The Complexity of the Lower Envelope of Collections of Various Geometric Shapes

Carlos Alegría[1], Anna Brötzner[2], Bengt J. Nilsson[2], Christiane Schmidt[3], and Carlos Seara[4]

1 **Dipartimento di Ingegneria, Università Roma Tre, Italy**
`carlos.alegria@uniroma3.it`

2 **Department of Computer Science and Media Technology, Malmö University, Malmö, Sweden**
`{anna.brotzner,bengt.nilsson.TS}@mau.se`

3 **Department of Science and Technology, Linköping University, Campus Norrköping, Sweden**
`christiane.schmidt@liu.se`

4 **Department of Mathematics, Universidad Politécnica de Catalunya, Spain.**
`carlos.seara@upc.edu`

Abstract

We study the problem of determining the complexity of the lower envelope of a collection of n geometric objects. For collections of rays; unit length line segments; and collections of unit squares to which we apply at most two transformations from translation, rotation, and scaling, we prove a complexity of $\Theta(n)$. If all three transformations are applied to unit squares, then we show the complexity becomes $\Theta\big(n\alpha(n)\big)$, where $\alpha(n)$ is the slowly growing inverse of Ackermann's function.

1 Introduction

Consider a set of n line segments (segments for short) in the plane. It is known that the complexity of their lower envelope is at least $\Omega\big(n\alpha(n)\big)$, where $\alpha(n)$ is the extremely slowly growing inverse of Ackermann's function. The lower bound was proved by Wiernik and Sharir [4], a matching upper bound of $O\big(n\alpha(n)\big)$ was proved by Hart and Sharir [1], and an $O(n \log n)$ time and $O\big(n\alpha(n)\big)$ space algorithm to find such a lower envelope was described by Hershberger [2]. The motivation of this work is to determine under which geometric properties of a given set of n geometric objects we can ensure that their lower envelope has a tight complexity, e.g., linear or $\Theta\big(n\alpha(n)\big)$.

For many of the results, we will make extensive use of Observation 1.1.

▶ **Observation 1.1.** *Let S_1 and S_2 be two sets of n_1 and n_2 planar geometric objects whose lower envelopes have complexity $O\big(f_1(n_1)\big)$ and $O\big(f_2(n_2)\big)$, respectively.*

If any pair of objects in the set $S_1 \cup S_2$ intersect at most $O(1)$ times, then the union of the lower envelopes of S_1 and S_2 has complexity $O\big(f_1(n) + f_2(n)\big)$, where $n = n_1 + n_2$.

The observation follows by merging the two sequences of intervals generated by the corresponding two envelopes, since any two objects will appear at most a constant number of times in the lower envelope where they intersect. Thus, the complexity becomes as stated.

40th European Workshop on Computational Geometry, Ioannina, Greece, March 13–15, 2024.
This is an extended abstract of a presentation given at EuroCG'24. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

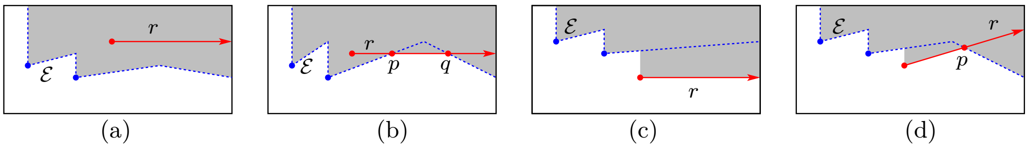

Figure 1 Illustrating the proof of Theorem 2.1.

2 Collections of Rays

▶ **Theorem 2.1.** *The lower envelope of a set of n rays has a complexity of $\Theta(n)$.*[1]

Proof. Given a ray r, let $s(r)$ and $\ell(r)$ denote the starting point and the supporting line of r, respectively. Without loss of generality, we assume in the following that no ray is vertical, since the lower envelope of any such ray is a single point.

To see the lower bound, consider n rays that move horizontally to the right, with starting points $(1, n), (2, n-1), \ldots, (n, 1)$. The lower envelope has vertices at each integer x-coordinate from 1 to n and therefore complexity $\Omega(n)$.

To see the upper bound, we argue as follows. Let R be the subset of rays that have no point to the left of their starting point and let L be the subset of remaining rays. We show next that the complexity of the lower envelope of R is $O(n)$. By symmetry (mirroring the rays in L along the line $x = 0$), the complexity of the lower envelope of L is also $O(n)$. The upper bound follows by combining these results and Observation 1.1.

