

Circumscribing Polygons and Polygonizations for Disjoint Line Segments

Cite this Accepted Manuscript (AM) as Accepted Manuscript (AM) version of Hugo A. Akitaya, Matias Korman, Oliver Korten, Mikhail Rudoy, Diane L. Souvaine, Csaba D. Tóth, Circumscribing Polygons and Polygonizations for Disjoint Line Segments, Discrete & Computational Geometry <https://doi.org/10.1007/s00454-021-00355-8>

This AM is a PDF file of the manuscript accepted for publication after peer review, when applicable, but does not reflect post-acceptance improvements, or any corrections. Use of this AM is subject to the publisher's embargo period and AM terms of use. Under no circumstances may this AM be shared or distributed under a Creative Commons or other form of open access license, nor may it be reformatted or enhanced, whether by the Author or third parties. See here for Springer Nature's terms of use for AM versions of subscription articles: <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>

The Version of Record of this article, as published and maintained by the publisher, is available online at: <https://doi.org/10.1007/s00454-021-00355-8>. The Version of Record is the version of the article after copy-editing and typesetting, and connected to open research data, open protocols, and open code where available. Any supplementary information can be found on the journal website, connected to the Version of Record.

Circumscribing Polygons and Polygonizations for Disjoint Line Segments*

Hugo A. Akitaya[†] Matias Korman[‡] Oliver Korten[§] Mikhail Rudoy[¶]
 Diane L. Souvaine[§] Csaba D. Tóth^{‡||}

Abstract

Given a planar straight-line graph $G = (V, E)$ in \mathbb{R}^2 , a *circumscribing polygon* of G is a simple polygon P whose vertex set is V , and every edge in E is either an edge or an internal diagonal of P . A circumscribing polygon is a *polygonization* for G if every edge in E is an edge of P . We prove that every arrangement of n disjoint line segments in the plane has a subset of size $\Omega(\sqrt{n})$ that admits a circumscribing polygon, which is the first improvement on this bound in 20 years. We explore relations between circumscribing polygons and other problems in combinatorial geometry, and generalizations to \mathbb{R}^3 . We show that it is NP-complete to decide whether a given graph G admits a circumscribing polygon, even if G is 2-regular. Settling a 30-year old conjecture by Rappaport, we also show that it is NP-complete to determine whether a geometric matching admits a polygonization.

1 Introduction

Reconstruction of geometric objects from partial information is a classical problem in computational geometry. In this paper we revisit the problem of reconstructing a simple polygon (alternatively, a triangulated simple polygon) P when some of its edges have been lost. Given a set V of n points in the plane, a polygonization of V is a simple polygon P whose vertex set is V . It is easy to see that, unless all points are collinear, V has a simple polygonization. The number of polygonizations is exponential in n , and there has been extensive work on determining the minimum and maximum number of polygonizations for n points in general position as a function of n (see [8, 20] for the latest upper and lower bounds, and [12] for a survey on this and related problems).

A natural generalization of this problem is to augment a given planar straight-line graph¹ (PSLG) $G = (V, E)$ into a simple polygon or a Hamiltonian PSLG. In particular, three variants have been considered: A simple polygon P on vertex set V is a *polygonization* if every edge in E is an edge of P ; a *circumscribing polygon* if every edge in E is an edge or an internal diagonal in P ;

*A preliminary version of this paper appeared as [2]. Research supported in part by the NSF awards CCF-1422311 and CCF-1423615. Akitaya was partially supported by NSERC.

[†]Department of Computer Science, University of Massachusetts Lowell, Lowell, MA, USA.

[‡]Siemens Electronic Design Automation, OR, USA

[§]Department of Computer Science, Tufts University, Medford, MA, USA

[¶]MIT CSAIL, Cambridge, MA, USA. Now at Google Inc., Cambridge, MA, USA.

^{||}Department of Mathematics, California State University Northridge, Los Angeles, CA, USA.

¹A more accurate term would be *plane* straight-line graph, but we decided to use planar straight-line graph as it is a much more commonly used term in the literature.

and a *compatible Hamiltonian polygon* if every edge in E is an edge, an internal diagonal, or an external diagonal in P .

Hoffmann and Tóth [11] proved that every plane straight-line matching admits a compatible Hamiltonian polygon, unless all segments are collinear, in which case no such polygon exists. Grünbaum [10] constructed an arrangement of six disjoint segments that does not admit a circumscribing polygon; see Fig. 14 (an earlier construction of size 16 is in [22]). However, a circumscribing polygon is known to exist when (i) each segment has at least one endpoint on the boundary of the convex hull [14], or (ii) no segment intersects the supporting line of any other segment [15]. Pach and Rivera-Campo [17] proved in 1998 that every set of n disjoint segments contains a subset of $\Omega(n^{1/3})$ segments that admits a circumscribing polygon; no nontrivial upper bound is known.

Rappaport [18] proved that it is NP-complete to decide whether G can be augmented into a simple polygon. In the reduction, G consists of disjoint paths, and Rappaport conjectured that the problem remains hard even if G is a perfect matching (i.e., disjoint line segments in the plane). In the special case that G is perfect matching and every segment has at least one endpoint on the boundary of the convex hull, then an $O(n \log n)$ time algorithm can compute a polygonization (or report that none exists [19]). If S is a set of $n \geq 3$ parallel chords of a circle, then neither S nor any subset of three or more segments from S admits a polygonization (so the analogue of the problem of Pach and Rivera-Campo [17] has a trivial answer in this case). In a related result, Ishaque et al. [13] proved that n disjoint line segments in general position, where n is even, can be augmented to a 2-regular PSLG (i.e., a union of disjoint simple polygons).

Our Results:

- We prove that every set of n disjoint line segments in general position contains a subset of $\Omega(\sqrt{n})$ segments that admit a circumscribing polygon (Theorem 1 in Sect. 2). This is the first improvement over the previous bound of $\Omega(n^{1/3})$ [17] in the last 20 years.
- While we do not have any nontrivial upper bound for circumscribing polygons proper, we relate that problem to the extensibility of disjoint line segments to disjoint rays. For every $n \in \mathbb{N}$, we construct a set of n disjoint line segments in the plane such that the size of any subset extensible to disjoint rays is $O(\sqrt{n})$ (Sect. 3).
- We prove that it is NP-complete to determine whether a given set of disjoint cycles in the plane admits a circumscribing polygon (Theorem 13 in Sect. 4). The reduction is from Hamiltonian paths in 3-connected cubic planar graphs.
- We prove that it is NP-complete to determine whether a given set of disjoint line segments admits a polygonization (Theorem 20 in Sect. 5). This settles a 30-year old conjecture by Rappaport [18] in the affirmative.

We conclude with a few open problems and three-dimensional generalizations in Sect. 6.

Further Related Previous Work

Hamiltonicity has fascinated graph theorists and geometers for centuries. Some planar graph results hold for PSLGs as well (i.e., planar graphs with a fixed straight-line embeddings). Hamiltonicity is NP-complete for planar cubic graphs [9], but can be solved in linear time in 4-connected planar

graphs [6], and all 4-connected triangulations (i.e., edge-maximal planar graphs) are Hamiltonian [23]. In terms of augmentation, a non-Hamiltonian triangulation cannot be augmented to a Hamiltonian planar graph by adding edges or vertices. However, Cardinal et al. [4, Theorem 5] proved that every planar graph on n vertices can be transformed into a Hamiltonian planar graph by subdividing at most $\lfloor (n-3)/2 \rfloor$ edges, with one vertex each, and by adding new edges. See also the surveys [7, 16] on Hamiltonicity of planar graphs and their applications.

2 Large Subsets with Circumscribing Polygons

For every integer $n \geq 2$, let $f(n)$ be the maximum integer such that every set of n disjoint segments in the plane in general position contains a subset of $f(n)$ segments that admit a circumscribing polygon. Pach and Rivera-Campo [17] proved that $f(n) = \Omega(n^{1/3})$. By building up on this result, we improve the bound to $\Omega(\sqrt{n})$.

Theorem 1. *Every set of $n \geq 2$ disjoint line segments in the plane in general position contains a subset of $\Omega(\sqrt{n})$ segments that admit a circumscribing polygon.*

The remainder of this section is dedicated to proving this statement. We start with a brief overview of our approach for a set S of n line segments.

- In Sect. 2.1 we show how to select a subset of $\Omega(\sqrt{n})$ segments to be included in a circumscribing polygon (cf. Fig. 1). The main property of the selected segments is that they are stabbed by few vertical lines and their slopes are monotonically increasing or decreasing when ordered along their intercepts along those lines. Ultimately, our circumscribing polygon will contain at least a quarter of the selected segments.
- Our initial candidate polygon will be the convex hull of the selected segments. In most cases, though, the convex hull is not a circumscribing polygon. Thus, we introduce four elementary operations, called **ChopWedges**, **BuildCap**, **Dip**, and **ShearDip**. These operations each make local changes to the candidate solution and increase the number of segment endpoints visited by the solution. In Sect. 2.2, we describe all four operations and describe the invariants of the polygons that we maintain in the course of the algorithm.
- With the invariants in place, we can describe an algorithm. It successively invokes the four elementary operations and returns a circumscribing polygon for a constant fraction of the $\Omega(\sqrt{n})$ selected segments. The algorithm proceeds in four phases and is described in Sect. 2.3.

2.1 Selecting a Subset of Segments

Let S be a set of $n \geq 2$ disjoint line segments in the plane. We may assume without loss of generality that none of the segments is vertical, and all segment endpoints have distinct x -coordinates. For a subset $S' \subseteq S$, a *halving line* is a vertical line ℓ such that the number of segments in S' strictly contained in the left and right open halfplanes bounded by ℓ differ by at most one. In particular, each halfplane contains at most $|S'|/2$ segments from S' .

We partition S recursively as follows. Find a halving line ℓ for S , and recurse on the nonempty subsets of segments lying in each open halfplane determined by ℓ . Denote by T the recursion tree, which is a binary tree of depth at most $\log n$. We denote by $V(T)$ the set of nodes of T , and

by $V_i(T)$ the set of nodes at level i of T for $i = 0, 1, \dots, \lfloor \log n \rfloor$. Associate each node $v \in V(T)$ to a halving line ℓ_v and to the subset $S_v \subseteq S$ of segments that intersect ℓ_v without intersecting the halving lines associated with any ancestor of v . This defines a partition of S into subsets S_v , $v \in V(T)$.

For every $v \in V(T)$, sort the segments in S_v by the y -coordinates of their intersections with the line ℓ_v ; and let $Q_v \subseteq S_v$ be a maximum subset of segments that have monotonically increasing or decreasing slopes. By the Erdős–Szekeres theorem, we have $|Q_v| \geq \sqrt{|S_v|}$ for every $v \in V(T)$. For a refined analysis, we consider the union of the sets Q_v for $v \in V_i(T)$ for $i = 0, \dots, \lfloor \log n \rfloor$, and then take one such union of maximal cardinality.

We need some additional notation. For every $v \in V(T)$, let $n_v = |S_v|$ and $m_v = |Q_v|$. For every integer $i = 0, 1, \dots, \lfloor \log n \rfloor$, let \mathcal{S}_i (resp., \mathcal{Q}_i) be the union of S_v (resp., Q_v) over all vertices $v \in V_i(T)$. Let $\nu_i = |\mathcal{S}_i|$ and $\mu_i = |\mathcal{Q}_i|$. By definition, we have $n = \sum_{i=0}^{\lfloor \log n \rfloor} \nu_i$. Let $M = \max\{\mu_i : 0 \leq i \leq \lfloor \log n \rfloor\}$. We claim that

$$M \geq \frac{\sqrt{n}}{2}. \quad (1)$$

By the Erdős–Szekeres Theorem, we have $m_v \geq \sqrt{n_v}$ for every $v \in V(T)$. Since $n_v \leq n/2^i$ for every $v \in V_i(T)$, then $m_v \geq \sqrt{n_v} = n_v/\sqrt{n_v} \geq n_v/\sqrt{n/2^i} = \sqrt{2^i/n} \cdot n_v$. Summation over all $v \in V_i(T)$ yields $M \geq \mu_i \geq \sqrt{2^i/n} \cdot \nu_i$, which in turn gives $\nu_i \leq M\sqrt{n/2^i}$. Summation over all $i = 0, \dots, \lfloor \log n \rfloor$ now gives $n = \sum_{i=0}^{\lfloor \log n \rfloor} \nu_i \leq M\sqrt{n} \sum_{i=0}^{\lfloor \log n \rfloor} 2^{-i/2} \leq 2M\sqrt{n}$, hence $M \geq \sqrt{n}/2$, which proves (1).

Let $i^* \in \{0, 1, \dots, \lfloor \log n \rfloor\}$ be an index where $M = \mu_{i^*}$, and put $\hat{S}_0 = \mathcal{Q}_{i^*}$. By construction, $\hat{S}_0 = \bigcup \{Q_v : v \in V_{i^*}(T)\}$. We further partition \hat{S}_0 into two subsets as follows. Let $V_{i^*}^<$ (resp., $V_{i^*}^>$) be the set of nodes in $V_i(T)$ such that the slopes in Q_v monotonically increase (resp., decrease). Let \hat{S}_1 be the larger of $\bigcup \{Q_v : v \in V_{i^*}^<\}$ and $\bigcup \{Q_v : v \in V_{i^*}^>\}$, breaking ties arbitrarily. Note that $|\hat{S}_1| \geq \sqrt{n}/4$. We may assume, by a reflection in the y -axis if necessary, that $\hat{S}_1 = \bigcup \{Q_v : v \in V_{i^*}^>\}$; see Fig. 1 for an example.

