Alternating Paths through Disjoint Line Segments*

Michael Hoffmann and Csaba D. Tóth**

Institute for Theoretical Computer Science
ETH Zürich, CH-8092 Zürich, Switzerland
{hoffmann,toth}@inf.ethz.ch

Abstract. We show that every segment endpoint visibility graph on n disjoint line segments in the plane admits an alternating path of length $\Theta(\log n)$, answering a question of Bose. This bound is optimal apart from a constant factor. We also give bounds on the constants hidden by the asymptotic notation.

Keywords: Computational geometry, visibility graphs.

1 Introduction

Consider a set S of n disjoint obstacles, represented by line segments, in the Euclidean plane. A *mobile agent* wishes to visit a maximal number of vertices (i.e., segment endpoints) under various constraints. More specifically, the agent may move along straight line segments between any two vertices, but it must not cross any of the obstacles from S (although it may walk along them from one endpoint to the other).

Similarly to the Euclidean *traveling salesman problem* (ETSP), for which it is known that the optimal solution consists of a *simple* circuit, we restrict the agent to walk along simple paths. But in contrast to ETSP, it is not quite obvious that there always exists a Hamiltonian circuit for the case of segment obstacles; this property was shown only recently [2].

On the other hand, there are sets of line segments for which there is no *circumscribing* Hamiltonian polygon (Figure 1(a)), that is, a polygon whose vertices are the segment endpoints and whose closure contains all the segments [5].

In this paper, we consider *alternating paths*, that is polygonal paths where every other segment is one of the obstacle segments (Figure 1(b)). (Note that *at most* every other segment of such a path can be from S, since the obstacles are disjoint.)

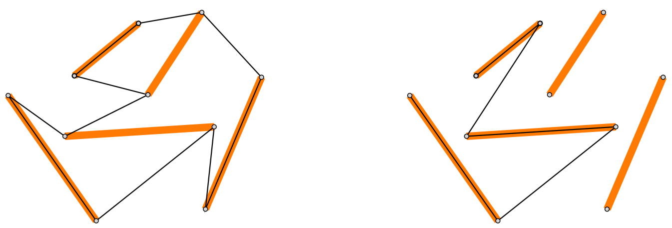

(a) Circumscribing polygon.

(b) Alternating path.

Fig. 1. Example: Obstacle segments.

* A preliminary version of this work was presented at the 18th European Workshop on Computational Geometry (Warsaw, 2002).

** Supported by the joint Berlin-Zürich graduate program "Combinatorics, Geometry, and Computation", financed by the German Science Foundation (DFG) and ETH Zürich.

It is known that there are sets of segments which do not admit an alternating Hamiltonian polygon. Even more, it is NP-complete to decide whether a given set of line segments has this property, if line segments are allowed to intersect at their endpoints [3]. But for some special cases, it can be computed efficiently whether an alternating Hamiltonian circuit exists [4]. So, what is the maximal number of vertices that can be visited by a simple alternating path [1, 6]? We prove the following lower bound.

Theorem 1. *For any set S of n disjoint closed line segments in the plane, there is a simple alternating path visiting $2\lceil \log_2(n+2)\rceil - 2$ vertices.*

Apart from a constant factor, this is best possible:

Theorem 2. *For any $n_0 \in \mathbb{N}$, there exists an $n > n_0$ and a set S of n disjoint closed line segments in the plane, such that S does not admit a simple alternating path visiting more than $\frac{12}{\log_2 3}\log_2 n - 17 < 7.57\log_2 n$ vertices.*

It is easy to turn our proof into an $O(n\log n)$ algorithm to compute an alternating path of length $2\log_2(n+2) - 2$. Note, however, that this path is not necessarily the longest alternating path for the given set of line segments. The optimization problem might have much larger complexity.

2 Preliminaries on Segment Endpoint Visibility Graphs

Consider a set S of n disjoint closed line segments in the plane and denote by V the set of the $2n$ segment endpoints. The *segment endpoint visibility graph* $\text{Vis}(S) = (V, E_S \cup E_{Vis})$ is defined on the vertex set V as follows. Two vertices $u, v \in V$ are connected by a

- *segment edge*, if and only if the corresponding line segment $\overline{uv}$ is in S,

and u and v are connected by a

- *visibility edge*, if and only if the corresponding line segment $\overline{uv}$ does not cross any segment from S.