For the rays in R, our proof makes use of the following observation.

▶ **Observation 2.2.** *Given two rays in R that intersect, the ray that lies above the other after their intersection point will never again be included in the lower envelope after that point.*

Consider the set of rays in R and sort them by the x-coordinate of the endpoint in order from left to right. By greedily inserting a ray r in this order into the lower envelope $\mathcal{E}$ of the previously added rays we argue that the number of intersection points in the lower envelope can increase by at most two. We have the following cases.

1. The starting point of r lies above $\mathcal{E}$ and r never intersects it, then r is never seen from below and the lower envelope does not change; see Figure 1(a),
2. The starting point of r lies above $\mathcal{E}$ and the ray intersects it at p, then p is a vertex of the new lower envelope. If r does not intersect the current lower envelope to the right of p, then r is the only object seen to the right of p. If r intersects $\mathcal{E}$ again at q, by Observation 2.2, the ray r will not be included in the lower envelope again, and furthermore, since two rays can only intersect once, between p and q, there must be at least one vertex of $\mathcal{E}$. Hence, either one vertex is added to $\mathcal{E}$ or two are added and at least one vertex must also be removed; see Figure 1(b),
3. The starting point of r lies below $\mathcal{E}$ and r never intersects $\mathcal{E}$, then a vertex in the lower envelope is introduced at the x-coordinate of the starting point of r and the ray is the only object seen to the right of this point; see Figure 1(c),
4. The starting point of r lies below $\mathcal{E}$ and r intersects it at p, then a new vertex in the lower envelope is introduced at the x-coordinate of the starting point of r and another one at p. By Observation 2.2, r can never appear again in the lower envelope; see Figure 1(d).

Thus, at most two new vertices are introduced to the lower envelope when we insert a ray. ◀

[1] This result was stated without a proof by Sharir and Agarwal in [3, page 112].

3 Collections of Line Segments with Unit Length

▶ **Theorem 3.1.** *The lower envelope of a set of n unit length segments has complexity $\Theta(n)$.*

Proof. To prove the lower bound, consider the lower envelope of the n unit length segments $[(1,0),(2,0)],[(3,0),(4,0)],\ldots,[(2n-1,0),(2n,0)]$ with $2n$ vertices, establishing the claim.

To prove the upper bound, consider a square grid covering the plane whose cells have side length 3/5. We denote by $S_{i,j}$ the grid cell at row i and column j. Let $L_{i,j}$ be the set of line segments obtained by intersecting the input set of line segments with the region bounded by $S_{i,j}$, and let $n_{i,j}$ be the number of line segments in $L_{i,j}$.

The lower envelope of $L_{i,j}$ has complexity $O(n_{i,j})$, since the input segments have length 1 and the grid cells have side length 3/5. The segments of $L_{i,j}$ have either a single endpoint or no endpoints in the interior of $S_{i,j}$. Observe that the subset of segments of $L_{i,j}$ with no endpoints behave as lines inside $S_{i,j}$ and each such segment has at most one connected piece on the lower envelope of $L_{i,j}$. Hence, the lower envelope of this subset has linear complexity [3]. On the other hand, note that the subset of segments of $L_{i,j}$ with a single endpoint behave as rays inside $S_{i,j}$. In particular, Observation 2.2 holds. Therefore, by Theorem 2.1 the lower envelope of this subset has also linear complexity. Combining these observations with Observation 1.1, we conclude that the lower envelope of $L_{i,j}$ has complexity $O(n_{i,j})$.

Let $S_j^{(k)} = \bigcup_{i=-\infty}^{\infty} S_{3i+k,j}$, for $k \in \{0,1,2\}$, be the union of the cells in a grid column that are three cells apart. We say that $S_j^{(k)}$ is the k^{th} *sub-strip* of the j^{th} grid column. We denote by $L_j^{(k)} = \bigcup_{i=-\infty}^{\infty} L_{3i+k,j}$ the set of line segments resulting from intersecting the input set of line segments with $S_j^{(k)}$. Let the number of segments of each set $L_j^{(k)}$ be $n_j^{(k)}$.