2.2 Algorithm Invariants and Elementary Operations

Pach and Rivera-Campo [17] proved that an arrangement of disjoint line segments admits a circumscribing polygon if they are (i) stabbed by a vertical line, and (ii) have monotonically increasing or decreasing slopes (in particular, each Q_v , $v \in V(T)$, admits a circumscribing polygon). In contrast, we construct a circumscribing polygon for the *union* of all Q_v , $v \in V_{i^*}(T)$, separated by vertical lines. Note that our construction will be a circumscribing polygon for a large subset $\hat{S}_2 \subseteq \hat{S}_1$ (not the whole set \hat{S}_1).

For ease of presentation, we introduce new notation for $\hat{S}_1 = \bigcup \{Q_v : v \in V_{i^*}^>\}$; see Fig. 1. Denote by $\ell_1, \ell_2, \dots, \ell_r$ the halving lines $\{\ell_v : v \in V_{i^*}^>\}$ sorted from left-to-right, and let Q_1, Q_2, \dots, Q_r be the corresponding sets in $\{Q_v : v \in V_{i^*}^>\}$. Denote by L_i , $i = 1, \dots, r-1$, the vertical lines that separate Q_i and Q_{i+1} . Refer to Fig. 1.

Overview. We construct a circumscribing polygon for a subset $\hat{S}_2 \subseteq \hat{S}_1$ incrementally, while maintaining a polygon in the segment endpoint visibility graph of \hat{S}_1 . We use a machinery developed in [11], with several important new elements. Initially, let P be the boundary of $\text{conv}(\hat{S}_1)$, where conv denotes the convex hull. Intuitively, think of polygon P as a rubber band, and stretch it




Figure 1: A set $\hat{S}_1 = \bigcup \{Q_v : v \in V_{i^*}^>\}$ of 25 line segments for $r = 4$; and $P = \text{conv}(\hat{S}_1)$.

successively to visit more segment endpoints from \hat{S}_1 , maintaining the property that all segments in \hat{S}_1 remain in the closed polygonal domain of P . A key invariant of P will be that if P visits only one endpoint of some segment in \hat{S}_1 , then we can stretch it to visit the other endpoint (a strategy previously used in [11, 14]). This tool allows us to produce a circumscribing polygon for a subset of \hat{S}_1 . We ensure that P reaches an endpoint of *at least a quarter* of the segments in \hat{S}_1 . To do this, we use the fact that each set Q_v , $v \in V_{i^*}^>$, is sorted along the halving lines in decreasing order by slope, and we ensure that P reaches the left endpoint of at least a quarter of the segments (later, we stretch P to visit the right endpoints). At the end, we define \hat{S}_2 to be the set of segments in \hat{S}_1 visited by P (i.e., we discard the remaining segments lying in the interior of P).

We maintain a polygon with the properties listed in Definition 2 below. There are a few important features to note: P is not necessarily a simple polygon in intermediate steps of the algorithm: it may be a weakly simple polygon that does not have self-crossings; it has clearly defined interior and exterior; and it can have repeated vertices. Specifically, each vertex can repeat at most twice (i.e., multiplicity at most 2), and if its multiplicity is 2, then one occurrence is a reflex vertex and the other is convex. Furthermore, all such reflex vertices can be removed simultaneously by suitable shortcuts (cf. property (F5) below) to obtain a simple polygon. We need to be very careful about reflex vertices in P : for each reflex vertex in P , we ensure either that it will not become a repeated vertex later, or that if it becomes a repeated vertex, its reflex occurrence can be removed by a suitable shortcut.

Invariants. As in [11], we maintain a weakly simple polygon, called a *frame* (defined below). A weakly simple polygon is a closed polygonal chain $P = (v_1, \dots, v_k)$ in counterclockwise order such that, for every $\varepsilon > 0$, displacing the vertices by at most ε can produce a simple polygon. Denote by \hat{P} the union of the interior and the boundary of P . A weakly simple polygon may have repeated vertices. Three consecutive vertices (v_{i-1}, v_i, v_{i+1}) define an interior angle $\angle(v_{i-1}, v_i, v_{i+1})$, or $\angle v_i$, which is either convex (less than or equal to 180°) or reflex (more than 180°).

The following definitions summarizes the properties that we maintain for a polygon P . It is based on a similar concept in [11]: we do not allow segments to be external diagonals (cf. (F2)) and relax the conditions on the possible occurrences of reflex vertices. Reflex vertices play an important role. We distinguish two types of reflex vertices: A reflex vertex v of a frame P is **safe** if the (unique) line segment in S incident to v subdivides the reflex angle $\angle v$ into two convex angles;

otherwise v is **unsafe** (see Fig. 2 for examples).

Definition 2. A weakly simple polygon $P = (v_1, \dots, v_k)$ is called frame for a set S of disjoint line segments in the plane, if

- (F1) every vertex of P is an endpoint of some segment in S ;
- (F2) \hat{P} , the union of the interior and the boundary of P , contains every segment in S ;
- (F3) every vertex of P has multiplicity at most 2;
- (F4) if a vertex of P has multiplicity 2, say $v_i = v_j$, then one of $\angle v_i$ or $\angle v_j$ is convex (and the other is reflex);
- (F5) if (v_i, \dots, v_j) is a maximal chain of reflex vertices of P that each have multiplicity 2, then $(v_{j+1}, v_j, \dots, v_i, v_{i-1})$ is a simple polygon that is disjoint from the interior of P (cf. Fig. 2).




Figure 2: A frame $P = (v_1, \dots, v_{18})$ for ten disjoint line segments (orange). The closed region \hat{P} is shaded gray. The vertices of multiplicity 1 (resp., 2) are marked with full (resp., empty) dots. The reflex vertex v_{13} is safe, all other reflex vertices (v_4, v_7, v_8, v_9 , and v_{17}) are unsafe.

Elementary Operations. Let S be a set of disjoint line segments in general position, and let P be a frame. For S we define four elementary operations that each transform P into a new frame. The first operation is the “shortcut” that eliminates reflex vertices of multiplicity 2, and increases the area of the interior. The remaining three operations each increase the number of vertices of the frame (possibly creating vertices of multiplicity 2) and decrease the area of its interior. For shortest path and ray shooting computations, we consider the line segments in S and the current frame P to be obstacles. For a polygonal path (a, b, c) that does not cross any segment in S , we define the *convex arc* $\text{carc}(a, b, c)$ to be the shortest polygonal path between a and c that is homotopic to (a, b, c) .

Operation 1 ($\text{ChopWedges}(P)$). Refer to Fig. 3(c-d). Input: a frame P . Action: While there is a vertex of multiplicity 2, do: let (v_i, \dots, v_j) be a maximal chain of reflex vertices of P that each have multiplicity 2, and replace the path $(v_{i-1}, v_i, \dots, v_j, v_{j+1})$ in P by a single edge $v_{i-1}v_{j+1}$.

Operation 2 ($\text{BuildCap}(P, \varrho, b)$). Refer to Fig. 3(a-c). Input: a frame P , an orientation $\varrho \in \{-1, +1\}$, and vertex b is of multiplicity 1 in P such that b the endpoint of $ab \in S$, where a is not a vertex of P . Let c be the neighbor of b in polygon P in orientation ϱ (where $\text{ccw} = 1, \text{cw} = -1$), and assume $\angle abc$ is convex. Action: Replace the edge bc of P with the polygonal path $ba + \text{carc}(a, b, c)$.




Figure 3: (a) A frame P . (b) $P := \text{BuildCap}(P, 1, v_1)$. (c) $P := \text{BuildCap}(P, 1, b)$. (d) $P := \text{ChopWedges}(P)$.




Figure 4: (a) A frame. (b) The result of operation $\text{Dip}(P, a, b)$. (c) A frame. (d) The result of operation $\text{ShearDip}(P, a, x)$.

Operation 3 ($\text{Dip}(P, a, b)$). Refer to Fig. 4 (a-b). Input: a frame P , a segment $ab \in S$ such that neither a nor b is a vertex of P , and the ray \overrightarrow{ba} hits an edge of P that is not a segment in S . Action: Assume that \overrightarrow{ba} hits edge uv of P at the point x . Replace the edge uv of P with the polygonal path $\text{carc}(u, x, a) + \text{carc}(a, x, v)$.

Operation 4 ($\text{ShearDip}(P, a, x)$). Refer to Fig. 4 (c-d). Input: a frame P , a segment endpoint a in the interior of P , and a point x in the interior of an edge of P that is not a segment in S such that ax does not cross P or any segment in S . Action: Let uv be the edge of P that contains x . Replace the edge uv of P with the polygonal path $\text{carc}(u, x, a) + \text{carc}(a, x, v)$.

Operations 1 and 2 have been previously used in [11, 14]; it was shown that if *all* reflex vertices in P have been created by **BuildCap** operations, then (F5) automatically holds [11, Sect. 2]. It is also not difficult to see that if *all* reflex vertices in P have been created by **Dip** operations, then (F5) holds. However, this property does not extend to a mixed sequence of **BuildCap** and **Dip** operations, and certainly not for **ShearDip** operations. We maintain (F5) by a careful application of these operations, using the fact that each set Q_i , $i = 1, \dots, r$, is stabbed by a vertical line.

Note that Operations 1–4 can only increase the vertex set of the frame (the **ChopWedges** operation decreases the multiplicity of repeated vertices from 2 to 1, but maintains the same vertex set). Initially, $P = \partial \text{conv}(S)$, and so all vertices of $\text{conv}(S)$ remain vertices in P in the course of the algorithm. In particular, the leftmost and rightmost segment endpoint in S are always vertices in P (with multiplicity 1 by (F4)). These vertices subdivide P into an *upper arc* and a *lower arc*. As a convention, the leftmost (resp., rightmost) vertex is part of the lower arc (upper arc). When our algorithm invokes the **BuildCap** operation at a vertex v , we use $\text{BuildCap}(P, \varrho(v), v)$.

We can now justify the distinction between safe and unsafe reflex vertices.

Lemma 3. *Let v be a reflex vertex of multiplicity 1 in a frame P such that v is safe. Then after any sequence of the above four operations, v becomes a convex vertex or remains a safe reflex vertex of the frame, in both cases the multiplicity of v remains 1.*

Proof. Each operation creates at most one new reflex vertex, which has multiplicity 1; and possibly many convex vertices along the convex arcs, which may have multiplicity 1 or 2. However, each point can be an interior vertex of at most one convex arc. Consequently, the multiplicity of a vertex v can possibly increase from 1 to 2 if it is first a reflex vertex of multiplicity 1, and then visited for a second time (by a convex arc) as a convex vertex; see Fig. 3 (c) and Fig. 4 (d) for examples. If v is a *safe* reflex vertex, then it cannot be an interior vertex of a convex arc, and so its multiplicity cannot increase from 1 to 2. Since each operation decreases the interior of the frame, the interior angle of the frame at v can only decrease. Consequently, v either becomes a convex vertex or remains a safe reflex vertex. \square

Additional Invariants for Vertical Lines. In this section, we assume that our instance S is the set $\hat{S}_1 = \bigcup_{i=1}^r Q_i$ constructed in Sect. 2.1. The first two phases of our algorithm will maintain the following additional properties (F6)–(F7) for frames of \hat{S}_1 . We make this distinction because we use Definition 2 and the elementary operations introduced here for a different class of input in Sect. 3.

- (F6) The vertical line L_i , $i \in \{1, \dots, r-1\}$, crosses P exactly twice: once in the upper arc and once in the lower arc.
- (F7) The line ℓ_i , $i \in \{1, \dots, r\}$, crosses the lower arc (resp., upper arc) of P an odd number of times, and the y -coordinates of the intersection points with the lower arc (resp., upper arc) are monotonically decreasing in a counterclockwise traversal of P .




Figure 5: A set \hat{S}_1 of 25 line segments and the frame P at the end of Phase 1.

2.3 Algorithm Description

With the operations and invariants described above we can proceed to describe our algorithm. Our algorithm has four phases described below. We define a default orientation for every (current and

future) vertex v of the frame P : if v is in the lower arc, then $\varrho(v) = 1$, otherwise $\varrho(v) = -1$. (We shall make some exceptions to this default assignment after Phase 1.)

Phase 1: Left Endpoints. Initially, the frame P is the boundary of the convex hull $\text{conv}(\widehat{S}_1)$. In the first phase of our algorithm, we use operations **BuildCap** and **Dip** as follows:

1. Initialize $P := \partial\text{conv}(\widehat{S}_1)$.
2. While condition (a) or (b) below is applicable, do:
 - (a) If there exists a segment $ab \in S$ such that the left endpoint a is a vertex of P , but the right endpoint is not, then set $P := \text{BuildCap}(P, \varrho(a), a)$.
 - (b) Else, if there exists a segment $ab \in Q_i$ for some $i \in \{1, \dots, r\}$ such that a is the left endpoint, a lies in the interior of P , and \vec{ba} hits an edge uv where $uv \notin \widehat{S}_1$, and the left endpoint of uv is an endpoint of some segment in Q_i , then set $P := \text{Dip}(P, a, b)$.
3. Return P , and terminate Phase 1.

An example is shown in Fig. 5. First we show that Phase 1 returns a frame. Instead of addressing property (F5) directly, we establish the stronger property (R1) defined below.

Lemma 4. *All operations in Phase 1 maintain (F1)–(F7) and the following property.*

(R1) *If v_i is an unsafe reflex vertex of P , then v_i is the right endpoint of a segment in \widehat{S}_1 , vertices v_{i-1} and v_{i+1} are convex and the triangle (v_{i+1}, v_i, v_{i-1}) is disjoint from the interior of P .*

Proof. Phase 1 applies a sequence of **BuildCap** and **Dip** operations. Both operations automatically maintain (F1) and (F2). Both operations create at most one new reflex vertex in P (at vertex a and b , respectively). In both cases, the reflex vertex has not been previously a vertex, so every reflex vertex has multiplicity 1 at the time it is created. A reflex vertex may be visited in a subsequent operation as a convex vertex, but at most once. This establishes (F3) and (F4).

