

Realization Spaces of Arrangements of Convex Bodies

Michael Gene Dobbins¹ · Andreas Holmsen² ·
Alfredo Hubard³

Received: 4 May 2015 / Revised: 19 January 2017 / Accepted: 28 January 2017 /
Published online: 15 March 2017
© Springer Science+Business Media New York 2017

Abstract We introduce *combinatorial types* of planar arrangements of convex bodies, extending order types of point sets to arrangements of convex bodies, and study their realization spaces. Our main results witness a trade-off between the combinatorial complexity of the bodies and the topological complexity of their realization space. First, we show that every combinatorial type is realizable and its realization space is contractible under mild assumptions. Second, we prove a universality theorem that says the restriction of the realization space to arrangements polygons with a bounded number of vertices can have the homotopy type of any primary semialgebraic set.

Keywords Convexity · Oriented matroid · Chirotope · Realization space

Mathematics Subject Classification 52A10 · 52C40 · 52C30 · 52C35

Editor in Charge: Günter M. Ziegler

Michael Gene Dobbins
mdobbins@binghamton.edu

Andreas Holmsen
andreash@kaist.edu

Alfredo Hubard
alfredo.hubard@u-pem.fr

- ¹ Department of Mathematical Sciences, Binghamton University, Binghamton, NY 13902-6000, USA
- ² Department of Mathematical Sciences, KAIST, 291 Daehak-ro Yuseong-gu, Taejeon 34141, South Korea
- ³ Université Paris-Est Marne-la-Vallée, 77454 Marne-la-Vallée Cedex 2, France

1 Introduction

We introduce a generalization of order type that provides a framework to study planar arrangements of convex sets and their convex dependencies. The notion we introduce is closely related to wiring diagrams [8] or primitive sorting networks [20]. It is also related to double pseudoline arrangements introduced by Pocchiola and Habert [16] and double allowable sequences introduced by Goodman and Pollack [12]. These related notions have applications in the study of visibility, transversal, and separation properties of convex sets [3, 15, 18, 24, 25]. Combinatorial type, the generalization of order type studied here, was fundamental to the authors' work on generalizations of the Erdős–Szekeres theorem to arrangements of convex sets in the plane [4, 5]. In this paper, we address the relevant realizability questions.

Two indexed point sets $P = \{p_1, \dots, p_n\}$ and $Q = \{q_1, \dots, q_n\}$ in the plane are said to have the same order type when for every triple (i, j, k) , the orientation of the triples p_i, p_j, p_k and q_i, q_j, q_k coincides. Equivalently, a point set P corresponds to a dual family P^* of oriented great circles in the sphere by projective duality, and point sets P and Q have the same order type when the families P^* and Q^* subdivide the sphere in the same way. That is, when there is a self-homeomorphism of the sphere that sends each cell of P^* to a corresponding cell of Q^* and preserves orientations.

More generally, we say that a sign function $\chi : \mathcal{L}^3 \rightarrow \{+, 0, -\}$ is an *order type* when it satisfies the axioms of rank 3 acyclic chirotopes [2, p. 126], [20, Chap. 10]. Specifically, χ is alternating, satisfies the Grassman–Plücker Relations, is acyclic (a restatement of Radon's Partition Theorem in terms of orientations), and is not identically zero. Order types that satisfy $\chi(i, j, k) \neq 0$ for any i, j, k distinct are called *simple*, and are equivalent to *uniform* rank 3 acyclic chirotopes and to Donald Knuth's CC-systems. Many geometric properties of point sets (such as convex dependency, transversal, and separation properties) depend solely on the order type of the points and not on the actual coordinates, and these axioms have been used to prove theorems and design algorithms involving such properties [1, 2, 20, 24, 25].

Not every order type can be realized by points in the plane, or dually by oriented great circles in the sphere. However, by the Topological Representation Theorem, any order type can be realized by a pseudocircle arrangement [7]. That is in the case of simple order types, by a family of simple oriented closed curves on the sphere such that each pair of curves intersect at exactly two points, any other curve separates these two points, and some point on the sphere is to the left of every curve. We may alternatively define a simple order type to be any subdivision of the sphere by such a pseudocircle arrangement $\mathcal{S} = \{S_1, \dots, S_n\}$. The orientations $\chi(i, j, k)$ can be recovered from this subdivision by the order S_i, S_j, S_k appear on the boundary of the region to the left of all three pseudocircles.

