On Erdős–Szekeres-Type Problems for k-convex Point Sets

Martin Balko[1](✉), Sujoy Bhore[2], Leonardo Martínez Sandoval[3], and Pavel Valtr[1]

[1] Department of Applied Mathematics, Charles University, Prague, Czech Republic
balko@kam.mff.cuni.cz

[2] Algorithms and Complexity Group, TU Wien, Vienna, Austria
sujoy.bhore@gmail.com

[3] Institut de Mathématiques de Jussieu – Paris Rive Gauche (UMR 7586), Sorbonne Université, Paris, France
leomtz@im.unam.mx

Abstract. We study Erdős–Szekeres-type problems for k-convex point sets, a recently introduced notion that naturally extends the concept of convex position. A finite set S of n points is *k-convex* if there exists a spanning simple polygonization of S such that the intersection of any straight line with its interior consists of at most k connected components. We address several open problems about k-convex point sets. In particular, we extend the well-known Erdős–Szekeres Theorem by showing that, for every fixed $k \in \mathbb{N}$, every set of n points in the plane in *general position* (with no three collinear points) contains a k-convex subset of size at least $\Omega(\log^k n)$. We also show that there are arbitrarily large 3-convex sets of n points in the plane in general position whose largest 1-convex subset has size $O(\log n)$. This gives a solution to a problem posed by Aichholzer et al. [2].

We prove that there is a constant $c > 0$ such that, for every $n \in \mathbb{N}$, there is a set S of n points in the plane in general position such that every 2-convex polygon spanned by at least $c \cdot \log n$ points from S contains a point of S in its interior. This matches an earlier upper bound by Aichholzer et al. [2] up to a multiplicative constant and answers another of their open problems.

The project leading to this application has received funding from European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No. 678765. M. Balko and P. Valtr were supported by the grant no. 18-19158S of the Czech Science Foundation (GAČR). M. Balko and L. Martínez-Sandoval were supported by the grant 1452/15 from Israel Science Foundation. M. Balko was supported by Center for Foundations of Modern Computer Science (Charles University project UNCE/SCI/004). S. Bhore was supported by the Austrian Science Fund (FWF) under project number P31119. L. Martínez Sandoval was supported by the grant ANR-17-CE40-0018 of the French National Research Agency ANR (project CAPPS). This research was supported by the PRIMUS/17/SCI/3 project of Charles University.

C. J. Colbourn et al. (Eds.): IWOCA 2019, LNCS 11638, pp. 35–47, 2019.
https://doi.org/10.1007/978-3-030-25005-8_4

Keywords: Convex position · Point set · k-convex point set · k-convex polygon

1 Introduction

A set of points in the plane is in *convex position* if its points are vertices of a convex polygon. We say that a planar point set is in *general position* if it does not contain a collinear triple of points. A classical result by Erdős and Szekeres [6], called the *Erdős–Szekeres Theorem*, states that every set of n points in the plane in general position contains a set of $\Omega(\log n)$ points in convex position. Moreover, this result is asymptotically tight, with the strongest bounds given in the papers [7,10,15]. The Erdős–Szekeres Theorem, published in 1935, was one of the starting points of both discrete geometry and Ramsey theory. Since then, numerous variants of this result have been studied.

For example, in 1978, Erdős [5] asked about the growth rate of the smallest integers $h(m)$, $m \geq 3$, such that every set P of at least $h(m)$ points in the plane in general position contains an *m-hole* in P, that is, m points in convex position with no point of P in the interior of their convex hull. It is easy to show that $h(3) = 3$ and $h(4) = 5$ and Harborth [9] proved $h(5) = 10$. After this, the question about the existence of the numbers $h(m)$ was settled in two phases. First, in 1983, Horton [11] showed that there are arbitrarily large sets of points with no 7-holes, proving that $h(m)$ does not exist for $m \geq 7$. Around 25 years later, Gerken [8] and Nicolás [13] independently proved that every sufficiently large set of points in the plane in general position contains a 6-hole. In particular, $h(m)$ exists if and only if $m \leq 6$.

In this paper, we study variants of these classical problems for so-called k-convex point sets, a notion that was recently introduced by Aichholzer et al. [2] and that naturally extends the concept of convex position. We also address further open problems about k-convex point sets posed in [2].

Throughout the paper, we consider only finite sets of points in the plane in general position. We use ∂S to denote the boundary of a simple polygon S. For a line segment s, we use $\overline{s}$ to denote the supporting line of s. A line ℓ *crosses* ∂S at a point v if ℓ passes through v from the interior of S to the outside of S. All logarithms in the paper are base two.

