On k-Convex Polygons[☆]

Oswin Aichholzer[a], Franz Aurenhammer[a], Erik D. Demaine[b], Ferran Hurtado[c], Pedro Ramos[d,*], Jorge Urrutia[e]

[a] *Institute for Software Technology, University of Technology, Graz, Austria,* {oaich,thackl,apilz,bvogt}@ist.tugraz.at

[b] *Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, USA,* edemaine@mit.edu

[c] *Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Spain,* Ferran.Hurtado@upc.edu

[d] *Departamento de Matemáticas, Universidad de Alcalá, Madrid, Spain,* pedro.ramos@uah.es

[e] *Instituto de Matemáticas, Universidad Nacional Autònoma de Mèxico,* urrutia@matem.unam.mx

Abstract

We introduce a notion of k-convexity and explore polygons in the plane that have this property. Polygons which are k-convex can be triangulated with fast yet simple algorithms. However, recognizing them in general is a 3SUM-hard problem. We give a characterization of 2-convex polygons, a particularly interesting class, and show how to recognize them in $O(n \log n)$ time. A description of their shape is given as well, which leads to Erdős-Szekeres type results regarding subconfigurations of their vertex sets. Finally, we introduce the concept of generalized geometric permutations, and show that their number can be exponential in the number of 2-convex objects considered.

Keywords: Convexity, Visibility, Transversal theory
2000 MSC: 52A35, 52C99

[☆] Partially supported by the FWF Joint Research Program 'Industrial Geometry' S9205-N12, Projects MEC MTM2006-01267, DURSI 2005SGR00692, Project Gen. Cat. DGR 2009SGR1040, DGR 2009SGR-1040, MEC MTM2009-07242, MEC MTM2008-04699-C03-02, and the bilateral Spain-Austria program 'Acciones Integradas' ES 01/2008 and HU2007-0017.

[*] Corresponding author

1. Introduction

The notion of convexity is central in geometry. As such, it has been generalized in many ways and for different reasons. In this paper we consider a simple and intuitive generalization of convexity, which to the best of our knowledge has not been worked on. It leads to an appealing class of polygons in the plane with interesting structural and algorithmic properties.

A set in $\mathbb{R}^d$ is convex if its intersection with every straight line is connected. This definition may be relaxed to directional convexity or *D-convexity* [16, 24], by considering only lines parallel to one out of a (possibly infinite) set D of vectors. A special case is *ortho-convexity* [29], where only horizontal and vertical lines are allowed. For any fixed D, the family of D-convex sets is closed under intersection, and thus can be treated in a systematic way using the notion of *semi-convex* spaces [31], which is sometimes appropriate for investigating visibility issues. The *D-convex hull* of a set M is the intersection of all D-convex sets that contain M. If D is a finite set, this definition of a convex hull may lead to an undesirably sparse structure—an effect which can be remedied by using a stronger, functional (rather than set-theoretic) concept of D-convexity [24].

***k**-Convex Sets.* We consider a different generalization of convexity for 2-*dimensional* sets: We say that a set M is 2-dimensional if for every $p \in M$ there exists a set M' homeomorph to a closed disk and such that $p \in M' \subset M$. All sets we will consider in this paper regarding k-convexity are 2-dimensional compact sets in the Euclidean plane and therefore don't have isolated points or 1-dimensional components. We say that M is *k-convex (with respect to transversal lines)* if there exists no straight line that intersects the interior of M, this is, $M \setminus \partial M$, in more than k connected components. Throughout the paper we will use the term *k-convex*, for short[1]. Note that 1-convexity refers to convexity in its standard meaning.

It is also often said that a set is convex if every two of its points *see each other*. To reformulate k-convexity in terms of visibility let us introduce the concept of *counting crossings*. Let s be a line or an open line segment. We count the number of crossings of s with ∂M in the following way: Let t be a connected component of $s \cap \partial M$; then t is either a point or a segment. The

[1] We face notational ambiguity. The term 'k-convex' has, maybe not surprisingly, been used in different settings, namely, for functions [27], for graphs [5], and for discrete point sets [35, 21]. Also, the concept of k-point convexity [34] has later been called k-convexity in [7].

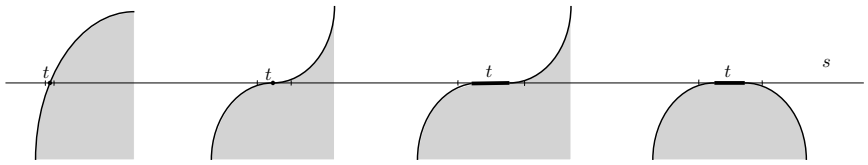

a) 1 crossing b) 1 crossing c) 1 crossing d) 0 crossings

Figure 1: Different cases for the counting of crossings.

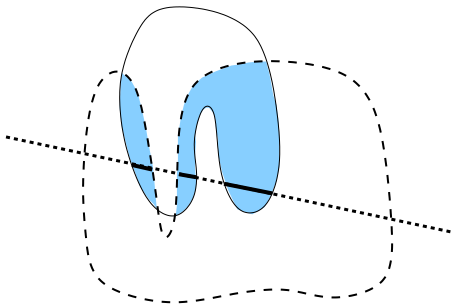

Figure 2: Intersecting two 2-convex sets

component t is counted as one crossing if and only if in every neighborhood of t in s there are points from both the interior and the exterior of M. The first case accounts for the simplest case of crossing, as happens with two segments that share exactly one point interior to the two of them (Figure 1.a,b), while the second case occurs when s supports ∂M along a segment that behaves as an inflection (Figure 1.c). Using the concept of crossing we see that (compact and 2-dimensional) k-convex sets are precisely those whose boundary can be crossed by a line at most $2k$ times.

Now, we can express k-convexity in terms of visibility: call two points $x, y \in M$ *k-visible* if the number of crossings between the open segment xy and ∂M is at most $2(k-1)$. Thus a set is k-convex if and only if any two of its points are mutually k-visible. We find interesting this way of articulating the notion, as applications of this concept may arise from placement problems for modems that have the capacity of transmitting through a fixed number of walls [2, 14].

Unlike directional convexity, k-convexity fails to show the intersection property: The intersection of k-convex sets is not k-convex in general (for fixed k). Figure 2 gives an example. For $k \geq 2$, a k-convex set M may be

disconnected, or if connected, its boundary may be disconnected. In this paper, we will restrict attention (with an exception in Section 4) to simply connected sets in two dimensions, namely, simple polygons in the plane.

There are two notions of planar convexity that appear to be close to ours. One is *k-point* convexity [34, 7]: A closed connected set $M \subset \mathbb{R}^2$ is k-point convex if for any k points in M, at least one of the line segments they span is contained in M. Thus 2-point convex sets are precisely the convex sets. The other is *k-link convexity* [23]: A simple polygon P is k-link convex if, for any two points in P, the geodesic path connecting them inside P consists of at most k edges. The 1-link convex polygons are just the convex polygons. While there is a relation between k-convexity and the former concept (as we will show in Section 2), the latter concept is totally unrelated.

We will study basic properties of k-convex polygons, in comparison to existing polygon classes and convexity concepts in Section 2. This offers an alternative to the approach in [4] to define 'realistic' polygons as those being guardable (visible) by at most k guards. We prove that given a simple polygon P, the problem of finding the smallest k such that P is k-convex (equivalently, to find the stabbing number of P) is 3SUM-hard. On the other hand, a recognition algorithm that runs in $O(n^2)$ time for a polygon with n vertices is easy to obtain. Interestingly, k-convex polygons can be triangulated, by a quite simple method, in $O(n \log k)$ time. An $O(nk)$ time complexity is achieved in [4] for k-guardable polygons.

The first nontrivial value, $k = 2$, deserves particular attention. Already in this case, a novel class of polygons is obtained. A characterization of 2-convex polygons is given in Section 3. It leads to an $O(n \log n)$ time algorithm for recognizing such polygons. Note that 2-convex polygons add to the list of special classes of polygons [12, 4] that allow for simple $O(n)$ time triangulation methods. We also provide a qualitative description of their shape, which implies an Erdős-Szekeres type result, namely, that every 2-convex polygon with n vertices contains a subset of at least $\sqrt{n}$ vertices in convex position, and that its vertex set can be decomposed into at most $2\sqrt{2n}$ subsets in convex position.

