

On k -Convex Point Sets

Oswin Aichholzer^{a,1}, Franz Aurenhammer^b, Thomas Hackl^{a,2}, Ferran Hurtado^{c,3}, Alexander Pilz^{a,4,*}, Pedro Ramos^{d,5}, Jorge Urrutia^{e,6}, Pavel Valtr^f, Birgit Vogtenhuber^{a,1}

^a*Institute for Software Technology, Graz University of Technology, Austria.*

^b*Institute for Theoretical Computer Science, Graz University of Technology, Austria.*

^c*Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Spain.*

^d*Departamento de Matemáticas, Universidad de Alcalá, Madrid, Spain.*

^e*Instituto de Matemáticas, Universidad Nacional Autónoma de México.*

^f*Department of Applied Mathematics and Institute for Theoretical Computer Science (ITI), Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic.*

Abstract

We extend the (recently introduced) notion of k -convexity of a two-dimensional subset of the Euclidean plane to finite point sets. A set of n points is

*Corresponding author.

Email addresses: `oaich@ist.tugraz.at` (Oswin Aichholzer), `auren@igi.tugraz.at` (Franz Aurenhammer), `thackl@ist.tugraz.at` (Thomas Hackl), `hurtado@upc.edu` (Ferran Hurtado), `apilz@ist.tugraz.at` (Alexander Pilz), `pedro.ramos@uah.es` (Pedro Ramos), `urrutia@matem.unam.mx` (Jorge Urrutia), `valtr@kam.mff.cuni.cz` (Pavel Valtr), `bvogt@ist.tugraz.at` (Birgit Vogtenhuber)

¹Partially supported by the ESF EUROCORES programme EuroGIGA - ComPoSe, Austrian Science Fund (FWF): I 648-N18.

²Supported by the Austrian Science Fund (FWF): P23629-N18 ‘Combinatorial Problems on Geometric Graphs’.

³Partially supported by ESF EUROCORES programme EuroGIGA, CRP ComPoSe: MICINN grant EUI-EURC-2011-4306, by projects MICINN MTM2009-7242 and MINECO MTM2012-30951, and by project Gen. Cat. DGR 2009SGR1040.

⁴Recipient of a DOC-fellowship of the Austrian Academy of Sciences at the Institute for Software Technology, Graz University of Technology. Part of this work was done while A.P. was visiting the Departamento de Matemáticas, Universidad de Alcalá, Madrid, Spain.

⁵Partially supported by MEC grant MTM2011-22792 and by the ESF EUROCORES programme EuroGIGA, CRP ComPoSe, under grant EUI-EURC-2011-4306.

⁶Partially supported by Consejo Nacional de Ciencia y Tecnología, CONACYT (Mexico) grant number CB-2012-01-0178379.

considered k -convex if there exists a spanning (simple) polygonization such that the intersection of any straight line with its interior consists of at most k disjoint intervals. As the main combinatorial result, we show that every n -point set contains a subset of $\Omega(\log^2 n)$ points that are in 2-convex position. This bound is asymptotically tight. From an algorithmic point of view, we show that 2-convexity of a finite point set can be decided in polynomial time, whereas the corresponding problem on k -convexity becomes NP-complete for any fixed $k \geq 3$.

Keywords: convexity, polygonization, Erdős-Szekeres problem
2010 MSC: 52C10, 68U05

1. Introduction

One of the most important concepts in geometry is convexity. Convex objects enjoy many interesting, elegant, and useful properties and are usually much ‘easier’ to deal with combinatorially and computationally. Various generalizations and relaxations of convexity have been considered in the literature. In particular, and relevant for the present paper, the notion of k -convexity was recently introduced in [1] for what are called the *2-dimensional* subsets of the (Euclidean) plane. This family encompasses simple polygons in particular, which is sufficient for this work; in this case we say that a simple polygon is *k-convex* if the intersection of any straight line with the polygon consists of no more than k disjoint intervals. Clearly 1-convexity refers to the usual convexity. Among other results, several properties of k -convex planar *polygons* are presented in [1]. Here, we generalize this work by extending k -convexity to *finite* point sets in the plane. A point set S is said to be k -convex if there exists a k -convex polygon whose vertex set is S . In other words, we are interested in point sets which admit a k -convex polygonization.

The introduction of k -convexity of polygons [1] follows the common approach to generalize the notion of convexity under various aspects. The concept of convexity of 2-dimensional domains (such as convex polygons) has been translated to point sets in convex position in the field of combinatorial geometry. One major and ubiquitous result on the size of convex point sets is the combination of an upper and lower bound by Erdős and Szekeres [9, 10], which we will refer to as the *Erdős-Szekeres Theorem*.

Theorem 1 (Erdős-Szekeres). *Any point set of n points in general position contains a subset of points of size $\Omega(\log n)$ that is in convex position, and there are sets where this bound is tight.*

As for point sets in convex position, our work extends the concept of k -convexity of polygons to k -convexity of point sets, motivated by the importance of Erdős-Szekeres-type results in combinatorial and computational geometry. A comparable line of research is taken by Arkin et al. [3], who consider the minimum number of reflex vertices among all simple polygonizations of a point set.

We devote Section 2 to a precise definition of our generalization. The closely related concepts of stabbing number and j -stabber (of a polygon or a geometric graph) are also discussed. In Section 3, we present basic properties of point sets regarding k -convexity, giving lemmas that are used throughout the paper. For example, it is shown that every set of n points is $O(\sqrt{n})$ -convex, and that every subset of a k -convex point set is k -convex itself. Moreover, the union of a k -convex set and a j -convex set is $(k+j+1)$ -convex. We then state our main combinatorial result, which can be seen as a variant of the Erdős-Szekeres Theorem for 2-convexity: every point set of cardinality n contains a subset of $\Omega(\log^2 n)$ points that are in 2-convex position (Theorem 5). This result is best possible, in the asymptotic sense, as there exist sets of n points such that the largest 2-convex subset has size $O(\log^2 n)$. We also provide two algorithmic results. In Section 5, a polynomial-time algorithm is given for deciding whether a point set is 2-convex. By contrast, in Section 6 the problem of deciding k -convexity is proved to be NP-complete, for all $k \geq 3$. A list of open problems related to k -convexity of point sets is given in Section 7.

2. Preliminaries

Throughout this paper, let S denote a finite set of points in the plane. Finite point sets are assumed to be in general position; i.e., they do not contain collinear triples. A *geometric graph* is a graph where each vertex is represented by a point, and each edge by a straight line segment between the points representing its end vertices. Only geometric graphs will be considered in this paper. A graph G is a *spanning* graph of a point set S if S is the union of the endpoints of the edges of G . We omit stating the point set if it is clear from the context. Given any sequence $\langle s_1, \dots, s_n \rangle$ of n distinct points,

let s_i be connected to s_{i+1} by an edge, for $1 \leq i < n$. The resulting graph is called a *polygonal chain*. If we add an edge between s_n and s_1 , we obtain a *polygonal cycle*. A polygonal chain or cycle is *simple* if no two edges share a point in their relative interiors. A *simple polygon* is the closed finite region bounded by a simple polygonal cycle. The cycle forming the boundary of a simple polygon P is denoted by ∂P . The vertices and edges of P are the vertices and edges of ∂P , respectively. We call a vertex of P *convex* if the angle greater than π between its incident edges of ∂P is in the exterior, and *reflex* otherwise (recall that we assume general position of the vertices).

In the following, all lines are straight lines. We follow the definitions of [1]. Suppose a line ℓ has non-empty intersection with ∂P . At each component of the intersection, ℓ either crosses ∂P , or locally supports P along the component (which is either an edge or a single vertex of P). A line is a *j -stabber of P* if it crosses ∂P at least j times. The *stabbing number* of P is the largest number of crossings between ∂P and a line. Let ℓ be a line that intersects the interior of P in exactly k connected components. Since all vertices are in general position, there exists a perturbation of ℓ that is a $2k$ -stabber of P . Therefore a polygon is k -convex but not $(k - 1)$ -convex if and only if its stabbing number is $2k$.

Given a point set S , let \mathcal{P} be the set of all *polygonizations* of S , i.e., the set of all simple polygons whose vertex set is exactly S . If \mathcal{P} contains at least one polygon that is k -convex, then we call S a *k -convex point set*.

Several of our proofs transform non-simple polygonal cycles to simple ones. Let G be a (possibly self-intersecting) geometric graph. A line ℓ is a *j -stabber of G* if it crosses at least j edges of G . Note that the degenerate cases where a line passes through a vertex of G can be disregarded due to the same perturbation arguments as for the stabbing number of a simple polygon, resulting in a consistent definition of the stabbing number of geometric graphs. We can therefore define that G is a *k -convex graph* if G has stabbing number at most $2k$. Now any simple k -convex polygonal cycle C is the boundary of a k -convex polygon, and for every k -convex polygon P , ∂P is a k -convex polygonal cycle. For the sake of brevity, we will sometimes refer to a line ℓ as a *local j -stabber* if ℓ is a j -stabber of a subgraph that is clear from the context.

3. Basic Properties of k -Convex Sets

In this section, we establish several basic properties of point sets regarding their k -convexity. A first and quite useful result shows that also any *non-simple* polygonal cycle can be taken to determine the degree of convexity of a set of points.

Lemma 1. *Let C be a spanning (non-simple) k -convex polygonal cycle of a set S of points. Then S is k -convex.*

Proof. We show that C can be transformed into a simple polygonal cycle without increasing its stabbing number. Let e_i and e_j be two crossing edges of C . Remove e_i and e_j from C and replace them by two different edges e_l and e_m spanned by the end points of e_i and e_j such that we obtain a spanning polygonal cycle C' . Note that e_l and e_m do not cross, as they lie on the boundary of the convex quadrilateral formed by the end points of e_i and e_j . Moreover, every line intersecting one (or both) of e_l or e_m also intersects one of e_i or e_j (or both, respectively). Thus C' is also k -convex. In addition, the sum of lengths of e_l and e_m is strictly smaller than the sum of lengths of e_i and e_j , which implies that C' is shorter than C .

As long as the polygonal cycle is non-simple, we repeat the described process. In every step cycles get shorter, and there is only a finite number of cycles for S . Therefore we finally obtain one that is simple and defines a spanning k -convex polygon. Thus, by definition, S is k -convex. \square

Theorem 2. *Any set S of n points is $O(\sqrt{n})$ -convex, and this bound is tight.*

Proof. It is known that for any set S of n points there exists a crossing-free spanning tree T which has stabbing number $O(\sqrt{n})$ [6]. We obtain a (non-simple) spanning polygonal cycle C of S by traversing T in preorder and connecting the points accordingly. (Whenever we pass by a vertex which has already been used, we ignore it, except the starting vertex, where we close the polygon.) In this way we pass all edges of T twice (there and back again) and take several shortcuts. Thus any line intersects C at most twice as often as T . As for any shortcut which is stabbed by a line, the same line also stabs one branch of the tree at the shortcut. By Lemma 1, S is $O(\sqrt{n})$ -convex.