We say that two line segments cross if they have at least one common point in the relative interior of both segments. Let E_S denote the set of segment edges, and E_{Vis} the set of visibility edges. Figure 2(a) shows an example where the visibility edges are shown as dotted lines.

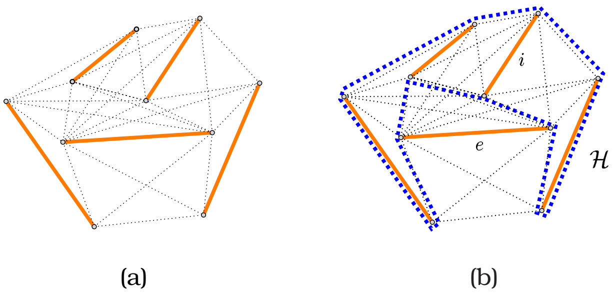

Fig. 2. Segment endpoint visibility graph and a Hamiltonian circuit.

Note that the graph $\text{Vis}(S)$ is defined in geometric terms; hence, there is an associated embedding into the Euclidean plane. In our terminology, a path in $\text{Vis}(S)$ is simple, if it corresponds to a simple polygonal path in the Euclidean plane. Observe that by our definition below an alternating path is always simple.

Definition 1. *A simple path* $p = (v_1\, v_2, \ldots, v_k)$ *in Vis*(S) *is called* alternating path *if it consists of segment edges and visibility edges in alternating order, or formally, if* $v_{2i-1}\, v_{2i} \in E_S$ *for every* $i = 1, \ldots, \lfloor k/2 \rfloor$ *or* $v_{2i}\, v_{2i+1} \in E_S$ *for every* $i = 1, \ldots, \lfloor (k-1)/2 \rfloor$.

It was shown recently that every segment endpoint visibility graph is Hamiltonian [2]. Moreover, if not all segments are collinear, then Vis(S) contains a simple Hamiltonian circuit $\mathcal{H}$. The circuit $\mathcal{H}$ is not necessarily alternating, it may contain several visibility edges in a row (see Figure 2(b) for an example). $\mathcal{H}$ can possibly consist of visibility edges only.

A segment $s \in S$ which is not in $\mathcal{H}$ is necessarily a diagonal of $\mathcal{H}$, which we call a *segment diagonal*. In Figure 2(b), for example, segment i is an internal segment diagonal, while segment e is an external segment diagonal of $\mathcal{H}$. For a graph G, denote by $V(G)$ the vertex set of G, and by $E(G)$ the edge set of G. For two subgraphs G and H of Vis(S), define the graph $G \cup H$ by $V(G \cup H) := V(G) \cup V(H)$ and $E(G \cup H) := E(G) \cup E(H)$. Denote by $\mathcal{D}$ the Hamiltonian circuit $\mathcal{H}$ together with all its segment diagonals.

Proposition 1. *Every vertex of* $\mathcal{D}$ *has degree* 2 *or* 3*. If deg*$(v) = 2$ *for a vertex* v*, then no segment diagonal is incident to* $\mathcal{H}$ *at* v*, therefore* v *is incident to a visibility edge and a segment edge. If deg*$(v) = 3$ *then* v *is incident to a segment diagonal and to two visibility edges along* $\mathcal{H}$. □

Observe that $\mathcal{D}$ is planar. Hence, a simple path in the abstract graph $\mathcal{D}$ always corresponds to a simple path in its planar embedding.

3 The Lower Bound

We show in the next lemma, that one can build an alternating path from any segment edge to any vertex in $\mathcal{D}$. The proof of Theorem 1 then follows by elementary arguments.

Lemma 1. *For any directed segment edge* $\overrightarrow{e} = \overrightarrow{(e_0, e_1)} \in E(\mathcal{D})$ *and any vertex* $f \in V$*,* $\mathcal{D}$ *contains a directed alternating path from* $\overrightarrow{e}$ *to* f.

We define a distance function d on the vertex set as follows. For any $v \in V$, $v \neq e_0$, let $d(v)$ be the length of the shortest (not necessarily alternating) path connecting v and f along $\mathcal{H}$ that does not pass through e_0. (Such a path always exists, since $\mathcal{H}$ is a circuit.) If $e_0 = f$, let $d(e_0) := 0$, else $d(e_0) := \infty$. Next, we orient all visibility edges in $\mathcal{D}$ such that they are directed towards the vertex with smaller value $d(\cdot)$. Two examples are depicted in Figure 3. Note that we do not consider $\mathcal{D}$ as directed graph, the orientation induced by $d(\cdot)$ is merely an aid to construct paths.

