Segment endpoint visibility graphs are Hamiltonian

M. Hoffmann and Cs. D. Tóth*
Institute for Theoretical Computer Science
ETH Zürich, CH-8092 Zürich, Switzerland
{hoffmann,toth}@inf.ethz.ch

Abstract

A question of Mirzaian is answered in the affirmative and it is shown that the *segment endpoint visibility graph* of any set of disjoint line segments in the plane admits a simple Hamiltonian polygon.

1 Introduction

The *segment endpoint visibility graph* is defined for n disjoint line segments in the plane. The vertices are the $2n$ segment endpoints and two vertices a and b are connected by an edge if and only if the corresponding open line segment ab is either one of the given segments or it does not intersect any given (closed) segment (see Figure 1 for an example). Note that this is different from the *segment visibility graph*, where vertices correspond to segments and an edge connects two vertices if and only if some points of the two segments are mutually 'visible'.

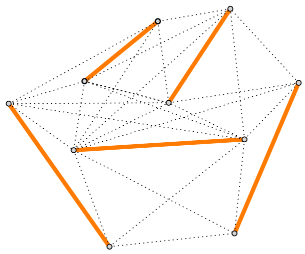

Figure 1: A segment endpoint visibility graph.

Visibility graphs of disjoint objects or vertices/sides of polygons are fundamental structures in computational geometry [11, 3]. They have applications in shortest path computation, motion planning, art gallery problems, but also in VLSI design, and computer graphics. The characterization and recognition of visibility graphs are also of

*The second author acknowledges support from the joint Berlin-Zürich graduate program "Combinatorics, Geometry, and Computation", financed by the German Science Foundation (DFG) and ETH Zürich.

independent interest. Visibility concerning disjoint line segments in the plane is basic, and problems for more complex objects can often be reduced to or approximated by this structure.

Previous works and main theorem. Segment endpoint visibility graphs have been subject of extensive research. The number of edges [16], the computational complexity [18, 13, 7, 9], storage space [1, 8], and on-line updates [6] have been studied for this class of graphs over the past decade.

We are interested in the following problem posed by Mirzaian [10]. What is the maximal number $f(n)$ such that any segment endpoint visibility graph on n segments has a cycle corresponding to a simple closed polygon in the plane of size $f(n)$? He conjectured that $f(n) = n$, or at least $f(n) = \Omega(n)$. The stronger statement can be formulated as follows. Every set of disjoint line segments admits a *Hamiltonian polygon*, that is, there exists a simple closed polygonal paths whose vertices are exactly the endpoints of the line segments and whose sides correspond to edges of the segment endpoint visibility graph.

Mirzaian [10] proved that there is a Hamiltonian polygon, if the line segments are *convexly independent*, that is, if every line segment has at least one endpoint on the boundary of the convex hull $int(conv(\bigcup L))$.

O'Rourke and Rippel [12] proved that there is a Hamiltonian polygon, if the supporting line of no line segment crosses any other line segment. (Two segments cross, iff they have a common point which lies in the interior of both segments.) Urabe and Watanabe [17] gave a construction where L does not always have a Hamiltonian polygon which contains $\bigcup L$, i.e. a *circumscribing polygon*. Earlier, Rappaport [14] showed that L does not always have a Hamiltonian polygon in which every line segment of L is a side (i.e. *alternating polygon*). We prove the conjecture of Mirzaian in its stronger form.

Theorem 1 *For any set L of at least two disjoint line segments in general position, there exists a Hamiltonian polygon.*

Segments are in *general position* if there are no three collinear segment endpoints.

Proof technique. We build a Hamiltonian polygon recursively, starting from the convex hull $conv(\bigcup L)$. In each step, the polygon is extended to go through more segment endpoints. As a first approach, our polygon P is extended to include both endpoints of the segments whose one endpoint is already on the convex hull, providing a new proof of Mirzaian's theorem [10]. The resulting polygon, in order to avoid crossings with segments, may pass through one endpoint of some segments which were in the interior of the convex hull. Our Algorithm 1 recursively extends P to include the second endpoint of segments if one endpoint is already a vertex in P; and, as a result, all remaining segments are in the interior of P.

The most prominent but rather technical part of the proof is to provide coherent operations that extend P to segments in $int(P)$. To be able to apply the defined operations, polygon P must satisfy a chiseled set of properties (see Definition 3) in every step. Once P attains certain properties, a simple induction completes the proof.

Every step of the algorithm and every operation relies only on elementary geometry, like ray shooting, convex hull, or sorting angles. It is straightforward to give an $O(n \log n)$ algorithm based on our proof.

The rest of the paper is organized as follows. In Section 2, we prove Theorem 1 by induction. The key lemma of the proof, Lemma 2, is proved algorithmically in three phases. Section 3 gives some basic operations of our algorithm, Section 4 provides a new proof of the theorem of Mirzaian [10] and provides the first phase of our algorithm. The two last phases and the complete algorithm are discussed then in Sections 5–6.

2 Proof of Theorem 1

Given a set L of disjoint line segments in the plane, denote by V_L the set of segment endpoints from L. P is a simple polygon if it is a closed region of the plane enclosed by a simple closed curve ∂P of line segments. Let $V(P)$ denote the cyclic sequence of vertices of P along its boundary ∂P. If P is a simple polygon, ∂P does not cross any segment of L, and the set of vertices of P is V_L, then P is a Hamiltonian polygon. The following lemma is crucial in our argument, it establishes Theorem 1 by a simple induction.

Lemma 2 *For a set L of disjoint line segments in general position and a side yz of $conv(\bigcup L)$, there is a simple polygon P with vertices from V_L and a set S of pairwise non-overlapping simple polygons within P satisfying the following properties.*

(L1) yz is a side of P;

(L2) for each $\ell \in L$, $V(P)$ includes either both endpoints of ℓ or $\ell \subset int(P)$;

(L3) for each $\ell \in L$, either $\exists D \in S$ with $\ell \subset int(D)$ or $\ell \cap int(D) = \emptyset$ for all $D \in S$;

(L4) every $D \in S$ is convex;

(L5) each polygon $D \in S$ has a common side with P which is different from yz.