Each operation **BuildCap** $(P, \varrho(a), a)$ maintains (F6): it stretches P from one endpoint to the other endpoint of a segment $ab \in \widehat{S}_1$; both endpoints lie between two consecutive lines L_{i-1} and L_i . We show that (F6) is also maintained by each operation **Dip** (P, a, b) . To see this, recall that a is the left endpoint of a segment $ab \in Q_i$, and so the point x where ray \vec{ba} hits the frame P is to the left of a . The algorithm applies operation **Dip** (P, a, b) if x lies on an edge uv of P , and the left endpoint of uv , say u , is an endpoint of some segment in Q_i . Therefore, a , b , and x are all between the lines L_{i-1} and L_i . Thus, $\text{carc}(u, x, a)$ does not cross these lines. If v is between L_{i-1} and L_i , then so is $\text{carc}(a, x, v)$; otherwise $\text{carc}(a, x, v)$ crosses the same vertical lines L_j , $j \geq i$, that edge uv crossed before the operation, each at most once.

Each operation **BuildCap** $(P, \varrho(a), a)$ maintains (F7): Assume the operation introduces a new pair of crossings with the vertical line ℓ_i at $ab \cap \ell_i$ and $\text{carc}(a, b, c) \cap \ell_i$, where c is the neighbor of a in orientation $\varrho(a)$. If a is on the lower arc, then $\varrho(a) = 1$ by convention, and $\text{carc}(a, b, c) \cap \ell_i$ is below $ab \cap \ell_i$. Otherwise a is on the upper arc, then $\varrho(a) = -1$ and $\text{carc}(a, b, c) \cap \ell_i$ is above $ab \cap \ell_i$. In both cases, (F7) is maintained.

We show that (R1) implies (F5). Note that initially, P has no reflex vertices, so all reflex vertices are created by an operation in Phase 1. By Lemma 3, the multiplicity of safe vertices of P remains 1. Let (v_i, \dots, v_j) be a maximal chain of reflex vertices that each have multiplicity 2 in P .

By (R1), $i = j$, and the chain consists of a single unsafe reflex vertex v_i . (R1) further implies that (v_{i+1}, v_i, v_{i-1}) is a simple polygon disjoint from the interior of P .

It remains to show that all operations in Phase 1 maintain (R1). Each operation creates at most one new reflex vertex. First, consider a reflex vertex created by an operation $\text{Dip}(P, a, b)$. It creates a reflex vertex at a , which is safe by construction, and it never becomes an unsafe reflex vertex by Lemma 3. It follows that every unsafe reflex vertex of P in the course of Phase 1 has been created by a **BuildCap** operation.




Figure 6: The left endpoint of segment $ab \in \widehat{S}_1$ is in the lower arc of P . (b) The result of operation $\text{BuildCap}(P, 1, a)$. (c) The result of a subsequent operation $\text{Dip}(P, a', b')$.

Second, consider a reflex vertex created by an operation $\text{BuildCap}(P, \varrho(a), a)$. Without loss of generality, assume that a is in the lower arc of P , hence $\varrho(a) = 1$. Refer to Fig. 6. The operation creates an unsafe reflex vertex at b , which is the right endpoint of $ab \in \widehat{S}_1$. Vertex b is incident to two edges of the frame: ab and bc' , where c' is first vertex of $\text{carc}(b, a, c)$ after b . Vertex a cannot be an unsafe reflex vertex of P by (R1), since it is the left endpoint of segment $ab \in \widehat{S}_1$. Segment ab remains an edge of the lower arc P in the remainder of Phase 1, since neither operation modifies such edges of P . If c' is an interior vertex of $\text{carc}(b, a, c)$, then it is a convex vertex of P , and remains convex. Assume that $c' = c$, and note that c cannot be an unsafe reflex vertex or else it would have been created in a previous **BuildCap** operation, which would imply that $bc \in \widehat{S}_1$, contradicting $ab \in \widehat{S}_1$. At the end of operation $\text{BuildCap}(P, \varrho(a), a)$, the polygon $\text{conv}(a, b, c')$ is empty and lies in the exterior of P , so vertex b satisfies (R1) at that time.

We show that if an unsafe reflex vertex v_i of P satisfies (R1) in the course of Phase 1, it continues to satisfy (R1) in the remainder of Phase 1. We may assume that $v_i = b$ is the right endpoint of $ab \in \widehat{S}_1$, $v_{i-1} = a$, and $v_{i+1} = c'$. As noted above, ab remains an edge of P in Phase 1, and neither a nor c' can become an unsafe reflex vertex. Since c' is on the lower arc of P , we have $\varrho(c') = \varrho(a)$, and subsequent **BuildCap** operations cannot modify bc' . However, a subsequent **Dip** operation may modify bc' . Suppose $\text{Dip}(P, a', b')$ is the first such operation, and ray $\overrightarrow{b'a'}$ hits bc' at x' . This operation replaces bc' by $\text{carc}(b, x', a')$ and $\text{carc}(a', x', c')$, where a' is a safe reflex vertex, and all interior vertices of $\text{carc}(b, x', a')$ and $\text{carc}(a', x', c')$ are convex. Denote by a'' the neighbor of b in $\text{carc}(b, x', a')$. We need to show that $\text{conv}(a, b, a'')$ lies in the exterior of P after the operation, or equivalently, that a'' is to the right of the line through ac' .

By condition (b) in Phase 1, we have $a'b' \in Q_i$. Therefore, a' is on the left side of ℓ_i . Since the segments in Q_i are sorted in decreasing order by slope, both a' and x' are below the line through ab , consequently $\text{carc}(b, x', a')$ lies in the triangle formed by the lines through ab , $a'b'$, and ax' . The line through $a'b'$ crosses both bc' and ac' , hence a'' is to the right of the line through ac' , as required. This

completes the proof that (R1) is maintained. \square

The next lemma helps identify the segments in Q_i , $i \in \{1, \dots, r\}$, whose left endpoints are *not* in P .

Lemma 5. *Let P be the frame returned by Phase 1. Let $i \in \{1, \dots, r\}$, and let $Q_i = \{a_j b_j : j = 1, \dots, |Q_i|\}$ be sorted in increasing order by the y -coordinates of $a_j b_j \cap \ell_i$. If a_j is a vertex in the lower (resp., upper) arc of P , then so is $a_{j'}$ for all $j' < j$ (resp., $j' > j$).*

Proof. Assume that a_j is a vertex in the lower arc of P (the argument is analogous when a_j is in the upper arc). Suppose, for contradiction, that the lower arc of P does not contain $a_{j'}$ for some $1 \leq j' < j$. By (F2), $a_{j'}$ is in the interior of P , and so $\overrightarrow{b_{j'} a_{j'}}$ hits a segment in \widehat{S}_1 or an edge of P . We claim that it hits an edge uv of P such that $uv \notin \widehat{S}_1$ and the left endpoint of uv is an endpoint of some segment in Q_i . It follows that the algorithm would have applied $\text{Dip}(P, a_{j'}, b_{j'})$ and P would contain $a_{j'}$, contradicting the assumption.

To prove the claim, note that Q_i is sorted by slope and so $\overrightarrow{b_{j'} a_{j'}}$ does not cross any segment in Q_i . Since $j' < j$, point a_j is above the line through $b_{j'} a_{j'}$. By (F2), the lower arc crosses ℓ_i at some point p_i below $\ell_i \cap a_{j'} b_{j'}$. Let γ be the portion of the lower arc of P between a_j and p_j . By (F7), the ray $\overrightarrow{b_{j'} a_{j'}}$ crosses γ . By (F6), γ is in the right halfplane bounded by L_{i-1} . Since all segments in \widehat{S}_1 between L_{i-1} and L_i are in Q_i , the ray $\overrightarrow{b_{j'} a_{j'}}$ hits some edge uv of P , where $uv \notin \widehat{S}_1$. The edge uv can cross neither $a_j b_j$ nor γ (it may be an edge of γ), therefore the left endpoint of uv is to the left of ℓ_i , hence it is an endpoint of a segment in Q_i . \square

By Lemma 5, the line segments in \widehat{S}_1 whose left endpoints are not in P form a continuous interval. That is, for every $i \in \{1, \dots, r\}$, there is a set of consecutive indices $M_i \subseteq \{1, \dots, |Q_i|\}$ (possibly $M_i = \emptyset$ or $M_i = \{1, \dots, |Q_i|\}$) such that $j \in M_i$ if and only if the left endpoint of $a_j b_j$ is *not* in P . Let $I = \{i \in \{1, \dots, r\} : M_i \neq \emptyset\}$, $S'_i = \{a_j b_j : j \in M_i\}$ for all $i \in I$, and $S' = \bigcup_{i \in I} S'_i$.

Phase 2: Middle Segments. In Phase 2, we use the **ShearDip** and **Dip** operations to reach the left endpoints of *at least quarter* of the segments in S' , followed by **BuildCap** operation to reach the right endpoints of those segments if necessary. For all $i \in I$, let $a_{t(i)}$ be the leftmost left endpoint of a segment in the set S'_i , and denote this segment by $a_{t(i)} b_{t(i)}$. Refer to Fig. 5. We show (in Lemma 6 below) that a **ShearDip**($P, a_{t(i)}, x_i$) operation can stretch the frame to reach $a_{t(i)}$ from some suitable point x_i on an edge of P . We would like to choose all points x_i on the same (upper or lower) arc of P . We decide which arc we use by comparing the number of segments in S'_i above and below $a_{t(i)} b_{t(i)}$ for all $i \in I$. The set of segments above and below are A and B , respectively, defined as follows:

$$A = \bigcup_{i=1}^r A_i, \quad \text{where } A_i = \{a_j b_j : j \geq t(i), j \in M_i\},$$

$$B = \bigcup_{i=1}^r B_i, \quad \text{where } B_i = \{a_j b_j : j \leq t(i), j \in M_i\}.$$

If $|A| \geq |B|$, then we reach the vertices $a_{t(i)}$ from the upper arc for all $i \in I$; otherwise we reach them from the lower arc. Without loss of generality, assume that $|A| \geq |B|$.

The operations $\text{ShearDip}(P, a_{t(i)}, x_i)$ may create unsafe reflex vertices at $a_{t(i)}$, $i \in I$. We need to ensure that $a_{t(i)}$ will not become an interior vertex of any carc created by subsequent BuildCap operations. We can protect a vertex $a_{t(i)}$, $i > 1$, from a subsequent carc by ensuring that edges of the lower arc that cross the line L_{i-1} do not change in phases 2-3 of the algorithm. We can guarantee this property for only some of the edges of the lower arc by setting the orientations $\varrho(v)$ for vertices of the frame that are right endpoints of segments in \widehat{S}_1 . Let f_1, \dots, f_h be the edges of the lower arc that cross any of the lines L_1, \dots, L_{r-1} , ordered from left to right; note that $h \leq r-1$ as edge may cross several vertical lines. Now let

$$J = \{i \in I : L_{i-1} \text{ crosses } f_j \text{ where } j \text{ is odd}\}.$$

If $|\bigcup_{i \in J} A_i| \geq |\bigcup_{i \in I \setminus J} A_i|$, then we apply $\text{ShearDip}(P, a_{t(i)}, x_i)$ for all $i \in J$; otherwise for all $i \in I \setminus J$. Without loss of generality, assume that $|\bigcup_{i \in J} A_i| \geq |\bigcup_{i \in I \setminus J} A_i|$. We now assign orientations to the endpoints of the some edges f_j as follows. If v is the right (resp., left) endpoint of an edge f_j , where j is odd, then let $\varrho(v) = -1$ (resp., $\varrho(v) = 1$). Recall that the default orientation for all other vertices is $\varrho(v) = 1$ if v is on the lower arc of P , and $\varrho(v) = -1$ otherwise.

Before presenting Phase 2 of the algorithm, we need to specify the points x_i for all $i \in J$. Consider the vertical upward ray from $a_{t(i)}$, and let u_i be the first point on the ray that lies in the upper arc of P or on a segment in Q_i . If u_i is in an edge of the upper arc, but not in a segment in Q_i , then let $x_i := u_i$. Otherwise, u_i lies in some segment $ab \in Q_i$, which is either an edge or an internal diagonal of P . Let a_0a be the edge on the upper arc that precedes a (in clockwise order). The following lemma shows that we can choose a point x_i in the interior of a_0a .

Lemma 6. *There exists a point x_i in the interior of a_0a such that $a_{t(i)}x_i$ does not cross P or any segment in \widehat{S}_1 .*

Proof. We first claim that the interior of the triangle $\Delta(a_{t(i)}u_i a)$ is disjoint from segments in \widehat{S}_1 . Suppose, to the contrary, that segment $a'b' \in \widehat{S}_1$, with left endpoint a' , intersects the interior of $\Delta(a_{t(i)}u_i a)$. By construction, $a_{t(i)}b_{t(i)}, ab \in Q_i$, then $\Delta(a_{t(i)}u_i a)$ lies between L_{i-1} and L_i , and $a'b' \in Q_i$. This means that $a'b'$ crosses ℓ_i . The triangle $\Delta(a_{t(i)}u_i a)$ is strictly left of ℓ_i . Therefore $a'b'$ must cross the boundary of $\Delta(a_{t(i)}u_i a)$. However $a'b'$ can cross neither $a_{t(i)}u_i$ (by construction) nor au_i (as $au_i \subset ab$). Hence $a'b'$ must cross the edge $a_{t(i)}a$ of $\Delta(a_{t(i)}u_i a)$, and in particular the left endpoint a' is in the interior of $\Delta(a_{t(i)}u_i a)$. Since $a_{t(i)}$ is the leftmost left endpoint of a segment in Q_i that is not a vertex in P , we know that a' is a vertex of P . By (R1), a' is a convex or safe reflex vertex in P .