In Section 4, we turn our attention to general k-convex sets. We give observations on the union and intersection properties of such sets, and elaborate on an attempt to generalize the notion of geometric permutations from convex sets to k-convex sets. In contrast to the $O(m)$ bound in [11] on the number of geometric permutations of m convex sets, it turns out that the number of generalized geometric permutations can be exponential in m, al-

ready for 2-convex sets. For 2-convex polygons, the maximum number of generalized geometric permutations is $\Theta(n^2)$, if n denotes the total number of their vertices.

Various open questions are raised by the proposed concept of k-convexity. We list those which seem most interesting to us, along with a brief discussion of our results, in Section 5.

2. k-Convex Polygons

2.1. Basic properties

We start with exploring some basic properties of k-convex polygons, and compare them to existing polygon classes and related concepts.

Let P be a simple polygon, and denote by n the number of vertices of P. Here and hereafter we assume that P does not have two consecutive edges that are collinear. We define the *stabbing number* of a polygon as the largest possible number of crossings between the boundary of the polygon and a straight line[2]. Therefore, a polygon is k-convex if and only if its stabbing number is at most $2k$, and our observations on k-convexity can be formulated for polygons in terms of their stabbing numbers.

The kernel of a simple polygon P is the set of points that see all the polygon. Its generalization to k-convexity shows that 2-convexity is already significantly more complex than standard convexity. The *k-kernel of P*, denoted as $M_k(P)$, is the set of points from which the entire polygon P is k-visible. Note that P is k-convex if and only if $P = M_k$. While M_1 is known to be a convex set which is computable in $O(n)$ time [22], M_2 may have $\Omega(n^2)$ complexity: If we consider the 'spike' in Figure 3(a), the wedge between the lines pq and pr is not part of M_2. If we arrange such spikes along the boundary of a rectangle, as in Figure 3(b), we get a quadratic number of disconnected areas which are part of the 2-kernel. Therefore, any algorithm computing M_2, or trying to check 2-convexity via the comparison of P and M_2 would have $\Omega(n^2)$ complexity in the worst case.

There is also no immediate relation to *star-shaped* polygons, i.e., polygons P with $M_1 \neq \emptyset$. Figure 4 shows a polygon on the left hand side which is star-shaped but only $\frac{n}{2}$-convex. On the right hand side, we see a polygon which is 2-convex but not star-shaped. Visually, 2-convexity

[2]This definition is very close to the concepts of *crossing number* and *stabbing number* of a set of segments defined in [36, 15], but slightly different from. We ellaborate more on this difference at the end of Section 2.2.

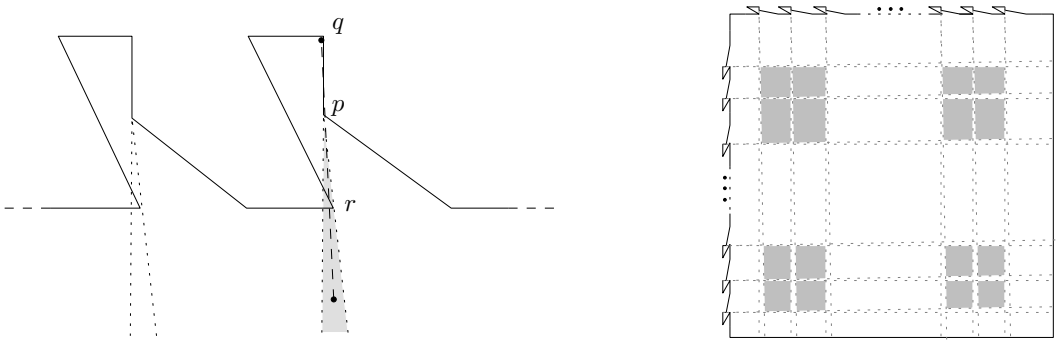

(a) (b)

Figure 3: (a) Points in the gray region are not in the 2-kernel. (b) 2-kernel with a quadratic number of connected components.

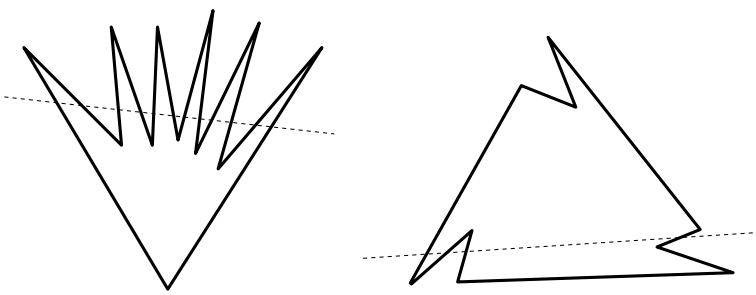

Figure 4: Star-shaped versus 2-convex

seems to be closer to convexity than is star-shapedness. Note that cutting a 2-convex polygon with any straight line leaves (at most) three parts, each being 2-convex itself. This is not true in general for star-shaped polygons.

There is no relation between 2-convexity and link distance [23]: The link distance in a 2-convex polygon may well be $\Theta(n)$ and, conversely, a polygon which is 2-link convex (such that any two of its points are at link distance 2 or less) may fail to be k-convex for sublinear k. The star-shaped polygon in Figure 4 (left) is an example.

There is, interestingly, a relation to *k-point* convexity as defined in [34]. A polygon P is called k-point convex if for any k points $p_1, \ldots, p_k$ in P, at least one of the closed segments $p_i p_j$ belongs to P. Every k-point convex polygon P is $(k-1)$-convex. To verify this, we prove that if P is not

$(k-1)$-convex, then P is not k-point convex: Because P is not $(k-1)$-convex, there exists a line L which intersects $P \setminus \partial P$ in at least k components. If we select a point in each component, it is clear that none of the segments defined by them is inside P, and therefore P is not k-point convex. However, no implication exists in the other direction. For example, the 2-convex polygon in Figure 5 fails to be k-point convex for $k < \frac{n}{3}$. Also, any k-point convex polygon can be expressed as the union of m convex polygons, where m depends (exponentially) on k but is independent of the polygon size n; see [7]. Such a property is not shared by k-convex polygons, as can be seen from the 2-convex polygon in Figure 5.

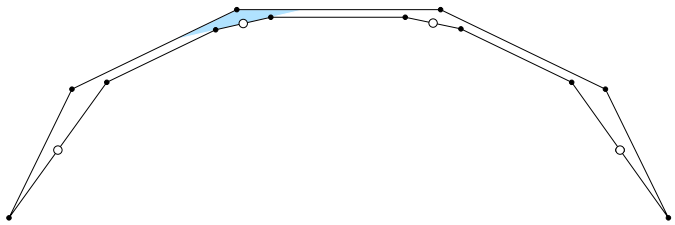

Figure 5: Every white dot requires a different guard

The class of k-convex polygons also differs from the class of *k-guardable* polygons defined in [4]. It is known that any simple polygon with n vertices can be guarded with at most $\lfloor \frac{n}{3} \rfloor$ guards [26, 33]. The example in Figure 5 shows that this number of guards can be already necessary for 2-convex polygons. This is one of the reasons why most tools developed to study guarding problems of polygons are not very useful in the study of modem illumination problems [2].

Pseudo-triangles are polygons with exactly three convex vertices, joined by three reflex side chains. Any pseudo-triangle is 2-convex: If a straight line crosses a side chain twice, then it can cross each of the remaining two side chains at most once (Figure 6, left). That is, the stabbing number of a pseudo-triangle is four or less. In the same way as a triangulation defines a partition of the underlying domain into convex polygons, any pseudo-triangulation [30] or any pseudo-convex decomposition [3] gives a partition into 2-convex polygons. It is an open problem (see Problem 5 in the final section) whether it is possible to subdivide a polygon with n vertices into a sublinear number of 2-convex polygons. If Steiner points are disallowed, then a 2-convex partition may have to consist of $\Theta(n)$ parts; see Figure 6 (right).