The proof of Lemma 2 is the subject of the remaining sections. The outline of the proof is as follows. We construct the polygon P and the dissection S of P in three phases. We start with $P = conv(\bigcup L)$ which is a simple polygon satisfying (*L1*); and $S = \{P\}$. In each phase, the polygon P and set S are modified such that $V(P)$ never decreases and we do not loose property (*L1*). After the first phase, the polygon has property (*L2*) as well. A simple dissection by diagonal segments assures (*L3*). After the second and third phase resp., we assure properties (*L4*) and (*L5*), too.

Proof. (of Theorem 1) We prove by induction the following statement. For a set L of at least two disjoint line segments in general position and for any fixed side yz of the polygon $conv(\bigcup L)$, there is a Hamiltonian polygon H such that yz is a side of H.

The statement holds for $|L| = 2$. Suppose it holds for all L' with $1 < |L'| < |L|$.

Consider the simple polygon P and the set S of polygons described in Lemma 2. If both endpoints of every segment are in $V(P)$, then the statement holds. If there is a segment ℓ whose neither endpoint is in $V(P)$, then by properties (*L2*) and (*L3*), the segment ℓ is in the interior of some $D \in S$. By property (*L5*), D has a common side ab with P which is not yz. Let $C(D) = conv(L \cap int(D))$. $C(D)$ has a side cd such that $ac \in E(G)$ and $bd \in E(G)$. If $|L \cap int(D)| = 1$ then cd is the unique segment in $int(D)$. If $|L \cap int(D))| > 1$ then, by induction, there is a Hamiltonian polygon $H(D)$ for $L \cap int(D)$ with side cd. Replace side ab of P by the path $ac \cup (\partial H(D) \setminus cd) \cup db$ for each $D \in S$ that contains segments from L. The resulting polygon H is a Hamiltonian polygon. $\square$

3 Basic definitions and operations

Our goal is to find a simple polygon satisfying the conditions of Lemma 2. In order to construct such a polygon, we run an algorithm which, in each step, makes local changes to our polygon, that is, replaces one edge by a path or two consecutive edges by one edge.

This algorithm, however, leads out from the family of simple polygons. We need to define a slightly wider class of polygons such that the boundary of the polygons do not have self-crossings but may have self-intersections.

P is a *polygon* if it is a simply-connected closed region in the plane which is the image of the unit disc D under a continuous mapping ρ, its boundary ∂P is a closed curve which is the image of the curve ∂D under ρ and consists of line segments; the endpoint of these segments are the *vertices* of the polygon. $V(P)$ is, as above, the cyclic sequence of vertices along ∂P and the *sides* of the polygon are the segments connecting two consecutive vertices of $V(P)$. The concept of *frame* polygons defined below has some other useful properties as well. All through our algorithm, we make sure that our intermediate polygons belong to this class.

Definition 3 *For a set L of disjoint line segments, a polygon P is called **frame**, if*

(F1) $V_L \subset P$ and $V(P) \subset V_L$;

(F2) ∂P does not any segment from L;

(F3) every vertex v of P appears at most twice in $V(P)$;

(F4) if $V(P) = (\dots abc \dots dbe \dots)$, then both $\angle eba$ and $\angle cbd$ are convex (possibly $a = e$ or $c = d$);

(F5) if $a \in V(P)$, and $b \in int(P)$ for some $ab \in L$, then the vertex a appears once in $V(P)$ and it is a convex vertex.

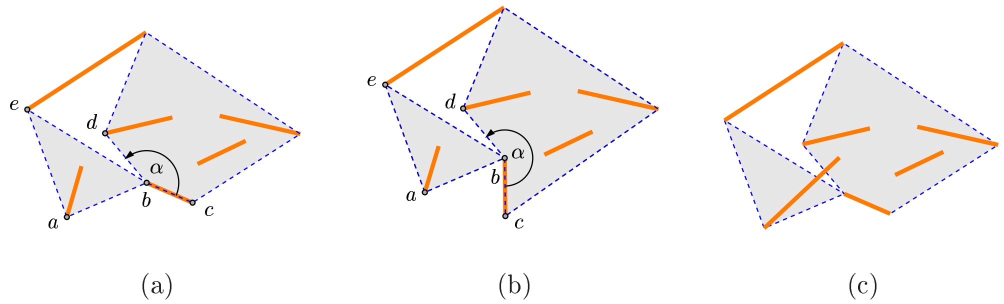

Figure 2: Examples for (non-)frame polygons.

For example, Figure 2(a) shows a frame polygon, while the polygons in Figure 2(b) $(\alpha > \pi)$ and 2(c) (crosses a segment) are not frames.

Proposition 4 *Let P be a frame for some set L of segments, and let b be a vertex visited twice by $V(P)$. Then at least one of the occurrences of b in $V(P)$ must be reflex.*

Proof. Let $V(P) = (\dots, a, b, c \dots, d, b, e, \dots)$ such that $\alpha := \angle cba$ is convex. According to property (*F3*), b can occur in $V(P)$ at most twice, and the general position assumption assures that no three points are collinear. Consider the rays $\overrightarrow{ba}$, $\overrightarrow{bc}$, $\overrightarrow{bd}$, and $\overrightarrow{be}$ (with possibly $a = d$ or $c = e$).

From property (*F4*) we know that e lies to the left of (or on) the oriented line $\overrightarrow{ab}$, and similarly, that d lies to the right of (or on) the oriented line $\overrightarrow{cb}$. AS ∂P has no self-crossing, the circular order of these rays around b is either $(\overrightarrow{ba}, \overrightarrow{bc}, \overrightarrow{bd}, \overrightarrow{be})$ or $(\overrightarrow{ba}, \overrightarrow{bd}, \overrightarrow{be}, \overrightarrow{bc})$. This second case can be excluded, because it would mean that one of the segments $\overline{bc}$ or $\overline{be}$ lies (at least partially) in the interior of P, whereas both are supposed to be part of ∂P. Thus $\angle ebd > 2\pi - \angle cba > \pi$. $\square$

Our first objective is to ensure property (*L2*). The method is really simple: start with the convex hull of L; whenever there is a line segment ℓ whose one endpoint is in $V(P)$ but the other is not, we extend the polygon locally to visit the other endpoint as well. This extension can be done in two different ways, which will be determined by an orientation defined as follows.

