Holes in 2-convex point sets*

Oswin Aichholzer† Martin Balko‡ Thomas Hackl† Alexander Pilz§ Pedro Ramos¶
Birgit Vogtenhuber† Pavel Valtr‡

Abstract

Let S be a finite set of n points in the plane in general position. A *k-hole* of S is a simple polygon with k vertices from S and no points of S in its interior. A simple polygon P is *l-convex* if no straight line intersects the interior of P in more than l connected components. Moreover, a point set S is *l-convex* if there exists an l-convex polygonalization of S.

Considering a typical Erdős-Szekeres type problem we show that every 2-convex point set of size n contains a convex hole of size $\Omega(\log n)$. This is in contrast to the well known fact that there exist general point sets of arbitrary size that do not contain a convex 7-hole. Further, we show that our bound is tight by providing a construction for 2-convex point sets with holes of size at most $O(\log n)$.

1 Introduction

Let S be a set of n points in the plane in general position, i.e., S does not contain a collinear point triple. A *k-hole* of S is a simple polygon whose k vertices are a subset of S and whose interior does not contain any point of S. Erdős [4] asked for the smallest integer $h(k)$ such that every set of $h(k)$ points in the plane contains at least one convex k-hole. Here, we consider this question for a restricted class of point sets.

A simple polygon P with boundary ∂P is *l-convex* if there exists no straight line that intersects the interior of P in more than l connected components [1]. We call a line that intersects ∂P in a finite set of at least j points a *j-stabber*; for an l-convex polygon, there cannot be a $(2l+1)$-stabber. Clearly, a convex polygon is 1-convex. In [2], the notion of l-convexity was transcribed to finite point sets. A point set S is *l-convex* if there exists a polygonalization $P(S)$ of S such that $P(S)$ is an l-convex polygon. Note that an l-convex polygon or point set is also $(l+1)$-convex. In this paper, we consider the following problem: What is the smallest number $f(k)$ such that any 2-convex point set of size $f(k)$ contains a convex k-hole?

Similar problems (for different generalizations of convexity) have also been considered, see e.g. [7, 8]. It has been shown that $h(k)$ is finite for $k \leq 6$, see e.g. [3] for details. For general point sets Horton [6] showed that there exist sets of arbitrary size that do not contain a convex 7-hole, that is, $h(7)$ is not bounded. In contrast we show that every 2-convex point set of size n contains a convex hole of size $\Omega(\log n)$, implying that $f(k)$ is bounded for any $k > 0$ (Section 3). Further, we show that our bound is tight by providing a construction for 2-convex point sets with holes of size at most $O(\log n)$ (Section 4). Due to space constraints, most proofs are omitted.

2 Properties of 2-convex polygons

We follow the definitions used in [1] and [2]. A *pocket* of a simple polygon P is a maximal chain on the boundary of P not containing any vertices of $\mathrm{CH}(P)$ except for its endpoints. For 2-convex polygons, the following is known about the structure of the pockets.

Lemma 1 ([1], Lemma 12) *Let $K = \langle p_0, \ldots, p_t \rangle$ be a pocket of a 2-convex polygon between two extreme points p_0 and p_t. Then K can be partitioned into three chains $C_1 = \langle p_0, p_1, \ldots, p_r \rangle$, $C_2 = \langle p_{r+1}, \ldots, p_s \rangle$, and $C_3 = \langle p_{s+1}, \ldots, p_t \rangle$ for $0 \leq r \leq s < t$, such that all vertices in C_1 and C_3 are convex vertices of P, while all vertices in C_2 are reflex.*

We call the segment p_0p_t the *lid* of the pocket. If C_2 is empty, the pocket consists solely of a convex hull edge. Otherwise, we call the edges p_rp_{r+1} and p_sp_{s+1} the two *inflection edges* of the pocket. Consider the (convex) polygons defined by C_1, C_2, and C_3, respectively. The next lemma follows from the proof of Lemma 12 in [2].

Lemma 2 ([2]) *The interior of a convex polygon defined by C_1, C_2, or C_3 does not intersect ∂P.*

*Research supported by OEAD project CZ 18/2015 and by project no. 7AMB15A T023 of the Ministry of Education of the Czech Republic. O.A., A.P., and B.V. supported by ESF EUROCORES programme Euro-GIGA - ComPoSe, Austrian Science Fund (FWF): I648-N18. M.B. and P.V. supported by grant GAUK 690214 and by project CE-ITI no. P202/12/G061 of the Czech Science Foundation GAČR. T.H. supported by Austrian Science Fund (FWF): P23629-N18.