2 Preliminaries

In 2012, Aichholzer et al. [1] introduced the following natural extension of convex polygons. For a positive integer k, a simple polygon S with vertices in general position is *k-convex* if no straight line intersects S in more than k connected components. This notion has been later transcribed to finite point sets [2]. A finite set P of points in the plane in general position is *k-convex* if P is a vertex set of a k-convex polygon. In other words, P is k-convex if there exists a spanning simple polygonization of P such that the intersection of any straight line with its interior consists of at most k connected components. It can be shown that a simple polygon S with vertices in general position is k-convex if and only if every line not containing a vertex of S intersects the boundary of S in at most $2k$ points (see Lemma 2).

The notion of k-convexity for point sets satisfies several natural properties. A point set is in convex position if and only if it is 1-convex. Clearly, for each $k \in \mathbb{N}$, every k-convex point set is $(k+1)$-convex. Aichholzer et al. [2, Lemma 2] showed that every subset of a k-convex point set is also k-convex. It is known that every set of n points is k-convex for some $k = O(\sqrt{n})$ and this bound is tight up to a multiplicative constant in the worst case [2, Theorem 2]. Some further results about k-convex polygons and k-convex point sets can be found in [1–3].

Erdős–Szekeres-type questions were among the first problems about k-convex point sets considered in the literature. Aichholzer et al. [2] showed that every set of n points in general position contains a 2-convex subset of size at least $\Omega(\log^2 n)$ [2, Theorem 5] and that this bound is tight up to a multiplicative constant. This result led the authors to pose the following problem.

Problem 1 ([2, Open problem 4]). Let k and n be positive integers. Find the maximum integer $g(k,n)$ such that every set of n points contains a k-convex set of size $g(k,n)$.

Using this notation, their result gives $g(2,n) = \Theta(\log^2 n)$ and the Erdős–Szekeres Theorem gives $g(1,n) = \Theta(\log n)$. No nontrivial bounds were known for $g(k,n)$ with $k \geq 3$.

In a slightly different direction, it was shown that every 2-convex polygon with n vertices contains a 1-convex subset of at least $\lceil \sqrt{n}/2 \rceil$ vertices and that this bound is tight [1, Theorem 14]. In [2], the authors considered related variants of this result and posed the following problem.

Problem 2 ([2, Open problem 3]). Let j, k and n be positive integers. Find the maximum integer $f(k,n)$ such that every k-convex set of n points contains a 1-convex subset of size $f(k,n)$. More generally, find the maximum integer $f(k,j,n)$ such that every k-convex set of size n contains a j-convex subset of size $f(k,j,n)$.

By definition, $f(k,n) = f(k,1,n)$ for all k and n. With this notation, the result by Aichholzer et al. [1, Theorem 14] gives $f(2,n) = f(2,1,n) = \Theta(\sqrt{n})$. We trivially have $f(1,n) = f(1,1,n) = n$ for every n. Since every set of n points is $(c\sqrt{n})$-convex for some constant $c > 0$ [2, Theorem 2], the Erdős–Szekeres Theorem gives $g(1,n) = f(k,n) = f(k,1,n) = \Theta(\log n)$ for each $k \geq c\sqrt{n}$. By the previous results, we also know that, for $k \geq c\sqrt{n}$, we have $g(2,n) = f(k,2,n) = \Theta(\log^2 n)$ and $g(j,n) = f(k,j,n)$ for each $j \in \mathbb{N}$.

For a point set P, a 2-convex polygon with vertices from P is *empty* in P if it contains no point of P in the interior. Concerning the question of Erdős about m-holes in point sets, Aichholzer et al. [2, Theorem 3] showed that every set P of n points in general position contains a 2-convex polygon that is empty in P and has size at least $\Omega(\log n)$. Using the tightness of the Erdős–Szekeres Theorem, they also proved that there are arbitrarily large point sets P of n points with no empty 2-convex polygon in P of size at least $c \cdot \log^2 n$ for some constant c. There is a gap between these two bounds and thus the authors posed the following problem.

Problem 3 ([2]). Close the gap between the $\Omega(\log n)$ and $O(\log^2 n)$ bounds for the size of empty 2-convex polygons in point sets of size n.

Let us also remark that it was shown by Aichholzer et al. [3] that every 2-convex point set of size n contains an m-hole for $m = \Omega(\log n)$ and that this bound is tight up to a multiplicative constant in the worst case.

The list of problems about k-convex point sets posed by Aichholzer et al. [2] contains several other interesting open questions.

3 Our Results

First, we prove the following extension of the Erdős–Szekeres Theorem for k-convex point sets.

Theorem 1. *Let k be a fixed positive integer. Then*

$$g(k, n) = \Omega(\log^k n).$$

That is, for every $n \in \mathbb{N}$, every set of n points in the plane in general position contains a k-convex subset of size at least $\Omega(\log^k n)$.

Note that Theorem 1 extends the result of Aichholzer et al. [2, Theorem 5] about the existence of large 2-convex point sets in general sets of n points. Unfortunately, we do not have matching upper bounds on the function $g(k, n)$. It follows from the proof of Theorem 13 in [3] that $g(k, n) = O(k\sqrt{n})$ for every $k \geq 3$.