Definition 5 *For a frame polygon P, let an* ***orientation*** *$u(P)$ of P be a partial function $u : V(P) \to \{-1, 1\}$ defined for every $a \in V(P)$ such that $ab \in L$ and $b \in int(P)$.*
For an $a \in V(P)$, denote by a^{+1} and a^{-1} the next vertex of $V(P)$ in counterclockwise and clockwise direction respectively.

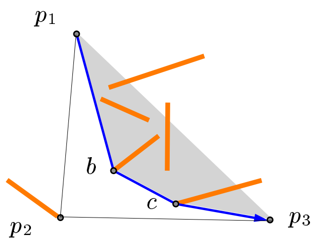

Figure 3: Example: $carc(p_1, p_2, p_3) = (p_1, b, c, p_3)$.

Definition 6 *Consider a polygonal arc (p_1, p_2, p_3) such that no segment of L crosses (p_1, p_2, p_3). Then define $carc(p_1, p_2, p_3)$ to be the shortest polygonal arc from p_1 to p_3 such that there is no segment endpoint in the interior of the closed polygonal curve $carc(p_1, p_2, p_3) \cup (p_3, p_2, p_1)$. (See the example in Figure 3.)*

It is easy to see that $carc(p_1, p_2, p_3) \cup (p_3, p_2, p_1)$ forms a simple polygon, moreover it is a *pseudo-triangle* where all internal vertices of $carc(p_1, p_2, p_3)$ are reflex.

Operation 1 (Build_cap(P, u, a)) *(Figure 4)*
Input: a frame polygon P, an orientation u, and a vertex $a \in V(P)$ such that $b \notin V(P)$, for $ab \in L$.

Let $c = a^{u(a)}$. Obtain P' from P by replacing the edge ac by the path $ab \cup carc(b, a, c)$. Set $u(p) := u(a)$ for all p on $carc(b, a, c)$.
Output: (P', u).

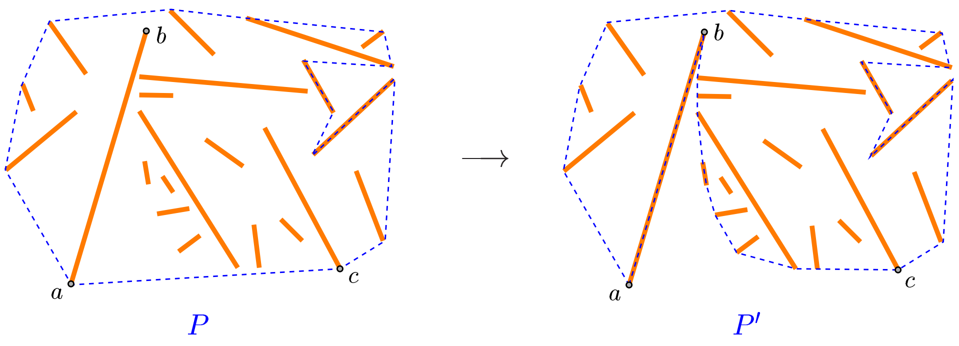

Figure 4: Build_cap(P, +1, a).

Observe that *Build_cap* produces exactly one reflex vertex in P', namely at b. Note also that P' is not necessarily simple, since some of the vertices from $carc(b, a, c)$ might already have been visited in $V(P)$.

Definition 7 *Let $k \in \mathbb{N}$ and $a, b_1, b_2, \ldots, b_k, c$ be a sequence of consecutive vertices of $V(P)$ such that b_i, $i = 1, \ldots k$, are reflex vertices and $conv(ab_1 \ldots b_k c) \cap L = \emptyset$. Then the sequence $b_1, b_2, \ldots, b_k$ is called* ***cap****.*
A reflex vertex of $V(P)$ that is not a cap is called ***anti-cap****.*

Proposition 8 *The output P' of* Build_cap *is a frame polygon.*

Proof. We have to check properties (*F1*)–(*F5*). (*F1*) and (*F2*) follow directly from the definition of *carc* and from the fact that the input polygon P is a frame.

Let $carc(b, a, c) = (b = p_0, \ldots, p_k = c)$ for some $k \in \mathbb{N}$. Operation Build_cap inserted vertices $p_0, \ldots, p_{k-1}$ into $V(P')$. Here $b \in V(P')$ is visited once, and the other endpoint of $ab \in L$ is also in $V(P)$; and $p_1, \ldots, p_{k-1}$ are inserted as convex vertices. This immediately implies property (*F5*) and (*F4*).

For (*F3*), we argue by contradiction. Suppose that p_i, $i \in \{1, \ldots\ k-1\}$, is already visited twice at the input P of Build_cap. By (*F4*), the angular domain around p_i intersects $int(P)$ in two convex angles. So by definition, p_i cannot be in $carc(b, a, c)$.
□

Operation 2 (Both_endpoints(P, u))
Input: a frame polygon P and an orientation $u(P)$.
As long as there exists an $a \in V(P)$ such that $ab \in L$ and $b \notin V(P)$,
let $(P, u) \leftarrow$ *Build_cap*(P, u, a).
Output: (P, u).

Proposition 9 Both_endpoints *does not create a new anti-cap in P. Sequences of consecutive caps form one cap.*

Proof. *Build_cap* produces exactly one new reflex vertex: b. Let $carc(b, a, c) = (b = p_0, \ldots, p_k = c)$ for some $k \in \mathbb{N}$. b is a cap because $\triangle(a\,b\,p_1) \cap L = \emptyset$ by construction.

By property (*F5*), vertex a is convex in $V(P)$ before and after *Build_cap*. In fact, all the vertices $a, p_1, \ldots, p_{k-1}$ are convex in $V(P)$. Hence, there is nothing more to show, if $k > 1$. So let us consider the case $k = 1$, that is $carc(b, a, c) = bc$. Suppose also that c is a reflex vertex of $V(P)$, which is part of a cap $c = c_1, c_2, \ldots, c_k$. Denote the other neighbor of c_k by d. If $\angle bcd < \pi$, then c is not a cap anymore. Otherwise we have $conv(abcd) \cap L = \emptyset$, because P is a frame polygon, $conv(ac_1c_2 \ldots c_kd) = \emptyset$, and $ab \in L$. Hence $b, c_1, c_2, \ldots, c_k$ is a cap. $\square$

Knowing that every sequence of double occurrences in list $V(P)$ corresponds to a cap (i.e., a wedge), it is easy to make a simple polygon from P by the following operation.

