ISRAEL JOURNAL OF MATHEMATICS 113 (1999), 341–379

ON VISIBILITY AND COVERING BY CONVEX SETS*

BY

JIŘÍ MATOUŠEK AND PAVEL VALTR

Department of Applied Mathematics, Charles University
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
e-mail: {matousek,valtr}@kam.mff.cuni.cz

ABSTRACT

A set $X \subseteq \mathbb{R}^d$ is n-convex if among any n of its points there exist two such that the segment connecting them is contained in X. Perles and Shelah have shown that any closed $(n+1)$-convex set in the plane is the union of at most n^6 convex sets. We improve their bound to $18n^3$, and show a lower bound of order $\Omega(n^2)$. We also show that if $X \subseteq \mathbb{R}^2$ is an n-convex set such that its complement has λ one-point path-connectivity components, $\lambda < \infty$, then X is the union of $O(n^4+n^2\lambda)$ convex sets. Two other results on n-convex sets are stated in the introduction (Corollary 1.2 and Proposition 1.4).

1. Introduction

1.1 REVIEW OF RESULTS. Let X be a set in the d-dimensional Euclidean space $\mathbb{R}^d$. Let $\gamma(X)$ denote the minimum (cardinal) number k such that X can be expressed as the union of k convex sets (not necessarily disjoint ones). For a subset $Y \subseteq X$, we write $\gamma_X(Y)$ for the minimum number of convex subsets of X whose union contains Y.

* Research supported by Charles University grants GAUK 99/158 and 99/159, and by U.S.–Czechoslovak Science and Technology Program Grant No. 94051. Part of the work by J. Matoušek was done during the author's visits at Tel Aviv University and The Hebrew University of Jerusalem. Part of the work by P. Valtr was done during his visit at the University of Cambridge supported by EC Network DIMANET/PECO Contract No. ERBCIPDCT 94-0623.
Received May 12, 1996

We investigate the possibility of a "local" characterization of $\gamma(X)$. Namely, suppose that $\gamma_X(P) \leq k$ for any at most c-point subset $P \subseteq X$. Can then $\gamma(X)$ be bounded by some function of c and k (for a given fixed dimension d)? In this paper we obtain a positive answer for $d = 2$ (Corollary 1.2 below). Before we get to this and other new results, let us introduce some terminology and notation and review related results from the literature.

Given $X \subseteq \mathbb{R}^d$, we define the **invisibility graph** of X as the graph $\mathcal{I}(X)$ with vertex set X and with a pair u, v of vertices connected by an edge iff the segment uv is not completely contained in X. We write $\chi_X(Y)$ and $\omega_X(Y)$ for the chromatic number and for the clique number of the subgraph of $\mathcal{I}(X)$ induced by Y, respectively, and we let $\chi(X) = \chi_X(X)$, $\omega(X) = \omega_X(X)$. A set X with $\omega(X) < n$ is traditionally called n**-convex**.

Clearly, we always have $\gamma(X) \geq \chi(X) \geq \omega(X)$. Much effort has been devoted to bounding $\gamma(X)$ in terms of $\omega(X)$. In general, there is no such bound, since there exist planar sets X with $\omega(X) = 3$ but $\gamma(X)$ arbitrarily large. An example is a disk with m one-point holes punctured very close to the boundary and forming the vertices of a regular convex m-gon; see Fig. 1 (this is a modification of Example 1 in [KG71]). The set in this example has many one-point components of the complement. To exclude such a bad behavior, the authors investigating this problem have mostly assumed that X is closed. Valentine [Val57] proved that closed planar sets with $\omega(X) \leq 2$ satisfy $\gamma(X) \leq 3$. Eggleston [Egg74] proved for compact planar sets that if $\omega(X)$ is finite then $\gamma(X)$ is finite as well. Breen and Kay [BK76] were first to show that for planar closed sets X, $\gamma(X)$ is bounded by a function of $\omega(X)$; their bound was exponential. Perles and Shelah [PS90] obtained a polynomial bound $\gamma(X) \leq \omega(X)^6$. Kojman, Perles and Shelah [KPS90] study an infinite version of the problem (which we do not treat here). Concerning lower bounds, Breen and Kay [BK76] reproduce an example, due to Perles, of a closed planar set X with $\omega(X) = n$ and $\gamma(X) \geq \text{const}.n^{3/2}$ (for infinitely many values of n).

For dimension $d \geq 4$, there is no local characterization of $\gamma(X)$, since one may give, for any integers c, K, an example of a set $X \subseteq \mathbb{R}^4$ (in fact, a closed, contractible, and polyhedral set) with $\gamma(X) \geq K$, but $\gamma_X(P) \leq 2$ for any at most c-point subset $P \subseteq X$; see Kojman, Perles and Shelah [KPS90] for a stronger (infinite) version of this example. For $d = 3$, the problem seems to be wide open; we are aware of no positive or negative results. The literature on n-convexity and related problems is quite extensive, and we refer to the above quoted papers for more references.

New results: In this paper we improve the upper bound of Perles and Shelah [PS90] and extend it to a considerably more general class of sets. For a set $X \subseteq \mathbb{R}^2$, we call any path-connected component of $\mathbb{R}^2 \setminus X$ a **hole of** X. For $n \geq 1$, we define $s(n)$ as the maximum of $\gamma(S)$ over all $(n+1)$-convex star-shaped sets $S \subseteq \mathbb{R}^2$. Our upper bound for general sets will be expressed in terms of $s(n)$. The best upper bound we can currently prove for $s(n)$ is $O(n^2)$ (see Section 4), but we conjecture that $s(n)$ is in fact linear.

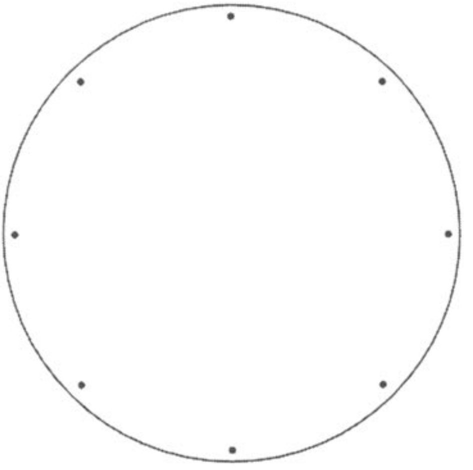

Figure 1. An example of a set X with $\omega(X) \leq 3$, $\gamma(X)$ large.

THEOREM 1.1:

(i) *For $n \geq 2$, any closed $(n+1)$-convex set in $\mathbb{R}^2$ is a union of fewer than $18n^3$ convex sets.*

(ii) *Any $(n+1)$-convex set $X \subseteq \mathbb{R}^2$ with λ one-point holes, $\lambda < \infty$, is a union of at most $O((n^2+\lambda)s(n)) \leq O(n^4 + \lambda n^2)$ convex sets.*

Using Theorem 1.1, one can establish a "local" characterization of $\gamma(X)$ for an arbitrary planar point set in the above mentioned sense:

COROLLARY 1.2: *For any $k \geq 2$ there exist integers K and c such that for an arbitrary set $X \subseteq \mathbb{R}^2$ with $\gamma(X) \geq K$ there is an at most c-point subset $P \subseteq X$ with $\gamma_X(P) \geq k$.*

We also present lower bound examples; for instance, we exhibit closed sets for which $\gamma(X)$ is quadratic in $\omega(X)$, and even in $\chi(X)$.

THEOREM 1.3: *For any two integers $n \geq 3$ and $\lambda \geq 0$, there exists an $(n+1)$-convex set $X \subset \mathbb{R}^2$ with λ one-point holes such that $\gamma(X) \geq c(n^3 + \lambda n)$, for* *an absolute constant $c > 0$. For $n \geq 1$, there exists a closed $(n+1)$-convex set $X \subset \mathbb{R}^2$ with $\chi(X) \leq n$, $\gamma(X) \geq c'n^2$.*

We remark without proof that the bound $c(n^3 + \lambda n)$ in Theorem 1.3 can be still improved to $c(n^3 + \lambda s(n))$, which is tight up to a constant factor for $\lambda = \Omega(n^2)$. As we were informed by Prof. Perles, he earlier obtained an example of a closed $(n+1)$-convex set with $\gamma \geq cn^2$, by modifying his example mentioned above. It is easy to show that $\gamma(X) \geq c\lambda^{1/3}$ for any planar set X with λ one-point holes. Hence, Theorem 1.1 shows, roughly speaking, that n-convex sets with many one-point holes are the only "bad" n-convex sets in the plane with $\gamma(X)$ much bigger than n.

In the proof of Theorem 1.1 we establish an auxiliary result which may be of independent interest. Call a set $X \subseteq \mathbb{R}^2$ **polygonal** if it is bounded and its boundary can be partitioned into finitely many points and open segments, in such a way that each of these parts is either contained in X or disjoint from X. The following proposition basically shows that in the proofs of upper bounds on $\gamma(X)$ in terms of $\omega(X)$ we may restrict our attention to polygonal sets only.

PROPOSITION 1.4: *Let $X \subseteq \mathbb{R}^2$ be a set with λ one-point holes, $\lambda < \infty$, such that $\omega(X)$ is finite and $\gamma(X) \geq K$ for some integer K. Then there exists a polygonal set $Y \subseteq X$ with at most λ one-point holes such that $\omega(Y) \leq \omega(X)$ and $\gamma(Y) \geq K$. If X is closed or star-shaped then also Y may be taken closed or star-shaped, respectively.*

Next, we introduce a technical notion, a certain weakening of n-convexity. A set X is called **weakly n-convex** if there exists a finite subset $H \subseteq X$ such that $X \setminus H$ is n-convex and each $h \in H$ is a one-point hole of $X \setminus H$. Let us remark that there indeed exist weakly n-convex sets that are not n-convex. The idea is to choose an n-point set V in a suitably general position in the plane, and form a set B by placing one point in the interior of each segment uv with $u, v \in V$. Then we have $\omega(\mathbb{R}^2 \setminus B) \geq n$, as is witnessed by the set V of mutually invisible points, but under suitable assumptions on the choice of V and B, it can be shown that $\omega(\mathbb{R}^2 \setminus B \setminus V) < n$. Since we do not need such an example and a rigorous proof is not quite simple, we omit it.

A **pseudotrapezoid** is a path-connected set $T \subseteq \mathbb{R}^2$, such that the intersection of T with any vertical line is connected (it may be empty). Our proof of Theorem 1.1 also relies on the following proposition.

PROPOSITION 1.5:

(i) *Every polygonal weakly $(n+1)$-convex pseudotrapezoid is a union of at most $8s(n)$ convex sets.*

(ii) *Every closed polygonal* $(n+1)$*-convex pseudotrapezoid is a union of at most* $4n$ *convex sets.*

There is a polynomial-time algorithm that decomposes any polygonal set $X \subset \mathbb{R}^2$ with λ holes into at most $O(\omega(X)^4 + \omega(X)^2\lambda)$ convex subsets. It can be generalized also to non-polygonal sets with a "nicely" defined boundary. One such algorithm is based on our proof of Theorem 1.1 and on the proof [BK76] of Theorem 1.6 stated below.

1.2 Open questions.

Dimension 3: One challenging open question is to clarify the possibility of local characterization of $\gamma(X)$ in our sense in dimension 3. Taking the Cartesian product of $\mathbb{R}$ with the set shown in Fig. 1 gives sets in $\mathbb{R}^3$ with $\omega(X) = 3$ but $\gamma(X)$ arbitrarily large. Thus, in $\mathbb{R}^3$ one has to exclude at least sets with holes of dimensions 0 and 1 when bounding $\gamma(X)$ in terms of $\omega(X)$.

Chromatic number: Another problem is to investigate the possibility of bounding the chromatic number $\chi(X)$ in terms of $\omega(X)$ and of bounding $\gamma(X)$ in terms of $\chi(X)$ for arbitrary planar sets X (one can restrict attention only to sets of type $X = K \setminus H$, where K is a convex compact set and H is a set whose every path-connected component is a single point — as can be derived from our proof of Theorem 1.1). Example 1 in Section 6 shows a planar set with $\gamma(X)$ exponentially large in $\chi(X)$. We conjecture, however, that $\gamma(X)$ can be bounded by a function of $\chi(X)$ for arbitrary planar sets X. It would mean (because of the set in Fig. 1; see Example 1 in Section 6) that $\chi(X)$ cannot be bounded by a function of $\omega(X)$.

Star-shaped sets: Recall that $s(n)$ denotes the maximum of $\gamma(S)$ for an $(n+1)$-convex star-shaped set $S \subseteq \mathbb{R}^2$. As was remarked in the introduction, we only know $s(n) = O(n^2)$, but we conjecture that $s(n) = O(n)$. If true, this conjecture would imply that the bound in Theorem 1.1(ii) is tight up to a constant factor. Let us emphasize that for *closed* star-shaped sets, $\gamma(S) \leq 2\omega(S)$ follows from the following precise result:

Theorem 1.6 (Breen and Kay [BK76]): *Let* S *be a closed star-shaped set with finite* $\omega(S)$*, which is star-shaped with respect to a point lying on the boundary of a closed halfplane containing the set* S*. Then* $\gamma(S) = \omega(S)$.

The following example shows that the situation is not so simple for sets which are not necessarily closed. Fig. 2 shows a polygon S with two points b and f

missing from its boundary (the points a, b, d are collinear, as well as the points d, f, g). It is easy to check that this set is star-shaped with respect to the point o and has no 3 visually independent points. On the other hand, $\gamma(S) \geq \chi(S) \geq 3$, since the points a, d, g, e, c in this order form the vertices of an induced 5-cycle in the invisibility graph of S. This proves $s(2) > 2$. A similar construction shows $s(n) \geq \lfloor \frac{3}{2} n \rfloor$ for any $n \geq 2$.