Like simple order types, combinatorial types are finite combinatorial objects that can be associated to families of geometric objects and serve as an invariant for properties that can be defined from their dual objects. Namely, they are associated to arrangements of convex bodies that satisfy certain genericity conditions. Combinatorial types will be defined precisely in Sect. 2, but for now we describe their invariant properties in terms of a duality for convex bodies that is analogous to projective duality for points in the plane. The *dual support curve* A^* of a convex body A in the plane, is the graph of

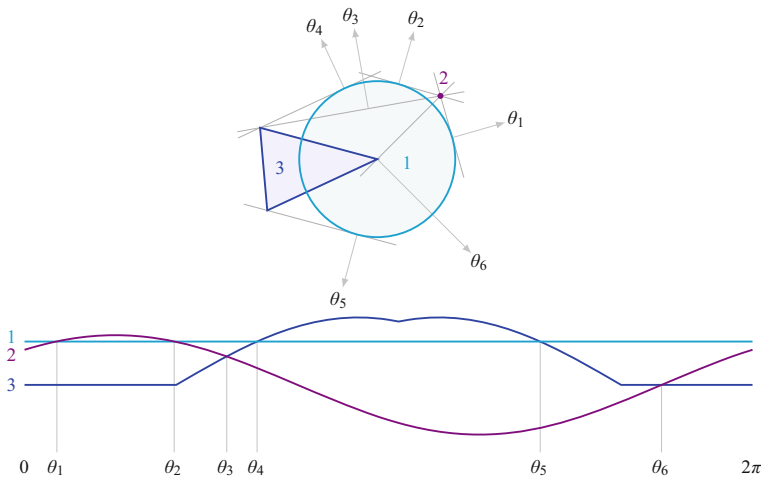


Fig. 1 Top: an arrangement \mathcal{A} and its common supporting tangents. Bottom: its dual support system \mathcal{A}^*

its support function $h_A : \mathbb{S}^1 \rightarrow \mathbb{R}^1, h_A(\theta) := \max_{p \in A} \langle \theta, p \rangle$ on the cylinder $\mathbb{S}^1 \times \mathbb{R}^1$, where \mathbb{S}^1 is the unit circle and $\langle \cdot, \cdot \rangle$ is the standard inner product; see Fig. 1. In this way, every arrangement $\mathcal{A} = \{A_1, \dots, A_n\}$ corresponds to the dual support system $\mathcal{A}^* = \{A_1^*, \dots, A_n^*\}$ of curves on the cylinder $\mathbb{S}^1 \times \mathbb{R}^1$ given by the graphs of the functions $\{h_{A_1}, \dots, h_{A_n}\}$. In the other direction, not all functions $h : \mathbb{S}^1 \rightarrow \mathbb{R}^1$ are support functions, but we do have the following sufficient conditions due to Blaschke [13, Lem. 2.2.3].

Proposition 1.1 *If $h : \mathbb{S}^1 \rightarrow \mathbb{R}^1$ is C^2 -smooth and $h + h'' > 0$, then h is the support function of a convex planar curve with curvature bounded by $\frac{1}{h+h''}$. Hence, by adding a sufficiently large constant to a finite family of smooth functions, we can ensure the family is the dual support system of an arrangement of convex bodies.*

Essentially, the combinatorial type of an arrangement of bodies $\text{ct}(\mathcal{A})$ encodes the subdivision of the cylinder $\mathbb{S}^1 \times \mathbb{R}^1$ by the dual support curves \mathcal{A}^* . The definition of combinatorial type is more technical than its intuitive significance so we postpone a detailed definition to Sect. 2. The advantage of the technical definition is that it is completely combinatorial and does not make reference to geometric objects (neither arrangements of convex bodies nor systems of curves), but only to equivalence classes of certain sequences of permutations (see 2.2). The following theorem gives the main invariant property of combinatorial types. One can take the equivalence relation on arrangements provided by this theorem as an alternative topological definition of combinatorial type.

Theorem 1.2 *Two arrangements of convex bodies \mathcal{A} and \mathcal{B} have the same combinatorial type if and only if their dual systems \mathcal{A}^* and \mathcal{B}^* are related by a self-homeomorphism of the cylinder that preserves orientation and $+\infty$.*

Here, we say that a self-homeomorphism $\phi : \mathbb{S}^1 \times \mathbb{R}^1 \rightarrow \mathbb{S}^1 \times \mathbb{R}^1$ preserves $+\infty$ when for y sufficiently large the second coordinate of $\phi(\theta, y)$ is positive for all θ .




