

THE ERDŐS–SZEKERES PROBLEM FOR NON-CROSSING CONVEX SETS

MICHAEL GENE DOBBINS, ANDREAS HOLMSEN AND
ALFREDO HUBARD

Abstract. We show an equivalence between a conjecture of Bisztriczky and Fejes Tóth about families of planar convex bodies and a conjecture of Goodman and Pollack about point sets in topological affine planes. As a corollary of this equivalence we improve the upper bound of Pach and Tóth on the Erdős–Szekeres theorem for disjoint convex bodies, as well as the recent upper bound obtained by Fox, Pach, Sudakov and Suk on the Erdős–Szekeres theorem for non-crossing convex bodies. Our methods also imply improvements on the positive fraction Erdős–Szekeres theorem for disjoint (and non-crossing) convex bodies, as well as a generalization of the partitioned Erdős–Szekeres theorem of Pór and Valtr to families of non-crossing convex bodies.

§1. Introduction.

1.1. *The happy ending theorem.* In 1935, Erdős and Szekeres proved the following foundational result in combinatorial geometry and Ramsey theory[†].

THEOREM (Erdős–Szekeres [9]). *For every integer $n \geq 3$ there exists a minimal positive integer $f(n)$ such that any set of $f(n)$ points in the Euclidean plane, in which every triple is convexly independent, contains a convexly independent subset of size n .*

Here *convexly independent* means that no point is contained in the convex hull of the others. Determining the precise growth of the function $f(n)$ is one of the longest-standing open problems of combinatorial geometry, and has generated a considerable amount of research. For history and details, see [1, 26] and the references therein. Two proofs are given in [9], one of which shows that $f(n) \leq \binom{2n-4}{n-2} + 1$, and in [10] Erdős and Szekeres give a construction showing that $f(n) \geq 2^{n-2} + 1$.

CONJECTURE (Erdős–Szekeres). *We have $f(n) = 2^{n-2} + 1$.*

This conjecture has been verified for $n \leq 6$ [9, 33], while for $n \geq 7$ the best known upper bound is $f(n) \leq \binom{2n-5}{n-2} + 1 \sim 4^n / \sqrt{n}$, which is due to

Received 30 April 2013.

MSC (2010): 52C10, 52C30 (primary).

[†] Paul Erdős colloquially referred to this as the “happy ending theorem” as it led to the marriage of George Szekeres and Esther Klein.




Figure 1: Four points in a topological plane. Each pair of points determines a unique pseudoline which contains their connecting pseudosegment (thickened). Their convex hull is the shaded region and shows that the points are not convexly independent.

Tóth and Valtr [35]. Asymptotically this is the same as the bound given by Erdős and Szekeres in their seminal paper.

1.2. *Generalized configurations.* It was observed by Goodman and Pollack [17] that the Erdős–Szekeres theorem extends to so-called *generalized configurations*, i.e. point sets in a topological affine plane [18, 20, 22]. One may consider this as a finite configuration of points in the plane where each pair of points is contained in a unique *pseudoline* in such a way that the resulting set of pseudolines forms a *pseudoline arrangement* [15]. This underlying pseudoline arrangement induces a convexity structure on the point configuration in a natural way: each pair of points of the configuration bounds a unique *pseudosegment* contained in the associated pseudoline. The complement of the pseudosegments determined by all the pairs of the configuration is a collection of open connected regions, one of which is unbounded. The *convex hull* of the configuration is the complement of the unbounded region; see Figure 1. It turns out that many basic theorems of convexity hold in this more general setting; for instance, a set of points is convexly independent if and only if every four of its points are convexly independent [7] (which is commonly called Carathéodory’s theorem).

Generalized configurations have a purely combinatorial characterization and there are several equivalent axiom systems which define them. Other names for generalized configurations found in the literature are *uniform rank 3 acyclic oriented matroids* [6] or *CC-systems* [25].

THEOREM (Goodman–Pollack [17]). *For every integer $n \geq 3$ there exists a minimal positive integer $g(n)$ such that any generalized configuration of size $g(n)$, in which every triple is convexly independent, contains a convexly independent subset of size n .*

It should be noted that this is a proper generalization of the Erdős–Szekeres theorem since most generalized configurations are not realizable by points and straight lines [11, 19]. By containment it follows that $f(n) \leq g(n)$.

CONJECTURE (Goodman–Pollack). *We have $f(n) = g(n)$.*

Since the proof of Tóth and Valtr does not use any metric properties, it can be extended to generalized configurations. This will be done in §3. We therefore have $g(n) \leq \binom{2n-5}{n-2} + 1$ for $n \geq 7$. Also, the computer aided proof of Szekeres and Peters [33] confirming that $f(6) = 17$ actually encodes generalized configurations, and it follows that $g(n) = f(n) = 2^{n-2} + 1$ for all $n \leq 6$.

1.3. *Mutually disjoint convex bodies.* In a different direction, initiated by Bisztriczky and Fejes Tóth, the Erdős–Szekeres theorem was generalized to families of compact convex sets in the plane (which we call *bodies* for brevity). A family of bodies is *convexly independent* if no member is contained in the convex hull of the others.

THEOREM (Bisztriczky–Fejes Tóth [3]). *For any integer $n \geq 3$ there exists a minimal positive integer $h_0(n)$ such that any family of $h_0(n)$ pairwise disjoint bodies in the Euclidean plane, in which every triple is convexly independent, contains a convexly independent subfamily of size n .*

This reduces to the Erdős–Szekeres theorem when the bodies are points, but was somewhat more complicated to establish in general. The added complexity is reflected in the original upper bound $h_0(n) \leq t_n(t_{n-1}(\dots t_1(cn)\dots))$, where t_n is the n th tower function. The upper bound was later reduced to $\binom{2n-4}{n-2}^2 + 1$ by Pach and Tóth in [28] (see also reference [24]). By containment we have $f(n) \leq h_0(n)$.

CONJECTURE (Bisztriczky–Fejes Tóth). *We have $f(n) = h_0(n)$.*

This conjecture has been verified for $n \leq 5$ [3, 4].

1.4. *Non-crossing convex bodies.* The disjointness hypothesis in the Bisztriczky–Fejes Tóth theorem was relaxed by Pach and Tóth, who showed that an Erdős–Szekeres theorem also holds for families of *non-crossing* bodies. A pair of bodies A and B are non-crossing if the set $A \setminus B$ is simply connected. Generically this means that A and B have precisely two *common supporting tangents* (we make this precise in §2.1).

THEOREM (Pach–Tóth [29]). *For any integer $n \geq 3$ there exists a minimal positive integer $h_1(n)$ such that any family of $h_1(n)$ non-crossing bodies in the Euclidean plane, in which every triple is convexly independent, contains a convexly independent subfamily of size n .*

The original upper bound on $h_1(n)$ was improved to a doubly exponential function in [24]. Recently Fox, Pach, Sudakov, and Suk [13] obtained the upper bound $h_1(n) \leq 2^{cn^2 \log n}$, for some absolute constant $c > 0$. See also [5, 34] for related work.