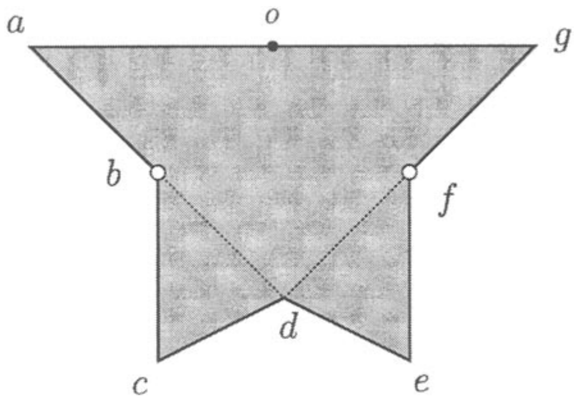

Figure 2. A star-shaped set S with $\omega(S) = 2$ and $\gamma(S) = 3$.

Closed and simply connected sets: There are examples showing that one may have $\chi(X) \geq \lfloor \frac{3}{2}\omega(X) \rfloor$ for closed (star-shaped) polygons (see Fig. 4 in [KG71]), but we have only $\chi(X) = O(\omega(X)^3)$ from Theorem 1.1 as an upper bound for closed sets in the plane. For planar simply connected sets, $\chi(X) = \gamma(X)$ is easy to see and we conjecture $\gamma(X) = O(\omega(X))$. Let us remark that $\gamma(X) \leq \omega(X)s(\omega(X))$ holds for planar simply connected sets: indeed, consider a maximum clique P in the invisibility graph of X, and for each $p \in P$, let S_p be the portion of X visible from p. Each S_p is star-shaped, and the S_p together cover X since P was a maximum clique.

Local conditions for small $\gamma(X)$*:* Examples show that the number K in Corollary 1.2 has to be larger than k. (For instance, consider the union of the n triangles ca_ia_{i+2}, $i = 1, \ldots, n$, where n is an odd integer and a_1 through a_n are vertices of a regular n-gon $a_1a_2\ldots a_n$ with center c.) However, our proof of Corollary 1.2 gives K exponential in k (and $c = 2k - 2$) which is probably not necessary. It would be interesting to prove Corollary 1.2 with K polynomial (or even linear) in k, with c being a function of k.

We remark that some natural (and geometrically defined) classes of graphs were investigated where the chromatic number can be bounded in terms of the clique number (for instance, the circle graphs; see e.g. [Gya87]). These results have a slightly different flavor than ours, however, since the classes of graphs considered there have certain forbidden induced subgraphs, while in our case, any finite graph can be an induced subgraph of the invisibility graph of a (closed) set in the plane.

A similar question as for covering a set by convex subsets can be asked for covering by star-shaped subsets. This question is one of so-called *art gallery problems* (see e.g., [O'R87]). Is it true that if any set of points, up to some bounded size, in an art gallery can be guarded by at most k guards, then the whole gallery can be guarded by at most $f(k)$ guards?

1.3 Organization of the paper, notation and auxiliary results.

The paper is organized as follows. In Section 2 we derive Corollary 1.2 from Theorem 1.1. Section 3 is devoted to the proof of the upper bound on $\gamma(X)$ in terms of $\omega(X)$ (Theorem 1.1) which relies on Propositions 1.5 and 1.4 proved in Sections 4 and 5, respectively. Finally, Section 6 provides examples showing the lower bounds in Theorem 1.3.

The paper is written so that Sections 2–6 may be read more or less independently. For this purpose, we introduce some conventions here. From now on, we deal only with planar sets.

For a point $p \in \mathbb{R}^2$, we write $x(p)$ for the x-coordinate of p, similarly for $y(p)$. For a set $S \subseteq \mathbb{R}^2$, define $x(S) = \{x(a); a \in S\}$ (the projection on the x-axis). For a set $S \subseteq \mathbb{R}^2$, $\overline{S}$ denotes the closure of S, ∂S the boundary of S, and $\operatorname{conv} S$ the convex hull of S.

The **lower envelope** $L(S)$ of S is defined as the portion of the boundary of S visible from below, that is, the set

$$L(S) = \{(a, b) \in \mathbb{R}^2; a \in x(S), b = \inf\{y; (a, y) \in S\}\}.$$

The **upper envelope** is defined analogously.

By a **segment** we always mean a line segment, sometimes we allow a segment to be a single point. We say that a point $p \in \mathbb{R}^2$ **sees** a point $q \in \mathbb{R}^2$ within (or in) a set $X \subseteq \mathbb{R}^2$ if the closed segment pq is fully contained in X. Points $p_1, \ldots, p_k$ are called **visually independent** if no two of them see each other.

By a **component** of a set, we always mean a component of path-connectivity. As is well-known, any two points of the same path-connected components can be connected by a homeomorphic (not just continuous) image of the unit interval.

In Sections 4 and 5 we use the following "compactness" result:

THEOREM 1.7 (Lawrence, Hare and Kenelly [LHK72]): *Let $X \subseteq \mathbb{R}^d$ be a set with $\gamma(X) \geq K$, for some integer K. Then there is a finite set $P \subseteq X$ with $\gamma_X(P) \geq K$.*

Here are some useful observations:

OBSERVATION 1.8: *Let C be any convex set in the plane. Then the following holds:*

(i) *For any point x in the plane, the set $C \setminus \{x\}$ is a union of 2 disjoint convex sets. In particular,*

$$\gamma(C \setminus \{x\}) \leq 2.$$

(ii) *For any finite set F of points in the plane, the set $C \setminus F$ is a union of $|F|+1$ pairwise disjoint convex sets. In particular,*

$$\gamma(C \setminus F) \leq |F| + 1.$$

Proof: (i) is obvious. (ii) follows from (i) by induction on $|F|$. ∎

OBSERVATION 1.9: *Let X be an $(n+1)$-convex set. Then, for any convex set C, $X \cap C$ has at most n components. In particular, the intersection of X with any line consists of at most n disjoint segments.*

Proof: Choosing a point in each component of $X \cap C$ yields a visually independent set in X. ∎

OBSERVATION 1.10: *For any $X \subseteq \mathbb{R}^2$, we have $\omega(\overline{X}) \leq \omega(X)$.*

Proof: If $x, y \in \overline{X}$ do not see each other in $\overline{X}$, there is a $z \notin \overline{X}$ on the segment xy; then also a small neighborhood of z is disjoint from Y, and hence any point x' close enough to x sees no point y' close enough to y within Y. Hence a visually independent set in $\overline{X}$ yields a visually independent set in X. ∎

OBSERVATION 1.11: *If a planar set X has k one-point holes which are vertices of a convex k-gon, then*

$$\gamma_X(P) \geq \frac{k}{2} + 1,$$

for some $(k+1)$-point set P, $P \subset X$.

Proof: Suppose that $h_1, \ldots, h_k$ are k one-point holes which are vertices of a convex k-gon. Fix a point $o \in X$ in the interior of their convex hull, and choose a point $x_i \in X$ on the line oh_i, very close to h_i and on the opposite side of h_i

than o. It is easy to check that the convex hull of any 3 of the x_i's contains some h_j, and hence $\gamma_X(\{o, x_1, \ldots, x_k\}) \geq \frac{1}{2}k + 1$. ∎

2. Proof of Corollary 1.2

Given $k \geq 2$, we consider a set $X \subseteq \mathbb{R}^2$ with $\gamma(X) \geq K$, with $K = K(k)$ sufficiently large, and we want to exhibit an at most c-point subset $P \subseteq X$ with $\gamma_X(P) \geq k$, where c is also a function of k. We give a short proof, which gives $c = 2k - 2$ and is very wasteful in the value of K.

A theorem of Erdős and Szekeres [ES35] says that if P is any sufficiently large set of points in general position in the plane (i.e., no 3 points are collinear), then it contains m points in convex position (i.e., vertices of a convex m-gon), where the required size of P depends on m. An easy corollary of this theorem says that for any k there is an integer $f(k)$ such that any set of at least $f(k)$ points in the plane contains $k-1$ collinear points or $2k-3$ points in convex position. We prove Corollary 1.2 with $K = O(k^4 + k^2 \cdot f(k))$ and $c = 2k - 2$. By Theorem 1.1, we get that any k-convex set $X \subseteq \mathbb{R}^2$ with at most $f(k)$ one-point holes is a union of fewer than K convex sets.

Let $X \subseteq \mathbb{R}^2$ be a set with $\gamma(X) \geq K$. We assume that $\omega(X) < k$, since otherwise the corollary obviously holds. It follows from Theorem 1.1 that X has at least $f(k)$ one-point holes, since otherwise $\gamma(X) < K$. Then, by the definition of $f(k)$, X has $k-1$ one-point holes on a line or $2k-3$ one-point holes in convex position. In the first case we obtain a contradiction with $\omega(X) < k$, in the latter case the corollary follows from Observation 1.11. ∎

3. Upper bound for polygonal sets

In this section, we show that Theorem 1.1 holds for polygonal sets. An application of Proposition 1.4 then completes the proof of Theorem 1.1. Our proof of Theorem 1.1 also relies on Proposition 1.5 (proved in Section 4). We first prove Theorem 1.1 for closed sets and for sets with no one-point holes, and then, in the end of the section, we indicate the changes in the proof needed to obtain Theorem 1.1 also for sets with one-point holes.

The proof has a somewhat similar structure as the proof of Perles and Shelah [PS90] showing $\gamma(X) \leq n^6$ for closed n-convex sets $X \subseteq \mathbb{R}^2$. We also apply Dilworth's theorem on partially ordered sets. Moreover, we use a vertical decomposition similar to some standard techniques in computational geometry (see, for instance, Mulmuley [Mul94]).

Let X be a polygonal $(n+1)$-convex set with no one-point holes. We may assume that X is connected since the bounds in Theorem 1.1 are super-additive functions of n and λ. Let L be a set of segments forming the boundary of X. We assume that no segment in L is vertical, and that the set of all end-points of segments of L contains no pair of points on a vertical line. If this is not the case, then we can simply achieve it by a slight rotation of X.

3.1 NUMBER AND STRUCTURE OF HOLES. Let H be the set of all holes in X. We partition H into two sets H_{seg} and H_{nonseg}, where H_{seg} is the set of segments in H, and H_{nonseg} contains all the other holes in X (recall that we assume X has no one-point holes).

3.1.1 Bounding the number of local extremes of holes

First we investigate the structure of H_{nonseg}. We put

$$X' = X \cup \bigcup H_{seg}.$$

We define points which can be thought of as locally leftmost and locally rightmost points of the complement of X'.

Definition 3.1: A point $e \in \mathbb{R}^2$ is called **left-extremal** for X' if there exists a neighborhood U of e and a component K of $(\mathbb{R}^2 \setminus X') \cap U$ such that e is in the closure of K, and $x(e) \le x(p)$ for all $p \in K$. The **right-extremal** points are defined analogously. Let E denote the set of all (left- or right-) extremal points.

Let $e \in E$. Note that e may or may not lie in X'. An **(e, X')-cone** is defined as a maximal cone C with apex e such that the relative interior of the intersection of C with some neighborhood of e is fully contained in X'. Similarly, an **$(e, \setminus X')$-cone** is defined as a maximal cone C with apex e such that the relative interior of the intersection of C with some neighborhood of e is fully contained in the complement of X'. An **e-cone** is then any (e, X')- or $(e, \setminus X')$-cone. Some of the e-cones may be only semilines emanating from e. For every $e \in E$, e-cones cover the plane, and their relative interiors are pairwise disjoint, which gives a natural circular order on them. With respect to this order, e-cones are alternately (e, X')- and $(e, \setminus X')$-cones. An e-cone is called the **top e-cone** if it contains the semiline emanating from e upwards, and the **bottom e-cone** if it contains the semiline emanating from e downwards. Since we assume that no boundary segment of X is vertical, the top and the bottom e-cones are unique (they may be equal) and have nonempty interiors for every $e \in E$. The **degree** $d(e)$ of a point $e \in E$ is defined as the number of $(e, \setminus X')$-cones that are neither top nor bottom ones. See Fig. 3 for an illustration.

LEMMA 3.2:

$$\sum_{e \in E} d(e) \leq 2n^2.$$

Proof: Let $\mathcal{C}$ be the set of all $(e, \setminus X')$-cones, $e \in E$, that are neither top nor bottom ones. We have to show that $|\mathcal{C}| \leq 2n^2$. The vertex of a cone $C \in \mathcal{C}$ is the leftmost or the rightmost point of C. Let $\mathcal{C}_0$ be the set of cones in $\mathcal{C}$ for which the vertex is the leftmost point; by symmetry, it suffices to show that $|\mathcal{C}_0| \leq n^2$.

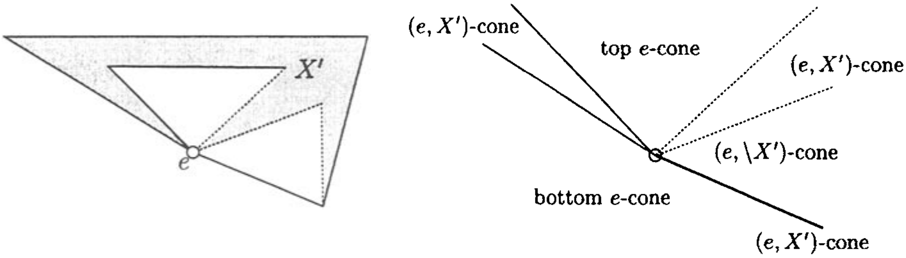

Figure 3. An extremal point of X' and its cones.