Fig. 2 Three generic points and their support curves

Equivalently, if $C \subset \mathbb{S}^1 \times \mathbb{R}^1$ is an oriented curve that wraps once around the cylinder in the counter-clockwise direction, then so is its image $\phi(C)$.

In the case of points, the duality that we defined through support functions is the usual projective duality renormalized to be on the cylinder. Consequently, two generic point sets have the same order type if and only if they have the same combinatorial type. Specifically, a point in the plane can be represented in homogeneous coordinates by a line in \mathbb{R}^3 , and its dual support curve is the intersection of the orthogonal complement of this line with the cylinder embedded in \mathbb{R}^3 . The same relationship holds between a body in the plane represented by a cone in \mathbb{R}^3 and the body's dual support curve represented by its polar cone.

1.1 Realizing Order Types

Although combinatorial types of arrangements are more general than simple order types, we associate an order type to the class of arrangements defined by the following local condition on triples of bodies. We say a triple of bodies is *orientable* when it has the combinatorial type of three generic points (Fig. 2), and we say an arrangement is orientable when it consists of at least three bodies and every triple is orientable. In the case of an orientable arrangement \mathcal{A} , every triple $\{A_i, A_j, A_k\} \subset \mathcal{A}$ contributes a single connected arc to the boundary of $\text{conv}(A_i, A_j, A_k)$, and we define the *orientation* of an ordered triple (A_i, A_j, A_k) to be positive when the bodies appear counter-clockwise in the given order on the boundary, and to be negative otherwise. Grünbaum implicitly observes that the cyclic orderings of the triples of an orientable arrangement form an order type in his discussion on planar arrangements of simple closed curves [14, Sect. 3.3].

Most order types cannot be realized by points, and in fact, it is NP-hard to decide if a given order type can [30]. Having a notion of combinatorial type allows us to approach questions regarding realizability by bodies rather than points [17]. The smallest *non-realizable* order type is the Non-Pappus Configuration, a configuration of nine elements that violates Pappus's Theorem [21, 29]. Pach and Tóth showed that the Non-Pappus Configuration can be realized by an arrangement of segments in the plane [26]. Figure 3 shows a non-realizable order type that can be realized by triangles, Goodman and Pollack's "Bad Pentagon" [9], and the authors conjecture that this order type cannot be realized by segments. In contrast to point sets, we show that any order type, and in fact any combinatorial type, can be realized by an arrangement of bodies.

Theorem 1.3 *The orientations of the triples of any orientable arrangement is a simple order type. Two orientable arrangements have the same order type if and only if they*




Fig. 3 Two realizations of the Bad Pentagon. *Left*: a realization in a topological plane [9]. *Right*: a realization by convex sets in the Euclidean plane

have the same combinatorial type. And, every simple order type can be realized by an orientable arrangement.

If we bound the complexity of the bodies, then some simple order types can no longer be realized. Indeed, we show that simple order types can always be realized by k -gons, but may require k to be arbitrarily large.

Theorem 1.4 *Let k_n be the smallest integer for which every simple order type on n elements can be realized by an arrangement of k_n -gons. There are constants $c_1, c_2 > 0$ such that*

$$c_1 \frac{n}{\log n} \leq k_n \leq c_2 n^2.$$

1.2 Realization Spaces

An old conjecture of Ringel claimed that given two point sets with the same order type, one point set can be continuously deformed to the other while maintaining the order type [29]. This naturally leads to the study of *realization spaces* of order types, the set of all families of points with a fixed order type modulo projectivities. The conjecture can then be restated as, any non-empty realization space is connected. Ringel’s conjecture was disproved in the early eighties, and the strongest result in this direction is Mnev’s Universality Theorem [22,23], which states that for any primary semialgebraic set \mathcal{Z} , there exists an order type whose realization space is homotopy equivalent to \mathcal{Z} . Recall that a primary semialgebraic set is the set of common solutions to some finite list of polynomial equations and strict inequalities in several real variables. This has led to a growing body of work [2,19,27,28,31,32].