†Institute of Software Technology, Graz University of Technology. [oaich|thackl|bvogt]@ist.tugraz.at.

‡Department of Applied Mathematics and Institute for Theoretical Computer Science (ITI), Charles University, [balko|valtr]@kam.mff.cuni.cz.

§Institute of Theoretical Computer Science, ETH Zürich, alexander.pilz@inf.ethz.ch.

¶Departamento de Física y Matemáticas, Universidad de Alcalá, pedro.ramos@uah.es.

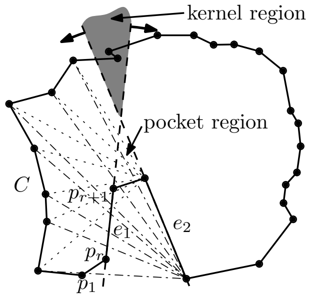

Figure 1: The order of the vertices defined by the inflection edges of a pocket ([2, Figure 9], relabeled). The gray wedge is the kernel region.

Lemma 3 ([2], Lemma 10) *Let P be a 2-convex polygon and let e_1 and e_2 be the inflection edges of a pocket K directed from the convex to the reflex vertex, with the vertices defined as in Lemma 1. Without loss of generality, p_r is left of e_2, i.e., $e_1 = p_r p_{r+1}$ and $e_2 = p_{s+1} p_s$. Let C be the part of ∂P defined by the vertices that are to the left of e_2 and not part of the pocket (starting at p_1, the left endpoint of the lid of K). Then the order of the points in C along ∂P is the same as the radial order around any point p on e_2. An analogous statement holds for any point on e_1 and the points of ∂P to the right of e_1.*

See Figure 1 for an illustration (taken from [2, Figure 9]). The *kernel region* of the pocket K with non-empty C_2 is the region that is to the left of e_1, to the right of e_2, and, if $r+1 \neq s$, to the left of $p_{r+1}p_s$. Observe that, for a star-shaped 2-convex polygon, the kernel of the polygon is the intersection of the kernel regions of all the pockets.

3 The lower bound

Let S be a 2-convex point set in the plane in general position and let P be a 2-convex polygon that is a polygonalization of S. In this section, we prove the following.

Theorem 4 *Every 2-convex point set of size n contains a convex k-hole for $k \in \Omega(\log n)$.*

Let us first sketch the proof: If P has a large pocket, Lemma 2 implies the existence of a large k-hole. When P has no large pocket, we will use Lemma 5 to find a large set $Q \subset S$ of points in convex position. If Q forms a hole in S, we are done. Finally, if Q does not form a hole in S, we will use Lemma 7 and Lemma 10 to find a big enough convex hole.

Lemma 5 *Let m be the size of the largest pocket in S. Then there exists a point p (probably not in S) s.t. there is a sequence σ of $\lceil \frac{n}{3m} \rceil - 1$ points of S that are separated by a line from p, and their order around p matches the order along ∂P, where they appear consecutively.*

Proof. Suppose first that P is star-shaped and let $p \notin S$ be a point in the kernel of P. Consider any half-plane H defined by a line through p that contains $\lceil \frac{n}{2} \rceil$ points of S. The radial order of the points in $S \cap H$ around p is the same as the order along P.

Suppose now that P is not star-shaped, i.e., its kernel is empty. The kernel of P is determined by the intersection of the kernel regions of all the pockets. A non-empty kernel region is the intersection of two half-planes defined by inflection edges (as discussed in [2]). By Helly's theorem [5], we know that, if the kernel of P is empty, there exists a triple of inflection edges such that the intersection of the half-planes (partly) defining their kernel regions is empty. (Similar to [2, Lemma 11].) This means that there exists at least one inflection edge e of a pocket K such that the open half-plane H defined by e that contains K also contains at least $\lceil n/3 \rceil$ points of S. Due to Lemma 3, the radial order of the points in $S \cap H$ and not on K around any point p on e is the same as their order along ∂P. Hence, there is a sequence of at least $\left\lceil \frac{\lceil n/3 \rceil - (m-2)}{m-2} \right\rceil \geq \lceil \frac{n}{3m} \rceil - 1$ points along ∂P that are consecutive in the order of all points of S around p (not containing a point of K and linearly separated from p by the supporting line of an edge of K). ☐

In the previous proof, when P is star-shaped, the point p was not part of S. However, we can define a point set S' consisting of p and $S \cap H$. Then, it is easy to see that there is a 2-convex polygonization P' of S' in which p sees all the points in the order as they appear along $\partial P'$. Any convex k-hole of S' is a convex $(k-1)$-hole or a convex k-hole of S. Thus, for simplicity, we will assume that $p \in S$.