For a cone $C \in \mathcal{C}_0$, let $e(C)$ denote the vertex of C, and let $s^+(C)$ be the semiline emanating from $e(C)$ that bounds C from above. For $\varepsilon > 0$, let $q^+(C, \varepsilon)$ be the point on $s^+(C)$ at distance ε from $e(C)$, and let $p^+(C, \varepsilon)$ be either the point $q^+(C, \varepsilon)$ (if $q^+(C, \varepsilon) \in X$) or the point that lies at distance ε^3 vertically above $q^+(C, \varepsilon)$ (otherwise). The semiline $s^-(C)$ and the points $q^-(C, \varepsilon)$ and $p^-(C, \varepsilon)$ are defined similarly, with top and bottom interchanged. We observe that for all sufficiently small $\varepsilon > 0$ we have $q^+(C, \varepsilon), q^-(C, \varepsilon) \in X$ (since X is polygonal); in the sequel we always assume that ε is so small that this holds for all the cones C.

We define a relation $\prec$ on $\mathcal{C}_0$. Let $C_1, C_2 \in \mathcal{C}_0$ be two cones with $x(e(C_1)) < x(e(C_2))$. We distinguish two cases:

- If $e(C_2)$ does not lie on the line $s^+(C_1)$, $C_1 \prec C_2$ holds iff $p^+(C_1, \varepsilon)$ sees $p^+(C_2, \varepsilon)$ within $\mathbb{R}^2 \setminus (C_1 \cup C_2)$ for all sufficiently small $\varepsilon > 0$.
- If $e(C_2)$ lies on the line $s^+(C_1)$, $C_1 \prec C_2$ holds iff $p^+(C_1, \varepsilon)$ sees $p^+(C_2, \varepsilon)$ within X for all sufficiently small $\varepsilon > 0$.

(Fig. 4 illustrates these two cases). These are all cases when $C_1 \prec C_2$ holds, and $C_1 \preceq C_2$ holds if $C_1 \prec C_2$ or $C_1 = C_2$. One can check that $\preceq$ is a partial ordering on $\mathcal{C}_0$.

If $|\mathcal{C}_0| > n^2$, Dilworth's theorem [Dil50] guarantees that $(\mathcal{C}_0, \preceq)$ contains a chain of size $n+1$ or an antichain of size $n+1$.

First, suppose that $C_1, \ldots, C_{n+1}$ is an antichain in $(\mathcal{C}_0, \preceq)$. Then, for $\varepsilon > 0$ small enough, the points $p^+(C_1, \varepsilon), \ldots, p^+(C_{n+1}, \varepsilon)$ are visually independent in X (this works even if $e(C_i) = e(C_j)$ for some i, j), which contradicts $\omega(X) = n$.

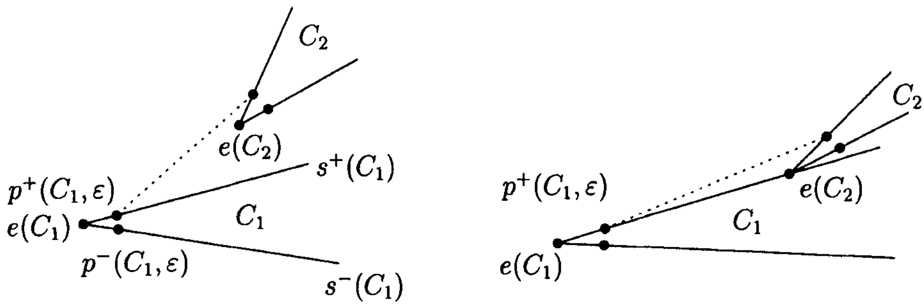

Figure 4. Illustration to the definition of $\prec$ on cones.

Next, let $C_1, \ldots, C_{n+1}$ form a chain in $(\mathcal{C}_0, \preceq)$. Then we claim that the points $p^-(C_1, \varepsilon), \ldots, p^-(C_{n+1}, \varepsilon)$ form a visually independent set in X (for small $\varepsilon > 0$). The reason is that if $C_i \prec C_j$, then a hole h with left extreme $e(C_i)$ obscures the view of $p^-(C_i, \varepsilon)$ to $p^-(C_j, \varepsilon)$. If $e(C_j) \in s^+(C_i)$, we have to use the assumption that h is not a segment. In both cases (chain and antichain) we thus get a contradiction to $\omega(X) = n$. Hence $|\mathcal{C}_0| \leq n^2$. ∎

3.1.2 Bounding the number of segment holes

We use a similar reasoning to examine the size and structure of H_{seg}. We say that two segments h_1 and h_2 in H_{seg} are **brothers** if they lie on a common line ℓ and $\mathrm{conv}(h_1 \cup h_2)$ is disjoint from the closure of every hole of X not fully contained in ℓ. Certainly, the brother relation is an equivalence relation (on H_{seg}). A **brotherhood** is an equivalence class in this equivalence relation.

LEMMA 3.3: *The brother equivalence partitions H_{seg} into at most n^2 brotherhoods, each consisting of at most $n-1$ brothers.*

Proof: Each brotherhood consists of at most $n-1$ brothers since otherwise the line containing the brotherhood contains at least $n+1$ visually independent points of X. Let H_{repr} be the set consisting of the leftmost segments of the brotherhoods. We want to show $|H_{repr}| \leq n^2$.

For each segment $h \in H_{repr}$, let $p(h)$ be its mid-point. For $\varepsilon > 0$, denote by $p^+(h,\varepsilon)$ the point at distance ε above $p(h)$, and by $p^-(h,\varepsilon)$ the point at distance ε below $p(h)$ (thus, $p^+(h,\varepsilon) = p(h) + (0,\varepsilon)$ and $p^-(h,\varepsilon) = p(h) + (0,-\varepsilon)$). As before, for small $\varepsilon > 0$ the points $p^+(h,\varepsilon)$, $p^-(h,\varepsilon)$ lie in X.

We define a partial ordering $\preceq$ on H_{repr}. Let $h_1, h_2 \in H_{repr}$ be segments with $x(p(h_1)) < x(p(h_2))$. We define

- If h_1 and h_2 lie on distinct lines, $h_1 \prec h_2$ holds iff $p^+(h_1,\varepsilon)$ sees $p^+(h_2,\varepsilon)$ within $\mathbb{R}^2 \setminus (h_1 \cup h_2)$ for every sufficiently small $\varepsilon > 0$.
- If h_1 and h_2 lie on a common line, $h_1 \prec h_2$ holds iff $p^+(h_1,\varepsilon)$ sees $p^+(h_2,\varepsilon)$ within X for every sufficiently small $\varepsilon > 0$.

Certainly, $\preceq$ is a partial ordering on H_{repr}. If $h_1,\dots,h_{n+1}$ is an antichain in $(H_{repr},\preceq)$ then the points $p^+(h_1,\varepsilon),\dots,p^+(h_{n+1},\varepsilon)$ are visually independent in X for every sufficiently small $\varepsilon > 0$. If $h_1,\dots,h_{n+1}$ is a chain then it is easily checked that $p^-(h_1,\varepsilon),\dots,p^-(h_{n+1},\varepsilon)$ are visually independent in X for every sufficiently small $\varepsilon > 0$ (for segments lying on a common line we have to use the fact that any two segments in H_{repr} belong to different brotherhoods). Thus, $|H_{repr}| \le n^2$ follows from Dilworth's theorem as in the proof of Lemma 3.2. ∎

3.2 Proof of Theorem 1.1 for sets with no one-point holes. It appears to be much easier to prove Theorem 1.1 for polygonal sets with no segment holes, which gives Theorem 1.1 for closed sets. This we show in Paragraph 3.2.1. Then, in Paragraph 3.2.2, we refine the proof and obtain Theorem 1.1 for sets with no one-point holes.

3.2.1 The first decomposition step; Theorem 1.1 for sets with no segment holes

We consider the set $X' = X \cup \bigcup H_{seg}$. It is connected since X is connected. We erect vertical segments through all extremal points $e \in E$, extending them both upwards and downwards until they hit a non-segment hole. More formally, for each extremal point $e \in E$, let $v(e)$ be the maximal vertical segment (possibly one-point) passing through e contained in X' except for (possibly) the point e. (Note that $v(e)$ is a single point, e, if both the top and the bottom e-cones are $(e, \setminus X')$-cones.) We denote by X'' the set obtained by deleting all the sets $v(e) \cap X'$ $(e \in E)$ from X' (we remark that $v(e) \cap X'$ is a union of at most two segments). Recall that pseudotrapezoids are defined as path-connected sets whose intersection with any vertical line is connected.

LEMMA 3.4: *The set X'' has at most $4n^2+1$ components; each component is a polygonal $(n+1)$-convex pseudotrapezoid. The set X' is a disjoint union of at most $4n^2$ convex sets and at most $4n^2+1$ $(n+1)$-convex pseudotrapezoids.*

Proof: If the sets $v(e)\cap X'$ are deleted from X' one by one, then the deletion of $v(e)\cap X'$ always results in splitting one component into at most $d(e)+2$ components. Thus, the resulting set X'' has at most

$$1+\sum_{e\in E}(d(e)+1)\le 1+\sum_{e\in E}2d(e)\le 4n^2+1$$

components (by Lemma 3.2). Obviously, the components of X'' are pseudotrapezoids. They are $(n+1)$-convex since X was $(n+1)$-convex. The lemma easily follows since each set $v(e)\cap X'$ is a union of at most two segments. ∎

Proof of Theorem 1.1(i): If X is a closed polygonal $(n+1)$-convex set then $X=X'$ is a disjoint union of at most $4n^2$ convex sets and at most $4n^2+1$ pseudotrapezoids (by Lemma 3.4), and, consequently for $n>2$, it is a union of $4n^2+4n(4n^2+1)<18n^3$ convex sets (by Proposition 1.5(ii)). For arbitrary closed $(n+1)$-convex set, the same bound now follows from Proposition 1.4. ∎

3.2.2 Second step of the decomposition

In this step, we treat the segment holes of X. Below we vary an auxiliary set Y, initially set to $Y=X''$, so that we end up with a set Y' with $\gamma(Y')=O(n^2s(n))$ whose union with certain $O(n^3)$ segments is the set X. This gives $\gamma(X)=O(n^2s(n))$, and Theorem 1.1 then follows quite easily.

After we set $Y=X''$, we vary Y as follows: one by one and in an arbitrary order, we extend each segment $h\in H_{seg}$ at both ends until each end hits either the closure of the complement of the (current) set Y or the relative interior of a segment of H_{seg}, and we delete the extension h' of h thus obtained from Y. (Thus, h' is the largest segment containing h whose relative interior intersects neither $\overline{\mathbb{R}^2\setminus Y}$ nor the relative interior of a segment of H_{seg} different from h.) Denote the obtained set, after all segments of H_{seg} have been treated, by Y'. It is not difficult to see that all components of Y' are pseudotrapezoids. See Fig. 5 for an illustration.

LEMMA 3.5: *The set Y' has at most $3n^3 + 2n^2$ components. All but at most $10n^2 + 2$ of them are convex.*

Proof: For a component C of Y', let $p(C)$ be a point of $\overline{C}$ with the smallest x-coordinate. If there are more than one such point, we choose $p(C)$ as the bottommost one of them. For each C, $p(C)$ is either the left end-point of h' for some $h \in H_{seg}$, or a crossing point of a segment $h \in H_{seg}$ with a segment $v(e), e \in E$, or the left end-point of the lower envelope of a component in X''. One can check that this statement has a "multiset" version in the case when some points $p(C)$ coincide. More precisely, if we define three multisets M_1, M_2, M_3 by

$M_1 = \{\text{the leftmost point of } \overline{h'}; h \in H_{seg}\}$,

$M_2 = \{\text{the crossing point of } h \text{ and } v(e); h \in H_{seg}, e \in E, h \text{ crosses } v(e)\}$,

$M_3 = \{\text{the leftmost point of } \overline{L(D)}; D \text{ a component in } X''\}$,

then the multiplicity of any point p in $M_1 \cup M_2 \cup M_3$ is at least the number of components C in Y' with $p(C) = p$. Therefore, the number of components in Y' is at most $|M_1| + M_2| + |M_3|$.

Clearly, $|M_1| = |H_{seg}| \leq n^3 - n^2$ (by Lemma 3.3) and $|M_3| \leq 4n^2 + 1$ (by Lemma 3.4). Since $|E| \leq 2n^2$ and each segment $v(e)$ is crossed by at most $n - 1$ segments $h \in H_{seg}$ (by Observation 1.9), we have $|M_2| \leq 2n^2(n-1)$. It follows that the number of components in Y' is at most $|M_1| + M_2| + |M_3| \leq 3n^3 + n^2 + 1 < 3n^3 + 2n^2$.

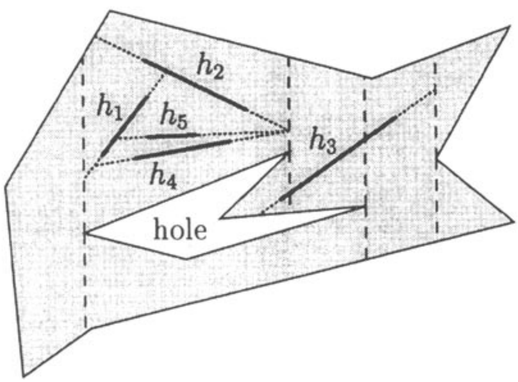

Figure 5. Extending the segment holes one by one.

