On Inducing Polygons and Related Problems ⋆

Eyal Ackerman[1], Rom Pinchasi[2], Ludmila Scharf[1], and Marc Scherfenberg[1]

[1] Institute of Computer Science, Freie Universität Berlin, Takustr. 9, 14195 Berlin, Germany. {eyal,scharf,scherfen}@mi.fu-berlin.de.
[2] Mathematics Department, Technion—Israel Institute of Technology, Haifa 32000, Israel. room@math.technion.ac.il.

Abstract. Bose et al. [1] asked whether for every simple arrangement $\mathcal{A}$ of n lines in the plane there exists a simple n-gon P that *induces* $\mathcal{A}$ by extending every edge of P into a line. We prove that such a polygon always exists and can be found in $O(n \log n)$ time. In fact, we show that every finite family of curves $\mathcal{C}$ such that every two curves intersect at least once and finitely many times and no three curves intersect at a single point possesses the following Hamiltonian-type property: the union of the curves in $\mathcal{C}$ contains a simple cycle that visits every curve in $\mathcal{C}$ exactly once.

1 Introduction

Arrangements of lines in the plane are among the most studied structures in Combinatorial and Computational Geometry (see, e.g., [4, 5]). Every set of straight-line segments S naturally *induces* an arrangement of lines, simply by extending every segment in S into a line. Bose et al. [1] asked the following natural question.

Problem 1. Does every *simple* arrangement $\mathcal{A}$ of n lines contain a simple n-gon that induces $\mathcal{A}$?

An arrangement of lines is *simple* if every pair of lines intersects, and no three lines intersect at a single point. A polygon (resp., curve) is *simple* if it is non-self-intersecting. Fig. 1(a) shows a simple arrangement of six lines and a simple hexagon that induces this arrangement.

Problem 1 remained open until now, though a few partial results were obtained. In [1] it was shown that a simple arrangement $\mathcal{A}$ of n lines contains a subarrangement of $m \geq \sqrt{n-1}+1$ lines that has an inducing simple m-gon, and that $\mathcal{A}$ always has an inducing simple n-path (a polygonal chain consisting of n line segments), which can be constructed in $O(n^2)$ time. Recently, the third and fourth authors [8] showed that an inducing n-path can be constructed in $O(n \log n)$ time, and that there always exists an inducing simple $O(n)$-gon, which can be found in $O(n^2)$ time.

Our main result is an affirmative answer to Problem 1.

⋆ Research by Eyal Ackerman was supported by a fellowship from the Alexander von Humboldt Foundation. Research by Rom Pinchasi was supported by the Israeli Science Foundation (grant No. 938/06).

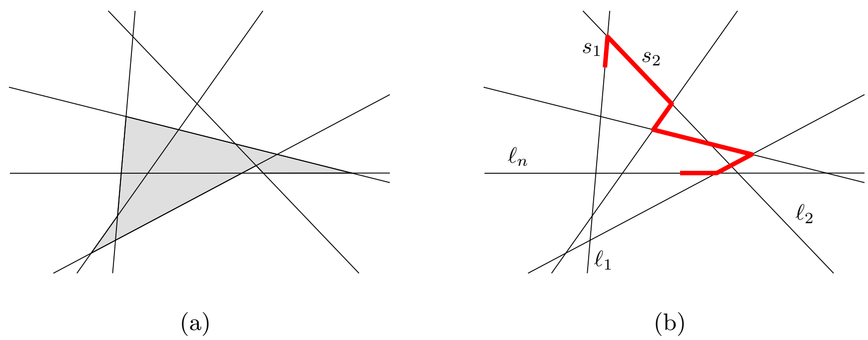

Fig. 1. An inducing simple n-gon and n-path.

Theorem 1. *For every simple arrangement $\mathcal{A}$ of $n > 2$ lines in the plane there is a simple n-gon that induces $\mathcal{A}$. Given the set of n lines that form $\mathcal{A}$, such a polygon can be constructed in $O(n \log n)$ time.*

We give two different constructive proofs for the existence of an inducing simple n-gon. The first proof is short and elegant and yields a non-optimal but polynomial-time algorithm for finding such a polygon. The second proof yields an $O(n \log n)$-time algorithm. It is based on a simple idea, however, it involves several case distinctions and is, thus, quite technical.

During our quest for a solution to Problem 1, we proved the following interesting fact.