The known bounds are summarized as follows:

$$\begin{aligned}
2^{n-2} + 1 &\leq f(n) \leq g(n) \leq \binom{2n-5}{n-1} + 1 && (\text{for } n \geq 7), \\
f(n) &\leq h_0(n) \leq \binom{2n-4}{n-2}^2 + 1, \\
h_0(n) &\leq h_1(n) \leq 2^{cn^2 \log n}, \\
2^{n-2} + 1 &= f(n) = g(n) && (\text{for } n \leq 6), \\
2^{n-2} + 1 &= h_0(n) && (\text{for } n \leq 5).
\end{aligned}$$

1.5. *Our results.* In this paper we make considerable improvements on $h_0(n)$ and $h_1(n)$ by establishing the following theorem.

THEOREM 1.1. *The Erdős–Szekeres problems for generalized configurations and for families of non-crossing bodies are equivalent. In other words, $g(n) = h_1(n)$.*

As a consequence we obtain the following bounds:

$$\begin{aligned}
2^{n-2} + 1 &\leq f(n) \leq h_0(n) \leq h_1(n) = g(n) \leq \binom{2n-5}{n-2} + 1 && (\text{for } n \geq 7), \\
2^{n-2} + 1 &= f(n) = h_0(n) = h_1(n) = g(n) && (\text{for } n \leq 6).
\end{aligned}$$

Here is the idea of the proof. For the lower bound we use the fact that a generalized configuration has a dual representation as a marked pseudoline arrangement, i.e. a wiring diagram [12, 14]. Using this representation we show that every generalized configuration can be represented by a family of non-crossing bodies in the Euclidean plane. This shows that $g(n) \leq h_1(n)$.

To establish the reverse inequality we start with a family of bodies and consider its dual system of support curves drawn on the cylinder $\mathbb{S}^1 \times \mathbb{R}^1$. This system of curves induces a cell complex which encodes the convexity properties of the family. We show how to modify this complex by elementary operations, similar to those of Habert and Pocchiola [23] and Ringel [32], while maintaining control of the convexity properties of the family. The process ends with a complex induced by a family representing a generalized configuration. The details are given in §2.

Our proof actually provides a general method for reducing families of non-crossing bodies to generalized configurations. This method can be applied to the multitude of Erdős–Szekeres-type results previously proven separately for point sets, then for families of bodies. Therefore we can generalize the results of the papers [2, 27, 30, 31] to families of non-crossing convex bodies, as well as improve on the bounds for the case of families of pairwise disjoint convex bodies. This will be discussed in §4.

Remark 1.2. It should be noted that the condition that the bodies are convex is not strictly necessary. The theorems also hold for any family of compact sets $\{A_1, \dots, A_n\}$ provided we impose conditions on the family $\{\text{conv}(A_1), \dots, \text{conv}(A_n)\}$. However, no real generality is gained by this formulation, so we restrict ourselves to families of convex bodies to make our statements simpler.

§2. Proof of Theorem 1.1.

2.1. *Duality.* We call a compact convex subset of \mathbb{R}^2 a *body*. For a body B , its support function $h_B: \mathbb{S}^1 \rightarrow \mathbb{R}^1$ on the unit circle is defined as

$$h_B(\theta) := \max_{x \in B} \langle \theta, x \rangle.$$

The *dual* of a body B is the graph of its support function drawn on the cylinder $\mathbb{S}^1 \times \mathbb{R}^1$, i.e.

$$B^* := \{(\theta, h_B(\theta)) : \theta \in \mathbb{S}^1\}.$$

We implicitly assume that the unit circle is oriented, and the dual curves should therefore be thought of as *directed* curves. It is important to make a consistent choice of orientation throughout, so we fix the positive orientation to be the *counter-clockwise direction*.

Notice that for each $\theta \in \mathbb{S}^1$ the value $h_B(\theta)$ measures the oriented distance from the origin to the directed *supporting tangent* of B with direction $\theta + \pi/2$, which has the body B on its left side. For instance (identifying \mathbb{S}^1 with the interval $[0, 2\pi)$) the dual of a point distinct from the origin is a sine curve, while the dual of a disk centered at the origin is the graph of a constant function. Note that if $h_B(\theta) \geq 0$ for all θ , then the origin is contained in the body B ; see Figure 2(a).

We use the term *system* when referring to a finite collection of curves on $\mathbb{S}^1 \times \mathbb{R}^1$ which are graphs of continuous functions $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^1$. In this way every family $\mathcal{A} = \{B_1, \dots, B_n\}$ is associated with its dual system $\mathcal{A}^* = \{B_1^*, \dots, B_n^*\}$. Notice that a body is uniquely determined by its support function (see [21, §2.2]), and consequently a family is uniquely determined by its dual system. Moreover, if a pair of dual curves intersect, that is, if $h_A(\theta) = h_B(\theta)$ for some $\theta \in \mathbb{S}^1$, then the bodies A and B have a *common supporting tangent* in the direction $\theta + \pi/2$.

A family is *generic* if the following hold.

- For any pair of bodies B_1 and B_2 of the family with common supporting tangent ℓ , the intersection $A_1 \cap A_2 \cap \ell$ is empty.
- No triple of bodies share a common supporting tangent.

A standard perturbation argument shows that the optimal values for $h_1(n)$ can be attained by generic families. We therefore assume henceforth that all families are generic. The conditions above imply the following for the dual system.

Observation 2.1. For the dual system of a generic family, each pair of curves intersects transversally and no three curves intersect in a common point.




Figure 2: (a) The support function $h_B(\theta)$ measures the distance between the origin and the directed supporting tangent in the direction $\theta + \pi/2$. (b) If a body appears on the convex hull, then it has a supporting tangent which also supports the convex hull of the union of the members of the family.

The *upper envelope* of a system given by functions $\{\gamma_i\}$ is the graph of the function $\max_i \gamma_i$. The following observation implies that the convexity properties of an family can be determined by its dual system.

Observation 2.2. A generic family is convexly independent if and only if every curve appears on the upper envelope of its dual system.

To see why this holds, notice that a family \mathcal{A} is convexly independent if and only if for any body $B \in \mathcal{A}$, the convex hull of the union of the members of \mathcal{A} can be supported by a supporting tangent of B which is disjoint from every member of $\mathcal{A} \setminus \{B\}$. If α is the direction of such a supporting tangent of B , then $h_B(\alpha - \pi/2) > h_A(\alpha - \pi/2)$ for every $A \in \mathcal{A} \setminus \{B\}$; see Figure 2(b).

2.2. Review of generalized configurations. In the sequel it will be useful to recall the duality between generalized configurations and pseudoline arrangements. This is a combinatorial extension of the classical projective duality between points and lines. From a combinatorial point of view this is very similar to our duality for families of bodies, and the connection is crucial for relating generalized configurations to families of bodies.

We start by recalling the notion of the *allowable sequence* of a set of points in the plane. We will assume that the point set is in a strongly general position, meaning no three points are collinear and no two lines determined by the points are parallel[†]. Let P be a set of n labelled points in strongly general position in the plane and consider a generic directed line l_1 . If we project the points orthogonally onto the line l_1 , then the direction of l_1 induces a linear ordering of the points which we record as a permutation $\pi_1 = \pi_1(P)$. As the line rotates counter-clockwise about a fixed point the ordering will change each time the line becomes orthogonal to a direction determined by a pair of points in P . This

[†] The general theory developed by Goodman and Pollack [14, 16] does not require this assumption, but for us it is no loss of generality.

results in a periodic sequence of permutations

$$\dots, \pi_1, \pi_2, \dots, \pi_{n(n-1)}, \pi_1, \pi_2, \dots$$

which is called the allowable sequence of P . Notice that the allowable sequence satisfies the following properties.

