume that $r \geq 1$. First, let us decompose the straight skeleton $S(P)$ for P into two subgraphs $S^c(P)$ and $S^r(P)$, being parts of $S(P)$ traced out by the convex and by the reflex wavefront vertices, respectively. By construction, $S^r(P)$ is a forest. To make it more precise, if no vertex events occur during the propagation, $S^r(P)$ consists of r edges of $S(P)$ incident to the reflex vertices of P ; otherwise, at least one tree from $S^r(P)$ has three or more edges.

Next, let us split $S^c(P)$ at the nodes, at which the reflex wavefront vertices vanish (Fig. 3a,b). We shall be left with a decomposition of $S^c(P)$ into $k \leq r+1$ connected subgraphs M_1, M_2, \dots, M_k . We claim that for any $i, 1 \leq i \leq k$, M_i is a part of the medial axis for a convex polygon. Strictly speaking, there are infinitely many such polygons; of course, we would like to retrieve one with the least computational effort. Below we shall formalize our intent.

3.2 Extraction of the Boundary Chains

For any edge e of P , we define its corresponding *cell* $C(e)$ to be the face of the partition of P induced by $S(P)$, which is adjacent to e . Equivalently, $C(e)$ is the region swept in the propagation by the portion of the linear wavefront originating from e .

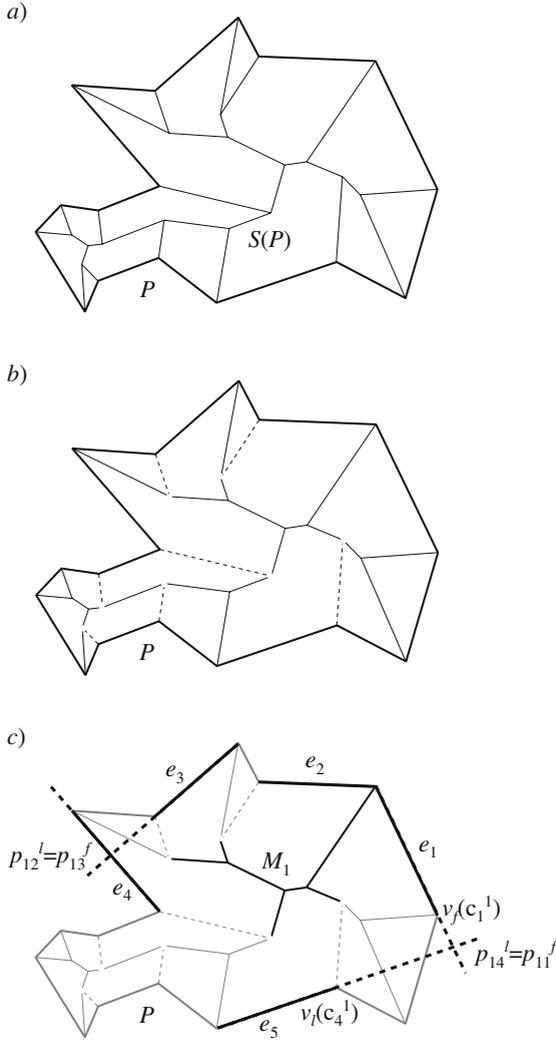


Fig. 3. a) A simple polygon P and its straight skeleton $S(P)$. b) Decomposition of the part $S^c(P)$ of the straight skeleton traced out by the convex vertices. The edges of $S(P)$ traced out by the reflex vertices are shown dashed. c) For the fragment M_1 , $E_1 = \{e_1, e_2, e_3, e_4, e_5\}$; $C_1 = \{c_1^1, c_2^1, c_3^1, c_4^1\}$, where the chain c_1^1 is formed of e_1 and e_2 , and each of c_2^1 , c_3^1 , and c_4^1 consists of a single edge - of e_3 , e_4 , and e_5 , respectively. By prolonging the edges as shown in dotted lines, we obtain the rays $\bar{f}(c_1^1)$, $\bar{l}(c_2^1)$, $\bar{f}(c_3^1)$, and $\bar{l}(c_4^1)$, respectively; here $p_{14}^l = p_{11}^f = \bar{l}(c_4^1) \cap \bar{f}(c_1^1)$, and $p_{12}^l = p_{13}^f = \bar{l}(c_2^1) \cap \bar{f}(c_3^1)$.

Consider any M_i . Since $S(P)$ is a tree, M_i is a tree as well. The embedding of M_i in the plane induces a cyclic order of its leaves. For any consecutive pair of leaves, when walking from one of them to the other along the edges of M_i , we follow the boundary of some cell. Moreover, it can be easily verified that for any two such pairs of leaves, the corresponding cells must be different. These cells are also cyclically ordered, in compliance with the ordering of the leaves.