We also address Problem 2. Using a variant of the sets defined by Erdős and Szekeres, we provide asymptotically tight estimates on the function $f(k, n)$ in the case $k \geq 3$.

Theorem 2. *There is a constant c such that, for every positive integer n, there are 3-convex sets of n points in the plane in general position with no 1-convex subset of size larger than $c \cdot \log n$.*

More precisely, for every $t \geq 3$, there is a 3-convex set of 2^{t-2} points in the plane in general position with no 1-convex subset of size t.

Thus $f(k, n) = O(\log n)$ for every integer k with $k \geq 3$. It follows from the Erdős–Szekeres Theorem that this bound is asymptotically tight. That is, $f(k, n) = \Theta(\log n)$ for $k \geq 3$. Therefore Theorem 2 asymptotically settles the first part of Problem 2. The statement in the second sentence of Theorem 2 implies the more precise bound $f(3, n) \leq \lceil \log(n) + 1 \rceil$. It also shows that the corresponding best known bound in the Erdős–Szekeres Theorem can be achieved by 3-convex sets. A famous conjecture of Erdős and Szekeres [7] states that this bound is tight for general sets. If true, the conjecture of Erdős and Szekeres together with our result would give the precise values $f(k,n) = \lceil \log(n) + 1 \rceil$ for any $n, k \geq 3$.

In the proof of Theorem 2 we define planar point sets which might be of independent interest. We call them *(combinatorial) Devil's staircases.*

Aichholzer et al. [1, Theorem 14] showed that every 2-convex point set of size n contains a 1-convex subset of size at least $\Omega(\sqrt{n})$. Their result together with Theorem 2 gives the following estimate on the function $f(k,2,n)$ for $k \geq 3$.

Corollary 1. *There is a constant c such that, for every positive integer n, there are 3-convex sets of n points in the plane in general position with no 2-convex subset of size larger than $c \cdot \log^2 n$.*

In particular, $f(k,2,n) = O(\log^2 n)$ for every integer k with $k \geq 3$. Aichholzer et al. [2, Theorem 5] also showed that every set of n points contains a 2-convex subset of size at least $\Omega(\log^2 n)$. Thus the bound from Corollary 1 is tight up to a multiplicative constant, settling the second part of Problem 2 in the case $j = 2$. The second part of Problem 2 remains open for $j \geq 3$.

Concerning empty 2-convex polygons in general sets of n points, we show that so-called *Horton sets* do not contain large empty 2-convex polygons. More specifically, we derive the following bound.

Theorem 3. *There is a constant $c > 0$ such that, for every positive integer n, there are sets of n points in the plane in general position that contain no empty 2-convex polygon on at least $c \cdot \log n$ vertices.*

The upper bound from Theorem 3 matches the earlier lower bound [2, Theorem 3] up to a multiplicative constant. In other words, Theorem 3 yields a solution to Problem 3.

Aichholzer et al. [2] proved that, for all positive integers k and l, the union of a k-convex point set T and an l-convex point set S is $(k+l+1)$-convex. Moreover, they showed that if a k-convex polygonization of T and an l-convex polygonization of S intersect, then $T \cup S$ is $(k+l)$-convex. Aichholzer et al. [2] found a set of 10 points that is not 2-convex and is a union of two 1-convex sets, showing that the first bound is tight for $k = l = 1$. Besides the case $k = l = 1$, no matching bound is known and Aichholzer et al. asked [2, Open problem 2] whether there are examples for general integers k and l such that the union of a k-convex point set and an l-convex point set is not $(k+l)$-convex. We prove the following almost matching bound.

Proposition 1. *For all positive integers k and l, there are point sets T_k and S_l such that T_k is k-convex, S_l is l-convex, and $T_k \cup S_l$ is not $(k+l-1)$-convex.*

It follows from the proof of Proposition 1 that the bound $k+l$ by Aichholzer et al. [2] on the convexity of a union of a k-convex point set with an l-convex point set with intersecting polygonizations is tight in the worst case. The proof also gives an explicit construction, for every $k \in \mathbb{N}$, of a k-convex set that is not $(k-1)$-convex. Such an example seemed to be missing in the literature.

The proofs of the second part of Theorems 2, 3, and Proposition 1 are in the full version of the paper.

4 Proof of Theorem 1

For a fixed positive integer k, we show that every set of n points in the plane in general position contains a k-convex subset of size at least $\Omega(\log^k n)$. We first state two auxiliary statements. The first one, the *Erdős–Szekeres Lemma*, is a classical result proved by Erdős and Szekeres [6].