To see that the bound is tight, we show that there are point sets where the best polygonization is not k -convex for any k in $o(\sqrt{n})$. To this end, let S be the n points of a $\sqrt{n} \times \sqrt{n}$ grid, slightly perturbed to be in general position. Let L be a set of $\sqrt{n} - 1$ horizontal and $\sqrt{n} - 1$ vertical lines which can be

drawn between the different rows and columns to separate the grid points. Then any edge of an arbitrary polygonization P of S intersects at least one line of L . Assign each edge of P to one of the lines in L which it intersects. This way on average each line in L gets assigned $\frac{n}{2\sqrt{n}-2}$ edges of P . Thus, by the pigeonhole principle, there is at least one line in L which intersects $\Omega(\sqrt{n})$ edges of P ; that is, P is not k -convex for any k in $o(\sqrt{n})$. \square

Lemma 2. *Every subset of a k -convex point set is k -convex.*

Proof. Consider a k -convex polygonization of the given set. Whenever we remove a vertex p of this polygonal cycle to obtain the required subset, we replace the two edges incident to p by the direct connection of the two neighbors of p , transforming the initially simple cycle into a (possibly) non-simple cycle. As any line that intersects the newly introduced edge also has to intersect one of the two removed edges, this guarantees that the obtained polygonal cycle also has an at most $2k$ -stabber. By Lemma 1, the result follows. \square

Note that Lemma 2 implies that adding points can never transform a point set that is not k -convex into a k -convex set.

Lemma 3. *Let S be a k -convex point set and let T be a j -convex point set. Then the union $S \cup T$ is $(k+j+1)$ -convex. Let P_S be a k -convex polygonization of S , and P_T be a j -convex polygonization of T . If P_S and P_T intersect, then $S \cup T$ is $(k+j)$ -convex.*

Proof. Consider P_S and P_T . Obviously, we can remove one edge from each, and connect the end points of the resulting paths with two “new” edges, in such a way that we obtain a (possibly non-simple) polygonal cycle with vertex set $S \cup T$. As we add only two edges, any line gets at most two more intersections, and thus the polygonization is $(k+j+1)$ -convex. By Lemma 1, the first statement follows.

If P_S and P_T intersect, then we take a crossing pair of edges to be removed. Similar to the proof of Lemma 1, replacing this pair of edges by two non-crossing edges to obtain the polygonization of $S \cup T$ does not increase the stabbing number. The second statement follows. \square

It is possible to construct two convex sets such that their union is not 2-convex. One way is to take two concentric regular n -gons—the smallest such example consists of two concentric regular 5-gons, resulting in a set of




Figure 1: Two concentric regular 5-gons. The resulting set of 10 points is not 2-convex.

10 points which is 3-convex, but not 2-convex (see Figure 1). So far we do not have tight examples for higher degrees of convexity.

Proposition 1. *Every set of n points contains a $\lfloor \frac{3k}{2} \rfloor$ -convex set of cardinality $\Theta(k \log n)$ for any $k \in o(\frac{n}{\log n})$.*

Proof. By the Erdős-Szekeres Theorem [9, 10] we know that any set of n points contains convex chains of size $\Omega(\log n)$. If we repeatedly remove k convex chains of logarithmic size from a point set, we end up with a set of size

$$\begin{aligned} & n - \Theta(\log n + \log(n - \log n) + \log(n - \log(n - \log n)) + \dots) \\ & = n - \Theta(\log(n \cdot (n - \log n) \cdot (n - \log(n - \log n)) \cdot \dots)) \end{aligned}$$

This implies that we have $n - O(\log(n^k)) = n - O(k \log n)$ points left. When choosing $k \in o(\frac{n}{\log n})$, there is still a linear number of points left, and therefore all our k vertex-disjoint chains are of size $\Theta(\log n)$. We connect them with k additional edges to form a (possibly non-simple) polygonal cycle. Any line intersects each chain at most twice and, in addition, possibly all connecting edges. So the stabbing number of the constructed polygon is bounded by $3k$. By Lemma 1, we obtain a set of $\Theta(k \log n)$ points which is $\lfloor \frac{3k}{2} \rfloor$ -convex. \square

Note that this can be seen as a warm-up result since, for $k > \log n$, Theorem 2 provides a stronger result, and we will show that we always have a 2-convex subset of size $\Omega(\log^2 n)$ (Theorem 5).

Lemma 4. *If all point sets of cardinality n are k -convex, then all point sets of cardinality $n + i$ are $(k + \lceil \frac{i}{2} \rceil)$ -convex.*

Proof. Take a k -convex polygonization P of the n points. Connect the i points to a chain and replace some edge of P with the chain to obtain a (non-simple) cycle on $n+i$ points. Consider the added chain, which contains $i+1$ edges. A $2k$ -stabber of P intersects, in the worst case, all edges of that chain. If i is odd, then $i+1$ is even and the resulting cycle is $(k + \lceil \frac{i}{2} \rceil)$ -convex. If i is even, then $i+1$ is odd and therefore the added chain is intersected an odd number of times. Since the resulting cycle has to be intersected an even number of times, this means that the former $2k$ -stabber crosses $2k-1$ edges of P , i.e., it crossed the removed edge. Thus the cycle is $(k + \lceil \frac{i}{2} \rceil)$ -convex, and by Lemma 1 the result follows. \square

Define $s(k)$ to be the cardinality of the smallest point set which is not k -convex. The following result shows the asymptotic behavior of $s(k)$.

Proposition 2. $s(k) = \Theta(k^2)$, more specifically $s(k) \leq (\lfloor 2k+1+2\sqrt{k^2-k} \rfloor)^2$.

Proof. That $s(k) = \Theta(k^2)$ follows directly from Theorem 2. The specific upper bound follows from the $\Omega(\sqrt{n})$ -convex example used in the proof of Theorem 2. Let $n = m^2$ be the number of points. Then on one hand, $s(k) \leq m^2$, and on the other hand, $2k \geq \frac{m^2}{2(m-1)}$. Solving the resulting quadratic equation for m gives the claimed bound. \square

For small, constant k we obviously have $s(1) = 4$, and from the order type database [2] we can observe that every point set with at most 9 points is 2-convex. Moreover there exist sets of 10 points which are not 2-convex, but 3-convex. Thus $s(2) = 10$. From Lemma 4 and the fact that every set of 9 points is 2-convex, it follows that every set of at most 11 points is 3-convex; i.e., $s(3) \geq 12$.

4. On the Size of 2-Convex Subsets

In this section we present Erdős-Szekeres-type results on the size of 2-convex subsets, and on empty 2-convex polygons, which all point sets of a given cardinality must contain.

Let us first recall basic definitions and facts about 2-convex polygons given in [1]. An edge of a simple polygon P is called an *inflection edge* if it joins a convex and a reflex vertex of P . An *inflection line* is the supporting line of an inflection edge. A line ℓ is an *inner tangent* if it supports the boundary of the polygon at two nonconsecutive reflex vertices such that




Figure 2: Lines stabbing a polygon. An inflection edge (left) is contained in ℓ ; its perturbation ℓ' is a 6-stabber. An edge with two convex vertices (middle) is contained in ℓ ; ℓ' is a 4-stabber. An inner tangent and its perturbation is shown to the right.

there are points interior to the polygon in each of the three intervals in which these two vertices split the line. See Figure 2. The following result shows the interrelation between 2-convexity, inner tangents, and stabbers.

Lemma 5 ([1, 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.*

4.1. The Erdős-Szekeres Theorem Revisited

The Erdős-Szekeres Theorem [9] states that there exists a least integer $ES(k)$, for any $k \geq 3$, such that any set of at least $ES(k)$ points in the plane in general position contains k points in convex position. In their seminal paper [9], Erdős and Szekeres proved that $ES(k) = O(4^k/\sqrt{k})$. Although a lot of work has been done on this problem over the last 75 years, and $ES(k)$ is conjectured to be $\Theta(2^k)$, Erdős and Szekeres's original result has not been improved asymptotically. A variation of the Erdős-Szekeres Theorem that was later suggested by Erdős in 1974 [8] is to look for convex holes in point sets; i.e., convex subsets whose convex hull contains no point from the set in its interior. The existence of convex k -holes in large enough point sets was soon established for $k = 4$ (Erdős [8]) and $k = 5$ (Harborth [14]), but only recently has a proof been obtained for empty convex hexagons [13, 17]. In the opposite direction, Horton proved that there are sets of any size that do not contain an empty convex heptagon [16].

It is natural to ask similar questions for 2-convex subsets and empty 2-convex polygons. The following theorem shows that the situation becomes quite different once we go beyond 1-convexity. Let $CH(S)$ denote the convex hull of a point set S .

Theorem 3. *Every set S of n points in general position contains a subset of size $\Omega(\log n)$ that is the vertex set of a 2-convex polygon which does not contain points from S in its interior.*

Proof. Take any convex subset C of S of size $\Omega(\log n)$ [9, 10, 20]. Let $S' = S \cap (\text{CH}(C) \setminus \partial\text{CH}(C))$ be the set of points of S in the interior of $\text{CH}(C)$. Select an edge pq of $\partial\text{CH}(C)$. Then $\text{CH}(S' \cup \{p, q\})$ is the convex hull of p , q , and the points in the interior of $\text{CH}(C)$. Concatenating the polygonization $\partial\text{CH}(C) \setminus pq$ at the point p (by adding p) with $\partial\text{CH}(S' \cup \{p, q\}) \setminus pq$, and closing the chain at the point q (by adding q) results in a crescent-shaped polygonization which is 2-convex and empty of points of S . \square

Let us recall a result from [1] about convex subsets of vertices in 2-convex polygons.

Theorem 4 ([1]). *Every 2-convex polygon with n vertices has a subset of at least $\sqrt{n}/2$ points in convex position.*

Since the Erdős-Szekeres Theorem states that there are point sets with only $O(\log n)$ points in convex position, Theorem 4 implies the following result:

Corollary 1. *There are sets of n points such that the largest 2-convex subset consists of at most $O(\log^2 n)$ points.*

Note that for this upper bound we do not require the 2-convex polygon to be empty. In fact, the following theorem, which is proved in Section 4.2, tightens the lower bound in that case.

Theorem 5. *Every set S of n points in general position contains a 2-convex subset of size $\Omega(\log^2 n)$.*

The remaining open question is to close the gap between $\Omega(\log n)$ and $O(\log^2 n)$ for the size of empty 2-convex polygons.

4.2. Proof of Theorem 5

We will now prove Theorem 5 using Erdős-Szekeres-type results, such as the Monotone Subsequence Theorem, the Cap/Cup Theorem, and the Erdős-Szekeres Theorem. Throughout the proof, all polygons are simple.

Several variants of the Erdős-Szekeres Theorem have been considered. One of them, the *Convex Clustering* problem, will be the basis of our next result.

Definition 1. *Let \mathcal{C} be a set of k subsets C_i of S , $1 \leq i \leq k$, such that for each C_i there exists a straight line l_i (“linear separator”) separating C_i from $\mathcal{C} \setminus C_i$. We call \mathcal{C} a convex clustering of S with k clusters C_i .*

Note that in the literature (e.g. [4, 18]) there exists the term “convex k -clustering” which refers to a convex clustering into k clusters such that the clusters have equal size. In this section however, the clusters of a convex clustering may have different sizes.

Bárány and Valtr [4] have shown that every set of n points (n large enough) in general position contains a convex k -clustering of size at least $\varepsilon_k \cdot n$. The constant ε_k , which shrinks doubly exponential with k in the original proof, was later improved by Pór and Valtr [18]. Although the theorem is stated in [18] for the total size of the clusters, they really prove a bound for the size of each cluster, which is better suited for our purposes. The following theorem summarizes these results.

Theorem 6 ([4, 18]). *For any $k \geq 3$, every finite set S of $n \geq \text{ES}(k)$ points in the plane in general position contains a convex k -clustering C_1, \dots, C_k for which $|C_i| \geq 2^{-32k} \cdot n$, for every $i = 1, \dots, k$. ($\text{ES}(k)$ is the function from the Erdős-Szekeres Theorem [9].)*

In Section 4.2.1, we present a construction of a 2-convex polygon P^* of size at least $\frac{\text{ld}^2 n}{128}$, for sufficiently large n . To this end, we need the so-called Cap/Cup Theorem.