- (1) Any two consecutive terms π_i and π_{i+1} differ by reversing the order of two adjacent elements.
- (2) In any $\binom{n}{2}$ consecutive terms of the sequence each pair of elements of P switches exactly once.

It is an immediate consequence that for all i , the permutation $\pi_{i+\binom{n}{2}}$ is the reverse of π_i . Every allowable sequence determines a periodic sequence of ordered switches, that is, rather than writing down each permutation, we only record which ordered pair switches between consecutive permutations in the sequence. The convention is to record the order of the pair *before* the switch, for instance, the consecutive pair of permutations

$$\dots, (\dots ab \dots), (\dots ba \dots), \dots$$

will be recorded as the ordered switch ab , and consequently, $\binom{n}{2}$ steps later we get the ordered switch ba . It turns out that an allowable sequence is determined by its sequence of ordered switches, which is shown in [16, Proposition 2.6]. For instance, the following half-period of a sequence of ordered switches

$$\dots, dc, ac, bc, ad, bd, ba, \dots$$

uniquely determines the following sequence of permutations

$$\dots, \underline{badc}, \underline{bacd}, \underline{bcad}, \underline{cbad}, \underline{cbda}, \underline{cdba}, \underline{cdab}, \dots$$

More generally, any sequence of permutations which satisfies properties (1) and (2) is called an *allowable sequence*, but not every such sequence can be obtained from a set of points in the plane by the procedure described above. This is where generalized configurations come in to play: for an ordered pair of points, a and b , of a generalized configuration, consider the directed pseudoline which first passes through a then through b , and label the point where it intersects the distinguished line at infinity by the ordered pair ab (thus the antipodal point is labelled ba). In this way we obtain a cyclic sequence of the ordered pairs of points which is antipodal in the sense that a half-period after the term ab we get the reverse pair ba . It turns out that this sequence of ordered pairs is precisely the sequence of ordered switches of an allowable sequence, and [18, Theorem 4.4] shows that *every allowable sequence* can be obtained in this way.

Clearly an allowable sequence is uniquely determined by any one of its half-periods, which can be encoded by a so-called *wiring diagram*, resulting in the dual pseudoline arrangement. Let π_1, \dots, π_N be the permutations of some half-period of the allowable sequence and S_i the ordered switch from π_i to π_{i+1} (where π_{N+1} is the reverse of π_1). Construct the wiring diagram as follows.




Figure 3: (a) A generalized configuration with the labelling on the line at infinity. (b) The dual wiring diagram corresponding to the half-period of ordered switches (dc, ac, bc, ad, bd, ba) .

Start with horizontal “wires” going from left to right, labelled by the elements of the permutations in the order in which they appear in the permutation π_1 from bottom to top. Apply the switch S_1 by crossing the wires corresponding to the elements appearing in the switch S_1 . After all switches have been applied, each pair of wires will have crossed precisely once, and we arrive at the reverse of the initial permutation; see Figure 3.

2.3. Non-crossing and orientable families. A pair of bodies B_1, B_2 is *non-crossing* if $B_1 \setminus B_2$ is simply connected, or equivalently, if B_1 and B_2 have precisely two common supporting tangents. Notice that the dual curves of a non-crossing pair of bodies will meet in precisely two crossing points.

A triple of bodies B_1, B_2, B_3 is *orientable* if every pair is non-crossing and $\text{conv}(B_i \cup B_j) \setminus B_k$ is simply connected for all choices of distinct i, j, k , or equivalently, the convex hull of $B_1 \cup B_2 \cup B_3$ is supported by exactly three of the common supporting tangents determined by the pairs B_i, B_j . A *non-crossing family* is one in which each pair of bodies is non-crossing, and an *orientable family* is one in which each triple of bodies is orientable.

Each member of an orientable triple contributes a single connected arc to the boundary of its convex hull, so traversing the boundary of its convex hull in the counter-clockwise direction will impose a cyclic ordering of the triple. Notice that in the dual system of an oriented triple, each curve appears precisely once on the upper envelope in the same cyclic order as the one we get by traversing the convex hull of the bodies; see Figure 4.

It is easily verified that the set of cyclic orderings of all triples of an orientable family satisfy the chirotope axioms of a rank 3 uniform oriented matroid [6, Definition 3.5.3], or equivalently the axioms of a CC-system (see [25, §1])[†]. This means that for every orientable family \mathcal{A} , there exists a generalized configuration \mathcal{P} and a bijection $\phi : \mathcal{A} \rightarrow \mathcal{P}$ which preserves the cyclic ordering of every triple. (The cyclic ordering of a triple in a generalized

[†] Grünbaum implicitly makes this observation in his discussion on planar arrangements of simple curves in [22, §3.3].




Figure 4: (a) A typical orientable triple with the three common supporting tangents which support the convex hull. Between consecutive common supporting tangents the boundary of the convex hull consists of a boundary arc of one of the bodies, which induces a cyclic ordering of the bodies. (b) The dual system of an orientable triple. The cyclic order in which the curves appear on the upper envelope (when traversed from left to right) coincides with the cyclic order in which we meet the bodies when traversing the boundary of the convex hull in the counter-clockwise order.

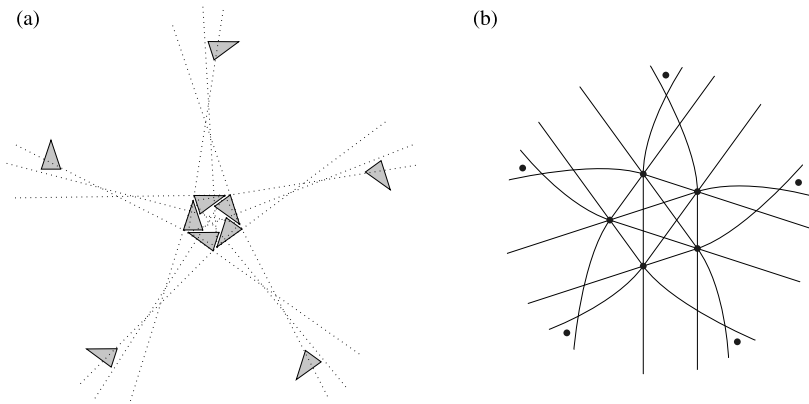


Figure 5: A family of convex bodies (a) and a realization of this family by a generalized configuration (b). This configuration is based on Goodman and Pollack’s “bad pentagon” [16] and can not be realized by points and straight lines.

configuration is defined, as before, by traversing the boundary of its convex hull in the counter-clockwise direction and reading off the cyclic order in which we meet the points.) See Figure 5.

When there exists such an order-preserving bijection as described above, we say that the family is *realizable* by the generalized configuration, and we may also say that the generalized configuration is realizable by the family. Our discussion above implies the following lemma.

LEMMA 2.3. *Every orientable family is realizable by a generalized configuration.*

We now establish the converse of Lemma 2.3.

LEMMA 2.4. *Every generalized configuration is realizable by an orientable family.*

What is of importance to us is that the convex independencies can be determined from the set of cyclic orderings of the triples. (See for instance [6, §9.1], [15, §5.2], or [25, §11].) Therefore Lemmas 2.3 and 2.4 imply the following.