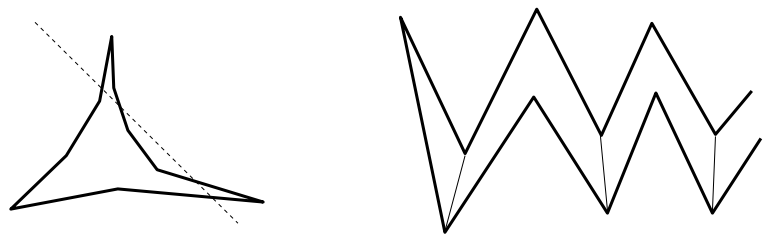

Figure 6: Pseudo-triangle and 2-convex partition

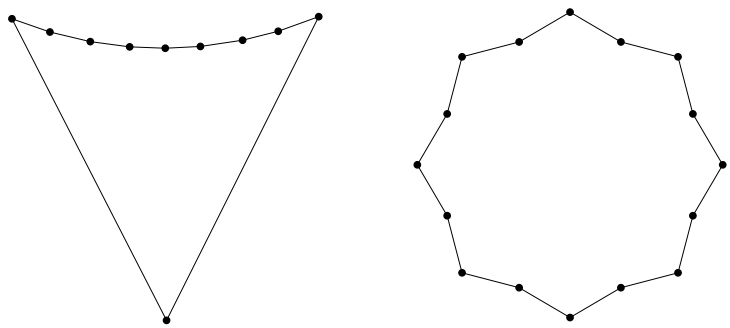

Figure 7: A polygon with many convex pieces (left) and many pockets (right).

Let us mention that the following natural questions for k-convex polygons are easy to answer: (i) decompose a k-convex polygon into few convex pieces (in terms of k), (ii) give a bound on the number of pieces of the convex hull minus the polygon. In both cases the answer is unrelated to k: for the first one, observe the polygon in Figure 7 (left). It is 2-convex yet requires $n-2$ convex pieces, which is obviously tight because every polygon can be triangulated. For the second one, the polygon in Figure 7 (right) is also 2-convex and has $\frac{n}{2}$ "pockets", which is tight because two consecutive points on the convex hull are at least two edges apart if they are not adjacent on the polygon boundary.

2.2. Recognition complexity

Let us start our considerations by observing that the stabbing number of a polygon with n vertices can be easily found in $O(n^2)$ time, as follows. The standard duality transform maps each edge of the polygon to a double wedge consisting of lines through a common point, not including the vertical; the two lines that bound this dual wedge correspond to the endpoints of the

primal segment. In the primal, any point inside a wedge is a line that stabs the segment. Therefore, the primal line that would stab most segments is in the dual a point that belongs to as many double wedges as possible—a maximum depth problem that can easily be solved by constructing the arrangement and then traversing its cells.

The obtained $O(n^2)$ time bound is essentially tight, as in this section we prove that finding the stabbing number of a polygon, or, equivalently, finding the smallest k for which the polygon is k-convex, is a 3SUM-hard problem. This family of problems is widely believed to have an $\Omega(n^2)$ lower bound for the worst case runtime [18, 20]. We start by giving the following result, which follows directly from Theorem 4.1 in [18].

Lemma 1. *For every integer a, let us consider the point $p_a = (a, a^3)$ on the cubic $y = x^3$. Then if a, b, c are distinct integers, p_a, p_b, p_c are collinear if and only if $a + b + c = 0$.*

Proof. The points p_a, p_b, p_c are collinear if and only if the determinant

$$\begin{vmatrix} a & a^3 & 1 \\ b & b^3 & 1 \\ c & c^3 & 1 \end{vmatrix} = (b-a)(c-a)(c-b)(a+b+c)$$

vanishes, which, the numbers being different, happens exactly when $a + b + c = 0$. □

We next show that the points p_x with $x \in \mathbb{Z}$ on the cubic $y = x^3$ can be replaced by infinitesimally small vertical segments S_x with upper endpoint p_x, such that three of them can be stabbed by a single line if and only if their three upper endpoints are collinear.

Lemma 2. *Let m, a, b, c, M be five integers such that $m < a < b < c < M$. Let $\varepsilon = \frac{1}{6(M-m)}$ and let s_t be the (vertical) segment with endpoints $p_t = (t, t^3)$ and $p'_t = (t, t^3 - \varepsilon)$. Then s_a, s_b and s_c can be stabbed by a single line if and only if p_a, p_b, and p_c are collinear.*

Proof. Assume that the points p_a, p_b, p_c are not collinear, and let us take three points $q_a = (a, a^3 - \varepsilon_1), q_b = (b, b^3 - \varepsilon_2), q_c = (c, c^3 - \varepsilon_3)$, with $0 \leq \varepsilon_1, \varepsilon_2, \varepsilon_3 \leq \varepsilon$, i.e., three points on the segments s_a, s_b and s_c. The points q_a, q_b, q_c would be collinear if and only if the determinant

$$\begin{vmatrix} a & a^3 - \varepsilon_1 & 1 \\ b & b^3 - \varepsilon_2 & 1 \\ c & c^3 - \varepsilon_3 & 1 \end{vmatrix} = \begin{vmatrix} a & a^3 & 1 \\ b & b^3 & 1 \\ c & c^3 & 1 \end{vmatrix} + \begin{vmatrix} a & -\varepsilon_1 & 1 \\ b & -\varepsilon_2 & 1 \\ c & -\varepsilon_3 & 1 \end{vmatrix} =$$

$$= \underbrace{(b-a)(c-a)(c-b)(a+b+c)}_{z} + \underbrace{\varepsilon_1(b-c) - \varepsilon_2(a-c) + \varepsilon_3(a-b)}_{\delta}$$

is 0, but this is impossible because by Lemma 1, z is an integer different from 0, which cannot become 0 by the addition of δ, because

$$|\varepsilon_1(b-c)-\varepsilon_2(a-c)+\varepsilon_1(a-b)| \leq |\varepsilon_1(b-c)|+|\varepsilon_2(a-c)|+|\varepsilon_3(a-b)| \leq \varepsilon \cdot 3(M-m) = \frac{1}{2}.$$

□

Theorem 3. *The problem of finding the stabbing number of a polygon is 3SUM-hard.*

Proof. The 3SUM problem is defined as follows: Given a set S of n integers, do there exist three elements $a, b, c \in S$ such that $a+b+c = 0$? We will prove below that this problem can be reduced in $O(n \log n)$ time to the problem of computing the stabbing number of an n-gon. In other words, using the notation in [18],

$$\text{3SUM} \lll_{O(n \log n)} \text{stabbing number of a polygon.}$$

Let $x_1, \ldots, x_n$ be the input integers. We proceed to the reduction by steps.

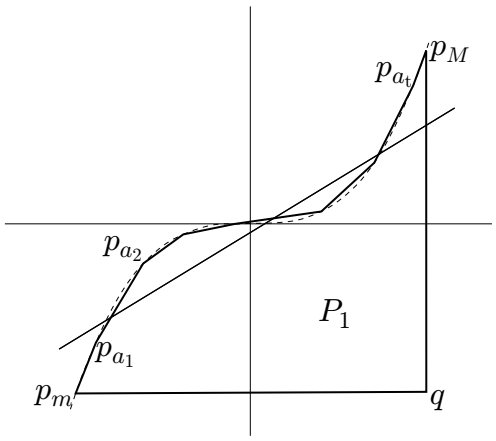

Figure 8: Polygon P_1 with vertices on the cubic $y = x^3$ (dashed). The scale of the axis is not 1 : 1, to make the figure visible.

Step 1. Sort the input numbers; let $y_1 \leq y_2 \leq \cdots \leq y_n$ be the resulting list, $\mathcal{L}$. This step is done in $O(n \log n)$ time.

Step 2. If 0 appears thrice in $\mathcal{L}$, exit with a sum of three numbers in the input being 0, otherwise continue. The step is completed in linear time.

Step 3. If $a \neq 0$ appears at least twice in $\mathcal{L}$, check whether $-2a \in \mathcal{L}$. If so, exit with a sum of three numbers in the input being 0, otherwise continue. The step is completed in $O(n \log n)$ time, as binary search can be used in the sorted list (in fact, $O(n)$ time is sufficient by scanning the list from left to right and from right to left, in a coordinated simultaneous advance).

Step 4. Remove multiples from the list $\mathcal{L}$ so that each number appears exactly once. This requires linear time. Let $a_1 < a_2 < \cdots < a_t$ (where $t \leq n$) be these numbers.

Step 5. Define $m = a_1 - 1$, $M = a_t + 1$, and $q = (M, m^3)$. Now let us consider the polygon P_1 whose vertices, described clockwise, are $p_m p_{a_1} p_{a_2} \ldots p_{a_t} p_M q$, where $p_x = (x, x^3)$, as in the preceding lemma (Figure 8). Observe that the stabbing number of P_1 is 4, and that the polygon can be constructed in $O(n)$ time.