Definition 2 (Cap/Cup). *A (finite) set X is a cap with respect to some point r if $X \cup \{r\}$ is in convex position. X is a cup with respect to r if no triple of X forms a convex 4-set together with r .*

In their seminal paper [9], Erdős and Szekeres studied the following problem: How many points (in general position) are necessary to guarantee that at least k of them form either a k -cup or a k -cap (with respect to a given point)? They proved a lower bound of $\binom{2k-4}{k-2} + 1$ on the number of necessary points. For details, see also [20].

4.2.1. Construction of P^*

In this section we give an overview of our construction of the 2-convex polygon P^* . Note that we will slightly refine the construction in Section 4.2.3 to obtain our bound.

Let \mathcal{C} be a convex clustering of S with k clusters. Note that the linear separability condition imposes a cyclic order on the clusters, and let C_1, \dots, C_k be the corresponding counterclockwise (ccw.) order. Without loss of generality, let C_k be the smallest cluster. Choose an arbitrary point r of C_k (we

do not consider the other points of C_k). Let l_i be a line that separates the cluster C_i from the other clusters. For $1 \leq i < k$, order the points of C_i ccw. around r . Let M_i be a monotone (sub-)sequence of the ordered points in C_i , decreasing with respect to the distance to l_i . Let $P_i \subseteq M_i$ be a cap or cup with respect to r , and let \overline{P}_i be the polyline resulting from connecting the points of P_i in ccw. order around r .




Figure 3: Overview of the construction of polygon P^* .

Consider the edges $e_{i,i+1}$, $0 \leq i < k$, that connect the last point of P_i with the first point of P_{i+1} , with $P_0 = P_k = \{r\}$. We call these edges *cluster-connecting edges* and direct them from P_i to P_{i+1} . The polylines \overline{P}_i together with the cluster-connecting edges form a simple polygon P^* . See Figure 3 for a schematic overview of the construction.

4.2.2. P^* is 2-Convex

As all points of P^* are ordered ccw. around r , P^* also has this orientation. Additionally, we let any line l which does not contain r be directed with respect to the same order around r . That is, r always lies to the left of l ; see Figure 4 (left). Using this orientation of lines, we say that the line l^+ intersects l_i in a *positive angle* and the line l^- intersects l_i in a *negative angle* with the lines l_i , l^+ , and l^- as shown in Figure 4 (right).

Before proving the 2-convexity of P^* , we show a structural property of the constructed polygon.

Lemma 6. *For every P_i , $0 < i < k$, let \overline{P}^i be the part of $\partial(P^*)$ from the first vertex of P_{i+1} up to r (in ccw. order around r). Then the points of \overline{P}^i are sorted in ccw. order around each point of P_i .*




Figure 4: Left: Line l is directed such that it intersects ray a before ray b with respect to the order around r . Right: Line l^+ intersects the line l_i in a *positive angle*; l^- intersects l_i in a *negative angle* with respect to the orientation of the lines.

Proof. Consider Figure 5. If $i = k - 1$, then $\overline{P^i}$ consists only of the point r and the lemma is trivially true. Let P_j be some cap/cup “after” P_i , $i < j < k$, and let q be the first point of P_{j+1} . For each vertex v of P_j , the next vertex w (possibly $w = q$) has to be next in ccw. order around r and “below” the line parallel to l_j and going through v . (Recall that P_j is a decreasing sequence with respect to l_j .) It is easy to see that w will also be after v in a cyclic (ccw.) order around any vertex in P_i (in fact, around any vertex in the gray area depicted in Figure 5). Since this line of argument holds for any j , $i < j < k$, regardless of whether $w \in P_j$ or $w \in P_{j+1}$, the lemma follows. \square

To prove that P^* is 2-convex, we use Lemma 5. First we prove that P^* has no inner tangents. Inner tangents cannot exist within a single cap/cup (the end points of a cap cannot be used for an inner tangent, as by construction the first point of a cap is always convex in P^*). Thus assume that there exists an inner tangent t between P_i and P_j . Without loss of generality, let $i < j$. Observe that t has a negative angle to l_j (as each line from P_i to P_j has a negative angle to l_j). Recall that P_j is a decreasing sequence with respect to l_j , and that both $e_{j-1,j}$ and $e_{j,j+1}$ have to intersect l_j . Thus each inner tangent to P_j has a positive angle to l_j . This is a contradiction to the negative angle of t .

It remains to show that P^* has no 4-stabber containing an inflection edge. Because of the construction of P^* , the first point of every cap/cup is always convex. The same is true for r . Consider some cap/cup P_i . Only the first and last edge of P_i and the edge $e_{i,i+1}$ can be inflection edges; see Figure 6. (The edge $e_{i-1,i}$ will be considered with P_{i-1} .) Let l_e be the supporting line of such an inflection edge e .




Figure 5: Consistent ccw. orientation in the proof of Lemma 6.

This leads to 4 cases: (1) $e = e_{i,i+1}$, (2) e is the last edge of a cup, (3) e is the last edge of a cap, and (4) e is the first edge of a cup. Again see Figure 6. Note that the first edge of a cap cannot be an inflection edge.

For cases (1) and (2) observe that the only points of P^* on the right side of the supporting line of $e_{i,i+1}$ are points of P_i . Further, if P_i is a cup, then all points of P_i that are not points of e lie to the right of l_e . Thus, apart from e , l_e can only intersect $e_{i-1,i}$. Therefore, l_e is a 2-stabber. If P_i is a cap, then P_i and $e_{i-1,i}$ lie in convex position. Thus, l_e is a 2-stabber as it can intersect either $e_{i-1,i}$ or one edge of P_i .

For cases (3) and (4) observe that the edge $e_{i,i+1}$ is to the right of l_e , and r is to the left of l_e . Thus we only have to consider the part of $\partial(P^*)$ that starts with the first point of P_{i+1} and ends with r . By Lemma 6, this part of $\partial(P^*)$ is in ccw. order around any point of P_i and thus l_e intersects there only once. Therefore, l_e is also a 2-stabber for cases (3) and (4).

As the supporting lines of all inflection edges are 2-stabbers and no inner tangent exists, we can conclude the following theorem:

Theorem 7. *Any polygon P^* constructed as described in Section 4.2.1 is 2-convex.*




Figure 6: The different cases for inflection edges at a cap/cup P_i .

4.2.3. Concerning the size of P^*

For a first simple proof of the asymptotic bound in Theorem 5 let S be a set of $n = 2^{36k}$ points. By Theorem 6, S contains a convex k -clustering C_1, \dots, C_k with $|C_i| \geq \frac{n}{2^{32k}} = 2^{4k}$ for $1 \leq i \leq k$. We fix a point r in C_k and remove the remaining points of C_k .

For each $i = 1, \dots, k-1$ we apply the Erdős-Szekeres Monotone-Subsequence Theorem [9] on C_i and obtain a subset M_i of C_i of size (at least) $\sqrt{|C_i|} = 4^k$ such that the points of M_i (listed in ccw. order as seen from r) have either increasing or decreasing distances to a line l_i separating the cluster C_i from the other clusters. At least $\frac{k}{2}$ of these sets M_i have, say, decreasing distances. (If there are more increasing sequences we change the order around r from counterclockwise to clockwise. Then the formerly increasing sequences become decreasing and the rest of the construction remains unchanged.)

For each of these $\frac{k}{2}$ sets, we apply the Erdős-Szekeres Cap/Cup Theorem [9], obtaining always a cap or cup (with respect to r) of size at least $k = \frac{\text{ld} n}{36}$. For the other sets M_i , we fix $P_i \subset M_i$ of size 1. The union of these caps/cups (and the sets of size 1) is a 2-convex subset of S of size at least $\frac{k^2}{2} = \frac{\text{ld}^2 n}{2 \cdot 36^2}$. Theorem 5 follows.

Using more sophisticated arguments, we can improve the constants in the construction to $8k^2 - 4k = \frac{\text{ld}^2 n}{128} - \frac{\text{ld} n}{8}$. To this end, we let S be of size $n = 2^{32k}$ for the next two lemmas, with $k \geq 3$. Further, we need a result implicitly

contained in the proof of Theorem 4 in [18]. We state the needed result in the next lemma and reproduce the relevant parts of the proof from [18].

Lemma 7 (Implicit result from [18]). *Let S be any set of n points in general position, with $n = 2^{32k}$ and $k \geq 3$. There exists a convex clustering \mathcal{C} of S with $2k$ clusters such that $\prod_{i=1}^{2k} (|C_i|) \geq \left(\frac{n}{2^{16k}}\right)^{2k}$.*

Proof. Let S' be a subset of 4^{4k} randomly and uniformly chosen points of S , and let S_{4k} be a subset of $4k$ randomly and uniformly chosen points of S' . From the best known upper bounds [9, 20] for the Erdős-Szekeres Theorem, it follows that for any set of 4^{4k} points in general position there exists a subset of $4k$ points in convex position. Thus S_{4k} is in convex position with probability at least $1/\binom{4^{4k}}{4k}$. Clearly, every subset of $4k$ points of S is chosen for S_{4k} with equal probability. Therefore, the number of $(4k)$ -subsets in convex position contained in S is at least $\frac{\binom{n}{4k}}{\binom{4^{4k}}{4k}}$.

For any subset S_{4k} in convex position, let $S_{2k} \subset S_{4k}$ be a subset of $2k$ points in convex position. If the points of S_{4k} , sorted in ccw. order, alternately belong to S_{2k} and $S_{4k} \setminus S_{2k}$, we say that S_{2k} *supports* S_{4k} . Clearly, S_{4k} is supported by two different subsets.

Since S has $\binom{n}{2k}$ subsets of size $2k$ there exists a subset S_{2k}^* that supports at least

$$\begin{aligned} 2 \cdot \frac{\binom{n}{4k}}{\binom{4^{4k}}{4k} \cdot \binom{n}{2k}} &> \frac{(n-4k)^{2k} \cdot \frac{(2k)!}{(4k)!}}{\binom{4^{4k}}{4k}} > \frac{(n-4k)^{2k} \cdot (2k)!}{(4^{4k})^{4k}} > \\ \frac{(n-4k)^{2k} \cdot (k+1)^k}{(2^{16k})^{2k}} &= \left(\frac{(n-4k) \cdot \sqrt{k+1}}{2^{16k}} \right)^{2k} \stackrel{k \geq 3}{\geq} \\ \left(\frac{(n-4k) \cdot 2}{2^{16k}} \right)^{2k} &= \left(\frac{n}{2^{16k}} + \frac{(2^{32k} - 8k)}{2^{16k}} \right)^{2k} > \left(\frac{n}{2^{16k}} \right)^{2k} \end{aligned}$$

subsets of S of $4k$ points in convex position.

Let p_1, \dots, p_{2k} be the points of such a subset S_{2k}^* , listed in ccw. order. Then for all $1 \leq i \leq 2k$, the three lines spanned by $p_{i-1}p_i$, $p_i p_{i+1}$, and $p_{i+1} p_{i+2}$, respectively (where p_α is actually $p_{((\alpha-1) \bmod 2k)+1}$), define a (possibly infinite) region T_i ; see Figure 7. Let the cluster C_i be the subset of points of S that lie inside T_i . Since each cluster C_i lies completely inside the region T_i , it is easy to see that there exists a linear separator l_i for each of these clusters.




Figure 7: The subset S_{2k}^* with points p_1, \dots, p_{2k} listed in ccw. order, with the corresponding regions T_i .