COROLLARY 2.5. *The Erdős–Szekeres problems for generalized configurations and for orientable families are equivalent.*

This, however, is not enough to prove Theorem 1.1 since there exist families which are not orientable. These will be dealt with in the next section.

Remark 2.6. In view of Observations 2.1 and 2.2 it should be clear that, for our purpose, the precise geometric information of the dual system is not of major importance. What matters is the intersection pattern of the dual curves which allows us to determine the upper envelopes of any sub-system. In the figures below the dual systems will be represented by “schematic diagrams” similar to wiring diagrams.

Proof of Lemma 2.4. Let \mathcal{P} be a generalized configuration. Consider its dual wiring diagram \mathcal{W} which encodes some half-period of the allowable sequence of \mathcal{P} , as described in §2.2. We can view \mathcal{W} as a system of curves drawn on the Möbius strip, each pair crossing once with all crossing points distinct. Extending \mathcal{W} to its double cover, we obtain a system of curves \mathcal{F} on the cylinder $\mathbb{S}^1 \times \mathbb{R}^1$, each pair crossing twice, with all crossing points distinct.

We now notice that the cyclic ordering of a triple of \mathcal{P} is encoded in \mathcal{F} by the order in which the corresponding triple of curves appear on the upper envelope (of the triple) in the double cover. To see why this happens it suffices to consider a triple of points, say, with cyclic order (a, b, c) . A full period of the corresponding sequence of ordered switches will be

$$ab, ac, bc, ba, ca, cb$$

giving us the allowable sequence

$$\dots, abc, bac, bca, cba, acb, abc, \dots$$

The wires that appear on the upper envelope of a full period of the double cover of the corresponding wiring diagram are the last entries of each permutation; see Figure 6.

We now interpret the curves of \mathcal{F} as graphs of functions. We will see that they can be modified into graphs of support functions without changing the intersection pattern of the system.

Clearly the curves of \mathcal{F} can be approximated by the graphs of C^2 -smooth functions $f: \mathbb{S}^1 \rightarrow \mathbb{R}^1$, without changing the crossing patterns and the upper envelopes.

It is known that if f is a 2π -periodic real C^2 -smooth function such that $f(t) + f''(t) > 0$ holds for all t , then f is the support function of a convex body. (See for instance [21, Lemma 2.2.3].) Hence, for any 2π -periodic real C^2 -smooth




Figure 6: The cyclic ordering of each triple of points (a) corresponds to the order in which the wires appear on the upper envelope of the double cover of the wiring diagram (b).

function f , there exists a constant c_0 such that $f + c$ is the support function of a body for all $c > c_0$. Consequently, there is a common constant we can add to each of the (smoothened) functions which makes it the dual system of a family of bodies. \square

2.4. *Weak maps.* Let \mathcal{A} and \mathcal{B} be families of bodies with $|\mathcal{A}| = |\mathcal{B}|$. A bijection $\phi : \mathcal{A} \rightarrow \mathcal{B}$ is a *weak map* if $\phi^{-1}(\mathcal{B}')$ is convexly independent for every convexly independent subfamily $\mathcal{B}' \subset \mathcal{B}$. The inequality $h_1(n) \leq g(n)$ is a consequence of the following lemma.

LEMMA 2.7. *For every non-crossing family \mathcal{A} , in which every triple is convexly independent, there exists a weak map $\phi : \mathcal{A} \rightarrow \mathcal{B}$ where \mathcal{B} is an orientable family.*

Lemma 2.7 is trivial if $|\mathcal{A}| \leq 2$, so we may assume \mathcal{A} consists of at least three bodies. The dual system \mathcal{A}^* induces a cell complex $\mathcal{C}(\mathcal{A})$, homeomorphic to $\mathbb{S}^1 \times [0, 1]$, and the weak map ϕ will be defined in terms of elementary operations on $\mathcal{C}(\mathcal{A})$. Since every triple of \mathcal{A} is convexly independent, there are two types of triples to consider: those that are *orientable* and those that are *non-orientable*. The orientable triples were discussed in §2.3. The dual of a non-orientable triple \mathcal{T} is characterized by one of its support curves appearing two distinct times on the upper envelope of \mathcal{T}^* . In $\mathcal{C}(\mathcal{T})$ this corresponds to a pair of disjoint triangular cells whose top edges are both contained in the same support curve. Notice that in the non-orientable case these are the only triangular cells, while in the orientable case every cell is triangular.

An equivalent way of distinguishing the two types of triple is by considering the cyclic order in which a curve intersects the other two in the dual system. We call the cyclic sequence (x, y, x, y) *alternating*, and the cyclic sequence (x, x, y, y) *separating*.

Observation 2.8. Let \mathcal{T} be a convexly independent triple of bodies with dual system \mathcal{T}^* .

- (1) If \mathcal{T} is orientable, then for any $\gamma \in \mathcal{T}^*$ the cyclic order in which γ intersects the curves of $\mathcal{T}^* \setminus \{\gamma\}$ is alternating.
- (2) If \mathcal{T} is non-orientable, then for any $\gamma \in \mathcal{T}^*$ the cyclic order in which γ intersects the curves of $\mathcal{T}^* \setminus \{\gamma\}$ is separating.

To see why this holds we refer the reader to Figure 7.




Figure 7: The two types of triples of bodies, orientable (a) and non-orientable (b). Notice that the cyclic order in which we meet the bodies when traversing the convex hull corresponds to the order in which we meet the curves when transversing the upper envelope. In the orientable case (a) any curve meets the others alternately; for instance, the cyclic order in which a meets b and c is (b, c, b, c) . In the non-orientable case a meets b and c in the cyclic order (b, c, c, b) .




Figure 8: A non-orientable triple (a) and the orientable triple (b) obtained after applying a triangle flip.

A non-orientable triple \mathcal{T} is related to an orientable one by an elementary operation called a *triangle flip*, which is defined by “flipping” the orientation of one of the two triangular cells of $\mathcal{C}(\mathcal{T})$. Notice that a triangle flip defines a weak map from a non-orientable triple to an orientable one; see Figure 8.

We deduce Lemma 2.7 from the following lemma.

LEMMA 2.9. *If \mathcal{A} is not orientable, then $\mathcal{C}(\mathcal{A})$ contains a triangular cell bounded by the support curves of a non-orientable triple.*

Proof of Lemma 2.7. If \mathcal{A} is not orientable, Lemma 2.9 implies that we can apply a triangle flip to $\mathcal{C}(\mathcal{A})$ obtaining a new cell complex \mathcal{C}' . We may assume \mathcal{C}' is the dual system of a family \mathcal{A}' (as shown in the proof of Lemma 2.4). This induces a weak map $\phi' : \mathcal{A} \rightarrow \mathcal{A}'$. Since a triangle flip reduces the number of non-orientable triples, Lemma 2.7 follows by induction. \square

A few technical terms are needed for proving Lemma 2.9. Let \mathcal{T} be a non-orientable triple. The *top edges* of the two triangular cells of $\mathcal{C}(\mathcal{T})$ belong to the same support curve, called the *top curve*, which appears twice on the upper envelope of \mathcal{T}^* . When \mathcal{T} is a non-orientable triple belonging to a larger




Figure 9: (Colour online) The triple $\mathcal{T}^* = \{a, b, c\}$ bounds two zones (shaded) and the top curve is b . Neither of the zones of \mathcal{T}^* are empty, but the left one is free.

family \mathcal{F} , the triangular cells of $\mathcal{C}(\mathcal{T})$ may no longer be cells in $\mathcal{C}(\mathcal{F})$, so instead we refer to these open triangular regions as the *zones* of \mathcal{T}^* . When we say that \mathcal{T}^* bounds a zone, it is implicit that \mathcal{T} is non-orientable. A zone is called *empty* if no curve of the system intersects its interior, and is called *free* if no curve intersects its top edge; see Figure 9.