No two grid cells $S_{3i+k,j}$ and $S_{3i'+k,j}$, with $i \neq i'$, in sub-strip $S_j^{(k)}$, contain a common segment since they are more than one unit apart vertically. The lower envelope of a sub-strip $S_j^{(k)}$ is the lower envelope of the lower envelopes for the squares included in $S_j^{(k)}$ and therefore any vertex in the lower envelope in a square is either included or excluded in the lower envelope of $S_j^{(k)}$, whereby the complexity is $\sum_{i=-\infty}^{\infty} O(n_{3i+k,j}) = O\left(n_j^{(k)}\right)$; see Figure 2(a).

We now consider the lower envelope of sub-strips that are three squares apart horizontally. We denote the union of such sub-strips by $S^{(k,l)} \stackrel{\text{def}}{=} \bigcup_{j=-\infty}^{\infty} S_{3j+l}^{(k)}$, for $l \in \{0,1,2\}$. Let $L^{(k,l)} = \bigcup_{j=-\infty}^{\infty} L_{3j+l}^{(k)}$ be the set of line segments resulting from intersecting the input set of line segments with $S^{(k,l)}$. Let the number of segments of each set $L^{(k,l)}$ be $n^{(k,l)}$.

Any two sets $L_{3j+l}^{(k)}$ and $L_{3j'+l}^{(k)}$, with $j \neq j'$, must have empty intersection since the corresponding sub-strips $S_{3j+l}^{(k)}$ and $S_{3j'+l}^{(k)}$ are more than one unit apart. Hence, no segment occurs in the two sub-strips whereby the lower envelope of the segments in $S^{(k,l)}$ has complexity $\sum_{j=-\infty}^{\infty} O\left(n_{3j+l}^{(k)}\right) = O\left(n^{(k,l)}\right) \subseteq O(n)$; see Figure 2(b).

The complexity of the lower envelope of $\bigcup_{\substack{0 \le k \le 2 \\ 0 \le l \le 2}} L^{(k,l)}$, i.e., the whole domain, is then

$$\sum_{\substack{0 \le k \le 2 \\ 0 \le l \le 2}} O\left(n^{(k,l)}\right) \subseteq \sum_{\substack{0 \le k \le 2 \\ 0 \le l \le 2}} O(n) = O(n), \tag{1}$$

using Observation 1.1 since we are summing over nine linear sized subsets. ◀

3.1 Segments Traced by Moving Points with Constant Speed

Let P be a set of n points in the plane, each moving at the same constant speed along a different line. The points start simultaneously moving at an instant $t = 0$, so at any given

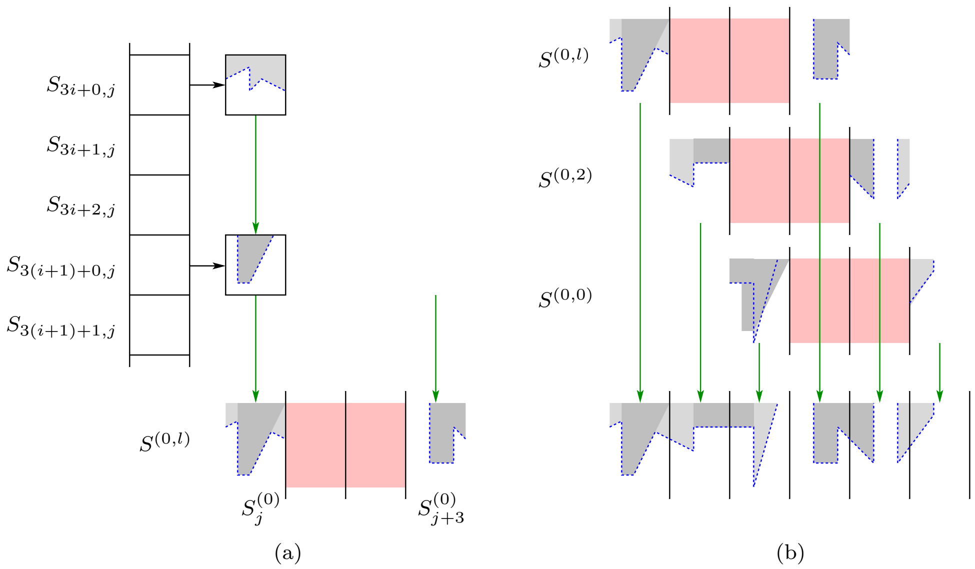

Figure 2 Illustrating the proof of Theorem 3.1.

instant $t > 0$, the points have traced a set L_t of n line segments, all with equal length; see Figure 3. By Theorem 3.1, and combining Corollary 3.2 with the algorithm by Hershberger [2], we obtain the following result.