Lemma 1 ([6]). *For every $n \in \mathbb{N}$, every sequence of $(n-1)^2+1$ real numbers contains a non-increasing or a non-decreasing subsequence of length at least n.*

The main idea of the proof of Theorem 1 is inspired by the approach of Aichholzer et al., who proved the lower bound $\Omega(\log^2 n)$ for the case $k = 2$ [2, Theorem 5]. The key ingredient of the proof is the so-called *Positive Fraction Erdős–Szekeres Theorem* proved by Bárány and Valtr [4]. We use a version of the theorem that was used by Suk [15] and that is based on a bound proved by Pór and Valtr [14] (Theorem 4 below). Before stating it, we first introduce some notation.

A set of points in the plane with distinct x-coordinates is a *cup* if the points lie on the graph of a strictly convex function. Similarly, it is a *cap* if the points lie on the graph of a strictly concave function. Given a cap or a cup $C = \{c_1, \ldots, c_l\}$ with points of C ordered according to the increasing x-coordinates, the *support* of C is the collection of open regions $T_1, \ldots, T_l$, where each T_i is the region outside of the convex hull of C bounded by the line segment $c_i c_{i+1}$ and by the lines $\overline{c_{i-1}c_i}$ and $\overline{c_{i+1}c_{i+2}}$, where $c_0 = c_l$, $c_{l+1} = c_1$, and $c_{l+2} = c_2$; see part (a) of Fig. 1. The *base* of each region T_i is the line segment $c_i c_{i+1}$ and we call the line $\overline{c_i c_{i+1}}$ the *base line* of T_i.

Theorem 4 ([4,14,15]). *Let $l \geq 3$ be an integer and P be a finite set of points in the plane in general position such that $|P| \geq 2^{32l}$. Then there is a set C of l points of P such that C is a cap or a cup and the regions $T_1, \ldots, T_{l-1}$ from the support of C satisfy $|T_i \cap P| \geq \frac{|P|}{2^{32l}}$ for every $i \in \{1, \ldots, l-1\}$.*

Let k be a fixed positive integer and let P be a set of n points in the plane in general position with n sufficiently large with respect to k. A curve C in the plane is *x-monotone* or *y-monotone* if every vertical or horizontal line, respectively, intersects C in at most one point. We proceed by induction on k and show that there is a k-convex subset Q of P of size at least $\Omega(\log^k n)$ such that the polygonization of Q has the boundary formed by a union of an x-monotone curve and one edge. We make no serious effort to optimize the constants.

First, Theorem 1 for $k = 1$ follows from the Erdős–Szekeres Theorem [6], as stated at the beginning of the introduction, because each set of points in convex

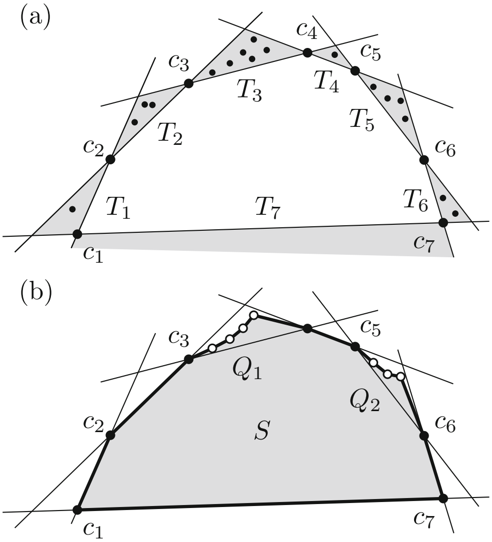

Fig. 1. (a) An illustration of the statement of the Positive Fraction Erdős–Szekeres Theorem (Theorem 4). (b) A construction of the polygon S for $k = 2$. In this example, we have $l = 7$. The described procedure gives $m = 2$, $i_1 = 3$, and $i_2 = 5$, because σ_3, σ_5 are non-decreasing.

position is a union of a cap and a cup that intersect only in two points. This finishes the base case.

Now, for the induction step, assume $k \geq 2$. Without loss of generality we assume that no two points of P have the same x-coordinate. By Theorem 4 applied with $l = \lfloor \log n/64 \rfloor$, there is a set $C = \{c_1, \ldots, c_l\}$ of l points from P such that C is a cap or a cup and the regions $T_1, \ldots, T_{l-1}$ from the support of C satisfy $|T_i \cap P| \geq \frac{|P|}{2^{32l}} \geq \sqrt{n}$. Let $\prec$ be the ordering of the points from P according to their increasing x-coordinates. Note that $c_1 \prec c_2 \prec c_3 \prec \cdots \prec c_l$. By symmetry, we assume that C is a cap.