Each non-convex component of Y' is bounded in part by a non-vertical part of the boundary of a component of X''. The boundary of each component of X'' consists of the upper envelope, the lower envelope, and at most two vertical segments. Parts of the lower envelope of a component C of X'' bound at most $1 + g_\ell(C)$ components of Y', where $g_\ell(C)$ is the number of segments h', $h \in H_{seg}$

(recall that h' denotes the extension of the segment h), that intersect C and have an end-point in an inner point of the lower envelope of C. Similarly for the upper envelope of C. For each brotherhood B with the leftmost and rightmost segments denoted respectively by h_l and h_r, among all end-points of segments $h', h \in B$, only the left end-point of $h_l{}'$ and the right end-point of $h_r{}'$ can lie in the lower or upper envelope of a component of X''. By Lemmas 3.4 and 3.3, it follows that Y' has at most $2(4n^2+1)+2n^2 = 10n^2+2$ non-convex components. ∎

Using Proposition 1.5(i), which is proved in the next section, we can now finish the proof of Theorem 1.1. Each of at most $10n^2+2$ non-convex components of Y' is an $(n+1)$-convex pseudotrapezoid. Thus, we can apply Proposition 1.5(i) for each of them. It follows from Lemma 3.5 that Y' is a union of at most $O(n^2 s(n))$ convex sets so that each point of Y' lies in at most $O(s(n))$ of these convex sets.

The set Y' can be obtained from the set X by deleting at most $O(n^3)$ segments and points. Thus, the set X is a union of fewer than $O(n^3+n^2 s(n)) = O(n^2 s(n))$ convex sets so that no point lies in more than $O(s(n))$ of these convex sets. This completes the proof of Theorem 1.1(ii) when X has no one-point holes.

3.3 Extensions for sets with λ one-point holes. The plan of the proof is as follows. Let H_1 be the set of the one-point holes of X. We want to cover the set $\tilde{X} = X \cup H_1$ by $O(n^2 s(n))$ convex sets in such a way that no point is covered by more than $O(s(n))$ sets. Then X can be covered by $O((n^2+\lambda)s(n))$ convex sets by Observation 1.8(ii).

We consider the weakly $(n+1)$-convex set $\tilde{X} = X \cup H_1$ as suggested above, and go through the proof in this section with $\tilde{X}$ in the role of X. Lemmas 3.2 and 3.3 are proved in the same way, we only observe that by varying ε, we can get infinite families of disjoint $(n+1)$-cliques from the arguments with Dilworth's theorem, hence they work under the weak $(n+1)$-convexity assumption as well. In Lemma 3.4, we obtain polygonal weakly $(n+1)$-convex pseudotrapezoids instead of polygonal $(n+1)$-convex pseudotrapezoids. In Lemma 3.5, we needed $(n+1)$-convexity to conclude that each vertical segment $v(e)$ only intersects $\leq n-1$ segment holes, and again, if it were not the case, we would get an infinite family of disjoint $(n+1)$-cliques. Hence we can cover $\tilde{X}$ by $O(n^2 s(n))$ convex sets so that each point in the plane is covered at most $O(s(n))$ times, and Theorem 1.1(ii) then follows from Observation 1.8(ii). This completes the proof of Theorem 1.1, since its first part was already proved in Paragraph 3.2.1. ∎

4. Covering pseudotrapezoids by convex sets

In this section, we establish Proposition 1.5, stating that each polygonal weakly $(n+1)$-convex pseudotrapezoid is a union of at most $8s(n)$ convex sets and that each closed polygonal $(n+1)$-convex pseudotrapezoid is a union of at most $4n$ convex sets. In the end of this section, we show a quadratic upper bound on $s(n)$, which is a consequence of Theorem 1.6.

COROLLARY 4.1: *Every $(n+1)$-convex star-shaped set is a union of at most $2n^2$ convex sets (i.e., $s(n) \leq 2n^2$).*

Let us call a polygonal pseudotrapezoid T a **hedge** if its lower or upper envelope is formed by a segment (or possibly a point) which is fully contained in T or in the complement of T. Proposition 1.5 is a consequence of the following three lemmas.

LEMMA 4.2: *Every polygonal weakly $(n+1)$-convex pseudotrapezoid T is a union of $(n+1)$-convex hedges $T_1, \ldots, T_k$ such that*

$$\sum_{i=1}^{k} \omega(T_i) \leq 4n.$$

Moreover, if T is closed then such hedges $T_1, \ldots, T_k$ may be found closed as well.

LEMMA 4.3:

(i) *Every hedge T with finite $\omega(T)$ is a union of at most $s(\omega(T))$ convex sets.*
(ii) *Every closed hedge T with finite $\omega(T)$ is a union of $\omega(T)$ closed convex sets.*

LEMMA 4.4: *The function $s(n)$ is superadditive, i.e.,*

$$s(n_1+n_2) \geq s(n_1) + s(n_2).$$

Lemma 4.2 relies on the following observation.

OBSERVATION 4.5: *Each weakly m-convex hedge is m-convex.*

Proof: Let T be a weakly m-convex hedge. Without loss of generality, let the upper envelope of T be a segment fully contained in T or outside T. Let H be a finite subset of T such that $T \setminus H$ is m-convex and each $h \in H$ is a one-point hole of $T \setminus H$. If P is a set of visually independent points in T, then we obtain a set P', $|P'| = |P|$, of visually independent points in $T \setminus H$ by taking P and pushing each point lying in $P \cap H$ slightly downwards within T. ∎

Proof of Lemma 4.2: For $t \in \mathbb{R}$, let $\ell(t)$ be the vertical line $\{x = t\}$. Let $a = \inf(x(T))$ and $b = \sup(x(T))$. We set either $T_a = \emptyset$ (if $a \in x(T)$) or

$T_a = \overline{T} \cap \ell(a)$ (otherwise). Similarly, we set either $T_b = \emptyset$ (if $b \in x(T)$) or $T_b = \overline{T} \cap \ell(b)$ (otherwise). Further, we set $T' = T \cup T_a \cup T_b$.

There is a finite increasing sequence $a = x_0 < x_1 < \cdots < x_r = b$ such that, for $0 \le i < j \le r$, at least one point on $\ell(x_i)$ sees at least one point on $\ell(x_j)$ within T' if and only if $j = i + 1$. To obtain such a sequence, one can take a sequence of the x-coordinates of all end-points of segments forming the boundary of T, and remove, one by one, "redundant" members from it as long as there is any (x_i is redundant in a sequence $a = x_0 < x_1 < \cdots < x_r = b$ if a point on $\ell(x_{i-1})$ sees a point on $\ell(x_{i+1})$ within T'). The lines $\ell(x_1), \ldots, \ell(x_{r-1})$ cut T into r pseudotrapezoids which we denote by $P_1, \ldots, P_r$ (from left to right). For each i, let w_i be the largest integer such that P_i isn't weakly w_i-convex. Then,

$$\sum_{i=1}^{r} w_i = \sum_{i \text{ odd}} w_i + \sum_{i \text{ even}} w_i \le n + n = 2n. \tag{1}$$

Let $i \in \{1, 2, \ldots, r\}$. We cut P_i into two weakly $(w_i + 1)$-convex hedges by an arbitrary segment $s \subseteq T'$ with one end-point on the line $\ell(x_{i-1})$ and the other one on $\ell(x_i)$. Each of the two hedges is $(w_i + 1)$-convex by Observation 4.5. Thus, P_i is partitioned into hedges such that the sum of the clique numbers of their invisibility graphs is at most $2w_i$. The lemma now follows from (1). ∎

Proof of Lemma 4.3: Let H be an $(n + 1)$-convex hedge. Without loss of generality, let the upper envelope of H be a segment which is fully contained in H or in the complement of H. We define (see Fig. 6)

$$T = \{(x, y) \in \mathbb{R}^2; \text{there exists } y_0 \le y \text{ with } (x, y_0) \in H\}.$$

Let v be the point at infinity common to all vertical lines (see Fig. 6). Fix a projective transformation π such that the set $S = \pi(T) \cup \{\pi(v)\}$ is star-shaped

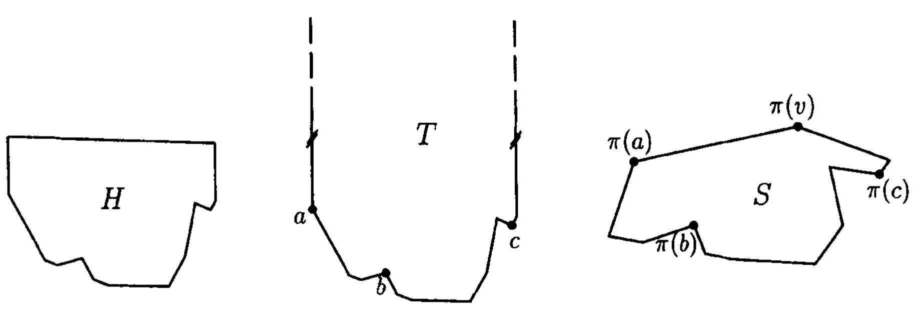

Figure 6. The sets H, T, and S.

with respect to the point $\pi(v)$. It is not difficult to see that $\omega(H) = \omega(T) = \omega(S)$ and $\gamma(H) = \gamma(T) = \gamma(S) \le s(\omega(S)) = s(\omega(H))$. This gives (i).

Now, let H be a closed $(n+1)$-convex hedge. Then the above defined set S is closed. Certainly, the set S is star-shaped with respect to a point lying on the boundary of a closed halfplane containing the set S. By Theorem 1.6, S is the union of at most n convex sets. Part (ii) of the lemma easily follows. ∎

Proof of Lemma 4.4: For $i = 1, 2$, let S_i be an $(n_i + 1)$-convex star-shaped set in $\mathbb{R}^2$ with $\gamma(S_i) = s(n_i)$. By Theorem 1.7, there is a finite set $X_i \subseteq S_i$ with $\gamma_{S_i}(X_i) = s(n_i)$. Translate S_i so that the origin is a kernel of S_i. We may assume that $o \notin X_i$ since $\gamma_{S_i}(X_i) = \gamma_{S_i}(X_i \setminus \{o\})$.

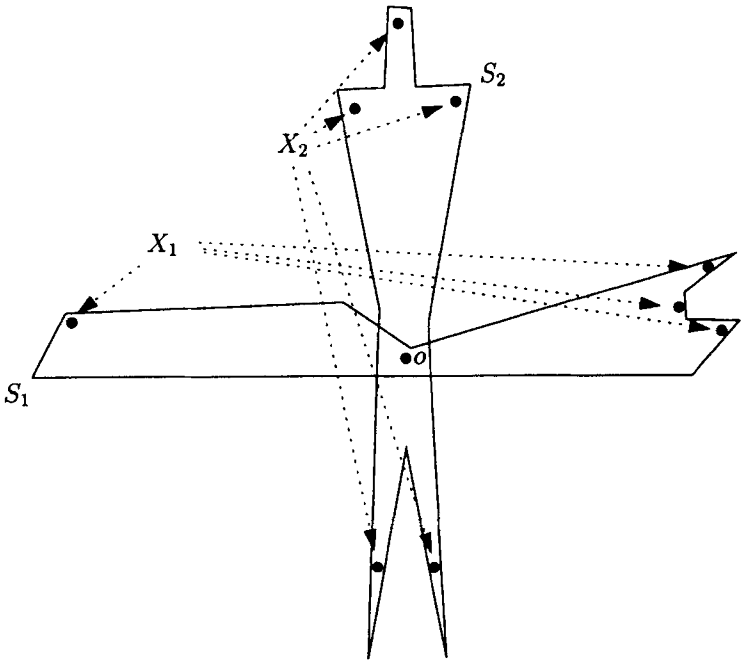

Figure 7. The sets S_1, S_2, X_1, X_2.

Without loss of generality, assume that no point of X_1 lies on the y-axis and shrink the y-coordinates of the points in S_1 so that S_1 becomes a very narrow set around the x-axis. Similarly, assume that no point of X_2 lies on the x-axis and shrink the x-coordinates of the points in S_2 (see Fig. 7). If S_1 and S_2 are sufficiently narrow then no convex set may contain both a point of X_1 and a point of X_2. Hence $\gamma(S_1 \cup S_2) \ge \gamma_{S_1 \cup S_2}(X_1 \cup X_2) = s(n_1) + s(n_2)$. Since $S_1 \cup S_2$ is an $(n_1 + n_2 + 1)$-convex set, the lemma follows. ∎

Proof of Proposition 1.5: Part (ii) easily follows from Lemmas 4.2 and 4.3(ii). For the proof of part (i), suppose that T is a polygonal weakly $(n+1)$-convex pseudotrapezoid. By Lemma 4.2, T is a union of $(n+1)$-convex hedges $T_1, \ldots, T_k$ such that

$$\sum_{i=1}^{k} \omega(T_i) \leq 4n.$$

Observe that we may easily partition the index set $\{1, \ldots, k\}$ into 8 subsets $I_1, \ldots, I_8$ so that

$$\sum_{i \in I_j} \omega(T_i) \leq n$$

for each $j = 1, \ldots, 8$. Then, by Lemmas 4.3(i) and 4.4, each of the 8 sets $\bigcup_{i \in I_j} T_i$ is a union of at most $s(n)$ convex sets. Part (i) of Proposition 1.5 follows. ∎

Proof of Corollary 4.1: Let S be an $(n+1)$-convex star-shaped set. By Observation 1.10, $\overline{S}$ is $(n+1)$-convex as well. From Theorem 1.6 we get that $\overline{S}$ is a union of at most $2n$ convex sets C_i. Certainly, $S = \bigcup_i (S \cap C_i)$. Since $\operatorname{int}(C_i) \subseteq S \cap C_i \subseteq C_i$, the set $\operatorname{int}(S \cap C_i) = \operatorname{int}(C_i)$ is convex. Therefore, $\gamma(S \cap C_i) = \gamma_{S \cap C_i}(\partial(S \cap C_i))$. Observation 1.9 yields that $\gamma_{S \cap C_i}(\partial(S \cap C_i)) \leq n$. Thus, $\gamma(S \cap C_i) \leq n$ for each i, and consequently $\gamma(S) \leq 2n \cdot n = 2n^2$. ∎

5. Polygonal sets suffice

In this section we prove Proposition 1.4.