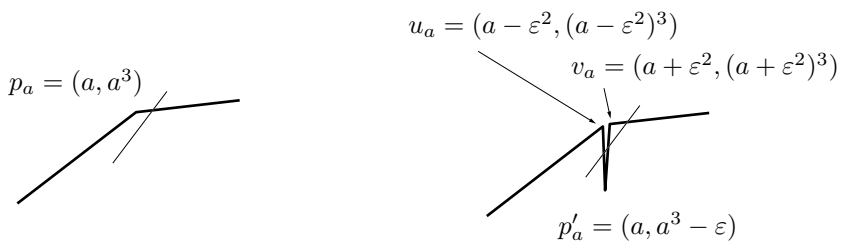

Figure 9: Polygon P_2 with vertices on the cubic $y = x^3$ replaced by infinitesimal vertical slots.

Step 6. Next we modify the polygon P_1 to become the polygon P_2 whose vertices, described clockwise, are $p_m u_{a_1} p'_{a_1} v_{a_1} u_{a_2} p'_{a_2} v_{a_2} \ldots u_{a_t} p'_{a_t} v_{a_t} p_M q$, where $p'_x = (x, x^3 - \varepsilon), u_x = (x - \varepsilon^2, (x - \varepsilon^2)^3), v_x = (x + \varepsilon^2, (x + \varepsilon^2)^3))$; see Figure 9. This polygon can be constructed in $O(n)$ time, and in the vicinity of the point p_{a_i} its stabbing number changes locally from 1 to 3; see Figure 10. Let s_{a_i} be the segment joining u_{a_i} to p'_{a_i}. Therefore, the stabbing number of P_2 is 10 if and only if three of the segments s_{a_i} can be simultaneously stabbed, and 8 otherwise, as two of those segments can always be stabbed.

Step 7. Compute the stabbing number of polygon P_2 using the best possible algorithm.

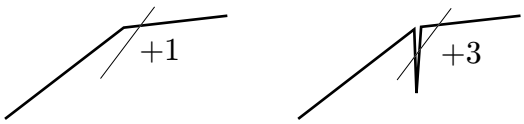

Figure 10: From P_1 to P_2 the stabbing number locally increases from 1 to 3.

Step 8. If the stabbing number of polygon P_2 is 10, conclude that there were three numbers in the initial input such that their sum is 0; if the stabbing number of P_2 is 8, conclude that there are no such three numbers.

The correctness of Step 8 is a direct consequence of Lemma 1 and Lemma 2. As all steps but Step 7 have overall complexity $O(n \log n)$, we conclude that Step 7, the computation of the stabbing number of a polygon, is a 3SUM-hard problem, as claimed. □

As an immediate consequence we obtain:

Corollary 4. *The problem of computing the smallest k such that a given polygon is k-convex is 3SUM-hard.*

Also, if we replace Step 7 in the proof of Theorem 3 by checking whether or not polygon P_2 is 4-convex, we get:

Corollary 5. *The problem of deciding whether a given polygon is* 4*-convex is 3SUM-hard.*

The stabbing number of a polygon defined at the beginning of Section 2.1 and studied in this paper is related to, but different from, the similar concept of the stabbing number of a set of segments that has appeared many times in the literature. Let us elaborate more on the latter concept and see also some implications of our results. The *stabbing number of a set of (closed) segments* is defined in [15] as the maximum number of segments that can be intersected by a line, and can be computed for n given segments in $O(n^2)$ time using the algorithm described at the beginning of Section 2.2. Notice that with this definition the boundary of a convex polygon with more than three vertices is a set of segments having stabbing number 4, given by any line joining a pair of nonadjacent vertices, while the stabbing number of the polygon itself, as defined in this paper, would be 2. In fact, by considering simple polygons that have many collinear vertices that are not adjacent it is possible to make these two numbers quite apart, therefore the two concepts should not be confused. However, for the polygon constructed in the proof

of Theorem 3 both numbers coincide, because an intersection of a line with the cubic $y = x^3$ yields 3 intersections that can be made 3 crossings, and reversely, when the point is replaced by the gadget for input points with integer coordinates used in the proof, and the cubic line by the polygon boundary. Therefore, we also get the following consequence:

Corollary 6. *The problem of finding the stabbing number of a set of segments is 3SUM-hard.*

2.3. Fast triangulation

Triangulating a simple polygon in $o(n \log n)$ time with a simple method is a challenging open problem. For k-convex polygons, this can be achieved, because we can sort the vertices of a k-convex polygon P in any given direction (say, x-direction) in $O(kn)$ time: Simply scan around ∂P and use insertion sort, starting each time from the place where the abscissa of the previous vertex has been inserted. After insertion, any x_j takes part in at most $2k - 1$ comparisons, because otherwise the vertical line $x = x_j$ would intersect P in more than k components. After sorting the vertices, a simplified plane sweep method can be used to build a vertical trapezoidation [10, 17] (and then a triangulation) of P. Only trivial data structures are needed, because each vertex can be processed, with a brute-force approach, in time $O(k)$ (by the k-convexity of P). We conclude:

Proposition 7. *Any k-convex polygon can be triangulated in $O(kn)$ time and $O(n)$ space.*

Using suitable data structures, a faster yet still implementable algorithm is possible, as we show next. Call a polygon *k-monotone* if every vertical line intersects the interior of the polygon in at most k intervals. (This property is implied by k-convexity, and is equivalent to x-monotonicity for $k = 1$.) Actually, we do not even need the polygon to be simple, we just need a sequence of x-coordinates such that every other x-coordinate comes between at most $2k$ consecutive pairs of x-coordinates.

Lemma 8. *The vertices of any k-monotone polygon can be x-sorted in $O(n \log(2 + k))$ time.*

Proof. We use a binary insertion sort, in which we add the points in order along the polygon into a balanced binary search tree. The binary search tree has the *dynamic finger property*: inserting an element that has rank r different from the previously inserted element costs only $O(\log(2 + r))$ time. (For example, splay trees [32] and Brown & Tarjan finger trees [8] both have

this property.) Once the elements are inserted into the binary search tree, we simply perform a linear-time in-order traversal to extract them in sorted order.

Now we find the bound on the total cost of the insertions. When we insert an element of rank difference r from the previously inserted element, we can charge this cost to r points formed from projecting (vertically) all vertices onto all edges of the polygon. There are at most $O(nk)$ such points of projection, so the total number of charges is at most $O(nk)$. Thus the total insertion cost is $O(\sum_{i=1}^{n} \log(2 + r_i))$ where $\sum_{i=1}^{n} r_i = O(nk)$. Such a sum is maximized when the r_i are all roughly equal, which means that they are all $\Theta((nk)/n) = \Theta(k)$. Therefore the total cost is at most $O(n \log(2 + k))$. $\square$

To see that this bound is optimal in the comparison model, consider the case in which the polygon is a comb with k tines. Then sorting the abscissae is equivalent to merging k sorted sequences of length n/k, and this problem has $\Theta(n \log k)$ time complexity[3].

Lemma 8 yields a fast triangulation method for general k-convex polygons. We first sort the vertices of the k-convex polygon P in a fixed direction, in $O(n \log k)$ time. Again, a plane sweep is used to compute a triangulation of P. As the intersection of P with the sweep line is of complexity $O(k)$ only, by the k-convexity of P, each of the n vertices of P can be processed in $O(\log k)$ time during the sweep.

Theorem 9. *Any k-convex polygon can be triangulated in $O(n \log k)$ time and $O(n)$ space.*

3. Two-Convex Polygons

3.1. Characterization

In this section we give a characterization of 2-convex polygons that allows their recognition in time $O(n \log n)$, and a description of their structure that will be used later in several of our results.

We observe that k-convexity is a property that may be lost by small perturbations on the positions of the vertices of a polygon. For example, small changes in the positions of some vertices in the polygon of Figure 4

[3] n distinct numbers can be arranged in $\Omega(k^n)$ different ways into k sorted lists. Hence, any comparison tree to sort n numbers given as input in k sorted lists has height $\Omega(\log k^n) = \Omega(n \log k)$.