Theorem 2. *For every simple arrangement $\mathcal{A}$ of n non-vertical lines in the plane there is an x-monotone n-path that induces $\mathcal{A}$.*

Note that the first part of Theorem 1 can also be phrased as follows: Every arrangement of lines contains a simple cycle (i.e., a closed curve) that visits every line exactly once. To be more precise, we say that a curve x *visits* another curve y if their intersection contains a point in which they neither cross nor touch. A simple curve visits y *exactly once* if it visits y and their intersection is connected. The first part of Theorem 1 is then equivalent to saying that every simple line arrangement contains a simple (polygonal) cycle that visits every line exactly once. We also have the following generalization of Theorem 1.

Theorem 3. *Let $\mathcal{C}$ be a finite family of $n > 2$ simple curves in $\mathbb{R}^3$, such that every pair of curves in $\mathcal{C}$ intersects at least once and at most finitely many times, and no three curves intersect at the same point. Then $\bigcup_{C \in \mathcal{C}} C$ contains a simple cycle that visits every curve exactly once.*

The rest of this paper is organized as follows. A first proof for the existence of an inducing simple n-gon is given in Section 2. This proof is then extended in Section 3 to a proof of Theorem 3. In Section 4 we describe a different and more efficient way of finding an inducing simple n-gon. Due to space limitations, we only sketch the idea of the proof and omit most of the details, which can be found in the full version of this paper. Theorem 2 is proved in Section 5, while Section 6 contains some concluding remarks.

2 First proof of the existence of an inducing simple n-gon

Let $\mathcal{A}$ be a simple arrangement of n lines in the plane. We begin by constructing a simple path that visits every line in $\mathcal{A}$ exactly once. This is done in a way similar to the construction of an inducing path in [1]. Consider an arbitrary intersection point of two lines, denote these lines by ℓ_1 and ℓ_2. Walk a short distance on ℓ_1 toward its intersection point with ℓ_2. Remove ℓ_1, and walk on ℓ_2 in a direction that contains at least one intersection point, until reaching the first intersection point. Let ℓ_3 be the other line that determines this intersection point. Remove ℓ_2 and repeat the same process for ℓ_3 and so on and so forth, until reaching a line that has no additional intersection points. Finally, walk a short distance on this line in some direction. See Fig. 1(b) for an example.

Since every pair of lines intersects, no line is missed and this process results in a path that induces every line in $\mathcal{A}$. Denote this path by Q. The lines in $\mathcal{A}$ are denoted by $\ell_1, \ell_2, \ldots, \ell_n$ according to the order they are visited by Q. Denote by s_j the segment of ℓ_j on Q. Assume to the contrary that Q is self-intersecting. Then there are two intersecting segments, s_i and s_j, such that $i \leq j-2$. But this is a contradiction to the definition of ℓ_{i+1} as the first line, different from $\ell_1, \ldots, \ell_i$, that we encounter while walking along ℓ_i.

Observe that Q lies in one of the two half-planes determined by ℓ_n. Indeed, otherwise ℓ_n would have crossed Q, contradicting the definition of ℓ_n as the last line in $\mathcal{A}$ we encounter while creating the path Q. We assume without loss of generality that ℓ_n coincides with the x-axis and that Q lies in the half-plane above ℓ_n. For convenience, denote the line ℓ_n by ℓ.

We call a simple inducing n-path of $\mathcal{A}$ *rooted above* ℓ if it lies in the closed half-plane above ℓ and ℓ includes an extreme segment of the path. As we have just seen, there is at least one simple inducing n-path rooted above ℓ, namely Q. For such a path W we denote by $q_1(W), \ldots, q_{n-1}(W)$ the $n-1$ internal vertices of the path starting from $q_1(W)$ on ℓ. We denote by $\ell(W)$ the line through $q_{n-1}(W)$ that includes the other extreme segment of W. Denote by $\ell^-(W)$ the half-line of $\ell(W)$ that consists of all points with y-coordinates smaller than the y-coordinate of $q_{n-1}(W)$. We denote by $q_n(W)$ the topmost (also first) intersection point of $\ell^-(W)$ with $\ell \cup [q_1(W)q_2(W)] \cup \ldots \cup [q_{n-2}(W)q_{n-1}(W)]$. (Here, $[ab]$ denotes the line segment connecting point a to point b.)