Definition 10 *A sequence $a, b_1, b_2, \ldots, b_k, c$ of consecutive vertices from a frame polygon P is called* ***wedge****, if $b_1, b_2, \ldots, b_k$ is a cap of P such that each b_i, $i = 1, 2, \ldots, k$, is visited twice by $V(P)$.*

Operation 3 (Chop_wedges(P)) *(Figure 5)*
Input: a frame polygon P.
As long as there is a wedge $a, b_1, b_2, \ldots, b_k, c$,

- *Obtain P' from P by replacing the path $a, b_1, b_2, \ldots, b_k, c$ by the single edge ac.*
- *Let $P \leftarrow P'$.*

Output: P.

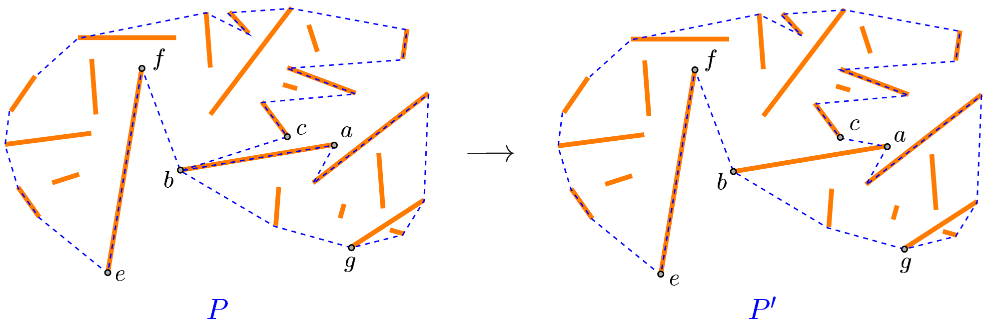

Figure 5: Chopping a wedge.

4 Extremally situated segments and more

In this section, we describe a simple algorithmic proof for the case of extremally situated segments. The procedure then serves as a base step for solving the general problem.

Algorithm 1
Input: a set L of disjoint line segments and an orientation u for $conv(\cup L)$.

- *Put* $P_0 \leftarrow conv(\cup L)$.
- *Let* $P_1 \leftarrow$Both_endpoints(P_0, u).
- *Let* $P_2 \leftarrow$Chop_wedges(P_1).

Output: P_2.

Proposition 11 *The output P_2 of Algorithm 1 is a simple frame polygon with property (L2).*

Proof. Property (*L2*) follows from the loop condition in *Both_endpoints* and the fact that *Chop_wedges* does not alter the set of visited vertices. P is simple because, by Proposition 4, for every vertex that is visited twice by $V(P_1)$, at least one occurrence is reflex. Proposition 9 tells us that every sequence of consecutive reflex vertices in $V(P_1)$ form a cap, and thus all repetitions in $V(P)$ are deleted by *Chop_wedges.* □

Corollary 12 *[10] If the line segments of L are extremally situated, then Algorithm 1 outputs a Hamiltonian polygon for any orientation u for $conv(\cup L)$.* □

Note that we did not make any use of the orientation u for the proof of Corollary 12. We could simply run Algorithm 1 with a uniform orientation $u \equiv 1$. But in this case we cannot guarantee that a prescribed side yz of $conv(\cup L)$ is a side of the output polygon, as required in (*L1*). (See Figure 9 for an example.)

Let the orientation u_{yz} of $conv(\cup L)$ be such that $u_{yz}(y) = -1$ and $u_{yz}(v) = 1$ for any other vertex of $conv(\cup L)$.

Proposition 13 *If Algorithm 1 is applied to L with orientation u_{yz}, then the output P_2 is a simple frame polygon satisfying properties (L1) and (L2).*

Proof. yz is a side of $P_0 = conv(\cup L)$, and none of the *Build_cap* operations replaces yz by something else. Moreover, both y and z remain convex vertices throughout *Both_endpoints.* Since *Chop_wedges* does only cut off edges adjacent to reflex vertices, the edge yz remains part of P_2 as well. □

Proposition 14 *If Algorithm 1 is applied to L with orientation u_{yz}, then the output P_2 has at most one cap with exactly two reflex vertices (double-cap), and all other caps contain exactly one reflex vertex.*

Proof. An operation *Build_cap*(P, u, a) creates exactly one new reflex vertex, namely at b where $ab \in L$. Let $c := a^{u(a)}$. The vertex a is convex in both P and P', hence we can have two consecutive reflex vertices only if $carc(b, a, c) = bc$ and c is a reflex vertex of P. Assuming this scenario, the reflex vertex c is created in a previous operation *Build_cap*$(\tilde{P}, \tilde{u}, d)$ such that in $\tilde{P}$ we had $a^{\tilde{u}(a)} = d$, $d^{\tilde{u}(d)} = a$ and $carc(c, d, a) = ca$.

This already implies that there is no cap of three consecutive vertices. The orientation u changes from anti-clockwise to clockwise exactly once, since *Build_cap* does not induce alternations in the orientation. Thus, there is at most one *double cap.* □

5 Dissecting P

For every frame polygon P, we can define a set $Diss(P)$ of non-overlapping simple polygons which together cover $int(P)$. The frame P is not necessarily simple, because it may have multiple vertices at *wedges*. Call a diagonal ab of P a *segment diagonal* if $ab \in L$ and $ab \subset P$. The heads of wedges and the segment diagonals dissect P into non-overlapping simple polygons. Let $Diss(P)$ be the set of these polygons. Observe that $Diss(P)$ satisfies property (*L3*).

Unfortunately, the polygons of $Diss(P)$ are not necessarily convex. The easy idea to obtain a set of convex polygons by dissecting elements $D \in Diss(P)$ is the following. At every reflex vertex b of every $D \in Diss(P)$, draw consecutively rays starting from b dissecting $\angle_D b$ into two convex angles until the ray hits the boundary of D or a previously drawn ray. If no ray crosses a segment of $L \cap int(D)$, then they dissect D into non-overlapping convex regions satisfying properties (*L1*), (*L3*), and (*L4*). The resulting partition depends on the order in which the rays are drawn, but any order would do at this point. If any of the rays crosses a segment of L, then such a partitioning would not grant (*L3*) of Lemma 2. For this case we introduce two new basic operations which extend the frame polygon by new segments.