If S_{4k} is a subset supported by S_{2k}^* , then $S_{4k} = S_{2k}^* \cup \{x_1, \dots, x_{2k}\}$, where $x_i \in C_i$. Thus, S_{2k}^* supports at most $\prod_{i=1}^{2k} (|C_i|)$ subsets of S of $4k$ points in convex position. This proves

$$\prod_{i=1}^{2k} (|C_i|) \geq \#(\text{subsets supported by } S_{2k}^*) \geq \left(\frac{n}{2^{16k}}\right)^{2k}. \quad \square$$

Lemma 8. *The size of the polygon P^* is at least $\frac{\text{ld}^2 n}{128} - \frac{\text{ld} n}{8}$.*

Proof. Recall the construction of P^* described in Section 4.2.1 and that $n = 2^{32k}$. From Lemma 7 we know that for each S there exists a convex clustering \mathcal{C} with $2k$ clusters C_i , such that $\prod_{i=1}^{2k} (|C_i|) \geq \left(\frac{n}{2^{16k}}\right)^{2k}$. For the construction of P^* we lose the smallest cluster C_{2k} as we pick only point r from it. We get

$$\prod_{i=1}^{2k-1} (|C_i|) \geq \left(\frac{n}{2^{16k}}\right)^{2k-1}.$$

From the Erdős-Szekeres Monotone Subsequence Theorem [9], we know that for each C_i there exist increasing (IS_i) and decreasing (DS_i) sequences (with respect to l_i and in ccw. order around r) such that $|IS_i| \cdot |DS_i| \geq |C_i|$. We get

$$\prod_{i=1}^{2k-1} (|IS_i|) \cdot \prod_{i=1}^{2k-1} (|DS_i|) = \prod_{i=1}^{2k-1} (|IS_i| \cdot |DS_i|) \geq \prod_{i=1}^{2k-1} (|C_i|)$$

$$\geq \left(\frac{n}{2^{16k}}\right)^{2k-1}, \text{ and thus}$$

$$\max \left\{ \prod_{i=1}^{2k-1} (|IS_i|), \prod_{i=1}^{2k-1} (|DS_i|) \right\} \geq \left(\frac{n}{2^{16k}}\right)^{\frac{2k-1}{2}}.$$

In every cluster C_i , we take the biggest increasing sequence and the biggest cap or cup therein. These caps/cups form the polygon P_I as described for P^* in Section 4.2.1. Similarly, taking the biggest decreasing sequence and the biggest cap or cup therein from every C_i results in a polygon P_D . For P^* we take the larger of P_I or P_D .

Thus P^* is of size at least

$$\begin{aligned} & \max \left\{ \sum_{i=1}^{2k-1} \left(\frac{\text{ld } |IS_i|}{2}\right), \sum_{i=1}^{2k-1} \left(\frac{\text{ld } |DS_i|}{2}\right) \right\} = \\ & \frac{1}{2} \cdot \max \left\{ \text{ld} \left(\prod_{i=1}^{2k-1} (|IS_i|) \right), \text{ld} \left(\prod_{i=1}^{2k-1} (|DS_i|) \right) \right\} = \\ & \frac{1}{2} \cdot \text{ld} \left(\max \left\{ \prod_{i=1}^{2k-1} (|IS_i|), \prod_{i=1}^{2k-1} (|DS_i|) \right\} \right) \geq \frac{1}{2} \cdot \text{ld} \left(\left(\frac{n}{2^{16k}}\right)^{\frac{2k-1}{2}} \right) = \\ & \frac{2k-1}{4} \cdot \text{ld} \left(\frac{2^{32k}}{2^{16k}} \right) > \frac{2k-1}{4} \cdot 16k = \\ & 8k^2 - 4k = \frac{\text{ld}^2 n}{128} - \frac{\text{ld } n}{8}. \quad \square \end{aligned}$$

From this argumentation, Theorem 5 again follows, this time with improved constants.

5. Deciding 2-Convexity of Point Sets

In this section we turn our attention to algorithmic aspects of 2-convexity. We study the problem of deciding whether a point set is 2-convex and show that if a 2-convex polygonization of a point set exists, it can be constructed in polynomial time.

5.1. Preliminaries

We give some preliminary definitions; see Figure 8 for an accompanying illustration. Unless stated otherwise, the edges of a polygon are considered

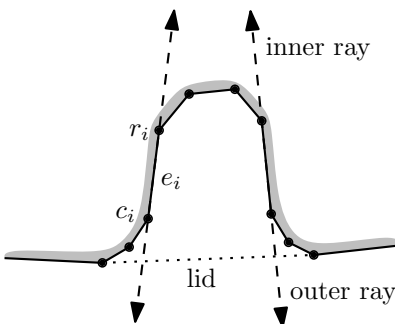


Figure 8: A pocket of a 2-convex polygon. The interior of the polygon is gray. The dotted edge illustrates the lid. The two dashed lines illustrate the two inflection lines of the pocket, consisting of the inner ray, the inflection edge, and the outer ray.

to be directed counterclockwise around the polygon, and all polygonal chains are simple.

Definition 3. A lid of a polygonization of S is an edge of $\text{CH}(S)$ (not necessarily part of the polygonization).

Definition 4. A pocket of a polygon is the polygonal chain between the first and second end-vertex of a lid. A pocket consisting solely of the lid is called a trivial pocket.

Definition 5. The kernel of a simple polygon P is the set of points p such that any ray starting at p crosses ∂P exactly once. If the kernel is not empty, then P is star-shaped.

Definition 6. For an inflection edge e_i , let c_i and r_i denote the convex and reflex vertex of e_i respectively. We partition an inflection line into the inflection edge and two rays; the inner ray, starting at r_i and the outer ray, starting at c_i .

Lemma 9 ([1, Lemma 12]). Given a 2-convex polygonization P of S , let $C = \langle p_1, p_2, \dots, p_t \rangle$ be the chain of vertices that connects (counterclockwise) two consecutive vertices p_1, p_t on $\text{CH}(S)$ (i.e., C defines a pocket). Then the vertices of the chain can be partitioned into three chains $C_1 = \langle p_1, \dots, p_u \rangle$, $C_2 = \langle p_{u+1}, \dots, p_s \rangle$, $C_3 = \langle p_{s+1}, \dots, p_t \rangle$, such that all the elements in C_1 and C_3 are convex vertices of P , while all the elements in C_2 are reflex.

Hence, each non-trivial pocket in a 2-convex polygonization has exactly one pair of inflection edges. The chain of reflex vertices in a pocket of a 2-convex polygon is called the *reflex chain*.

Finally, note that when we say that a point is to the left of a directed edge, we mean that it is to the left of the directed supporting line of that edge.

5.2. Outline of the Algorithm

Recognizing 2-convexity of a point set S can be done in polynomial time if it has a star-shaped 2-convex polygonization. A brute-force approach would be to consider all $\Theta(n^4)$ cells of the arrangement of lines spanned by two points of S as part of the potential kernel. For each choice, the resulting star-shaped polygon can be constructed and checked for 2-convexity in $O(n \log n)$ time [1], resulting in an $O(n^5 \log n)$ algorithm. Hence, for the remainder of this section, we assume that the point set S does not have a 2-convex star-shaped polygonization.

Suppose we have fixed a non-trivial pocket that is part of the 2-convex polygonization. Consider any line ℓ that crosses the pocket exactly twice in such a way that ℓ intersects the pocket in exactly two points (i.e., ℓ does not contain an edge of the pocket). The two crossing points partition ℓ into a segment and two rays. Each of these rays crosses ∂P exactly once, since otherwise ℓ would be a 6-stabber. The key observation for the algorithm, which will be proven formally in Lemma 10, is that if we rotate ℓ in such a way that it always crosses the pocket twice, the order in which ℓ traverses points not in the pocket is the same as the order of these points along ∂P . We look for a triple of pockets that give us the order for all points and show that if the point set has a 2-convex polygonization, but no star-shaped polygonization, such a triple must exist. The polygonization is found by iterating over all triples. Instead of choosing a polygonization of a pocket, we only consider the $O(n^4)$ possible inflection edges for each lid. We show that the choice of the inflection edges suffices to find a 2-convex polygonization, if one exists.

5.3. Observations and Lemmas

Since no inflection line can be a 4-stabber (see Lemma 5), we make the following observations for 2-convex polygons.

Observation 1. *Consider the pair of inflection lines of a pocket. The lines must not cross any other part of this pocket.*

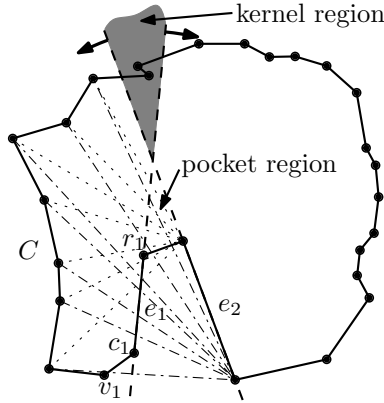


Figure 9: The order of the vertices defined by the inflection edges of a pocket. The gray wedge is the kernel region.

This immediately implies the next observation.

Observation 2. *For the pair of inflection lines of any pocket, an intersection between them occurs either at both the inner or both the outer rays.*

Consider the 2-convex polygon drawn in Figure 9. Any line that passes through a pocket twice can only pass through ∂P two more times. In particular, if such a line also passes through a point of the point set not in the pocket, it separates the neighbors of that point along ∂P . Therefore, the order in which the points appear along the polygonization is constrained by the pocket. We formalize this in the following lemma; see Figure 9 for an accompanying illustration.

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. Without loss of generality, c_1 is left of e_2 . 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 v_1 , the left endpoint of the lid of K). Then the order of the points in C is the same as the radial order around any point p on e_2 . This also holds for any point on e_1 and the points of ∂P to the right of e_1 .*

Proof. We claim that a ray r starting at $p \in e_2$ and contained in the left halfplane of e_2 cannot cross C more than once. Otherwise, consider the supporting line ℓ of r . If p is an extreme point, slightly perturb ℓ such that it

crosses ∂P twice in a neighborhood of p . In any case, ℓ (or its perturbation) crosses the pocket twice, once through e_2 and once to the left of e_2 . There is another crossing with ∂P to the right of e_2 . Crossing C more than once would make ℓ a 6-stabber, contradicting 2-convexity. Therefore the order around p is the same as the order in C . \square

Again let $e_1 = c_1r_1$ and $e_2 = c_2r_2$ be the two inflection edges of a pocket in counterclockwise order. Further, let $\mathcal{H}^-(ab)$ and $\mathcal{H}^+(ab)$ be the closed half-planes to the left and to the right, respectively, of the directed line through the points a and b . We associate two regions to each pocket; see again Figure 9. The *kernel region* of the pocket is the intersection of $\mathcal{H}^-(c_1r_1)$, $\mathcal{H}^+(c_2r_2)$, and, if $r_1 \neq r_2$, $\mathcal{H}^-(r_1r_2)$. For a trivial pocket (i.e., only a convex hull edge), the kernel region is the closed half-plane to the left of it. Analogously, the *pocket region* is the intersection of the half-planes $\mathcal{H}^+(c_1r_1)$, $\mathcal{H}^-(c_2r_2)$, and $\mathcal{H}^-(r_1r_2)$ if $r_1 \neq r_2$; if $r_1 = r_2$, then the pocket region is the empty set. Lemma 10 tells us that once we know one pocket, the remaining polygonization is fixed except for the points of S in the kernel region. The most sophisticated part of our proof will be concerned with determining the pocket (which is, as we will see, relevant when there are no points of S in the kernel region but in the interior of the pocket region).

5.3.1. Pocket Triples

For the next lemma, we need a strong result by Helly.