Let $z_1, \ldots, z_m$ denote all the intersection points in $\mathcal{A}$ indexed in any way such that $i < j$ if the y-coordinate of z_i is smaller than the y-coordinate of z_j. We then define for every j, $Y(z_j) = j$.

For a simple inducing n-path W rooted above ℓ, let $Y(W) = \sum_{i=1}^{n} Y(q_i(W))$. Consider the simple inducing n-path W rooted above ℓ such that $Y(W)$ is minimum. If $q_n(W)$ lies on ℓ, then observe that the vertices $q_1(W), \ldots, q_n(W)$ define

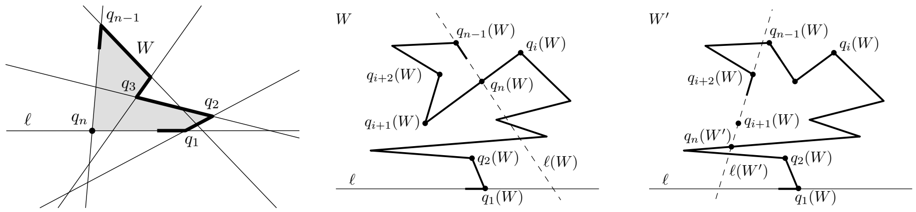

(a) q_n lies on ℓ (b) q_n lies on W (c) W'

Fig. 2. The paths W and W'

a simple inducing closed n-path of $\mathcal{A}$ (see Fig. 2(a)). Assume therefore that $q_n(W)$ is the intersection point of $\ell(W)$ with the segment $[q_i(W)q_{i+1}(W)]$ for some $1 \leq i \leq n-2$ (see Fig. 2(b)). Then we define W' as the path whose internal vertices are

$$q_1(W), \ldots, q_i(W), q_n(W), q_{n-1}(W) \ldots, q_{i+2}(W),$$

and hence $\ell(W')$ is the line through $q_{i+1}(W)$ and $q_{i+2}(W)$. Observe that W' is a simple inducing n-path rooted above ℓ. We have $Y(W') < Y(W)$ because $q_n(W')$ has a smaller y-coordinate than the y-coordinate of $q_{i+1}(W)$ (see Fig. 2(c)). We have thus reached a contradiction to the minimality of W. $\square$

Remark. The proof of Theorem 1, presented above, yields an algorithm with running time polynomial in n. This is because $Y(W)$ is always smaller than n^3 and this gives a bound on the number of iterations going from W to W' required to find a simple inducing closed n-path for $\mathcal{A}$.

3 Proof of Theorem 3

Let $\mathcal{C}$ be a family of n simple curves in $\mathbb{R}^3$, such that every pair of curves in $\mathcal{C}$ intersects at least once and at most finitely many times, and no three of the curves meet at a point. We will show that $\bigcup_{C \in \mathcal{C}} C$ contains a simple closed path that visits every curve in $\mathcal{C}$ exactly once.

The proof is a modification of the argument in the proof of Theorem 1. We first find a simple path Q that visits every curve in $\mathcal{C}$ exactly once, exactly in the same way that was described in Section 2, applied this time to $\mathcal{C}$. Let c be a curve in $\mathcal{C}$ containing the last segment of Q thus constructed. As we observed in the case of lines, c does not meet Q at any point outside the segment of Q contained in c.

A simple (oriented) path W that visits every curve in $\mathcal{C}$ exactly once will be called *rooted in* c if c is the first curve visited by W. Clearly, Q is an example for such a path.

For a path W, as above, we denote by $q_1(W), \ldots, q_{n-1}(W)$ the $n-1$ internal vertices of the path starting from $q_1(W)$ on c. For $i = 1, \ldots, n-2$ we denote by $s_i(W)$ the segment of W whose vertices are $q_i(W)$ and $q_{i+1}(W)$, these will be called the *internal segments* of W. We denote by $c(W)$ the curve in $\mathcal{C}$ that passes through $q_{n-1}(W)$ and contains the last segment of W.