▶ **Corollary 3.2.** *For any fixed* $t > 0$*, the lower envelope of* L_t *has a complexity of* $\Theta(n)$ *and can be computed in* $O(n \log n)$ *time and* $O(n)$ *space.*

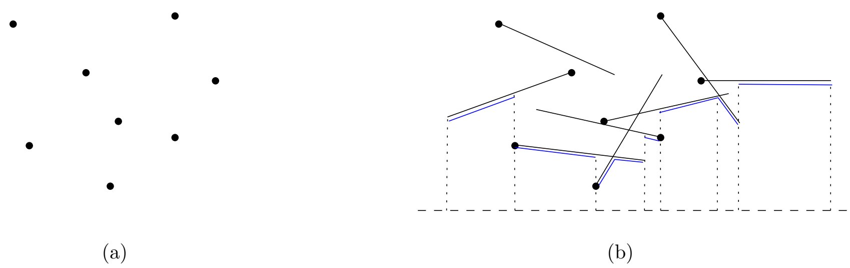

Figure 3 (a) A set P of eight points in the plane at an instant $t = 0$. (b) The lower envelope of the segments traced by moving the points of P along linear trajectories, at an instant $t > 0$.

4 Collections of Unit Squares Under Linear Transformations

We consider unit squares and allow rotation, translation, and scaling. We settle the complexity of the lower envelope for all possible combinations of these transformations, see Table 1. More specifically, we consider n copies of the unit square with corners at $[0, 0]$, $[0, 1]$, $[1, 0]$, and $[1, 1]$, and apply a subset of the transformations to these n unit squares.

If we do not allow scaling, each square has four unit-length segments, out of which at most two appear on the lower envelope. Combining Theorem 3.1 and Observation 1.1, we get

▶ **Corollary 4.1** (Cases 2, 5, and 6)**.** *The lower envelope of a set of* n *unit squares that can be rotated and/or translated has a complexity of* $\Theta(n)$*.*

Case	Rotation	Translation	Scaling	Complexity
1	$\times$	$\times$	$\times$	$\Theta\big(n\alpha(n)\big)$
2	$\times$	$\times$		$\Theta(n)$
3	$\times$		$\times$	$\Theta(n)$
4		$\times$	$\times$	$\Theta(n)$
5	$\times$			$\Theta(n)$
6		$\times$		$\Theta(n)$
7			$\times$	$\Theta(1)$

Table 1 Complexity of the lower envelope of unit squares under various linear transformations.

On the other hand, if we only allow scaling, we can only achieve constant complexity.

▶ **Lemma 4.2** (Case 7). *The lower envelope of a set of n unit squares that can be scaled has a complexity of $\Theta(1)$.*

▶ **Lemma 4.3** (Case 4). *The lower envelope of a set of n unit squares that can be translated and scaled has a complexity of $\Theta(n)$.*

Proof. For the lower bound, consider the lower envelope of the n axis-aligned squares with base edges $[(1,0),(2,0)],[(3,0),(4,0)],\dots,[(n-1,0),(n,0)]$. Since such a lower envelope has $2n$ vertices, we have established the $\Omega(n)$ bound.

For the upper bound, we split the segments of the squares into two groups: All n squares have the same rotation, hence, for each square either a single horizontal line segment (if the squares are axis-aligned) or two line segments with two coinciding slopes for all squares may appear on the lower envelope. We consider the line segments with non-negative and those with negative slope separately. Let L^+ and L^- be the set of line segments with non-negative and negative slope, respectively. (L^- may be empty.) All line segments in L^+ have the same slope, hence, no two line segments from the set can intersect. We order the segments in L^+ by their left endpoint and insert them one after another. When we insert a new line segment $\ell_i \in L^+$, we introduce at most two vertices to the lower envelope:

1. If the left endpoint of ℓ_i is to the right of all previously inserted segments, we introduce two new vertices to the lower envelope.
2. If the left endpoint of ℓ_i is to the left of some right endpoints of previously inserted line segments, but its right endpoint is to the right of all previously inserted segments, we introduce at most two vertices to the lower envelope: both of ℓ_i's endpoints or its right endpoint (the other vertex is in that case an already introduced endpoint of a segment).
3. If the complete segment ℓ_i is within the interval of x-coordinates of previously inserted segments, no left endpoint of a segment $\ell_j, j < i$ can be to the right of ℓ_i's left endpoint, thus, depending on the y-coordinates of the so far introduced segments, we introduce either ℓ_i's endpoints, its right endpoint, or no point to the lower envelope.