Consider a maximal subchain of P that intersects the interior of $\Delta(a_{t(i)}u_i a)$. As the frame cross neither $a_{t(i)}u_i$ nor $au_i \subset ab$, both endpoints of such a subchain are on the side $a_{t(i)}a$ of $\Delta(a_{t(i)}u_i a)$. Consequently, its farthest vertex from $a_{t(i)}a$ is a reflex vertex of P . Without loss of generality, this reflex vertex is a' . Since $a'b'$ crosses the side $a_{t(i)}a$ of $\Delta(a_{t(i)}u_i a)$, it cannot subdivide the reflex angle of P into two convex angles. Hence, a' cannot be a convex or safe reflex vertex of P . This contradiction completes the proof of the claim.

It follows that the edges of P are disjoint from the interior of $\Delta(a_{t(i)}u_i a)$. In particular, edge a_0a is disjoint from the interior of $\Delta(a_{t(i)}u_i a)$. By (R1), a is either a convex vertex or a safe reflex vertex of P . In either case the angular domain $\angle a_0ab$ is convex, hence it contains $a_{t(i)}a$. As no three segment endpoints are collinear, there exists a point x_i in the interior of a_0a (in some neighborhood of a) such that $a_{t(i)}x_i$ crosses neither P nor any segment in \widehat{S}_1 . \square

Phase 2 proceeds as follows:

1. For all $i \in J$, do:

set $P := \text{ShearDip}(P, a_{t(i)}, x_i)$.

2. While condition (a) or (b) below is applicable, do:

(a) If there exists a segment $ab \in \widehat{S}_i$ such that the left endpoint a is a vertex of P , but the right endpoint is not, then set $P := \text{BuildCap}(P, \rho(a), a)$.

(b) Else if there exists a segment $ab \in Q_i$ for some $i \in \{1, \dots, r\}$ such that ab lies in the interior of P , a is the left endpoint, and \vec{ba} hits an edge uv where $uv \notin \widehat{S}_i$, and the left endpoint of uv is an endpoint of some segment in Q_i , then set $P := \text{Dip}(P, a, b)$.

3. Return P , and terminate Phase 2.

We show that Phase 2 maintains a frame. Instead of addressing (R1), as in Phase 1, we use properties (R2)–(R3) below to establish (F5).

Lemma 7. *All operations in Phase 2 maintain (F1)–(F7) and the following properties:*

(R2) *every vertex $a_{t(i)}$, $i \in J$, has multiplicity 1 in P ;*

(R3) *if v_i is an unsafe reflex vertex of P and $v_i \notin \{a_{t(j)} : j \in J\}$, then v_i is the right endpoint of a segment in \widehat{S}_1 , vertices v_{i-1} and v_{i+1} are convex or have multiplicity 1 and the triangle (v_{i+1}, v_i, v_{i-1}) is disjoint from the interior of P .*

Proof. The **ShearDip** operations can be performed independently. They each maintain (F1)–(F4) and (F6)–(F7) of the frame. We show that they also maintain (F5). The **ShearDip** operations introduce reflex vertices at $a_{t(i)}$ for all $i \in J$. Importantly, these operations do not modify the lower chain of P . By Lemma 4, P has property (R1) at the beginning of Phase 2. Consequently, after performing all **ShearDip** operations in Phase 2, (R1) holds for all unsafe reflex vertices with the possible exception of $a_{t(i)}$, $i \in J$, which have multiplicity 1. This implies (R2)–(R3), hence (F5).

The **while** loop of Phase 2 is similar to Phase 1, but we need to prove that the multiplicity of every vertex $a_{t(i)}$, $i \in J$, remains 1. Specifically, we show that the changes incurred by **BuildCap** and **Dip** operations in the **while** loop are limited to regions that do not contain the vertices $a_{t(i)}$. For all $i \in J$, let L_i^- be a polyline that consists of $x_i a_{t(i)}$ and a vertical line from $a_{t(i)}$ to the lower arc of P . Operation **ShearDip** $(P, a_{t(i)}, x_i)$ creates two convex arcs incident to $a_{t(i)}$, on opposite sides of L_i^- : The interior vertices of the **carc** on the left of L_i^- are to the left of $a_{t(i)}$, and by (R1), they cannot be left endpoints of segments in Q_i . Therefore this **carc** is not involved in subsequent **BuildCap** or **Dip** operations for any segment in Q_i . That is, the impact of these operations is confined to the region to the right of L_i^- and left of L_j^- , where $j > i$ is the next index in J (if it exists). This further implies that the multiplicity of $a_{t(i)}$ remains 1 in the remainder of Phase 2, hence (R2) is maintained.

Analogously to the proof of Lemma 4, all **BuildCap** and **Dip** operations in the **while** loop maintain (F1)–(F4) and (F6)–(F7). They also maintain (R1) for unsafe reflex vertices other than $a_{t(i)}$, $i \in J$, which implies (R2). The combination of (R2) and (R3) yields (F5). \square




Figure 7: The set \hat{S}_1 from Fig. 5, and frame P at the end of Phase 2. Vertex v has multiplicity 2.

Lemma 8. *At the end of Phase 2, P visits both endpoints of all segments in $\bigcup_{i \in J} A_i$. Consequently, P visits both endpoints of at least quarter of the segments in \hat{S}_1 .*

Proof. The proof of the first claim is analogous to the proof of Lemma 5. To prove the second claim, note that $a_{t(i)}$ is a vertex of the upper arc of P for all $i \in J$. Consequently $a_{j'}$, for all $j' > t(i)$, is also part of the upper arc. The **BuildCap** operations in the while loop of Phase 2 ensure that if a left endpoint of a segment $a_{j'}b_{j'} \in Q_i$ is a vertex of P , then so is the right endpoint. \square

Phase 3: Right Endpoints. At the end of Phase 2, P visits both endpoints of every segment in $\bigcup_{i \in J} A_i$. However, it may visit only one endpoint of some segments in $\bigcup_{i \in I \setminus J} A_i$ and $B \setminus A$. In this phase, we use **BuildCap** operations to ensure that P visits *both* endpoints of these segments. Recall that we have already specified the orientation $\varrho(v)$ for every vertex v of the frame P , depending on whether v is the left or the right endpoint of a segment in \hat{S}_1 . Phase 3 proceeds as follows.

1. While condition (a) below is applicable, do
 - (a) If there exists a segment $ab \in \hat{S}_i$ such that one endpoint, say b , is a vertex of P , but the other endpoint is not, then set $P := \text{BuildCap}(P, \varrho(b), b)$.
2. Return P , and terminate Phase 3.

At the end of Phase 3, we obtain a frame P that contains, for each segment, either both endpoints or neither endpoint; see Fig. 7. Some vertices may have multiplicity 2, but the multiplicity of the special vertices $a_{t(i)}$, for all $i \in J$, remains 1.

Lemma 9. *All operations in Phase 3 maintain (F1)–(F5) and (R2) for P .*

Proof. Phase 3 applies only **BuildCap** operations, which always maintain (F1)–(F4).

We establish (R2). Similarly to Phase 2, the impact of any **BuildCap**($P, \varrho(b), b$) operation, where b is a vertex of the upper arc, is confined to the region to the right of L_i^- and left of L_j^- , where $j > i$ is the next index in J (if it exists). We show that the same holds if b is a vertex of the lower arc of P . Note that Phase 2 modifies only the upper arc of P , and so the lower arc at the beginning of Phase 3 is the same as at the end of Phase 1. Suppose, to the contrary, that Phase 3 modifies an edge that crosses L_i^- , for some $i \in J$. Let **BuildCap**($P, \varrho(b), b$) be the first

such operation. Letting c be the neighbor of b in the frame in $\varrho(b)$ orientation, $\text{BuildCap}(P, \varrho(b), b)$ replaces edge bc of P with the polygonal path $ba + \text{carc}(a, b, c)$. By the choice of $a_{t(i)}$, if a segment in \hat{S}_1 crosses the vertical line through $a_{t(i)}$, it is in Q_i and both of its endpoints are vertices of the frame at the end of Phase 1. Consequently, b cannot be the left endpoint of a segment in Q_i , and so bc crosses the vertical line L_{i-1} . At the beginning of Phase 3, L_{i-1} crosses only one edge of the lower arc by (F6), say $f_{j'}$, where j' is odd by the definition of J . However, the BuildCap operations in Phase 3 do not modify edge $f_{j'}$ due to carefully chosen orientation $\varrho(b)$. Indeed, if b is to the left of L_{i-1} , then $\varrho(b) = -1$, otherwise $\varrho(b) = 1$. In either case, both b and c are on the same side of L_{i-1} , and bc does not cross L_{i-1} . This completes the proof that Phase 3 maintains (R2).

It remains to handle (F5). Phases 1–3 use only operations BuildCap , Dip , and ShearDip , each of which creates at most one reflex vertex. The reflex vertices created by Dip operations are safe their multiplicities does not increase by Lemma 3. The reflex vertices created by ShearDip operations have multiplicity 1 by (R2). Consider operation $\text{BuildCap}(P, \varrho(b), b)$, where $ab \in \hat{S}_1$ and c is the neighbor of b in the frame in $\varrho(b)$ orientation. It creates a reflex vertex at a ; the two neighbors of a are b and c' , where c' is the first vertex of $\text{carc}(a, b, c)$ adjacent to a (possibly $c' = c$). Edge ab remains an edge of P until the end of Phase 3, since $ab \in \hat{S}_1$, and our operations do not modify such edges. Vertex b cannot be a reflex vertex created by a prior BuildCap operation. Assume that after operation $\text{BuildCap}(P, \varrho(b), b)$, vertex a is part of a chain of two or more unsafe reflex vertices created by BuildCap operations. Then $\text{carc}(a, b, c) = ac$ and c is a reflex vertex created by a previous operation $\text{BuildCap}(P', \varrho(d), d)$, where $cd \in \hat{S}_1$ and $\varrho(d) = -\varrho(b)$. As noted above, segment cd remains an edge of the frame, and d is not a reflex vertex created by a BuildCap operation. It follows that a chain C of reflex vertices created by BuildCap operations has at most two vertices. If C has only one vertex, then C satisfies (F5) at that time, and possible Dip operations in Phases 1–2 maintain this property due to (R3) in Lemma 7. If C has length 2, say $C = (a, c) = (v_i, v_{i+1})$, then one of the reflex vertices was added to chain in Phase 3, and subsequent BuildCap operations do not modify the chain $(b, a, c, d) = (v_{i-1}, v_i, v_{i+1}, v_{i+2})$ of the frame. Assume, without loss of generality, that the reflex vertex v_i is created after v_{i+1} in a BuildCap operation. Before the operation, the triangle $(v_{i+2}, v_{i+1}, v_{i-1})$ is disjoint from the interior of P . Then the quadrilateral $(v_{i+2}, v_{i+1}, v_i, v_{i-1})$ is also disjoint from the interior of P after the operation, hence at the end of Phase 3. This confirms that Phase 3 maintains (F5). \square

Lemma 10. *Let P be the frame at the end of Phase 3. If one endpoint of a segment in \hat{S}_1 is a vertex in P , then so is the other endpoint.*

Proof. The claim trivially holds when the **while** loop terminates. \square

Phase 4: Obtaining a Simple Polygon. In the last phase of our algorithm, we set $P := \text{ChopWedges}(P)$. This is a valid operation by (F5). The resulting frame P is a simple polygon whose vertex set is the same as at the end of Phase 3. By Lemma 10, if one endpoint of a segment in \hat{S}_1 is a vertex in P , then so is the other endpoint. Consequently, P is a circumscribing polygon for a set of segments in \hat{S}_1 , which we denote by \hat{S}_2 . By Lemma 8, we have $|\hat{S}_2| \geq |\hat{S}_1|/4$. This completes the proof of Theorem 1.

3 Disjoint Segments versus Disjoint Rays

In this section, we give two sufficient conditions for an arrangement of disjoint segments to admit a circumscribing polygon. Both conditions involve extending the segments.

- (C1) A set S of n disjoint line segments is *extensible to rays* if there exists a set R of n disjoint rays, each of which contains a segment from S ; see Fig. 8 (left).
- (C2) A set S of n disjoint line segments admits *escape routes* if there exists an ordering and orientation of the segments $S = \{a_i, b_i : i = 1, \dots, n\}$ such that if we shoot a ray from b_i in direction $\overrightarrow{a_i b_i}$ for $i = 1, \dots, n$, then each ray either goes to infinity or intersects a previous ray (without crossing any segment in S); see Fig. 8 (right).




Figure 8: Left: an arrangement of disjoint segments extensible to rays. Right: an arrangement of disjoint segments that is not extensible to rays, but admits escape routes.

Clearly, (C1) implies (C2), but the converse is false in general. We can test (C1) in $O(n \log n)$ time. Indeed, there are two possible extensions for a segment, which can be encoded by a Boolean variable, and pairwise disjointness can be expressed by a 2SAT formula. However, we do not know whether condition (C2) can be tested efficiently. Here we show that (C1) and (C2) each imply the existence of a circumscribing polygon.

Theorem 11. *If S is a set of disjoint line segments satisfying (C2), then there is a circumscribing polygon for S .*

Proof. By (C2), we may assume that $S = \{a_i b_i : i = 1, \dots, n\}$ such that if we shoot a ray from b_i in direction $\overrightarrow{a_i b_i}$ for $i = 1, \dots, n$, then each ray either goes to infinity or intersects a previous ray. We call the part of the ray $\overrightarrow{a_i b_i}$ from b_i to the first point where it intersects a previous ray or $\partial \text{conv}(S)$ the *extension* of segment $a_i b_i$.

Given the ordering and orientation of the segments in S , we construct a circumscribing polygon using the following algorithm, using the operations **BuildCap** and **Dip** introduced in Sect. 2; see Fig. 9 for an example.

1. Initialize $P := \partial \text{conv}(S)$.
2. For $i = 1$ to n :

If b_i is not a vertex of P , then set $P := \text{Dip}(P, b_i, a_i)$.