Let s be a portion of a curve in $\mathcal{C}$. We define $|s|$ as the number of intersection points of pairs of curves in $\mathcal{C}$ that lie on s. Finally, we define

$$Y(W) = f(|s_1(W)|, \ldots, |s_{n-2}(W)|),$$

where $f(x_1, \ldots, x_{n-2})$ is a strictly monotone increasing function of the lexicographic order of $(x_1, \ldots, x_{n-2})$.[3]

Consider the simple path W that is rooted in c and visits every curve in $\mathcal{C}$ exactly once, such that $Y(W)$ is minimum. Let p be an intersection point of $c(W)$ and c. Let $q_n(W)$ be the intersection point of $c \cup s_1(W) \cup \ldots \cup s_{n-2}(W)$ and the portion of $c(W)$ between $q_{n-1}(W)$ and p that is closest to $q_{n-1}(W)$ along the curve $c(W)$.

If $q_n(W)$ lies on c, then observe that the vertices $q_1(W), \ldots, q_n(W)$ define a simple closed path that visits every curve in $\mathcal{C}$ exactly once. Assume therefore that $q_n(W)$ is an intersection point of $c(W)$ with $s_i(W)$ for some $1 \leq i \leq n-2$. Let s' denote the portion of $s_i(W)$ delimited by $q_i(W)$ and $q_n(W)$. Let s'' denote the portion of $c(W)$ delimited by $q_n(W)$ and $q_{n-1}(W)$. Then we define W' as the path rooted in c whose internal segments are

$$s_1(W), \ldots, s_{i-1}(W), s', s'', s_{n-2}(W), s_{n-3}(W) \ldots, s_{i+2}(W),$$

and $c(W')$ is the curve containing the segment $s_{i+1}(W)$.

Observe that W' is a simple path rooted on c that visits every curve in $\mathcal{C}$ exactly once. It immediately follows that $Y(W') < Y(W)$, because $s_j(W') = s_j(W)$ for $j = 1, \ldots, i-1$ while it is easy to see that $|s_i(W')| < |s_i(W)|$ as $s_i(W') = s' \subset s_i(W)$ and $q_{i+1}(W)$ is an intersection point in $s_i(W) \setminus s_i(W')$. We have thus reached a contradiction to the minimality of W. □

Remarks. (1) Because Theorem 3 is stated in $\mathbb{R}^3$, geometry actually does not play any role here. We may conclude the same result for "combinatorial curves" that "intersect" finitely many times, as long as there is a total order on the set of intersection points in each curve.
(2) The result in Theorem 3 is valid also if the curves in $\mathcal{C}$ are not simple and have self-crossings. In this case we repeat the proof and ignore self-intersections of curves. Finally, when obtaining the resulting closed path we observe that self-intersections of the closed path result only from loops in the path. These loops can easily be canceled.

[3] For example, $f(11, 0, 6, \ldots) > f(6, 9, 5, \ldots) > f(6, 9, 4, \ldots)$.

4 Finding an inducing simple n-gon efficiently

Let $\mathcal{A}$ be a simple arrangement of n lines in the plane. We incrementally construct a polygon inducing $\mathcal{A}$ by starting with the boundary of a cell of $\mathcal{A}$. In every construction step the polygon is extended using a part of the boundary of the cell containing it. We assume that $n > 4$, since, combinatorially, there is only one arrangement of size three and one of size four and their inducing polygons can be easily found.

We start with a so-called critical point p, i.e., p is the first intersection point on both lines g_1 and g_2 containing it. The initial polygon P_0 is then the boundary of the only bounded face incident to p, see Fig. 3.

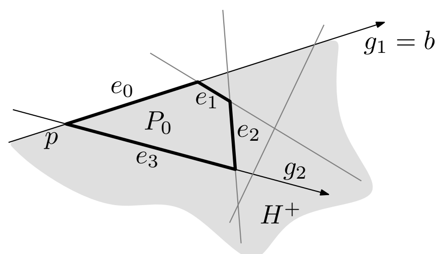

Fig. 3. Initialize P_0 to be the boundary of the bounded face of $\mathcal{A}$ incident to a critical point p.

Let P_i denote the polygon constructed in step i, and $|P_i|$ its number of edges. Denote by $\mathcal{A}_i$ the arrangement of all the lines except the ones induced by P_i. We maintain the following invariants throughout the construction of the polygons P_i.

Property 1.
1. P_i is a simple polygon;
2. P_i induces $|P_i|$ lines of the arrangement $\mathcal{A}$; and
3. P_i is contained in an unbounded face of the arrangement $\mathcal{A}_i$.

