Linear Axis for Planar Straight Line Graphs

Kira Vyatkina

Department of Mathematics and Mechanics
Saint Petersburg State University,
28 Universitetsky pr., Stary Peterhof,
Saint Petersburg 198504, Russia,
Email: kira@meta.math.spbu.ru

Abstract

A linear axis is a straight line skeleton for a polygonal shape. The concept of a linear axis ε-equivalent to the medial axis has been introduced and studied for simple polygons and for those with holes. In this paper, we generalize the notions of a linear axis and of ε-equivalence to the case of planar straight line graphs. We show that for some graphs, a linear axis ε-equivalent to the medial axis does not exist, for any $\varepsilon > 0$. However, if the graph vertices are in general position, a sought linear axis does exist for any $\varepsilon > 0$, and can be computed in $O(n \log n)$ time in the absence of certain correlations in the graph structure.

Keywords: Linear axis, medial axis, straight skeleton.

1 Introduction

Skeletons for planar shapes have been studied for over forty years, and remain attractive for researches, – in particular, due to their notably wide applicability in such areas as computer vision, image processing, shape retrieval, and many others. For an important case of polygonal figures, there exist three types of skeletons: a medial axis, a straight skeleton, and a linear axis.

For a simple polygon P, its *medial axis* $M(P)$ is a locus of the centers of the maximal inscribed discs. It is closely related to the Voronoi diagram $Vor(P)$ of P: in case P is convex, the two structures are identical; otherwise, $M(P)$ can be obtained from $Vor(P)$ by discarding the edges of the latter incident to the reflex vertices of P (Fig. 1a). (A vertex v of P is called *reflex* if the interior angle of P at v is greater than π; however, throughout our reasoning, we shall need to treat *degenerate* vertices with internal/external angle of π as reflex ones.)

Alternatively, the medial axis can be defined in terms of the wavefront propagation model. A *uniform wavefront* initially coincides with the boundary ∂P of P; at time $t > 0$, it consists of the inner points of P at distance t from ∂P (Fig. 1a). Thus, during the propagation, the uniform wavefront moves inside the polygon, and finally vanishes; the medial axis is thereby traced out by the wavefront vertices.

The medial axis is generally recognized as a useful shape descriptor, because of its capability to capture key visual features of a shape. But in presence of reflex vertices in P, its medial axis contains parabolic arcs, which sometimes is regarded as a disadvantage from the practical point of view.

A *straight skeleton* was introduced by Aichholzer et. al (1995). As suggested by the name, all its edges are straight line segments. The straight skeleton $S(P)$ for P results from a *linear wavefront* propagation. As in the above case, at start of the propagation, the linear wavefront is identical to ∂P. In the process, the wavefront edges move inside the polygon at equal speed, all along remaining parallel to themselves, and the straight skeleton is traced out by the wavefront vertices (Fig. 1b).

A *linear axis* is due to Tănase & Veltkamp (2004). The main idea behind is to insert at some reflex vertices of P a few *hidden edges* of zero length. Those edges are oriented so that for each reflex vertex v of P, the internal angles at all the vertices of the modified polygon P', which reside at the same point as v, are equal. Then the linear wavefront propagation is applied to P', and from the obtained straight skeleton $S(P')$, the edges incident to the reflex vertices of P' are removed (Fig. 1c). (This definition slightly differs from the one originally given by Tănase & Veltkamp (2004); see the next Section for details.) The resulting graph is the *linear axis* $L^k(P)$ for P, which corresponds to the sequence of hidden edges $k = (k_1, \ldots, k_r)$ inserted at the reflex vertices $v_1, \ldots, v_r$ of P, where $r \geq 0$ is the number of the reflex vertices of P, and $k_i \geq 0$, for $1 \leq i \leq r$. The larger are k_i, for $1 \leq i \leq r$, the better the linear wavefront for P' approximates the uniform wavefront for P at the same moment. Thus, the linear axis can approximate arbitrarily well the medial axis, while all the edges of the former are straight line segments. The topological quality of the approximation is captured by the notion of ε-equivalence (Tănase 2005), and can be controlled by choosing an appropriate value of the underlying parameter $\varepsilon > 0$.

The concept of the medial axis applies to arbitrary planar and higher-dimensional shapes. The notion of the straight skeleton was extended to the case of planar straight line graphs by Aichholzer & Aurenhammer (1996), and very recently – to the case of three-dimensional polyhedra by Barequet et. al (2008). The linear axis for polygons with holes has been studied by Trofimov & Vyatkina (2007).

In this work, we perform the next step in exploring the linear axis, and address the case of planar straight line graphs. We show that, in contrast to polygons, there are graphs, for which a linear axis ε-equivalent for the medial axis does not exist, for any $\varepsilon > 0$. However, this can occur only if the graph vertices are not in general position; if they are – that is, no three vertices are collinear, – the desired linear axis does exist for any $\varepsilon > 0$, and can be computed in $O(n \log n)$, assuming absence of certain correlations in the graph structure, which will be formally described later. We explain how to extend the algorithm han-

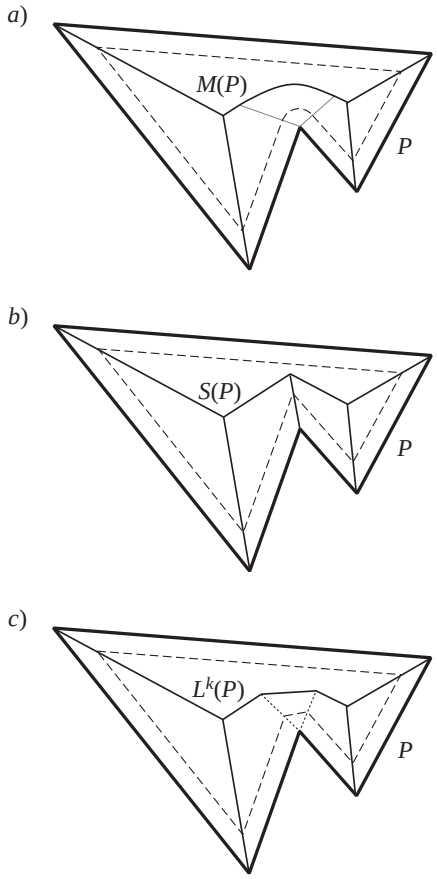

Figure 1: A simple polygon P is depicted bold. a) The medial axis $M(P)$ is shown solid; the two edges of the Voronoi diagram $Vor(P)$, which are not part of $M(P)$, are indicated gray; the uniform wavefront soon after the propagation starts is depicted dashed. b) The straight skeleton $S(P)$ is shown solid; the linear wavefront soon after the propagation starts is depicted dashed. c) One hidden edge is inserted at the reflex vertex of P; the corresponding linear axis $L^k(P)$ is shown solid, and the respective linear wavefront soon after the propagation starts is depicted dashed. The traces of the reflex wavefront vertices, which are not part of the linear axis, are marked dotted.

dling the polygons to work with the unbounded face of a graph.

In the next Section, we introduce the notions involved in our reasoning, and establish properties of a linear axis for the unbounded face. Section 3 addresses achievability of ε-equivalence: first, we prove existence of a linear axis ε-equivalent to the medial axis for any polygon; then, we state a necessary and sufficient condition of its existence for a planar straight line graph, and finally, we provide an example of such graph, for which any linear axis is not ε-equivalent to the medial axis, for any $\varepsilon > 0$. For the graphs, the vertices of which are in general position, computation of a linear axis ε-equivalent to the medial axis is discussed in Section 4. We conclude with a few remarks on the developed algorithm.

2 Basic Definitions and Properties

2.1 Event Taxonomy

Let us briefly recall how the linear propagation flows. For simplicity of exposition, let us consider the process of constructing the straight skeleton $S(P)$ for a simple polygon P; the same arguments can be easily applied to other cases.

Initially, the linear wavefront coincides with the boundary of P; throughout the propagation, the wavefront edges move inside at a unit speed, thereby remaining parallel to themselves. In the process, three types of events may occur, at which the wavefront structure changes.

At an *edge event*, a wavefront edge shrinks to zero. At a *split event*, a reflex vertex collides with an edge, thereby splitting a wavefront component into two, and any of those continues shrinking in the same manner. At a *vertex event*, two reflex vertices meet at the same point, thereby giving rise to a new reflex vertex in the wavefront.

This classification is very much alike to the one proposed by Eppstein & Erickson (1999). The only distinction is how a vertex event is introduced. Eppstein & Erickson (1999) defined it as a meeting of two or more reflex vertices, *and nothing else*, at the same point, and claimed that unlike edge or split events, a vertex event could lead to appearance of a new reflex vertex in the wavefront. However, we point out that a new reflex vertex can also emerge at two simultaneous edge events, when two wavefront edges adjacent at a convex vertex, and each incident to a separate reflex vertex, vanish at once (Fig 2). Thus, at that very moment, two reflex vertices meet, and two edges annihilate at the same point.

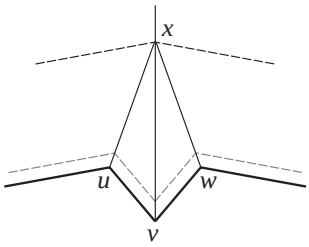

Figure 2: A portion of the linear wavefront at the start of the propagation (bold), soon after the start (gray dashed), and at the moment when the edges uv and vw simultaneously vanish (black dashed). When those edges annihilate, the node x of the straight skeleton is generated, and a new reflex vertex appears in the wavefront.

The second important property of a vertex event highlighted by Eppstein & Erickson (1999) is that it leads to appearance of a node in the straight skeleton, which has degree at least four, and cannot be replaced by several nodes of degree three using standard perturbation techniques. The same will hold, however, for a complex event illustrated in Fig. 2, what provides an extra argument for the former to be classified as a vertex event. On the other hand, if two or more reflex vertices meet at the same point, but no new reflex vertex is thereby created in the wavefront, we can always interpret such event as a sequence of edge and split events that occur at the same place, with a zero time interval between any two consecutive ones.

Having taken into account all these observations, we modified the definition of a vertex event as given above.

Annihilation of a wavefront component comprises at least three edge events that occur simultaneously at the same place. The propagation process continues until all the wavefront components vanish; the vertices of the wavefront thereby trace out the edges of $S(P)$. Figure 3 illustrates different types of events.