Now retrieve all the edges of P , such that the boundaries of their cells contribute to M_i ; denote the resulting set by E_i (Fig. 3c). From the above discussion, it follows that $|E_i|$ equals the number of the leaves in M_i . Let the edges in E_i inherit the cyclic order of the cells. For any two consecutive edges $e, e' \in E_i$, their cells $C(e)$ and $C(e')$ share an edge of M_i incident to a leaf. If the leaf corresponds to a convex vertex of P , then e and e' share this vertex. Otherwise, the leaf corresponds to an inner node u of $S(P)$ adjacent to a reflex vertex of P . In this case, e and e' can be (but not necessarily are) adjacent only if E_i consists solely of e and e' . To see this, suppose e and e' are adjacent at vertex v of P . Then $C(e)$ and $C(e')$ must share the edge of $S(P)$ incident to v . On the other side, any two cells can share at most one edge of $S(P)$. Therefore, the edge of M_i shared by $C(e)$ and $C(e')$ must be uv . It follows immediately that v is convex, M_i consists of a single edge uv , and e and e' are the only two edges in E_i .

Thus, we conclude that the edges from E_i together compose one or a few disjoint convex chains cut out of the boundary of P . Let $\mathcal{C}_i = \{c_1^i, \dots, c_{m_i}^i\}$ denote the set of those chains; observe that m_i equals the number of the leaves of M_i that correspond to the inner nodes of $S(P)$. Denote by $f(c_j^i)$ and $l(c_j^i)$ the first and the last edge of the chain c_j^i , respectively; assume that c_j^i is traversed from $f(c_j^i)$ to $l(c_j^i)$ when walking counterclockwise along the boundary of P , where $1 \leq j \leq m_i$. If c_j^i consists of a single edge, then $f(c_j^i) = l(c_j^i)$. Let $v_f(c_j^i)$ and $v_l(c_j^i)$ denote the first and the last vertex of c_j^i , respectively. Without loss of generality, suppose that the chains in \mathcal{C}_i are enumerated in such a way that $f(c_{d+1}^i)$ follows $l(c_d^i)$ in the cyclic order of the edges from E_i , where $1 \leq d < m_i$. To unify the notation, let $c_{m_i+1}^i = c_1^i$, and let $c_0^i = c_{m_i}^i$.

Lemma 1. *For any j , $1 \leq j \leq m_i$, at least one of $v_l(c_j^i)$ and $v_f(c_{j+1}^i)$ is reflex.*

Proof. Consider the edge (u, x) of M_i shared by $C(l(c_j^i))$ and $C(f(c_{j+1}^i))$. Assume that when walking along (u, x) from u to x , we follow counterclockwise the boundary of $C(l(c_j^i))$; then u is necessarily a leaf of M_i , which corresponds to an inner node u' of $S(P)$.

If u' appeared as a result of a vertex event, then for any of the four cells incident to u' , its generative edge is incident to a reflex vertex of P , which is adjacent with u' in $S(P)$. In particular, this holds for the edges $l(c_j^i)$ and $f(c_{j+1}^i)$. Otherwise, u' has degree three; therefore, precisely three cells meet at u' . Two of those are $C(l(c_j^i))$ and $C(f(c_{j+1}^i))$. On the other hand, u' is adjacent to a reflex vertex r of P , and the cells of the both edges of P incident to r are incident to u' . This implies that r is incident either to $l(c_j^i)$ or to $f(c_{j+1}^i)$.

To prove our claim, it remains to demonstrate that u' can be adjacent neither to the first vertex of $l(c_j^i)$, nor to the last vertex of $f(c_{j+1}^i)$. By symmetry, it

suffices to show the first statement. To this end, recall that $C(l(c_j^i))$ is a simple polygon. Let v denote the first vertex of $l(c_j^i)$. Note that $v_l(c_j^i)$ is encountered immediately after v , and x – immediately after u' , when walking counterclockwise along the boundary of $C(l(c_j^i))$. Consequently, u' and v are non-adjacent vertices of $C(l(c_j^i))$, and thus, they cannot be adjacent in $S(P)$.

Corollary 1. *For any j , $1 \leq j \leq m_i$, one of the following three possibilities occurs in the propagation:*

- $v_l(c_j^i)$ is a reflex vertex that runs into $f(c_{j+1}^i)$;
- $v_f(c_{j+1}^i)$ is a reflex vertex that runs into $l(c_j^i)$;
- $v_l(c_j^i)$ and $v_f(c_{j+1}^i)$ are both reflex and collide.

3.3 Edge Delineation

Consider any tree M_i , where $1 \leq i \leq k$. We shall analyze how the edges of M_i could have been traced out.

Let us first assume that M_i has precisely one edge g . Then E_i consists of two edges e_1 and e_2 , which can represent either a single or two separate chains comprising C_i . In the former case, g is incident to the convex vertex of P shared by e_1 and e_2 , and its delineation is terminated by a sticking event, at which a reflex vertex incident to a wavefront edge emanating from one of e_1 and e_2 , runs into an edge originating from the other one.