Let $\phi \subseteq S^3$ be the ternary relation representing the cyclic order of the vertices of P as they appear on the boundary of P traversed in counterclockwise direction. That is, a triple (u, v, w) of points of S is in ϕ if we can trace u, v, w in this order along the boundary of P in counterclockwise direction. For $u, w \in S$, a *(closed) interval* $[u, w]$ *from u to w* in ϕ is the set $\{v \in S : (u, v, w) \in \phi\} \cup \{u, w\}$. Note that the intervals $[u, w]$ and $[w, u]$ are in general distinct. Each point $u \in S$ defines a linear order $<_u$ on $S \setminus \{u\}$ where $x <_u y$ if and only if $(x, y, u) \in \phi$.

Note that vertices of a pocket $K = \langle p_0, \ldots, p_t \rangle$ of P induce a closed interval $[p_0, p_t]$ in ϕ. Consequently, ϕ induces a cyclic order of pockets of P. We choose an arbitrary pocket K_0 of P and use $K_0, \ldots, K_{m-1}$ to denote this cyclic order where m is the number of pockets of P. In the rest of the section, the indices of pockets are always taken modulo m.

For $r, s \in \{0, \ldots, m-1\}$, we use $[K_r, K_s]$ to denote the interval consisting of pockets $K_r, K_{r+1}, \ldots, K_s$.

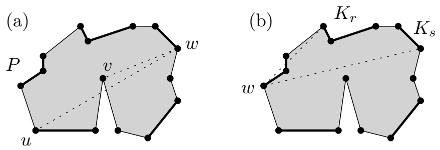

Figure 2: (a) An example of a reversed triple (u, v, w). (b) The point w controls the interval $[K_r, K_s]$.

The *length* of $[K_r, K_s]$ is the number of pockets in $[K_r, K_s]$. A *subinterval* of $[K_r, K_s]$ is any interval that can be obtained from $[K_r, K_s]$ by deleting the first i and the last j consecutive pockets of $[K_r, K_s]$ for some $i, j \in \mathbb{N}_0$.

We say that a triple $(u, v, w) \in \phi$ is *reversed* if the triangle with the vertices u, v, w traced in this order is oriented clockwise.

For an interval $[K_r, K_s]$, a point v from $S \setminus (\cup_{i=r-1}^{s+1} K_i)$ *controls* $[K_r, K_s]$ if the following conditions are satisfied:

(i) There is no reversed triple (x, y, v) with x and y contained in distinct pockets of $[K_r, K_s]$,

(ii) $\mathrm{CH}(\cup_{i=r}^{s} K_i)$ contains no point of $S \setminus (\cup_{i=r}^{s} K_i)$,

(iii) $\mathrm{CH}(\cup_{i=r}^{s} K_i \cup \{v\})$ contains no point of $S \setminus (\cup_{i=r}^{s} K_i)$ except of vertices of pockets containing v.

We note that Condition (i) especially implies that there is no reversed triple (x, y, v) with x and y being vertices of pockets in $[K_r, K_s]$ and x or y being a convex hull vertex. Hence, if v controls $[K_r, K_s]$, then v also controls every subinterval of $[K_r, K_s]$. Further, Condition (i) implies that v is linearly separable from $[K_r, K_s]$.

Lemma 6 *Let (u, v, w) be a reversed triple of points in S and let ab be the lid of the pocket K of v s.t. $(a, v, b) \in \phi$. If $\overline{uw}$ separates v from ab, then the order $<_v$ is the same as the radial order around v for $[u, a]$ and for $[b, w]$.*

Proof. We prove the statement for $[u, a]$, as the argument for $[b, w]$ follows by symmetry. Let C be the part of ∂P defined by the interval $[u, a]$. Since $\overline{uw}$ separates v from ab and thus intersects K twice, its only intersection with C is at u. Hence, any line through v crossing C has exactly one ray starting at v crossing C. Suppose there exists a line ℓ through v s.t. the ray r crossing C has a crossing with C where it enters P. We claim that a perturbation of ℓ is a 6-stabber of P, contradicting 2-convexity. Let r' be the complement of r on ℓ.

Suppose first that r enters the interior of P at v. Then r intersects ∂P in at least three points other than v. Since ab is separated from v by $\overline{uw}$, r' crosses ∂P in a point not on the pocket K. Thus, if r' leaves P at v, then ℓ is a 6-stabber. If r' does not leave P at v, then ℓ supports ∂P at v, in which case there is a perturbation of ℓ that is a 6-stabber.