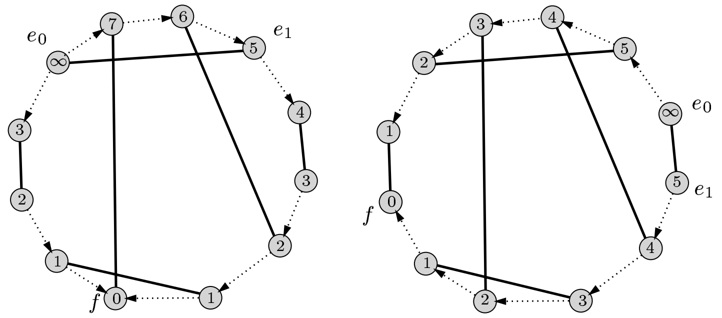

Fig. 3. Orientation and distances.

Proposition 2. *With respect to the orientation according to $d(\cdot)$, every vertex of degree 3 in $\mathcal{D}$ is incident to at least one outgoing and at most one incoming visibility edge, except for f which might be incident to two incoming visibility edges.* $\square$

With help of this orientation, we can try to build an alternating path in $\mathcal{D}$ starting from $\overrightarrow{e}$ and directed towards f as follows.

Algorithm 1 $((v_0, v_1) \in E(\mathcal{D}),\ X \subseteq V)$

$P \leftarrow (v_0, v_1)$.
while *P is simple and has not reached any vertex from X* ***do***
  $(u, v) \leftarrow$ last edge of P.
  if *$deg_{\mathcal{D}}(v) = 2$* ***then*** *append the other ($\neq u$) neighbor of v in $\mathcal{D}$ to P.*
  elsif *$(u, v) \in E_{Vis}$* ***then*** *append the unique segment edge incident to v in $\mathcal{D}$ to P.*
  else *append a visibility edge outgoing from v to P.*
od
return *P.*

To check that Algorithm 1 is well defined, refer to Propositions 1 and 2. When called with $((e_0, e_1), \{f\})$, the algorithm either terminates by reaching f, or when the path P reaches a vertex for the second time. One of these conditions is surely met after finitely many steps, since the graph $\mathcal{D}$ is finite. If the path does not reach f, we are left with a path connected to a cycle, which looks like a balloon with a cord attached to it. Let us derive a more formal – and slightly more general – description for this type of configuration.

Definition 2. *A subgraph G of Vis(S) is called* walkable *from a vertex $v \in V(G)$, if for any vertex $u \in V(G)$, $u \neq v$, there is an alternating path within G from v to u whose edge incident to u is a segment edge.*

Note that, in particular, a graph consisting of a single vertex or two vertices connected by a segment edge always form a walkable subgraph.

Definition 3. *The union $B = G \cup P$ of two subgraphs of Vis(S) with $V(G) \cap V(P) = \{v\}$ and $E(G) \cap E(P) = \emptyset$ is called* balloon, *if G is walkable from v, and $P = (v = v_0, v_1, \ldots, v_k = u)$, $k \in \mathbb{N}$, is an alternating path in Vis(S), such that $(v, v_1) \in E_S$. We call src$(B) := u$ the* source, *hrt$(B) := v$ the* heart, *bdy$(B) := G$ the* body, *and cor$(B) := P$ the* cord *of B. See Figure 4 for an example.*

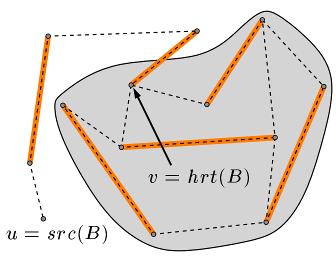

Fig. 4. A balloon B; the body (shaded) is walkable from v.

Proposition 3. *The path P computed by Algorithm 1$((v_0, v_1), X)$ either reaches a vertex from X, or forms a circuit, or forms a balloon.*

Proof. If the algorithm does not reach a vertex from X, it terminates with a path

$$P = (v_0, v_1, \ldots, v_s, v_{s+1}, \ldots, v_k = v_s),$$

for some $s \in \mathbb{N}_0$. If $s = 0$, P forms a circuit. Otherwise, we claim that both (v_{k-1}, v_k) and (v_s, v_{s+1}) are visibility edges.