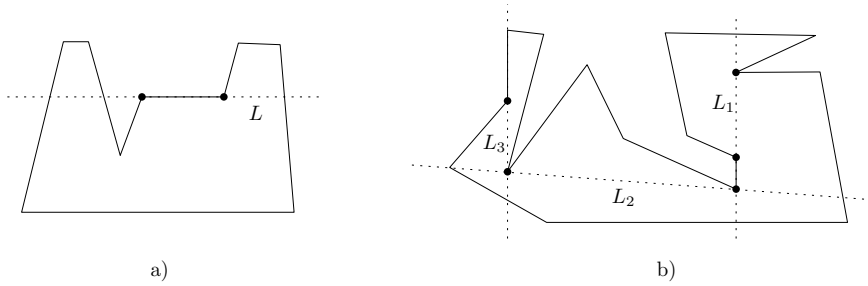

Figure 11: a) L is a 4-stabber. b) L_1 and L_2 are inner tangents, but L_3 is not.

(right), could yield a 2- or 3-convex polygon. As a consequence, stabbing a polygon along its edges will not, in most cases, give enough information for deciding its k-convexity.

We will use the following terminology to describe different relative positions between a line L and a polygon P:

- L is called a *j-stabber* of P if it crosses ∂P at least j times (see Figure 11.a for an example of a 4-stabber).
- L is an inflection line if it contains an inflection edge of P.
- L is tangent to P at vertex v if it passes through v without crossing ∂P.
- L is an inner tangent if it is tangent to two nonconsecutive reflex vertices of the polygon, and there are points interior to the polygon in each of the three intervals in which these two points split the line (see Figure 11.b).

Lemma 10. *A simple polygon P is 2-convex if and only if P has no inner tangent, and no inflection line that can be infinitesimally perturbed to a 6-stabber.*

Proof. The 'only if' implication is obvious, because an inner tangent can also be infinitesimally perturbed to a 6-stabber.

To prove the 'if' implication, assume that P is not 2-convex. Then there exists a 6-stabber L of P. Assuming that L is not vertical, let $c_1, \ldots, c_6$ be the six left-most crossings of L with ∂P, ordered left to right.

Two types of crossing pattern arise, according to whether polygonal chains between c_2 and c_3 and between c_4 and c_5 are on the same side of L

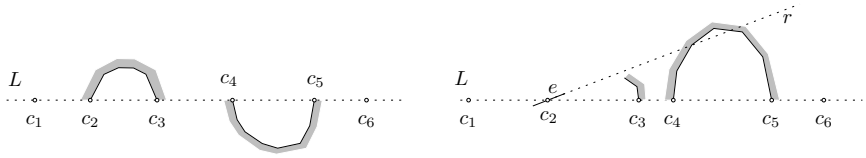

Figure 12: Patterns for a 6-stabber (the interior of the polygon is shaded).

or not (see Figure 12). If the chains are on different sides, then the geodesic path between c_2 and c_5 contains a portion of an inner tangent. If the chains are on the same side, consider the vertices of chains $c_2 \cdots c_3$ and $c_4 \cdots c_5$ above L; we can assume without loss of generality that the furthest point to L is on the chain $c_4 \cdots c_5$, as in Figure 12 (right). Let e be the edge containing c_2 and let r be the line defined by e. If r does not intersect the chain $c_4 \cdots c_5$ we are done because then the geodesic path between c_2 and c_5 contains the portion of an inner tangent. Otherwise, consider the endpoints of e. If one is reflex and the other one is convex, then we are done because r is an inflection line that can be infinitesimally perturbed to a 6-stabber. If both endpoints are reflex, then we move along the boundary of the polygon clockwise (away from c_3). For each new edge, we have the same three possibilities as with e: if the line containing the edge does not cross the chain $c_4 \cdots c_5$ there is an inner tangent, while if the line cross the chain we have finished if the edge is an inflection edge and we have to proceed if both endpoints are reflex. Clearly, this process has to end at some point. Finally, if both endpoints of e are convex, the same argument applies, but this time we have to move in counterclockwise direction (towards c_3).

□

3.2. Recognition

Suppose that we want to decide if a polygon P is 2-convex. (Assume that P is not convex; the problem is trivial, otherwise.) Our recognition algorithm is based on Lemma 10 and proceeds in two steps. In the first one, for each reflex vertex u we consider the four rays defined by the adjacent edges. If any of these rays intersects ∂P more than once, then a 6-stabber exists. This can be checked in $O(n \log n)$ time using the result in [9] to process the interior of the polygon and its pockets (because we compute two intersections for at most one ray). In this step we can also take care of degenerate situations, when one of these rays contains another vertex of the polygon: it can be easily checked in constant time whether the ray can be perturbed to a 6-stabber, or not.

For the second step, we observe that the set of lines tangent to the reflex vertex v (and not containing the edges adjacent to v) is mapped under the standard duality transform to an open segment (if the vertical line is not tangent to v) or to the complement of a segment in a line (if the vertical line is tangent to v). Therefore, the existence of an inner tangent (not found in step 1), reduces to the problem of checking if there exists an intersection in a set of lines segments, which can be solved in $O(n \log n)$ time and $O(n)$ space (see Theorem 7.9 in [28]).

We have then the following result:

Theorem 11. *Deciding if a simple polygon P with n vertices is 2-convex can be done in $O(n \log n)$ time and $O(n)$ space.*

3.3. Shape structure

We have given a geometric characterization of 2-convex polygons, in Subsection 3.1. The present subsection aims at giving a qualitative description of their shape. Recall that the polygons we consider do not have two consecutive collinear edges. Therefore, the inner angle at vertex u is either smaller or bigger than π. In the first case, the vertex is convex, while in the second case it is reflex.

Lemma 12. *Let P be a 2-convex polygon. Let $C = p_0 p_1 \dots p_t$ be the chain of vertices along P that connects (counterclockwise) two consecutive vertices p_0, p_t on the convex hull $CH(P)$. Then C can be partitioned into three chains $C_1 = p_0 p_1 \dots p_r$, $C_2 = p_{r+1} \dots p_s$, and $C_3 = p_{s+1} \dots p_t$, 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.*

Proof. If C_2 is empty, the lemma is obviously true, so we assume that the chain C contains at least one reflex vertex. Suppose that the line L defined by p_0 and p_t is horizontal, and that P lies below it. Let p_k be the last point of C to be hit when we move L down, in a parallel sweep. Observe that p_k is necessarily a reflex vertex. Let us recall that an inflection edge is adjacent to a reflex vertex and to a convex vertex and that, according to Lemma 10, 2-convex polygons do not have 4-stabbers containing an inflection edge. We are going to see that there is at most one inflection edge in the chain $p_0 \dots p_k$. Of course, the same applies to the chain $p_k \dots p_t$. Let $p_i p_{i+1}$ be the first inflection edge starting from p_0 (see Figure 13.a) and assume that there exists a convex vertex p_j in the chain $p_{i+1} \cdots p_k$. Then we have that $p_{j-1} p_j$ is an inflection edge supporting a 4-stabber: the ray $p_j p_{j-1}$ intersects the polygon at least once, while the ray $p_{j-1} p_j$ intersects the chain $p_k \dots p_t$ and therefore intersects the polygon at least twice. $\square$

Observe that the chains C_1 and C_3 might be singletons in some cases, and that C_2 might be empty. However, the generic aspect of a pocket and the shape of a 2-convex polygon are as shown in Figure 13.b). The next result follows directly:

Corollary 13. *If the convex hull of a 2-convex polygon P has k vertices, then the boundary of P can be decomposed into at most k convex chains and k reflex chains.*

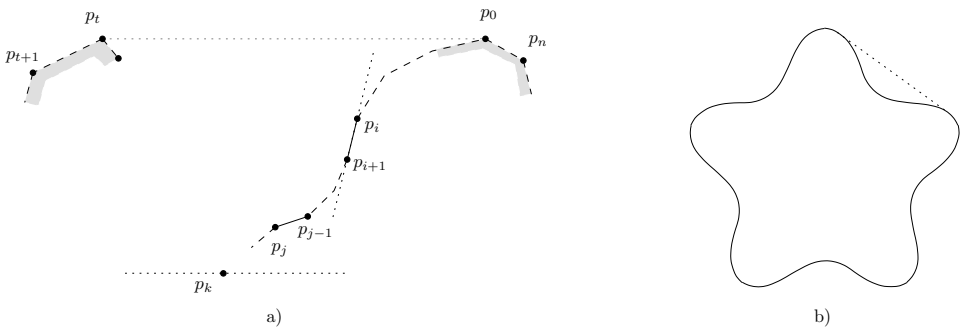

Figure 13: Illustration for the proof of Lemma 12.