The main objective of this paper is to extend the study of realization spaces to arrangements of bodies of a fixed combinatorial type and exhibit a trade-off between

the combinatorial complexity of the bodies and the topological complexity of their realization space. The first indication of this trade-off may be observed from Theorems 1.3 and 1.4, which imply that for general convex bodies the realization space of any simple order type is non-empty, but this fails for k -gons. We prove two contrasting results. First, we show in Theorem 2.4 that Ringel’s intuition is correct in this generalized context: the realization space of any combinatorial type satisfying some mild assumptions is contractible; that is, it is non-empty and has no holes of any dimension. In particular, the set of arrangements (modulo planar rotations) of convex bodies realizing any fixed simple order type is contractible, and therefore connected. Second, we show in Theorem 2.5 that if the bodies are restricted to polygons with at most k vertices, then Mnev’s Theorem generalizes.¹ Specifically, we show that for every k and every primary semialgebraic set \mathcal{Z} there is a combinatorial type whose k -gon realization space is homotopy equivalent to \mathcal{Z} .

The proof of Theorem 2.4 provides an explicit deformation retraction to a standard arrangement, which will be defined for each combinatorial type. The proof of Theorem 2.5 depends on Mnev’s Theorem. Specifically, the proof uses a reduction from the case of k -gons to the case of points by constructing an arrangement of k -gons where the only obstruction to realizability is that certain vertices must always have the same fixed order type.

1.3 Relationship to Double Pseudoline Arrangements

Pocchiola and Habert introduced an extension of chirotopes to arrangements of convex sets based on a similar notion of duality to what is presented here, called double pseudoline arrangements [16]. The essential difference is that the dual double pseudoline of a convex set is defined as the quotient of the dual support curve by the \mathbb{Z}_2 action on the cylinder $(\theta, y) \sim (-\theta, -y)$. Instead of a curve that wraps monotonically once around the cylinder, the dual double pseudoline is a curve that wraps monotonically twice around the Möbius strip. This leads to an extended notion of chirotopes that provides information about arrangements of convex sets which combinatorial types do not distinguish, such as disjointness and visibility. On the other hand, combinatorial types distinguish convex position of subarrangements and are simpler in certain respects that are crucial to the analysis in [4, 5] and the results of this paper.

1.4 Organization of the Paper

Section 2 gives definitions, states the main theorems of the paper, and clarifies some issues that were treated vaguely in the introduction. Section 3 deals with realizing order types, Theorems 1.3 and 1.4; this section conveys the overall theme of the paper, but in a less technical setting, and involves ideas similar to what will be used in the rest of the paper. Section 4 proves contractibility, Theorem 2.4. Section 5 proves the topological invariance of combinatorial type, Theorem 1.2. While Theorem 1.2 is a more

¹ Note that Mnev’s Theorem is more specific as it deals with stable equivalence.

fundamental result, it appears later in the text as it depends crucially on Theorem 2.4. Section 6 gives the universality construction for convex k -gons, Theorem 2.5. Finally, Sect. 7 ends with some remarks and open problems.

2 Preliminaries and Main Theorems

2.1 Genericity Assumptions

A *common supporting tangent* of a pair of bodies is a *directed line* tangent to each body such that both bodies are on its *left* side. In the dual, this corresponds to an intersection between two support curves. We say that a pair of bodies intersect *transversally* when no point of intersection is contained in a common supporting tangent. In the dual this corresponds to a pair of curves in the cylinder that cross at each point of intersection; that is, for a pair of curves that are respectively the graphs of functions f_1, f_2 , the function $f_1 - f_2$ has only isolated zeros and changes sign at each zero. By an arrangement we mean a *finite indexed* non-empty collection of compact convex sets, which we call *bodies*, that satisfy following genericity conditions:

- Each pair of bodies intersect transversally.
- No three bodies share a common supporting tangent.
- There are finitely many common supporting tangents.

Analogously, by a *system* we mean a finite indexed collection of closed curves in the cylinder $\mathbb{S}^1 \times \mathbb{R}^1$ that are monotonic in the first coordinate and satisfy following genericity conditions:

- Each pair of curves cross at each point of intersection.
- No three curves share a common point of intersection.
- There are finitely many crossings.

2.2 Combinatorial Type

Let \mathfrak{S}_m be the symmetric group on m elements and $[m] = \{1, \dots, m\}$. Given $i \in [m - 1]$, the *adjacent transposition* $\tau_i \in \mathfrak{S}_m$ is the permutation interchanging the i 'th and $i + 1$ 'st entries,

$$\tau_i(x_1, \dots, x_m) = (x_1, \dots, x_{i-1}, x_{i+1}, x_i, x_{i+2}, \dots, x_m).$$

Let $H(\tau_i) = i$ denote the height of an adjacent transposition. A *swap sequence* $\sigma : [N] \rightarrow \mathfrak{S}_m$ is any sequence of adjacent transpositions such that $\sigma_N \circ \dots \circ \sigma_1$ is the identity permutation.