For every odd i with $3 \leq i < l - 1$, we apply Lemma 1 to the sequence of distances of points from $T_i \cap P$ to the base line of T_i, ordered by $\prec$. For each such sequence, we obtain a non-increasing or a non-decreasing subsequence σ_i of length at least $\sqrt{|T \cap P_i|} \geq n^{1/4}$. By the pigeonhole principle, there are $m \geq (l-3)/4$ subsequences $\sigma_{i_1}, \ldots, \sigma_{i_m}$ with odd indices $1 < i_1 < \cdots < i_m < l - 1$ such that all these sequences are non-increasing or all non-decreasing. By symmetry, we may assume that $\sigma_{i_1}, \ldots, \sigma_{i_m}$ are all non-decreasing. In the other case we would proceed analogously, considering the ordering of $\prec^{-1}$. For every $j \in \{1, \ldots, m\}$, let P_j be the set of points from $T_{i_j} \cap P$ that determine the distances in σ_{i_j}. In particular, the distances of the points of P_j to the base line of T_{i_j} are non-decreasing in $\prec$ and $|P_j| \geq n^{1/4}$.

By the induction hypothesis applied to each set P_j, there is a $(k-1)$-convex subset Q_j of P_j of size at least $c\log^{k-1}(n^{1/4}) = \frac{c}{4}\log^{k-1} n$ for some constant $c = c(k-1) > 0$ such that some $(k-1)$-convex polygonization of Q_j is formed by a union of an x-monotone curve O_j and one edge. Observe that $c_{i_j} \prec q \prec c_{i_j+1}$ for every $q \in Q_j$, as $1 < i_j < l-1$.

We construct a polygonization of the set $Q = C \cup \bigcup_{j=1}^{m} Q_j$ by connecting the first and the last vertex of O_j in $\prec$ to c_{i_j} and c_{i_j+1}, respectively, with a line segment for each $j \in \{1,\ldots,m\}$. We then add the line segments c_ic_{i+1} for each $i \in \{1,\ldots,l-1\} \setminus \{i_1,\ldots,i_m\}$ and the line segment c_1c_l; see part (b) of Fig. 1. Since $c_1 \prec \cdots \prec c_l$ and $c_{i_j} \prec q \prec c_{i_j+1}$ for all $j \in \{1,\ldots,m\}$ and $q \in Q_j$, the resulting closed piecewise linear curve is a boundary of a simple polygon S with the vertex set $Q \subseteq P$. Moreover, the boundary of S is formed by a union of an x-monotone curve and the edge c_1c_l. Note that

$$|Q| > \sum_{j=1}^{m} |Q_j| \geq m\frac{c}{4}\log^{k-1} n \geq \frac{\lfloor \log n/64 \rfloor - 3}{4}\frac{c}{4}\log^{k-1} n = \Omega(\log^k n)$$

for n sufficiently large with respect to k.

It remains to prove that the polygon S is k-convex. We start with the following simple observation that restricts the set of lines we have to check.

Lemma 2. *For every $k \in \mathbb{N}$, a simple polygon S with vertices in general position is k-convex if and only if every line not containing a vertex of S intersects ∂S in at most $2k$ points.*

Proof. First, if S is k-convex, then each line ℓ intersects S in at most k connected components. If l contains no vertex of S then each such a component is a line segment with endpoints in ∂S and with interior contained in the interior of S. Thus ℓ intersects ∂S in at most $2k$ points.

On the other hand, if S is not k-convex, then there is a line ℓ that intersects S in more than k connected components. We say that a component of $S \cap \ell$ is *regular* if it contains no vertex of S. Suppose for simplicity that ℓ is horizontal. Since the vertices of S are in general position, at most two components are not regular. Every regular component intersects ∂S in exactly two points. It follows that if all components are regular then ℓ intersects ∂S in at least $2k+2$ points.

Suppose now that there is a unique component A containing one or two vertices of S. Then either moving ℓ a little bit up or moving it a little bit down turns the component A into one or more regular components and the other components remain regular. Consequently, the perturbed line ℓ contains no vertex of S and intersects ∂S in at least $2k+2$ points.

Finally, suppose that there are two non-regular components A and B, each containing exactly one vertex of S. Then A can be turned into a regular component either by slightly perturbing ℓ arbitrarily in such a way that it passes above the point $A \cap \mathrm{vert}(S)$, where vert(S) denotes the vertex set of S, or by slightly perturbing it arbitrarily in such a way that it passes below the point $A \cap \mathrm{vert}(S)$. A similar statement holds for the component B. We claim that a suitable slight perturbation of ℓ turns each of the components A and B into a regular component. Indeed, it is sufficient to move ℓ a little bit up or down or to rotate it slightly clockwise or counterclockwise around the middle point of the segment connecting the points $A \cap \text{vert}(S)$ and $B \cap \text{vert}(S)$. Thus there is a perturbation of ℓ such that the resulting line does not contain a vertex of S and intersects ∂S in at least $2k+2$ points. This finishes the proof of Lemma 2.

By Lemma 2, it suffices to show that every line ℓ not containing a vertex of S intersects ∂S in at most $2k$ points. Since such a line ℓ intersects ∂S in an even number of points, it actually suffices to show that it intersects ∂S in at most $2k+1$ points. Every edge of ∂S is contained in the closure $\text{cl}(T_i)$ of some T_i. Since ℓ intersects at most two regions $\text{cl}(T_i)$, it suffices to prove the following claim.