5.1 Convex cells of S

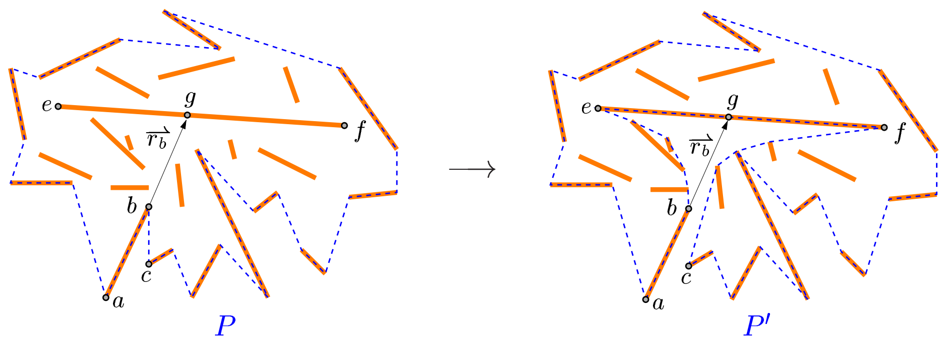

Figure 6: *Extend_cap*$(P, u, S, b, c, \vec{r_b})$ to a segment ef.

Definition 15 *Consider a polygonal arc* (p_1, p_2, p_3, p_4) *such that no* $\ell \in L$ *crosses* (p_1, p_2, p_3, p_4). *Then define* $marc(p_1, p_2, p_3, p_4)$ *as the shortest polygonal arc from* p_1 *to* p_4 *such that there is no segment endpoint in the interior of the closed polygonal curve* $marc(p_1, p_2, p_3, p_4) \cup (p_4, p_3, p_2, p_1)$.

Similarly to *carc*, the polygon $marc(p_1, p_2, p_3, p_4) \cup (p_4, p_3, p_2, p_1)$ has reflex vertices at internal vertices of $marc(p_1, p_2, p_3, p_4)$, but $marc(p_1, p_2, p_3, p_4)$ is not necessarily simple, p_1 or p_4 may occur twice on the arc. (See Figure 7 for two examples.)

Operation 4 (Extend_cap$(P, u, S, b, c, \overrightarrow{r_b})$) *(Figure 6)*
Input: a frame P along with an orientation u, a set S of non-overlapping polygons in P, a reflex vertex b of some $D \in S$ that is visited exactly once by $V(P)$, a vertex c,

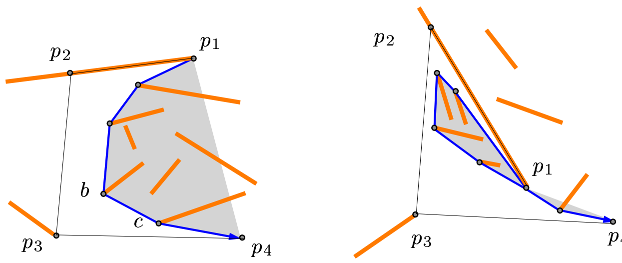

(a) (b)

Figure 7: $marc(p_1, p_2, p_3, p_4)$ for convex and concave quadrilateral $p_1p_2p_3p_4$.

and a ray $\overrightarrow{r_b}$ emanating from b.
Preconditions: P satisfies property (*L2*), bc is a common side of D and $V(P)$, $\overrightarrow{r_b}$ cuts $\angle cba$ into two convex angles, c is a convex vertex of $V(P)$, and ray $\overrightarrow{r_b}$ hits segment $ef \subset int(P)$ at point g. We may suppose that c and f are on the same side of the supporting line of $\overrightarrow{r_b}$.
Operation: Obtain P' from P by replacing the edge bc by the path $carc(b, g, e) \cup ef \cup marc(f, g, b, c)$. Set $u(\cdot) := -1$ for all interior vertices of $carc(b, g, e)$, and $u(\cdot) := 1$ for all interior vertices of $marc(f, g, b, c)$.
Output: (P', u).

Remark 16 There are two variants of *Extend_cap*, depending on whether c precedes or follows c in $V(P)$. We have only described the latter above, the former variant is completely symmetric.

Proposition 17 *The output P' of* Extend_cap *is a frame polygon.*

Proof. Observe that $marc(f, g, b, c)$ is a simple polygonal arc, i.e., f and c occur only once on the arc (as endpoints). It obviously holds if quadrilateral $fgbc$ is a convex polygon. If $fgbc$ has a reflex vertex at c then s is still not visited once more since c is a convex vertex of $V(P)$; if f is a reflex vertex of $fgbc$ then f might be visited twice by $marc(f, g, b, c)$. (*F1*) and (*F2*) of Definition 3 then follow from the definition of *carc* and *marc*.

For internal vertices of $carc(b, g, e)$ and $marc(f, g, b, c)$, one can argue analogously as in Proposition 8. Hence, we have to consider b, c, e, and f only.

Since neither e nor f has been visited by $V(P)$ before, (*F3*) is obvious. For (*F4*) note that b is visited once only in both P and P', and that the angles in (*F4*) for c can only decrease. Finally, observe that *Extend_cap* creates only two new reflex vertices in $V(P')$: e and f. Since the segment ef becomes part of $\partial P'$, (*F5*) follows. □

Proposition 18 Extend_cap *creates at most one new anti-cap.*

Proof. *Extend_cap* makes b a convex vertex and keeps c convex. It creates, however, two new reflex vertices: e and f. We will show that at least one of e or f is a cap in $V(P)$. Denote by d the second vertex of $carc(e, g, b)$, and by h the second vertex of

$marc(f, g, b, c)$ (possibly $d = b$ or $h = c$), and let v be the intersection of the supporting line of de and hf (Figure 8). Since the edges ed and fh do not cross by definition, we have $d \in \overline{ve}$ or $h \in \overline{vf}$. In the first case $df \in E(G)$ and f is a cap, and in the second case $he \in E(G)$ and e is a cap. □

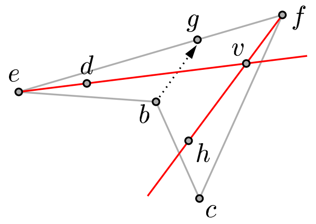

Figure 8: Illustration for Proposition 18.