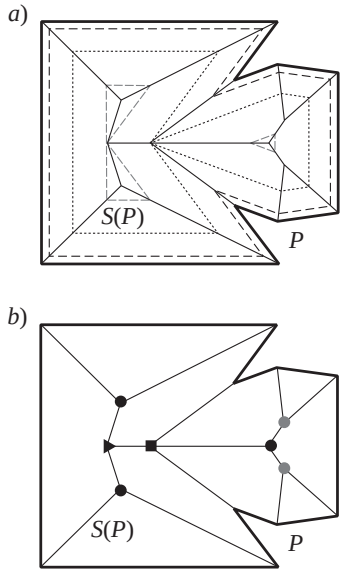

Figure 3: A simple polygon P and its straight skeleton $S(P)$. During the linear wavefront propagation, one vertex event, one split event, and several edge events occur. a) The linear wavefront is depicted at three different times: soon after the propagation starts (dashed), when the vertex event occurs (dotted), and when the split event occurs (dashed gray). b) For each node of $S(P)$, the type of event that produced it is indicated: square – the vertex event, triangle – the split event, gray circle – an edge event, at which a single wavefront edge vanishes, black circle – three simultaneous edge events, at which a wavefront component annihilates.

2.2 Linear Axis

Let G be a planar straight line graph without isolated vertices. Denote by $\mathcal{F}$ the set of faces of the planar subdivision induced by G.

Any bounded face $F \in \mathcal{F}$ can be interpreted as a polygon, in the boundary ∂F of which, certain degeneracies may occur: a connected component of ∂F can contain coinciding edges, and edges of zero length. The connected components of ∂F are in one-to-one correspondence with the boundary cycles of F.

To obtain a connected component of ∂F from a boundary cycle β of F, we first walk around β, thereby constructing a closed polyline P_β in the plane. If β is not simple, then P_β contains at least one pair of identical edges. Any such pair corresponds to a bridge in G, which is incident to a single face, – and namely, to F. If such bridge b is incident to a hanging node, then the two edges of P_β corresponding to b are found adjacent in P_β; in this case, we insert between them a zero-length edge orthogonal to b (Fig. 4). Having thus eliminated all the pairs of adjacent identical edges in P_β (if any), we get the required connected component of ∂F. In particular, an isolated edge of G will produce a degenerate rectangle bounding a (degenerate) hole in the surrounding face.

Subsequently, the definition of the linear axis for polygons can be applied to define the linear axis $L^k(F)$ for F; here $k = (k_1, \ldots, k_r)$ denotes the sequence of hidden edges inserted at the reflex vertices $v_1, \ldots, v_r$ of F, and $r \geq 0$ is the number of the reflex vertices of F under the above convention. Let ∂F^k denote the polygon obtained from ∂F after insertion of the hidden edges.

Definition 1. The linear axis $L^k(F)$, corresponding to the sequence k of hidden edges, is obtained from the straight skeleton $S(\partial F^k)$ by removing from it the edges incident to the reflex vertices of ∂F^k.

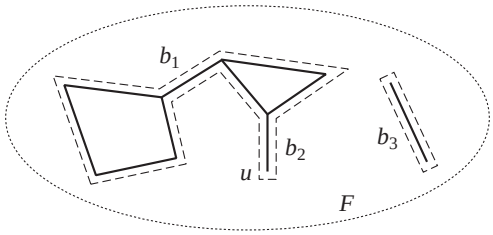

Figure 4: Interpretation (dashed) of the two inner boundary cycles of F. Two connected components of G, which give rise to those boundary cycles, are depicted bold. There are three bridges b_1, b_2, and b_3; b_2 is incident to a hanging node u, and b_3 is an isolated edge.

Tănase & Veltkamp (2004) initially defined the linear axis to be the trace of the convex vertices of the wavefront emanating from ∂F^k, obtained during the linear wavefront propagation. Besides, they established a convention stating that at least one hidden edge must be inserted at any reflex vertex of the initial wavefront, which has an internal angle of at least $3\pi/2$; the goal was to thereby assure the connectivity of the linear axis. Indeed, this would prevent us from appearance of new reflex vertices in the wavefront at vertex events as defined by Eppstein & Erickson (1999); consequently, it was assumed that any edge of $S(\partial F^k)$ traced out by a reflex wavefront vertex, would be incident to a reflex vertex of ∂F^k. However, in this way we cannot avoid vertex events of the kind illustrated in Fig. 2, and thus, some edges of $S(\partial F^k)$ delineated by reflex wavefront vertices still may be incident to two inner nodes of $S(\partial F^k)$. Deletion of such edges would cause the linear axis to be disconnected.

Because of that, we alter the definition of the linear axis, and explicitly require the removal from the underlying straight skeleton $S(\partial F^k)$ of the edges incident to the reflex vertices of ∂F^k. Subsequently, we do not need to request insertion of hidden edges at the vertices of the initial wavefront having the internal angles of at least $3\pi/2$.

It is straightforward to extend the above definition to the case of the unbounded face $F_\infty \in \mathcal{F}$.

To define a linear axis $L^k(G)$ for G, we need first to specify a sequence k_F of hidden edges for each face $F \in \mathcal{F}$.

Definition 2. $L^k(G) = \cup_{F \in \mathcal{F}} L^{k_F}(F)$, where k is a sequence of hidden edges formed by concatenating the sequences k_F, for all $F \in \mathcal{F}$, appropriately matched to the set of reflex vertices of all the faces $F \in \mathcal{F}$ (Fig. 5).

The properties of the linear axis for polygons have been studied by Tănase & Veltkamp (2004), and by Trofimov & Vyatkina (2007); they will obviously hold for the case of the bounded faces of G.

Let us now consider a linear axis $L^{k_\infty}(F_\infty)$ and the medial axis $M(F_\infty)$ for the unbounded face F_∞, where k_∞ is a sequence of hidden edges associated with F_∞. The edges and the reflex vertices of F_∞ will be referred to as *sites*; with each site S, we associate a portion of the linear and of the uniform wavefront emanating from it, which we call the *linear* and the *uniform offset* of S, respectively. For an edge e and its incident reflex vertex v, their offsets are separated by the wavefront point lying on the perpendicular to e at v (Fig. 6). The regions swept by the linear and by the uniform offset of S during the propagation

are called the *linear* and the *Voronoi cell* of S, respectively, and are denoted by $LC(S)$ and by $VC(S)$, respectively (Fig. 5).

As in the bounded case, the following property of the cells will hold.

Lemma 1. *For any site S, the cells $LC(S)$ and $VC(S)$ are connected and simply connected.*

Proof. During the propagation, the wavefront edges move continuously; moreover, they can either split or vanish, but neither merge nor reappear. These observations imply the claim. □

Observe that $L^{k\infty}(F_\infty)$ and $M(F_\infty)$ can contain half-infinite edges; we shall interpret any such edge as being incident to one finite and one infinite node.

Moreover, a linear axis for the unbounded face may consist of several connected components, regardless of the number of hidden edges associated with each reflex vertex. However, the following property will hold.

Lemma 2. *Any connected component of $L^{k\infty}(F_\infty)$ contains at least one half-infinite edge.*

Proof. Suppose for contradiction that some connected component L of $L^{k\infty}(F_\infty)$ has only finite edges. In case L contain cycles, the latter bound a number of simple polygons in the plane; let us denote those polygons by $P_1, \ldots, P_h$, where $h \geq 0$.

Let $\mathcal{S}_L$ denote the set of all the sites, the linear cells of which are incident to L and lie outside any cycle of L. Observe that for any site $S \in \mathcal{S}_L$, the edges of $LC(S)$ that belong to L represent a path in L. If $\partial LC(S)$ contains an edge e incident to a reflex vertex of ∂F_∞, then this vertex is either S itself, or a site S' incident to S (which is then an edge). In the latter case, S' necessarily belongs to $\mathcal{S}_L$, since the second node incident to e (and to $LC(S')$) must be contained in L.

Having glued together the linear cells of all the sites from $\mathcal{S}_L$ and all the polygons P_i, for $1 \leq i \leq h$, we shall obtain a simple polygon P_L, any edge of which is an edge of $LC(S)$ for some $S \in S_L$. From the above discussion, it follows any such edge is necessarily a site. But it means that L lies inside a bounded face of G, which is a contradiction. □

Lemma 3. *Inside each face of $L^{k\infty}(F_\infty)$, there lies a separate connected component of ∂F_∞.*

Proof. Any face f of $L^{k\infty}(F_\infty)$ represents a union of a number of linear cells and, possibly, some bounded faces of G. If for some site S, $LC(S) \subset f$, then, in particular, $S \subset f$. Moreover, ∂f cannot cross ∂F_∞ (though can touch the latter at its convex vertices; see Fig. 5a). We conclude that inside f, there must lie at least one connected component of ∂F_∞.

Now suppose for contradiction that some face f of $L^{k\infty}(F_\infty)$ contains inside at least two connected components of ∂F_∞; let us denote those components by $L_1, L_2, \ldots, L_k$, where $k \geq 2$.

Let us glue together the linear cells of all the sites contained in L_1, and all the simple polygons bounded by the cycles of L_1 (if any). As a result, we obtain a connected polygonal domain P_1 (possibly multiconnected and/or unbounded). For any $S \subset f$, also $LC(S) \subset f$; therefore, $P_1 \subset f$. Moreover, $f \setminus P_1 \neq \emptyset$, since for any site $S' \subset L_j$, for some j, $2 \leq j \leq k$, we have $LC(S') \subset f$ and $LC(S') \nsubseteq P_1$. Finally, observe that any edge g of P_1 separates the linear cells of some two sites, one belonging, and the other one not belonging to L_1, and thus, g is contained in $L^{k\infty}(F_\infty)$.

Consider any point $p \in P_1$, and any point $q \in f \setminus P_1$. Since $p, q \in f$, there must exist a path from p to q, which entirely lies inside f. However, any such path should leave P_1 at some moment, thereby intersecting an edge of P_1, which is necessarily an edge of $L^{k\infty}(F_\infty)$. This implies a contradiction, and the claim follows. □

The same property holds for the medial axis $M(F_\infty)$; this can be proved analogously.