3. For $i = 1$ to n :

If a_i is not a vertex of P , then set $P := \text{BuildCap}(P, 1, b_i)$.

4. $\text{ChopWedges}(P)$.

5. Return P .

We show that polygon P maintains (F1)–(F5) from Definition 2 (note that (F6)–(F7) are no longer applicable); and it also maintains the following invariant:

(F8) After operation $\text{Dip}(P, b_i, a_i)$ (for $i = 1, \dots, n$), the extension of segment $a_i b_i$ lies in the exterior of P .

Indeed, operations BuildCap and Dip always maintain (F1)–(F4). For $i = 1, \dots, n$, (F8) follows from (C2) and the definition of the Dip operation; and it guarantees the correctness of subsequent Dip operations.

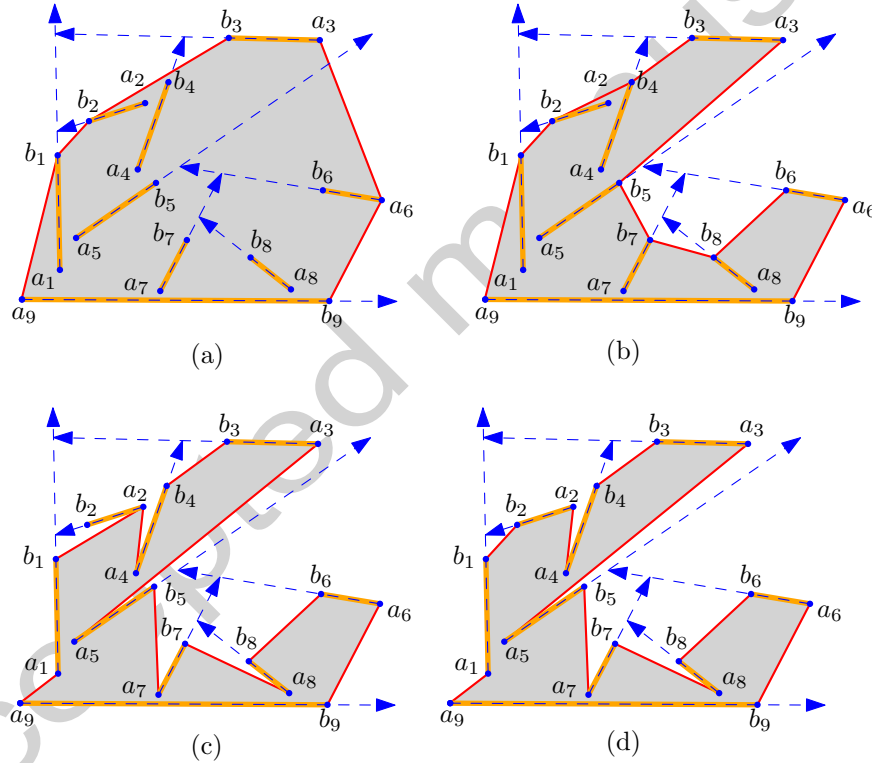


Figure 9: (a) An arrangement of nine segments that admit escape routes from Fig. 8 (right), and $P = \partial \text{conv}(S)$. (b) Polygon P after the for loop of Dip operations. (c) Polygon P after the for loop of BuildCap operations. (d) The circumscribing polygon for S after the ChopWedges operation.

For (F5), note that each $\text{Dip}(P, b_i, a_i)$ operation creates a unique reflex vertex at b_i , and it is safe. At the end of the for loop of Dip operations, the vertices of P are either convex vertices or safe reflex vertices. It remains to consider the for loop of BuildCap operations.

Each $\text{BuildCap}(P, 1, b_i)$ operation creates a unique reflex vertex at a_i , which is unsafe and incident to both $a_i b_i$ and some edge $a_i c_i$. Since $a_i b_i \in S$, subsequent operations $\text{BuildCap}(P, 1, b_j)$, $i < j$, modify neither $a_i b_i$ nor $a_i c_i$. Consequently, b_i and c_i are each convex or safe reflex vertices of P before operation $\text{BuildCap}(P, 1, b_i)$. As the operation can only decrease the interior angles at existing vertices, b_i and c_i remain convex or safe reflex vertices in the remainder of the algorithm. This means that the maximum chain of unsafe reflex vertices of P that contain a_i consists of a single vertex a_i . As the triangle $\text{conv}(a_i, b_i, c_i)$ is in the exterior of P after the operation $\text{BuildCap}(P, 1, b_i)$, (F5) is maintained, as required. By (F5), the final ChopWedges operation is correct, and the algorithm returns a circumscribing polygon for S . \square

The above results link the circumscribing polygon problem to the problem of extending line segments to rays. We now give an upper bound for the latter problem that seems to imply that the lower bound of the former problem (Theorem 1) is tight.

Lemma 12. *For every $n \in \mathbb{N}$, there is a set S of n disjoint line segments in the plane such that the cardinality of every subset $S' \subseteq S$ that admits an escape route is $|S'| \leq 2\lceil\sqrt{n}\rceil - 1$.*

Proof. First, assume that $n = k^2$ for some $k \in \mathbb{N}$ (we will consider the other case later). Consider the circle C of unit radius centered at the origin, and place k points a_1, \dots, a_k in clockwise order along the arc of C lying in the first quadrant (refer to Fig. 10). For every $i \in \{1, \dots, k\}$, create a line segment $a_i b_i$ of length 4 tangent to C such that a_i is its left endpoint; let $S_0 = \{a_i b_i : i = 1, \dots, k\}$. Finally, for every $i \in \{1, \dots, k\}$, let C_i be a circle of unit radius passing through b_i and tangent to $a_i b_i$. We create $k - 1$ additional segments that are almost parallel to $a_i b_i$. That is, the difference in slopes between any two of them is at most ε (for some ε small enough so that none of the segments cross). We make these segments all of equal length (four units) and all tangent to C_i . Let S_i represent the k segments tangent to C_i , $i \in \{1, \dots, k\}$. This completes the construction for a set of segments $S = \bigcup_{i=1}^k S_i$. Notice that if we pick a sufficiently small $\varepsilon > 0$ no two segments will cross.




Figure 10: Our lower bound construction for $k = 4$.

Let $Q \subset S$ be a subset of segments that admit escape routes. We claim that $|Q| < 2k$. Indeed, let $i_1 < \dots < i_k$ be the indices such that $Q \cap S_{i_j} \neq \emptyset$ (for all $j \leq k$).

First note that $Q \cap S_{i_j}$ can contain only one segment for all $j < k$: the left extension of each segment in $Q \cap S_{i_j}$ must hit any segment in $Q \cap S_{i_{j+1}}$, thus all segments in $Q \cap S_{i_j}$ must be directed

to the right. However, if $|Q \cap S_{i_j}| \geq 2$, then the right extensions of one of these segments hits another segment in the same set. Therefore $|Q \cap S_{i_j}| = 1$ (for all $j \leq k$). Overall, we have k different groups S_i , each with k segments. Any subset that admits an escape route contains at most one segment from all but one group and can potentially contain all segments of the last one. Since there are k groups, the result follows.

To complete the proof it remains to consider the case in which n is not a perfect square. The key trick is that our construction is hereditary (that is, even if we remove a few segments from our construction, the same upper bound holds). Thus, given $n \in \mathbb{Z}$, apply the construction described above for $k = \lceil \sqrt{n} \rceil$. This construction will have $\lceil \sqrt{n} \rceil^2 \geq n$ segments. Remove segments arbitrarily until only n remain, and apply the same argument to any subset that admits an escape route. This will give an upper bound of $2k - 1 = 2\lceil \sqrt{n} \rceil - 1$ as claimed. \square

4 Hardness for Circumscribing Polygons

In this section we prove that it is NP-hard to decide whether a given set of disjoint line segments admits a circumscribing polygon. We start by proving NP-hardness for circumscribing a set of disjoint straight-line cycles (Sect. 4.1), and then modify this reduction so that it works for disjoint line segment, as well (Sect. 4.2).

4.1 Hardness for Disjoint Straight-Line Cycles

Garey et al. [9, p. 713] proved that HAMILTONIAN PATH IN 3-CONNECTED CUBIC PLANAR GRAPHS (HP3CPG) is NP-complete. They reduce from 3SAT, and their reduction produces 3-connected planar graphs in which any Hamiltonian path connects two possible vertex pairs. This implies that the problem remains NP-complete if one endpoint of a Hamiltonian path is given. Since the graph is 3-regular, the problem remains NP-complete if the first (or last) edge of a Hamiltonian path is given. We call this problem HAMILTONIAN PATH IN 3-CONNECTED PLANAR CUBIC GRAPHS WITH START EDGE (HP3CPG-SE): Given a 3-connected cubic planar graph $G = (V, E)$ and an edge $uv \in E$, decide whether G has Hamiltonian path whose first edge is uv . We reduce HP3CPG-SE to deciding whether a given PSLG admits a circumscribing polygon.

Theorem 13. *It is NP-complete to decide whether a given PSLG admits a circumscribing polygon, even if the PSLG is regular-degree-2.*

Proof. Let $G = (V, E)$ and $uv \in E$ be an instance of HP3CPG-SE. We construct a PSLG $\hat{G} = (\hat{V}, \hat{E})$ in four steps below, and then show that G has a Hamiltonian circuit containing uv if and only if \hat{G} admits a circumscribing polygon.

Construction of a PSLG $\hat{G} = (\hat{V}, \hat{E})$. Let $G = (V, E)$ be a 3-connected cubic planar graph, $uv \in E$, and let $n = |V|$. We construct a PSLG $\hat{G} = (\hat{V}, \hat{E})$ as follows.

Step 1. We first modify G in the neighborhood of vertex u , and then specify a straight line embedding as follows. Refer to Fig. 11. Subdivide each of the three edges incident to u into a path of length 2, where uv is subdivided into (u, u', v) . Apply a $Y\Delta$ -transform at u , which creates a triangular face incident to u' , denoted $\Delta_u = \Delta(abu')$. Let G_1 be the resulting graph. Note that G_1

is a 3-connected cubic planar graph, and G contains a Hamiltonian path starting with uv if and only if G_1 contains a Hamiltonian path starting with (a, b, u', v) .

We construct an embedding of G_1 in which the outer face is Δ_u . Using a result by Chambers et al. [5], G_1 admits a straight-line embedding in \mathbb{R}^2 such that (i) Δ_u is the outer face, (ii) every interior face is strictly convex, (iii) every vertex has integer coordinates, and (iv) the coordinates of the vertices are bounded by a polynomial in n , which does not depend on G_1 . Such embedding can be computed in polynomial time. (The method in [5] improves Tutte's spring embedding [21], which may require coordinates exponential in n .)




Figure 11: (a) A 3-connected cubic graph G with special edge uv . (b) Graph G_1 after a $Y\Delta$ -transform around vertex u . (c) A convex embedding of G_1 such that the outer face is Δ_u . (d) Rotating the supporting line of internal edges. (e) Thickening the edges into corridors.

Step 2. We call an edge of G_1 *external* if it bounds the outer face, and *internal* otherwise. Let δ be the *feature ratio* of G_1 , i.e., the minimum distance between a vertex and a nonincident edge. Delete the internal edges adjacent to a and b ; and denote by G_2 the resulting PSLG.

Step 3. In this step, we split each internal vertex of degree 3 in G_2 into three vertices via a $Y\Delta$ -transform, creating a new triangular face $\Delta(w_1w_2w_3)$. Refer to Fig. 12. To do so, simultaneously rotate the supporting line of every internal edge about the midpoint of the edge in general, and about the endpoint u' in case of edge $u'v$, by the same angle θ . Since w has degree 3, the supporting lines of edges incident to w will intersect at three different points: w_1 , w_2 , and w_3 . We choose the angle θ to be the maximum rotation so that w_1 , w_2 , and w_3 are still within $\delta/2$ distance from the original location of w for all internal vertices w in G' . Let C_w be the circle obtained by dilating the inscribed circle of $\Delta(w_1w_2w_3)$ by $1/2$ from its center. Place a new edge w_4w_5 on the horizontal diameter of C_w . Denote by G_3 the resulting PSLG.

Step 4. We define visibility for a PSLG as follows: Two points in the plane are *visible* if the line segment between them is disjoint from the interior of opaque faces, where we consider the triangles obtained by $Y\Delta$ -transforms to be *transparent*, and all other faces *opaque*. Opaque faces are shown gray in Figs. 11 (d) and 12 (b). Simultaneously shrink the opaque faces of G_3 using their straight skeleton by half of the minimum amount that would create a new visibility relation between a pair of vertices. We assume that vertices w_1 , w_2 , and w_3 remain incident to the opaque faces in which they are strictly convex corners: in particular, each vertex belongs to a unique opaque face. Note that w_4 and w_5 are only visible from w_1 , w_2 , and w_3 . We call the pair of parallel edges that used to be collinear a *corridor*, and the triangle $\Delta(w_1w_2w_3)$ a *chamber* (which is no longer a face). Denote by $\hat{G} = (\hat{V}, \hat{E})$, the resulting PSLG. This concludes the construction of \hat{G} .




Figure 12: A chamber incident to three corridors and a segment in the chamber that sees only three other vertices.