Corollary 19 *If $g = e$ in operation* Extend_cap*, then f is a cap of P'.*

5.2 Common side for each $D \in S$ and P

So far we have a frame P and a dissection S of P fulfilling Properties ($L1$)–($L4$). Unfortunately, P does not always have property ($L5$), as can be seen in Figure 9. Some further modification of the polygon P can, however, assure ($L5$) as well.

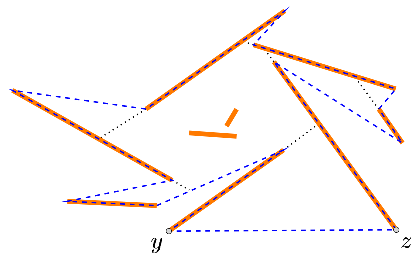

Figure 9: No common side with the frame.

Operation 5 (Mend_cap(P, u, S, b, $\overrightarrow{r_b}$, cd)) *(Figure 10)*
Input: a frame P with an orientation u, a set S of non-overlapping convex sets in P, a cap b that is visited exactly once by $V(P)$, a ray $\overrightarrow{r_b}$ emanating from b, and the segment cd of ∂P hit by the ray $\overrightarrow{r_b}$.
Preconditions: P satisfies property ($L2$), cd is a common side of $V(P)$ and some $D \in S$, d is a reflex vertex and c is a convex vertex of $D \in S$, $\overrightarrow{r_b}$ cuts the reflex angle at b into two convex angles, c is a convex vertex of $V(P)$. Let q be the point where $\overrightarrow{r_b}$ hits cd.
Operation: Obtain P' from P by replacing the edge cd by the path $carc(c, q, b) \cup carc(b, q, d)$. Set $u(.) := -1$ for all interior vertices of $carc(c, q, b)$ and $u(.) := 1$ for all interior vertices of $carc(b, q, d)$.
Output: (P', u).

Proposition 20 *The output P' of* Mend_cap *is a frame polygon.*

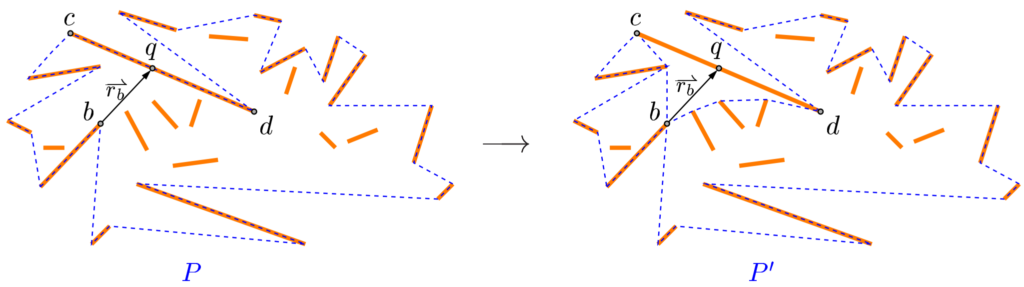

Figure 10: Mending a cap.

Proof. We have to check properties (*F1*)–(*F5*). (*F1*) and (*F2*) are obvious from the definition of *carc*. For internal vertices of convex arcs, one can argue analogously as in Proposition 8. Hence, we have to consider b, c and d only.

The vertex b was visited once only by $V(P)$, hence it is visited twice by $V(P')$. Since c and d are not revisited in $V(P')$ compared to $V(P)$, (*F3*) follows. For (*F5*) note that *Mend_cap* does neither create any new reflex vertex in $V(P)$ nor does it revisit vertices from $V(P)$, except for the fact that b appears now twice as a reflex vertex in $V(P')$. But since b was already a reflex vertex in $V(P)$, the conditions of (*F5*) do not apply. Finally, (*F4*) is clearly true for b, since for both visiting paths, the adjacent vertices are on different sides of the line through bq. For both c and d, the angles mentioned in (*F4*) can only decrease. Hence, (*F4*) holds for all vertices in $V(P')$. □

Proposition 21 Mend_cap *creates at most one new anti-cap.*

Proof. The operation does not create any new reflex vertices, so only the two existing reflex vertices b and d might become anti-caps. But by the definition of *carc*, the new occurrence of b in $V(P')$ is a cap. □

6 Algorithm and its analysis

Algorithm 2

Input: a set L of disjoint line segments and a side yz of $conv(\cup L)$.

P	$\leftarrow$	$conv(\cup L)$.	*(frame polygon)*
S	$\leftarrow$	$\{P\}$.	*(dissection)*
(a, b, c)	$\leftarrow$	$\emptyset$.	*(vertex + adjacent reflex vertex + adjacent vertex)*
u	$\leftarrow$	u_{yz}.	*(orientation)*

Repeat until every $D \in S$ is convex in step c below.

a) $(P, u) \leftarrow$ Both_endpoints(P, u).

b) Update S by replacing each $D \in S$ by $Diss(D)$.

c) If every $D \in S$ is convex, then $P \leftarrow$ Chop_wedges(P) *and exit.*

d) If $(a, b, c) = \emptyset$, then

(1) If there is a double-cap k, l in some $D_b \in S$, then $(a, b, c) \leftarrow (k, l, m)$, where m is the other ($\neq k$) neighbor of l in ∂D_b; and $\overrightarrow{r_b} \leftarrow \overrightarrow{ab}$.