Lemma 3. *The following two conditions are satisfied.*

(i) For every i, $|\ell \cap \partial S \cap \text{cl}(T_i)| \leq 2k$.
(ii) If ℓ intersects two different regions T_α and T_β then $|\ell \cap \partial S \cap \text{cl}(T_\alpha)| \leq 1$ or $|\ell \cap \partial S \cap \text{cl}(T_\beta)| \leq 1$.

Proof. We first prove part (i) of Lemma 3. If $i \notin \{i_1, \ldots, i_m\}$ then $\text{cl}(T_i)$ contains at most one edge of ∂S. Thus, we have $|\ell \cap \partial S \cap \text{cl}(T_i)| \leq 1 < 2k$ in this case. Otherwise $i = i_j$ for some $j \in \{1, \ldots, m\}$, and then $|\ell \cap \partial S \cap \text{cl}(T_i)| \leq 2k$, since ℓ intersects O_j in at most $2k-2$ points and it intersects each of the two remaining edges of ∂S contained in $\text{cl}(T_i)$ at most once. Part (i) of Lemma 3 follows.

To show part (ii) of Lemma 3, assume that, say, $1 \leq \alpha < \beta \leq l$. If $\beta = \alpha + 1$ then α or β is even and thus $|\ell \cap \partial S \cap \text{cl}(T_\alpha)| \leq 1$ or $|\ell \cap \partial S \cap \text{cl}(T_\beta)| \leq 1$, as required. Similarly, we have $|\ell \cap \partial S \cap \text{cl}(T_\beta)| \leq 1$ if $\beta = l$ and thus we assume $\beta < l$.

Assume now that $\beta \geq \alpha + 2$. Then ℓ intersects the bases of T_α and of T_β. We claim that $|\ell \cap \partial S \cap \text{cl}(T_\alpha)| \leq 1$. This is obvious if $\alpha \notin \{i_1, \ldots, i_m\}$.

Assume now that $\alpha = i_j$ for some $j \in \{1, \ldots, m\}$. Let $c_{i_j} = q_1 \prec q_2 \prec \cdots \prec q_{s-1} \prec q_s = c_{i_j+1}$ be the points from $Q_j \cup \{c_{i_j}, c_{i_j+1}\}$. We use x to denote the intersection point of ℓ and the base of T_{i_j}. Let ℓ^+ be the open half-plane determined by ℓ containing q_s; see Fig. 2.

Let q_t be a point from $Q_j \cap \ell^+$. The distance of the point q_t to the base line of T_{i_j} is at most as large as such a distance for q_{t+1} by the choice of Q_j. Thus the point q_{t+1} does not lie in the strip between the base line of T_{i_j} and the line ℓ' parallel to this line containing q_t. Since $q_t \in \ell^+$, the line ℓ intersects ℓ' to the left of q_t. It then follows from $i_j = \alpha < \beta < l$ that the intersection of ℓ with ℓ' is to the left of x and thus ℓ intersects the vertical line containing q_t below q_t. Since $q_t \prec q_{t+1}$, the point q_{t+1} is thus separated from ℓ by ℓ' and the vertical line that contains q_t; see Fig. 2. In particular, $q_{t+1} \in \ell^+$ and ℓ does not intersect the edge $q_t q_{t+1}$.

Since the vertices along O_j are ordered according to $\prec$, it follows that at most one edge of S in $\text{cl}(T_{i_j})$ intersects ℓ and we have $|\ell \cap \partial S \cap \text{cl}(T_{i_j})| \leq 1$, which completes the proof of part (ii) of Lemma 3 and thus also the proof of Theorem 1.

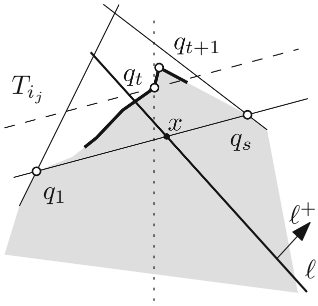

Fig. 2. An illustration of the proof of the fact $|\ell \cap \partial S \cap \text{cl}(T_{i_j})| \leq 1$.

5 Proof of the First Part of Theorem 2

In this section, we construct a 3-convex set of n points with the largest 1-convex subset of size at most $O(\log n)$. Together with the Erdős–Szekeres Theorem, this gives $f(k,n) = f(k,1,n) = \Theta(\log n)$ for all positive integers $k \geq 3$ and n and asymptotically settles the first part of Problem 2. Our example, the so-called Devil's staircase, has a very simple structure and may be of independent interest for reasons discussed in the introduction. In full version of the paper, we also give another example, in which we get a more precise bound described in the second part of Theorem 2. Our examples are modifications of the construction used by Erdős and Szekeres [6] to show the asymptotic tightness of the Erdős–Szekeres Theorem.