Equivalence. We claim that an HP3CPG-SE instance G with edge uv is positive if and only if \hat{G} admits a circumscribing polygon. Assume G admits a Hamiltonian path starting with edge uv . Then G_1 contains a Hamiltonian path P_1 that starts with (a, b, u', v) . We obtain a circumscribing polygon \hat{P} of \hat{G} incrementally as follows. Initialize \hat{P} to be the convex hull of \hat{G} . Refer to Fig. 13 (a). For each degree-2 vertex w in P_1 , add to \hat{P} both edges of each of the two corridors corresponding to the edges in P_1 incident to w . Without loss of generality, assume that one edge from each of the two corridors is incident to vertex w_3 . If w corresponds to a chamber, add the path (w_1, w_5, w_4, w_2) to \hat{P} (which connect the other two edges of the corridors). Refer to Figs. 13 (b-c). For the last vertex w in P_1 , if w corresponds to a chamber, w.l.o.g., let w_2 and w_3 be the endpoints of the corridor that corresponds to the edge in P adjacent to w . Add to \hat{P} the path $(w_3, w_4, w_5, w_2, w_1)$. Else, add the edge connecting the endpoints of such corridor. By construction, every opaque face is in the interior of \hat{P} . Every edge of \hat{G} is either an edge of P_1 or a corridor in the interior of \hat{P} . Hence, \hat{P} is a circumscribing polygon.




Figure 13: A circumscribing polygon \hat{P} at endpoints of corridors and a solution to the reduction in Fig. 11. The shaded area is in the interior of \hat{P} .

Now assume that \hat{G} admits a circumscribing polygon \hat{P} . Note that the interior of an opaque face must be in the interior of \hat{P} , or else some edge bounding the opaque face would be an external diagonal. Let I be the (closed) convex hull of a corridor (recall that a corridor consists of two parallel edges of \hat{P}). The corridors are constructed so that all visibility edges that intersect I are induced by the four vertices of I . Furthermore, the number of edges of \hat{P} in I is exactly 0 or 2, or else one of the adjacent opaque faces would be in the exterior of \hat{P} (in case \hat{P} has three edges in I) or such edges would induce a cycle (in case \hat{P} has four edges in I). If I contains two edges of \hat{P} , then either these are the edges of the corridor, or they are adjacent. If they are adjacent, one of

them is a diagonal of I and the other is an edge of the corridor; we call such an adjacent pair of edges a *spike*.

We define the *degree of a chamber* as the number of incident corridors in which both edges are edges of \hat{P} . We claim that the degree of every chamber is 1 or 2. To prove the claim, consider a chamber in which the vertices are labelled w_1, \dots, w_5 as in Fig. 12(c). Notice that w_4 and w_5 are visible from only three other vertices, and so \hat{P} must contain a path connecting two of the three vertices through the interior of the chamber. Assume, without loss of generality, that such a path is (w_2, w_4, w_5, w_1) as in Fig. 13(a). It follows that the degree of a chamber cannot be three, or else w_1 and w_2 would have degree 3 in \hat{P} . Since the path (w_2, w_4, w_5, w_1) is adjacent to the rest of the circuit \hat{P} , it must have edges intersecting the convex hull I of some corridor incident to the chamber. If the degree of a chamber is zero, it can only be adjacent to spikes. Then (w_2, w_4, w_5, w_1) and the adjacent spikes form a cycle that does not span all vertices of \hat{G} , which is a contradiction. This completes the proof of the claim.

Consider a chamber of degree two, with vertices labeled as in Fig. 12(c). Without loss of generality, \hat{P} contains the corridors incident to w_1 and w_3 , and w_2 and w_3 as in Fig. 13(a). Neither of the edges in the corridor incident to w_2 and w_3 can be in \hat{P} or else either w_1 or w_2 would have degree 3 in \hat{P} . Consequently, \hat{P} intersects each chamber of degree 2 exactly twice: once by the single-vertex path (w_3) and once by (w_2, w_4, w_5, w_1) . Then, such chamber cannot be adjacent to any spikes. Now consider a chamber of degree one such that \hat{P} contains the corridor incident to w_1 and w_3 as in Fig. 13(b). Then, \hat{P} contains a path from w_2 to w_3 containing the edge w_4w_5 (as in Fig. 13(b)), or a path from w_1 to w_3 containing the edge w_4w_5 . In the first case, it is possible that the chamber is incident to a spike (w_1, v, w_2) where v is not a vertex in the chamber. In the latter, w_2 is not visited locally by the subset of \hat{P} intersecting the chamber and, therefore, must be visited by a spike adjacent to an adjacent chamber. Since these are the only possible occurrences of spikes, we can modify any solution by shortcutting spikes of the form (w_1, v, w_2) to edges w_1w_2 in the first case, and, in the latter case, modifying \hat{P} to traverse the chamber as in Fig. 13(b). Hence, from now on we may assume that \hat{P} does not contain any spikes.

For a circumscribing polygon \hat{P} of \hat{G} , we define a subgraph P_1 of G_1 as follows: If both edges of a corridor belong to \hat{P} , then let the edge corresponding to the corridor be in P_1 . Since the degree of each chamber in \hat{G} is 1 or 2, every interior vertex of G_1 has degree 1 or 2 in P_1 , hence P_1 is a union of paths. A path that is not adjacent to the corridor incident to u' will induce a cycle. Hence, there can be at most one path. Since every chamber must be visited by \hat{P} , this path visits every vertex except for a and b . Then, we can easily obtain a Hamiltonian path of G_1 starting with (a, b, u', v) ; and a Hamiltonian path starting with edge uv in the original graph G . \square

4.2 Hardness for Disjoint Line Segments

In this section, we modify the constructions from Sect. 4.1 and extend our hardness result to the case of disjoint line segments. In the previous reduction, disjoint cycles are used to trace out the convex faces of an embedding of a graph \hat{G} . To simulate a convex cycle with a geometric matching, we create *turn gadgets* that simulate the corners of a convex polygon.

Fig. 14 depicts an arrangement of six disjoint line segments, which we call a *diode*, that, by itself, does not admit a circumscribing polygon [10]. The visibility relationships within a diode that guarantee this property are the following.

(V1) A point in $\{d_5, d_6\}$ ($\{d_9, d_{10}\}$) can see at most three points outside of this set, namely d_3, d_4 ,




Figure 14: A diode

and d_7 (d_7 , d_{11} , and d_{12}).

(V2) Point d_1 (d_2) does not see d_3 (d_{11}).

Since polygonizations are affine invariants, a nondegenerate affine image of a diode does not admit a polygonization, either.

Using two affine copies of the diode as subconfigurations we can now construct a turn gadget. Fig. 15 shows this construction, but, for clarity, the positions of the points are not accurate. We first describe the important visibility properties in the gadget, and then show how to obtain coordinates for the points in order to achieve them. We refer to the points in the left diode as d_1, \dots, d_{12} as in Fig. 14, and to the points in the right diode as d'_1, \dots, d'_{12} . We refer to points in $\{p_1, q_1, a, b\}$ as *access points*, and to points in the two diodes, c_1, c_2, c_3 , or c_4 as *interior points*. We construct the gadget so that

(V3) interior points can only see other interior points and access points;

(V4) d_2 (d'_2) sees p_1 (q_1), but not q_1 (p_1);

(V5) c_3 and c_4 can only see c_1 , c_2 , d_2 , and d'_2 ;

(V6) a point in $\{d_3, \dots, d_{12}\}$ ($\{d'_3, \dots, d'_{12}\}$) can see at most three points outside of this set, namely d_1 , d_2 , and a (d'_1 , d'_2 , and b);

(V7) a point in $\{d_9, \dots, d_{12}\}$ ($\{d'_9, \dots, d'_{12}\}$) can see at most three points outside of this set, namely d_1 , d_7 , and d_8 (d'_1 , d'_7 , and d'_8).

Construction of the Turn Gadget. Each turn gadget is constructed by replacing a vertex w_i of a chamber $\Delta(w_1 w_2 w_3)$ as shown in Fig. 16. We simultaneously perform the steps below for every chamber vertex w_i . We describe the construction for w_1 ; refer to Fig. 16 (a). Let ε be one tenth of the feature ratio of the graph \widehat{G} . Let a' and b' be points on the edges incident to w_1 at distance ε from w_1 . If we consider w_1 transparent, by construction, the only vertex that can (partially) see the left side of segment $a'b'$ is w_2 . Let e be the intersection of the supporting line of $w_1 w_2$ and $a'b'$. By construction, e must lie in the interior of the segment from the midpoint of $a'b'$ to b' . We shrink the two edges incident to w_1 creating new endpoints p_1 and q_1 at distance $\delta > 0$ from w_1 . We choose





Figure 15: Turn gadget

$\delta > 0$ to be the maximum value so that, for every turn gadget, if any new endpoint (at distance δ from some vertex in \hat{V}) outside of a turn gadget sees a point $c \in ab$, then $|ce| \leq \min \{|ea'|, |eb'|\}/3$ (in Fig. 16 (b), δ is the radius of the dotted circles). At each turn gadget, the upper bound for δ can be expressed as a bounded-degree rational function of the coordinates of the gadget, therefore the minimum over all turn gadgets is also bounded-degree rational function over the coordinates of the vertices in \hat{V} .


 Figure 16: Turn gadget construction. For clarity, ε is not shown in proportion.

Assume, without loss of generality, that $a'b'$ is horizontal and w_1 lies above its supporting line. Add a segment $ab \subset a'b'$ by shrinking $a'b'$ so that a (b) is at distance at most $\min \{|ea'|, |eb'|\}/3$ from a' (b'). By construction, there is a triangle T_1 (T_2) on top of ab whose interior can only be seen by p_1 , q_1 , a , and b . Let T'_1 be the maximum triangle in T_1 so that it has a vertical side, i.e., whose supporting line is perpendicular to ab , so that p_1 is on or to the right of such supporting line; and define T'_2 analogously. As long as we construct the diodes in the interior of T'_1 and T'_2 , (V3) is satisfied and all remaining necessary visibility constraints (constraints (V1)–(V2) and (V4)–(V7) become local, i.e., the positions of the points depend only on other points of the gadget). Since the constraints translate to linear inequalities, where each line is spanned by two vertices, we can compute the coordinates of all the interior points of the turn gadget in polynomial time. Figure 16 (c) illustrates linear constraints by gray lines.

Properties of the Turn Gadget. The following four lemmas establish the key properties of a turn gadget. We show that if a set of disjoint line segments S contains a turn gadget, then in every circumscribing polygon for S , there is a polygonal chain between p_1 and q_1 that visits all vertices

of the gadget, and no other vertices (Lemma 17). We begin with a lemma that handles only one diode in the turn gadget.

Lemma 14. *Let S be a set of disjoint line segments that contains a turn gadget with properties (V1)–(V7). Every circumscribing polygon P for S contains a path \mathcal{P} on its boundary with one endpoint at a and the other at d_1 or d_2 and whose interior vertices are exactly $\{d_3, \dots, d_{12}\}$.*

Proof. Refer to Fig. 17 for examples of \mathcal{P} . Let \mathcal{B} be a the minimal path on the boundary of P containing $\{d_3, \dots, d_{12}\}$ such that its endpoints are in $\{d_1, d_2, a\}$. By (V6), \mathcal{B} exists and must contain only points in the left diode and a . Further, the two endpoints of \mathcal{B} cannot be d_1 and d_2 or else one could obtain a circumscribing polygon of a diode by removing a (if present) from \mathcal{B} , since the triangle $\Delta(ad_i d_j)$ formed by a and any pair of vertices in the diode visible from a is empty. However, as noted above, a diode does not admit a circumscribing polygon [10]. It follows that one endpoint of \mathcal{B} is a . Let v be the vertex in \mathcal{B} adjacent to a . Note that $v \notin \{d_1, d_2\}$, else we could omit a from \mathcal{B} , contradicting the minimality of \mathcal{B} . Then $\mathcal{B} = \mathcal{P}$ as required. \square




Figure 17: Examples of \mathcal{P} . The interior (exterior) of P is shown in gray (white).

Recall that every segment in S is either an edge or an internal diagonal of a circumscribing polygon P . The union of P and S forms a plane graph that subdivide the plane into faces. Clearly, each face is either in the interior or the exterior of P . We define the *left side* (resp., *right side*) of a directed segment $ab \in S$ as the face of this graph on the left (resp., right) of ab .

Lemma 15. *Let S be a set of disjoint line segments that contains a turn gadget satisfying (V1)–(V7). In every circumscribing polygon P for S , the left side of the segment $d_2 d_1$ is in the exterior of P , hence $d_2 d_1$ is an edge of P .*

Proof. Let \mathcal{P} be the path described in Lemma 14 from a to d_1 or d_2 , in which d_3, \dots, d_{12} are interior vertices. Suppose, for contradiction, that the left side of $d_2 d_1$ is in the interior of P . The left side of $d_2 d_1$ contains $\Delta(d_1 d_2 d_7)$, as no visibility edge crosses this triangle by (V2). Both edges of \mathcal{P} incident to d_7 are in the same closed halfplane bounded by the line through $d_7 d_8$, or else $d_7 d_8$ would be an external diagonal of P . Assume first that both edges are in the closed halfplane on the right of $d_7 d_8$. In particular, d_7 is adjacent to neither d_5 nor d_6 . By (V1), d_5 and d_6 must be adjacent to d_3 and d_4 , respectively, and $d_5 d_6$ is an edge of P . In particular $f = (d_3, d_4, d_6, d_7)$ is a face in $P \cup S$. Furthermore, since no edge in \mathcal{P} separates $\{d_5, d_6\}$ from $\Delta(d_1 d_2 d_7)$, the left side of $d_5 d_6$ is in the interior of P . This implies that the f is in the exterior of P . Now $d_3 d_4$ must be an edge of P , or it would be an external diagonal, and so P contains all edges of the face f , which is a contradiction. When both edges incident to d_7 are in the closed half plane on the left of $d_7 d_8$, the argument is analogous by the symmetry of the diode. \square

Lemma 16. *Let S be a set of disjoint line segments that contains a turn gadget satisfying (V1)–(V7). In every circumscribing polygon P for S , the left side of ab is in the exterior of P , hence ab is an edge of P .*

Proof. Suppose, for contradiction, that the left side of ab is in the interior of P . Let \mathcal{P} be the path described in Lemma 14 from a to d_1 or d_2 . Note that both d_1d_2 and $d'_1d'_2$ are edges of P by Lemma 15.