Theorem 8 (Helly's Theorem [15], [21, p. 70]). *Let F be a finite family of convex sets in \mathbb{R}^n containing at least $n + 1$ members. A necessary and sufficient condition that all the members of F have a point in common is that every $n + 1$ members of F have a point in common.*

Lemma 11. *If a point set S admits a 2-convex polygonization P that is not star-shaped, then there exist three pockets of P that completely determine P .*

Proof. The kernel of a polygon is determined by the intersection of all half-planes to the left of the edges. For each pocket, the kernel region defines this intersection for all the edges of that pocket. Therefore the kernel of P is determined by the intersection of all the kernel regions. Since P is not star-shaped, its kernel is empty. Thus, due to Helly's Theorem, there must exist a triple of pockets such that the intersection of their kernel regions is empty. Since the order in the polygonization is now determined for all vertices due to Lemma 10, the result follows. \square

Checking all triples of possible pockets and the consistency of the implied orders clearly gives us a 2-convex polygonization if one exists. There may, however, exist an exponential number of pocket candidates for any lid. But there are only $O(n^4)$ possible inflection edges per lid. For every pair of inflection edges we distinguish two cases:

- If the kernel region contains points of S , we show that the inflection edges completely determine the pocket; we can then check every pocket triple according to Lemma 11.
- If the kernel region does not contain any point of S , the pocket is not defined for the part in the pocket region; however, any valid pocket with these two inflection edges determines the whole remaining polygonization, and we show how to find such a pocket in polynomial time, if one exists.

We first prove the case with a non-empty kernel. The second case is more involved and is handled in Section 5.3.2.

Lemma 12. *Given only the lid and the inflection edges of an unknown pocket in a 2-convex polygonization, the convex vertices of that pocket are determined.*

Proof. Let v_1 , e_1 , c_1 , and r_1 be defined as in the proof of Lemma 10. If $v_1 \neq c_1$, then there is a triangular region t defined by $\mathcal{H}(c_1r_1)$, $\mathcal{H}^+(v_1c_1)$, and the closed half-plane to the left of the lid. Due to the characterization of 2-convex polygons in Lemma 9, the convex chain between v_1 and c_1 is defined by the convex hull of the points in t (after removing the edge v_1c_1). The second convex chain can be determined symmetrically. \square

Lemma 13. *Suppose a pair of inflection lines of a polygonization P of a point set S defines a kernel region containing points of S . Then the corresponding pocket is determined by the inflection edges.*

Proof. What is left after Lemma 12 is to determine the vertices of the reflex chain. Obviously, all vertices of the reflex chain must be in the pocket region. We claim that all points in the pocket region are in the reflex chain. Suppose there is a point p inside the pocket region that does not belong to the reflex chain. Then the part of the polygonization other than the pocket must pass through p . If it enters and leaves the pocket region through the same




Figure 10: A mighty pocket. The empty dots depict points of T , of which s dominates v .

inflection line, the inflection line is a local 3-stabber, which means that there would exist a 6-stabbing perturbation of the inflection line. Otherwise, each inflection line is traversed once. Since there are still points of S in the kernel region, the polygonization crosses at least one inflection line again, and thus one of the inflection lines could be perturbed to a 6-stabber. Therefore, all points in the pocket region belong to the reflex chain. \square

5.3.2. Mighty Pockets

The more complicated case arises if there are no points of S in the kernel region, but some points in the pocket region. Note that this case may occur when either the inner or outer rays cross, but we do not need to distinguish between these two possibilities. Let $T \subset S$ be the subset of points in the pocket region. Recall that the points in T are the ones for which we do not know the position in a 2-convex polygonization of S . We now have to split T into the vertices of the reflex chain of the pocket and the rest, which then define the part of the 2-convex polygonization that passes through the pocket region but is not part of the pocket. We call the latter the *opposite chain*. It follows from Lemma 10 that after we have correctly split T , the whole 2-convex polygonization is determined. We call such a pocket *mighty*.

Consider two points $s \in S, v \in T \subset S$ and the inflection edges $e_1 = c_1r_1$ and $e_2 = c_2r_2$. Suppose the triangle r_1r_2s contains v . We then say that s *dominates* v . See Figure 10 for an illustration. Note that s might not be an element of T . Nevertheless, we have to check the dominance in order to get a subset of T that contains only non-dominated vertices; as soon as we have decided which of these points should be part of the opposite chain, we know the order in which all points appear along ∂P .




Figure 11: Possible conflicts: an inner tangent ℓ (a) between two vertices on the opposite chain; (b) between a vertex on the reflex chain and a vertex of the opposite chain; and (c) a 4-stabbing inflection line ℓ .

Lemma 14. *If s dominates v , then v has to belong to the reflex chain, and s cannot be part of the reflex chain.*

Proof. Having s in the reflex chain would contradict the chain's reflexivity. The points s and v have a different radial order around r_1 and r_2 ; if none of them were in the reflex chain, these different orders would contradict Lemma 10. \square

Note that the polygonization is already determined for all points not in T . From an algorithmic point of view, this polygonization needs to be checked for 2-convexity, and its pockets may also determine some points that have to belong to the reflex chain and that must go to the opposite chain. A conflict implies that such a polygonization does not exist.

So far, we might not have decided the position of all points of T in the polygonization. There exist configurations with $|T| \in \Theta(n)$ in which any point can be put either to the opposite or the reflex chain, resulting in an exponential number of 2-convex polygonizations. On the other hand, there exist configurations that do not allow a 2-convex polygonization at all. Also note that the two chains might not be linearly separable. In the following lemmas we develop a constructive approach for finding a 2-convex polygonization, if one exists, in polynomial time. More precisely, we try to find a polygonization with the given inflection edges having the smallest possible number of vertices on the reflex chain. See Figure 11 for some illustrations of possible conflicts.

Let an *intermediate* polygonization be a polygonization that fulfills the following two properties:

- The radial order of the points not on the mighty pocket around any point on an inflection edge of the mighty pocket is the same as on the polygonization (in conformance with Lemma 10).

- All points contained in the reflex chain of the mighty pocket have to be in the reflex chain in every 2-convex polygonization of the underlying point set with the chosen inflection edges of the mighty pocket.

In particular, the first property implies that the sub-chains consisting of points not in T are the same in any intermediate polygonization.

The basic idea of the algorithm is to apply a greedy approach. We build an intermediate polygonization with as few points on the reflex chain of the mighty pocket as possible. If it is not a 2-convex polygonization, we find further points that have to be on the reflex chain of the mighty pocket in every 2-convex polygonization of the underlying point set with the chosen inflection edges of the mighty pocket. Then we iterate on the new intermediate polygonization until a 2-convex polygonization is found or there is an unresolvable conflict. We start by showing some properties of intermediate polygonizations.

Let us first consider possible inner tangents in an intermediate polygonization. In the following, let ℓ be an inner tangent defined by the vertices t and t' in an intermediate polygonization (we call t and t' the *contact points* of ℓ). We will see later that the point set only allows a 2-convex polygonization using the two inflection edges of the mighty pocket if both t and t' are points in T . However, to obtain this result, we make no assumptions on t and t' . Lemma 14 gives us the following property.

Corollary 2. *In an intermediate polygonization, the two inflection edges of the mighty pocket are on the same side of an inner tangent ℓ .*

Obviously, we have two different types of inner tangents, one where both t and t' are not part of the mighty pocket in the intermediate polygonization, and one where one point is on the mighty pocket and the other is not. For both types, the following result holds.

Lemma 15. *Let p be any point on an inflection edge of the mighty pocket and, without loss of generality, let the inflection edges of the mighty pocket be below the inner tangent ℓ . Suppose, without loss of generality, that t is not part of the mighty pocket. Then the two neighbors of t are above ℓ and the line through p and t crosses ∂P at t .*

Proof. Recall that there are no points of S in the kernel region of the mighty pocket. The result immediately follows from the fact that the order of all points with respect to any point p on the inflection edges is determined (as




Figure 12: The inflection edges of the mighty pocket are on the other side of ℓ from the neighbors of the contact points (left). The contrary case would disrespect the order induced by the inflection edges (middle and right). The gray regions depict parts of the polygon's interior.

stated in Lemma 10). If the neighbors of t (a contact point not at the mighty pocket) would also be below ℓ , the ray starting at p passing through t would not leave the polygon at t , which contradicts the order determined by the mighty pocket (see Figure 12). As shown in Lemma 10, the supporting line of the ray would be at least a 6-stabber, as it would have to leave the polygon at another point. Similarly, if the neighbors of t are above ℓ but the ray does not leave the polygon at t , there is as well a contradiction with the order of the polygon. \square

We will use the previous lemma to show that both t and t' have to be in the mighty pocket in the final 2-convex polygonization, if there is one. If we do not have an inner tangent but the intermediate polygonization is not 2-convex, then there has to be a 4-stabbing inflection line. The next lemma will allow us to assume a certain structure of the pockets when there is no inner tangent.

Lemma 16. *Consider an intermediate polygonization that contains a pocket with more than two inflection edges. Let e_1 and e_2 be the first and the last inflection edge, respectively, when traversing the pocket. Then either it is the case that one of the two corresponding inflection lines crosses that pocket at least two more times, or there also exists an inner tangent with both tangency points contained in that pocket.*

Proof. The proof is similar to that of Lemma 12 in [1]. Without loss of generality, we assume that the lid of the pocket is horizontal, and the polygon is below it. Let e_1 be the first inflection edge encountered when traversing the pocket counterclockwise. Let p_k be the point of the pocket with lowest




Figure 13: A conflict induced by the inflection line ℓ . The vertices in C are separated from the rest of their pocket by ℓ . If the mighty pocket is involved in the conflict (b), there also exists an inner tangent (dotted).

y -coordinate; p_k is obviously reflex. Let the line ℓ be the supporting line of e_1 . If ℓ crosses the pocket another time we are done. Otherwise, rotate ℓ clockwise keeping it supporting the chain between e_1 and p_k . If there is a convex vertex between e_1 and p_k , then we will find an inner tangent having a contact point at the pocket. The same argument holds for the other side with e_2 . \square

We consider now the situation where there is no inner tangent. Since the radial order of all points not in the mighty pocket is fixed for any intermediate polygonization, the relative position of an inflection edge $e_i = c_i r_i$ and the mighty pocket is determined. We formalize this in the following two lemmas.

Lemma 17. *Consider an intermediate polygonization without inner tangents but containing an inflection edge e_i supported by a 4-stabbing inflection line ℓ . Then ℓ cannot intersect the mighty pocket.*

Proof. Let x_1, \dots, x_k be the sequence of points where the inner ray of e_i crosses ∂P . Now suppose x_j and x_{j+1} are the two crossing points of ℓ with the mighty pocket (see Figure 13 (b)). Again, any ray starting at a point p on an inflection edge of the mighty pocket crossing e_i has to leave the polygon through e_i , which follows from the order induced by the inflection edges of the mighty pocket, as already handled in the proof of Lemma 15. Consider the shortest path inside the intermediate polygon from the convex vertex of e_i to x_{j+1} , the second intersection point of the inner ray with the mighty pocket. This path has at least one left turn and one right turn. Therefore one of the edges of the path would define an inner tangent; a contradiction. \square

Intuitively, if we want to “repair” a situation in which a 4-stabbing inflection line occurs, we move the points that are “cut off” by the inflection

line to the mighty pocket (if possible). We argue about the position of such points in the following lemma.

Lemma 18. *Consider an intermediate polygonization without inner tangents but containing an inflection edge e_i supported by a 4-stabbing inflection line ℓ . The first two crossings of the inner ray of e_i with ∂P partition it into two sub-chains. Among these two chains, let C be the one that does not contain e_i (see Figure 13). Then C is on the same side of ℓ as the inflection edges of the mighty pocket.*

Proof. This again follows directly from Lemma 10, with similar arguments as in the proof of Lemma 17. Without loss of generality, let the mighty pocket be to the right of $e_i = c_i r_i$. Let x_1 and x_2 be the first two crossing points of the inner ray of e_i with ∂P in the order as they occur along the ray. The inner ray leaves and then enters the polygon at these points. (Intuitively, ℓ “cuts off” C at x_1 and x_2 .) Suppose the points of C are to the left of e_i . From Lemma 10 we know that ∂P turns left at r_i . However, the shortest path from c_i to x_2 inside P has to turn right again before reaching x_2 (at some point of C). Hence, there is an inner tangent, a contradiction. \square

We have now obtained enough insight into the structure of the intermediate polygonization to state the main lemmas for assigning the points in T to a chain. For both of the following lemmas, recall the invariant that, for a given choice of inflection edges of the mighty pocket, all points that are in the reflex chain of the mighty pocket in an intermediate polygonization have to be there in any 2-convex polygonization of the underlying point set.