Fix an index set \mathcal{L} of size n . A *swap pair* (ρ, σ) on \mathcal{L} is a bijection $\rho : [n] \rightarrow \mathcal{L}$ together with a swap sequence $\sigma : [N] \rightarrow \mathfrak{S}_n$. We define an equivalence relation ($\overset{\text{swap}}{\sim}$) on swap pairs as follows. Let $(\rho', \sigma') \overset{\text{swap}}{\sim} (\rho, \sigma)$ when (ρ', σ') can be obtained from (ρ, σ) by performing any sequence of the following two *elementary operations*:




Fig. 4 A system with its swap pair (ρ, σ) and its incidence sequence $\overline{\rho\sigma}$. Note that systems are drawn as viewed from outside the cylinder, so counter-clockwise is to the *right*

– a *cyclic shift*

$$\rho' = \tau_{\sigma_1}(\rho), \quad \sigma'_i = \sigma_{(i+1 \bmod N)},$$

– an *independent transposition*

$$\rho' = \rho, \quad \sigma' = \tau_i(\sigma) \quad \text{provided} \quad |H(\sigma_i) - H(\sigma_{i+1})| > 1.$$

A *combinatorial type* Ω on \mathcal{L} is the equivalence class $\Omega = \{(\rho', \sigma') : (\rho', \sigma') \stackrel{\text{swap}}{\sim} (\rho, \sigma)\}$ of a swap pair (ρ, σ) .

To define the combinatorial type of a system \mathcal{S} , we order the crossings of \mathcal{S} lexicographically in $\mathbb{S}^1 \times \mathbb{R}^1$ where \mathbb{S}^1 is ordered according to the standard parametrization by the half-open interval $(0, 2\pi]$. Let ρ be the order of the indices of each curve from bottom to top before the first crossing of the system. Let σ be the swap sequence corresponding to each crossing. That is, let $H(\sigma_i)$ be 1 plus the number of curves below the i 'th crossing of \mathcal{S} . Observe that the sequence $\sigma_i \circ \dots \circ \sigma_1(\rho)$ for $i = 0, 1, \dots, N$ records the order of the curves in a sweep of the cylinder. The *combinatorial type* $\text{ct}(\mathcal{S})$ of a system \mathcal{S} is the equivalence class of its swap pair (ρ, σ) . The combinatorial type of an arrangement \mathcal{A} is that of its dual system, and by slight abuse of notation, we write $\text{ct}(\mathcal{A}) = \text{ct}(\mathcal{A}^*)$ (Fig. 4).

The *incidence sequence* $\overline{\rho\sigma} : [N] \rightarrow \mathcal{L}^2$ of a swap pair (ρ, σ) records the ordered pair of indices transposed by the action of each swap,

$$\overline{\rho\sigma}_i = (x_{H(\sigma_i)+1}, x_{H(\sigma_i)}) \quad \text{where} \quad x = \sigma_{i-1} \circ \dots \circ \sigma_1(\rho).$$

Note that the incidence sequence of equivalent swap pairs have the same multi-set of entries.

The *layers* of a system are the connected components of the union of curves of the system. Analogously, the layers of a combinatorial type are the connected components of the graph on \mathcal{L} defined by its incidence sequence. The *depth* of a combinatorial type is the number of layers excluding isolated vertices, and the depth 1 case is called *non-layered*.

Remark 2.1 Orientable combinatorial types are always non-layered.

2.3 Realization Spaces

The *full realization space* $\mathcal{R}(\Omega)$ of a combinatorial type Ω is defined by

$$\mathcal{R}(\Omega) := \{A \in \mathcal{U}_\Omega : \text{ct}(A) = \Omega\},$$

where \mathcal{U}_Ω is the set of all arrangements of bodies indexed by \mathcal{L} . The Hausdorff metric d_H on compact subsets of \mathbb{R}^2 induces a metric on $\mathcal{R}(\Omega)$ by taking the maximum distance between bodies having the same index. That is, for $A = \{A_i\}_{i \in \mathcal{L}}$ and $B = \{B_i\}_{i \in \mathcal{L}}$,

$$d(A, B) = \max_{i \in \mathcal{L}} d_H(A_i, B_i).$$

Remark 2.2 The map that takes a convex body to its support function is an isometry from the space of convex bodies with the Hausdorff metric to the space of support functions on \mathbb{S}^1 with the supremum metric.