A point set D is *deep below* a point set U if the following two conditions are satisfied.

(i) Every point of D lies strictly below each line determined by two points of U, and
(ii) every point of U lies strictly above each line determined by two points of D.

We say that a set S of 2^t points in the plane in general position is a *(combinatorial) Devil's staircase*[1] if S satisfies one of the following two conditions.

(ES1) Either $t = 1$ and the set S consists of two points (x_1, y_1) and (x_2, y_2) with $x_1 < x_2$ and $y_1 < y_2$, or
(ES2) $t \geq 2$ and the set S admits a partition $S = X \cup Y$, where X and Y are both Devil's staircases with 2^{t-1} points. Moreover, X is deep below Y and every point of Y has larger x-coordinate than any point of X.

Let $\{p_1, \ldots, p_n\}$ be the points of a Devil's staircase X_t of size $n = 2^t$ for some $t \in \mathbb{N}$, sorted by increasing x-coordinates. We let $p_1 = (x_1, y_1)$ and $p_n = (x_n, y_n)$ and we define the set $Z_t = X_t \cup \{q\}$ with $q = (x_n, y_1)$.

[1] We chose this name, since the set resembles Cantor function, which is also known under the name Devil's staircase [16].

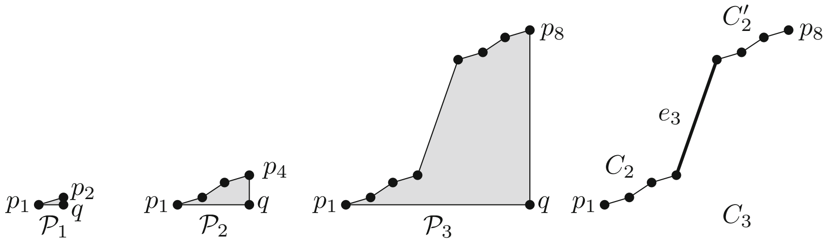

Fig. 3. The polygonizations $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_3$, and the curve C_3.

Now, we show that the set Z_t is 3-convex. To do so, we consider the following polygonization $\mathcal{P}_t$ of Z_t. Let C_t be an x-monotone piecewise-linear curve formed by the line segments $p_i p_{i+1}$ for each $i \in \{1, \dots, n-1\}$. Note that, by Properties ES1 and ES2, the chain C_t is also y-monotone. The polygonization $\mathcal{P}_t$ of Z_t is then the polygon whose boundary consists of C_t and the two line segments $p_1 q$ and $p_n q$; see Fig. 3. The polygon $\mathcal{P}_t$ is simple, since C_t has both coordinates increasing if we traverse it from p_1 to p_n. We now prove that $\mathcal{P}_t$ is a 3-convex polygon.

Lemma 4. *Any line ℓ intersects C_t at most five times. Furthermore, if ℓ is non-vertical and passes above the rightmost point of C_t, then it intersects C_t at most four times.*

Proof. We proceed by induction on t. The case $t \leq 2$ is trivial, thus we assume $t \geq 3$. Since X_t is a Devil's staircase, there is a partition $X_t = X_{t-1} \cup X'_{t-1}$ such that X_{t-1} and X'_{t-1} are Devil's staircases of size 2^{t-1}, X_{t-1} lies deep below X'_{t-1}, and X_{t-1} is to the left of X'_{t-1}. Let C_{t-1} and C'_{t-1} be the x- and y-monotone piecewise-linear curves formed by X_{t-1} and X'_{t-1}, respectively.

Let ℓ be a line. Since the points of X_t are in general position, we can, due to Lemma 2, assume that ℓ does not contain a vertex of X_t in the rest of the proof. First, observe that Property ES2 implies that every line that intersects at least two edges of C_{t-1} lies below X'_{t-1}. Similarly, Property ES2 implies that every line intersecting at least two edges of X'_{t-1} is above X_{t-1}. Thus we can assume that ℓ does not intersect both C_{t-1} and C'_{t-1}. Otherwise ℓ intersects both curves at most once and, since C_t contains only a single edge e_t besides C_{t-1} and C'_{t-1}, the line ℓ intersects C_t at most three times.

Since we assume that ℓ does not intersect C_{t-1} or C'_{t-1}, we may also assume that it intersects the other of these two sets at least twice. This will imply restrictions on ℓ.

We assume first that ℓ intersects C_{t-1} at least twice and show that the statement of the lemma is then satisfied. In this case ℓ passes below the rightmost point of C_t, and we only have to show that it intersects C_t at most five times. This is indeed the case because if ℓ passes below the rightmost point of C_{t-1} then it does not intersect e_t and the statement follows from the inductive hypothesis.

If ℓ passes above the rightmost point of C_{t-1} then it intersects C_{t-1} at most four times by the inductive hypothesis and consequently it intersects the curve $C_t = C_{t-1} \cup C'_{t-1} \cup e_t$ at most five times.