We may restate Lemma 2.9 as follows.

If \mathcal{A} contains a non-orientable triple, then $\mathcal{C}(\mathcal{A})$ contains an empty zone.

For the proof we will consider a minimal counter-example. It is, however, easier to handle free zones rather than empty ones, so we first establish the following claim.

CLAIM 2.10. *If $\mathcal{C}(\mathcal{A})$ contains a free zone, then $\mathcal{C}(\mathcal{A})$ contains an empty zone.*

Proof. We assume without loss of generality that Z_0 is a free zone bounded by a, b, c where b is the top curve. Let w_1, \dots, w_k denote curves that intersect Z_0 . We first make some simple observations (these will also be of use later).

- (1) *Each w_i intersects Z_0 in a single connected arc. We may assume w_i enters Z_0 by crossing curve c and exits Z_0 by crossing curve a .*

To see why this holds assume for contradiction that w_i intersects Z_0 in more than one connected arc. Up to symmetry we may then assume w_i enters Z_0 by crossing c and then immediately crosses c again. In this case, whenever w_i is above curve c , it is also below curve b . Consequently only b and c appear on the upper envelope of the triple b, c, w_i , which contradicts our assumption that every triple of \mathcal{A} is convexly independent.

- (2) *The triangular region in Z_0 bounded by a, w_i, c is a zone.*

If the region in question is not a zone, then w_i should meet the curves a and c alternately. Therefore, after w_i exits Z_0 by crossing a it should meet curve c before it meets curve a again. Arguing as in (1), this would imply that only curves b and c appear on the upper envelope of the triple b, c, w_i .

- (3) *Distinct curves w_i and w_j cross at most once inside Z_0 .*

Suppose w_i and w_j are distinct curves which cross twice inside Z_0 . We may assume the w_i enters Z_0 above w_j , which implies that w_i also exits Z_0 above w_j . Therefore the only time w_j is above w_i it will be below b .




Figure 10: (Colour online) (a) The zone Z is bounded by a, b and c (shaded). After w leaves Z and crosses a it enters a digon bounded by curves a and b , so it must cross one of them again before crossing c . (b) If the next crossing of w is with a , then a, w, b bounds a zone (shaded). (c) If the next crossing of w is with b then w, b, c bounds a zone (shaded).

Consequently, the upper envelope of b, w_i, w_j consists only of curves b and w_i .

We also need the following observation concerning zones that are not free.

Observation 2.11. Let Z be a zone bounded by a, b, c where b is the top curve. Suppose w enters Z by crossing c and exits Z by crossing b , then proceeds to cross a . Then one of the triples a, w, b or w, b, c bounds a zone with top vertex at the crossing point between b and w on the top edge of Z .

To see why this holds we notice that after w leaves Z and crosses a , it enters a digon bounded by curves a and b . This means that the next curve that w crosses must be a or b . If w exits the digon by crossing curve a , then w crosses curves a and b in a separating cyclic sequence, which implies that the triple a, w, b is non-orientable, and therefore bounds a zone. Otherwise, w exits the digon by crossing curve b , which implies that w crosses curves b and c in a separating cyclic sequence, and consequently w, b, c is non-orientable, and bounds a zone; see Figure 10.

For the proof of Claim 2.10 we proceed by induction on k , the number of curves which intersect Z_0 . If $k = 0$, then Z_0 is an empty zone, so assume $k > 0$. Start at the top left corner of Z_0 at the crossing between b and c . Move on the boundary of Z_0 along c and stop at the first crossing we encounter. Assume that this is the crossing between c and w_k . This crossing is the top corner of a zone $Z_k \subset Z_0$ (bounded by a, w_k, c by (2) above). Move into the interior of Z_0 along curve w_k (the top edge of Z_k) and stop at the first crossing we encounter. If this is the crossing between w_k and a , we may apply the induction step, since then Z_k is a free zone with less than k intersecting curves. So assume that this is a crossing between w_k and w_{k-1} . By Observation 2.11 (with $b = w_k$ and $w = w_{k-1}$) this crossing is the top vertex of a zone Z_{k-1} . If $Z_{k-1} \subset Z_k$ (i.e. bounded by w_{k-1}, w_k, c), then Z_{k-1} is free and we are done by induction, so assume Z_{k-1} is bounded by curves a, w_{k-1}, w_k . Now proceed along curve w_{k-1}




Figure 11: (Colour online) Starting at top left corner of Z_0 (light shade) move along the boundary until we meet the first crossing. This is the top corner of a zone bounded by a , w_k and c . Proceed along w_k until we meet the next crossing. By Observation 2.11 one of the two dark shaded regions must be a zone.

(the top edge of Z_{k-1}) and repeat the process. In general, we proceed from the left top vertex of the zone Z_j along the curve w_j (the top edge of Z_j) and stop at the first curve we meet. If the first curve we meet is a , then we are done because Z_j is a free zone intersected by fewer than k curves. Otherwise we meet curve w_{j-1} , and by Observation 2.11 this crossing point is the top vertex of a zone Z_{j-1} . If $Z_{j-1} \subset Z_j$ then we are done, or else we can repeat by proceeding along w_{j-1} (the top edge of Z_j). This process simply amounts to moving along the upper envelope of the curves a, c, w_1, \dots, w_k within the zone Z . By (3) above, each pair w_i and w_j cross at most once within Z , so each curve can appear on the upper envelope at most once (inside Z). Since there are only finitely many curves w_i the process must eventually end, either with a free zone $Z_i \subset Z_{i+1}$, or in the case where a is the first curve which w_i meets after appearing on the upper envelope, which implies that w_{i+1}, w_i, a bounds a free zone $Z_i \subset Z$. In either case we eventually reach a free zone which is crossed by less than k curves, completing the proof of Claim 2.10; see Figure 11. \square

We are in position to complete the proof of Lemma 2.9.

Proof of Lemma 2.9. Suppose \mathcal{A} is a minimal counter-example. Then $\mathcal{C}(\mathcal{A})$ contains zones, but no empty ones, and any proper subfamily $\mathcal{A}' \subset \mathcal{A}$ is either orientable or the complex $\mathcal{C}(\mathcal{A}')$ has at least one empty zone. We will reach a contradiction by showing that $\mathcal{C}(\mathcal{A})$ contains a free zone.

Assume first that any curve we delete from the lower envelope of \mathcal{A}^* (defined as the graph of the function $\min_i \gamma_i$) destroys all non-orientable triples. Then the lower envelope consists of exactly two curves a and c , and a triple is non-orientable if and only if it includes both of these curves. To see this, note that if there were three curves on the lower envelope, then these form an orientable triple, so for any non-orientable triple there is a curve on the lower envelope not belonging to it. Let Z be a zone bounded by a, b, c . Some curve w should intersect the top edge of Z or else it is free. Then, by Observation 2.11, one of the triples a, b, w or b, c, w is non-orientable, contradicting our initial assumption.