Depending on context, it may be convenient to regard realizations of a fixed combinatorial type as “the same” when they are related by a projective transformation. Let $A \overset{\text{proj}}{\sim} B$ when they are related by an admissible projectivity; that is, an invertible projective transformation π such that $\pi(A_i) = B_i$ for all $i \in \mathcal{L}$ and π is bounded and preserves orientation on the convex hull of $\bigcup A$. The (*projective*) *realization space*, which we may simply call the “realization space”, is the quotient space

$$\tilde{\mathcal{R}}(\Omega) := \mathcal{R}(\Omega) / \overset{\text{proj}}{\sim}.$$

By a *k-gon* we mean a convex polygon with *at most* k vertices. The *full k-gon realization space* is given by

$$\mathcal{R}_k(\Omega) := \{A \in \mathcal{R}(\Omega) : \forall i \in \mathcal{L}, A_i \text{ is a } k\text{-gon}\}.$$

Similarly, we have the (*projective*) *k-gon realization space* $\tilde{\mathcal{R}}_k(\Omega) := \mathcal{R}_k(\Omega) / \overset{\text{proj}}{\sim}$.

Remark 2.3 By considering the direction of some line ℓ passing through an arrangement A , the admissible projectivities of A factor into a contractible set of projective transformations fixing the direction of ℓ and rotations of the plane $SO(2)$. As such, $\mathcal{R}(\kappa)$ is homotopic to $\tilde{\mathcal{R}}(\kappa) \times \mathbb{S}^1$.

We now state the main theorems on realization space of arrangements of convex bodies.

Theorem 2.4 *If Ω is a non-layered combinatorial type (in particular, if Ω corresponds to a simple order types), then its realization space $\mathcal{R}(\Omega)$ is contractible. Moreover, if Ω has depth $d > 1$, then $\tilde{\mathcal{R}}(\Omega)$ is homotopic to a product of $d - 1$ circles.*

Theorem 2.5 *For every primary semialgebraic set \mathcal{Z} and every positive integer k , there exists a combinatorial type Ω such that its k -gon realization space $\tilde{\mathcal{R}}_k(\Omega)$ is homotopic to \mathcal{Z} .*

3 Realizing Order Types

3.1 Orientability and Order Types

We start by showing a bijective correspondence between orientable arrangements and simple order types.

Proof of Theorem 1.3 The Topological Representation Theorem gives a bijective correspondence between simple order types χ and equivalence classes of uniform acyclic pseudocircle arrangements on the sphere \mathbb{S}^2 . We extend this bijection to equivalence classes of orientable arrangements by mapping the dual support systems to pseudocircle arrangements. Note that a positive orientation on the sphere is chosen; if we were to forget this orientation, then each equivalence class of pseudocircle arrangements would correspond to the pair $\{\chi, -\chi\}$.

Let $\mathcal{A} = \{A_1, \dots, A_n\}$ be an orientable arrangement. Let $\phi: \mathbb{S}^1 \times \mathbb{R}^1 \rightarrow \mathbb{S}^2$ be the compactification of the cylinder defined by adding one point $p_{+\infty}$ at $+\infty$ and another point $p_{-\infty}$ at $-\infty$ where $(\theta, y) \rightarrow p_{\pm\infty} \in \mathbb{S}^2$ as $y \rightarrow \pm\infty$. The image of the dual support system $\phi(\mathcal{A}^*) = \{\phi(A_1^*), \dots, \phi(A_n^*)\}$ is now a uniform acyclic pseudocircle arrangement; $\phi(A_i^*)$ is a uniform pseudocircle arrangement by the genericity assumptions and $\phi(\mathcal{A}^*)$ is acyclic since the region of cylinder that is above every curve of \mathcal{A}^* is now to the left of every curve of $\phi(\mathcal{A}^*)$.

For the other direction, let $\mathcal{S} = \{S_1, \dots, S_n\}$ be a uniform acyclic pseudocircle arrangement, and let $p_{+\infty}$ be a point to the left of each curve and $p_{-\infty}$ be a point to the right of each curve. Pseudocircle arrangements can be swept, meaning a path γ_t from the point $p_{-\infty}$ to the point $p_{+\infty}$ can be continuously deformed while keeping its end-points fixed such that it passes through all other points on the sphere exactly once returning to its original position and it always intersects each pseudocircle at exactly one point [10, Thm. 2.9]. This defines a homeomorphism from $\mathbb{S}^2 \setminus \{p_{+\infty}, p_{-\infty}\}$ to the cylinder such that the image of each pseudocircle is the graph of a function $f_i: \mathbb{S}^1 \rightarrow \mathbb{R}^1$. Each of the f_i can then be approximated by a smooth function h_i while preserving their order above each point on the circle,

$$\forall \theta \in \mathbb{S}^1, f_{i_1}(\theta) \leq \dots \leq f_{i_n}(\theta) \Leftrightarrow h_{i_1}(\theta) \leq \dots \leq h_{i_n}(\theta).$$