The unbounded face of $\mathcal{A}_i$ containing P_i is denoted by $C^{(i)}$ and its by $R^{(i)}$. Define the orientation of the two initial lines g_1 and g_2 in direction from p towards the remaining intersection points. Without loss of generality we can assume that all intersection points of g_2 lie in the positive half-plane of g_1, denoted by $H^+(g_1)$, as in Fig. 3.

For every construction step we maintain a so-called *base line* $b^{(i)}$. Intuitively, the base line will be the line that determines the direction in which P_i is extended. For P_0 the base line is $b^{(0)} = g_1$. The edges of P_i are labeled in the following way: the edge contained in the base line $b^{(i)}$ is the edge $e_0^{(i)}$. In counter-clockwise order we enumerate with negative indices the edges contained in the previous base lines $e_{-1}^{(i)}, \dots, e_{-m}^{(i)}$, where $e_{-m}^{(i)}$ is contained in the first base line $b^{(0)}$. These edges are referred to as *base edges*. It can be shown that the base edges form a connected concave chain in P_i. The remaining *non-base edges* are enumerated in clockwise order with positive indices $e_1^{(i)}, \ldots, e_k^{(i)}$, where $e_1^{(i)}$ is incident to $e_0^{(i)}$ and $e_k^{(i)}$ is incident to $e_{-m}^{(i)}$.

A line containing an edge $e_j^{(i)}$ is denoted by $l_j^{(i)}$ and the intersection point of two lines $l_j^{(i)}, l_m^{(i)}$ by $x_{j,m}^{(i)}$. We define the orientation of base edges in clockwise direction and the orientation of non-base edges in counter clockwise direction with respect to the polygon P_i. For each line $l_j^{(i)}$ its orientation is defined by the orientation of the edge $e_j^{(i)}$. The part of $l_j^{(i)} \setminus e_j^{(i)}$ oriented in positive (negative) direction of $l_j^{(i)}$ is called positive (negative) half-line and is denoted by $l_j^{(i)+}$ ($l_j^{(i)-}$), respectively. For simplicity we will omit the index (i) if all identifiers refer to the same step i, and will use the index in order to distinguish between different steps.

We maintain the following properties for base lines, and non-base lines, respectively.[4]

Property 2. All intersection points of a base line l_j, $j \leq 0$, with $\mathcal{A}_i$ lie in the positive half-line, i.e., $l_j^- \cap \mathcal{A}_i = \emptyset$ for $j \leq 0$. The base edges $e_0, e_{-1}, \ldots, e_{-m}$ form a concave chain in P_i, and every non-base edge is contained in the union of the positive half-planes (i.e., half-planes to the right of the oriented line) of the base lines $l_0, l_{-1}, \ldots, l_{-m}$.

Property 3. The intersection of a non-base line l_j with a non-base edge e_k is empty, for $k > j+1$.

For the line l_1 it would be helpful to have an even stronger property:

Property 4. The intersection of l_1 with P_i is exactly the edge e_1. That is, l_1 supports P_i.

The idea of the extension step is to extend the polygon P_i in direction of the base line by modifying the edges e_0 up to at most e_3 and adding a part of the boundary $R^{(i)}$ to the new polygon P_{i+1}. In every extension step we remove a chain of edges from the polygon P_i, and attach a simple polygonal chain to the open ends. Thus, if the added chain does not intersect the unchanged part of P_i, the polygon P_{i+1} is simple.

Depending on the combinatorial configuration of the lines l_1, l_2, l_3, the chain of base edges, and the boundary R, one of several extension construction steps is taken, until all lines of $\mathcal{A}$ are induced by P_j, for some j. The inducing polygon for $\mathcal{A}$ is then $P = P_j$.

The first case distinction is whether the negative half-line of l_1 intersects the boundary R. If it does, Case 1 applies.

[4] Property 3 can be violated in a special case that is considered in the full version of the paper.

Case 1 [$l_1^- \cap R \neq \emptyset$]: The edge e_1 is replaced by the part of l_1^- from $x_{1,2}$ to its intersection with R. The edge e_0 is extended until the intersection of b and R. Finally, we add the segment of R between these two intersection points, see Fig. 4(a). The base line for P_{i+1} remains unchanged $b^{(i+1)} = b^{(i)}$.