Finally, we give the following characterization to the shape of the linear cells.

Lemma 4. *For any edge S_i, $LC(S_i)$ is a polygon monotone with respect to the line through S_i. For any reflex vertex S_j, $LC(S_j)$ is a star-shaped polygon having S_j in its kernel.*

The Voronoi cells have similar properties, though their boundaries can comprise parabolic segments. Correctness of the both statements follows from analysis of the wavefront propagation process.

2.3 ε-Equivalence

In this subsection, we adapt the terminology introduced by Tănase & Veltkamp (2004), and by Trofimov & Vyatkina (2007), to the case of the unbounded face F_∞ of G.

2.3.1 ε-Clusters

Let ε be a positive real constant.

Definition 3. Let v_i and v_j be two adjacent nodes of $M(F_\infty)$ generated by triples of sites S_k, S_i, S_l, and S_k, S_j, S_l, respectively. The Voronoi edge between v_i and v_j is an *ε-edge* if $d(v_i, S_j) < (1+\varepsilon)d(v_i, S_i)$ or $d(v_j, S_i) < (1+\varepsilon)d(v_j, S_j)$ (Fig. 7).

A Voronoi edge that is not an ε-edge is called a *non-ε-edge*. In particular, any Voronoi edges incident to a hanging node is a non-ε-edge.

Definition 4. A path between two nodes of $M(F_\infty)$ is an *ε-path* if it consists only of ε-edges.

Definition 5. Let v and w be two nodes of $M(F_\infty)$. We say that w is an *ε-neighbour* of v, if v and w are connected by an ε-path.

Let $N_\varepsilon(v)$ denote the set of all ε-neighbours of v.

Definition 6. For any node v of $M(F_\infty)$, the set $C(v) = \{v\} \cup N_\varepsilon(v)$ of nodes of $M(F_\infty)$ is called an *ε-cluster* of v.

Observe that any hanging node of $M(F_\infty)$ constitutes a separate ε-cluster, for any $\varepsilon > 0$.

2.3.2 Geometric Graphs

Definition 7. A *geometric graph* (V, E) is a set in $\mathbb{R}^2$ that consists of a finite set V of vertices, any of which is either a finite point or an artificial point at infinity, and a finite set E of arcs being mutually disjoint, simple curves, any of which connects two vertices of V.

Let us associate with the linear axis $L^{k\infty}(F_\infty)$ a geometric graph $(V_{L^{k\infty}}, E_{L^{k\infty}})$ as follows. Any vertex from $V_{L^{k\infty}}$ corresponds to a node of $L^{k\infty}(F_\infty)$ having degree either one or at least three. Any arc from $E_{L^{k\infty}}$ corresponds to a path p in $L^{k\infty}(F_\infty)$ between two nodes, neither of which has degree two, while any inner node of p has degree two.

In other words, when passing from the linear axis to its geometric graph, we eliminate topologically less

significant nodes of degree two of the former, by iteratively merging the two incident edges of any such node into a single one.

In particular, the finite hanging vertices of $(V_{L^{k\infty}}, E_{L^{k\infty}})$ reside at the convex vertices of F_∞, and the infinite ones are matched to the infinite nodes of $L^{k\infty}(F_\infty)$. Any vertex from $V_{L^{k\infty}}$ has the same degree in the geometric graph as its counterpart – in $L^{k\infty}(F_\infty)$.

The geometric graph (V_M, E_M) of the medial axis $M(F_\infty)$ is defined analogously.

For any two vertices $v_1, v_2 \in V_M$, and for any $\varepsilon > 0$, we shall say that v_1 and v_2 are connected by an ε-path, if the corresponding statement holds for the counterparts of v_1 and v_2 among the nodes of $M(F_\infty)$. Subsequently, the notions of an ε-neighbor and an ε-cluster can be applied to (V_M, E_M). Whenever this does not lead to ambiguities, we shall use the same notation for those terms relative to (V_M, E_M) as introduced for $M(F_\infty)$ in the previous paragraph.

2.3.3 Matching the Axes

In the further, we shall interpret any node of degree $d \geq 4$ of either axis 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 (Fig. 8). The same convention applies to the vertices of the geometric graphs.

Definition 8. $M(F_\infty)$ and $L^{k\infty}(F_\infty)$ are *ε-equivalent* if there exists a bijection $f : V_M \to V_{L^{k\infty}}$ such that:

1) $u \in V_M$ is finite $\Leftrightarrow$ $f(u)$ is finite;

2) $u \in V_M$ corresponds to a convex vertex p of F_∞ $\Leftrightarrow$ $f(u)$ corresponds to p;

3) $\forall v_i, v_j \in V_M$ with $v_j \notin N_\varepsilon(v_i)$, $\exists$ an arc in E_M connecting v_i and v_j $\Leftrightarrow$ $\exists$ an arc in $E_{L^{k\infty}}$ connecting $f(v'_i)$ and $f(v'_j)$, where $v'_i \in C(v_i)$ and $v'_j \in C(v_j)$.

In particular, condition 1 requires ε-equivalent axes $M(F_\infty)$ and $L^{k\infty}(F_\infty)$ to have the same number of infinite nodes. Later we shall demonstrate that in this case, $M(F_\infty)$ and $L^{k\infty}(F_\infty)$ will also have the same numbers of nodes of degree three, and thus, one can indeed ask for a bijection between the vertices of the geometric graphs.

Suppose that $M(F_\infty)$ and $L^{k\infty}(F_\infty)$ are ε-equivalent. Let us collapse ε-clusters in (V_M, E_M), and glue together their images under f in $(V_{L^{k\infty}}, E_{L^{k\infty}})$; the resulting graphs will be isomorphic. Therefore, a demand of ε-equivalence between the axes may be viewed as a relaxation of a requirement of isomorphism between their geometric graphs.

The definition of ε-equivalence between the axes for any bounded face of G can be obtained from the above one by dropping condition 1, which will then trivially hold.

Definition 9. $M(G)$ and $L^k(G)$ are *ε-equivalent* if for each $F \in \mathcal{F}$, $M(F)$ and $L^{k_F}(F)$ are ε-equivalent.

3 Attainability of ε-Equivalence

Let us start with the following Lemma, which assures consistence of the definition of ε-equivalence.

Lemma 5. *If $L^{k\infty}(F_\infty)$ and $M(F_\infty)$ have the same number of infinite nodes, then they also have the same number of nodes of degree three.*

Proof. Let us denote by n_{inf} the number of infinite nodes in either axis, by n_{cv} – the number of convex vertices of F_∞, and by f – the number of connected components of ∂F_∞. Recall that n_{cv} and f equal the number of finite hanging nodes and the number of faces in either axis, respectively (see Lemma 3).

Consider the linear axis $L^{k\infty}(F_\infty)$; denote by n_2^L, n_3^L, and n_e^L the number of its nodes having degree two and three, and the number of its edges, respectively. Let us glue together all the infinite nodes of $L^{k\infty}(F_\infty)$, thereby introducing a single artificial node. The resulting graph L^* is planar; it has $n_{cv} + n_2^L + n_3^L + 1$ nodes, n_e^L edges, and f faces. Applying to L^* Euler formula, we get

$$(n_{cv} + n_2^L + n_3^L + 1) - n_e^L + f = 2.$$

On the other hand,

$$n_{cv} + 2n_2^L + 3n_3^L + n_{inf} = 2n_e^L.$$

Combining the two equations, we derive

$$n_3^L = n_{cv} - n_{inf} + 2f - 2.$$

Having applied a similar reasoning to $M(F_\infty)$, we shall obtain the same expression for the number of its nodes of degree three, which implies the claim. □

For a polygon, given any $\varepsilon > 0$, one can always compute a linear axis ε-equivalent to the medial axis (Tănase 2005, Tănase & Veltkamp 2004, Trofimov & Vyatkina 2007). Moreover, the opposite task can be easily completed, as it will be shown in the next Lemma.

Lemma 6. *Let P be a polygon. For any sequence k of hidden edges associated with P, there exists $\varepsilon_0 > 0$, such that for any $\varepsilon \geq \varepsilon_0$, $L^k(P)$ is ε-equivalent to $M(P)$.*

Proof. The geometric graphs (V_M, E_M) and (V_{L^k}, E_{L^k}) have the same number of hanging vertices, which equals the number of convex vertices of P. Consequently, they have the same number of vertices of degree three (this can be proved analogously to Lemma 5).

Suppose that for some $\varepsilon_0 > 0$, all the inner vertices of (V_M, E_M) fall into the same ε-cluster. Then a bijection f required by the definition of ε-equivalence can be obtained by arbitrarily mapping the inner vertices of (V_M, E_M) to distinct inner vertices of (V_{L^k}, E_{L^k}), and by mapping to each other the hanging vertices of either geometric graph, which correspond to the same convex vertex of P. Moreover, if we specify such ε_0, then the same property will hold for any $\varepsilon > \varepsilon_0$.

It remains to find such ε_0. To this end, for any inner node v of $M(P)$, consider the distance $d(v, P)$ from v to the boundary of P, and let $d_{\min} = \min_v d(v, P)$. Denote by D_P the diameter of P, and let $\varepsilon_0 = D_P/d_{min}$. For this choice of ε_0, and for any two adjacent inner nodes v_i and v_j of $M(P)$, (v_i, v_j) is an ε-edge, since $d(v_i, S_j) < D_P < d(v_i, S_i) + D_P \leq (1 + D_P/d_{min})d(v_i, S_i) = (1 + \varepsilon_0)d(v_i, S_i)$. Consequently, all the inner nodes of $M(P)$ form a single ε-cluster, and so do the inner vertices of (V_M, E_M). □

In the case of planar straight line graphs, a special attention should be paid to the half-infinite edges of the axes.

Lemma 7. *Let G be a planar straight line graph, and let k be a sequence of hidden edges associated with G. Then there exists $\varepsilon_0 > 0$, such that for any $\varepsilon \geq \varepsilon_0$, $L^k(G)$ is ε-equivalent to $M(G)$, if and only if $L^k(G)$ and $M(G)$ have the same number of infinite nodes.*

Proof. The "only if" case follows immediately from the definition of ε-equivalence. Let us now consider the "if" case.

