Segment endpoint visibility graphs are Hamiltonian

Michael Hoffmann* Csaba D. Tóth*

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

Visibility graphs represent an important tool in computational geometry. They have applications in VLSI design, computer graphics, motion planning, and art gallery problems. Visibility concerning disjoint line segments in the plane is fundamental, and problems for more complex objects can often be reduced to or approximated by this structure.

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, iff the corresponding line segment ab is either a given segment or does not cross any given segment. Two segments cross, iff they have a common point which lies in the interior of both segments. Note that this graph is different from the *segment visibility graph*, where vertices correspond to segments and an edge connects two vertices iff some points of the two segments are mutually 'visible'.

We are interested in a problem posed by Mirzaian [2], who conjectured the following. 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 [2] proved that there is a Hamiltonian polygon, if the line segments are *extremally situated*, that is, if none of the given line segments are in the interior of the convex hull $int(conv(\bigcup L))$.

O'Rourke and Rippel [3] proved that there is a Hamiltonian polygon, if the supporting line of no line segment crosses any other line segment. Urabe and Watanabe [4] gave a construction where L does not always have a Hamiltonian polygon which contains $\bigcup L$, i.e. a *circumscribing polygon*. We prove the conjecture of Mirzaian.

[1]Department of Computer Science, ETH Zurich, CH-8092 Zurich, Switzerland, {`hoffmann,toth`}`@inf.ethz.ch`.

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

The segments are in *general position* if there is no three collinear segment endpoints. In Section 2, we prove Theorem 1 by induction. The key lemma of the proof, our Lemma 2, is proved algorithmically in three phases. The sketch of the algorithm and the basic operations can be found in Sections 3–4.

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. Let P denote a closed connected polygonal domain (shortly *polygon*). Let $V(P)$ be the cyclic list of the vertices of P along the boundary in counterclockwise order, and let ∂P denote the closed polygonal path visiting one-by-one the vertices in $V(P)$. A polygon P is simple if ∂P has no self-intersections. If P is simple and every vertex of $V(L)$ appears in $V(P)$, 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 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 Both endpoints

Our goal is to find a simple polygon satisfying the properties of Lemma 2. In order to exhibit one such polygon, we run an algorithm which, in each step, makes local changes to our polygon, that is replace 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 polygons do not have self-crossings but may have self-intersections at vertices. 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 cross itself nor any segment from L;

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

(F4) if $V(P) = (\ldots abc \ldots dbe \ldots)$, 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.

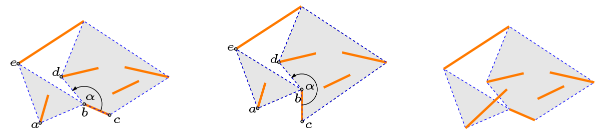

(a) (b) (c)

Figure 1: Examples for (non-)frame polygons.

For example, Figure 1(a) shows a frame polygon, while the polygons in Figure 1(b) ($\alpha > \pi$) and 1(c) (crosses a segment) are not frames.

Our fist 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, 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 4 *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, resp., direction.

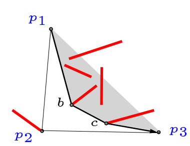

Figure 2: Convex arc example.

Definition 5 *Consider a polygonal arc (p_1, p_2, p_3) such that no $\ell \in L$ 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 $carc(p_1, p_2, p_3) \cup (p_1, p_2, p_3)$. (See the example in Figure 2.)*

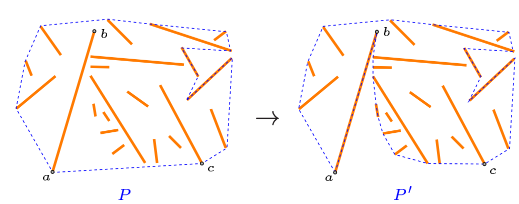

Figure 3: Building a cap.

Operation 1 (Build_cap$(P, u, a, u(a))$)
Input: a frame P, its orientation u, a vertex $a \in V(P)$ such that $b \notin V(P)$, for $ab \in L$, and $u(a) \in \{-1, 1\}$.
Let $c = a^{u(a)}$. Replace one edge by a path: $ac \to ab \cup carc(b, a, c)$. Let $u(p) = u(a)$ for all $p \in carc(b, a, c)$.
Output: (P, u). (See Figure 3.)

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

Definition 6 *Let $k \in \mathbf{N}$ and $a, b_1, b_2, \ldots, b_k, c$ be a sequence of consecutive vertices of a polygon P such that b_i, $i = 1, \ldots k$, are reflex vertices of $V(P)$ and $conv(ab_1 \ldots b_k c) \cap L = \emptyset$. Then the sequence $b_1, b_2, \ldots, b_k$ is called a **cap**.*
*A reflex vertex of $V(P)$ that is not a cap is called an **anti-cap**.*

We can now assure property (*L2*).

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

Proposition 7 *As a result of operation* Both_endpoints, *there is no new anti-cap, and sequences of consecutive caps form one cap.*

Knowing that every double occurrence in the sequence $V(P)$ corresponds to a cap, it is easy to make a simple polygon from P by the following operation. (See Figure 4.)