Lemma 19. *Let t and t' be two tangency points of an inner tangent ℓ in an intermediate polygonization. Then both t and t' have to be in the reflex chain in any 2-convex polygonization of the underlying point set (with the given inflection edges of the mighty pocket).*

Proof. See Figure 14. We know that any point that is not fixed is either at the opposite chain or at the reflex chain of the mighty pocket in any 2-convex polygonization, if one exists.

First, suppose that neither t nor t' is part of the reflex chain. Then we know due to Lemma 15 that their neighbors are on the other side of the inner tangent ℓ , say above it, and that any point p on one of the inflection edges of the mighty pocket t defines a line that separates the two neighbors of t . The same holds for p and t' . Let ℓ' be a perturbation of ℓ that is a




Figure 14: Moving a neighbor of t to the reflex chain either preserves the conflict or results in a non-simple polygonization.

6-stabber. To prevent ℓ from being an inner tangent and moving neither t nor t' to the reflex chain, one would have to get rid of some of the edges adjacent to them in some way. Suppose we repolygonize the point set with a neighbor of t now being part of the reflex chain (which is the only way of getting rid of an edge). Both ℓ and ℓ' then would have to cross the reflex chain twice. However, this can only happen either to the left of pt , between pt and pt' , or to the right of pt' , and the reflex chain cannot pass through the rays starting at p more than once. Hence, at most two edges adjacent to t and t' can be removed. However, ℓ' now crosses the reflex chain twice and therefore remains a 6-stabber. Thus, at least one of t or t' must be part of the reflex chain.

Suppose now, without loss of generality, that t' is part of the reflex chain. Again, let the two inflection edges of the mighty pocket be below ℓ and t' be to the right of t . For any point p on the inflection edges, the supporting line of p and t separates the neighbors of t . Therefore only the right edge adjacent to t can be removed by changing the reflex chain while keeping t on the opposite chain, but since the number of times ℓ' crosses the boundary of the polygon is even, ℓ' remains a 6-stabber. Hence, also t has to be part of the reflex chain. The lemma follows. \square

Lemma 20. *Consider an intermediate polygonization without inner tangents but containing an inflection edge e_i supported by a 4-stabbing inflection line ℓ . Let C be a part of a pocket that is separated by ℓ from the polygonization (as in Lemma 17). Then the points of C must be part of the reflex chain of the mighty pocket in any 2-convex polygonization of the point set (with the given inflection edges of the mighty pocket).*

Proof. See Figure 13 (a). Due to Lemma 16 we can assume that the inflection edge e_i that causes the conflict is the first or the last inflection edge

encountered when traversing its pocket. Due to Lemma 17, we know that C is not part of the reflex chain. Suppose we do not want all of the points in C be part of the reflex chain of the mighty pocket. Again, let ℓ' be a perturbation of ℓ that is a 6-stabber. We can now use exactly the same line of argument as in the proof of Lemma 19; the reflex chain would have to pass through ℓ' , but we can only remove at most two edges crossed by ℓ' . The only difference is that the conflict might be resolved after just adding the points of C to the reflex chain of the mighty pocket, but not when adding only the reflex vertex of e_i (and keeping some points of C on the opposite chain). The lemma follows. \square

These lemmas now immediately imply an algorithm for finding a valid reflex chain for a 2-convex polygonization with a given pair of inflection edges for a mighty pocket (i.e., the pair of inflection edges defines a kernel region not containing any point of S). We start with an intermediate polygonization that includes all dominated points on the reflex chain. If we find an inner tangent (in $O(n \log n)$ time), then we add both vertices involved to the reflex chain. If there is no inner tangent but a 4-stabbing inflection line, we add the points of C of Lemma 20 to the reflex chain. During any step we know that all points in the reflex chain have to be there. Hence, we either arrive at a 2-convex polygonization after adding $O(n)$ points, or we cannot change the position of a point, which means that there is no 2-convex polygonization of the underlying point set with the given pair of inflection edges of the mighty pocket. This implies the following lemma.

Lemma 21. *Whether two inflection edges can be completed to a mighty pocket in a 2-convex polygonization using the points of T can be decided in $O(n^2 \log n)$ time.*

5.4. Putting Things Together

The overall algorithm for checking 2-convexity of a point set is the following.

1. Check whether there is a star-shaped 2-convex polygonization by creating the arrangement of all lines defined by two points of the set S . Radially sort the points around a pivot in each cell inside $\text{CH}(S)$ and check all the resulting polygonizations for 2-convexity.

2. For each convex hull edge (i.e., a lid), iterate over all possible inflection edges that have no points of S in the kernel region. Try to construct a mighty pocket giving a 2-convex polygon (see Section 5.3.2).
3. If there is no mighty pocket, check all triples of lid/inflection-edge combinations having points of S in their kernel regions (see Section 5.3.1).

Theorem 9. *2-convexity of a point set can be decided in time polynomial on the size of the point set.*

While we achieved our goal of showing that the problem is solvable in polynomial time, the approach we propose is far from being efficient. Clearly, checking all triples of pocket candidates is the most time-consuming step. There are $O(n^{12})$ choices for the inflection edge combinations. For each pair of inflection edge candidates, there are at most two possible lids, and these can be stored beforehand for every inflection edge candidate. Since we can also store the radial order of the point set around each point of the set, we only need linear time to check whether the orders induced by the inflection edge candidates are compatible. This approach leads to a running time of $O(n^{13})$.

6. Deciding k -Convexity of Point Sets

The algorithm shown in the previous section is quite involved, but has polynomial running time. A natural next step is to consider algorithmic properties when the degree of convexity is increased. This section shows NP-completeness of the problem of deciding whether a point set in the plane allows a 3-convex polygonization. The proof can easily be adapted for any higher degree of convexity.

For ease of presentation we first consider the setting where some edges of the polygonization are fixed and then extend the result to point sets without any fixed edges.

6.1. Fixed Edges

Our proof of the following proposition can be seen as purely instructional, as it is not directly used for showing NP-hardness of the problem without fixed edges. The goal is to give the general idea of the construction, and to address the parts we have to alter when no edges are fixed later.

Proposition 3. *Let S be a set of points in the plane and let E be a set of edges with $E \subset S \times S$. Suppose there exists a polygonization of S that contains all edges of E . Then it is NP-complete to decide whether there exists such a polygonization that is 3-convex.*

Note that the problem is in NP as the k -convexity of a polygon can be decided in quadratic time [1]. Further note that E is required to allow polygonizations of S , as otherwise the problem would be at least as hard as the NP-complete problem of deciding the existence of a polygonization of a set of line segments [19], rendering the result meaningless.

The NP-completeness is shown by reducing 3SAT [12, p. 259] to the problem. We build gadgets using fixed edges that represent the variables, literals, and clauses of a 3SAT formula and show that there exists a 3-convex polygonization if and only if the given formula is satisfiable. We refer to a *literal* as the occurrence of a variable within a single clause (negated or unnegated). Hence, a literal occurs only once in a formula.

For any given 3SAT formula ϕ , let V_ϕ be the set of its variables, L_ϕ the set of its literals, and C_ϕ its set of clauses. Further, let T be a temporary point set in convex position consisting of three disjoint sets T_V , T_L , and T_C (we will later replace them by other points) in which each point corresponds to a variable, literal, or clause of ϕ , respectively. Place the points of T in convex position such that the points of each group are consecutive on the convex hull boundary of T . Further, every triple of points in $T_V \times T_L \times T_C$ should define a triangle that is “roughly equilateral” (this latter informal requirement is intended to ease the presentation of the construction). The literal points are sorted by the variable they represent and unnegated literals of a variable are encountered before the negated literals when traversing the points of T_L on $\partial \text{CH}(T)$ counterclockwise. Between each consecutive pair of the same class, place another temporary point. The set of these points is called T_S . Let $T' = T \cup T_S$.

In the final construction each point in T' is replaced by a corresponding gadget. In order to obtain a valid reduction we have to ensure that we can place the points in polynomial time, in particular, the coordinates of all points need to have a representation that is polynomial in the size of ϕ . One way to do this is to select all the points that are on the convex hull of the final construction among the dense rational points on the unit circle (Canny et al. [5] provide appropriate algorithmic tools), and those from inner points that are adjacent to these points in the final construction on a smaller

circle (where the difference in the radii of the circle depends on the number of gadgets needed). The reader can observe throughout the description of the gadgets that the remaining points can be chosen with respect to the arrangement of the supporting lines of these points.

6.2. Gadgets

The basic idea is that the information in the construction is transported by lines that are potential 8-stabbers, between the variable gadgets and the literal gadgets, and between the literal gadgets and the clause gadgets. We call sets of such lines common to two gadgets a *beam*⁷. More precisely, a beam is defined by the union of potential 8-stabbers through a gadget pair. Hence, the beams “transport” the truth assignment of variables.

We introduce the gadgets by describing their intended behavior. We then show that the gadgets actually have to behave in the intended way. Note that the graphical representations of the gadgets are sketches.




Figure 15: Placement of some gadget (gray), replacing a temporary point t , in order to prevent a line passing through three gadgets.

Every gadget replaces a point $t \in T$ and therefore some of its parts are in extreme position. The gadgets need to be “small” enough such that there is a line through the two edges incident to t that separates the gadget from the remaining domain; see Figure 15. This, in connection with the construction of the gadgets, will ensure that there exists no 8-stabber through gadgets of the same class.

6.2.1. Variables

The variable gadget is shown in Figure 16. The dotted lines in Figure 16 are part of beams leading to literals of the variable, one to an unnegated literal and the other to a negated literal. Note that several beams pass through a variable in this way, one for each literal of the variable. The intended behavior of the variable gadget with assignment “true” is that no line to the

⁷Culberson and Reckhow [7] use a similar terminology.

unnegated literals is a local 3-stabber but the lines to the negated ones are local 3-stabbers, and vice-versa (see Figure 16 (b) and (c), respectively).




Figure 16: A variable gadget: its fixed edges (solid), its intended polygonizations (dashed) for true (b) and false (c), and two potential 8-stabbers in it (dotted).

6.2.2. Literals

The literal gadgets relate each clause to the variables contained in the respective clause. The literals of a variable are placed on neighboring points of T , the unnegated literals below the supporting line of v_3v_4 , and the negated literals above. Figure 17 shows two literal gadgets x_i and $\neg x_j$ and their interaction with their variable x .

In the example in Figure 17, the variable is set to false, as indicated by the dashed edges v_2v_4 and v_1v_3 . Literals are defined by the fixed edges l_1l_3 , l_2l_4 , and l_5l_6 . Their beams are defined by the lines through l_1l_6 and l_3l_5 , and the lines through l_2l_6 and l_4l_5 . Let the former beam be the *clause beam* and the latter the *variable beam*. Note that the edges have to be short enough and need to be placed appropriately such that the variable beam is narrow enough to pass through v_3v_4 . (This can, e.g., be done by first choosing l_1, l_2, l_3 , and l_4 ; the arrangement of supporting lines of the points together with v_3, v_4 and points at the clause gadgets defines a convex region in which the edge l_5l_6 can be placed.)