5.1 Relative convexity and finitely generated sets. Let $Y \subseteq X \subseteq \mathbb{R}^2$. We say that Y is **convex relative to** X if for any two points $x, y \in Y$ such that the segment xy is contained in X, we also have $xy \subseteq Y$. For $A \subseteq X$, we define the **convex hull of** A **relative to** X, denoted by $\operatorname{conv}_X(A)$, as the intersection of all sets $Y \subseteq X$ that contain A and are convex relative to X. We say that X is **finitely generated** if $X = \operatorname{conv}_X(F)$ for some finite set $F \subseteq X$.

Let us remark that $\operatorname{conv}_X(A)$ can be obtained by the following iterative procedure: We let $A_0 = A$, and we define A_{i+1} as the union of all segments connecting pairs of points of A_i and lying fully in X. Then we have $\operatorname{conv}_X(A) = \bigcup_{i=0}^{\infty} A_i$ (both inclusions are straightforward to check).

From now on, let $X \subseteq \mathbb{R}^2$ be a set with λ one-point holes, $\lambda < \infty$, with $\omega(X)$ finite, and with $\gamma(X) \geq K$.

Choose a finite set $P \subseteq X$ with $\gamma_X(P) = K$ as in Theorem 1.7, and let $Y = \operatorname{conv}_X(P)$.

LEMMA 5.1:
(i) $\omega(Y) \leq \omega(X)$,
(ii) $\gamma_Y(P) \geq \gamma_X(P)$, *and*
(iii) Y *has at most* λ *one-point holes.*

Proof: (i) and (ii) are easy to check; we prove (iii). It suffices to show that every one-point hole in Y is a one-point hole in X. For contradiction, suppose that $\{h\}$ is a one-point hole in Y that is not a one-point hole in X. We have $h \in X$. Choose an arbitrary line ℓ through the point h. Since $\{h\}$ is a hole in Y, every segment of ℓ containing h contains infinitely many points of $Y \subseteq X$. Since Y is convex relative to X and $h \notin Y$, every open segment of ℓ containing h contains a point of the complement of X. Thus, points from X and outside X must alternate on ℓ in an arbitrarily small neighborhood of h, but this yields an arbitrarily large visually independent set in X — a contradiction. ∎

We are going to prove the following:

LEMMA 5.2: *Let $X \subseteq \mathbb{R}^2$ be finitely generated in the above defined sense, and suppose that $\omega(X)$ is finite. Then X is polygonal.*

From this Proposition 1.4 follows easily. Given X and K as in that Proposition, we find P as in Theorem 1.7, and we let $Y = \text{conv}_X(P)$. By Lemma 5.1, Y satisfies $\gamma(Y) \geq K$, $\omega(Y) \leq \omega(X)$, and has at most λ one-point holes; by Lemma 5.2 it is polygonal. If X is closed, we may take $\overline{Y}$ instead of Y. This is contained in X, has $\gamma(\overline{Y}) \geq \gamma_X(P) \geq K$, and it also satisfies $\omega(\overline{Y}) \leq \omega(Y) \leq \omega(X)$ by Observation 1.10. If X is star-shaped then we can ensure that Y is also star-shaped by inserting any point of the kernel of X to P. This proves Proposition 1.4. It remains to prove Lemma 5.2.

5.2 SLICING SETS INTO PSEUDOTRAPEZOIDS. Let X be a set with finitely many one-point holes, and let $n = \omega(X)$ be finite. In this section, we show that X can be decomposed into finitely many pseudotrapezoids. The proof is somewhat similar to the proof in Section 3; the difference is that on one hand, we can be wasteful in the number of pseudotrapezoids, but on the other hand, we have to deal with arbitrary sets rather than with polygonal sets.

In this section we will not need X to be finitely generated; we return to it later, when showing that the pseudotrapezoids arising in this decomposition must be polygonal.

We begin with several simple observations.

OBSERVATION 5.3: *Let $\pi \subset \mathbb{R}^2$ be a path, that is, a homeomorphic image of the interval $[0,1]$, let F be a closed set such that $F \cap \pi \neq \emptyset$. Then the intersection $F \cap \pi$ has the first and the last point (in the ordering along π).*

Proof: A standard compactness argument. ∎

OBSERVATION 5.4: *Let F be a (closed) square, let C be a component of $\mathbb{R}^2 \setminus X$. Then $F \cap C$ has at most $4n$ components.*

Proof: If C does not intersect ∂F there is nothing to prove. Otherwise any point of $C \cap F$ is connected to a point of $\partial F \cap C$ by a path contained in $F \cap C$ (by Observation 5.3). By Observation 1.9, ∂F is intersected by X in at most $4n$ intervals, and Observation 5.4 follows. ∎

Let E be the set of all extremal points for X according to Definition 3.1 (i.e., $e \in E$ if there exists a neighborhood U of e and a component K of $(\mathbb{R}^2 \setminus X) \cap U$ such that $e \in \overline{K}$, and $x(e)$ is either the smallest or the largest element in $x(\overline{K})$). Unlike in Section 3., parts of the boundary of X may be vertical, so E may very well be infinite. Note that E contains all one-point holes.

Let us remark that we may require the neighborhood U to be arbitrarily small. Indeed, given **some** neighborhood U as in the definition, we may take a small closed square $F \subseteq U$ with center e. Then by Observation 5.4, $K \cap F$ has only finitely many components, so e is in the closure of one of them.

We now aim at proving that the number of possible x-coordinates of extremal points is finite. First we need a combinatorial result (one used by Perles and Shelah too).

Definition 5.5: Consider a set S of m semilines $\sigma_1, \ldots, \sigma_m$ with end-points $p_1, \ldots, p_m$, respectively, such that $x(p_1) < x(p_2) < \cdots < x(p_m)$, and each σ_i is directed from p_i to the right. We call S **homogeneous** if for each $i = 1, 2, \ldots, m$, all the points p_j with $j > i$ lie all above and on σ_i or all below and on σ_i (see Fig. 8).

Given an integer k, let $g(k)$ be the smallest number such that any set S as above with $m \geq g(k)$ contains a homogeneous subset of size at least k.

A straightforward induction shows that $g(k) \leq 2^{k-1}$. Trivially, $g(1) = 2^0$. For $k > 1$, choose σ_1 as the first element of a homogeneous subset. One of the halfplanes determined by σ_1 contains at least 2^{k-2} of the remaining end-points, and (by induction) we can select the other $k-1$ elements of a homogeneous subset among their corresponding semilines. The proof of Perles and Shelah

[PS90] yields a polynomial bound on $g(k)$; for us, it is sufficient that $g(k)$ is always finite. We may return to investigating extremal points of X.

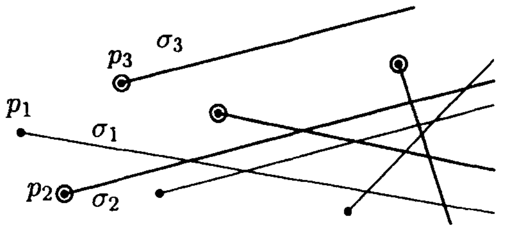

Figure 8. A set of semilines; one of the homogeneous subsets is marked.

LEMMA 5.6: *The set $x(E)$ of all x-coordinates of extremal points is finite.*

Proof: Let E_0 be the set of all points in E that are not one-point holes of X. Since X has finitely many one-point holes, it suffices to show that $x(E_0)$ is finite. Suppose it is not the case, and put $N = 2ng(n+1) + n$, where $g(n)$ is as in Definition 5.5. Fix N distinct x-coordinates in $x(E_0)$, $x_1 < x_2 < \ldots < x_N$. For each i, select an extremal point $e_i \in E_0$ with $x(e_i) = x_i$. Moreover, for each e_i choose a small square F_i centered at e_i and a component C_i of $(\mathbb{R}^2 \setminus X) \cap F_i$, such that F_i and C_i witness that e_i is extremal, as in Definition 3.1. The squares F_i are taken so small that their vertical projections on the x-axis do not overlap.

We now classify the indices i into 3 types, according to the shape of the component C_i. Type 1 are indices i with C_i a vertical segment, type 2 are indices i with C_i a non-vertical segment, the remaining indices are of type 3 (note that since $e_i \in E_0$, C_i is never a point). Now at least one type will have sufficiently many indices (where 'sufficiently many' is specified below), and we distinguish the corresponding cases.

CASE 1. THERE ARE AT LEAST $n+1$ INDICES OF TYPE 1 (=VERTICAL SEGMENTS). For each such i, draw a horizontal right oriented segment of length ε from the mid-point of C_i, where $\varepsilon > 0$ is a suitably small number. This segment certainly contains a point of X (otherwise it would be a part of C_i), call it q_i. If ε is small enough, then no 2 points q_i, q_j see each other in X (for $i < j$, the segment C_j blocks the view from q_j to q_i), a contradiction to $\omega(X) \leq n$.

CASE 2. THERE ARE AT LEAST $ng(n+1)$ INDICES OF TYPE 2 (= NON-VERTICAL SEGMENTS). Let I be the set of these indices. Call two segments C_i, C_j $(i, j \in I)$ **equivalent** if they lie on the same line. Each equivalence class has at most n elements, by Observation 1.9. Pick one i from each equivalence class; this yields a set $I_1 \subseteq I$ of at least $g(n+1)$ indices.

For each C_i, $i \in I_1$, let ℓ_i denote the line containing C_i. Fix an interior point p_i of C_i such that no ℓ_j $(j \neq i)$ passes through it. Draw a segment of length ε with mid-point p_i and perpendicular to C_i, where $\varepsilon > 0$ is a suitably small number. By the same argument as in Case 1 above we get that this segment contains a point $q_i^- \in X$ below C_i and a point $q_i^+ \in X$ above C_i. By choosing ε small enough, we can guarantee that the segment C_i obscures the view of q_i^+ to any q_j^+ or q_j^- such that $j \in I_1$ and p_j lies in the halfplane below ℓ_i, and similarly q_i^- sees no point q_j^+, q_j^- with p_j above ℓ_i.

For each $i \in I_1$, let σ_i be the semiline on ℓ_i lying to the right of p_i, and consider the set $\{\sigma_i; i \in I_1\}$ of semilines. This is as in Definition 5.5, thus, there is a homogeneous subset $I_2 \subset I_1$ of $n+1$ indices. From each pair q_i^-, q_i^+, $i \in I_2$ we choose the point lying in the open halfplane containing none of the points p_j, $i < j \in I_2$ (and thus also none of the points q_j^+, q_j^-, $i < j \in I_2$). These $n+1$ points are visually independent in X, a contradiction.

Case 3. There are at least $2g(n+1)$ indices of type 3 (=non-segments). By symmetry, we may assume that at least $g(n+1)$ indices of type 3 correspond to left-extremal points e_i; let I be the set of such indices. We proceed similarly as in Cases 1 and 2, by defining suitable semilines σ_i for these points.

Let $i \in I$, and suppose that for all $j \in I, j < i$, the semilines σ_j have already been defined, in such a way that no σ_j passes through the point e_i.

Consider the component C_i within the square F_i. Since C_i is not a segment, we can choose points $r, r' \in C_i$ such that $x(r) > x(e_i)$, $x(r') \geq x(e_i)$, and the points e_i, r and r' are not collinear; see Fig. 9. Choose $\varepsilon > 0$ sufficiently small, and fix a point $p_i \in C_i$ at distance at most ε^3 from e_i (this is possible, since $e_i \in \overline{C_i}$). If ε is small enough, p_i is not collinear with r, r' either. Without loss of generality assume that the direction from r to r' around p_i is clockwise, as in Fig. 9.

Let σ_i be a semiline with end-point p_i in the middle third of the angle $\angle r' p_i r$, chosen so that it intersects no point e_j with $j > i$, $j \in I$.

Finally we define points q_i^+, q_i^-. To this end, let t^- be the line with slope $1/\varepsilon$ through p_i, and t^+ the line with slope $-1/\varepsilon$ through p_i. Let q_i^- be a point of X on t^- lying below p_i at distance at most ε from p_i but outside $\mathrm{conv}(C_i \cap t^-)$ (such a point exists, since the points of t^- lying from p_i downwards at distance $\geq \varepsilon$ lie on the left of e_i and therefore don't belong to C_i). Symmetrically we choose q_i^+ on t^+ at distance at most ε upwards from p_i.

We note that if ε is small enough then the squares F_j with $j > i$ lie on the right of t^+ and t^-. We claim that with small enough ε, q_i^- sees no point

s above σ_i, below t^-, and outside F_i. A similar claim will hold also for q_i^+. Once we know this, we may apply the definition of the number $g(n+1)$ (on the inductively constructed set $\{\sigma_i; i \in I\}$ of semilines) and conclude that $n+1$ visually independent points can be selected among the points q_i^+, q_i^-.

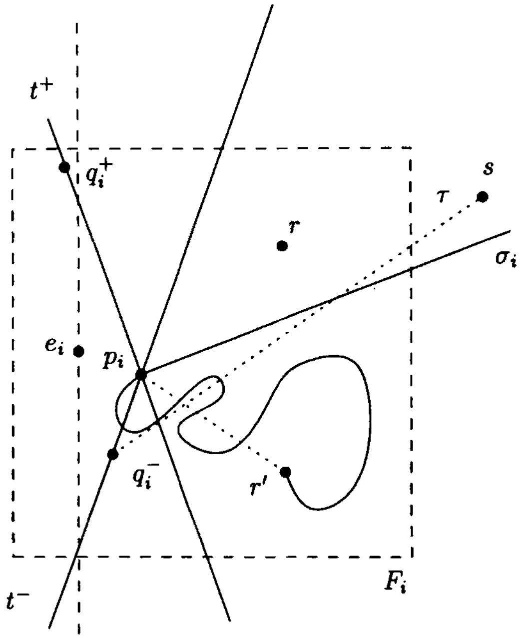

Figure 9. The situation in Case 3.