By Proposition 1.1, we may now assume the h_i are support functions of an arrangement \mathcal{A} . With this, $\phi(\mathcal{A}^*)$ is a pseudocircle arrangement of the same order type as \mathcal{S} .

An orientable arrangement \mathcal{A} now has a simple order type given by $\phi(\mathcal{A}^*)$, and the orientations of a triple A_i, A_j, A_k are given by the order $\phi(A_i^*), \phi(A_j^*), \phi(A_k^*)$ appear on the boundary of the region to the left of each curve. By Theorem 1.2 two orientable arrangements have the same combinatorial type if and only if they have the same order type. And, every simple order type can be realized by an orientable arrangement. \square

3.2 Counting Arrangements of Polygons

Here we give an upper bound on the number of combinatorial types that can be realized by k -gons, and bounds on the value of k_n needed to realize all simple order types on $[n]$ by k_n -gons. Our results are based on the following.

Theorem 3.1 (Goodman and Pollack [11]) *Let $t(n)$ denote the number of distinct order types (not necessarily simple) on $[n]$ that can be realized by points. For some constants c_1, c_2 ,*

$$2^{4n \log n + c_1 n} \leq t(n) \leq 2^{4n \log n + c_2 n}.$$

Theorem 3.2 (Felsner and Valtr [6]) *Let $t_{\text{ot}}(n)$ denote the number of distinct simple order types on $[n]$.*

$$2^{0.188n^2} \leq t_{\text{ot}}(n) \leq 2^{0.6571n^2}.$$

Theorem 3.3 *Let $t_k(n)$ denote the number of distinct combinatorial types on $[n]$ that can be realized by k -gons. For some constant c ,*

$$t_k(n) \leq 2^{ckn(\log n + \log k)}.$$

Proof Since the combinatorial type of an arrangement of n k -gons is determined by the order type of the kn vertices of the arrangement, we have the inequality, $t_k(n) \leq t(kn)$. Specifically, for every combinatorial type Ω that can be realized by k -gons, fix a realization by k -gons $\mathcal{A} = \{A^1, \dots, A^n\}$. Then, for each k -gon A^s , fix a labeling of vertices so that A^s is the convex hull of a point set $\{p_1^s, \dots, p_k^s\}$. Now associate the order type of the points $\{p_1^1, \dots, p_k^n\}$ to Ω . In this way, we define an injective map from the set of combinatorial types on $[n]$ that can be realized by k -gons to the set of order types on $\left\{\binom{1}{1}, \dots, \binom{n}{1}\right\}$ that can be realized by points. The upper bound on the number of combinatorial types now follows from Theorem 3.1. \square

Theorem 3.4 *A combinatorial type Ω with N common supporting tangents can be realized by N -gons.*

Proof Recall that the wiring digram of a sequence of adjacent transpositions σ is a family of polygonal paths that cross according to σ [8]. Informally, we will define an arrangement of N -gons by “wrapping” a wiring diagram of σ around a large circle.

Let $(\rho, \sigma) \in \Omega$ and $\pi_t = \sigma_t \circ \dots \circ \sigma_1$. Choose r sufficiently large and let $p_{i,t} = (2\pi t/N; r + \pi_t^{-1}(i))$ in polar coordinates, so that each point set $P_i = \{p_{i,1}, \dots, p_{i,N}\}$ for $i \in [n]$ is in convex position. For this, $r \geq (1 - \cos(2\pi/N))^{-1}$ would suffice. Let $A_i = \text{conv}(P_i)$. Since the support functions h_i of the A_i now satisfy

$$h_{\rho\pi_t^{-1}(1)}(\theta_t) < \dots < h_{\rho\pi_t^{-1}(n)}(\theta_t)$$

for $\theta_t = 2\pi t/N$ with only once crossing between θ_t and θ_{t+1} , the swap pair of the arrangement $\mathcal{A} = \{A_1, \dots, A_n\}$ is (ρ, σ) . Hence, \mathcal{A} is an arrangement of N -gons with combinatorial type Ω . \square