Indeed, a segment edge is included into the path directly after one of its incident vertices has been reached. Hence, (v_{k-1}, v_k) cannot be a segment edge, otherwise the algorithm would have selected v_{k-1} instead of v_{s+1} as successor of v_s. If $(v_s, v_{s+1}) \in E_S$, the preceding edge (v_{s-1}, v_s) is a visibility edge with $d(v_s) < d(v_{s-1})$ by construction. For the same reason, the edge (v_{k-1}, v_k) is directed towards $v_k = v_s$. Then Proposition 2 tells us that $v_s = f$, and the algorithm would have stopped there.

Proposition 1 implies that every second edge in P is a segment edge; in particular, if $(v_s, v_{s+1}) \in E_{Vis}$, then $(v_{s-1}, v_s) \in E_S$. Thus, every vertex on the path $(v_{s+1}, \ldots, v_{k-1})$ can be reached from v_s on an alternating path that ends with a segment edge: either via v_{s+1} or via v_{k-1} along P. Altogether, we have shown that the constructed path P forms a balloon with source v_0 and heart v_s. □

Proposition 4. *For any vertex v in the body of a balloon B, there is an alternating path from hrt(B) to v in B which starts with a visibility edge and ends with a segment edge.*

Proof. Since the segment edges are pairwise disjoint, there can only be one segment edge incident to any vertex. The segment edge incident to hrt(B) is part of cor(B) by definition; hence, there are only visibility edges incident to hrt(B) in bdy(B). The claim follows from the fact that bdy(B) is walkable. □

Definition 4. *A sequence $\mathcal{B} = (B_1, B_2, \ldots, B_\ell)$, $\ell \in \mathbb{N}_0$, is called* balloon-path *in $\mathcal{D}$, if it satisfies the following conditions.*

1. *For any i, $1 \le i \le \ell$, B_i is a balloon in Vis(S).*
2. *For any i, $1 \le i < \ell$ and any j, $i < j \le \ell$,*

$$V(B_i) \cap V(B_j) = \begin{cases} \{\mathit{src}(B_j)\} \subset V(\mathit{bdy}(B_i)) & \text{, if } j = i+1 \\ \emptyset & \text{, otherwise,} \end{cases}$$
$$E(B_i) \cap E(B_j) = \emptyset\,.$$

Denote $|\mathcal{B}| := \ell$, $V(\mathcal{B}) := \bigcup_{i=1}^{\ell} V(B_i)$, src$(\mathcal{B}) =$ src(B_1), and bdy$(\mathcal{B}) := \bigcup_{i=1}^{\ell}$ bdy(B_i).

We observe a few immediate consequences of this definition. Consider a balloon-path $\mathcal{B} = (B_1, B_2, \ldots, B_\ell)$ in $\mathcal{D}$.

Proposition 5. *For any i, $2 \le i \le \ell$, the edge incident to src(B_i) in B_i is a visibility edge.*

Proof. Since src$(B_i) \in V($bdy$(B_{i-1}))$ by definition, there is an alternating path from src(B_{i-1}) to src(B_i) in B_{i-1} that ends with a segment edge. There is exactly one segment edge incident to every vertex in Vis(S), and $E(B_{i-1}) \cap E(B_i) = \emptyset$. Thus, any edge incident to src(B_i) in B_i must be a visibility edge. □

Note that Proposition 5 implies that $|V(\text{cor}(B_i))| \ge 3$ for any $2 \le i \le \ell$.

Proposition 6.

(i) For every vertex $u \in V(\mathcal{B})$, $\mathcal{B}$ contains an alternating path from src(B_1) to u.
(ii) For every vertex $u \in$ bdy$(\mathcal{B})$, $\mathcal{B}$ contains an alternating path from src(B_1) to u that ends with a segment edge.