We may therefore assume that there is a curve w appearing on the lower envelope, and a triple a, b, c which bounds a zone Z where b is the top curve,




Figure 12: (Colour online) Consider w after it leaves Z . Case (1), (a): If w crosses b before c , then w, b, c bounds an empty zone contained in Z . If w crosses c before b , then w, a, c does not come from a convexly independent triple. Case (2), (b): If w crosses a before c , then w, a, c bounds a free zone below Z . If w crosses c before a , then b intersects w again after its two crossings with a , which implies that w, a, b does not come from a convexly independent triple.

and w is the *only* curve which meets the interior of Z . Furthermore w must cross the top edge of Z (if not Z is free, contradicting the assumption that \mathcal{A} was a counter-example). Since w is on the lower envelope it must also cross one of the edges of Z . Up to symmetry there are two cases that can occur; see Figure 12:

- (1) w is on the lower envelope, crosses a , then c (entering Z), and then b (leaving Z);
- (2) w is on the lower envelope, crosses c , then a (entering Z), and then b (leaving Z).

In both cases we consider the order in which w intersects the other curves after leaving Z .

Case (1) results in a zone bounded by w, b, c contained in Z . This zone must be empty, since w is the only curve that intersects Z . We use Observation 2.8 to argue that w, b, c comes from a non-orientable triple. Notice that after w leaves Z by crossing b it must cross b again before crossing c , or else only curves c and w appear on the upper envelope of the triple a, w, c . Therefore, the cyclic order in which w intersects curves b and c is separating.

Case (2) results in a zone bounded by a, w, c where a is the top curve. This zone is adjacent to Z along the curve a which implies that it is free (since w is the only curve intersecting the interior of Z). Again we use Observation 2.8 to argue that a, w, c comes from a non-orientable triple. This depends on the order in which curve w meets curves a and c after leaving Z . However, if w intersects c before a , then only curves b and w appear on the upper envelope of the triple a, w, b .

We can therefore conclude that there always exists a free zone. □

§3. *The upper bound for generalized configurations.* Here we show that $g(n) \leq \binom{2n-5}{n-2} + 1$ for all $n \geq 5$. The proof is dual to the proof given by Tóth and Valtr in [35], but more general since it is in terms of arbitrary wiring diagrams.

There are two basic wiring diagrams which we refer to throughout. A k -cup is the unique wiring diagram on k wires in which every wire appears on the upper envelope, and an l -cap is the unique wiring diagram on l wires in which every wire appears on the lower envelope; see Figure 13.




Figure 13: A 5-cup (a) and a 5-cap (b).

Observation 3.1. Let \mathcal{P} be a generalized configuration and W the wiring diagram corresponding to any half-period of the allowable sequence of \mathcal{P} . If W contains a k -cup (or k -cap), then the wires of the k -cup (or k -cap) correspond to a k -tuple of \mathcal{P} which is convexly independent.

It is important to note that the converse of Observation 3.1 does not hold. That is, \mathcal{P} may contain convexly independent subsets of size k , while no half-period of its allowable sequence contains a k -cup or k -cap. Also notice that every wiring diagram on three wires is a 3-cup or a 3-cap. In the following observation we assume that the wires are ordered by the way in which they intersect a vertical line at the start (left) of the half-period.

Observation 3.2. Suppose W is a k -cup and W' is an l -cap where x is the *bottom* wire of W and the *top* wire of W' . Then $W \cup W'$ contains a $(k + 1)$ -cup or an $(l + 1)$ -cap.

This observation is dual to the one appearing in [9]. It also holds if we switch the roles of the top and bottom wires. To see why the observation holds, consider the order in which the wire x intersects the other wires of W and of W' . If x intersects every wire of W before intersecting every wire of W' , then W can be extended to a $(k + 1)$ -cup. Otherwise, W' can be extended to an $(l + 1)$ -cap.

Observation 3.3. Every wiring diagram on $\binom{k+l-4}{k-2} + 1$ wires contains a k -cup or an l -cap.

The proof is the same as the one given by Erdős and Szekeres [9]. One proceeds by induction on $k + l$ and defines A to be the set of wires which are the bottom wire of some $(k - 1)$ -cup. By induction, it follows that $|A| \geq \binom{k+l-5}{k-3} + 1$, thus A contains a k -cup or an $(l - 1)$ -cup. In the latter case, Observation 3.2 implies the existence of a k -cup or an l -cup.

We now formulate the lemma which lies at the heart of the proof of Tóth and Valtr [35].

LEMMA 3.4. *Let \mathcal{P} be a generalized configuration with $|\mathcal{P}| = \binom{2n-5}{n-2}$ and $n \geq 5$. Let W be a wiring diagram corresponding to any half-period of the allowable sequence of \mathcal{P} . If no subset of n points of \mathcal{P} is convexly independent, then W contains an $(n - 1)$ -cup.*

Remark 3.5. Notice that if $|\mathcal{P}| = \binom{2n-5}{n-2} + 1$, then the conclusion of Lemma 3.4 is guaranteed by Observation 3.3, thus the extra effort which goes into proving Lemma 3.4 only pays off an improvement of a single point!




Figure 14: (Colour online) Sending the left crossings over to the right side through the line at infinity.

Before getting to the proof of Lemma 3.4, let us see why this implies $g(n) \leq \binom{2n-5}{n-2} + 1$. Let \mathcal{P} be a generalized configuration with $|\mathcal{P}| = \binom{2n-5}{n-2} + 1$ and let W be the wiring diagram of one of the half-periods of its allowable sequence. We insert a new wire x starting above every wire of W such that it first crosses the initial top wire w of W while w is on the upper envelope, and then follows wire w , slightly below, meeting the remaining wires of W in the same order as w does. In this way x crosses every wire of W precisely once.

The wire x separates the crossings of W into “left crossings” and “right crossings”, and the idea is to shift the left crossings over to the right side by sending them through the line at infinity. More precisely, consider the first crossing on the left side of x representing the ordered switch ij . At the right end of the wiring diagram, these wires will be adjacent, but in the opposite order. Delete the ij crossing from the left side and add the corresponding ji crossing on the right side at the end of the wiring diagram. The procedure can be repeated until all left crossings have been evacuated, after which we can apply a homeomorphism of the plane which maps the wire x to a vertical line. We now delete wire x and obtain a new wiring diagram W' on the same number of wires as W , with the property that the top wire of W' crosses every wire below before any other crossings occur; see Figure 14.

It is easily verified that the above operation does not change the orientation of any triples. That is, when we extend W' to its double cover, we obtain the full period of the allowable sequence of a generalized configuration \mathcal{P}' , and there will be a bijection $\varphi : \mathcal{P} \rightarrow \mathcal{P}'$ which preserves the cyclic order of every triple of points. In particular, a subset of \mathcal{P} is convexly independent if and only if its image under φ is convexly independent. (Recall the proof of Lemma 2.4.)

We may therefore assume that the top wire of W intersects every other wire of W before any other crossing occurs, and then consider the sub diagram $W_0 \subset W$ obtained by deleting the top wire from W . Now W_0 corresponds to the half-period of a subconfiguration $\mathcal{P}_0 \subset \mathcal{P}$ consisting of $\binom{2n-5}{n-2}$ points. If \mathcal{P}_0 contains n points which are convexly independent, then we are done. Otherwise, Lemma 3.4 implies that W_0 contains an $(n-1)$ -cup, which, together with the wire we deleted from W , forms an n -cup.