In the latter case, g could either appear entirely at once, or be traced out from one endpoint to another. The first possibility could be realized only if g emerged as a result of collision between two wavefront edges originating from e_1 and e_2 , respectively, attended by two simultaneous sticking events (see Section 2). Then, in particular, e_1 and e_2 must be two disjoint parallel edges lying not on the same line. Otherwise, g is traced out properly. It means that during the propagation, two wavefront edges emanating from e_1 and e_2 , respectively, become adjacent – through an event involving a reflex vertex – at a convex wavefront vertex w , which then traces out g . The delineation of g ends as in the above case.

If M_i has more than one edge, then it must have at least three edges, and any its edge incident to a leaf is also incident to an inner node.

Lemma 2. *Among the edges of M_i incident to a leaf, at most one was not traced out starting from the leaf.*

Proof. If M_i has a single edge, the claim trivially holds. Now assume that M_i has at least three edges.

Let z be any inner node of M_i . Observe that z must have appeared as a result of an edge event, and is incident to three edges of M_i . Two of those were delineated by the convex wavefront vertices incident to the vanishing edge; therefore, they were traced out towards z . The third one might have been traced out towards z as well, or delineated starting from z , or generated entirely at once (Fig. 4).

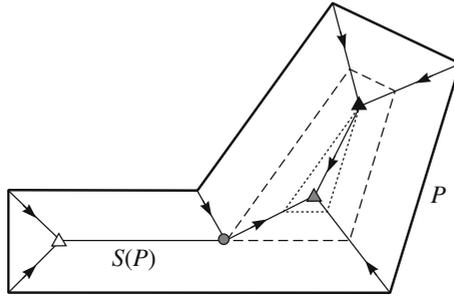


Fig. 4. A polygon P and its straight skeleton $S(P)$. Directions, in which the edges of $S(P)$ are traced out, are shown with arrows. Edge events of the three possible kinds are indicated with triangles. The horizontal edge of $S(P)$ is generated entirely at once; at the same moment, an edge event (white triangle) and a sticking event (gray circle) occur. The linear wavefront thereby becomes as depicted dash. Next, an edge event marked by the black triangle occurs, and the wavefront becomes as shown dotted. Finally, an edge event marked by the gray triangle occurs, and the wavefront vanishes.

First, suppose that M_i contains an edge (x, y) , which emerged from an edge collision and is incident to two inner nodes. Having removed this edge from M_i , we shall be left with two trees M_i^x and M_i^y containing the nodes x and y , respectively. Let us consider M_i^x . From the above observation, it follows that the two edges (x', x) and (x'', x) of M_i incident to x must have been traced out towards x . Having rooted M_i^x at x and traversed it in a breadth first (or depth first) order, at every step applying the same argument, we shall conclude that any edge of M_i^x incident to a leaf is delineated starting from that leaf. The same holds for M_i^y , and hence, for M_i .

In case M_i has an edge (u, x) , which emerged from an edge collision and is incident to a leaf u , a similar reasoning implies that any other edge of M_i incident to a leaf was traced out starting from the leaf.

Finally, assuming that all the edges of M_i have been delineated properly, let us suppose for contradiction that at least two edges (u, x) and (u', x') of M_i were traced out towards the leaves u and u' , respectively. Consider the path $(x_0 = x, x_1, \dots, x_h = x')$ between x and x' in M_i , where $h \geq 1$. Any node x_l is an inner node of M_i , where $1 \leq l \leq h$. Since the edge (u, x) was traced out from x to u , this implies that each edge (x_l, x_{l-1}) must have been traced out from x_l to x_{l-1} , for $1 \leq l \leq h$, and finally, the edge (u', x') must have been traced out from u' to x' , which is a contradiction.

Corollary 2. *At most one edge of M_i could have appeared as a result of an edge collision. If g is such an edge, then the delineation of M_i ended at the moment when g was generated.*

Corollary 3. *Let (u, x) be an edge of M_i incident to a leaf u . If (u, x) had been traced out from x to u , then the delineation of M_i ended at the moment when u was generated.*

For any edge $e \in E_i$, denote by p_e^i the path in M_i consisting of the edges that belong to $\partial C(e)$. Obviously, p_e^i connects two consecutive (with respect to the cyclic order) leaves of M_i . By analyzing the process of edge delineation, we shall prove the following lemma.

Lemma 3. *For any $e \in E_i$, p_e^i is a convex chain strictly monotone with respect to the line through e .*

Proof. Let u and w denote the two consecutive leaves of M_i being the endpoints of p_e^i , u preceding w in the counterclockwise order of the leaves. Without loss of generality, assume that e is horizontal, and $t_w \geq t_u \geq 0$, where t_w and t_u denote the time, at which the nodes w and u appeared, respectively. By definition of M_i , each of u and w either coincides with a convex vertex of P , or was generated at an event involving a reflex vertex, and any inner node of p_e^i must have been generated at an edge event.

Let us first consider a special case when M_i has a single edge g that appeared through an edge collision. Then so does p_e^i ; moreover, g is parallel to e , and thus, the claim holds.