Definition 8 *A polygonal path $ab_1b_2 \ldots b_k c$ is called a **wedge** of polygon P, 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 $V(P)$ and a and c are the other neighbors of b_1 and b_k, resp.*

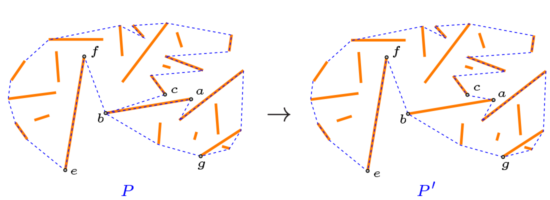

Figure 4: Chopping a wedge.

Operation 3 (Chop_wedges)
Input: Frame P.
As long as there is a wedge $ab_1b_2 \ldots b_k c$, replace a path by one side in P: $ab_1b_2 \ldots b_k c \to ac$.
Output: P. (See Figure 4.)

4 Extremal segments

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: Set of disjoint line segments L and an orientation u for $conv(\cup L)$.

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

Output: P_2.

Proposition 9 *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 Proposition 9. 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 property (*L1*). See Figure 5 for an example.

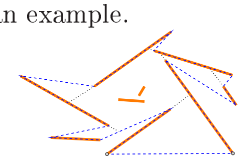
Figure 5: A pair of polygons.

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

Proposition 10 *The output of Algorithm 1 with orientation $u(yz)$ is a simple frame polygon satisfying properties (L1) and (L2).*

Proposition 11 *The output of Algorithm 1 with orientation $u(yz)$ has at most one cap with exactly two reflex vertices, and all other caps contain exactly one reflex vertex.*

5 Two more operations

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

But the polygons of $S'(P)$ are not necessarily convex. The easy idea to obtain a set of convex polygons by dissecting elements $D \in S'(P)$ is the following. Draw consecutively rays starting from every reflex vertex b of every $D \in S'(P)$ until it hits the boundary of D or a previously drawn ray such that the ray dissects $\angle_D b$ into two convex angles. If no ray crosses a segment of $L \cap int(D)$, then they dissect every D into non-overlapping convex regions satisfying properties ($L1$), ($L3$), and ($L4$). If any of the rays crosses a segment of L, then we can extend our polygon by the following basic operation.

5.1 Convex cells of S.

Operation 4 (Extend_cap$(P, u, S, b, c, \overrightarrow{r_b})$**)**
Input: a frame P, 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 $V(P)$, vertex c, such that bc is a common side of $V(P)$ and D_b, and a ray $\overrightarrow{r_b}$ emanating from b.
Let g be the point where B_b hits ef. We may suppose that a and e (c and f, resp.) are on the same side of B_b. Replace the edge bc of P by a path: $bc \to carc(bge) \cup ef \cup carc(fgbc)$. Let $u(v) = -1$ for all interior vertex v of $carc(bge)$ and $u(w) = 1$ for all interior vertices w of $carc(fgbc)$.
Output: (P, u). (See Figure 6.)

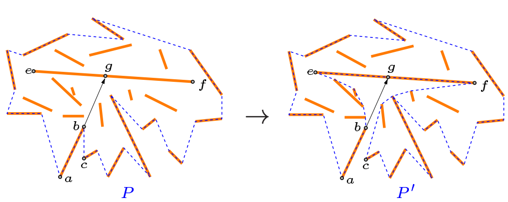

Figure 6: Extending a cap.

Proposition 12 Extend_cap *creates at most one new anti-cap.*

5.2 One side to each $D \in S$.

Unfortunately, P does not always have property ($L5$), as can be checked for the example in Figure 5. Some further modification of the polygon P can, however, assure ($L5$) as well.

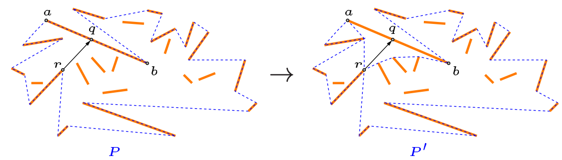

Figure 7: Mending a cap.

Operation 5 (Mend_cap$(P,\, u,\, S,\, b,\, \overrightarrow{r_b},\, cd)$**)**
Input: a frame P, an orientation u, s set of non-overlapping sets of polygons in P, a cap b that is visited exactly once by $V(P)$, a ray $\overrightarrow{rb}$, and cd is the segment of ∂P hit by $\overrightarrow{r_b}$ at point q. Replace edge ab of $V(P)$ by a path: $carc(cqb) \cup carc(bqd)$. Let $u(\cdot) = -1$ (resp., $u(\cdot) = 1$) for all interior vertices of $carc(cqb)$ (resp., $carc(bqd)$).
Output: (P, u). (See Figure 7.)

5.3 Sketch of main algorithm.

First observe that all four basic operations (presented above) preserve the frame property of P, and that no operation creates more than one new anti-cap.

Combining the operations properly allows to compute the dissection S of P into convex polygons, such that at any time during the algorithm

- every $D \in S$ has at least one side s in common with P, where s is neither part of a wedge nor the special side yz,
- no anti-cap of $V(P)$ is visited twice,
- and there is at most one anti-cap that is a reflex vertex of some $D \in S$.

The algorithm considers the reflex vertices of every $D \in S$ in turn, dissecting them into two convex angles using a ray, and applying *Extend_cap* or *Mend_cap* depending on what is hit by the ray.

For lack of space, we have to refer to the full paper for the description of the algorithm and its analysis.

References

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

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

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

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