For any bounded face $F \in \mathcal{F}$, let k_F denote the sequence of hidden edges induced by k; the corresponding sequence for the unbounded face F_∞ will be denoted by k_∞.

For any bounded face $F \in \mathcal{F}$, Lemma 6 implies existence of $\varepsilon_F > 0$, such that for any $\varepsilon \geq \varepsilon_F$, $L^{k_F}(F)$ is ε-equivalent to $M(F)$.

Now let us examine the unbounded face F_∞. For any inner node v of $M(F_\infty)$, consider the distance $d(v, G)$ from v to the graph G, and let $d_{\min} = \min_v d(v, G)$. Next, consider a circle that contains inside both G and all the finite nodes of $M(F_\infty)$, and let D_∞ denote its diameter. Finally, let $\varepsilon_{F_\infty} = D_\infty / d_{min}$. By a similar reasoning as given in the proof of Lemma 6, it can be shown that for any $\varepsilon \geq \varepsilon_{F_\infty}$, $L^{k_\infty}(F_\infty)$ is ε-equivalent to $M(F_\infty)$.

Let $\varepsilon_0 = \max_{F \in \mathcal{F}}$; this completes the proof. □

Before proceeding to the main result of this Section, let us prove a Lemma we shall need.

Lemma 8. *Let v be a reflex vertex of the linear wavefront, such that v lies strictly above its adjacent vertices u and w, and the edges (v, u) and (v, w) have finite slopes of different signs. Then the vertical projection of the speed of v is greater than 1.*

Proof. Let r_v^p and r_v^n denote a positively and a negatively directed horizontal rays with an endpoint at v, respectively. Without loss of generality, assume that u precedes v in the counterclockwise order of the wavefront vertices. Suppose that the angle α between (v, u) and r_v^n is smaller than the angle β between (v, w) and r_v^p (Fig. 9); the cases $\alpha > \beta$ and $\alpha = \beta$ are analogous. Then the bisector of the wavefront angle at v, along which v moves, has a positive slope.

Let v' denote the position, to which v will move in a unit time, assuming that it will not vanish earlier at any event. Then the speed of v equals $|vv'|$. Denote by y the projection of v' onto the vertical line through v; observe that y lies above the line l through v parallel to (v, u). Now it is sufficient to show that $|vy| > 1$.

Let z denote the intersection point of l and the segment vy. Note that the distance from v to l equals 1. Therefore, vz is a hypotenuse of a right triangle with a cathetus of length 1, which implies $|vz| > 1$. Consequently, $|vy| > |vz| > 1$, and the claim follows. □

Theorem 1. *There are planar straight line graphs, such that for any $\varepsilon > 0$, a linear axis ε-equivalent to the medial axis does not exist for them.*

Proof. To prove the claim, we present an example of a planar straight line graph G , such that the number of infinite nodes in any linear axis for its unbounded face F_∞ differs from that in $M(F_\infty)$; see Fig. 10a. Observe that $M(F_\infty)$ has two unbounded edges. We shall demonstrate that for any sequence k_∞ of hidden edges associated with F_∞, $L^{k_\infty}(F_\infty)$ has a single unbounded edge (Fig. 10b).

First, we point out that the linear offset of Y contains no horizontal edge, for any number k_Y of hidden edges inserted at Y. To see this, let us fix any $k_Y \geq 0$, and drop a perpendicular from Y on any edge e from the linear offset of the latter (Fig. 10c). Observe that it will make an angle $\theta_e = m/2(k_Y + 1)$ with the bisector ray of the interior angle of F_∞ at Y, for some integer m, $0 \leq m \leq k_Y + 1$. In particular, θ_e will always be rational. But for a horizontal edge e, it would have been equal to $\pi/10$, which is irrational. This contradiction justifies our assertion. By symmetry, the same property holds for Y'.

This implies that for any sequence k_∞, the linear offsets of Y and Y' will each have a unique topmost vertex, which we shall denote by y_0 and y_0', respectively (Fig. 10b). Note that both to y_0 and y_0', Lemma 8 applies; therefore, the vertical projection of the speed of either is greater than 1.

It follows that at any time $t > 0$, both y_0 and y_0' will lie above the linear offsets of the sites YU, UZ, Z, ZZ', Z', $Z'U'$ and $U'Y'$. (More precisely, this statement applies to those of the mentioned wavefront elements, which have not yet vanish.) Consequently, while both y_0 and y_0' exist, they split the wavefront into two parts that cannot interact. Since one of those parts initially contains no convex vertices, no edges of $L^{k_\infty}(F_\infty)$ can appear in the region it sweeps. Now let us restrict our attention to the other part of the wavefront, which comprises two convex vertices at the start of the propagation; let us denote it by $\mathcal{W}$.

Inside $\mathcal{W}$, neither split events nor vertex events can occur. In particular, this implies that neither y_0 nor y_0' can disappear before the linear offsets of all the sites YU, UZ, Z, ZZ', Z', $Z'U'$ and $U'Y'$ annihilate.

When the propagation starts, the two convex vertices u and u' of $\mathcal{W}$ originating from U and U', respectively, will start tracing out two arcs g and g' of the geometric graph $(V_{L^{k_\infty}}, E_{L^{k_\infty}})$, respectively.

At a contraction of an edge incident to a convex and a reflex vertex, the delineation of some edge of $L^{k^\infty}(F_\infty)$ will terminate, at a node of degree two, and the tracing of the next edge will start. Note that after such event, the number of the convex vertices in $\mathcal{W}$ will remain unchanged. If at some moment, two convex vertices become adjacent in $\mathcal{W}$, and later the edge incident to them vanishes, then a node of $L^{k^\infty}(F_\infty)$ having degree three will be generated. This node will correspond to a vertex of $(V_{L^{k_\infty}}, E_{L^{k_\infty}})$, at which g and g' meet. Afterwards, only one convex vertex will be present in the wavefront, and it will trace out a single unbounded arc of $(V_{L^{k_\infty}}, E_{L^{k_\infty}})$.

Now let us demonstrate that two convex vertices will necessarily become adjacent in $\mathcal{W}$. Consider $\mathcal{W}$ immediately after the propagation starts, and let z_0 denote the topmost (reflex) vertex from the linear offset of Z. The wavefront edge z_0x incident to z_0 on the right is horizontal, and x is the topmost vertex from the linear offset of Z'. During the propagation, the edge incident to x on the right may annihilate; then the horizontal edge will become incident on the right to a convex vertex, which we shall still denote by x. Let us track the evaluation of the portion of $\mathcal{W}$ between y_0 and x.

Observe that for a reflex vertex q adjacent to u, the bisector ray of the interior wavefront angle at q will intersect that of the interior angle at u, implying that the edge qu will annihilate at their intersection point, unless one of the two edges adjacent to qu vanishes before. Consequently, an edge event will necessarily occur between y_0 and x.

Let us examine such event. If z_0 is not incident to the vanishing edge, then u certainly is. In this case, the convex wavefront vertex, which we shall still denote by u, will be updated. Otherwise, either z_0 has been adjacent to u, and the edge uz_0 has annihilated, or x was a convex vertex, and the edge zx has annihilated. If, as a result, two convex vertices have become adjacent, then we are done. If not, then in the former case, the horizontal wavefront edge emanating from ZZ' has become adjacent to a convex wavefront vertex on the left, and we stop tracking evaluation of the wavefront portion under consideration. In the latter case, the topmost reflex vertex of the linear offset of Z has changed; now it is the one that was previously

adjacent to z_0 on the left. We shall appropriately update the notation, so that z_0 will again denote the topmost reflex vertex of the linear offset of Z, and x will be adjacent to z_0 on the right. Subsequently, we shall look for the next edge event between y_0 and x.

Since after each iteration, the number of edges under consideration is reduced by one, the process is guaranteed to terminate. If we are left with a convex vertex incident to the horizontal wavefront edge, let us repeat the same procedure for the portion of $\mathcal{W}$ between x' and y_0', where x' initially denotes the topmost vertex of the linear offset of Z. This will give us either an edge incident to two convex vertices, or an evidence of the fact that the horizontal edge will become adjacent to a convex vertex on the right as well.

To summarize, having completed the above procedure, we shall locate a wavefront edge e incident to two convex vertices (see Fig. 10b). It emanates from one of the sites UZ, Z, ZZ', Z', and $Z'U'$; therefore, neither y_0 nor y_0' can vanish before e. The wavefront vertices incident to e may undergo several updates, caused by an annihilation of edges from $\mathcal{W}$ adjacent to e, but they will clearly remain convex. Since the number of edges in $\mathcal{W}$ is finite, e will finally shrink to zero itself.

We conclude that $(V_{L^{k\infty}}, E_{L^{k\infty}})$ will have a single unbounded arc, and thus, $L^{k\infty}(F_\infty)$ will have a single unbounded edge. This proves the claim.

□

4 Linear Axis for the Unbounded Face

Let G be a planar straight line graph, no three vertices of which lie on the same line; consider any $\varepsilon > 0$. For any bounded face of G, a linear axis ε-equivalent to the medial axis can be computed from the latter in linear time, as described by Tănase & Veltkamp (2004), and by Trofimov & Vyatkina (2007). Now let us focus our attention on the unbounded face F_∞.

4.1 A Sufficient Condition of ε-Equivalence

For any edge e of $M(F_\infty)$, let us denote by S_1^e and S_2^e the two sites, the Voronoi cells $VC(S_1^e)$ and $VC(S_2^e)$ of which share e.

Within our reasoning, we shall use a notion of a barrier, which we borrow from Trofimov & Vyatkina (2007).

Definition 10. Let e be an edge of $M(F_\infty)$. For any point $c \in e$, the *barrier* b_c^e on the edge e is formed by the two segments connecting c with its closest points from S_1 and S_2, respectively (Fig. 11). The point c is the *center* of the barrier b_c^e.

Observe that for a barrier b_c^e with the endpoints $p_1 \in S_1^e$ and $p_2 \in S_2^e$, any interior point of the segment cp_i is an inner point of the cell $VC(S_i^e)$, for $i = 1, 2$.

In case we do not need to refer explicitly to the center of a barrier, we omit the respective subscript.