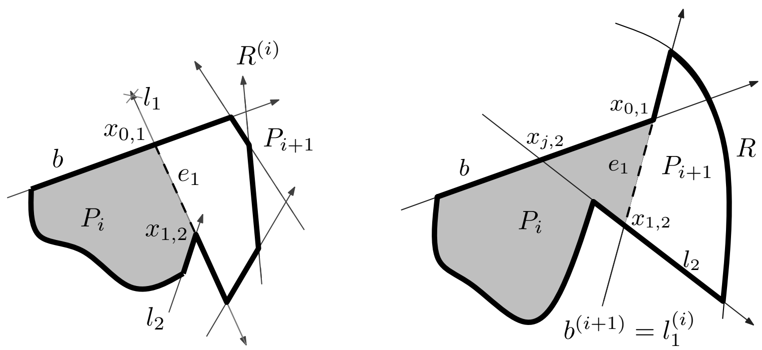

(a) Case 1 (b) Case 2

Fig. 4. Case 1 and 2: The polygon P_i is the shaded area. The identifiers refer to P_i, and the new polygon P_{i+1} is outlined by the bold black line.

The next distinction is whether l_2^+ intersects R:

Case 2 [$l_1^- \cap R = \emptyset$ and $l_2^+ \cap R \neq \emptyset$]: In this case e_1 is replaced by the part of l_1^+ from $x_{0,1}$ to its intersection with R. The edge e_2 is extended following the orientation of l_2 until the intersection of l_2^+ and R. Finally, we add the segment of R between these two intersection points, see Fig. 4(b). The new base line for the polygon P_{i+1} is now $b^{(i+1)} = l_1^{(i)}$.

It is easy to verify that all the above-mentioned properties are maintained when applying Cases 1 or 2. Due to space limitations we do not include the remaining and more complicated cases in this extended abstract, and refer the reader to the full version of the paper for those missing details.

Running Time. In the initialization step we need to find an intersection point of the arrangement that is the last point on both lines intersecting in it. Ching and Lee [3] showed that such points are a subset of the intersection points between two neighboring lines sorted by slope. Thus, the initialization can be performed in $O(n \log n)$ time by sorting the lines by slope, computing the intersection points of the neighboring lines and selecting the point with the maximum or minimum x-coordinate.

For the extension steps we consider the dual points of the lines of the arrangement, where the dual space π^* is defined as in [2]: The dual of a point

$p : (a, b)$ in the primal space is the line $p^* : f(x) = ax - b$ in π^*; the dual of a line $l : f(x) = ax + b$ in the primal space is the point $l^* : (a, -b) \in \pi^*$.

Let $\mathcal{A}^*$ denote the set of points in π^* dual to the lines of the arrangement $\mathcal{A}$. We will utilize the following property of the dual points: the points of the lower/upper convex hull of $\mathcal{A}^*$ are the duals of the lines in $\mathcal{A}$ that form the boundary of the upper/lower unbounded face of the arrangement.

For that purpose we can rotate the arrangement $\mathcal{A}$ such that the initial two lines g_1 and g_2 have the maximal and the minimal slope, the initial point $p = g_1 \cap g_2$ is a vertex of the lower unbounded face, and no line of $\mathcal{A}$ is vertical. Observe that p must be the only vertex of the lower unbounded face.

When the lines g_1 and g_2 are removed from $\mathcal{A}$ the point p is contained in the new lower unbounded face. Similarly, after every extension step the constructed polygon is contained in the lower unbounded face of the arrangement of the remaining lines.

In every extension step we need to determine the intersection points of a constant number of lines with the boundary of the lower unbounded face of the arrangement of the remaining lines and to update the boundary of the lower unbounded face after deleting some lines. Updating the boundary of the lower unbounded face corresponds to updating the upper convex hull of the dual point set. Using the dynamic convex hull data structure by Hershberger and Suri [6] updates of the upper convex hull of the point set can be performed in $O(\log n)$ time, that is $O(n \log n)$ time in total.

Intersection points of a line l with the boundary of the lower unbounded face correspond in dual space to lines through l^* that are tangent to the upper convex hull of the remaining points. These tangent lines can be found in $O(\log n)$ time.

Thus the total time complexity of the construction algorithm is $O(n \log n)$.

5 x-monotone inducing n-path: Proof of Theorem 2

In this section we show that every simple arrangement of n non-vertical lines, contains an inducing x-monotone n-path. Since the path is x-monotone, it is clearly simple. Suppose first that n is an even number. We sort the lines according to their slopes, and denote by A the set of the first $n/2$ lines in this order, and by B the rest of the lines. Initially, all the lines are unmarked. Pick the leftmost intersection point of two unmarked lines, one from A and one from B, then mark these lines. Continue to pick a total of $n/2$ points $p_1, p_2, \ldots, p_{n/2}$ this way. We will construct an x-monotone n-path through $p_1, p_2, \ldots, p_{n/2}$.