Otherwise, the first edge (u, x) of p_e^i must have been traced out starting from u . Consequently, as soon as u comes into existence, it becomes a convex vertex of the wavefront incident to two edges \hat{e} and \hat{e}_u of the latter, the first of which originated from e , and the second – from the edge $e_u \in E_i$, which precedes e in the cyclic order of the edges composing E_i . Observe that \hat{e} lies on the left of \hat{e}_u , if the latter is oriented towards u , and the bisector of the interior wavefront angle between \hat{e} and \hat{e}_u , on which (u, x) lies, is inclined to the right.

If the node x coincides with w , our claim follows immediately. Otherwise, x is generated at an edge event. First, suppose that \hat{e} thereby neither collapses, nor collides with another edge; then \hat{e}_u must shrink to zero. Note that \hat{e}_u must have been incident to two convex vertices straight before the event. Let \hat{e}_x denote the second wavefront edge adjacent to \hat{e}_u . At the event, \hat{e} and \hat{e}_x become adjacent at a convex wavefront vertex v , the interior angle at which is less than was the one between \hat{e} and \hat{e}_u ; therefore, the next edge (x, x') of p_e^i traced out by v will be inclined to the right as well, and have a smaller slope than that of (u, x) . The process continues until the node w is reached, or \hat{e} shrinks to zero, or the generation of the next node involves collision of \hat{e} with another edge.

In case we finally reach w , both convexity and strict x -monotonicity of p_e^i are guaranteed by the above reasoning. Otherwise, if \hat{e} shrinks to zero, then the node x_0 generated thereby is the topmost vertex of $C(e)$, at which the part of p_e^i traced out starting from u meets the one delineated starting from w . Since symmetric arguments apply to the second part of p_e^i , our claim holds in this case.

Finally, if the generation of the next node \bar{x} resulting from an edge event is also attended by a collision of \hat{e} with another edge, then a horizontal edge (\bar{x}, \bar{y}) of p_e^i must appear at that moment. Two cases are possible: either $\bar{y} = w$, or \bar{y} appears simultaneously due to another edge event. The former case is similar to the one when we reach w , having started the delineation of p_e^i from u . In the

latter case, the part of p_e^i traced out starting from w terminates at \bar{y} ; the rest is similar to the case when the two parts meet at the topmost vertex of $C(e)$.

3.4 Chain Prolongation

For an edge (u, x) of M_i incident to a leaf u , let $r_{u,x}$ denote the open ray with the endpoint u , collinear to (u, x) , such that $r_{u,x} \cap (u, x) = \emptyset$.

Observation 1. *Let (u, x) be the edge of M_i shared by the cells $C(l(c_j^i))$ and $C(f(c_{j+1}^i))$ for some j , $1 \leq j \leq m_i$; assume that u denotes a leaf. If (u, x) is traced out from u to x , the ray $r_{u,x}$ and the lines through $l(c_j^i)$ and $f(c_{j+1}^i)$ intersect at a common point. Otherwise, neither line intersects $r_{u,x}$.*

For any chain c_j^i , let us take the edge $l(c_j^i)$ and prolong it to infinity, thereby eliminating $v_l(c_j^i)$; denote the resulting ray by $\bar{l}(c_j^i)$. Similarly, let $\bar{f}(c_j^i)$ denote the ray obtained by prolonging $f(c_j^i)$ to infinity beyond the vertex $v_f(c_j^i)$ (see Fig. 3c).

For any j , $1 \leq j \leq m_i$, consider the edge (u, x) of M_i shared by $C(l(c_j^i))$ and $C(f(c_{j+1}^i))$. Assume that when walking along (u, x) from u to x , we follow counterclockwise the boundary of $C(l(c_j^i))$; then u is necessarily a leaf of M_i corresponding to an inner node of $S(P)$.

If the lines through $l(c_j^i)$ and $f(c_{j+1}^i)$ intersect the ray $r_{u,x}$ at a point p , let $p_{ij}^l = p_{i,j+1}^f = p$; otherwise, let p_{ij}^l and $p_{i,j+1}^f$ be two points at infinity lying on $\bar{l}(c_j^i)$ and on $\bar{f}(c_{j+1}^i)$, respectively.

Lemma 4. *Let c_j^i be a chain formed of at least two edges. Then $p_{ij}^f \in \bar{f}(c_j^i)$, and $p_{ij}^l \in \bar{l}(c_j^i)$.*

Proof. Let us demonstrate that $p_{ij}^l \in \bar{l}(c_j^i)$; the second statement is symmetric. Consider the edge (u, x) of M_i shared by $C(l(c_j^i))$ and $C(f(c_{j+1}^i))$. Assume that when walking along (u, x) from u to x , we follow counterclockwise the boundary of $C(l(c_j^i))$; then u is necessarily a leaf of M_i corresponding to an inner node of $S(P)$. Without loss of generality, suppose that (u, x) is vertical, and x lies above u ; then $C(l(c_j^i))$ locally lies to the left of (u, x) .

If (u, x) was traced out from x to u or generated entirely at once, then p_{ij}^l is infinite, and the claim holds by definition of p_{ij}^l . Otherwise, the lines through $l(c_j^i)$, $f(c_{j+1}^i)$, and (u, x) intersect at a common point below u (see Observation 1), and the polygon locally lies above the edges $l(c_j^i)$ and $f(c_{j+1}^i)$.