Suppose now that r leaves P at v. Since ab is an edge of the convex hull of S and r crosses C, r cannot cross ab. Hence, it enters P again at the pocket K, implying that there are at least four points other than v where r crosses ∂P. The fact that r' intersects ∂P in a point not on C makes ℓ a 6-stabber.

Therefore, there is no ray starting at v entering P at C, which completes the proof. ☐

Lemma 7 *Let K_i, K_j, and K_l be pockets in a sequence of pockets that is controlled by a point $p \in S$. Let (u, v, w) be a reversed triple of points from S such that u, v, and w are contained in K_i, K_j, and K_l, respectively. Then v controls the intervals $[K_{i+1}, K_{j-2}]$ and $[K_{j+2}, K_{l-1}]$, provided that $\overline{uw}$ separates v from the endpoints of K_j.*

Lemma 8 *Let $[K_r, K_{r+3d+3}]$ be an interval controlled by some point $p \in S$. Then there is a subinterval of $[K_r, K_{r+3d+3}]$ of length d controlled by a point of a pocket that is contained in $[K_r, K_{r+3d+3}]$.*

Let H be a hole in S. If H contains at most one point from every pocket of S, then H is *transversal*. We say that an interval $[K_r, K_s]$ of pockets *contains a hole H* if every vertex of H is contained in some pocket of the interval $[K_r, K_s]$. We call a hole H *nice*, if there is no reversed triple of vertices of H.

Lemma 9 *For every integer $k \geq 2$, let $[K_r, K_s]$ be an interval of pockets that contains a nice convex transversal $(k-1)$-hole. If a point p of S controls $[K_r, K_s]$, then there is a pocket K containing p such that the intervals $[K_r, K]$ and $[K, K_s]$ contain a nice convex transversal k-hole.*

First, we prove the following lemma and then we show how it implies Theorem 4.

Lemma 10 *For every positive integer k and every interval $[K_r, K_s]$ of pockets, if the length of $[K_r, K_s]$ is at least $2 \cdot 3^k - 2$ and $[K_r, K_s]$ is controlled by some point of S, then $[K_r, K_s]$ contains a nice convex transversal k-hole.*

Proof. We prove the lemma by induction on k. For $k = 1$, the lemma follows from the fact that every interval of length 1 contains a 1-hole. For the induction step, let $k > 1$. For $d := 2 \cdot 3^{k-1} - 2$, let $[K_r, K_s]$ be the interval of length at least $3d + 4 = 2 \cdot 3^k - 2$ that is controlled from some point of S. By Lemma 8, there is a point q contained in a pocket from $[K_r, K_s]$ such that q controls a subinterval $[K_i, K_j]$ of $[K_r, K_s]$ with length at least d. Using the induction hypothesis, it follows that $[K_i, K_j]$ contains a nice convex transversal

$(k-1)$-hole H. By Lemma 9, the hole H can be extended to a nice convex transversal k-hole contained in $[K_r, K_s]$. $\square$

Proof of Theorem 4. To show that Lemma 10 implies Theorem 4, we prove that in every 2-convex point set S of size n there is a convex k-hole for $k \geq \log n/3$, or we have an interval of length $\Omega(n/\log^3 n)$ that is controlled by a point from S. In the latter case we then apply Lemma 10 and obtain a convex k-hole with $k \geq c \log n$ for an absolute constant $c > 0$.

First, assume that there is a pocket $K = \langle p_0, \ldots, p_t \rangle$ in P with $t \geq \log n$ in P. By Lemma 1, the pocket K can be partitioned into three chains $C_1 = \langle p_0, p_1, \ldots, p_r \rangle$, $C_2 = \langle p_{r+1}, \ldots, p_s \rangle$, and $C_3 = \langle p_{s+1}, \ldots, p_t \rangle$ for $0 \leq r \leq s < t$, such that all vertices in C_1 and C_3 are convex in P, while all vertices in C_2 are reflex. Since K contains at least $\log n$ vertices, at least one of the chains C_1, C_2, and C_3 contains at least $\log n/3$ vertices. For some $i \in \{1, 2, 3\}$, let C_i be such a chain. By Lemma 2, the vertices of C_i are vertices of a convex k-hole for $k \geq \log n/3$. See Figure 3 (a).