Denote the lines that intersect at p_i by $a_i \in A$ and $b_i \in B$, $i = 1, 2, \ldots, n/2$. First, pick arbitrarily one of the lines that intersect at p_1, say a_1, walk a short distance on a_1 from a point left of p_1 to p_1, then walk a short distance on b_1 rightwards. Assume that we have built an x-monotone $2i$-path that goes a short distance rightwards beyond p_i and induces the lines $a_1, \ldots, a_i$ and $b_1, \ldots, b_i$. We will show how to extend it into an x-monotone $2(i+1)$-path that goes a short distance rightwards beyond p_{i+1} and induces the lines $a_1, \ldots, a_{i+1}$ and $b_1, \ldots, b_{i+1}$.

Observation 4 *The intersection of a_{i+1} (resp., b_{i+1}) and b_i (resp., a_i) is to the right of p_i.*

Proof. Otherwise this intersection point would be picked instead of p_i.

Consider the triangle with a vertex at p_i, an edge on the vertical line through p_{i+1}, an edge e_a on a_i, and an edge e_b on b_i (see Fig. 5).

Observation 5 *e_a (resp., e_b) is crossed by a_{i+1} or b_{i+1}.*

Proof. We consider two cases based on whether p_{i+1} is inside the wedge determined by a_i and b_i. Suppose that it is. Then a_{i+1} (resp., b_{i+1}) must cross either e_a and e_b. Suppose, w.l.o.g., that they both cross e_a (otherwise, we can reflect everything with respect to the x-axis). See Fig. 5(a). Then a_{i+1} must have a

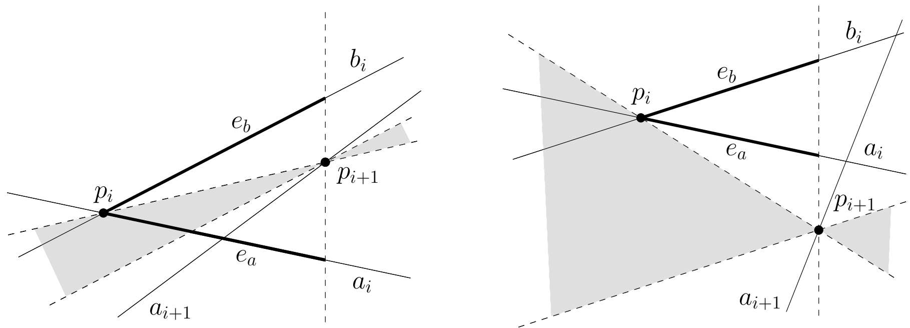

(a) p_{i+1} is inside the wedge determined by a_i and b_i.

(b) p_{i+1} is outside the wedge determined by a_i and b_i.

Fig. 5. An illustration for the proof of Observation 5. a_{i+1} cannot be in the shaded region.

larger slope than b_i, otherwise it will cross b_i to the left of p_i, contradicting Observation 4. This is of course impossible.

Suppose that p_{i+1} is outside the wedge determined by a_i and b_i. We can assume, w.l.o.g., that it is below the wedge, for otherwise we can reflect everything with respect to the x-axis. If a_{i+1} does not cross both e_a and e_b, then it must have a larger slope than b_i, or cross b_i to the left of p_i, which is impossible. See Fig. 5(b) for an illustration.

Now, suppose that the path built so far goes a short distance rightwards beyond p_i on e_a (resp., e_b). Then by Observation 5 there is a line $\ell \in \{a_{i+1}, b_{i+1}\}$ that crosses e_a (resp., e_b). Walk on e_a (resp., e_b) until the intersection point with ℓ, then walk on ℓ until p_{i+1}, and finally walk a short distance rightwards on the other line in $\{a_{i+1}, b_{i+1}\}$. The new path is an x-monotone $2(i+1)$-path that goes a short distance rightwards beyond p_{i+1} and induces the lines $a_1, \ldots, a_{i+1}$ and $b_1, \ldots, b_{i+1}$.