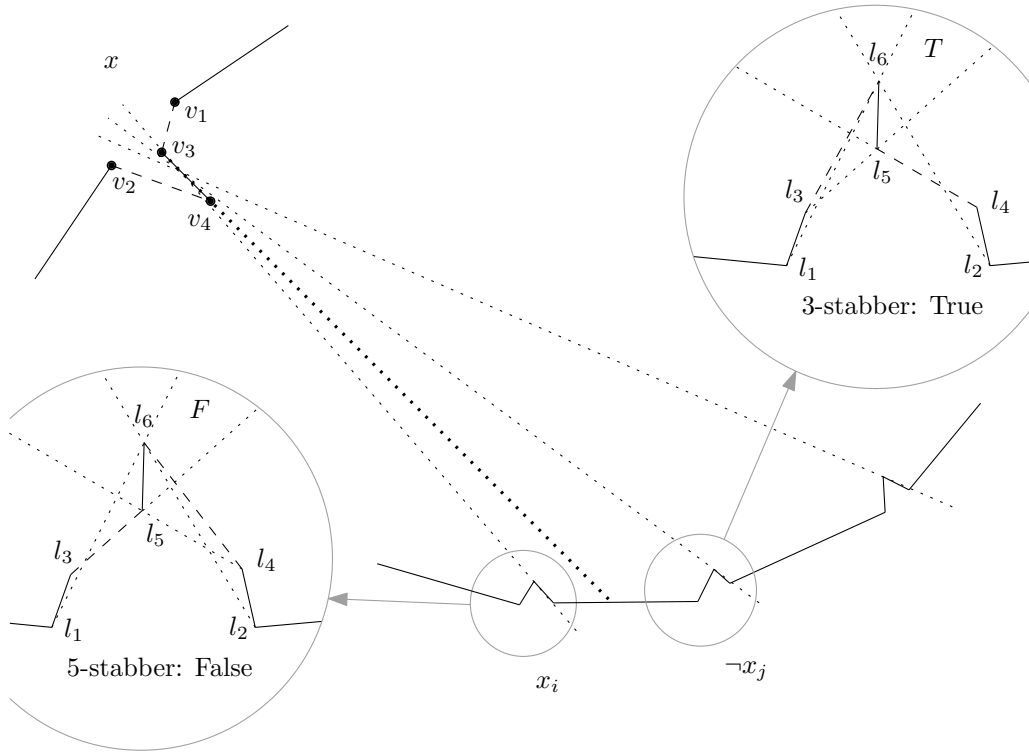


Figure 17: The interaction between literals and their variable.

Consider literal x_i . Its variable beam contains local 3-stabbers at the variable, which is therefore assigned to “false”. Hence, the polygonization of x_i is chosen such that the variable beam does not contain any local 5-stabbers at x_i . The variable beam of literal $\neg x_j$, however, contains no 3-stabbers at the variable. Therefore it can be polygonized the opposite way to x_i . This makes the clause beam of $\neg x_j$ contain no local 5-stabbers at the literal, whereas the clause beam of x_i already contains local 5-stabbers. We define a literal to be “false” if its clause beam contains local 5-stabbers, and true otherwise. The intended behavior of the literal construction is that an unnegated and a negated literal of the same variable cannot both be true (they may both be false but that obviously does not influence the satisfiability properties of the formula).

6.2.3. Clauses

We defined a clause beam to transport an assignment of “false” if it contains lines that are local 5-stabbers at the literal. This means that its lines cannot be local 3-stabbers at the clause. The clause gadget is constructed in a way that it allows a 3-convex polygonization if at least one of the beams does not contain any lines that are local 5-stabbers at the variable, which therefore can also be local 3-stabbers at the clause. We will later show that a 3-convex polygonization exists only if one of the beams of each clause transports “true”.




Figure 18: The clause gadget and its expected polygonizations (a–c). If no beam contains a local 3-stabber at a clause gadget then there has to be an 8-stabber, indicated by the dotted line and the bold-style segments (d).

The clause gadget is shown in Figure 18. The beams pass through fixed edges whose end vertices are placed along two flat arc segments (light gray). If one of the beams is true, a 3-convex polygonization can be done as shown in (a) to (c). Note that in (a) the point c_8 is below the line through c_1c_3 . To have no local 3-stabber in any of the three beams one could sequentially connect the edges, as shown in Figure 18 (d). However, this would introduce an 8-stabber, as depicted by the bold-style segments. The intended behavior of a clause gadget is that there is no 3-convex polygonization if all its literals are “false”. Observe that lines that are local 5-stabbers at the gadget (e.g., a line passing by close to c_8 in Figure 18 (b)) leave any polygonalization in a close neighborhood of the gadget if the two flat arcs on which we place the points are sufficiently close to each other. Hence, no line passing, say, through two clause gadgets can become an 8-stabber. (We give a short account on placing all points such that they have coordinates with a polynomial representation at the end of this section.)

6.3. Necessity of Satisfiability

By construction, such a set of edges allows a 3-convex polygonization whenever the formula is satisfiable. What remains to be shown is that a non-satisfiable formula prevents a 3-convex polygonization; i.e., that the gadgets behave in the intended way. The major difficulty in showing this is that the whole configuration needs to be considered. It is insufficient to inspect the gadgets only locally. We can, however, restrict our attention to the local behavior with the help of a construction we call a separator.

A *separator* gadget is constructed by slightly moving apart the two convex hull edges incident to a temporary point of T_s . The resulting gap is filled by edges as shown in Figure 19. A line in the beam is a local 5-stabber at the separator. Together with the antipodal edge through which the beam passes, such a line becomes a 6-stabber and therefore there cannot be any more edges crossing the separator beam. If we place a separator between all neighboring gadgets (see Figure 20), we may return to our local view.




Figure 19: A separator produces a beam of 6-stabbers between the two dotted lines.

The following simple observations are useful when proving the correct behavior of the gadgets.

Observation 3. *When walking along the boundary of a polygon, any intersection of half-planes is entered as many times as it is exited.*

Observation 4. *A polygonal chain connecting two points separated by a line crosses that line an odd number of times.*

Lemma 22. *Given any pair of lines through a variable gadget, one to an unnegated and one to a negated literal, one of these lines must be (at least) a local 3-stabber.*

Proof. Consider again Figure 16. As the clause is isolated by two separators, it contains a path from v_1 to v_2 , which means that the dotted stabbers are

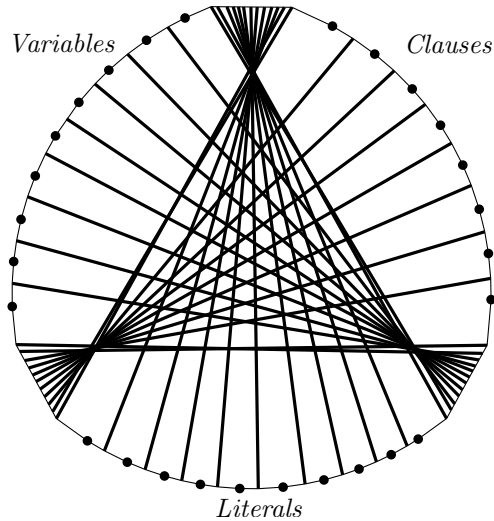


Figure 20: The whole configuration representing a formula. The dots denote the position of the gadgets of the formula, and the thick strokes denote the beams of the separators.

crossed locally an odd number of times. The two lines separate the plane into four regions. Let A be the one that contains v_3 . A is already entered (or left) by the edge v_3v_4 . This means that there has to be another edge leaving A , crossing one of the lines. As that line is crossed twice, it needs to be crossed at least a third time to result in an odd number of crossings. \square

Lemma 23. *Given two literal gadgets, one representing an unnegated and the other one a negated literal of the same variable, at least one of their clause beams contains lines that are local 5-stabbers at the literal.*

Proof. Take any pair of lines, of which one is contained in the clause beam and the other one is contained in the variable beam of the first literal. Arguing analogously to the proof of Lemma 22, it is obvious that one of the lines is a local 5-stabber. If it is in the clause beam, we are done. As it otherwise has to be in the variable beam, we know from Lemma 22 that the variable beam of the second literal contains local 3-stabbers at the variable. Hence, the clause beam of the other literal contains local 5-stabbers. \square

Note that two literals of the same variable might both be set to false, but this obviously does not impose a problem for the overall argumentation. Further note that the proofs of the previous two lemmas are kept quite general,

as we will use similar techniques when proving the correctness of gadgets for point sets.

Lemma 24. *There is no 3-convex polygonization with a clause gadget having all its literals set to false.*

Proof. Consider again Figure 18(d). As the gadgets are divided by separators and all beams contain local 5-stabbers when set to false, the beams define isolated regions. As each region contains only two points, the only choice is to draw an edge between each of these pairs. This, however, yields exactly the polygonal path shown in Figure 18(d). As this path creates an 8-stabber, the proof follows. \square

Proof of Proposition 3. As already discussed, the problem is in NP. For any given 3SAT formula ϕ we can construct a set of edges E in polynomial time representing ϕ . By construction, E has a 3-convex polygonization if ϕ is satisfiable. Lemma 23 shows that an unnegated and a negated literal of the same variable cannot both be true in a 3-convex polygonization of E , and Lemma 24 shows that if all literals of a clause are false, then there cannot be a 3-convex polygonization of E . This establishes that there is a 3-convex polygonization of E only if ϕ is true. \square

6.4. General Point Sets

The proof of Proposition 3 relies on fixed edges. In fact, we did not make use of any isolated points. To transfer the previous result to the domain of point sets, we need a way to force edges to a more or less fixed position.




Figure 21: A chain of at least ten points inside a triangle defined by any three lines implies an edge between two of the points.

Lemma 25. *Let R be any subset of a point set S contained in the triangle defined by three lines. If $|R| > 9$, then there has to be at least one edge between two points in R in any 3-convex polygonization of S .*

Proof. See Figure 21 for an accompanying illustration. Suppose that there is no edge between any two points in R . Then every edge incident to a point in R must cross at least one of the lines defining the triangle region. Every line may only be crossed six times for the polygonization to be 3-convex. As every point is incident to two edges, the bound follows by the pigeonhole principle. \square

Note that this bound may be tightened when considering that there has to be a path between the edges entering and leaving the triangular region. For ease of presentation, at least ten points are chosen.

Let such a subset of ten points along a flat arc segment be called a *bunch*. We will need that the supporting lines of the edges of the bunch lie within a given wedge. For a sufficiently flat arc segment, all the edges spanned by two points of the bunch will fulfill that property. Thus, we can suppose, without loss of generality, that the edge in the statement of Lemma 25 is similar in that sense to an edge which connects two successive points on the arc segment.

During the construction of the gadgets using fixed edges, the positions of the beams were fixed. We show that we can still guarantee the existence of the beams using bunches (but not their exact position). We demonstrate a construction that, by cascading bunches, allows us to place in a defined region any number of edges that all cross a common stabber.




Figure 22: The first two bunches define the width of the hyperbeam (a). Bunches at the intersection of each beam with an arc segment increase the stabbing number of the lines in the part of the beam (b).

The main idea is to replace the fixed edges with flat arc segments on which the bunches are placed. Let a *hyperbeam* be the union of potential beams.

A hyperbeam is defined by these arc segments and is directed in the same way as the beams in the previous section. Now consider the first two arc segments of a hyperbeam, as depicted in Figure 22(a). The length of these two segments defines the width of the hyperbeam, and they should therefore be sufficiently narrow. Place a bunch on each of these two segments. Any pair of edges, one on the first segment and the other one on the second, would define a beam within the hyperbeam defined by the two segments. Each of these beams intersects a part of the third segment. We now place a bunch at each of these intersections, which thus guarantees a local 3-stabber inside the hyperbeam (the intersections might intersect themselves, but it is only necessary that each intersection contains 10 points). This construction can be continued in the same way for all further segments. Hence we may now determine not the exact, but the approximate position of potential 8-stabbers.