Let z denote the vertex of P incident to $l(c_j^i)$ other than $v_l(c_j^i)$. If z lies on the left of the vertical line l through (u, v) , the claim holds. Otherwise, z must be a reflex vertex: if z was convex, the horizontal projection of its speed would be positive, and the wavefront edge originating from $l(c_j^i)$ would never cross l . Thus, $C(l(c_j^i))$ would lie entirely to the right of l , contradicting the fact that $C(l(c_j^i))$ locally lies to the left of (u, x) . But if z is reflex, then c_j^i must consist of a single edge, which is a contradiction.

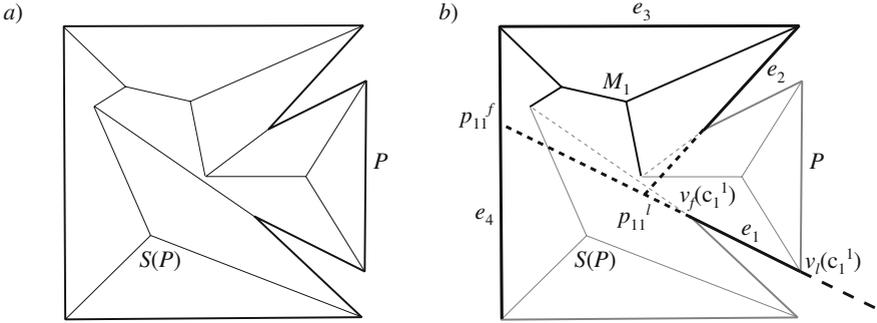


Fig. 5. a) A simple polygon P and its straight skeleton $S(P)$. b) The polygon P and its straight skeleton $S(P)$ are depicted gray; the edges of $S(P)$ incident to the reflex vertices of P are marked dotted. For the fragment M_1 of $S(P)$, $E_1 = \{e_1, e_2, e_3, e_4\}$; $C_1 = \{c_1^1, c_2^1\}$, where the chain c_1^1 consists solely of e_1 , and c_2^1 is formed of e_2, e_3 , and e_4 . By prolonging e_1 beyond $v_l(c_1^1)$ as shown dashed, we obtain the ray $\bar{l}(c_1^1)$. Here $p_{1,1}^f$ lies outside $\bar{l}(c_1^1)$; the chain \bar{c}_1^1 is the segment between $p_{1,1}^f$ and $p_{1,1}^l$, and $\bar{c}_1^1 \cap c_1^1 = \emptyset$.

If c_j^i is formed of a single edge, an analogous statement is not necessarily true (Fig. 5).

However, the following property will hold.

Lemma 5. *Let c_j^i be a chain formed of a single edge $(v_f(c_j^i), v_l(c_j^i))$. Then the vectors $\overrightarrow{p_{ij}^f p_{ij}^l}$ and $\overrightarrow{v_f(c_j^i) v_l(c_j^i)}$ have the same direction.*

Proof. If at least one of p_{ij}^l and p_{ij}^f is infinite, the claim obviously holds. Otherwise, it is implied by Lemma 3.

3.5 Formation of the Convex Polygons

For each j , $1 \leq j \leq m_i$, consider the chain c_j^i . If c_j^i is formed of at least two edges, construct a chain \bar{c}_j^i from c_j^i by adjusting the first and the last edge of the latter, so that they will terminate at p_{ij}^f and p_{ij}^l , respectively, if those are finite, or become unbounded, if the corresponding point is infinite. Otherwise, let \bar{c}_j^i be the segment between p_{ij}^f and p_{ij}^l (see Fig. 5); if one or both of p_{ij}^f and p_{ij}^l are infinite, \bar{c}_j^i will become a ray or a line, respectively. Let $\bar{c}^i = \cup_j \bar{c}_j^i$.

Lemma 6. *\bar{c}^i bounds a convex region in the plane.*

Proof. Applying Observation 1 and Lemma 2, we derive that \bar{c}^i can be either composed of two parallel lines or a polygonal chain (closed or open). In the former case, \bar{c}^i bounds an infinite strip in the plane, and the claim trivially holds.

Let us take M_i and prolong to infinity each its edge incident to a leaf, thereby eliminating all the leaves. Lemma 3 assures that no two neighbor unbounded

edges will intersect; consequently, the obtained tree-like structure \mathcal{M}_i induces a partition \mathcal{P}_i of the plane into $|E_i|$ unbounded convex regions.

Suppose that \bar{c}^i is a closed polygonal chain. It is easy to verify that the unbounded edges of \mathcal{M}_i pass precisely through the vertices of \bar{c}^i . Any edge \bar{c}_h of \bar{c}^i cuts away a convex polygon \bar{P}_h from a separate unbounded region of \mathcal{P}_i , where $1 \leq h \leq |E_i|$; in particular, no two such polygons overlap. Having glued all the polygons \bar{P}_h , for $1 \leq h \leq |E_i|$, along the corresponding bounded edges or parts of the unbounded edges of \mathcal{M}_i , we shall obtain a plane simply connected domain D^i bounded by \bar{c}^i . Our construction implies that any interior angle of D^i is less than π ; therefore, D^i is a convex polygon.