Suppose now that ℓ intersects C'_{t-1} at least twice. If ℓ passes above the rightmost point of C'_{t-1} then it intersects C'_{t-1} at most four times by the inductive hypothesis. Since ℓ passes above all points of C_{t-1}, it intersects $C'_{t-1} \cup e_t$ an even number of times, thus at most four times. If ℓ passes below the rightmost point of C'_{t-1} then it intersects C'_{t-1} at most five times by the inductive hypothesis. Since ℓ passes above all points of C_{t-1}, it intersects $C'_{t-1} \cup e_t$ an odd number of times, thus at most five times. This finishes the proof.

Consider a line ℓ containing no point of Z_t. Since ℓ intersects $\partial \mathcal{P}_t$ an even number of times, the first part of Lemma 4 implies that ℓ intersect $\partial \mathcal{P}_t$ at most six times. Lemma 2 then implies that Z_t is a 3-convex point set.

We now show that the largest 1-convex subset of Z_t contains at most $O(t) = O(\log n)$ points. We use an argument analogous to the one used by Erdős and Szekeres [6] (see also Matoušek [12, Sect. 3.1]). Every 1-convex set C of points with distinct x-coordinates is a union of a cup and a cap meeting exactly in the leftmost and the rightmost points of C. To prove the desired bound it is sufficient to show that a Devil's staircase X_t of size $n = 2^t$ contains no cup or cap having more than $t + 1 = \log{(n)} + 1$ points. A cup in X_1 contains at most two points. Due to the construction of Devil's staircase, every cup in $X_t = X_{t-1} \cup X'_{t-1}$ is either fully contained in one of the smaller Devil's staircases X_{t-1} or X'_{t-1} or it contains at most one point of X'_{t-1}. It follows by induction on t that a cup in X_t contains at most $t+1$ points. Analogously, every cap in X_t contains at most $t+1$ points. Thus, every 1-convex subset of X_t contains at most $2t = O(\log n)$ points.

Since any subset of a 3-convex point set is 3-convex [2, Lemma 2] and removing points from Z_t does not increase the size of the largest 1-convex subset, we obtain the first part of Theorem 2.

Acknowledgements. The authors would like to thank Paz Carmi for interesting discussions during the early stages of the research.

References

1. Aichholzer, O., Aurenhammer, F., Demaine, E.D., Hurtado, F., Ramos, P., Urrutia, J.: On k-convex polygons. Comput. Geom. **45**(3), 73–87 (2012)
2. Aichholzer, O., et al.: On k-convex point sets. Comput. Geom. **47**(8), 809–832 (2014)
3. Aichholzer, O., et al.: Holes in 2-convex point sets. Comput. Geom. **74**, 38–49 (2018)
4. Bárány, I., Valtr, P.: A positive fraction Erdős-Szekeres theorem. Discrete Comput. Geom. **19**(3, Special Issue), 335–342 (1998)
5. Erdős, P.: Some more problems on elementary geometry. Austral. Math. Soc. Gaz. **5**(2), 52–54 (1978)

6. Erdős, P., Szekeres, G.: A combinatorial problem in geometry. Compos. Math. **2**, 463–470 (1935)
7. Erdős, P., Szekeres, G.: On some extremum problems in elementary geometry. Ann. Univ. Sci. Budapest. Eötvös Sect. Math. **3–4**, 53–62 (1960/1961)
8. Gerken, T.: Empty convex hexagons in planar point sets. Discrete Comput. Geom. **39**(1–3), 239–272 (2008)
9. Harborth, H.: Konvexe Fünfecke in ebenen Punktmengen. Elem. Math. **33**(5), 116–118 (1978)
10. Holmsen, A.F., Mojarrad, H.N., Pach, J., Tardos, G.: Two extensions of the Erdős-Szekeres problem. Preliminary version: http://arxiv.org/abs/1710.11415 (2017)
11. Horton, J.D.: Sets with no empty convex 7-gons. Canad. Math. Bull. **26**(4), 482–484 (1983)
12. Matoušek, J.: Lectures on Discrete Geometry. Graduate Texts in Mathematics, vol. 212. Springer, New York (2002). https://doi.org/10.1007/978-1-4613-0039-7
13. Nicolás, C.M.: The empty hexagon theorem. Discrete Comput. Geom. **38**(2), 389–397 (2007)
14. Pór, A., Valtr, P.: The partitioned version of the Erdős-Szekeres theorem. Discrete Comput. Geom. **28**(4), 625–637 (2002)
15. Suk, A.: On the Erdős-Szekeres convex polygon problem. J. Am. Math. Soc. **30**(4), 1047–1053 (2017)
16. Thomson, B.S., Bruckner, J.B., Bruckner, A.M.: Elementary Real Analysis. Prentice-Hall, Upper Saddle River (2001)