Let $CH(\partial F_\infty)$ denote the convex hull of ∂F_∞. The following Lemma formalizes an observation that the half-infinite edges of $M(F_\infty)$ are in one-to-one correspondence with the edges of $CH(\partial F_\infty)$ not belonging to ∂F_∞.

Lemma 9. *Two sites S_1 and S_2 generate a (single) half-infinite edge of $M(F_\infty)$ if and only if S_1 and S_2 are two non-adjacent reflex vertices of F_∞, which correspond to a pair of adjacent vertices of $CH(\partial F_\infty)$.*

Proof. Let S_1 and S_2 be two non-adjacent reflex vertices of F_∞ that correspond to a pair of adjacent vertices of $CH(\partial F_\infty)$. Since the nodes of G are in general position, no interior point of the segment S_1S_2 belongs to ∂F_∞. This implies existence of an empty circle $\mathcal{C}$ passing through S_1 and S_2, the center p of which lies on the opposite side of S_1S_2 from $CH(\partial F_\infty)$. Obviously, p lies on the perpendicular bisector b of S_1S_2. Let $r_p \subset b$ denote the ray with an endpoint at p, directed away from $CH(\partial F_\infty)$. For any point $q \in r_p$, the circle centered at q that passes through S_1 and S_2 is empty; therefore, r_p is entirely contained in a Voronoi edge e shared by $VC(S_1)$ and $VC(S_2)$, which is thus necessarily half-infinite.

Arguing as above, we can associate a separate half-infinite edge of $M(F_\infty)$ with any edge of $CH(\partial F_\infty)$ not belonging to ∂F_∞. Now it remains to show that there are no other half-infinite edges in $M(F_\infty)$.

Let us consider the boundary of $CH(\partial F_\infty)$, and replace each its edge, which does not belong to ∂F_∞, with a barrier on the associated half-infinite Voronoi edge, centered at any interior point of the latter. As a result, we shall obtain a simple closed polygonal chain, which, by Jordan's theorem, bounds a simple polygon H_∞ in the plane (Fig. 12).

Let us show that apart from the infinite pieces of the edges, on which the barriers are imposed, neither entire edges of $M(F_\infty)$ nor their fragments can be found outside H_∞. Suppose for contradiction that there exists an edge e' of $M(F_\infty)$, different from any of those on which the barriers are imposed, such that $e' \cap H_\infty \neq \emptyset$. Consider the connected component M' of $M(F_\infty)$, which contains e'. Our construction implies that none of the edges of M' can intersect the boundary of H_∞; therefore, M' entirely lies outside H_∞. Since there are no sites outside H_∞, M' can contain no cycles (see Lemma 3); thus, M' is a tree. Then M' must have at least two leaves. Since outside or on the boundary of H_∞, there are no convex vertices of F_∞, none of the leaves can be finite. Consider a path l between any two infinite leaves in M'. Observe that l is infinite as well; therefore, it partitions the plane into two regions, any of which must contain at least one face of $M(F_\infty)$. But H_∞ cannot be traversed by l, and thus, all the sites lie on the same side of l. It follows that some face of $M(F_\infty)$ will contain no sites, which is a contradiction.

We conclude that any half-infinite edge of $M(F_\infty)$ was indeed encountered at the first step of our reasoning, and the claim follows.

□

Theorem 2. ***(A sufficient condition of ε-equivalence.)***

$L^{k\infty}(F_\infty)$ and $M(F_\infty)$ are ε-equivalent, if:

(i) for any non-ε-edge e of $M(F_\infty)$, the incident nodes of which belong to different ε-clusters, and neither coincides with a convex vertex of F_∞, there exists a point $c \in e$, such that $b_c^e \subset LC(S_1^e) \cup LC(S_2^e)$;

(ii) for any site S lying on $CH(\partial F_\infty)$, $LC(S)$ is unbounded.

Proof. Suppose that both conditions (i) and (ii) hold. Let $\mathcal{B}$ denote the respective set of barriers. Observe that any half-infinite edge of $M(F)$ is incident to an inner node, and thus, $\mathcal{B}$ contains a barrier for every such edge.

First, consider the barriers on the unbounded edges of $M(F_\infty)$, and form a simple polygon H_∞ as in the proof of Lemma 9. Then let $H_\infty^* = H_\infty \cap F_\infty$. By construction, any barrier from $\mathcal{B}$ lies either inside or on the boundary of H_∞^* (Fig. 13a).

Applying the arguments given by Trofimov & Vyatkina (2007), it can be shown that the barriers from $\mathcal{B}$ partition the interior of H_∞^* into a number of polygonal regions, each containing the same number of nodes of degree three of $M(F_\infty)$ and of $L^{k\infty}(F_\infty)$

(Fig 13a,b). And thus, each region will contain the same number of vertices having degree three of the geometric graphs (V_M, E_M) and $(V_{L^{k\infty}}, E_{L^{k\infty}})$.

Let us define a bijection f' between the finite vertices of the geometric graphs as follows. The hanging vertices, which correspond to the same convex vertex of F_∞, are mapped to each other. The vertices from any ε-cluster of (V_M, E_M), which have degree three, are mapped to distinct vertices of $(V_{L^{k\infty}}, E_{L^{k\infty}})$ having degree three, which lie in the same region of the partition induced by $\mathcal{B}$.

By a reasoning similar to the one developed by Trofimov & Vyatkina (2007), it can be demonstrated that f' satisfies condition (3) from the definition of ε-equivalence in respect of finite vertices, while conditions (1) and (2) are trivially met.

It remains to show that f' can be appropriately extended to the infinite vertices of the geometric graphs.

First, observe that the existence of the barriers, which satisfy condition (i), on the unbounded edges, assures that for any site S lying inside $CH(\partial F_\infty)$, the cell $LC(S)$ lies inside H_∞.

By condition (ii), for any site S lying on the boundary of $CH(\partial F_\infty)$, $LC(S)$ is unbounded. It follows that the linear cells of two such sites share a half-infinite edge if and only if those sites are encountered one after another when walking around the boundary of $CH(\partial F_\infty)$. Moreover, the linear cells of any two non-consecutive sites lying on the boundary of $CH(\partial F_\infty)$ cannot become adjacent outside $CH(\partial F_\infty)$, and thus, outside H_∞.

We conclude that outside H_∞, no three linear cells can meet, and therefore, no nodes of degree three of $L^{k\infty}(F_\infty)$ lie outside H_∞.

Next, observe that for any edge S_1 of F_∞ lying on the boundary of $CH(\partial F_\infty)$, and its incident reflex vertex S_2, their linear cells are necessarily separated by a single half-infinite edge e emanating from S_2 perpendicular to S_1. Since e is traced out by a reflex wavefront vertex, it is not part of $L^{k\infty}(F_\infty)$.

On the contrary, for any two non-adjacent reflex vertices S_1 and S_2 of F_∞ being consecutive sites in the boundary of $CH(\partial F_\infty)$, the half-infinite edge shared by $LC(S_1)$ and $LC(S_2)$ is necessarily present in $L^{k\infty}(F_\infty)$.

Together with Lemma 9, this implies that we can define a bijection f^* between the infinite nodes of $M(F_\infty)$ and $L^{k\infty}(F_\infty)$, by mapping to each other the infinite nodes generated in either axis by any two non-adjacent reflex vertices S_1 and S_2 of F_∞, such that the segment S_1S_2 is an edge of $CH(\partial F_\infty)$.

Naturally, f^* induced a bijection f'' between the infinite vertices of the geometric graphs (V_M, E_M) and $(V_{L^{k\infty}}, E_{L^{k\infty}})$. Having extended f' to the infinite vertices as prescribed by f'', we shall obtain a bijection f between V_M and $V_{L^{k\infty}}$, which satisfies conditions (1) and (2) from the definition of ε-equivalence.

To demonstrate that f also respects condition (3), it remains to verify that for any two corresponding infinite vertices V_M and $V_{L^{k\infty}}$, their (finite) adjacent vertices fall into the same region of the partition of H^*_∞ induced by $\mathcal{B}$. This can be done by means of the same arguments as proposed by Trofimov & Vyatkina (2007) for analysis of the (finite) arcs of the geometric graphs in the case of polygons with holes. □

4.2 The Algorithm

To construct a linear axis ε-equivalent to the medial axis for the unbounded face F_∞, we follow the same scheme as for polygons. At the first stage, a sequence k_∞ of hidden edges is computed, which guarantees ε-equivalence between the axes, and at the second stage, $L^{k\infty}(F_\infty)$ is reconstructed by appropriately adjusting $M(F_\infty)$.

To describe the algorithm, we shall need the notions of a conflicting pair of sites, of an impending vertex, and of an outer vertex. The definition of a conflicting pair given below is equivalent to that proposed by Trofimov & Vyatkina (2007).

For a site S, let $N_{VC}(S)$ denote the set of nodes of $M(F_\infty)$ incident to its Voronoi cell $VC(S)$.

Definition 11. Let e be a finite non-ε-edge of $M(F_\infty)$ incident to the nodes u and v. Two sites S_1 and S_2 form a *conflicting pair* for e if the following three conditions hold:
- neither $VC(S_1)$ nor $VC(S_2)$ is incident to e,
-$N_{VC}(S_1) \cap C(u) \neq \emptyset$ and $N_{VC}(S_2) \cap C(v) \neq \emptyset$,
- at least one of S_1 and S_2 is a reflex vertex.

Note that if in the above definition, either u or v coincides with a convex vertex of F_∞, no conflicting pairs of sites will exist for e.

Definition 12. Let e be a half-infinite edge of $M(F_\infty)$ incident to a finite node u. A reflex vertex S is *impending* for e if $VC(S)$ is not incident to e, and $N_{VC}(S) \cap C(u) \neq \emptyset$.

Definition 13. A reflex vertex S is *outer* if S lies on $CH(\partial F_\infty)$, and at least one edge of F_∞ incident to S is not contained in the boundary of $CH(\partial F_\infty)$.

Below we outline the algorithm for computation of a sequence k_∞ of hidden edges, such that $L^{k\infty}(F_\infty)$ is ε-equivalent to $M(F_\infty)$, for a given $\varepsilon > 0$. Here $k_\infty = \{k_1, \ldots, k_r\}$, where k_j is the number of hidden edges inserted at a reflex vertex S_j of F_∞, for $1 \leq j \leq r$, and r is the number of the reflex vertices of F_∞. Further we shall provide details on handling conflicting pairs, and impending and outer vertices.