(2) *Else let b be a reflex vertex of some $D_b \in S$, and let a and c be the adjacent (in ∂D_b) convex vertices, such that c is also adjacent to b in $V(P)$ (see Proposition 24); and $\overrightarrow{r_b} \leftarrow \overrightarrow{ab}$.*

e) *If ray $\overrightarrow{r_b}$ hits a segment $ef \subset int(D_b)$, then*

(1) *If the supporting line of ef crosses segment bc, then rotate $\overrightarrow{r_b}$ around b until it hits an endpoint of a segment $e'f' \subset int(D_b)$ or ∂D_b.*

(2) *If ray $\overrightarrow{r_b}$ still hits a segment $ef \subset int(D_b)$, then*

i, $(P, u) \leftarrow \text{Extend\_cap}(P, u, S, b, c, \overrightarrow{r_b})$.

ii, If Extend_cap *created an anti-cap h in P, then $b \leftarrow h$; $c \leftarrow$ a convex neighbor, and $a \leftarrow$ the other neighbor of a in $V(P)$;*

iii, else $(a, b, c) \leftarrow \emptyset$.

f) *If ray $\overrightarrow{r_b}$ hits ∂D_b at point g on side de, then*

(1) *Dissect D_b by $\overline{bg}$ and update S accordingly.*

(2) *If one of d or e (say e) is a reflex vertex of D_b, then*

i, If b is a cap, then $(P, u) \leftarrow \text{Mend\_cap}(P, u, S, b, \overrightarrow{r_b}, de)$.

ii, $b \leftarrow e$; update D_b; $a \leftarrow d$, $c \leftarrow$ other ($\neq d$) neighbor of e in ∂D_b.

(3) *Else $(a, b, c) \leftarrow \emptyset$.*

Output: (P, S).

Proposition 22 *Algorithm 2 terminates.*

Proof. If P is changed in step a, at least one segment endpoint is added to $V(P)$ that was not visited before. As no vertex ever leaves $V(P)$, these changes can only occur in a finite number of steps. Apart from this, either step f or step e is executed in every iteration. In the latter, $V(P)$ is augmented by a segment that was in the interior of P before. In the former, a reflex angle of a region $D_b \in S$ is destroyed, while no new reflex angle is added. Hence, after a finite number of iterations, every $D \in S$ will be convex and the algorithm terminates. □

Proposition 23 *Whenever* $\text{Mend\_cap}(P, u, S, b, \overrightarrow{r_b}, de)$ *or* $\text{Extend\_cap}(P, u, S, b, c, \overrightarrow{r_b})$ *are called in Algorithm 2, the vertex b is visited only once by $V(P)$.*

Proof. b is always chosen as a reflex vertex of some $D_b \in S$. If b would be visited twice by $V(P)$, it cannot be a reflex vertex of any $D \in S$ by property (*F4*). □

Proposition 24 *If a, b, c are three consecutive vertices in $V(P)$ where b is a cap and also a reflex vertex of some $D \in S$, then either ab or bc is a side of D during Algorithm 2.*

Proof. The side ab (or bc) is not a side of D if and only if the ray drawn from a previous anti-cap hit it. (If the side was hit from a cap b', *Mend_cap* would have been applied to b' and ab or bc would not have been an edge of P anymore.) Algorithm 2 is organized so that right after the ray from an anti-cap hits, say, side ab (step *f2*), it shoots a ray from b in the next step, such that from there on, b is no longer a reflex vertex of any set in S. But then, side bc is still a common side of $V(P)$ and D. □

Corollary 25 *In step f(2)i of Algorithm 2, segment de is a common side of D_e and $V(P)$ (hence all preconditions of* Mend_cap$(P, u, S, b, \overrightarrow{r_b}, de)$ *are satisfied).*

Lemma 26 *During Algorithm 2, the frame P never visits an anti-cap twice.*

Proof. An anti-cap can be created in two places only: in *Extend_cap* (step *e2*), or in *Mend_cap* (step *f(2)i*). In both cases, at most one anti-cap is created. Consider such an anti-cap e. $V(P)$ visits e exactly once directly after the operation that created it. In the next iteration, Algorithm 2 dissects the region $D \in S$ containing e along a ray emanating from e. Hence, e will appear as a convex vertex in any $D \in S$ from there on; and no operation ever revisits a convex vertex of any $D \in S$.

The only point where e could possibly be revisited by $V(P)$ is in the call to *Both_endpoints* (step *a*) immediately following the step where e became an anti-cap. We argue that the orientation u along *carc* and *marc* are set such that $V(P)$ cannot revisit e in any of the resulting *Build_cap* operations.

Consider operation *Extend_cap*, the argument is analogous for *Mend_cap*. First we show that *Both_endpoints* applied to vertices of $carc(b, g, e)$ does not revisit e. Recall that $u(k) = -1$, for all $k \in carc(b, g, e)$, and *Build_cap* preserves this orientations for all new vertices.

Denote the output of *Extend_cap* by P_0, and let e_* be the penultimate vertex on $carc(b, g, e)$. Let t be the point where $\overrightarrow{e_*e}$ hits ∂P; the segment et dissects P into two parts. We define recursively an arc $\varepsilon(k)$ for each vertex vertex k reached by *Both_endpoints* from $carc(b, g, e)$. Suppose that k is reached by a $carc(p'_1, p_1, p^1_1)$ in a step *Build_cap*$(P, -1, p_1)$. Let $\varepsilon(k)$ follow follow $carc(p'_1, p_1, p^1_1)$ from k to p_1, then follow $\varepsilon(p_1)$ to b. For every three consecutive vertices p_1, p_2, p_3 of an arc $\varepsilon(k)$, $\angle p_1p_2p_3$ is convex; and $\varepsilon(k) \subset P_0$.

If e is revisited by *Both_endpoints*, then $\varepsilon(e)$ is defined from e to b. Also, the angle $\angle(carc(b, g, e), \varepsilon(e))$ at e of $\gamma(e)$ is convex, since $e \in carc(p'_1, p_1, p^1_1)$. Hence $\varepsilon(e)$ cannot cross segment et, so cannot reach b. Contradiction.

For the case of $marc(f, g, b, c)$, observe that if $\overrightarrow{ef}$ hits side bc, then $g = e$ by the rotation of $\overrightarrow{r_b}$; and by Corollary 19, f is a cap of P. If $\overrightarrow{ef}$ does not hit bc, then by the above argument one can show that *Both_endpoints* applied to vertices of $marc(f, g, b, c)$ does not revisit f. □

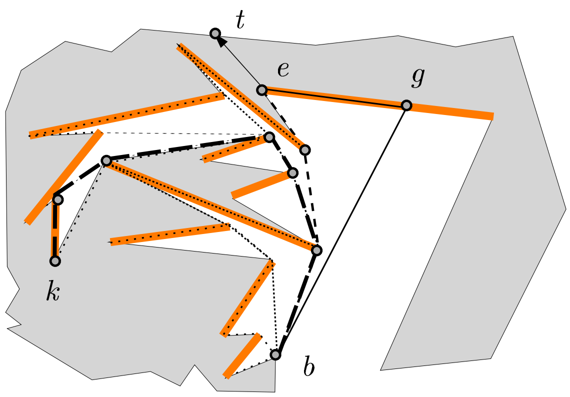

Figure 11: Illustration for Lemma 26.