Now let \bar{c}^i be an open polygonal chain. If M_i has a single edge, then \bar{c}^i consists of two half-infinite edges and bounds a wedge with an apex angle less than π , which is an unbounded convex polygon.

Otherwise, for some j , $1 \leq j \leq m_i$, the points p_{ij}^l and $p_{i,j+1}^f$ are both infinite. Let g denote the edge of M_i shared by $C(l(c_j^i))$ and $C(f(c_{j+1}^i))$; observe that g is incident to a leaf. Consider the unbounded edge \bar{g} of \mathcal{M}_i obtained from g . Either half-infinite edge of \bar{c}^i cuts away an unbounded convex polygon from a separate region of \mathcal{P}_i incident to \bar{g} ; either such polygon will retain \bar{g} in its boundary. Any finite edge of \bar{c}^i cuts away a bounded convex polygon from some region of \mathcal{P}_i , as in the previous case. Following a similar reasoning as above, we conclude that \bar{c}^i bounds an unbounded convex polygon.

Let Q_i denote the convex region bounded by \bar{c}^i (see Fig. 6). We shall refer to Q_i as to a convex polygon, either bounded or unbounded.

Now construct a tree-like structure \bar{M}_i from the tree M_i as follows. For each j , $1 \leq j \leq m_i$, consider the edge $g_j = (u_j, x_j)$ of M_i shared by $C(l(c_j^i))$ and $C(f(c_{j+1}^i))$; assume that it is traversed from u_j to x_j when walking counter-clockwise along the boundary of $C(l(c_j^i))$. Consequently, u_j is necessarily a leaf of M_i , which corresponds to an inner node of $S(P)$. Prolong g_j beyond u_j until the point p_{ij}^l , if p_{ij}^l is finite, and to infinity, otherwise. The reasoning carried out above implies correctness of the proposed construction.

Lemma 7. \bar{M}_i is the medial axis for Q_i .

Proof. Since Q_i is a convex polygon, its medial axis coincides with its Voronoi diagram. Therefore, it is sufficient to demonstrate that \bar{M}_i partitions the interior of Q_i into the Voronoi cells of its edges.

By construction, \bar{M}_i partitions Q_i into convex polygonal faces, each adjacent to a separate edge of Q_i . For an edge e of Q_i , let $F(e)$ denote its adjacent face of the partition.

Suppose for contradiction that for some edge e of Q_i , the face $F(e)$ contains a point x , for which the closest edge of Q_i is e' other than e . Let us drop a perpendicular from x onto e' ; note that its foot a must fall inside e' . Since x lies outside $F(e')$, the segment xa intersects the boundary $\partial F(e')$ of $F(e')$ at some point z belonging to \bar{M}_i . Such point is unique by convexity of $F(e')$.

Recall that $\partial F(e')$ is composed of e' and of pieces of bisectors between e' and some other edges of Q_i . If z is interior to some edge g of \bar{M}_i , let e'' denote the

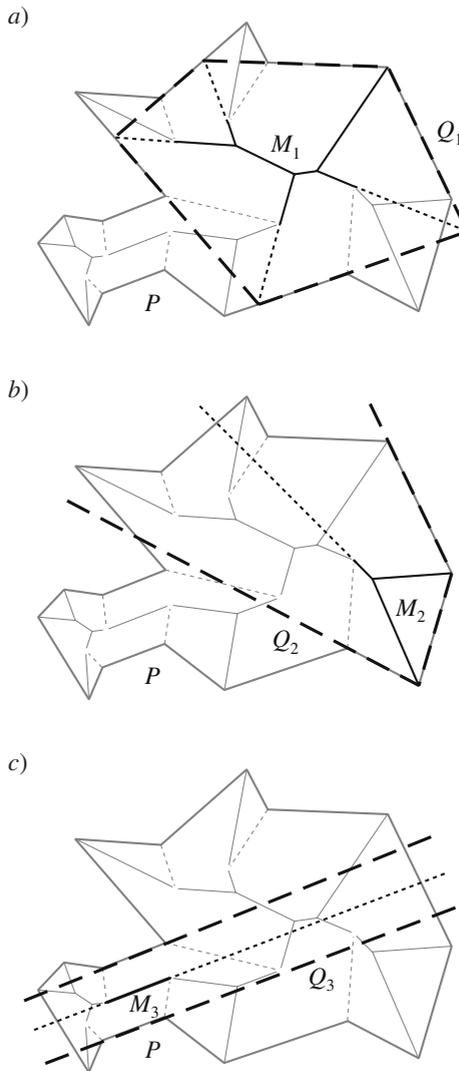


Fig. 6. A simple polygon P (bold gray) and its straight skeleton (gray); the edges of the latter incident to the reflex vertices of the former are depicted dotted gray. The convex region Q_i bounded by \bar{c}^i (dashed) can be of one of the three types: a) a convex polygon; b) an infinite convex region bounded by a chain, the first and the last edges of which are infinite; c) an infinite strip bounded by two parallel lines. For any Q_i , the corresponding subtree M_i of $S(P)$ is shown bold, and the parts of the edges of $M(Q_i)$ not belonging to M_i are shown dotted, where $1 \leq i \leq 3$.

edge of Q_i , such that the bisector of e' and e'' contains g . In particular, e'' and e may be the same edge. In case z coincides with a node u of \overline{M}_i , any edge other than e' , such that its adjacent face is incident to u , can be chosen as e'' .