Algorithm *ComputeHiddenEdges*(F_∞,ε)

Input: the unbounded face F_∞ of a planar straight line graph G, and a real constant $\varepsilon > 0$.
Output: a sequence k_∞ of hidden edges, such that the linear axis $L^{k\infty}(F_\infty)$ is ε-equivalent to the medial axis $M(F_\infty)$.

1. Compute $M(F_\infty)$.
2. For each reflex vertex S_j of F_∞:
   Let α_j be the size of the internal angle at S_j.
   /* Initialize the speed s_j of the vertex S_j. */
   $s_j := \frac{1}{\cos((\alpha_j - \pi)/2)}$
3. *ComputeConflictingSites*(ε).
4. For each conflicting pair (S_i, S_j) of sites for each finite non-ε-edge e, the incident nodes of which belong to different ε-clusters, and neither coincides with a convex vertex of F_∞:
   HandleConflictingPair(e, S_i, S_j).
5. *FindImpendingVertices*(ε).
6. For each impending vertex S_k for each half-infinite edge e:
   HandleImpendingVertex(e, S_k).
7. *FindOuterVertices*().
8. *HandleOuterVertices*().
9. For each reflex vertex S_j of F_∞:
   $k_j := \lceil \frac{\alpha_j - \pi}{2\cos^{-1}(1/s_j)} \rceil$.
10. *PreserveUnboundedness*(k_∞).

4.2.1 Handling Conflicting Pairs

When handling a conflicting pair (S_1, S_2) of sites for a finite non-ε-edge e, we bound the speed of any reflex vertex present in the pair as to assure the existence of a barrier on e separating the linear offsets of S_1 and S_2 throughout the propagation. This can be done by means of the technique developed by Tănase (2005). As shown by Trofimov & Vyatkina (2007), the existence of such barrier on the edge e for each conflicting pair for e will imply existence of a single barrier b^e, which will separate the linear offsets of any sites forming a conflicting pair for e, and satisfy condition (i) from the Theorem 2.

Consequently, having completed steps 3-4 of the algorithm, we shall be able to meet condition (i) for any finite non-ε-edge e, the incident nodes of which belong to different ε-clusters, and neither coincides with a convex vertex of F_∞, by inserting at each reflex vertex of F_∞ a number of hidden edges that allows to respect the derived bound on its speed.

4.2.2 Handling Impending Vertices

The technique by Tănase (2005) can be adapted to process the impending vertices. More precisely, her approach for treating a conflicting pair formed of a reflex vertex and an edge, with respect to a non-ε-edge generated by two reflex vertices, can be extended so as to handle an impending (reflex) vertex for a half-infinite Voronoi edge (generated by two reflex vertices). Applying the resulting method to a half-infinite edge e of $M(F_\infty)$ and its impending vertex S_k, and appropriately updating the bound on the speed of S_k if necessary, we assure existence of a barrier on e, which the linear offset of S_k cannot reach *from inside* H_∞.

Suppose that for any impending vertex for e, we have presented a suitable barrier. Then the one with the center farthest from the finite node of $M(F_\infty)$ incident to e, will be unreachable from inside H_∞ for all the impending vertices for e.

To summarize, having performed steps 5-6 of our algorithm, we can claim that on each half-infinite Voronoi edge e, a barrier b^e exists, which cannot be reached from inside H_∞ by the linear offset of any site $S \neq S_1^e, S_2^e$. To ensure that condition (i) from the Theorem 2 holds for b^e, we need to further demonstrate that for any reflex vertex S lying on the boundary of $CH(\partial F_\infty)$, its linear offset will not reach b^e *from outside* H_∞ as well. This will be accomplished when handling the outer vertices.

4.2.3 Handling Outer Vertices

Let us surround the graph G with a circle $\mathcal{C}$, which would also contain inside all the finite edges of $M(F_\infty)$, and the centers of all the barriers constructed at the previous steps (Fig. 14). For any site S lying on the boundary of $CH(\partial F_\infty)$, let $\overline{LC}(S)$ denote the part of its linear cell that lies inside $\mathcal{C}$, and let $\overline{e}_S \subset \mathcal{C}$ denote the circular arc, which is present in the boundary of $\overline{LC}(S)$.

Our goal is to assure that the barrier we have imposed on any half-infinite edge of $M(F_\infty)$ will never be reached *from outside* H_∞ by the linear offset of any reflex vertex S lying on the boundary of $CH(\partial F_\infty)$. To this end, it is sufficient to guarantee that for any outer vertex S, there will exist a point $c \in \overline{e}_S$, such that the segment Sc can never be reached by the linear offset of an outer vertex S', for which the segment SS' belongs to the boundary of $CH(\partial F_\infty)$. Here the segment Sc plays a role of a "simplified" barrier on the edge e_S, which should lie inside $LC(S)$.

To ensure the existence of such point $c \in \overline{e}_S$, we shall again apply a modified technique by Tănase (2005). In case the both vertices S' and S'', for which SS' and SS'' are contained in the boundary of $CH(\partial F_\infty)$, are outer, they should be handled analogously to a conflicting pair formed of two reflex vertices, with respect to the edge $\overline{e}_S$. Otherwise, precisely one of S' and S'' is outer; assume it is S'. Then S' should be handled as if, together with some imaginary segment, they constituted a conflicting pair for $\overline{e}_S$.

Thus, at step 8, we update the bound on the speed of any outer vertex, so as to finally establish the existence on any half-infinite Voronoi edge of a barrier satisfying condition (i) from Theorem 2. In particular, we may assume that the center of any such barrier lies outside $CH(\partial F_\infty)$; let us admit this assumption.

Next, for each reflex vertex of F_∞, we calculate the (minimum) number of hidden edges that allows to meet the currently obtained bound on its speed. This completes step 9.

4.2.4 Preserving Unboundedness

At step 10, we aim to assure the unboundedness of the linear cell of each site that lies on the boundary of $CH(\partial F_\infty)$. To this end, it is convenient to directly compute the number of hidden edges that should be inserted at each outer vertex, rather than to bound their speeds.

To meet our goal, it is sufficient to guarantee that in the linear offset of each outer vertex, there will be a (reflex) vertex, the trace of which will be unbounded.

Let S be an outer vertex of F_∞, with an internal angle $\alpha > \pi$. Consider the reflex vertices S_1 and S_2 of F_∞, which lie on the boundary of $CH(\partial F_\infty)$, such that the segments SS_1 and SS_2 are the edges of $CH(\partial F_\infty)$ (Fig. 15). At least one of S_1 and S_2 is necessarily outer. Raise perpendiculars l_1 and l_2 at S on the segments SS_1 and SS_2, respectively. Consider the rays $r_1 \subset l_1$ and $r_2 \subset l_2$ lying on the other side from $CH(\partial F_\infty)$ of the lines through SS_1 and SS_2, respectively. Let W_S denote the wedge in the plane defined by l_1 and l_2, and bounded by r_1 and r_2. By our non-degeneracy assumptions, the size γ of its apex angle satisfies $0 < \gamma < \pi$. Let r' and r'' be two rays with an endpoint at S, which lie inside W_S and partition its apex angle in proportion 2:1:2; assume that r' denotes the ray forming an angle of $2\gamma/5$ with r_1.

Observe that if we insert at S at least $k_S = \lceil 5(\alpha - \pi)/\gamma \rceil$ hidden edges, then at least five wavefront vertices emanating from S will move along the rays that fall inside W_S, and at least one of those rays will lie between r' and r''.

Let us update (if needed) the number of hidden edges associated with S, so that it would be at least k_S. Further, we perform the same operation for any other outer vertex.

Without loss of generality, assume that S_1 is an outer vertex encountered before S when walking counterclockwise along the boundary of $CH(\partial F_\infty)$, the segment S_1S is horizontal, and $CH(\partial F_\infty)$ lies below the line through S_1S.

Let e denote the half-infinite edge of $M(F_\infty)$ shared by $VC(S_1)$ and $VC(S)$. Consider a barrier b_c^e, which satisfies condition (i) of Theorem 2, and the center c of which lies above S_1S; the existence of such barrier has been assured before. Since $b_c^e \subset LC(S_1) \cup LC(S)$, it must intersect an edge of $L^{k_\infty}(F_\infty)$ shared by $LC(S_1)$ and $LC(S)$, at some point p. We conclude that at some moment, the linear offsets of S_1 and S are found adjacent at a convex vertex v, which resides at the point p. As the prop-

agation proceeds, one of the edges incident to v may shrink to zero; thereby, its incident reflex wavefront vertex (that belongs to the offset of either S_1 or S) will disappear; the convex vertex v will be updated, and the process will continue.

Recall that in the linear offset of S, there is a vertex w that moves along a ray r^* lying between r' and r''. If there are two or more such vertices, denote by w the one that minimizes the angle between r' and r^*. Thus, the vertex z, which precedes w in the clockwise order of the vertices from the linear offset of S, will move along the ray lying strictly between r_1 and r'.

We would like to preserve the vertex w in the linear wavefront throughout the propagation. Since we shall guarantee the existence of a wavefront vertex tracing out an unbounded ray, in the linear offset of *any* outer vertex, we may assume that w cannot vanish in a result of an interaction between the linear offset of S, and that of any site other than S_1 and S_2. Now let us determine the number of hidden edges for S_1, sufficient to guarantee that the edge zw will never be caused to vanish by the linear offset of S_1.

Let y_1 and y denote the topmost vertices of the linear offsets of S_1 and S, respectively. If the offset of S_1 contains a horizontal edge, let y_1 denote its leftmost incident vertex; if the offset of S contains a horizontal edge, let y denote its rightmost incident vertex. Observe that since at least two vertices from the linear offset of S move along the rays, which fall between r_1 and r', we have $y \neq z$.

In case y moves upwards at least as fast as y_1, zw obviously cannot disappear through an interaction between the offsets of S_1 and S. Otherwise, y_1 can never disappear in a result of such interaction, and thus, the edge y_1q incident to y_1 on the right cannot shrink to zero. Consequently, to achieve our goal, it is sufficient to assure that the angle ψ between the oriented segment $\overrightarrow{y_1q}$ and the positive direction of the x-axis, is not greater than the angle ϕ made by $\overrightarrow{zw}$ with the positive direction of the x-axis.