Let v and v' be the two neighbors of a in P that are in \mathcal{P} and not in \mathcal{P} , respectively. Then $v' \neq d_2$, otherwise $\mathcal{P} \cup (v, d_2, d_1)$ would be a closed polygon in P ; and $v' \neq d'_2$ otherwise the path (v, d'_2, d'_1) is part of P and the left side of ab is the same as the right side of $d'_2d'_1$, contradicting Lemma 15. It follows that the left side of ab and the right side of av are the same face of $P \cup S$, and so by assumption the right side of av is in the interior of P (cf. Fig. 17 (middle), where $v = d_8$).

We claim that the other endpoint of \mathcal{P} is d_1 (and not d_2). Suppose, to the contrary, that \mathcal{P} is a path from a to d_2 . The last edge of \mathcal{P} is d_1d_2 by Lemma 15. As \mathcal{P} traverses both av and d_1d_2 in the same direction, the interior of P is on the right side of both. This means that the left side of d_2d_1 is in the interior of P , which contradicts Lemma 15.

By a symmetric argument for the right diode, we conclude that there must be a subpath \mathcal{P}' on the boundary of P from b to d'_1 going only through vertices of the right diode. By (V5), c_3 and c_4 can only be adjacent to c_1 and c_2 . If the left side of ab is in the interior of P , then c_1c_2 is an external diagonal. This contradicts our initial assumption, and concludes the proof. \square

Lemma 17. *Let S be a set of disjoint line segments that contains a turn gadget satisfying (V1)–(V7). Every circumscribing polygon P for S must contain a chain between p_1 and q_1 whose interior vertices are exactly the vertices in the interior of the gadget and the edge ab ; furthermore, the right side of p_2p_1 and q_1q_2 are in the interior of P .*

Proof. By (V3), the interior points of a turn gadget can see only access points, i.e., $\{p_1, q_1, a, b\}$. By Lemma 14, a and b are each adjacent to some interior point and, by Lemma 16, ab is an edge of P . It follows that P contains the required chain. See Fig. 17 (right) for an example. Also by Lemma 16, the right side of every edge in the directed path from p_1 to q_1 is in the interior of P . If the right side of p_2p_1 (q_1q_2) is in the exterior of P , then p_2p_1 (q_1q_2) is an external diagonal, a contradiction. \square

Construction of the Reduction. Let \hat{G} be a PSLG constructed in the proof of Theorem 13 (Sect. 4.1). Recall that \hat{G} consists of line segments and opaque faces which are the interior of circuits. There are exactly two reflex vertices in the outer opaque face (the one that intersects the boundary of the convex hull of \hat{G}), and all other opaque faces are convex. We will replace the convex angles of all opaque faces as described in the construction of the turn gadget with two differences at a few special vertices, that we describe next.

- For vertices on the boundary of a chamber, the previous description is unchanged.
- For the four vertices in the convex hull (corresponding to vertices a , b , and u' in Fig. 18 (d)), there is no vertex w_2 , i.e., no point that initially sees the left side of ab . This only decreases the number of constraints to construct the turn gadget.
- For the two reflex vertices of the outer opaque polygon, we adjust the turn gadget as follows. Refer to Fig. 18 (a). Let w_1 and r_1 be the convex and reflex vertices of \hat{G} corresponding to

the vertex adjacent to b in the interior of the convex hull of G_1 . Shrink one of the edges incident to r_1 by $\delta > 0$, creating a new vertex r'_1 . We proceed analogously for the reflex vertex r_2 corresponding to the vertex adjacent to a in G_1 . The choice of which edge to shrink is made so that vertices r'_1 and r'_2 cannot see each other. We follow by replacing w_1 by a turn gadget as described before with e defined by the point to the left of ab visible from r_1 (r_2) by making w_1 transparent. We require an additional constraint making δ small enough so that no new visibility is created (in particular, r'_1 (r'_2) does not see the segment w_4w_5 in the center of adjacent chambers).




Figure 18: Reduction from the instance shown in Fig. 11 to disjoint segments. For clarity, only segment ab from each turn gadget is shown.

Lemma 18. *The set of disjoint line segments S constructed in the above reduction admits a circumscribing polygon if and only if the original instance of HP3CPG-SE is positive.*

Proof. Assume the instance of HP3CPG-SE is positive. As described in the proof of Theorem 13, we can construct a circumscribing polygon for the set of cycles created in the instance produced in Sect. 4.1. Given a polygonization of these cycles, Figs. 17(c) and 18(b-d) show how to transform this into a circumscribing polygon of S . For every turn gadget, replace each w_1 by its corresponding path from p_1 to q_1 shown in Fig. 17(c). Figures 18(b-c) show how to replace paths such as in Figs. 13(c-d) that traverse the two reflex vertices.

Conversely, assume that S admits a circumscribing polygon P , with edge set $E(P)$. By construction, S is composed of the boundaries of convex opaque polygons whose corners have been replaced with turn gadgets, one opaque polygon with two reflex angles modified as shown in Fig. 18(a), and a segment w_4w_5 in the center each chamber (as seen in Fig. 16(a)). The segments w_4w_5 in the chambers are unchanged from the instance described in Sect. 4.1, and due to Lemma 17, in any circumscribing polygon of S the turn gadgets must behave just as the corners of polygons in the reduction used to prove Theorem 13, i.e., we can consider the points in the gadget as a single point. Thus, it remains only to show that the two reflex angle gadgets also behave as the corners of a polygon in circumscribing polygon of S .

We first claim that if $r_1r_2 \notin S$, then $r_1r_2 \notin E(P)$. Refer to Fig. 18. Suppose, for contradiction, that $r_1r_2 \notin S$ and $r_1r_2 \in E(P)$. Note that r_1 and r_2 are the only reflex vertices of an opaque polygon in the graph \widehat{G} . Since $r_1r_2 \notin S$, the edge r_1r_2 partitions this opaque face into two faces, each of which has at least one additional (convex) vertex. The convex vertices of \widehat{G} have been

replaced by turn gadgets in the instance S . By Lemma 17, all vertices in turn gadgets are in subpaths of P whose endpoints are visible to neither r_1 nor r_2 . By Lemma 16, both sides of $r_1 r_2$ are in the interior of P , which is a contradiction. This proves the claim.

By construction, r_1 and r_2 see each other, but r'_1 and r'_2 do not. Further, by Lemma 17, points r_1 , r'_1 , r_2 , and r'_2 cannot be adjacent to points of a turn gadget other than those labeled p_1 and q_1 . We now show that if S admits a circumscribing polygon, then we may assume that it contains the edges $r_1 r'_1$ and $r_2 r'_2$. Suppose, for contradiction, that there is a circumscribing polygon P , but $r_1 r'_1 \notin E(P)$. Let s_1 and s'_1 , respectively, be the neighbors of r_1 and r'_1 in S and consider the adjacent turn gadget as shown in Fig. 18 (a). Since $r_1 r'_1 \notin E(P)$, vertex r'_1 (r_1) must be adjacent to at least one vertex in $\{p_1, p_2, q_1, q_2, s_1\}$ ($\{p_1, p_2, q_1, q_2, s'_1\}$). We distinguish cases:

- If $r'_1 q_2 \in E(P)$ or $r'_1 s_1 \in E(P)$, then this edge blocks visibility for r_1 to all possible neighbors (other than s_1), and r_1 cannot have two neighbors in P .
- If $r'_1 q_1 \in E(P)$, then this edge would block all but one visibility edge for r_1 , and would imply $r_1 q_2 \in E(P)$. Then the only possible second neighbor of r_1 in P is s_1 . In this case, we can modify P by expanding $r'_1 q_1$ to the path (r'_1, r_1, q_1) and shortcutting the path (s_1, r_1, q_2) to $s_1 q_2$.
- If $r'_1 p_1 \in E(P)$, then the right side of $p_1 r'_1$ is in the interior (exterior) of P because of Lemma 17 applied for the turn gadget between at p_1 (s'_1), which is a contradiction.

The symmetric claims also hold for neighbors of r_1 . By exclusion, the only remaining possibility is that P contains the paths (s'_1, r'_1, p_2) and (s_1, r_1, q_2) . However, since P must visit p_1 , this implies $p_2 p_1 \in E(P)$. Then the path $(s'_1, r'_1 p_2, p_1)$ in P contradicts Lemma 17 for vertices s'_1 and p_1 , as the interior of P should be on the same side of every edge of P . We have shown that if P is a circumscribing polygon for S , then $r'_1 r_1 \in E(P)$. By symmetry the same holds for r_2 and r'_2 , completing the proof that the reflex angle gadgets behave as the corners of a polygon in any circumscribing polygon. \square

As already discussed in the descriptions of the gadgets, the reduction can be computed in polynomial time and by Lemma 18 we have the following result.

Theorem 19. *It is NP-complete to decide whether a set of disjoint line segments admits a circumscribing polygon.*

5 Simple Polygonizations of Disjoint Segments

Rappaport [18] proved that it is NP-hard to decide whether a given PSLG $G = (V, E)$ admits a polygonization: The reduction from HAMILTONIAN PATH IN PLANAR CUBIC GRAPHS (HPPCG) produces an instance in which G is a union of disjoint paths, each edge in E is horizontal or vertical, and the vertices in V have integer coordinates, bounded by a polynomial in $n = |V|$. Rappaport [18] raised the question whether the problem remains NP-hard when G is a perfect matching (i.e., a set of disjoint line segments in the plane). In this section, we settle this problem (Theorem 20). Specifically, we modify Rappaport's reduction, and describe the *connection gadget* made of disjoint line segments that simulates a pair of adjacent line segments.

Description of the Connection Gadget. We later provide the details on how to construct the gadget. Refer to Fig. 19. Given a PSLG $G = (V, E)$, with a vertex $p_2 \in V$ of degree 2, incident to $p_1p_2, p_2p_3 \in E$, delete the edge p_1p_2 , and insert six new edges $p_1p'_2, p_4p_5, p_6p_7, p_8p_9, p_{10}p_{11}, p_{12}p_{13}$, and eleven new vertices p'_2 and $p_i, i = 3, \dots, 13$. Denote by $G' = (V', E')$ the resulting new PSLG. We choose the position of the new vertices close to p_2 so that:

- (i) the two small segments p_6p_7 and $p_{10}p_{11}$ are visible only from points p_4, p_5, p_8 , and p_9, p_{12}, p_{13} respectively;
- (ii) the union of the visibility regions of p_4, p_5, p_{12} , and p_{13} each contain only some of $p_2, p'_2, p_4, \dots, p_{13}$, but no other vertices.




Figure 19: (a) Two adjacent axis-parallel line segments p_1p_2 and p_2p_3 . (b) Connection gadget that simulates (a) using seven disjoint line segments. The polygonal path shown in black and blue line segments is $[p_1, p'_2, p_4, p_5, p_6, p_7, p_8, p_9, p_{10}, p_{11}, p_{12}, p_{13}, p_2, p_3]$. (c) The union of the visibility regions of the solid black points p_4 and p_{13} .

Construction of the Connection Gadget. We now describe the placement of the points of the connection gadget that ensures ((i))–((ii)). Let δ be the feature ratio of G , defined as the minimum distance between a vertex and a nonincident edge of G . Let α be the smallest angle between two adjacent edges in G , and set $\varepsilon = \delta \sin^{-1}(\alpha/(4n))$, where n is the number of vertices in G . We describe now the position of auxiliary lines. Refer to Fig. 20. The dashed line ℓ_1 connects the pair of points in $p_1p_2p_3$ that are ε apart from p_2 . This line will contain vertices p_9 and p_8 . Let W be the cone of angle α from p_2 placed so that its bisector coincides with the bisector of $\angle p_1p_2p_3$. Divide W into $2n$ cones with equal angle. By definition of ε , if we place a disk of radius ε at each vertex of G , each disk can intersect only two $\alpha/(2n)$ cones. By the pigeonhole principle, one of these cones w will contain no such disks. The visibility of p_2 inside w is either empty or the interior of an edge containing points that are more than ε apart from the endpoints. Let ℓ_2 and ℓ_3 be the supporting lines of the rays defining w , and let ℓ_4 be the bisector of w . Let a (resp., b) be the intersection point between ℓ_1 and ℓ_4 (resp., ℓ_2), and c be the intersection point of the circle of radius δ centered at p_2 and ℓ_3 that is in the boundary of w . We place vertex p'_2 at the intersection of ac and p_1p_2 . Let t be the intersection point of ℓ_2 and ac , and let d (resp., e) be the point in p_1p_2 (resp., p_2p_3) that is 2ε away from p_2 . Place p_8 (resp., p_9) at the intersection between ℓ_1 and dt (resp., ed'_2). We will place points p_4, \dots, p_7 and p_{10}, \dots, p_{13} in the triangle $\Delta(abt)$. This guarantees that ((ii)) is satisfied.

Let f be the midpoint of ab , and $p'_4p'_5$ (resp., $p'_{12}p'_{13}$) be the segment defined by the intersection between $\Delta(abt)$ and the line through f parallel to p_2p_3 (resp., p_1p_2). Define the segment p_4p_5 (resp.,




Figure 20: The figure is not drawn scale. (a) The construction of the dashed line ℓ_1 and the cone w . (b) The placement of p'_2 , and the triangle $\Delta(abt)$. (c) The placement of p_8 and p_9 .

$p_{12}p_{13}$) by scaling $p'_4p'_5$ (resp., $p'_{12}p'_{13}$) by $1/2$ anchored at its midpoint. Let $p'_6p'_7$ (resp., $p'_{10}p'_{11}$) be the segment contained in the perpendicular bisector of p_4p_5 (resp., $p_{12}p_{13}$) that is only visible to p_4 , p_5 , and p_8 (resp., p_9 , p_{12} , and p_{13}). Define the segment p_6p_7 (resp., $p_{10}p_{11}$) by scaling $p'_6p'_7$ (resp., $p'_{10}p'_{11}$) by $1/2$ anchored at its midpoint. The construction guarantees that (i) is satisfied. This concludes the construction of the connection gadget.