Next, let us drop a perpendicular from z onto e'' . Lemma 3 together with our construction imply that its foot b will fall inside e'' . But since, by triangle inequality, $|xb| < |xz + zb| = |xz + za| = |xa|$, it follows that x is closer to e'' than to e' , which contradicts our assumption.

Corollary 4. *M_i is part of $M(Q_i)$, and can be obtained from the latter by appropriately trimming its edges incident to the vertices of Q_i not being those of P , and the unbounded ones, if any exist.*

It is easy to see that each M_i is a *maximal* fragment of a medial axis, in a sense that it cannot be extended along the edges of $S(P)$ while remaining a part of the medial axis for any polygon.

Given P and $S(P)$, and assuming that the representation of the latter provides information on the partition of P induced by $S(P)$, it is straightforward to decompose $S^c(P)$ into the subtrees M_1, \dots, M_k , and to retrieve the corresponding sets of chains $\mathcal{C}_1, \dots, \mathcal{C}_k$. For any i , $1 \leq i \leq k$, the convex polygon Q_i can then be constructed from \mathcal{C}_i following the procedure described above.

We summarize our results in the next theorem.

Theorem 1. *Let P be a simple polygon. The subgraph $S^c(P)$ of the straight skeleton $S(P)$ for P , traced out by the convex vertices of the linear wavefront, can be uniquely partitioned into a set of maximal fragments of medial axes. Each of those fragments represents a pruned medial axis for a certain convex polygon. Both the partition and the corresponding set of convex polygons can be computed from P and $S(P)$ in linear time.*

4 General Cases

In this section, we generalize our results to the case of polygons with holes, and further to the case of planar straight line graphs. To this end, we demonstrate that in either case, by extracting from the straight skeleton its subgraph traced out by the convex vertices of the wavefront, and splitting it at the nodes, at which the reflex wavefront vertices vanish, we again obtain a forest. Subsequently, to any connected component of the forest, which is a tree, the previously developed reasoning fully applies. Special attention is paid to the unbounded edges of the straight skeleton for a planar straight line graph.

The only remark to be made is that in general cases, a split event causes either a break of a linear wavefront component into two (if the edge and the reflex vertex involved in the event belong to the same connected component of the wavefront), or a merge of two wavefront components (if the respective edge and vertex belong to different components). Similarly, a vertex event can result in merging two wavefront components. Though in the case of simple polygons, different wavefront components can never merge, this distinction does not affect the above discussion.

4.1 Polygons with Holes

Let P be a polygon with holes; let $S(P)$ denote the straight skeleton of P . Consider the subgraph $S^c(P)$ of $S(P)$ traced out by the convex vertices of the linear wavefront. Having split $S^c(P)$ at the inner nodes, at which the reflex wavefront vertices vanish, we obtain a partition of $S^c(P)$ into a number of connected components M_1, \dots, M_k .

Lemma 8. *For any $i, 1 \leq i \leq k, M_i$ is a tree.*

Proof. Suppose for contradiction that for some $i, 1 \leq i \leq k, M_i$ contains cycles. Let c denote any cycle of M_i ; note that c bounds a simple polygon P_c formed as a union of the cells and of the holes of P that lie inside c . In particular, the generative edge of any cell being part of P_c lies inside c . Thus, at least one connected component of ∂P falls inside c .

Any vertex of c is incident to a cell lying inside c , and to a cell lying outside c . It follows that the first generated vertex u of c must have appeared as a result of interaction between two different connected components of the linear wavefront, one emanating from inside c , and the other – from outside c . But such interaction must have involved a reflex wavefront vertex. Therefore, c must have been split at u , which is a contradiction.

4.2 Planar Straight Line Graphs

Let G be a planar straight line graph; let $S(G)$ denote the straight skeleton for G . It follows from the definition of $S(G)$ that for any face F of G , the restriction of $S(G)$ to the interior of F is the straight skeleton $S(F)$ for F . Consequently, $S(G)$ can be viewed as a union of the straight skeletons for all the faces of G .