The angle ϕ is bounded from below by $\gamma/5$. For k_1 hidden edges inserted at S_1, the angle ψ is bounded from above by $(\alpha_1 - \pi)/(k_1 + 1)$, where α_1 denotes the internal angle of F_∞ at S_1. Consequently, we shall require $(\alpha_1 - \pi)/(k_1 + 1) \leq \gamma/5$, which gives us $k_1 \geq 5(\alpha_1 - \pi)/\gamma - 1$. Let $k_{S_1,S} = \lceil 5(\alpha_1 - \pi)/\gamma \rceil - 1$.

From what has been said above, we derive the following claim. If we insert at S_1 at least $k_{S_1,S}$ hidden edges, then any vertex from the linear offset of S, which moves along a ray lying between r' and r'', will never vanish because of an interaction between the offsets of S_1 and S.

If S_2 is not outer, then its linear offset can never meet that of S. Otherwise, applying symmetric arguments, we shall obtain a number $k_{S_2,S}$ of hidden edges for S_2, which would assure that any vertex from the linear offset of S that moves along a ray lying between r' and r'', – in particular, w – can never be eliminated through an interaction between the linear offsets of S and S_2.

Next, we shall update (if necessary) the number of hidden edges inserted at S_1 and S_2, so that it would be at least $k_{S_1,S}$ and $k_{S_2,S}$, respectively.

Further we repeat the same procedure for any outer vertex S. Thereby, for any such S, we ensure the existence in its linear offset of a reflex wavefront vertex that will never meet the offsets of the (reflex) vertices S_1 and S_2 of F_∞, for which the segments S_1S and SS_2 belong to the boundary of $CH(F_\infty)$. As soon as we are done, we obtain a guarantee that the linear cell of any outer vertex is unbounded, which implies that for *any* site lying on the boundary of $CH(\partial F_\infty)$, its linear cell is unbounded.

4.2.5 Putting Everything Together

After we obtain the sequence k_∞ of hidden edges, which guarantees ε-equivalence between $L^{k_\infty}(F_\infty)$ and $M(F_\infty)$, the former can be reconstructed from the latter by a method similar to that proposed by Tănase (2005) for the case of simple polygons.

For any finite non-ε-edge of $M(F_\infty)$, each conflicting pair can be handled in constant time. Handling an impending vertex for any half-infinite edge of $M(F_\infty)$ requires $O(1)$ time as well. All the outer vertices can be processed at steps 7, 8, and 10 in total linear time. It follows that in case each ε-cluster of $M(F_\infty)$ has a constant size, the algorithm *Compute-HiddenEdges*(F_∞,ε) will run in linear time. Moreover, the same will hold for the algorithm, which computes $L^{k_\infty}(F_\infty)$ from $M(F_\infty)$.

We summarize the above observations, along with the results previously obtained by Tănase & Veltkamp (2004), by Tănase (2005), and by Trofimov & Vyatkina (2007), in the following Theorem.

Theorem 3. *Let G be a planar straight line graph on n vertices in general position, none of those being isolated, with a constant number of nodes in each ε-cluster of its medial axis $M(G)$. For a given $\varepsilon > 0$, a linear axis $L^k(G)$ ε-equivalent to $M(G)$ can be computed from the latter in linear time. If $M(G)$ is not pre-computed, the time complexity of the algorithm amounts to $O(n \log n)$.*

5 Conclusion

In this work, we have studied a linear axis for planar straight line graphs as opposed to that for polygons. A principal difference between the two cases is due to the presence of half-infinite edges in a linear axis for the unbounded face of a graph.

If the graph vertices are in general position, then for any $\varepsilon > 0$, one can compute a linear axis ε-equivalent to the medial axis with an asymptotically efficient algorithm. However, in presence of triples of almost collinear vertices in the unbounded face, the number of hidden edges thereby inserted at its reflex vertices can appear arbitrarily large (though constant in the total number of its vertices).

On the other hand, any potential practical applications are unlikely to require similarity between the two axes at infinity. Yet to achieve semblance of those inside a bounded domain of interest, one can choose a sufficiently large circle $\mathcal{C}$ at step 8 of our algorithm (see Section 4.2.3 for details), and omit the last step. The resulting number of hidden edges can be thus significantly reduced. Moreover, under such modification, no general position assumption on the graph vertices is required.

The motivation for this research was theoretical. Yet it would be interesting to find a practical application of the obtained results.

Acknowledgements

This work was completed with the support of Russian Foundation for Basic Research (grant 07-07-00268-a) and Human Capital Foundation.

The author thanks Franz Aurenhammer for pointing out that the possibility of generalizing earlier results to the case of planar straight line graphs should be investigated.

References

Aichholzer, O., Aurenhammer, F., Alberts, D. & Gärtner., B. (1995), A novel type of skeleton for

polygons, *The Journal of Universal Computer Science*, **1**, 752–761.

Aichholzer, O. & Aurenhammer, F. (1996), Straight skeletons for general polygonal figures, *in* '2nd Ann. Int'l. Computing and Combinatorics Conf. COCOON'96, Hong Kong', LNCS, Springer Verlag **1090**, 117–126.

Barequet, G., Eppstein, D., Goodrich, M. T. & Vaxman, A. (2008), Straight skeletons of three-dimensional polyhedra, *arXiv:0805.0022v1 [cs.CG]*, *to appear in* '16th Annual European Symposium on Algorithms'.

Eppstein, D. & Erickson, J. (1999), Raising roofs, crashing cycles, and playing pool: applications of a data structure for finding pairwise interactions, *Discrete and Computational Geometry* **22**(4), 569–592.

Tănase, M. (2005), Shape Decomposition and Retrieval, Ph.D., Utrecht University.

Tănase, M. & Veltkamp, R. C. (2004), Straight skeleton approximating the medial axis, *in* '12th Annual European Symposium on Algorithms', pp. 809–821.

Trofimov V. & Vyatkina, K. (2007), Linear Axis for General Polygons: Properties and Computation, *in* 'International Conference on Computational Science and Its Applications ICCSA 2008', LNCS, Springer **4705**, 122–135.

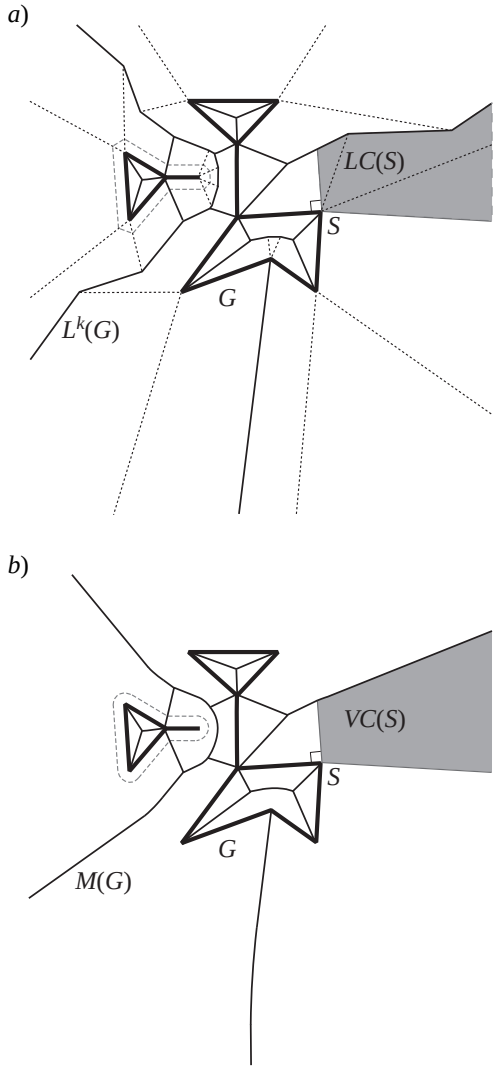

Figure 5: A planar straight line graph G is depicted bold; S is a reflex vertex of the unbounded face. a) A linear axis $L^k(G)$ is shown solid; one hidden edge is inserted at each reflex vertex. The traces of the reflex wavefront vertices are indicated dotted. The linear cell $LC(S)$ is shaded gray. For the left connected component of G, the portion of the linear wavefront emanating from it, which sweeps the unbounded face, is shown dashed gray. b) The medial axis $M(G)$ is shown solid; the Voronoi cell $VC(S)$ is shaded gray. For the left connected component of G, the portion of the uniform wavefront emanating from it, which sweeps the unbounded face, is shown dashed gray. Here, at each vertex of G incident to a bridge, two leaves of each axis coincide; however, we treat any such two leaves as two different non-adjacent nodes of the corresponding axis.

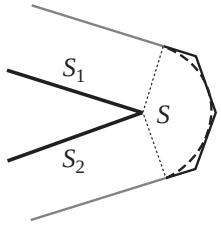

Figure 6: The reflex vertex S is incident to the edges S_1 and S_2; two hidden edges are inserted at S. The linear offset of S is shown solid black, and the unform one – dashed black. The linear and the uniform offsets of each of S_1 and S_2 are identical; their fragments are depicted solid gray. The offsets of S are separated from those of S_1 and S_2 by the wavefront points lying on the perpendiculars at S to S_1 and S_2, respectively; the perpendiculars are indicated dotted.

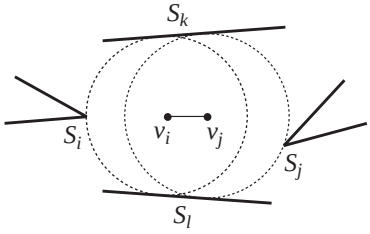

Figure 7: The sites S_i and S_j are reflex vertices; the sites S_k and S_l are edges. For a relatively small $\varepsilon > 0$, $v_i v_j$ is an ε-edge if S_i, S_k, S_j, and S_l are almost cocircular.