6.5. Point Set Separators

In the previous section, the use of separators allowed the correctness of the construction to be verified. Using bunches, we can create a similar construction (as in Figure 23) and therefore prevent other edges from crossing such a separator. However, paths from one side of the domain to the other could still use the edges of such a separator (as shown in Figure 24), since we can no longer say for sure which edges are adjacent. In order to prevent this we apply the following construction, shown in Figure 25. We place two separators (instead of one) indicated by bunches between every formula gadget. The separator beams are directed to the antipodal side of the polygon, intersecting each other. We then place a separator on the antipodal side between the beams. Recall that by using bunches we cannot exactly define the position of a separating line, but can assure its existence somewhere inside the hyperbeam. Hence, such a separator array again allows us to consider the formula gadgets only locally.

6.6. Adaption of the Gadgets

With the help of the bunches, the gadgets can be constructed, up to a certain extent, in a manner similar to the fixed-edges case. When placing the arc segments for the bunches like the fixed edges, a 3-convex polygonization is possible if the formula is satisfiable.

A variable and a literal gadget are shown in Figure 26. Instead of fixing edges, we cascade the bunches as described above. The proof of Lemmas 22




Figure 23: A separator constructed using points. The bold-style arc segments sketch the regions where bunches can be placed, and the dashed line sketches a possible sequence in the polygonization.




Figure 24: An example of an unintended behavior of a separator gadget.

and 23 can be applied directly to these gadgets, since one of the (many) edges on the innermost arc segment of the variable or the literal takes the role of v_3v_4 or l_5l_6 , respectively.

Showing the correctness of the adapted clause gadget is more involved. The sketch in Figure 27 accompanies the description. There, the gray regions depict the hyperbeams carrying the literal assignment. When points are placed along the dark arc segments, there obviously exists a 3-convex polygonization if at least one literal is true. Suppose, without loss of generality, that h is a horizontal line. Recall that a hyperbeam is a union of beams. Place at least one point in the interior of each beam on the arc segments between c_2 and c_3 , c_4 and c_5 , and c_6 and c_7 , respectively. Observe that all of these points are above h . Then place four points in the vicinity of the line h on the arc segment between c_6 and c_7 , and three further points in a similar manner between c_8 and c_9 , such that when sorted along the y -axis no point has its successor on the same arc segment.

Lemma 26. *There is no 3-convex polygonization of the point set with all literals of a point clause gadget set to false.*

Proof. The interesting case is the one where the hyperbeams contain lines that are local 5-stabbers at the corresponding literals. Since the beams may now only be crossed once, every region between two such 5-stabbers is entered

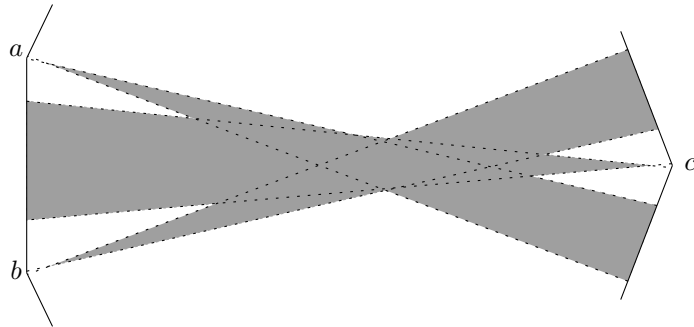


Figure 25: A separator array preventing unintended paths from passing through it.




Figure 26: A point set variable gadget (a) and a literal gadget (b). The gray arc segments denote the position of the bunches, the dotted lines denote a pair of potential stabbers, and the solid segments denote edges that have to exist somewhere within the bunches.

and left once. Consider the dashed line h , which should be an 8-stabber if the clause evaluates to false (as in Figure 27). It is essential for the behavior of the gadget that in this case h is crossed twice between two beams (Property 1). Further, it has to be crossed once in the region of c_1 (Property 2) and three times in the region of c_7 (Property 3).

Property 1: Observe that if all literals for the clause are set to false, there exists a beam that contains an infinite number of 5-stabbers at each of the dark arc segments. Since there is a point placed inside the region of each such beam, there is an infinite number of 5-stabbers through the neighborhood of such a point. This means that the path consisting of the two edges incident to this point has to cross each of these 5-stabbers in that neighborhood. Recall that all these points are placed above h . Therefore, the region between such 5-stabbers (containing the region between two hyperbeams) is entered and left above h by the path defining the overall polygonization. Since the path




Figure 27: A clause gadget. The points to the right ensure the existence of an 8-stabber if all literals are assigned false.

has to “fetch” the point below h , it has to cross h twice within that region.

Property 2: The line h is obviously crossed within the region of c_1 and c_2 , since the gadgets are isolated by separator arrays, and c_1 and c_2 are on different sides of h .

Property 3: The path enters the region above h (by the same arguments as used for Property 1) and leaves it below h , hence h is crossed an odd number of times. Suppose that the path is y -monotone through the points on the arcs. Then the path zig-zags through these seven points, provoking an 8-stabber. If the path is not monotone and crosses h only once, there is a vertex m with both edges leaving it in the same y -direction. Translate h to another horizontal line h' past m . The new line h' crosses the path twice in the vicinity of m , which means that it crosses the path three times. \square

Finally, observe that we can choose the points in such a way that all coordinates are rational and have numerators and denominators that are polynomial in the size of the input. The points on the convex hull of the construction can be selected from the (dense) set of rational points on the unit circle. The additional interior points for the gadgets can be placed inside convex regions defined by the supporting lines of a constant number of pairs of initial points. That is, they can be chosen inside the solution space of linear programs with a constant number of polynomial-sized constraints. Note that the bunches for the hyperbeams do not necessarily have to lie on an arc as in the sketches; they simply have to lie inside a triangle, as demanded by Lemma 25, without increasing the stabbing number in an unintended way.

From these arguments our final theorem follows.

Theorem 10. *It is NP-complete to decide whether a point set is 3-convex.*

7. Open Problems

In this paper we have introduced a new measure of convexity of sets of points in the plane. Due to the novelty of the approach, there are many related open questions. We present some of them.

1. What is the relation between k -convexity and the reflexivity [3] of point sets?
2. Are there examples for general k and j such that the union of a k -convex point set and a j -convex point set is not $(k + j)$ -convex? So far we only have examples for $k = j = 1$.
3. Find $f(k, n)$ such that every k -convex set of n points contains a subset in convex position of size $f(k, n)$. More generally, find $f(k, j, n)$ such that every k -convex set contains a j -convex subset of size $f(k, j, n)$.
4. Find $h(k, n)$ such that every set of n points contains a k -convex set of size $h(k, n)$.
5. Find $H(k, n)$ such that every set of n points contains an *empty* k -convex polygon of size $H(k, n)$ with vertices in the set. That is, there exists a k -convex hole of size $H(k, n)$ in the set. These two questions (4 and 5) are a generalization of the famous Erdős-type problems [9, 10] and from the Erdős-Szekeres Theorem we can construct a 2-convex hole of logarithmic size (see Theorem 3); that is, $H(2, n) = \Omega(\log n)$.
6. Find $g(k, n)$ such that every k -convex set of n points can be partitioned into $g(k, n)$ convex sets. Analogously, find $g(k, j, n)$ such that every k -convex set of n points can be partitioned into $g(k, j, n)$ j -convex sets, $j \leq k$.
7. A triangulation is a decomposition of the convex hull of a point set of size n into $2n - 2 - h$ triangles (where h is the number of points on the convex hull boundary). If we ask for a decomposition into convex polygons, it has been shown that $n - 3 + \lfloor \sqrt{2(n - 3)} \rfloor$ convex polygons are sometimes necessary [11]. What if we use general 2-convex polygons? Is it always possible to decompose a given planar point set with a sublinear number of 2-convex polygons?

Acknowledgements

We thank the anonymous reviewers for their valuable comments that helped to improve the presentation of the results.

References

- [1] O. Aichholzer, F. Aurenhammer, E.D. Demaine, F. Hurtado, P. Ramos, J. Urrutia, On k -convex polygons, *Comput. Geom.* 45 (2012) 73–87.
- [2] O. Aichholzer, F. Aurenhammer, H. Krasser, Enumerating order types for small point sets with applications, *Order* 19 (2002) 265–281.
- [3] E.M. Arkin, S.P. Fekete, F. Hurtado, J.S.B. Mitchell, M. Noy, V. Sacristán, S. Sethia, On the reflexivity of point sets, in: B. Aronov, S. Basu, J. Pach, M. Sharir (Eds.), *Discrete Comput. Geom.: The Goodman-Pollack Festschrift*, volume 25 of *Algorithms Combin.*, Springer-Verlag, 2003, pp. 139–156.
- [4] I. Bárány, P. Valtr, A positive fraction Erdős-Szekeres Theorem, *Discrete Comput. Geom.* 19 (1998) 335–342.
- [5] J.F. Canny, B.R. Donald, E.K. Ressler, A rational rotation method for robust geometric algorithms, in: *Symposium on Computational Geometry 1992*, Proc. ACM, pp. 251–260.
- [6] B. Chazelle, E. Welzl, Quasi-optimal range searching in space of finite VC-dimension, *Discrete Comput. Geom.* 4 (1989) 467–489.
- [7] J.C. Culberson, R.A. Reckhow, Covering polygons is hard, *J. Algorithms* 17 (1994) 2–44.
- [8] P. Erdős, Some more problems on elementary geometry, *Austral. Math. Soc. Gaz.* 5 (1978) 52–54.
- [9] P. Erdős, G. Szekeres, A combinatorial problem in geometry, *Compositio Math.* 2 (1935) 463–470.
- [10] P. Erdős, G. Szekeres, On some extremum problems in elementary geometry, *Ann. Univ. Sci. Budapest. Eötvös, Sect. Math.* 3/4 (1960) 53–62.

- [11] J. García-López, M. Nicolás, A counterexample about convex partitions, in: IV Jornadas de Matemática Discreta y Algorítmica 2004, Cercedilla, p. 213.
- [12] M.R. Garey, D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.
- [13] T. Gerken, Empty convex hexagons in planar point sets, *Discrete and Computational Geometry* 39(1–3) (2008) 239–272.
- [14] H. Harborth, Konvexe Fünfecke in ebenen Punktmengen, *Elemente Math.* 33 (1978) 116–118. In German.
- [15] E. Helly, Über Mengen konvexer Körper mit gemeinschaftlichen Punkten, *Jahresbericht Deutsch. Math. Vereinig.* 32 (1923) 175–176. In German.
- [16] J.D. Horton, Sets with no empty convex 7-gons, *Canad. Math. Bull.* 26 (1983) 482–484.
- [17] C. Nicolás, The empty hexagon theorem, *Discrete and Computational Geometry* 38(2) (2007) 389–397.
- [18] A. Pór, P. Valtr, The partitioned version of the Erdős-Szekeres Theorem, *Discrete Comput. Geom.* 28 (2002) 625–637.
- [19] D. Rappaport, Computing simple circuits from a set of line segments is NP-complete, *SIAM J. Comput.* 18 (1989) 1128–1139.
- [20] G. Tóth, P. Valtr, The Erdős-Szekeres Theorem: Upper bounds and related results, in: J. Goodman, J. Pach, E. Welzl (Eds.), *Combinatorial and Computational Geometry*, volume 55 of *Mathematical Sciences Research Institute Publications*, Cambridge University Press, 2005, pp. 557–568.
- [21] F.A. Valentine, *Convex Sets*, McGraw-Hill, Inc., New York, U.S.A., 1964.