The Erdős-Szekeres Theorem says that every set of n points in general position contains at least $\log n$ points that are in convex position, and that this value is asymptotically tight [13]. As every point set can be 'polygonized', one cannot expect a better value when the points are chosen from the set of vertices of an arbitrary polygon. However, when a point set is the set of vertices of a 2-convex polygon, we can improve this bound as follows.

Theorem 14. *Every 2-convex polygon with n vertices has a subset of $\lceil\sqrt{n/2}\rceil$ vertices in convex position. This bound is tight.*

Proof. By Corollary 13, the boundary of a 2-convex polygon with k vertices on its convex hull can be decomposed into at most k convex chains and k reflex chains. If $k \geq \lceil\sqrt{n/2}\rceil$, we are done, otherwise one of the $2k$ chains necessarily has size at least $\lceil\sqrt{n/2}\rceil$. The amoeba-like example in Figure 13.b), with $k = \lceil\sqrt{n/2}\rceil$ vertices in the convex hull and $2k$ chains of equal size shows that this bound is tight. □

We conclude this section with a consequence of the preceding theorem.

Corollary 15. *If an n-gon is 2-convex, then its vertices can be grouped into at most $2\sqrt{2n}$ subsets, each in convex position.*

Proof. Let $S(n)$ be the number of convex subsets needed to partition the vertex set of a 2-convex polygon with n vertices. We show that $S(n) \leq \alpha\sqrt{n}$ by induction over n. The induction base for $n = 3$ is obvious, and valid for any $\alpha \geq 1$. By Theorem 14 we find one convex subset of size at least $\lceil\sqrt{n/2}\rceil$ which is either the set of vertices of the convex hull, or one of the $2k$ chains mentioned in Corollary 13. In both cases, it is clear that if we delete the points and consider the new polygonization given by the original order of the points the remaining points define also a 2-convex polygon and we have $S(n) \leq 1 + S(\lfloor n - \sqrt{n/2}\rfloor) \leq 1 + \alpha\sqrt{\lfloor n - \sqrt{n/2}\rfloor}$, where the last inequality comes from the induction hypothesis. To prove the corollary it is sufficient to show that $1 + \alpha\sqrt{n - \sqrt{n/2}} \leq \alpha\sqrt{n}$. Standard manipulation shows that this is true for any $\alpha \geq 2\sqrt{2}$ and any $n \geq 1$.

□

4. General k-Convex Sets

The union or intersection of simple polygons may not be a polygon. In view of this fact, the issue of how the degree of convexity behaves with respect to these operations is not meaningful for this class of objects. In this section, we consider larger classes of sets in $\mathbb{R}^2$ for which these natural questions may be discussed. We first study some properties of compact 2-dimensional (not necessary polygonal) subsets of $\mathbb{R}^2$.

Lemma 16. *Given a k-convex set Q_1 and an m-convex set Q_2, the union $Q_1 \cup Q_2$ is a $(k+m)$-convex set, which is the maximum attainable value.*

Proof. The number of crossings of a line with the boundary of $Q_1 \cup Q_2$ can be at most $2k + 2m$. On the other hand, if Q_1 and Q_2 are disjoint and the line that gives k and m connected components, respectively, is the same, the value is achieved. □

Lemma 17. *Let Q_1 and Q_2 be, respectively, k-convex and m-convex sets such that the intersection is a 2-dimensional set. Then, $Q_1 \cap Q_2$ is a $(k + m - 1)$-convex set, which is the maximum attainable value.*

Proof. An oriented line will cross the boundary of Q_1 at most $2k$ times and the boundary of Q_2 at most $2m$ times. However, the first intersection point a does not contribute to the total number of crossings with $Q_1 \cap Q_2$ unless $a \in Q_1 \cap Q_2$, in which case it contributes only once instead of twice as crossing. The same happens with the last point, which gives the upper bound. An example proving the tightness appears in Figure 14(a). □

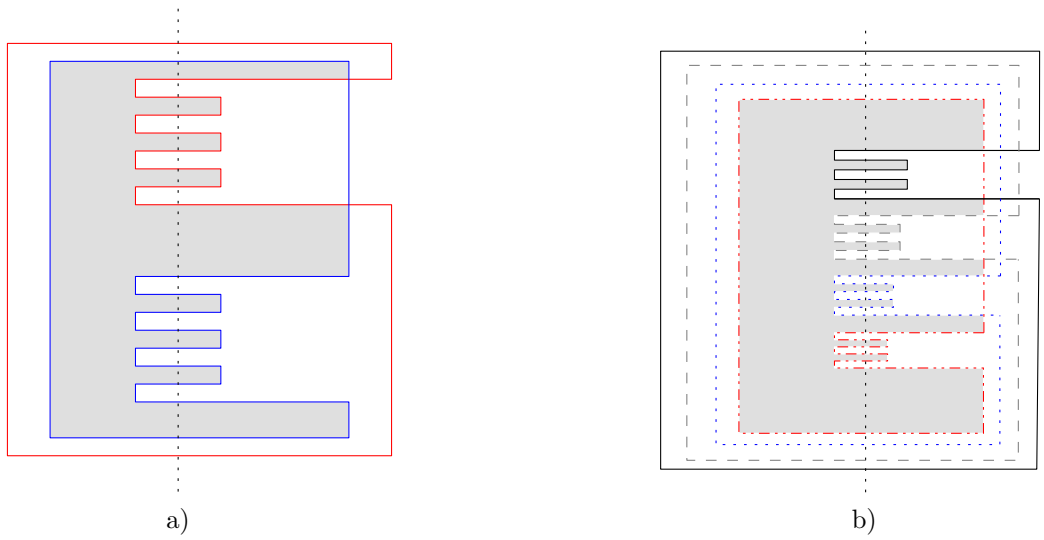

Figure 14: Intersecting k-convex sets

Corollary 18. *Let $Q_1, \ldots, Q_m$ be a family k-convex sets such that the intersection is a 2-dimensional set. Then, $\bigcap_{i=1}^{m} Q_i$ is a $(m(k-1)+1)$-convex set, which is the maximum attainable value.*

Proof. The upper bound follows from the preceding lemma, and a construction giving its tightness is shown in Figure 14(b). □

Theorem 19. *There is no Helly-type theorem for k-convex sets.*

Proof. We are constructing a family of m 2-convex sets such that any subfamily has nonempty intersection yet there is no point common to all of them. Let Q_m be a regular polygon with m edges $e_1, \ldots, e_m$ (refer to Figure 15). Let P_i^* be the polygonal chain obtained from the boundary of Q_m by removing edge e_i and an infinitesimal portion of e_{i-1} and e_{i+1}. Finally, let us give some slight thickness to the chain so it becomes a polygon P_i. Notice that the polygons $P_1, \ldots, P_m$ are 2-convex, the intersection $\bigcap_{i=1}^{m} P_i$ is clearly empty, while the intersection of every proper subfamily F is nonempty because it contains the intervals e_j for all those $P_j \notin F$. □

The preceding lemmas apply to k-convex sets in general, not only to sets with bounded description complexity. However, a significant difference appears in our next results, that are possibly the most natural to explore,

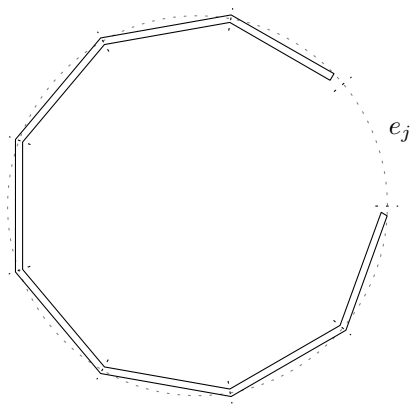

Figure 15: No Helly-type theorem for k-convex sets

because they involve transversal lines, which are precisely the main concept underlying the definition of k-convexity.

Let us recall [37] that given a family of sets $Q_1, \ldots, Q_m$, a line L is said to be a *transversal* of the family if L has a nonempty intersection with each of the sets. When the sets are convex, the ordering in which they are traversed (disregarding the orientation of the line) is called a *geometric permutation*, a topic that has received significant attention [37]. In particular, it has been proved that m compact disjoint convex sets admit at most $2m-2$ geometric permutations, which is tight [11].