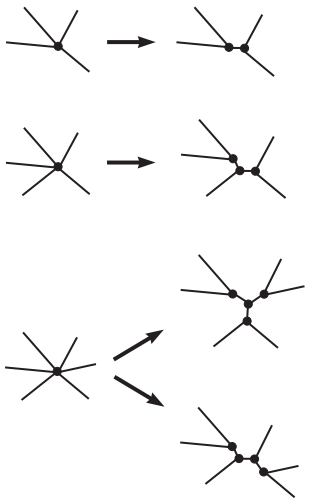

Figure 8: The interpretation of nodes of degree 4 and 5, and two possible interpretations of a node of degree 6.

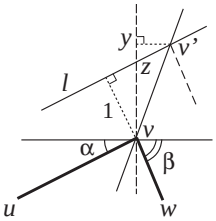

Figure 9: The reflex wavefront vertex v is adjacent to u and w; $\alpha < \beta$. In a unit time, v will move to v'. The line l is parallel to (u, v), and $|vy| > |vz| > 1$.

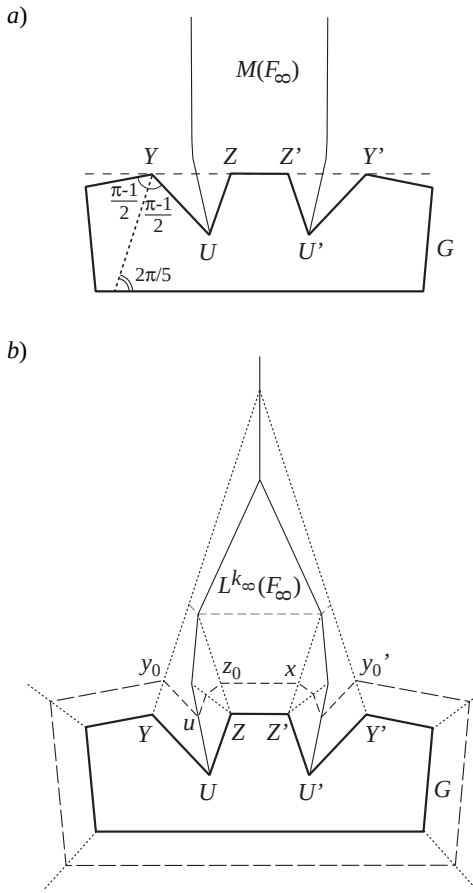

Figure 10: a) The graph G has a single connected component; the medial axis $M(F_\infty)$ for its unbounded face F_∞ has two half-infinite edges. b) One hidden edge is inserted at each of Z and Z'; the corresponding linear axis $L^{k_\infty}(F_\infty)$ has a single half-infinite edge. The topmost vertices y_0 and y_0' of the linear offsets of Y and Y', respectively, split the linear wavefront into two parts. One of those (short dashed) initially contains two convex vertices; in the text, it is denoted by $\mathcal{W}$. The other one (long dashed) contains no convex vertices. The convex wavefront vertex u emanates from U; z_0 and x are the topmost vertices of the offsets of Z and Z', respectively. In the propagation, the linear offset of ZZ' will become incident to two convex vertices; the portion $\mathcal{W}$ of the wavefront at that moment is depicted short dashed gray. c) Two hidden edges are inserted at Y. The linear offset of Y (solid black) comprises four edges; perpendiculars from Y onto those edges are shown dashed. The traces of the reflex vertices are indicated dotted, except for the one belonging to the bisector ray of the interior wavefront angle at Y, which is depicted solid gray. For the edge e from the linear offset of Y, the angle θ_e between the perpendicular dropped from Y onto e and the bisector ray of the interior wavefront angle at Y, equals 1/6.

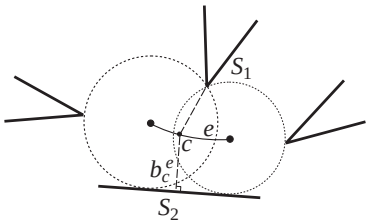

Figure 11: The edge e of $M(F_\infty)$ is generated by the sites S_1 and S_2, being a reflex vertex and a segment, respectively. For the point $c \in e$, the barrier b_c^e is depicted dashed.

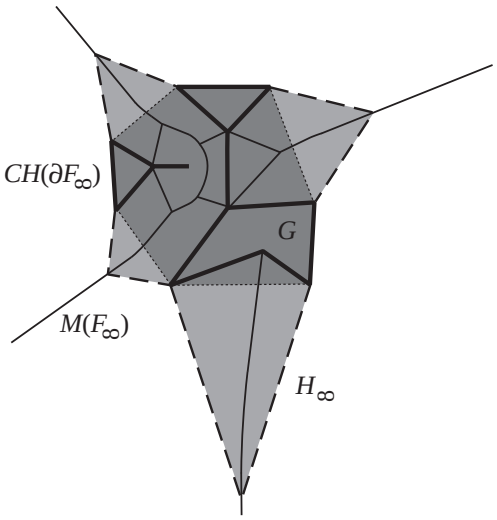

Figure 12: A planar straight line graph G (bold), and the medial axis $M(F_\infty)$ for its unbounded face F_∞. The convex hull $CH(\partial F_\infty)$ of ∂F_∞ is shown dark gray; the edges of the former not belonging to the latter are indicated dotted. The barriers on the half-infinite edges of $M(F_\infty)$ are depicted dashed, and the corresponding simple polygon H_∞ is shaded gray; in particular, the regions forming $H_\infty \setminus CH(F_\infty)$ are light gray.

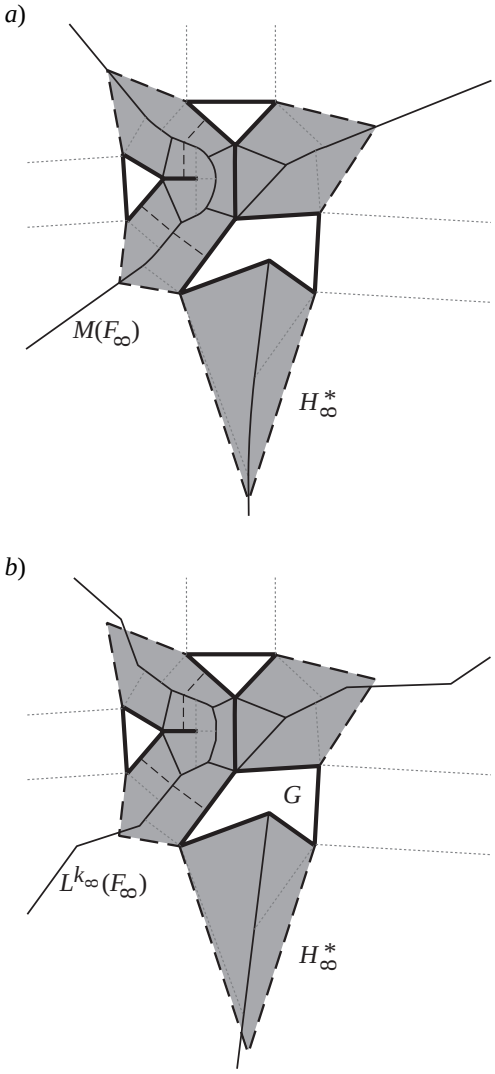

Figure 13: A planar straight line graph G is depicted bold. a) The medial axis $M(F_\infty)$ for the unbounded face F_∞ of G is shown solid; the edges of the Voronoi cells not being part of $M(F_\infty)$ are indicated dotted gray. For an appropriate value of ε, a barrier is imposed on each non-ε-edge of $M(F_\infty)$, the incident nodes of which belong to different ε-clusters, and neither coincides with a convex vertex of $M(F_\infty)$; the barriers are depicted dashed. The corresponding polygon H_∞^* is shaded gray. b) The linear axis $L^{k_\infty}(F_\infty)$ for the unbounded face F_∞ of G is shown solid; one hidden edge is inserted at each reflex vertex of F_∞. The edges of the linear cells not being part of $L^{k_\infty}(F_\infty)$ are indicated dotted gray. For any edge e of $M(F_\infty)$, on which a barrier b^e is imposed, b^e lies inside $LC(S_1^e) \cup LC(S_2^e)$. Inside each region of the partition of H_∞^* induced by the barriers, there lies the same number of nodes of degree three of either axis.

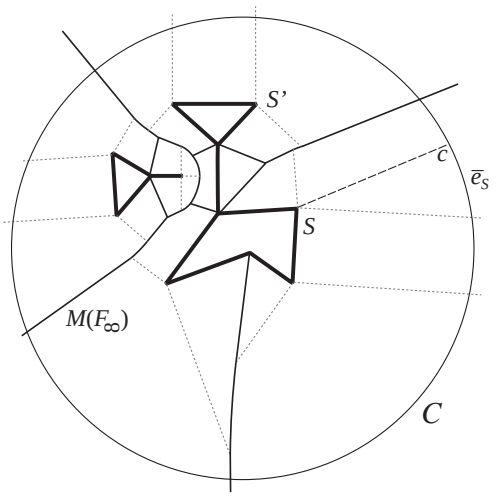

Figure 14: The circle $\mathcal{C}$ encloses the graph G, all the finite edges of $M(F_\infty)$, and the centers of all barriers constructed for the non-ε-edges of $M(F_\infty)$. The vertices S and S' are outer; SS' is present in the boundary of $CH(\partial F_\infty)$, but not in ∂F_∞. The restricted linear cell $\overline{LC}(S)$ is bounded by the arc $e_S \subset \mathcal{C}$. To assure that for a point $c \in e_S$, the segment Sc will lie inside $LC(S)$, we shall process S' as if together with some segment, it formed a conflicting pair for e_S, and update the number of hidden edges associated with S'.

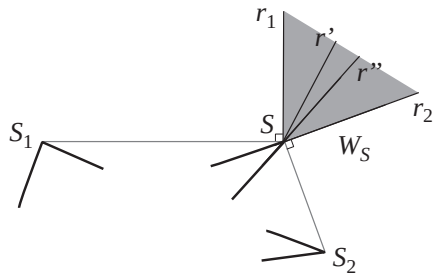

Figure 15: The vertex S is outer; for the reflex vertices S_1 and S_2, the segments S_1S and SS_2 are contained in the boundary of $CH(\partial F_\infty)$, but not in ∂F_∞. The rays r_1 and r_2 are perpendicular to S_1S and SS_2, respectively. The wedge W_S is grayed. The rays r' and r'' partition the apex angle of W_S in proportion 2:1:2.