Proof of Lemma 3.4. Let \mathcal{P} be a generalized configuration on $\binom{2n-5}{n-2}$ points which does not contain any convexly independent subset of size n . Fix a

half-period of the allowable sequence of \mathcal{P} and let W be its wiring diagram. By assumption, W contains no n -cap, and suppose for contradiction that W contains no $(n - 1)$ -cup.

Let A denote the set of wires which are the bottom wire of some $(n - 2)$ -cup and let B denote the remaining wires. If $|A| > \binom{2n-6}{n-3}$, then Observation 3.3 implies that A contains an $(n - 1)$ -cup, since we have assumed that there is no $(n - 1)$ -cup. But then Observation 3.2 implies that A contains an $(n - 1)$ -cup or an n -cap. Therefore $|A| \leq \binom{2n-6}{n-3}$. On the other hand, if $|B| > \binom{2n-6}{n-2}$, then B contains an $(n - 2)$ -cup or an n -cap, by Observation 3.3, which contradicts the definition of A . We conclude that $|A| = \binom{2n-6}{n-3}$ and $|B| = \binom{2n-6}{n-2}$.

Next, we claim that for every $a \in A$, the set $B \cup \{a\}$ contains an $(n - 2)$ -cup whose bottom wire is a . This follows from Observation 3.3 and the definition of A . Similarly, we claim that for every $b \in B$ the set $A \cup \{b\}$ contains an $(n - 1)$ -cup whose top wire is b . This follows from Observation 3.3. Indeed, if A contains an $(n - 1)$ -cup, then Observation 3.2 would imply the existence of an $(n - 1)$ -cup or an n -cap.

Now consider the set of pairs (a, b) with $a \in A$ and $b \in B$ such that a is the bottom wire and b the top wire of an $(n - 2)$ -cup or an $(n - 1)$ -cup, and choose the pair (a, b) whose crossing is the leftmost among all such pairs. We now deal with two separate cases depending on whether the pair (a, b) bounds an $(n - 2)$ -cup or an $(n - 1)$ -cup.

For the first case, suppose (a, b) bounds an $(n - 2)$ -cup whose wires appear on the upper envelope in order b, w, \dots, a . By the argument above, there exists an $(n - 1)$ -cup whose wires appear on the lower envelope in order a', w', \dots, b . Note that $a \neq a'$ or else \mathcal{P} contains a convexly independent subset of size $2n - 3$, so by our assumption b meets a before a' . To see this, extend to the full period and observe that every wire of the $(n - 2)$ -cup and the $(n - 1)$ -cup appear on the upper envelope. Next, observe that the starting point of a' is between the starting points of b and w . If not, then a' meets b before w , which implies that the wires of the $(n - 1)$ -cup, a', w', \dots, b together with w , correspond to a convexly independent subset of \mathcal{P} , which can be seen, again, by extending to the double cover. Finally, consider the order in which w' meets wires a' and w . If w' meets a' before w , then there is an $(n - 1)$ -cup whose wires appear in the order w', a', w, \dots, a . Otherwise, there is an n -cap whose wires appear in the order w, a, w', \dots, b ; see Figure 15(c).

The second case is similar. Suppose (a, b) bounds an $(n - 1)$ -cup whose wires appear on the lower envelope in order a, \dots, w, b . There exists an $(n - 2)$ -cup whose wires appear on the upper envelope in order b', \dots, w', a . As before, we must have $b \neq b'$, so by our assumption a meets b before b' . Moreover, the starting points of b' must be above the starting point of b , or else the wires of the $(n - 1)$ -cup a, \dots, w, b together with b' correspond to a convexly independent subset of \mathcal{P} (again, extend to the double cover). Finally, consider the order in which w' meets wires b and w . If w' meets b before w , then there is an $(n - 1)$ -cup whose wires appear in the order b, \dots, w', b, w . Otherwise, there is an n -cap whose wires appear in the order a, \dots, w, b, w' . See Figure 15(b). \square




Figure 15: First case (a): (a, b) bound an $(n - 2)$ -cup whose wires appear in order b, w, \dots, a . Second case (b): (a, b) bound an $(n - 1)$ -cap whose wires appear in order a, \dots, w, b .

§4. *Further generalizations.* A *convex k -clustering* is a disjoint union of point sets $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k$ of equal size such that all k -tuples (p_1, p_2, \dots, p_k) with $p_i \in \mathcal{S}_i$ are convexly independent. The cardinality of the \mathcal{S}_i is the *size* of the clustering. This notion naturally extends to families of bodies as well.

4.1. *The positive fraction version.* Bárány and Valtr gave the following generalization of the Erdős–Szekeres theorem, known as the *positive fraction Erdős–Szekeres theorem*.

THEOREM (Bárány–Valtr [2]). *For every integer $k > 3$ there exists a constant $c_k > 0$ such that any finite set \mathcal{S} in the Euclidean plane, in which every triple is convexly independent, contains a convex k -clustering of size $c_k |\mathcal{S}|$.*

The current best value for c_k is due to Pór and Valtr [30], and shows that $c_k \geq k \cdot 2^{-32k}$. Their argument can be repeated verbatim to hold for generalized configurations as well. By Lemmas 2.4 and 2.7 the positive fraction ES for generalized configurations is equivalent to the positive fraction ES for non-crossing families of bodies, and therefore holds with the same bound on c_k .[†]

THEOREM 4.1. *For every integer $n > 3$ there exists a constant $c_n > 0$ such that any non-crossing family \mathcal{A} , in which every triple is convexly independent, contains a convex n -clustering of size $c_n |\mathcal{A}|$.*

4.2. *The partitioned version.* Answering a question of Kalai, the positive fraction ES was further generalized by Pór and Valtr [30] with what is called the *partitioned Erdős–Szekeres theorem*.

THEOREM (Pór–Valtr [30]). *For every $k \geq 3$ there exist constants $p = p_k$ and $r = r_k$ such that for any finite set \mathcal{S} in the Euclidean plane, in which every triple is convexly independent, there is a subset \mathcal{S}' of size at most r such that $\mathcal{S} \setminus \mathcal{S}'$ can be partitioned into at most p convex k -clusterings.*

Extending this theorem to generalized configurations can be done in a more or less routine way. Essentially one needs to modify the proofs of Claims 1–3 in [30]. This can be done by replacing any “distance arguments” by “continuous sweep arguments” (see for instance [7, 20]). The remaining parts of the proof

[†] The positive fraction Erdős–Szekeres theorem was first established for families of pairwise disjoint bodies by Pach and Solymosi [27], and their method was subsequently improved by Pór and Valtr [31]. Our methods imply a substantial quantitative improvement, as well as relaxing the disjointness assumption.

of Pór and Valtr are combinatorial, and do not need further modification. By applying Lemmas 2.4 and 2.7 it follows that the fractional ES theorems for generalized configurations and for non-crossing families are equivalent. We therefore obtain the following theorem[†].

THEOREM 4.2. *For every $k \geq 3$ there exist constants $p = p_k$ and $r = r_k$ such that the following holds. For every non-crossing family \mathcal{A} in which every triple is convexly independent there is a subfamily \mathcal{A}' of size at most r such that $\mathcal{A} \setminus \mathcal{A}'$ can be partitioned into at most p convex k -clusterings.*

Remark 4.3. It is natural to ask whether the non-crossing condition can be further relaxed. In a subsequent paper [8] we show that this is indeed the case, confirming a conjecture of Pach and Tóth [29].