It remains to consider the case that n is an odd number. Let ℓ be the line with the median slope. Create a new line ℓ' that is a slightly rotated copy of ℓ such that its slope is slightly smaller than the slope of ℓ, and their intersection point is the leftmost intersection point in the arrangement $\mathcal{A} \cup \{\ell'\}$. Now continue as before, while choosing ℓ' as the first induced line. Finally, remove the segment of ℓ' from the constructed path.

Time complexity. An inducing x-monotone n-path can be found in $O(n^2)$ time as follows. First we construct the arrangement of lines. This can be done in $O(n^2)$ time [2]. Then we find the sets A and B in $O(n \log n)$ time. For every line in A we find its leftmost intersection point with a line from B. The first vertex of the path is the leftmost point among these points. The two lines defining this minimum point are removed from the arrangement while updating the minimum leftmost points for the other lines. This can be done in $O(n)$ time. The process of finding the next leftmost intersection point between a line from A and a line from B (among the remaining lines), removing the corresponding lines, and making appropriate updates is then repeated $O(n)$ times.

6 Concluding remarks

We proved in two different ways that every simple arrangement of n lines contains an inducing simple n-gon. The proof given in Section 2 actually works also for *pseudoline* arrangements. A pseudoline arrangement consists of a finite set of x-monotone curves, unbounded in both directions, such that every two curves intersect at exactly one point where they properly cross each other. It is enough to show that there is a partial order of the intersection points that lie above the pseudoline ℓ_n. Such an order can be derived from orienting every pseudoline toward its intersection point with ℓ_n. The proof then shows that there is a simple cycle that visits every pseudoline exactly once, and that such a cycle can be found in polynomial time. In fact, the proof also works for *pseudo-parabolas* (pseudo-parabolas are defined similarly to pseudolines, except that two curves cross exactly twice). Here, a partial order of the intersection points can be defined as in [7].

The second proof, given in Section 4, yields an $O(n \log n)$-time algorithm for finding an inducing simple polygon. We believe that this time complexity is the best possible, but leave it as an open question.

An inducing simple polygon need not be unique. It would be interesting to determine the maximum and minimum number of inducing simple n-gons of an arrangement of n lines. Fig. 6 shows an arrangement with exponentially many inducing simple n-gons.

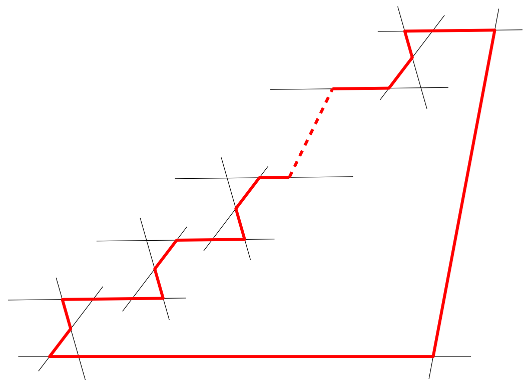

Fig. 6. n lines with exponentially many inducing simple n-gons. At every "step" of the "stairs" one can "climb" either from left or from right.

References

1. P. Bose, H. Everett, and S. Wismath, Properties of arrangement graphs, *Int. J. Comput. Geom. Appl.* **13** (2003), 447–462.
2. M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry.* Springer, Berlin, 3rd ed., 2008.
3. Y. T. Ching and D. T. Lee, Finding the diameter of a set of lines, *Pattern Recognition* **18**(3–4) (1985), 249–255.
4. S. Felsner, *Geometric Graphs and Arrangements. Some Chapters from Combinatorial Geometry*, Advanced Lectures in Mathematics, Friedr. Vieweg & Sohn, Wiesbaden, 2004.
5. B. Grünbaum, *Arrangements and spreads.* Conference Board of the Mathematical Sciences Regional Conference Series in Mathematics, No. 10. American Mathematical Society Providence, R.I., 1972.
6. J. Hershberger and S. Suri, Applications of a semi-dynamic convex hull algorithm, *BIT* **32**(2) (1992), 249–267.
7. A. Perlstein, *Problems in Combinatorial Geometry*, Ph.D. Thesis, Mathematics Department, Technion—Israel Institute of Technology, 2008.
8. L. Scharf and M. Scherfenberg, Inducing polygons of line arrangements, *Proc. 19th Int. Symp. on Algorithms and Computation* (ISAAC), *LNCS* **4059**, 508–520, 2008.