Theorem 20. *It is NP-complete to decide whether a set S of disjoint line segments admits a simple polygonization, even if S contains only segments with four distinct slopes.*

Proof. Membership in NP is proven in [18]. We reduce NP-hardness from finding polygonizations for a disjoint union of paths. Let $G = (V, E)$ be a PSLG produced by the reduction in [18]. Let $n = |V|$. We modify G by simultaneously replacing every vertex of degree 2 by a connection gadget (described above), and show that the resulting plane straight-line matching M admits a polygonization if and only if G does. Since each gadget is constructed independently, all coordinates can be described by bounded-degree rational functions as each is obtained by a constant number of intersections between lines and circles determined by G . Since G contains only axis-parallel edges, the construction produces a plane straight-line matching M , in which all edges have up to four distinct slopes. The reduction clearly runs in time polynomial in n .

We now show that M admits a polygonization if and only if G does. Note that the connection gadget places edges in the convex corner of a degree-2 vertex in G , and it does not block or create visibility between two leafs of G . By construction, if p is a leaf in G , then the set of other leaves visible from p remains same in M . Since G is max-degree-2, it remains to prove that, for every connection gadget, a polygonization of M must contain a chain of length 11 from p'_2 to p_2 that uses only edges of the connection gadget.

By (i) of the connection gadget, if a simple polygonization P of M exists, P must connect p_8 with p_6 or p_7 , and p_6 or p_7 to p_4 or p_5 ; otherwise, P would contain a cycle of length 4 and P would be disconnected. The same argument applies to vertices p_9, \dots, p_{13} . Figure 19 (c) shows the forced edges in a polygonization in red. By (ii), p'_2 must be adjacent to p_4 or p_5 , and p_2 must be adjacent to p_{12} or p_{13} , or else either P would contain a cycle of length 10 and P would be disconnected, or P would not be simple. Hence, each gadget behaves exactly like a degree-2 vertex of G , and M admits a polygonization if and only if G does. \square

6 Conclusions

Our results raise interesting open problems, among others, about circumscribing polygons in the plane (Sect. 6.1), and about higher dimensional generalizations (Sect. 6.2).

6.1 Geometric Matching with Few Slopes

As noted above, Grünbaum [10] constructed an arrangement of six disjoint segments in \mathbb{R}^2 that does not admit a circumscribing polygon; see Fig. 14. If all segments have the same slope (but they are not all collinear), then there always exists a circumscribing polygon. We conjecture that disjoint segments with two distinct slopes still admit a circumscribing polygon. Here we present negative instances with three slopes.

Proposition 21. *For every $n \geq 9$, there is a set of n disjoint segments of three different slopes that do not admit a circumscribing polygon.*

Proof. Consider the set $S = \{s_1, \dots, s_9\}$ of segments in Fig. 21, where $s_i = a_i b_i$, for $i = 1, \dots, 9$. Suppose, for the sake of contradiction, that P is a circumscribing polygon for S . Rappaport, Imai, and Toussaint [19, Lemma 2.1] proved that in every Hamiltonian simple polygon, the vertices of $\text{conv}(S)$ appear in the same counterclockwise order in P and $\text{conv}(S)$. Hence vertices a_1, a_2, b_2 , and b_1 appear in this ccw order in P . Denote by A and B , respectively, the polygonal path from a_1 to a_2 , and from b_2 to b_1 in ccw order along P . Since $a_1 b_1$ and $a_2 b_2$ are edges of $\text{conv}(S)$, they are also edges of P .




Figure 21: A set of nine disjoint line segments of slopes 0, 1, and ∞ , that do not admit a circumscribing polygon.

Consequently, the endpoints of s_3, \dots, s_7 are in A or B . By symmetry, we may assume that at least one endpoint of s_8 is in A . Since $a_1 a_2$ is also an edge of $\text{conv}(S)$, it does not cross any edge of A , and so it can complete the path A into a simple polygon. By [19, Lemma 2.1], the vertices of $\text{conv}(A)$ appear in the same order in both $\text{conv}(A)$ and A .

Note first that A contains some vertex to the right of s_8 ; otherwise, either A crosses segment $a_4 b_4$, or both a_4 and b_4 are vertices of P and so $a_4 b_4$ is an external diagonal of P . This means that A contains some vertices of segments s_5, s_6, s_9 (as b_1 and b_2 are in B). Note that all possible edges between the vertices to the left of $a_5 b_5$ and the endpoints of s_6 and s_9 cross $a_5 b_5$. If A contains any vertex to the right of $a_5 b_5$, then A would contain both a_5 and b_5 , and $a_5 b_5$ would be an external diagonal of P . Consequently A contains a_5 or b_5 , but no vertices to the right of $a_5 b_5$. If A contains b_5 without a_5 , then either A crosses $a_4 b_4$ or $a_4 b_4$ is an external diagonal of P . We conclude that A contains a_5 .

It follows that B contains b_1, b_2 , both endpoints of s_6 and s_9 , and possibly b_5 . Points b_1, b_2, a_6, b_6 , and an endpoint of s_9 are vertices of $\text{conv}(B)$. Since they appear in the same ccw order in $\text{conv}(B)$ and B , segment a_6b_6 is an external diagonal of P . This contradicts our initial assumption and completes the proof. \square

6.2 Higher Dimensions

Generalizations to higher dimensions are also of interest. For a set V of points in \mathbb{R}^3 , a *polyhedralization* is a polyhedron homotopic to a sphere whose vertex set is V .² It is known that every set of $n \geq 4$ points in general position admits a polyhedralization [1], and even a polyhedralization of bounded vertex degree [3].

For a set S of disjoint line segments in \mathbb{R}^3 , we define a *polyhedralization* as a polyhedron homotopic to a ball whose vertices are the segment endpoints, and every segment in S is either an edge or an (external or internal) diagonal. A polyhedralization *circumscribes* S if every segment in S is an edge or an internal diagonal. It is not difficult to see that an arrangement of disjoint segments in general position in \mathbb{R}^3 need not admit a polyhedralization.

Proposition 22. *For every $n \geq 4$, there is a set of n disjoint segments in \mathbb{R}^3 that do not admit a polyhedralization.*




Figure 22: Left: A set of four disjoint segments in \mathbb{R}^3 that do not admit a polyhedralization; perspective view (a) and view from above (b).

Proof. Let $s_0 = [(0, 0, -1), (0, 0, 1)]$ be a vertical segment along the z -axis. We construct s_1, \dots, s_{n-1} by taking line segments along the supporting lines of a regular $(n-1)$ -gon centered at the origin in the xy -plane, and then perturbing them to be in general position; see Fig. 22.

Suppose that these n segments admit a polyhedralization P . Then s_0 is the edge of some face. However, the triangle spanned by the two endpoints of s_0 and an endpoint of any segment s_i , $i \in \{1, \dots, n-1\}$, stabs segment s_{i-1} or s_{i+1} (where the subindex is taken modulo $n-1$). This contradicts our assumption that every segment is an edge of P . \square

We suspect that it is NP-hard to decide whether a set of n line segments in \mathbb{R}^3 admits a polyhedralization (or even a circumscribing polyhedralization). However, our proof techniques do not seem to extend to higher dimensions.

²We thank Joe Mitchell for introducing us to the high dimensional variations of this problem.

6.3 Open Problems

We conclude with a collection of open problems.

1. In Sect. 5 we established NP-hardness for the polygonization problem, even when the input consists of disjoint line segments with four distinct slopes. Is it NP-hard to decide whether n disjoint *axis-parallel* segments in the plane admit a polygonization? Is it NP-hard for segments of three possible directions?
2. Does every arrangement of disjoint axis-parallel segments in \mathbb{R}^2 , not all in a line, admit a circumscribing polygon?
3. Does every arrangement of disjoint line segments in \mathbb{R}^3 , not all in a plane, admit a circumscribing polyhedron?
4. We can decide in $O(n \log n)$ time whether n disjoint segments are extensible to disjoint rays (cf. Sect. 3). Can we decide efficiently whether they admit escape routes?
5. Let $f(n)$ be the maximum integer such that every set of n disjoint segments contains $f(n)$ segments that admit a circumscribing polygon. In Sect. 2, we prove a lower bound of $f(n) = \Omega(\sqrt{n})$. Is it possible that $f(n) = \Omega(n)$? Is there a nontrivial upper bound?
6. Let $g(n)$ be the maximum integer such that every set of n disjoint segments contains $g(n)$ segments that are extensible to disjoint rays. Theorem 11 implies $g(n) \leq f(n)$. Pach and Rivera-Campo [17] proved that $g(n) = \Omega(n^{1/3})$, and Lemma 12 gives $g(n) = O(\sqrt{n})$. What is the asymptotic growth rate of $g(n)$?
7. Let $h(n)$ be the maximum integer such that every set of n disjoint segments contains $h(n)$ segments that admit an escape route. Theorem 11 implies $g(n) \leq h(n) \leq f(n)$. We have $h(n) = \Omega(n^{1/3})$ and $h(n) = O(\sqrt{n})$. What is the asymptotic growth rate of $h(n)$?

References

- [1] Agarwal, P.K., Hurtado, F., Toussaint, G.T., Trias, J.: On polyhedra induced by point sets in space. *Discrete Appl. Math.* **156**(1), 42–54 (2008)
- [2] Akitaya, H.A., Korman, M., Rudoy, M., Souvaine, D.L., Tóth, C.D.: Circumscribing polygons and polygonizations for disjoint line segments. In: 35th International Symposium on Computational Geometry (Portland 2019). *Leibniz Int. Proc. Inform.*, vol. 129, # 9. Leibniz-Zent. Inform., Wadern (2019)
- [3] Barequet, G., Benbernou, N., Charlton, D., Demaine, E.D., Demaine, M.L., Ishaque, M., Lubiw, A., Schulz, A., Souvaine, D.L., Toussaint, G.T., Winslow, A.: Bounded-degree polyhedronization of point sets. *Comput. Geom.* **46**(2), 148–153 (2013)
- [4] Cardinal, J., Hoffmann, M., Kusters, V., Tóth, C.D., Wettstein, M.: Arc diagrams, flip distances, and Hamiltonian triangulations. *Comput. Geom.* **68**, 206–225 (2018)

- [5] Chambers, E.W., Eppstein, D., Goodrich, M.T., Löffler, M.: Drawing graphs in the plane with a prescribed outer face and polynomial area. *J. Graph Algorithms Appl.* **16**(2), 243–259 (2012)
- [6] Chiba, N., Nishizeki, T.: The Hamiltonian cycle problem is linear-time solvable for 4-connected planar graphs. *J. Algorithms* **10**(2), 187–211 (1989)
- [7] Di Giacomo, E., Liotta, G.: The Hamiltonian augmentation problem and its applications to graph drawing. In: 4th International Workshop Algorithms and Computation (WALCOM) (Dhaka 2010). *Lecture Notes in Comput. Sci.*, vol. 5942, pp. 35–46. Springer, Berlin (2010)
- [8] García, A., Noy, M., Tejel, J.: Lower bounds on the number of crossing-free subgraphs of K_N . *Comput. Geom.* **16**(4), 211–221 (2000)
- [9] Garey, M.R., Johnson, D.S., Tarjan, R.E.: The planar Hamiltonian circuit problem is NP-complete. *SIAM J. Comput.* **5**(4), 704–714 (1976)
- [10] Grünbaum, B.: Hamiltonian polygons and polyhedra. *Geombinatorics* **3**(3), 83–89 (1994)
- [11] Hoffmann, M., Tóth, C.D.: Segment endpoint visibility graphs are Hamiltonian. *Comput. Geom.* **26**(1), 47–68 (2003)
- [12] Hurtado, F., Tóth, C.D.: Plane geometric graph augmentation: a generic perspective. In: *Thirty Essays on Geometric Graph Theory*, pp. 327–354. Springer, New York (2013)
- [13] Ishaque, M., Souvaine, D.L., Tóth, C.D.: Disjoint compatible geometric matchings. *Discrete Comput. Geom.* **49**(1), 89–131 (2013)
- [14] Mirzaian, A.: Hamiltonian triangulations and circumscribing polygons of disjoint line segments. *Comput. Geom.* **2**(1), 15–30 (1992)
- [15] O’Rourke, J., Rippel, J.: Two segment classes with Hamiltonian visibility graphs. *Comput. Geom.* **4**(4), 209–218 (1994)
- [16] Ozeki, K., Van Cleemput, N., Zamfirescu, C.T.: Hamiltonian properties of polyhedra with few 3-cuts—a survey. *Discrete Math.* **341**(9), 2646–2660 (2018)
- [17] Pach, J., Rivera-Campo, E.: On circumscribing polygons for line segments. *Comput. Geom.* **10**(2), 121–124 (1998)
- [18] Rappaport, D.: Computing simple circuits from a set of line segments is NP-complete. *SIAM J. Comput.* **18**(6), 1128–1139 (1989)
- [19] Rappaport, D., Imai, H., Toussaint, G.T.: Computing simple circuits from a set of line segments. *Discrete Comput. Geom.* **5**(3), 289–304 (1990)
- [20] Sharir, M., Sheffer, A., Welzl, E.: Counting plane graphs: perfect matchings, spanning cycles, and Kasteleyn’s technique. *J. Combin. Theory Ser. A* **120**(4), 777–794 (2013)
- [21] Tutte, W.T.: Convex representations of graphs. *Proc. London Math. Soc.* **10**, 304–320 (1960)

- [22] Urabe, M., Watanabe, M.: On a counterexample to a conjecture of Mirzaian. *Comput. Geom.* **2**(1), 51–53 (1992)
- [23] Whitney, H.: A theorem on graphs. *Ann. of Math.* **32**(2), 378–390 (1931)

Accepted manuscript