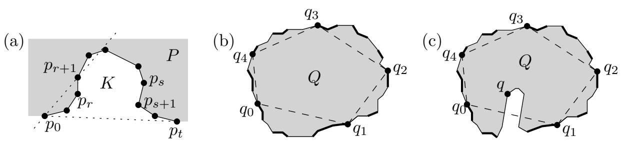

Figure 3: (a) A large pocket gives a large hole. (b) If no point of S interferes, then Q is a hole. (c) If there is a point inside Q, then we use Lemma 7 and apply Lemma 10.

In the rest of the proof we thus assume that every pocket of P contains less than $\log n$ vertices. In particular, there are more than $n/\log n$ pockets in P and $\mathrm{CH}(S)$ has more than $n/\log n$ vertices. By Lemma 5, there are at least $m := \left\lceil \frac{n}{3 \log n} \right\rceil - 1$ points that are "controlled" by a point p (that is not necessarily in S). We call these points the *initial interval*. However, by the discussion after Lemma 5 we can assume for the following that $p \in S$. Let $q_0, \ldots, q_{\log n - 1}$ be vertices of $\mathrm{CH}(S)$ traced in counterclockwise direction along the boundary of P in the initial interval such that the points in each interval $[q_i, q_{i+1}]$ for $i = 0, \ldots, \log n - 1$ (indices taken modulo $\log n$) form at least $m/\log^2 n$ pockets. Clearly, if the polygon Q with the vertices $q_0, \ldots, q_{\log n - 1}$ is a hole, then we are done; see Figure 3 (b). Otherwise there is a point q in the interior of Q and we have a reversed triple (q_i, q, q_j) for some $i, j \in \{0, \ldots, \log n - 1\}$. Let K, K', and K'' be pockets containing q_i, q, and q_j, respectively. The endpoints of K' are separated from q by $\overline{q_i q_j}$, as q_i and q_j are vertices of $\mathrm{CH}(S)$; See Figure 3 (c). By Lemma 7, the point q controls the interval of pockets that are between K and K' and between K' and K''. From the choice of Q, at least one of these intervals has length at least $m/(2 \log^2 n) = \Omega(n/\log^3 n)$. $\square$

4 An upper-bound construction

Theorem 11 *For any n there exists a 2-convex point set S of size n such that all convex holes it contains have size $O(\log n)$.*

Proof. The set is constructed recursively, following the idea shown in Figure 4. We define $S_i = L_i \cup R_i \cup \{c_i\}$, where L_i and R_i are flattened enough copies of S_{i-1}. For $i = 0$, we set $L_0 = R_0 = \emptyset$.

An empty convex hole K intersecting R_i cannot intersect both the left and right part of L_i, and this is true for every level in the recursion. Of course, an analogus statement is true if K intersects L_i. Therefore, $|K| = O(\log n)$. $\square$

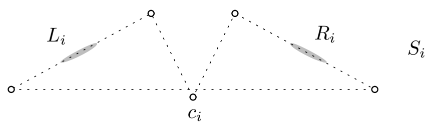

Figure 4: Recursive operation for the construction of an upper bound example.

Acknowledgments. This work was initiated during the ComPoSe Workshop on Algorithms using the Point Set Order Type held in March/April 2014 in Ratsch, Austria.

References

[1] O. Aichholzer, F. Aurenhammer, E. D. Demaine, F. Hurtado, P. Ramos, and J. Urrutia. On k-convex polygons. *Comput. Geom.*, 45(3):73–87, 2012.

[2] O. Aichholzer, F. Aurenhammer, T. Hackl, F. Hurtado, A. Pilz, P. Ramos, J. Urrutia, P. Valtr, and B. Vogtenhuber. On k-convex point sets. *Comput. Geom.*, 47(8):809–832, 2014.

[3] O. Aichholzer, R. Fabila-Monroy, H. Gonzalez-Aguilar, T. Hackl, M. A. Heredia, C. Huemer, J. Urrutia, P. Valtr, and B. Vogtenhuber. On k-gons and k-holes in point sets. *Comput. Geom.*, 48(7):528–537, 2015.

[4] P. Erdős. Some more problems on elementary geometry. *Austral. Math. Soc. Gaz.*, 5:52–54, 1978.

[5] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresber. Deutsch. Math.-Verein.*, 32:175–176, 1923. In German.

[6] J. D. Horton. Sets with no empty convex 7-gons. *Canad. Math. Bull.*, 26(4):482–484, 1983.

[7] P. Valtr. A sufficient condition for the existence of large empty convex polygons. *Discrete Comput. Geom.*, 28(4):671–682, 2002.

[8] P. Valtr, G. Lippner, and Gy. Károlyi. Empty convex polygons in almost convex sets. *Period. Math. Hungar.*, 55(2):121–127, 2007.