At the last step of Algorithm 2, *Chop_wedges* is applied. Lemma 26 assures that any vertex, that is visited twice by $V(P)$, is adjacent to a wedge that can be chopped off. That is, the output of Algorithm 2 is a simple frame polygon. To show that the output P and the partition S satisfy the properties of Lemma 2, it suffices to prove the following.

Lemma 27 *All through Algorithm 2, every $D \in S$ has a common side with P which is different from wedge edges and the special side yz of P.*

Proof. The statement holds for $conv(\cup L)$. It is enough to check that it remains true after each operation.

Build_cap, *Mend_cap*, or *Extend_cap* may dissect a region $D \in S$ into several regions if the arc $carc(p_1, p_2, p_3)$ or $marc(f, g, b, c)$ revisits caps, thereby reverting sides of D to wedge-edges. Still, in each new region $D' \subset D$, $carc(p_1, p_2, p_3)$ or $marc(f, g, b, c)$ have a common side with both D' and P.

In step *f1* of Algorithm 2, the region $D_b \in S$ is dissected into regions D_e and D_d by the ray $\overrightarrow{r_b}$, where b is a reflex vertex of both D_b and P. We have to check that our statement still holds for both D_e and D_d. Denote the neighbors of b in P by α and γ. (The neighbors of b in P might be different from those in D_b.)

If b is an anti-cap, then αb and $b\gamma$ are sides of D_e and D_d, respectively, since Algorithm 2 draws the ray $\overrightarrow{r_b}$ right after the path $\alpha b \gamma$ is created. αb and $b\gamma$ are clearly neither wedge edges nor equal to yz.

So suppose that b is a cap and, say, αb is not a side of D_d. This means that a previously drawn ray $\overrightarrow{r_{b'}}$ from a reflex vertex b' hits αb. Then b' is necessarily an anti-cap, since otherwise *Mend_cap* would have been applied to b' and αb would not have been a side of P any longer.

But then we can show that a side of $V(P)$ adjacent to b' is a common side of P and D_d. Anti-cap b' is also on the boundary of D_d, as the dissection by $\overrightarrow{r_b}$ immediately follows the dissection by $\overrightarrow{r_{b'}}$ in Algorithm 2 and no other ray (including $\overrightarrow{r_b}$) could hit $\overrightarrow{r_{b'}}$. Hence ∂D_d contains at least part of a side $a'b'$ of $V(P)$. Again, $a'b'$ is a common side of $V(P)$ and D_d, because no other ray (including $\overrightarrow{r_b}$) could hit $a'b'$ between these two iterations. □

References

[1] P. K. Agarwal, N. Alon, B. Aronov, and S. Suri, Can visibility graphs be represented compactly? *Proc. 10th ACM Symp. on Comput. Geom. (San Diego, 1993)*, and *Discrete Comput. Geom.* **12** (1994), 347–365.

[2] H. Alt and E. Welzl, Visibility graphs and obstacle-avoiding shortest paths, *Z. Oper. Res.* **32** (1988), 145–164.

[3] T. Asano, S. K. Ghosh, and T. C. Shermer, Visibility in the plane, in: *Handbook of Computational Geometry (J.R. Sack, J. Urrutia eds.)*,Elsevier, Amsterdam, 2000, 829–876.

[4] P. Bose, M. E. Houle, and G. Toussaint, Every set of disjoint line segments admits a binary tree, Algorithms and computation (Beijing, 1994), *Lecture Notes in Comput. Sci.*, Vol. 834, Springer, Berlin, 1994, 20–28.

[5] E. D. Demaine and J. O'Rourke, Open problems from CCCG'99, *Proc. 11th Canadian Conf. on Comput. Geom.*, (Vancouver, 1999).

[6] S. Ghali and A. J. Stewart, Maintenance of the set of segments visible from a moving viewpoint in two dimensions, In: *Proc. 12th ACM Symp. Comput. Geom.*, (Philadelphia, 1996), pp. V3–V4.

[7] S. K. Gosh and D. M. Mount, An output-sensitive algorithm for computing visibility graphs, *SIAM J. Comput.* **20** (1991), 888–910.

[8] H. Everett, C. T. Hoang, K. Kilakos, and M. Noy, Planar segment visibility graphs, *Comput. Geom.* **16** (2000), 235–243.

[9] M. Keil, D. M. Mount, and S. K. Wismath, Visibility stabs and depth-first spiralling on line segments in output sensitive time. *Internat. J. Comput. Geom. Appl.* **10** (2000), 535–552.

[10] A. Mirzaian, Hamiltonian triangulations and circumscribing polygons of disjoint line segments, *Comput. Geom.* **2** (1992), 15–30.

[11] J. O'Rourke. Visibility, in: *Handbook of Discrete and Computational Geometry (J.E. Goodman, J. O'Rourke, eds.)*, CRC Press, Boca Raton, 1997, 467–480.

[12] J. O'Rourke, J. Rippel, Two segment classes with Hamiltonian visibility graphs, *Comput. Geom.* **4** (1994), 209–218.

[13] M. H. Overmars and E. Welzl, New methods for computing visibility graphs, In: *Proc. 4th ACM Symp. Comput. Geom.* (Urbana-Champaign, 1988), 164–171.

[14] D. Rappaport, Computing simple circuits from a set of line segments is NP-complete,*SIAM J. Comput.* **18** (1989), 1128–1139.

[15] D. Rappaport, H. Imai, and G. Toussaint, Computing simple circuits from a set of line segments, *Discrete Comput. Geom.* **5** (1990), 289–304.

[16] X. J. Shen and H. Edelsbrunner, A tight lower bound on the size of visibility graphs, *Inform. Process. Lett.* **26** (1987), 61–64.

[17] M. Urabe and M. Watanabe, On a counterexample to a conjecture of Mirzaian, *Comput. Geom.* **2** (1992), 51–53.

[18] E. Welzl, Constructing the visibility graph for n line segments in $O(n^2)$ time. *Inform. Process. Lett.* **20** (1985), 167–171.