Hence, the lower envelope of the line segments in L^+ has complexity $O(n)$. An analogous argument yields the same bound for the lower envelope of line segments in L^-. Hence, with Observation 1.1, we yield the upper bound. ◀

▶ **Lemma 4.4** (Case 3). *The lower envelope of a set of n unit squares that can be scaled and rotated has complexity $\Theta(n)$.*

While in all previous cases, we could give a linear bound on the complexity of the lower envelope, this does not hold if we allow all three linear transformations.

▶ **Theorem 4.5** (Case 1). *The lower envelope of a set of n unit squares that can be rotated, translated, and scaled has complexity $\Theta\big(n\alpha(n)\big)$.*

Proof. For a collection of n line segments in the plane (none of which are vertical), Hart and Sharir [1] showed that the complexity of the lower envelope can be at most $O\big(n\alpha(n)\big)$. For each of the n squares at most two line segments appear on the lower envelope. Thus, the result by Hart and Sharir in combination with Observation 1.1 yields an upper bound of $O\big(n\alpha(n)\big)$ on the complexity of the lower envelope of the squares.

Let $G = \{\ell_1, \ldots, \ell_n\}$ be the set of line segments from the lower-bound construction of Wiernik and Sharir [4]. We construct a new set of line segments $G' = \{\ell'_1, \ldots, \ell'_n\}$ from G: for a large constant M, we substitute each endpoint (e_x, e_y) of a segment by the endpoint $(e_x, \frac{e_y}{M})$. For an example, see Figure 4. If any two segments ℓ_i and ℓ_j intersect in the point (p_x, p_y), then ℓ'_i and ℓ'_j intersect in the point $(p_x, \frac{p_y}{M})$. Thus, if ℓ_i is on the lower envelope of G for the interval (x_1, x_2), then ℓ'_i is on the lower envelope of G' for the same interval (x_1, x_2). Consequently, the complexity of the lower envelope of G' equals that of G.

We use the new ("nearly" horizontal) line segments as the base edge of our squares S, see Figure 4: Let ℓ'_i have endpoints (l^i_x, l^i_y) and (r^i_x, r^i_y), we construct a square s_i with side length $||\ell'_i||$. For each of these squares at most two edges are on the lower envelope.

For M being large enough, the square edges appearing on the lower envelope that do not stem from the line segments from the construction by Wiernik and Sharir have a very large absolute value of slope (that is, they are "nearly" vertical). We can choose M large enough such that for ℓ_i, ℓ_j, ℓ_k appearing on the lower envelope in that order, the vertical edges of s_i and s_k cannot block s_j (with higher y-coordinate) from appearing on the lower envelope: Let ℓ_j appear on the lower envelope of G within the interval $I_j = [r^i_x, l^k_x]$. We split I_j into three equal-length closed intervals I_{jl}, I_{jm} and I_{jr}, aiming that s_i and s_k may block at most I_{jl} and I_{jr}, respectively—which yields the claim. We consider the case that ℓ_i has negative slope (with positive slope, s_i does not block any of I_{jl}). For simplicity, assume that ℓ_j is horizontal (similar arguments hold in the other cases). We can consider ℓ_j, because if ℓ_j is not blocked in I_{jm} then ℓ'_j (with smaller y-coordinates) is not blocked either. Assume, we construct a square σ_i with ℓ_i as base edge (instead of ℓ'_i for s_i). Let e_i be σ_i's second edge that may appear on the lower envelope, and let $p^{ij} = (p^{ij}_x, p^{ij}_y)$ be the intersection point of e_i and ℓ_j. If $p^{ij} \in I_{jl}$, we are done for any value of M, hence, let $p^{ij} \in I_{jm} \cup I_{jr}$. The slope of e_i is $\frac{p^{ij}_y - r^i_y}{p^{ij}_x - r^i_x}$, the equivalent edge of s_i should have slope at least $\frac{p^{ij}_y - r^i_y}{r^i_x + 1/3(l^k_x - r^i_x) - r^i_x} = \frac{p^{ij}_y - r^i_y}{1/3(l^k_x - r^i_x)}$, such that we achieve $p^{ij} \in I_{jl}$. Hence, we need to choose $M \geq \frac{1/3(l^k_x - r^i_x)}{p^{ij}_x - r^i_x}$. By choosing M larger than these constraints for all pairs of line segments appearing consecutively on the lower envelope, we can ensure that while s_i and ℓ_i will not appear on the exact same interval on the lower envelope for S and G, the sequence of the s_i appearing on the lower envelope for S coincides with the sequence of the ℓ_i appearing on the lower envelope for G. Hence, the lower bound for line segments established by Wiernik and Sharir [4] translates to unit squares that can be rotated, translated and scaled. ◀