Let us consider now transversals of 2-convex sets. Notice that every object will appear at least once, but may appear twice on the transversal, which we consider as combinatorially different cases of the associated *generalized geometric permutation*. Formally, let $\mathcal{F}$ be a family of 2-convex sets and let L be a line intersecting all members of $\mathcal{F}$. The generalized geometric permutation induced by L is a list of labels, one for each connected component of the intersection of L with the sets of $\mathcal{F}$ (see Figure 16 for an illustration). The following is clearly the first natural question raised by this definition: What is the maximum number of generalized geometric permutations a family of m disjoint 2-convex sets may have?

Theorem 20. *The number of generalized geometric permutations of a set of m disjoint 2-convex polygons can be exponential in m.*

Proof. A *nose* of a object P is a zig-zag sequence of a reflex and a convex

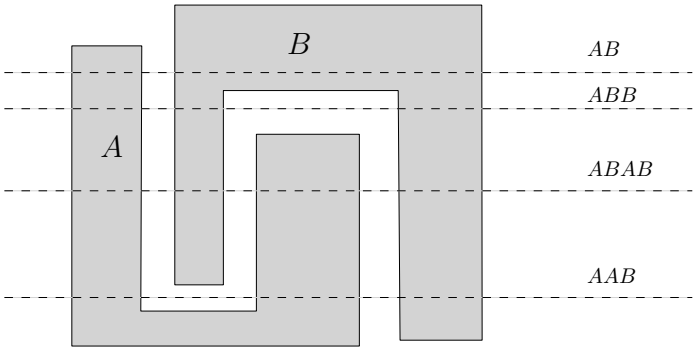

Figure 16: While two convex sets define a single geometric permutation, two 2-convex sets can define several generalized geometric permutations.

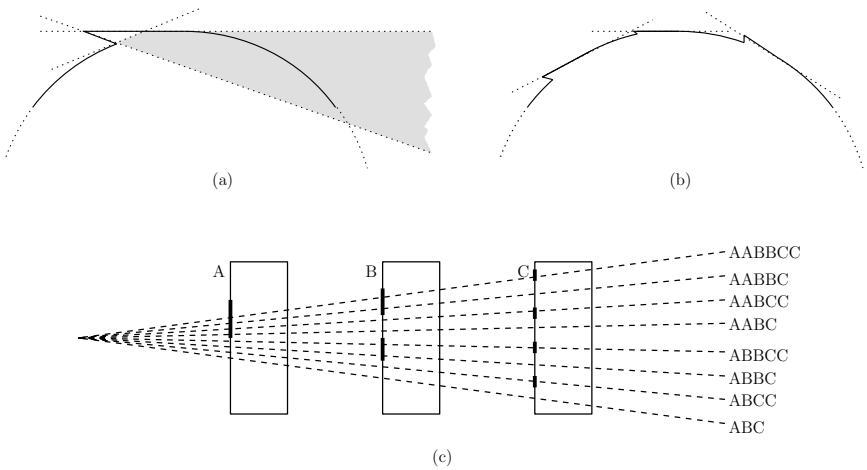

Figure 17: (a) and (b): 2-convex polygons can have an unbounded number of noses. (c): The number of generalized geometric permutations for a set of 2-convex polygons can be exponential.

vertex of the boundary of O as depicted in Figure 17(a). Consider a polygon with noses constructed in a way that they lie essentially in the direction of the boundary where they have been added to. The shaded area in Figure 17(a) indicates the region which is not intersected by a line tangent to one of the vertices of the nose, that is, a line e within this region intersects the nose only once. Thus we can iteratively construct further noses in this region without destroying the 2-convexity of O. Figure 17(b) shows an example where the principle shape of O is part of a disk. Observe that when the radius of the disk is large enough we can arrange an arbitrary number of flat noses such that O stays 2-convex.

Let R_i be an object which has the base shape of an axis-aligned rectangle, where the left side is actually part of a circle with sufficiently large radius and a center point far to the right of R_i. We place 2^{i-1} noses along this side, so that R_i stays 2-convex as described above. Next we arrange m objects R_1 to R_m from left to right, as depicted in Figure 17(c) for $m = 3$. We position the noses for each R_i in a regular way such that a rotating line (see the dashed lines in Figure 17(c)) intersects the noses in the same manner as the digit "1" shows up in the sequence of all 2^m binary numbers of length m. Thus, we get 2^m different generalized geometric permutations for this setting, as each object appears twice if the nose is intersected, but only once otherwise. □

Note that the complexity of the polygons in the preceding proof is exponential in m. The next theorem gives a tight bound for the number of generalized geometric permutations in terms of the total complexity of the polygons.

Theorem 21. *The maximum number of generalized geometric permutations of a set of* 2*-convex polygons with a total of* n *edges is* $\Theta(n^2)$.

Proof. As the standard duality transform maps each edge to a double wedge, the induced arrangement of $2n$ lines in the dual plane yields a quadratic number of cells that bound from above the number of possible ways of stabbing the set of objects.

To see that the bound is asymptotically tight, we give a construction using n 2-convex polygons; in fact, the simplest possible ones, namely, non-convex quadrilaterals. Let $Q_1^* = a_1^* b_1 c_1 d_1$ be the quadrilateral shown in Figure 18. Let h be a horizontal line through a_1^* and let $Q_i^* = a_i^* b_i c_i d_i$ be translates of Q_1 in such a way that all of them are pairwise disjoint and

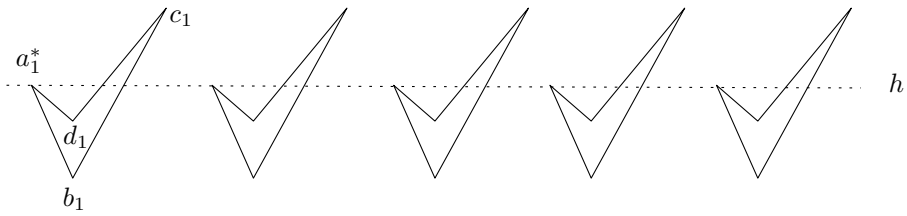

Figure 18: The number of generalized geometric permutations for a set of 2-convex polygons can be quadratic.

$a_1^*, a_2^*, \ldots, a_n^*$ appear in this order on h. Finally, let us perturb infinitesimally a_i^* to a_i in such a way that

a) points $a_1, a_2, \ldots, a_n$ are in general position,

b) for all i, j, k, with $i \neq j$, the line $a_i a_j$ leaves above the point c_k and below the point d_k.

For $i = 1, \ldots, n$, let Q_i be the quadrilateral with vertices $a_i b_i c_i d_i$. There are $\binom{n}{2}$ lines of the type $a_i a_j$; each of them leaves above and below a different set of points a_k, and is a transversal because it crosses all the segments $b_k c_k$ for every k. Now, if a_k is below the transversal, Q_k is intersected once, while if a_k is above the transversal, Q_k is intersected twice. Therefore, we have obtained $\binom{n}{2}$ generalized geometric permutations. $\square$

Observe that the two preceding theorems apply *mutatis mutandis* to k-convex sets, because 2-convex sets are also k-convex for $k \geq 3$.

5. Discussion and Open Problems

In this paper we have considered a new concept of generalized convexity. Moving from convexity to 2-convexity is seemingly a small change, as we are just accepting lines to intersect in at most two connected components instead of one. It is remarkable that this modest departure has strong consequences in the complexity of the new class of objects, as we have seen in this paper; obviously, even more when the degree of convexity is increased. Several open problems remain and many interesting questions can be raised. We list some of them below.

1) Can the recognition of 2-convex polygons be carried out in linear time, improving on the $O(n \log n)$ algorithm we provide?

2) Finding the smallest k such that a given polygon is k-convex is a 3SUM-hard problem. In particular, recognizing 4-convexity is already 3SUM-hard. We do not know whether the situation is the same for 3-convexity or whether a subquadratic time algorithm exists for this case.

3) Is it possible to generalize Theorem 14? For example, is it true that every k-convex polygon with n vertices has a large subset of vertices that are the vertices of a $(k-1)$-convex polygon?

4) Let us define the k-convex hull of a point set S as the smallest area polygon which is k-convex, has a subset $T \subset S$ as vertex set and every point in $S \setminus T$ is inside the polygon. Which is the complexity of computing this k-convex hull? Observe that for $k = 1$ this notion is the usual convex hull of a point set.