Acknowledgement. We are grateful to the anonymous referee for many useful comments which helped improve the exposition of this paper.

M. G. Dobbins was supported by NRF grant 2011-0030044 funded by the government of South Korea (SRC-GAIA) and BK21.

A. Holmsen was supported by Basic Science Research Program through the National Research Foundation of Korea funded by the Ministry of Education, Science and Technology (NRF-2010-0021048).

A. Hubard was supported by Fondation Sciences Mathématiques de Paris. A. Hubard would like to thank KAIST for their hospitality and support during his visit.

References

1. I. Bárány and G. Károlyi, Problems and results around the Erdős–Szekeres convex polygon theorem. In *Discrete and computational geometry (Tokyo, 2000) (Lecture Notes in Computer Science)*, Springer (Berlin, 2001), 91–105.
2. I. Bárány and P. Valtr, A positive fraction Erdős–Szekeres theorem. *Discrete Comput. Geom.* **19** (1998), 335–342.
3. T. Bisztriczky and G. Fejes Tóth, A generalization of the Erdős–Szekeres convex n -gon theorem. *J. Reine Angew. Math.* **395** (1989), 167–170.
4. T. Bisztriczky and G. Fejes Tóth, Nine convex sets determine a pentagon with convex sets as vertices. *Geom. Dedicata* **31** (1989), 89–104.
5. T. Bisztriczky and G. Fejes Tóth, Convexly independent sets. *Combinatorica* **10** (1990), 195–202.
6. A. Björner, M. Las Vergnas, B. Sturmfels, N. White and G. M. Ziegler, Oriented matroids. In *Encyclopedia of Mathematics and its Applications*, 2nd edn, Cambridge University Press (1999).
7. R. Dhandapani, J. E. Goodman, A. Holmsen, R. Pollack and S. Smorodinsky, Convexity in topological affine planes. *Discrete Comput. Geom.* **38** (2007), 243–257.
8. M. G. Dobbins, A. Holmsen and A. Hubard, Regular systems of paths and families of convex sets in convex position. *Trans. Amer. Math. Soc.* (to appear).
9. P. Erdős and G. Szekeres, A combinatorial problem in geometry. *Compositio Math.* **2** (1935), 463–470.
10. P. Erdős and G. Szekeres, On some extremum problems in elementary geometry. *Ann. Univ. Sci. Budapest. Eötvös Sect. Math.* **3–4** (1960/1961), 53–62.
11. S. Felsner and P. Valtr, Coding and counting arrangements of pseudolines. *Discrete Comput. Geom.* **46** (2011), 405–416.

[†] This theorem was announced by Pór and Valtr in [30] for the case of pairwise disjoint bodies, but their proof is complicated and appears only in an unpublished manuscript.

12. J. Folkman and J. Lawrence, Oriented matroids. *J. Combin. Theory Ser. B* **25** (1978), 199–236.
13. J. Fox, J. Pach, B. Sudakov and A. Suk, Erdős–Szekeres-type theorems for monotone paths and convex bodies. *Proc. London Math. Soc.* **105** (2012), 953–982.
14. J. E. Goodman, Proof of a conjecture of Burr, Grünbaum, and Sloane. *Discrete Math.* **32** (1980), 27–35.
15. J. E. Goodman, Pseudoline arrangements. In *Handbook of Discrete and Computational Geometry (Discrete Mathematics and Its Applications)*, CRC (Boca Raton, FL, 1997), 83–109.
16. J. E. Goodman and R. Pollack, On the combinatorial classification of nondegenerate configurations in the plane. *J. Combin. Theory Ser. A* **29** (1980), 220–235.
17. J. E. Goodman and R. Pollack, A combinatorial perspective on some problems in geometry. In *Proceedings of the Twelfth Southeastern Conference on Combinatorics, Graph Theory and Computing, Vol. I* (Baton Rouge, LA, 1981), 383–394.
18. J. E. Goodman and R. Pollack, Semispaces of configurations, cell complexes of arrangements. *J. Combin. Theory Ser. A* **37** (1984), 257–293.
19. J. E. Goodman and R. Pollack, Upper bounds for configurations and polytopes in \mathbf{R}^d . *Discrete Comput. Geom.* **1** (1986), 219–227.
20. J. E. Goodman, R. Pollack, R. Wenger and T. Zamfirescu, Arrangements and topological planes. *Amer. Math. Monthly* **101** (1994), 866–878.
21. H. Groemer, Geometric applications of Fourier series and spherical harmonics. In *Encyclopedia of Mathematics and its Applications*, Cambridge University Press (1996).
22. B. Grünbaum, *Arrangements and spreads*, American Mathematical Society (Providence, RI, 1972).
23. L. Habert and M. Pocchiola, LR characterization of chirotopes of finite planar families of pairwise disjoint convex bodies. *Discrete Comput. Geom.* **50** (2013), 552–648.
24. A. Hubard, L. Montejano, E. Mora and A. Suk, Order types of convex bodies. *Order* **28** (2011), 121–130.
25. D. E. Knuth, *Axioms and Hulls (Lecture Notes in Computer Science 606)*, Springer (Berlin, 1992).
26. W. Morris and V. Soltan, The Erdős–Szekeres problem on points in convex position—a survey. *Bull. Amer. Math. Soc. (N.S.)* **37** (2000), 437–458.
27. J. Pach and J. Solymosi, Canonical theorems for convex sets. *Discrete Comput. Geom.* **19** (1998), 427–435.
28. J. Pach and G. Tóth, A generalization of the Erdős–Szekeres theorem to disjoint convex sets. *Discrete Comput. Geom.* **19** (1998), 437–445.
29. J. Pach and G. Tóth, Erdős–Szekeres-type theorems for segments and noncrossing convex sets. *Geom. Dedicata* **81** (2000), 1–12.
30. A. Pór and P. Valtr, The partitioned version of the Erdős–Szekeres theorem. *Discrete Comput. Geom.* **28** (2002), 625–637.
31. A. Pór and P. Valtr, On the positive fraction Erdős–Szekeres theorem for convex sets. *European J. Combin.* **27** (2006), 1199–1205.
32. G. Ringel, Über Geraden in allgemeiner Lage. *Elem. Math.* **12** (1957), 75–82.
33. G. Szekeres and L. Peters, Computer solution to the 17-point Erdős–Szekeres problem. *ANZIAM J.* **48** (2006), 151–164.
34. G. Tóth, Finding convex sets in convex position. *Combinatorica* **20** (2000), 589–596.
35. G. Tóth and P. Valtr, The Erdős–Szekeres theorem: upper bounds and related results. In *Combinatorial and Computational Geometry (Mathematical Sciences Research Institute Publications 52)*, Cambridge University Press (Cambridge, 2005), 557–568.

Michael Gene Dobbins,
GAIA, Postech, Pohang,
South Korea
E-mail: dobbins@postech.ac.kr

Andreas Holmsen,
Department of Mathematical Sciences,
KAIST, Daejeon,
South Korea
E-mail: andreash@kaist.edu

Alfredo Hubard,
Département d’informatique,
École Normale Supérieur,
Paris,
France
E-mail: hubard@di.ens.fr