To establish the above claim for q_i^-, say, we may proceed as follows. We may assume that q_i^- is closer to σ_i than r', and so if q_i^- saw a point above σ_i, below t^-, and outside F_i, the connecting segment τ would intersect the segment $p_i r'$. But r' is connected to p_i by a path within C_i, and τ would have to intersect this path as well, a contradiction.

This finishes the discussion of Case 3. Since at least one of Cases 1, 2, and 3 must occur, Lemma 5.6 is proved. ∎

We now define a decomposition of X into a bounded number of pseudotrapezoids. Let $x_1 < x_2 < \cdots < x_m$ be the x-coordinates of all extremal points. Let v_i be the vertical line $\{x = x_i\}$, and let V_i be the open vertical strip between v_i and v_{i+1}, V_0 and V_m being halfplanes. By Observation 1.9, each $V_i \cap X$ has at most n components; let T be any such component.

Lemma 5.7: *T is a pseudotrapezoid.*

Proof: Suppose it is not the case; then there exists a vertical line v and points $p_1, b, p_2 \in v$ with $y(p_1) < y(b) < y(p_2)$ and $p_1, p_2 \in T$, $b \notin T$; see Fig. 10. Choose a path $\pi \subseteq T$ connecting p_1 to p_2. Let q_1 be the last point of π intersecting the part of v below b, and let q_2 be the first point on π after q_1 intersecting the part of v above b (these points exist, since $\pi \cap v$ is closed and by Observation 5.3).

We note that the portion π' of π between q_1 and q_2 lies completely on one side of v, say on the right. The union of π' with the segment q_1q_2 is a closed simple Jordan curve, which bounds an interior region. We let R denote the closure of this region.

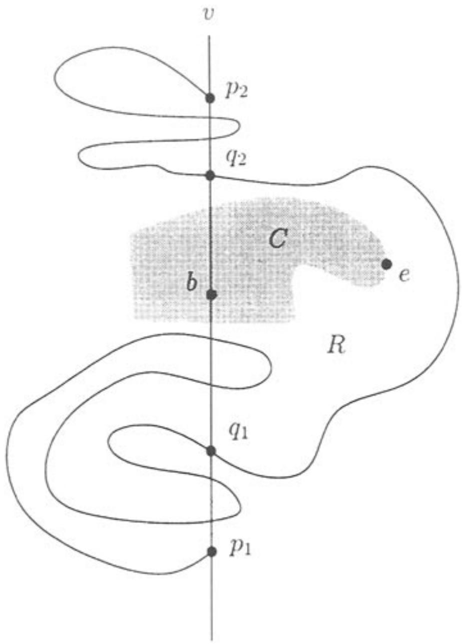

Figure 10. Illustration to the proof of Lemma 5.7.

Consider now the component C of $\mathbb{R}^2 \setminus X$ containing the point b. It only intersects the boundary of R in the segment q_1q_2. Since the intersection of C with the segment q_1q_2 has finitely many components (by Observation 1.9), $R \cap C$ has finitely many components too. It is easy to check that any point $e \in \overline{C \cap R}$ with $x(e) = \sup_{q \in C \cap R} x(q)$ is a right-extremal point for X according to Definition 3.1. This is a contradiction, since e is contained in the open vertical strip V_i containing T, and this strip was constructed so that it contains no extremal point. ∎

5.3 FINITELY GENERATED PSEUDOTRAPEZOIDS ARE POLYGONAL. Consider the subdivision of the plane into vertical lines $v_1, \ldots, v_m$ and the vertical strips $V_0, \ldots, V_{m+1}$. We now recall that the considered set X is finitely generated, i.e. there is a finite point set $P \subseteq X$ with $\operatorname{conv}_X(P) = X$. Let us add the vertical lines through all points of P to the v_i's. In this way, we may assume that the considered pseudotrapezoids in the decomposition of X contain no points of P. Consider one such pseudotrapezoid, T, arising as a component of the intersection of X with an open vertical strip V, which is bounded by vertical lines v and v' (one of them may be "at infinity", if V is a halfplane). Let $V(T)$ be the vertical strip containing all points $p \in \mathbb{R}^2$ with $x(p) \in x(T)$.

LEMMA 5.8: *The set $U \subset V(T)$ of points lying above T is convex. Similarly, the set of points lying below T is convex.*

Proof: We recall that X is the convex hull of P relative to X, and so it can be written as $X = \bigcup_{i=1}^{\infty} X_i$, where $X_0 = P$ and X_{i+1} is the union of all segments connecting pairs of points of X_i and lying fully in X.

Let $a_1, a_2 \in U$ be two points above T. We prove that no point of $X_i \cap T$ lies on or above the segment $a_1 a_2$; this will imply the convexity of U. For $i = 0$, this is clear since $P \cap T = \emptyset$. Suppose that this is true for some i, and let $z \in T \cap X_{i+1}$. By definition of X_{i+1}, there must be two points $z_1, z_2 \in X_i$ such that $z_1 z_2 \subseteq X$ and $z \in z_1 z_2$. We may assume $x(z_1) < x(z_2)$ and $x(a_1) \le x(z) \le x(a_2)$ (the other cases are easy to discuss); see Fig. 11. If $x(z_1) \in [x(a_1), x(a_2)]$, we let $z_1' = z_1$, otherwise we let z_1' be the intersection of $z_1 z_2$ with the vertical line $\{x = x(a_1)\}$. We claim that z_1' always lies below the segment $a_1 a_2$: If $z_1' = z_1$, this is clear from the inductive assumption, and if z_1' is on the vertical line through a_1, it follows from the fact that the segment $z z_1'$ is contained in X and that T is a pseudotrapezoid. Similarly we define z_2' as either z_2 or the intersection of $z_1 z_2$ with the vertical line $\{x = x(a_2)\}$, and we get that z_2' lies below $a_1 a_2$ as well. Therefore z is below $a_1 a_2$ too. ∎

The set T is thus bounded from above by a convex curve and from below by a concave curve. To establish Proposition 1.4, it remains to prove that both the lower and the upper envelopes of T consist of finitely many segments. If this doesn't hold for the lower envelope, say, then there are $n+1$ points $p_1, \ldots, p_{n+1}$ on the lower envelope of T such that each segment $p_i p_j$ intersects the interior of $\mathbb{R}^2 \setminus T$. For $\varepsilon > 0$, let $p_i(\varepsilon)$ be an arbitrary point of T in the ε-neighborhood of p_i. For $\varepsilon > 0$ small enough, the $n+1$ points $p_i(\varepsilon)$ are visually independent — a contradiction. Therefore T is polygonal. Thus, each $V_i \cap X$ is polygonal as well. Since each $v_i \cap X$ is also polygonal, Lemma 5.2 is proved. ∎

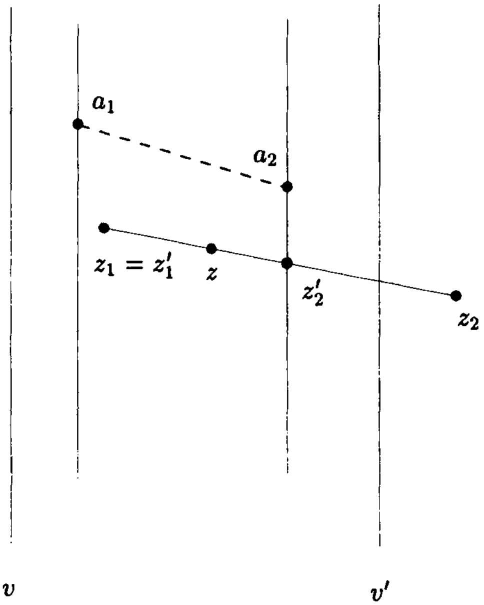

Figure 11. Illustration to the proof of the convexity of U.

6. Examples

In this section we describe examples of sets in the plane proving Theorem 1.3 (Examples 2–5) and an example showing a set $X \subseteq \mathbb{R}^2$ with $\gamma(X)$ exponentially large in $\chi(X)$ (Example 1). The sets in these examples have the following properties of interest:

- Example 1: sets with $\gamma(X)$ at least exponential in $\chi(X)$.
- Examples 2 and 3: closed sets with $\gamma(X) \geq c \cdot \chi(X)^2$.
- Example 4: sets with m one-point holes, such that $\gamma(X) \geq c \cdot m \cdot \omega(X)$.
- Example 5: sets with no one-point holes and with $\gamma(X) \geq c \cdot \omega(X)^3$.

Examples 4 and 5 give together the lower bound $\gamma(X) \geq c(\omega(X)^3 + \lambda \cdot \omega(X))$ in Theorem 1.1 (Example 4 for $\lambda \geq \omega(X)^2$, Example 5 for $\lambda \leq \omega(X)^2$). Examples 2 and 3 give the same, quadratic, lower bound on $\gamma(X)$ for closed sets. We present both of them since they might provide some insight to the problem of closing the gap between our lower and upper bounds on $\gamma(X)$ in terms of $\omega(X)$ (or in terms of $\chi(X)$) for closed sets X. In the proof of Theorem 1.1(i), we first bound the number of local extremes of $\mathbb{R}^2 \setminus X$ by $O(\omega(X)^2)$, and then we show that $\gamma(X)$ is at most $O(\omega(X))$ times larger. The examples show that none of the steps in this proof scheme can be improved independently of the other: The set in Example 3 has an asymptotically maximal number $\Theta(\omega(X)^2)$ of local extremes of $\mathbb{R}^2 \setminus X$, while the set in Example 2 has only $O(\omega(X))$ of them but the ratio of $\gamma(X)$ to the number of local extremes of $\mathbb{R}^2 \setminus X$ is linear in $\omega(X)$.

The variables m and m' appearing in the examples may be any integers bigger than 2. Since the constructions in Examples 2–5 are partially similar, we now introduce common notions listed in the left margin below.

Notation for Examples 2–5:

P, C, Q: Let $P = \{p_1, \dots, p_m\}$ be a set of vertices of a regular m-gon, and let C be the circle containing points of P in the counterclockwise order $p_1, \dots, p_m$. Let o be any point inside the convex hull of P that lies on no segment p_ip_j. (E.g., for odd m, it is most transparent to put o in the center of conv P.) For $i = 1, 2, \dots, m$, let q_i be an inner point of the segment p_io that lies close to p_i, so that the interior of any triangle $p_jp_kp_\ell$ contains at least one point q_i. Set $Q = \{q_1, \dots, q_m\}$ and assume that no two points of Q lie on a common vertical line.

C_i, $\mathcal{R}_i$, $\mathcal{R}$: For each point p_i, consider a circle C_i of large radius touching C at the point p_i, and take a set $\mathcal{R}_i$ of m disjoint skinny rectangles in a small neighborhood of p_i, so that one of the two longer sides of each of these rectangles is a chord of C_i, and the other one lies inside C_i (see Fig. 12).

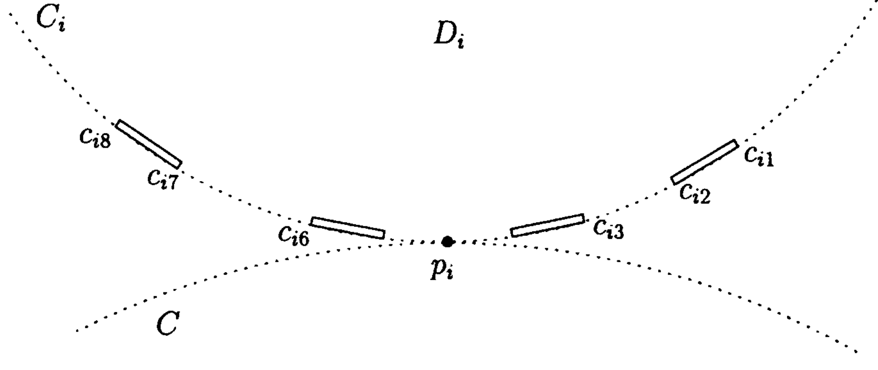

Figure 12. Rectangles in $\mathcal{R}_i$.

The rectangles in $\mathcal{R}_i$ are placed so close to p_i that any line containing a point of Q intersects at most one of the m sets $\mathrm{conv}(\bigcup_{R\in\mathcal{R}_i} R)$, $i = 1, \ldots, m$, and the diameter of each rectangle is taken much smaller than its distance to the closest other rectangle. Put $\mathcal{R} := \bigcup_{i=1}^{m} \mathcal{R}_i$.

D_i, R', R'', $\mathcal{D}$: For $i = 1, 2, \ldots, m$, let D_i be the interior of the disk bounded by C_i. For each rectangle $R \in \mathcal{R}_i$, the lines containing the longer sides of R partition D_i into three regions. Denote the closure of the region containing R by R', and denote the interior of the (small) region containing neither R nor the center of D_i by R''. The m^2 regions R' ($R \in \mathcal{R}$) are disjoint because of the small size of the rectangles $R \in \mathcal{R}$. Put $\mathcal{D} = \bigcup_{i=1}^{m} D_i$.