Proof. The statement is obvious for $u \in V(\text{cor}(B_1))$. Otherwise, there is by definition an alternating path in B_1 from $\text{src}(B_1)$ to $\text{hrt}(B_1)$ that ends with a segment edge. By Proposition 4, any vertex in $\text{bdy}(B_1)$ can be reached from $\text{hrt}(B_1)$ within B_1 on an alternating path starting with a visibility edge and ending with a segment edge. Since both paths can be concatenated to a single alternating path, we are done for the case that $u \in V(\text{bdy}(B_1))$. Otherwise, we can use the same argument for $\text{src}(B_2)$, that lies in $\text{bdy}(B_1)$ by definition, to construct an alternating path from u to $\text{src}(B_2)$ which ends with a segment edge. The claim follows by induction together with Proposition 5, since for $\ell > 1$, $(B_2, \ldots, B_\ell)$ forms again a balloon-path. □

Proposition 7. *For any vertex $v \in V(\mathcal{B})$, $v \neq \text{src}(B_1)$, the segment edge (u, v) incident to v is in $E(\mathcal{B})$.*

Proof. If $v \in V(\text{bdy}(\mathcal{B}))$, the claim follows from Proposition 6 (ii). So, let $v \in V(\text{cor}(B_i))$: if $v = \text{src}(B_i)$ and $i > 1$, we have $v \in V(\text{bdy}(B_{i-1}))$; if $v = \text{hrt}(B_i)$, it is $v \in V(\text{bdy}(B_i))$. In the remaining case, $\deg_{\text{cor}(B_i)}(v) = 2$. Since $\text{cor}(B_i)$ is an alternating path, one of the edges incident to v in $\text{cor}(B_i)$ must be a segment edge. □

We have now collected all tools to describe an algorithm to construct a balloon-path from $\vec{e}$ headed towards f, that will provide the proof of Lemma 1.

Algorithm 2 $((e_0, e_1) \in E(\mathcal{D}),\ f \in V)$

$\mathcal{B} \leftarrow ()$.
$(m, \mu) \leftarrow (e_0, e_1)$.
while $f \notin V(\mathcal{B})$ ***do***
  $P := (v_1 = m,\ v_2 = \mu,\ \ldots,\ v_k) \leftarrow$ *Algorithm 1*$((m, \mu),\ \{f\} \cup V(\mathcal{B}))$.
  $(B_1, \ldots, B_\ell) \leftarrow \mathcal{B}$.
  if $v_k \in V(B_i)$ *for some* $1 \leq i \leq \ell$ ***then*** $\mathcal{B} \leftarrow \Big(B_1, \ldots, B_{i-1},\ P \cup \bigcup_{j=i}^{\ell} B_j\Big)$.
  else $\mathcal{B} \leftarrow (B_1, \ldots, B_\ell, P)$.
  $m \leftarrow \min_{w.r.t.\, d(\cdot)}$ *bdy*$(B_{|\mathcal{B}|})$.
  $\mu \leftarrow$ *a vertex with* $(m, \mu) \in E(\mathcal{D})$ *and* $d(\mu) < d(m)$. *(cf. Proposition 2)*
od
return $\mathcal{B}$.

Note that $\min_{w.r.t.\, d(\cdot)} \text{bdy}(B_{|\mathcal{B}|})$ is not necessarily unique, because two nodes might have the same $d(\cdot)$ value; but in this case any of them will do.

Proposition 8. *At the beginning of any iteration of the loop in Algorithm 2, $\mathcal{B}$ forms a balloon-path with source e_0.*

Proof. The statement is trivial for the first iteration, since $\mathcal{B} = ()$. In the second iteration, Algorithm 1 is called with parameters $((e_0, e_1), \{f\})$. According to Proposition 3, if the path P does not reach f, it forms either a circuit or a balloon. The function $d(\cdot)$ is defined such that no visibility edge is directed towards e_0, unless $e_0 = f$. The segment edge incident to e_0, (e_0, e_1), is already in P from beginning; hence, it is not possible to revisit e_0 along a segment edge, either. Therefore, the path P returned by Algorithm 1 in the second iteration cannot be a circuit: it must form a balloon.

Assume $\mathcal{B} = (B_1, \ldots, B_\ell)$ is a balloon-path at the beginning of some iteration. Algorithm 1 returns a path $P = (v_1 = m,\ v_2 = \mu,\ \ldots,\ v_k)$ which, according to Proposition 3, either reaches a vertex from $\{f\} \cup V(\mathcal{B})$, or forms a balloon. Notice that if P forms a circuit, it necessarily reaches a vertex from $V(\mathcal{B})$, since $m = v_k \in V(\text{bdy}(\mathcal{B}))$.