5) Give combinatorial bounds and efficient algorithms for decomposing a polygon into k-convex subpolygons. This is a classical problem when convex subpolygons are considered [19]. The decomposition into pseudotriangles, a particular class of 2-convex polygons, has also been studied [3, 30]. However, the latter result might be improved by considering more general 2-convex polygons.

6) A k-convex decomposition of a set S of n points in the plane is a decomposition of its convex hull into k-convex polygons such that every point in S is a vertex of some of the polygons. For $k = 1$, a triangulation suffices, though it has been proved that if we allow arbitrary convex sets the number can be reduced [25]. On the other hand, it has been shown that there exist always a decomposition into exactly $n-2$ pseudotriangles, which are 2-convex polygons [30]. It is an intringuing open problem to decide whether this number can be reduced to sublinear if we allow arbitrary 2-convex polygons.

7) As mentioned in the introduction of the paper, applications of k-convex polygons arise in Art Gallery type problems. In this context, instead of illuminating a polygon, we want to *cover* the interior of a simple polygon with a set of j-modems, whose signal can cross the boundary of the polygon j times. Our last open problem is that of establishing bounds on the number of j-modems needed to cover k-convex polygons with n vertices. See [2, 6, 14] for some partial results on this problem.

Finally, let us mention that in this paper we have focused on k-convex *polygons*. It is natural to define a similar concept for finite point sets, namely,

being in *k-convex position*, where k is given by the smallest degree of convexity attained when *all* possible polygonizations of the point set are considered. This issue is considered in a companion paper [1].

Acknowledgements We would like to thank Thomas Hackl, Clemens Huemer, David Rappaport, Vera Sacristán and Birgit Vogtenhuber for sharing discussions on this topic, and to the anonymous referees, particularly to one of them whose sharp comments on basic definitions helped crucially to improve the presentation of this paper.

[1] O. Aichholzer, F. Aurenhammer, T. Hackl, F. Hurtado, P.A. Ramos, J. Urrutia, P. Valtr, B. Vogtenhuber. k-convex points sets. Manuscript, 2011.

[2] O. Aichholzer, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, J. Urrutia, and B. Vogtenhuber. Modem Illumination of Monotone Polygons. In Proc. 25th European Workshop on Computational Geometry EuroCG '09, Brussels, Belgium, 2009, 167-170.

[3] O. Aichholzer, C. Huemer, S. Kappes, B. Speckmann, C.D. Tóth. Decompositions, partitions, and coverings with convex polygons and pseudo-triangles. Graphs and Combinatorics 23 (2007) 481-507.

[4] G. Aloupis, P. Bose, V. Dujmovic, C. Gray, S. Langerman, B. Speckmann. Triangulating and guarding realistic polygons. Proc. 20th Canadian Conference on Computational Geometry, 2008.

[5] D. Artigas, M.C. Dourado, J.L. Szwarcfiter. Convex partitions of graphs. Electronic Notes in Discrete Mathematics 29 (2007) 147-151.

[6] B. Ballinger, N. Benbernou, P. Bose, M. Damian, E. D. Demaine, V. Dujmovic', R. Flatland, F. Hurtado, J. Iacono, A. Lubiw,, P. Morin, V. Sacristán, D. Souvaine, and R. Uehara. Coverage with k-Transmitters in The Presence of Obstacles. Proc. COCOA 2010 (Combinatorial Optimization and Applications, Hawaii, December 2010). Lecture Notes in Computer Science, Springer-Verlag, 2010, Volume 6509/2010, 1-15.

[7] M. Breen, D.C. Kay. General decomposition theorems for m-convex sets in the plane. Israel Journal of Mathematics 24 (1976) 217-233.

[8] M.R. Brown, R.E. Tarjan. Design and Analysis of a Data Structure for Representing Sorted Lists. SIAM Journal on Computing 9(3) (1980) 594-614.

[9] B. Chazelle, H. Edelsbrunner, M. Grigni, L. Guibas, J. Hershberger, M. Sharir, J. Snoeyink. Ray shooting in polygons using geodesic triangulations. Algorithmica 12 (1994) 54-68.

[10] B. Chazelle, J. Incerpi. Triangulation and shape-complexity. ACM Transactions on Graphics 3 (1984) 135-152.

[11] H. Edelsbrunner, M. Sharir. The maximum number of ways to stab n convex non-intersecting sets in the plane is $2n-2$. Discrete & Computational Geometry 5 (1990) 35-42.

[12] H. ElGindy, G.T. Toussaint. On geodesic properties of polygons relevant to linear time triangulation. The Visual Computer 5 (1989) 68-74.

[13] P. Erdős, G. Szekeres. A combinatorial problem in geometry. Compositio Mathematica 2 (1935) 463-470.

[14] R. Fabila-Monroy, Andres Ruis Vargas, and J. Urrutia. On Modem Illumination Problems. XIII Encuentros de Geometria Computacional, Zaragoza, España, June 29 - July 1 2009, 9-19.

[15] S.P. Fekete, M.E. Lübbecke, H. Meijer. Minimizing the stabbing number of matchings, trees, and triangulations. Proc. 15th Ann. ACM-SIAM Symp. on Discrete Algorithms, 2004, 430-439.

[16] E. Fink, D. Wood. Fundamentals of restricted-orientation convexity. Information Sciences 92 (1996) 175-196.

[17] A. Fournier, D.Y. Montuno. Triangulating simple polygons and equivalent problems. ACM Transactions on Graphics 3 (1984) 153-174.

[18] A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. Computational Geometry: Theory and Applications 5 (1995) 165-185.

[19] M. Keil. Polygon Decomposition in Handbook of Computational Geometry, J.R. Sack and J. Urrutia (Eds) (2000), 491-518.

[20] J. King. A Survey of 3SUM-Hard Problems. http://www.cs.mcgill.ca/~jking/papers/3sumhard.pdf

[21] G. Kun, G. Lippner. Large empty convex polygons in k-convex sets. Periodica Mathematica Hungarica 46 (2003) 81-88.

[22] D.T. Lee, F.P. Preparata. An optimal algorithm for finding the kernel of a polygon. Journal of the ACM 26 (1979) 415-421.

[23] A. Maheshwari, J.-R. Sack, H. Djidjev. Link distance problems. In: Handbook of Computational Geometry, J.-R. Sack and J. Urrutia (eds.), Elsevier, 2000, 519-558.

[24] J. Matoušek, P. Plecháč. On functional separately convex hulls. Discrete & Computational Geometry 19 (1998) 105-130.

[25] V. Neumann-Lara, E. Rivera-Campo, J. Urrutia. A note on convex decompositions of point sets in the plane. Graphs and Combinatorics 20 (2004) 223-231.

[26] J. O'Rourke. Art Gallery Theorems and Algorithms. Oxford University Press, 1987.

[27] T. Pennanen. Graph-convex mappings and K-convex functions. Journal of Convex Analysis 6 (1999) 235-266.

[28] F.P. Preparata and M. I. Shamos. Computational Geometry: An Introduction. Springer-Verlag, 1988.

[29] G.J.E. Rawlins, D. Wood. Ortho-convexity and its generalizations. In: Computational Morphology, G.T. Toussaint (ed.), North-Holland Publishing Co., Amsterdam, 1988, 137-152.

[30] G. Rote, F. Santos, I. Streinu. Pseudo-triangulations—a survey. Contemporary Mathematics 453 (2008) 343-410.

[31] S. Schuierer, D. Wood. Visibility in semi-convex spaces. Journal of Geometry 60 (1997) 160-187.

[32] D.D. Sleator, R.E. Tarjan. Self-adjusting binary search trees. Journal of the ACM 32(3) (1985), 652-686.

[33] J. Urrutia. Art Gallery and Illumination Problems. Handbook on Computational Geometry, Elsevier Science Publishers, J.R. Sack and J. Urrutia eds. pp. 973-1026, 2000.

[34] F.A. Valentine. A three-point convexity property. Pacific Journal of Mathematics 7 (1957) 1227-1235.

[35] P. Valtr. A sufficient condition for the existence of large empty convex polygons. Discrete & Computational Geometry 28 (2002) 671-682.

[36] E. Welzl. On spanning trees with low crossing numbers. In: Data Structures and Efficient Algorithms. B. Monien and T. Ottmann (eds.), Springer LNCS 594, 1992, 233-249.

[37] R. Wenger, Progress in geometric transversal theory. Contemporary Mathematics, B. Chazelle and J.E. Goodman, Eds., American Math Society (1999), 375-393.