5 Open Questions

The complexity for other geometric shapes, such as differently oriented parabolae, ellipses, and fat objects in general would be of great interest to settle. In particular, we are interested in knowing if the complexity can be proved with purely geometric arguments.

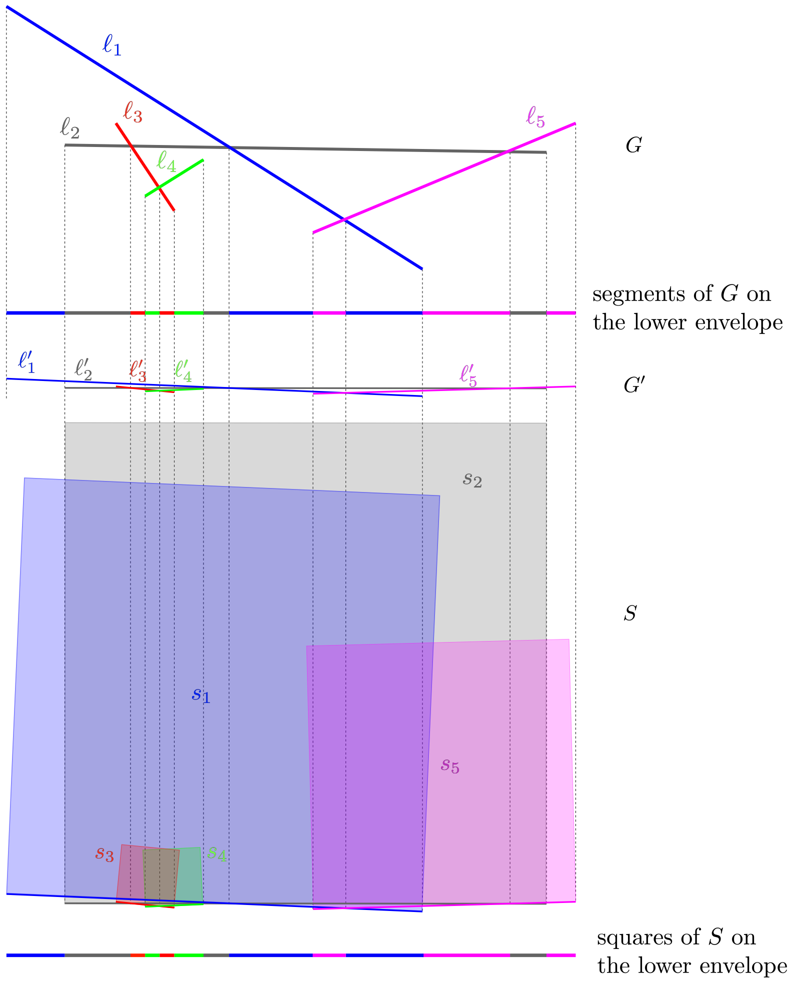

Figure 4 Example for the construction from the proof of Theorem 4.5. Top: Set of line segments $G = \{\ell_1, \ldots, \ell_5\}$; second to top: the sequence of segments of G appearing on the lower envelope; middle: new set of line segments $G' = \{\ell'_1, \ldots, \ell'_5\}$; second to bottom: set of squares $S = \{s_1, \ldots, s_5\}$; bottom: sequence of squares of S appearing on the lower envelope.

References

1 Sergiu Hart and Micha Sharir. Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes. *Combinatorica*, 6(2):151–177, 1986. doi:10.1007/BF02579170.

2 John Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Information Processing Letters*, 33(4):169–174, 1989. doi:10.1016/0020-0190(89)90136-1.

3 Micha Sharir and Pankaj K. Agarwal. *Davenport–Schinzel Sequences and their Geometric Applications*. Cambridge University Press, 1995.

4 Ady Wiernik and Micha Sharir. Planar realizations of nonlinear Davenport-Schinzel sequences by segments. *Discrete & Computational Geometry*, 3(1):15–47, 1988. doi:10.1007/BF02187894.