Let us first consider the case that P reaches a vertex $v \in V(\mathcal{B})$. We claim $(v_{k-1}, v_k) \in E_{Vis}$: If $v \neq e_0$, recall that by Proposition 7 the segment edge (v, w) incident to v lies in $\mathcal{B}$ as well. Thus, Algorithm 1 stops, if P reaches w. Similarly, the segment edge (e_0, e_1)

incident to e_0 is part of $\mathcal{B}$, and the algorithm stops, if P reaches e_1. Now there are two subcases to consider.

(1) $v_k \in V(\text{bdy}(B_i))$ (Figure 5(a)) There is by definition an alternating path from $\text{src}(B_i)$ to v_k that ends with a segment edge. Hence, any vertex $v_2, \ldots, v_{k-1}$ can be reached from $\text{src}(B_i)$ on an alternating path ending with a segment edge: either via v_k and P or via the balloon-path $(B_i, \ldots, B_\ell)$ to m and then P (cf. Proposition 6 (ii)). (Note that $v_k = m$, if P forms a circuit. But the argument still goes through; one can even argue that this case does never occur during the course of Algorithm 2.) Furthermore, the same argument can be applied to the vertices in $\bigcup_{j=i+1}^{\ell} V(\text{cor}(B_j))$. Thus, $P \cup \bigcup_{j=i}^{\ell} B_j$ is a balloon with source $\text{src}(B_i)$ and heart $\text{hrt}(B_i)$.

(2) $v_k \in V(\text{cor}(B_i))$ (Figure 5(b)) Let $\text{cor}(B_i) = (u_1 = \text{src}(B_i), \ldots, u_s = v_k, \ldots, u_r)$. Since all paths in $\mathcal{B}$ are constructed using Algorithm 1, all visibility edges in the cords of the balloons in $\mathcal{B}$ are oriented from the source to the heart of the balloon. Hence, we can argue as in Proposition 3 that $(u_{s-1}, u_s) \in E_S$. By the same reasoning as above, $P \cup \bigcup_{j=i}^{\ell} B_j$ is a balloon with source $\text{src}(B_i)$ and heart $\text{hrt}(B_i)$.

It remains to consider the else-branch, that is, the case that the path P that is constructed recursively by Algorithm 1 hits itself before reaching any vertex from $V(\mathcal{B})$ (Figure 5(c)). Then, by Proposition 3, P either reaches f or it forms a balloon with $\text{src}(P) = m$. If P reaches f, the algorithm terminates; otherwise, $(B_1, \ldots, B_\ell, P)$ forms again a balloon-path, since $m \in \text{bdy}(B_\ell)$. □

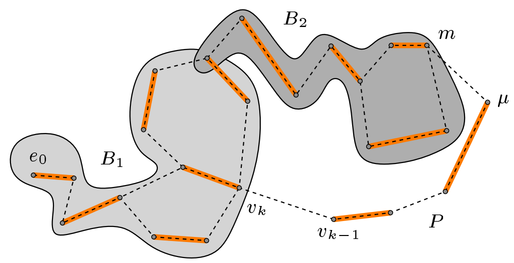

(a) P hits the body of a balloon.

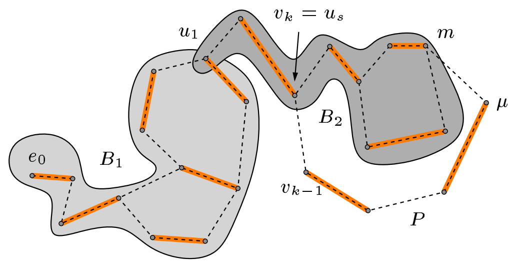

(b) P hits the cord of a balloon.

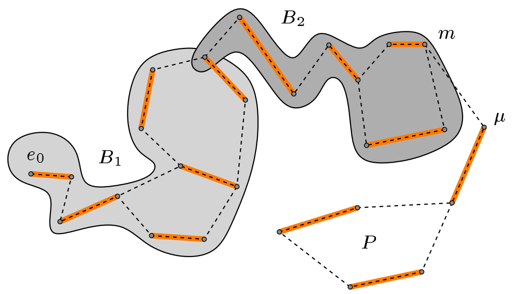

(c) P hits itself.

Fig. 5. Illustrations for Algorithm 2.