Any bounded face of G is a polygon, either simple or with holes. Thus, for each bounded face, its straight skeleton can be processed by means of the technique developed above. Let F_∞ denote the unbounded face of G . Note that $S(F_\infty)$ contains unbounded edges, any of which is incident to one finite and one infinite node. Consider the subgraph $S^c(F_\infty)$ of $S(F_\infty)$ traced out by the convex vertices of the linear wavefront, and split $S^c(F_\infty)$ at the inner nodes, at which the reflex wavefront vertices vanish. As a result, we obtain a number of connected subgraphs M_1, \dots, M_k of $S^c(F_\infty)$. Absence of cycles in any M_i , where $1 \leq i \leq k$, can be demonstrated in the same way as in the case of polygons with holes; therefore, each M_i can be handled as described in Section 3, unless it contains an unbounded edge.

Lemma 9. *M_i contains at most one unbounded edge.*

Proof. Observe that an unbounded edge of M_i must be traced out towards the infinite node, and apply a similar reasoning as in the last part of the proof of Lemma 2.

However, the case when M_i contains an unbounded edge is fully similar to that when all the edges of M_i are finite, and one of its edges incident to a leaf is

traced out towards the leaf. The only difference is that at the very last step, there is no need to clip the infinite edge of $M(Q_i)$ in order to obtain M_i .

Thus, we can generalize Theorem 1 as follows.

Theorem 2. *Let G be a planar straight line graph. The subgraph $S^c(G)$ of the straight skeleton $S(G)$ for G , traced out by the convex vertices of the linear wavefront, can be uniquely partitioned into a set of maximal fragments of medial axes. Each of those fragments represents a pruned medial axis for a certain convex polygon. Both the partition and the corresponding set of convex polygons can be computed from G and $S(G)$ in linear time.*

5 Conclusion

The principal objective of this research was to enhance understanding of the geometry of the straight skeleton. We restricted our attention to its subgraph traced out by the convex vertices of the linear wavefront, and claimed that by splitting it at the nodes, at which the reflex vertices of the wavefront vanish, we would obtain a set of pruned medial axis for certain convex polygons. Moreover, any subgraph we thus get is a maximal fragment of a medial axis embedded in the straight skeleton, in a sense that it cannot be extended along the straight skeleton edges, while remaining a piece of the medial axis for any polygon. Finally, we pointed out that, given a polygon or, more generally, a planar straight line graph, and its straight skeleton, we can easily retrieve this set of pruned medial axes along with the respective convex polygons in total linear time.

An interesting development of our work would be to speed up the computation of the straight skeleton by exploiting structural properties of the latter described and analyzed in this paper.

Acknowledgement

This research was supported by Russian Foundation for Basic Research (grant 07-07-00268-a).

References

1. Aichholzer, O., Aurenhammer, F.: Straight skeletons for general polygonal figures. In: Cai, J.-Y., Wong, C.K. (eds.) COCOON 1996. LNCS, vol. 1090, pp. 117–126. Springer, Heidelberg (1996)
2. Aichholzer, O., Aurenhammer, F., Alberts, D., Gärtner, B.: A novel type of skeleton for polygons. *J. Univ. Comp. Sci.* 1, 752–761 (1995)
3. Barequet, G., Goodrich, M.T., Levi-Steiner, A., Steiner, D.: Contour interpolation by straight skeletons. *Graphical Models (GM)* 66(4), 245–260 (2004)
4. Cheng, S.W., Vigneron, A.: Motorcycle graphs and straight skeletons. *Algorithmica* 47(2), 159–182 (2007)
5. Chin, F., Snoeyink, J., Wang, C.: Finding the medial axis of a simple polygon in linear time. *Discr. Comp. Geom.* 21(3), 405–420 (1999)

6. Demaine, E.D., Demaine, M.L., Lubiw, A.: Folding and cutting paper. In: Akiyama, J., Kano, M., Urabe, M. (eds.) JCDCG 1998. LNCS, vol. 1763, pp. 104–118. Springer, Heidelberg (2000)
7. Eppstein, D., Erickson, J.: Raising roofs, crashing cycles, and playing pool: applications of a data structure for finding pairwise interactions. *Discr. Comp. Geom.* 22(4), 569–592 (1999)
8. Tănase, M.: Shape decomposition and retrieval. Ph.D. Thesis, Utrecht Univ. (2005)
9. Tănase, M., Veltkamp, R.C.: Straight skeleton approximating the medial axis. In: *Proc. 12th Annu. European Symp. on Algorithms*, pp. 809–821 (2004)
10. Trofimov, V., Vyatkina, K.: Linear axis for general polygons: properties and computation. In: Gervasi, O., Gavrilova, M.L. (eds.) ICCSA 2007, Part I. LNCS, vol. 4705, pp. 122–135. Springer, Heidelberg (2007)
11. Vyatkina, K.: A Search for Medial Axes in Straight Skeletons. In: *Proc. 24th European Workshop on Comp. Geom., Nancy, France*, pp. 157–160 (2008)
12. Vyatkina, K.: Linear Axis for Planar Straight Line Graphs. In: *Proc. 15th Computing: Australasian Theory Symp., CRPIT 1994*, pp. 137–150 (2009)
13. Yap, C.K.: An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments. *Discr. Comp. Geom.* 2, 365–393 (1987)