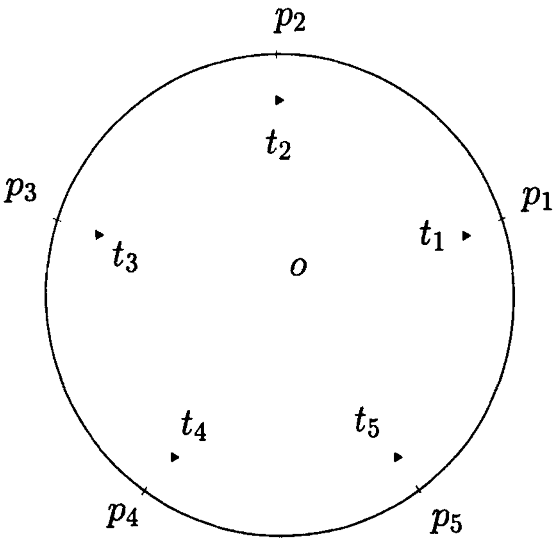

Figure 13. The triangles t_i.

t_i, Q': For each $i = 1, \ldots, m$, take an equilateral triangle t_i bounded from the left by a vertical segment whose mid-point is q_i (see Fig. 13). Define Q' as the union of the interiors of the m triangles t_i, $1 \le i \le m$. The triangles t_i are taken so small that any line passing through a point of Q' intersects at most one of the sets $\mathrm{conv}(\bigcup_{R\in\mathcal{R}_i} R)$, $i = 1, \ldots, m$, and that any vertical line crosses at most one triangle t_i.

$Z(m, m')$, p_{ij}, e_i: For every $i = 1, \ldots, m$, take $m' + 1$ points p_{ij}, $j = 0, \ldots, m'$, in a small neighborhood of p_i so that they appear in the clockwise order $p_{i0}, p_{i1}, \ldots, p_{im'}$ on the circle C_i. Define $Z(m, m')$ as the closed $(m'+1)$-convex $((m'+1)m)$-gon

$$p_{10} \cdots p_{1m'} p_{20} \cdots p_{2m'} \cdots\cdots p_{m0} \cdots p_{mm'}$$

(see Fig. 14). For $i = 1, \ldots, m$, e_i is defined as the point of intersection of lines $p_{i0}p_{i1}$ and $p_{in}p_{(i+1)0}$.

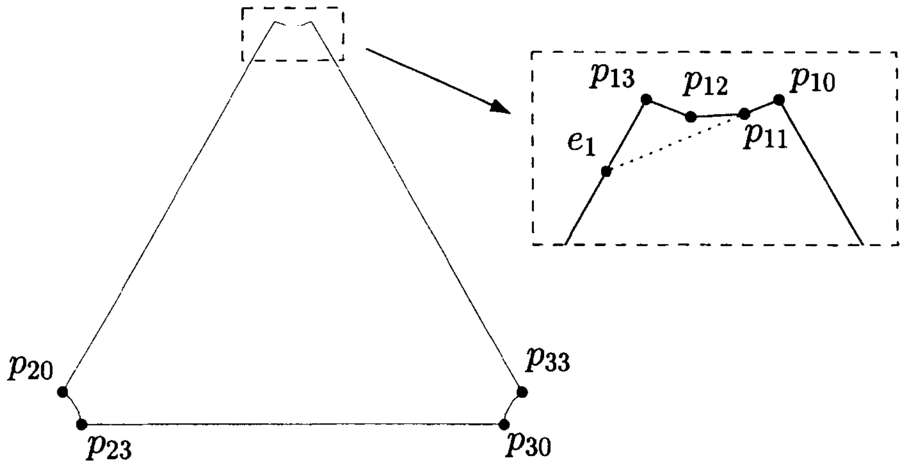

Figure 14. The set $Z(3, 3)$.

Before the examples we include an observation.

OBSERVATION 6.1: *Let C be any convex set in the plane. Then the following holds:*

(i) *Let S be a set of k vertical segments in the plane. Then the set $C \setminus S$ is a union of $k + 1$ pairwise disjoint convex sets. In particular,*

$$\gamma(C \setminus S) \leq k + 1.$$

(ii) *Let Q' be the union of the interiors of the m (open) triangles defined in the notation above. Then the set $C \setminus Q'$ is a union of $2m + 1$ pairwise disjoint convex sets. In particular,*

$$\gamma(C \setminus Q') \leq 2m + 1.$$

Proof: The first part certainly holds for $k \leq 1$, and it can be easily extended to an arbitrary k by induction. For the proof of the second part we refer to related Fig. 16. ∎

Example 1: We construct a set $X_1 \subset \mathbb{R}^2$, with

$$\omega(X_1) = 3, \qquad \chi(X_1) \leq m + 1,$$

$$\gamma(X_1) \geq \frac{\binom{m}{\lfloor m/2 \rfloor}}{2} + 1 = \Theta(2^m m^{-1/2}).$$

Set $h = \binom{m}{\lfloor m/2 \rfloor}$. Let D be a disk and let A be a set of h points $a_1, \ldots, a_h$ inside D, which are placed near the boundary of D and form a regular h-gon $a_1 a_2 \cdots a_h$ (see Fig. 1). We claim that the set $X_1 = D \setminus A$ has the required properties.

Obviously $\omega(X_1) = 3$, and $\gamma(X_1) \geq \frac{1}{2}h + 1$ follows from Observation 1.11. It remains to color the set X_1 by $m + 1$ colors. For $i = 1, \ldots, h$, let G_i be the small sector of D bounded by a chord g_i of D with center in a_i. For each $i = 1, \ldots, h$, we delete from G_i all points of g_i that lie on an arbitrary fixed side of a_i. Then $X_1 \setminus \bigcup_{i=1}^h G_i$ may be colored by one color.

We now take a set C of m colors, and denote the $\lfloor \frac{m}{2} \rfloor$-element subsets of C by $C_1, \ldots, C_h$ (in an arbitrary order). Let p be a point in $G_i, i \in \{1, \ldots, h\}$. If the line pa_i intersects no region $G_j, j \neq i$, then we color p by an arbitrary color of C_i. If the line pa_i intersects a region $G_j, j \neq i$, then we color p by an arbitrary color of $C_i \setminus C_j$. This gives a proper coloring of $\bigcup_{i=1}^h G_i$ by m colors. Thus, $\chi(X_1) \leq m + 1$. This concludes Example 1. ∎

Example 2: We construct a closed set $X_2 \subset \mathbb{R}^2$, with

$$\chi(X_2) \leq 3m, \qquad \gamma(X_2) \geq \frac{1}{2}m^2.$$

Consider the sets $Z(m, m')$ and Q' introduced in the notation before Observation 6.1. We claim that the set $X_2 = Z(m, m) \setminus Q'$ (see Fig. 15) has the required properties.

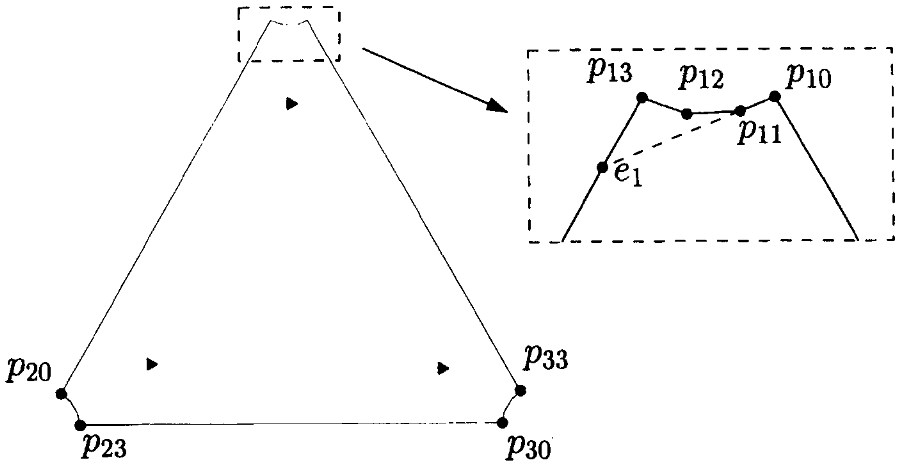

Figure 15. The set X_2 in Example 2.

The intersection of X_2 with the convex $(2m)$-gon $p_{10}e_1p_{20}e_2\cdots p_{m0}e_m$ can be colored by $2m+1$ colors according to Observation 6.1(ii), see Fig. 16, and the rest of X_2 may easily be colored by $m-1$ colors. So, $\chi(X_2) \le 3m$.

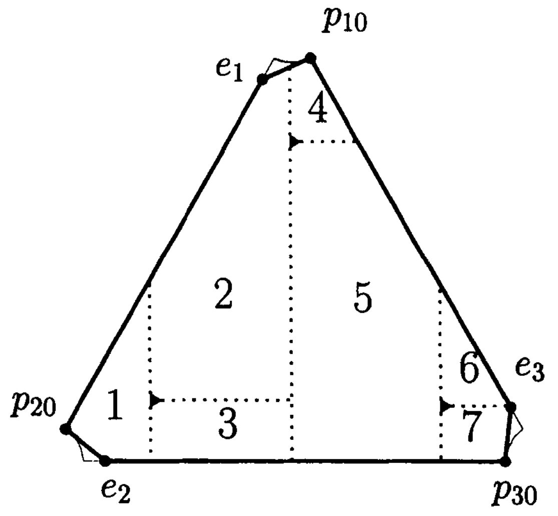

Figure 16. Coloring of a part of X_2 by $2m+1$ colors.

Let M be the (m^2)-element set of the mid-points of the segments $p_{i(j-1)}p_{ij}$ $(i,j=1,2,\ldots,m)$. No three points of M can be covered by one convex subset of X_2. Therefore, $\gamma(X_2) \ge \frac{1}{2}m^2$. This concludes Example 2. ∎

Example 3: *We construct a closed set $X_3 \subset \mathbb{R}^2$, with*

$$\chi(X_3) \le 6m + 5, \qquad \gamma(X_3) \ge \frac{1}{2}m^2.$$

Consider the set $\mathcal{R}$ of m^2 rectangles and the set Q' introduced in the notation before Observation 6.1, and replace each rectangle $R \in \mathcal{R}_i$, $i = 1, \dots, m$, by a triangle whose two vertices a, b are the two points in which R meets C_i and the third vertex is the mid-point of the side of R opposite to ab (see Fig. 17).

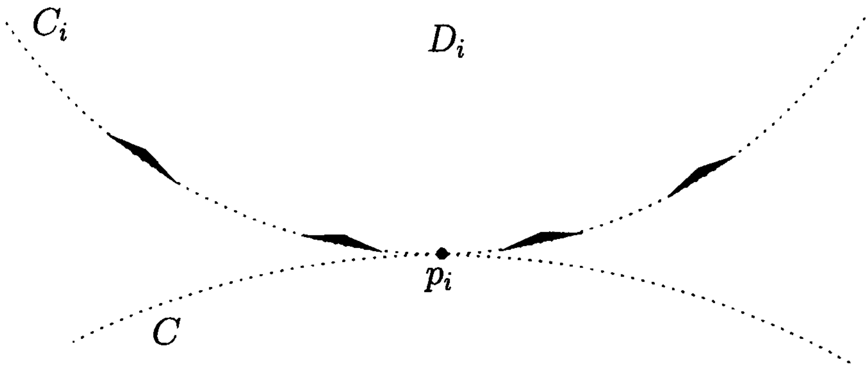

Figure 17. Triangles of T lying inside D_i.

Define T as the union of the interiors of these m^2 triangles. We claim that the set $X_3 = \mathbb{R}^2 \setminus (T \cup Q')$ has the required properties.

First we find a proper coloring of (the invisibility graph of) X_3 by $6m + 5$ colors.

For any $i = 1, \dots, m$, the intersection of X_3 with the set $(D_i \setminus (C_i \cup \bigcup_{R\in\mathcal{R}_i} R''))$ may be colored by two colors. Thus, $X_3 \cap (\mathcal{D} \setminus \bigcup_{R\in\mathcal{R}} R'')$ may be colored by $2m$ colors. The intersection of X_3 with $W = (\mathbb{R}^2 \setminus \mathcal{D}) \cup \bigcup_{R\in\mathcal{R}} R''$ can be colored by $4m + 5$ colors as follows.

The set $W \cap X_3$ consists of two components. The unbounded one can be colored by three colors. Denote the bounded one by B. Further, denote the point of intersection $C_i \cap C_{i+1} \cap B$ by c_{i+1} (if $i = m$, we let $m + 1 := 1$). The circle C_i contains $2m$ vertices of the rectangles $R \in \mathcal{R}_i$. We denote them by $c_{i1}, c_{i2}, \dots, c_{i(2m)}$ (in the clockwise order in which they appear on C_i; c_{i1} is the nearest one to c_i), see Fig. 12. Let d_i be the point of intersection of the segment $c_{i(2m)}c_{(i+1)1}$ with the line $c_{i1}c_{i2}$. With this notation, the set B may be covered by the convex $(2m)$-gon $G := c_{11}d_1c_{21}d_2 \dots c_{m1}d_m$, the m triangles $T_i := c_{i(2m)}c_{i+1}c_{(i+1)1}$ $(1 \le i \le m)$, and the m non-convex $(2m)$-gons $G_i := c_{i2}c_{i3} \cdots c_{i(2m)}d_i$ $(1 \le i \le m)$.

The set $B \cap G = G \setminus Q'$ can be colored by $2m + 1$ colors according to Observation 6.1(ii) (see Fig. 16 for a coloring of a similar set). The set $B \cap \bigcup_{i=1}^{m} T_i$ can be colored by three colors since each of the three sets $B \cap T_1$, $B \cap \bigcup_{\substack{2 \le i \le m \\ i \text{ even}}} T_i$, $B \cap \bigcup_{\substack{3 \le i \le m \\ i \text{ odd}}} T_i$ can be colored by one color. It is also not difficult to see that the union of the m sets $B \cap G_i$ $(1 \le i \le m)$ can be colored by $2m - 2$ colors. So the set $W \cap X_3$ may be colored by $3 + (2m + 1) + 3 + (2m - 2) = 4m + 5$ colors, and X_3 may be colored by $2m + (4m + 5) = 6m + 5$ colors.