Proof (of Lemma 1). We simply apply Algorithm 2 to (e_0, e_1) and f. If we can show that the algorithm always terminates, the claim follows from Propositions 8 and 6 (i). Strictly speaking, after termination $\mathcal{B} = (B_1, \ldots, B_\ell)$ will in general not form a balloon-path anymore; but $(B_1, \ldots, B_{\ell-1})$ is a balloon-path and B_ℓ is an alternating path whose one endpoint lies in $\mathrm{bdy}(B_{\ell-1})$ and is incident to a visibility edge in B_ℓ. Hence, the argument from Proposition 6 goes through.

To show termination, note first that no edge is ever discarded from $\mathcal{B}$, that is, $|E(\mathcal{B})|$ is monotonely increasing over the execution of the algorithm. Moreover, this increase is strict, since in every iteration at least the edge (m, μ) is added to $E(\mathcal{B})$. □

It might be worthwhile to note that we did not use anywhere the fact that $\mathcal{D}$ is planar. The proof of Theorem 1 is completed by the following elementary argument.

Proof (of Theorem 1). Take an arbitrary segment $(e_0, e_1) \in E_S$. We construct recursively a family $(T_i)_{i\in\mathbb{N}_0}$ of trees. Let T_0 be the tree with $V(T_0) := \{e_0, e_1\}$ and $E(T_0) := \{(e_0, e_1)\}$, and let $l(T_0) := \{e_1\}$. Build T_{i+1} from T_i by

$$V(T_{i+1}) := V(T_i) \;\cup\; \left\{ v \in V \setminus T_i \;\middle|\; (l, v) \in E(\mathcal{D}) \cap \begin{Bmatrix} E_S,\ i \text{ odd} \\ E_{Vis},\ i \text{ even} \end{Bmatrix} \text{for some } l \in l(T_i) \right\},$$

$$l(T_{i+1}) := V(T_{i+1}) \setminus V(T_i), \text{ and}$$

$$E(T_{i+1}) := E(T_i) \;\cup\; \left\{ (l, v) \in E(\mathcal{D}) \cap \begin{Bmatrix} E_S,\ i \text{ odd} \\ E_{Vis},\ i \text{ even} \end{Bmatrix} \;\middle|\; l \in l(T_i) \text{ and } v \in l(T_{i+1}) \right\}.$$

Note that T_i forms a tree of alternating paths starting with (e_0, e_1), that consist of $(i + 1)$ vertices. According to Lemma 1, there is some $k \in \mathbb{N}$ such that $V(T_k) = V$. Furthermore, we have $l(T_{i+1}) \leq l(T_i)$ for i odd, and $l(T_{i+1}) \leq 2\, l(T_i)$ for i even (cf. Proposition 1). Thus, for even i we have $|V(T_i)| \leq 2^{\frac{i+4}{2}} - 2$ (for odd i, it is $|V(T_i)| \leq 3 \cdot 2^{\frac{i+1}{2}} - 2$), which yields $k \geq 2\log(n + 2) - 4$. Since T_k contains an alternating path consisting of at least $k + 2$ vertices, the claim follows. □

4 Upper bound

Complementing the results from the previous section, we show here an asymptotically matching lower bound, that is, we construct sets $\mathcal{S}_k$, $k \in \mathbb{N}$, of disjoint line segments that do not have long alternating paths.

Proof (of Theorem 2). We construct the sets of segments $\mathcal{S}_k$, $k \in \mathbb{N}$, recursively as follows. All line segments are chords of a circle c. $\mathcal{S}_1$ consists of three segments arranged in a triangular fashion, i.e., such that $Vis_{\mathcal{S}_1} \cong K_6$. The endpoints of the chords partition c into arcs. $\mathcal{S}_k$ is obtained from $\mathcal{S}_{k-1}$ by inserting a sequence of three segments (i.e., a copy of $\mathcal{S}_1$) on every arc of c that is bounded by only one segment from $\mathcal{S}_{k-1}$. Figure 6 shows $\mathcal{S}_1$ and $\mathcal{S}_2$.

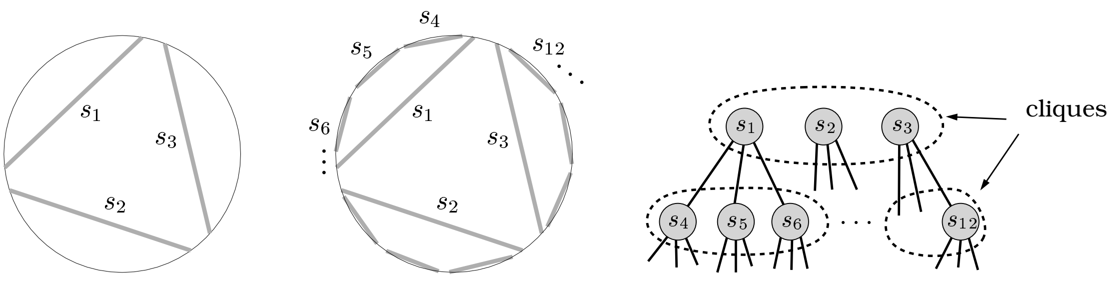

(a) $\mathcal{S}_1$. (b) $\mathcal{S}_2$. (c)

Fig. 6. The construction of $\mathcal{S}_k$.

The two endpoints of any segment in this construction are adjacent to the same set of segment endpoints in the visibility graph Vis(S_k). Hence, we can interpret $\mathcal{S}_k$ as complete ternary tree of depth $k-1$ where each vertex is formed by a clique of three segments (Figure 6(c)). Let λ_k be the length of a longest alternating path in $\mathcal{S}_k$, where the length of a path is defined as the number of vertices along the path. Since the longest simple path in a tree of depth $k-1$ has length $2\,k-1$ and since visiting a 3-clique of segments means visiting 6 vertices, we conclude that $\lambda_k = 12\,k-6$.

$\mathcal{S}_k$ contains exactly $n_k := 3^{k+1}-3$ vertices. Hence,

$$\lambda_k \;=\; 12\left(\log_3 n_k - 1 + \log_3 \frac{3^k}{3^k-1}\right) - 6 \;=\; \frac{12}{\log_2 3}\,\log_2 n_k - 18 + 12\,\log_3 \frac{3^k}{3^k-1}$$

and the claimed result follows, since the last term is less than one for $k \geq 3$. □

Note that an $\Omega(\log n)$ bound was already known by a construction due to Urrutia [6], but with a weaker constant coefficient.

5 Open questions

For our upper bound construction from Section 4 it is straightforward to find a Hamiltonian polygon. Visiting vertices along the circle c gives a *circumscribing* polygon (i.e., where all segment edges are sides or internal diagonals). We note here that it is considerably easier to establish Theorem 1 for sets of segments which admit a circumscribing polygon.

We note also that the maximal alternating paths are contained in the graph $\mathcal{D}$ for this example. Our algorithm from Section 3 always gives a path of maximum length, if the right starting edge is chosen.

We conclude with two open questions. Are there matching lower and upper bounds for the length of a longest alternating path that any set of n disjoint line segments has? We showed that it must be between $2\log_2(n+2)-2$ and $7.57\,\log_2 n - 17$.

Our approach, using only the abstract graph $\mathcal{D}$, can possibly lead to the solution. Therefore we formulate the following question: Let H and M be a Hamiltonian circuit and a complete matching on the same set V of $2n$ vertices. What is the longest simple path in the abstract graph $(V, H \cup M)$ in which every second edge belongs to M? For this problem, the best lower and upper bounds we know are the same as for alternating paths in the segments endpoint visibility graph.

References

1. Demaine, E. D., and O'Rourke, J. Open problems from CCCG'99. In *Electronic Proc. 11th Canadian Conf. on Comput. Geom.* (Vancouver, 1999).
2. Hoffmann, M., and Tóth, Cs. D. Segment endpoint visibility graphs are Hamiltonian. *Comput. Geom. Theory Appl.* (2002). to appear. `http://www.inf.ethz.ch/~toth/hamilton.ps.gz`.
3. Rappaport, D. Computing simple circuits from a set of line segments is NP-complete. *SIAM J. Comput. 18*, 6 (1989), 1128–1139.
4. Rappaport, D., Imai, H., and Toussaint, G. T. Computing simple circuits from a set of line segments. *Discrete Comput. Geom. 5*, 3 (1990), 289–304.
5. Urabe, M., and Watanabe, M. On a counterexample to a conjecture of Mirzaian. *Comput. Geom. Theory Appl. 2*, 1 (1992), 51–53.
6. Urrutia, J. Algunos problemas abiertos (in Spanish). In *Actas de los IX Encuentros de Geometría Computacional* (Girona, 2001).