We now show that $\gamma(X_3) \ge \frac{1}{2}m^2$. For each triangle t that is a component of T, let $g(t) \in X_3$ be the center of the longest side of t. The set X_3 is constructed so that no three of the m^2 points $g(t)$ may be covered by one convex subset of X_3. Therefore, $\gamma(X_3) \ge \frac{1}{2}m^2$. This is the end of Example 3. ∎

Example 4: We construct a set $X_4 \subset \mathbb{R}^2$ with m one-point holes, with

$$\omega(X_4) = m', \qquad \chi(X_4) \le m' + O(\log m),$$

$$\gamma(X_4) \ge \frac{m'm}{2}.$$

We put $X_4 = Z(m, m') \setminus Q$. Obviously $\omega(X_4) = m'$, since any visually independent set of points in X_4 lying in a small neighborhood of a point of P has size at most m', and any other visually independent set of points in X_4 has size at most 3.

The intersection of X_4 with the convex $2m$-gon $p_{10}e_1p_{20}e_2 \cdots p_{m0}e_m$ can be colored by $O(\log m)$ colors (see Example 1 for a coloring of a similar set), and the remaining part of X_4 can easily be colored by $m' - 1$ colors. Thus, $\chi(X_4) \le m' + O(\log m)$.

Let M be the $(m'm)$-element set of the mid-points of the segments $p_{i(j-1)}p_{ij}$ $(i = 1, 2, \ldots, m,\ j = 1, 2, \ldots, m')$. No three points of M can be covered by one convex subset of X_4. Thus, $\gamma(X_4) \ge m'm/2$. This concludes Example 4. ∎

Example 5: We construct a set $X_5 \subset \mathbb{R}^2$ with no one-point holes and with

$$\chi(X_5) \le 10m + 9, \qquad \gamma(X_5) \ge \frac{1}{2}m^3.$$

For each rectangle $R \in \mathcal{R}$ (see the notation before Observation 6.1), place m very short segments on each of the two longer sides of R, s, s', so that, after orthogonal projection on s, they become $2m$ disjoint segments of equal length, which are equidistantly distributed along the whole s and originate alternately from s and from s' (see Fig. 18).

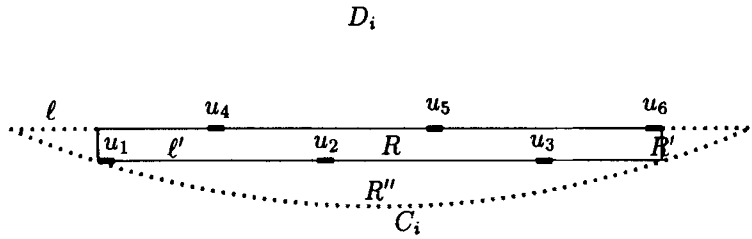

Figure 18. The segments u_i.

Moreover, assume that if we view R from the center of C, then the leftmost of the $2m$ segments lies on the nearest side of R. Define S as the union of all these $2m^3$ segments and of m short vertical segments with respective mid-points q_i. We claim that the set $X_5 = \mathbb{R}^2 \setminus S$ has the required properties.

First we show that $\gamma(X_5) \geq \frac{1}{2}m^3$. Denote the longer sides of R by $s(R)$ and $s'(R)$, so that $s'(R)$ is a chord of C_i. For each segment u of S lying on a segment $s(R), R \in \mathcal{R}$, let $g(u)$ be the center of the orthogonal projection of u to $s'(R)$ (see Fig. 19).

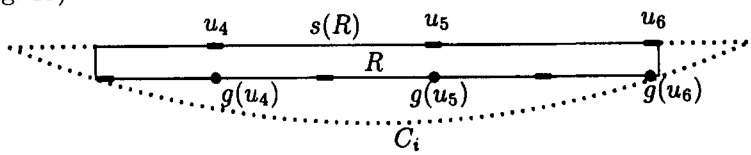

Figure 19. The points $g(u)$.

The set X_5 is constructed so that no three of the m^3 points $g(u)$ may be covered by one convex subset of X_5. Therefore, $\gamma(X_5) \geq \frac{1}{2}m^3$.

We now find a proper coloring of (the invisibility graph of) X_5 by $10m + 9$ colors, which is, however, quite complicated. We separately color the intersection of X_5 with the following three regions covering the plane:

- $S_1 := \mathcal{D}$ minus the $2m^2$ regions R', R'' $(R \in \mathcal{R})$.
- $S_2 :=$ the union of $(\mathbb{R}^2 \setminus \mathcal{D})$ with the m^2 regions R'' $(R \in \mathcal{R})$.
- $S_3 :=$ the union of the m^2 regions R' $(R \in \mathcal{R})$.

The set $S_1 \cap X_5 = S_1$ may be colored by m colors since each of the m sets $S_1 \cap D_i$ covering S_1 may be colored by one color.

The region S_2 coincides with the region W in Example 3. Therefore, $S_2 \cap X_5$ can be colored by $3m+5$ colors similarly as the set $W \cap X_3$ was colored by $4m+5$ colors in Example 3 (we save m colors since the set $X_5 \cap G$ can be colored by

$m+1$ colors according to Observation 6.1(i), while the set $X_3 \cap G = G \setminus Q'$ was colored by $2m+1$ colors according to Observation 6.1(ii).)

We now color the set $S_3 \cap X_5$ by $6m+4$ colors, which completes the coloring of $X_5 = (S_1 \cap X_5) \cup (S_2 \cap X_5) \cup (S_3 \cap X_5)$ by $m + (3m+5) + (6m+4) = 10m+9$ colors. Let $R \in \mathcal{R}_i$, where $i \in \{1, \dots, m\}$. Denote the longer sides of R by $s(R)$ and $s'(R)$ as above. Let $\ell(R)$ and $\ell'(R)$ be the intersections of R' with the lines containing $s(R)$ and $s'(R)$, respectively. Further, denote the segments of S lying in R by $u_1, u_2, \dots, u_{2m}$ so that $u_1, u_2, \dots, u_m$ lie on $s'(R)$, see Fig. 18. For $z \in \mathbb{R}^2$, let $|zu_i|$ be the distance of z from the segment u_i. We cover $R' \setminus (\ell'(R) \cup \ell(R))$ by $2m$ Voronoi regions

$$r_j(R) := \{z \in R' \setminus (\ell'(R) \cup \ell(R)); |zu_j| \le |zu_k| \text{ for each } k = 1, \dots, 2m\},$$

$j = 1, \dots, 2m$; see Fig. 20.

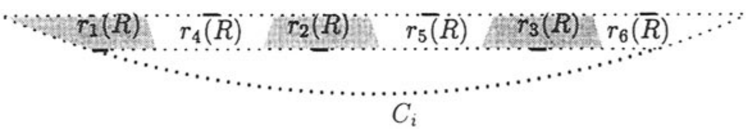

Figure 20. The regions $r_j(R)$.

For any $1 \le i \le m$, the m^2 regions $r_j(R)$ $(1 \le j \le m,\ R \in \mathcal{R}_i)$ may be colored by one color. So all the m^3 regions $r_j(R)$ $(1 \le j \le m,\ R \in \mathcal{R})$ may be colored by m colors. It is not difficult to see that the remaining m^3 regions $r_j(R)$ $(m+1 \le j \le 2m,\ R \in \mathcal{R})$ also may be colored by m colors, so that, for every $R \in \mathcal{R}$, the regions $r_{m+1}(R), \dots, r_{2m}(R)$ are colored by one color, and any two regions $r_j(R)$, $r_{\tilde{j}}(\tilde{R})$, $R, \tilde{R} \in \mathcal{R}_i$, $R \ne \tilde{R}$, are colored by different colors. Thus, the set $\bigcup_{R \in \mathcal{R}} R' \setminus (\ell(R) \cup \ell'(R))$ can be colored by $2m$ colors.

It remains to color the set $\bigcup_{R \in \mathcal{R}} (\ell(R) \cap X_5) \cup \bigcup_{R \in \mathcal{R}} (\ell'(R) \cap X_5)$ by $4m+4$ colors. We first color the set $\bigcup_{R \in \mathcal{R}} (\ell(R) \cap X_5)$ by $2m+2$ colors. Each set $(\ell(R) \cap X_5), R \in \mathcal{R}$, consists of $m+1$ segments which we denote by $\ell_1(R), \ell_2(R), \dots, \ell_{m+1}(R)$. Below, to each segment $\ell_j(R)$ we assign a list of two colors, which are allowed to be used on $\ell_j(R)$. We say that a point *sees* a segment v within X, if it sees every point of v in X. Let p be an arbitrary point on any segment $\ell_j(R), R \in \mathcal{R}_i, 1 \le i \le m, 1 \le j \le m+1$. Since the segments of S are very short, the point p doesn't see at most one of the $(m+1)^3 - (m+1)^2$ segments $\ell_{\tilde{j}}(\tilde{R})$ $(\tilde{R} \in \mathcal{R}_{\tilde{i}}, \tilde{i} \ne i, 1 \le \tilde{j} \le m+1)$ within $\mathbb{R}^2 \setminus (R \cap S)$. If the point p sees all the segments $\ell_{\tilde{j}}(\tilde{R})$, then we color it by one (arbitrarily chosen) of the two colors assigned to $\ell_j(R)$. If p doesn't see a segment $\ell_{\tilde{j}}(\tilde{R})$ $(\tilde{R} \in \mathcal{R}_{\tilde{i}}, \tilde{i} \ne i, 1 \le \tilde{j} \le m+1)$

within $\mathbb{R}^2 \setminus (R \cap S)$, then we color it by a color that is assigned to $\ell_j(R)$ and isn't assigned to the segment $\ell_{\tilde{j}}(\tilde{R})$. If this is possible for every p then it certainly gives a proper coloring. To make it possible, we only need to assign a pair of colors to each segment $\ell_j(R), R \in \mathcal{R}$, so that the pairs of colors assigned to any two segments $\ell_j(R), \ell_{\tilde{j}}(\tilde{R})$ $(R \in \mathcal{R}_i, \tilde{R} \in \mathcal{R}_{\tilde{i}}, \tilde{i} \neq i, j, \tilde{j} \in \{1, 2, \ldots, m+1\})$ are always different. This can be done by $2(m+1)$ colors as follows. Let C be a set of $2(m+1)$ colors. Take m perfect matchings $M_1, \ldots, M_m$ on the vertex set C that have mutually disjoint edge sets. (It is even not difficult to find $2m+1$ such matchings whose union is the complete graph on the vertex set C). Then, for every $i = 1, \ldots, m$ and for every $R \in \mathcal{R}_i$, we assign each of the $m+1$ pairs of colors joined by an edge in the matching M_i to exactly one of the $m+1$ segments $\ell_j(R)$, $j = 1, \ldots, m+1$. This gives the required assignment of colors. This completes the coloring of $\bigcup_{R \in \mathcal{R}} (\ell(R) \cap X_5)$ by $2m+2$ colors.

We now analogously color the set $\bigcup_{R \in \mathcal{R}} (\ell'(R) \cap X_5)$ by $2m+2$ colors. We first note that it is enough to color the set $L = \bigcup_{R \in \mathcal{R}_1} (\ell'(R) \cap X_5)$ by $2m+2$ colors, since each of the $m-1$ sets $\bigcup_{R \in \mathcal{R}_i} (\ell'(R) \cap X_5)$ $(2 \le i \le m)$ may be colored analogously as L using the same set of $2m+2$ colors. Let us fix $R_i \in \mathcal{R}_i$ for every $i = 1, \ldots, m$. The set L may be colored by $2m+2$ colors analogously as the set $\bigcup_{i=1}^{m} (\ell(R_i) \cap X_5)$ was colored above. This completes the proof that $\chi(X_5) \le 10m + 9$, and Example 5 ends here. ∎

ACKNOWLEDGEMENT: We would like to thank Prof. Micha Perles for interesting comments and for explaining some of his related results to us. A great help in preparing pictures for this paper was *Ipe (Integrated Picture Environment)*, a drawing editor written by Otfried Schwarzkopf. We also thank the referee for numerous valuable comments to the presentation.

References

[BK76] M. Breen and D. C. Kay, *General decomposition theorems for m-convex sets in the plane,* Israel Journal of Mathematics **24** (1976), 217–233.

[Dil50] R. P. Dilworth, *A decomposition theorem for partially ordered sets,* Annals of Mathematics **51** (1950), 161–166.

[Egg74] H. G. Eggleston, *A condition for a compact plane set to be a union of finitely many convex sets,* Proceedings of the Cambridge Philosophical Society **76** (1974), 61–66.

[ES35] P. Erdős and G. Szekeres, *A combinatorial problem in geometry,* Compositio Mathematica **2** (1935), 463–470.

[Gya87] A. Gyarfás, *Problems from the world surrounding perfect graphs*, Zastosowonia Matematyki **19** (1987), 413–441.

[KG71] D. Kay and M. Guay, *Convexity and certain property P_n*, Israel Journal of Mathematics **8** (1971), 39–52.

[KPS90] M. Kojman, M. Perles and S. Shelah, *Sets in a Euclidean space which are not a countable union of convex subsets*, Israel Journal of Mathematics **70** (1990), 313–342.

[LHK72] J. Lawrence, W. Hare, Jr., and J. Kenelly, *Finite unions of convex sets*, Proceedings of the American Mathematical Society **34** (1972), 225–228.

[Mul94] K. Mulmuley, *Computational Geometry: An Introduction Through Randomized Algorithms*, Prentice-Hall, Englewood Cliffs, NJ, 1994.

[O'R87] J. O'Rourke, *Art Gallery Theorems and Algorithms*, Oxford University Press, New York, 1987.

[PS90] M. Perles and S. Shelah, *A closed $(n+1)$-convex set in R^2 is a union of n^6 convex sets*, Israel Journal of Mathematics **70** (1990), 305–312.

[Val57] F. Valentine, *A three-point convexity property*, Pacific Journal of Mathematics **34** (1957), 1227–1235.