Proof of Theorem 1.4 The upper bound follows immediately from Theorem 3.4 and the observation that an orientable arrangement on $[n]$ always has $N = n(n - 1)$ common supporting tangents. For the lower bound, observe from Theorems 3.2 and 3.3 that for any fixed k , the number of simple order types grows faster than the number of combinatorial types that can be realized by k -gons. Specifically, there is some constant c_1 such that if $k \leq c_1 \frac{n}{\log n}$ then $t_k(n) < t_{\text{ot}}(n)$, which implies that some simple order type on $[n]$ cannot be realized by k -gons. \square

4 Contractibility

To show contractibility, we construct a standard arrangement of convex bodies for each combinatorial type by defining its dual support system. We then show that the full realization space $\mathcal{R}(\Omega)$ is equivariantly homotopic to a circle \mathbb{S}^1 by defining a deformation retraction to the subspace of rotated copies of the standard arrangement. By *equivariantly* homotopic we mean that the corresponding homotopy maps commute with $SO(2)$. We then pass to the (projective) realization space $\widetilde{\mathcal{R}}(\Omega)$ by identifying arrangements related by admissible projectivities. Since rotations are admissible projectivities, this defines a deformation retraction from $\widetilde{\mathcal{R}}(\Omega)$ to a point.

The deformation retraction from $\mathcal{R}(\Omega)$ to a circle will proceed in two steps; see Fig. 7. First in Lemma 4.2, we deform the support system of a given arrangement to a system of the same combinatorial type that depends only on the (angular) position of each crossing. We can then consider just the positions of the crossings and ignore the rest of the geometry of the system. Second in Lemma 4.5, we move the crossings to a set of standard positions that depend only on: the given combinatorial type and the position of a certain crossing that we fix. The set of possible standard systems we get in the end is parametrized by the position of this fixed crossing, which defines an embedding of the circle in $\mathcal{R}(\Omega)$. The first deformation retraction depends on the following remark.

Remark 4.1 For any pair of convex bodies A and B , $(A + B)^* = A^* + B^*$ with Minkowski addition on the left and addition of the support functions defining the curves on the right. And, for $t \geq 0$, $(tA)^* = t(A^*)$. Hence, the set of all support functions is a convex cone. That is, if h_1 and h_2 are support functions, then so is $t_1 h_1 + t_2 h_2$ for $t_i \geq 0$. Note however, that the set of dual support systems of a fixed combinatorial type is not a convex set.

4.1 Support Configurations

The *support configuration* of an arrangement \mathcal{A} indexed by \mathcal{L} is a labeled vector configuration $\text{sc}(\mathcal{A}) \subset \mathcal{L}^2 \times \mathbb{S}^1$ which contains a triple (i, j, θ) if bodies A_i, A_j have a common supporting tangent line ℓ that first meets A_i and then meets A_j and has outward normal vector θ . We say labels $(i, j), (i', j')$ are *disjoint* when $\{i, j\} \cap \{i', j'\} = \emptyset$. Note that a unit vector θ may appear in multiple elements of $\text{sc}(\mathcal{A})$ with disjoint labels.




Fig. 5 From top to bottom: (1) The dual support system of an arrangement \mathcal{A} of combinatorial type Ω ; compare with Fig. 1. (2) The system $\alpha^* \text{sc}(\mathcal{A})$ of Lemma 4.2, which depends only on the angular positions of the crossings of \mathcal{A}^* . (3) The system $\alpha^* W(\Omega, \theta, \delta)$ obtained by fixing the marked crossing of $\alpha^* \text{sc}(\mathcal{A})$, and rotating all other crossings clockwise. (4) The standard system $\alpha^* W(\Omega, \theta)$ of Lemma 4.5, which depends only on the angular position θ of the fixed crossing and Ω

Dually, $\text{sc}(\mathcal{A})$ corresponds to the crossings of \mathcal{A}^* . Specifically, $(i, j, \theta) \in \text{sc}(\mathcal{A})$ when curves A_i^* and A_j^* cross at θ with A_i^* crossing downward and A_j^* crossing upward. That is, the respective support functions f_i, f_j of A_i, A_j are equal at θ and $f_j - f_i$ is increasing at θ (Fig. 5).

Observe that the support configuration of an arrangement determines the combinatorial type of that arrangement. For a given combinatorial type Ω , we will define its *support configuration space* $\mathcal{V}(\Omega)$, which will turn out to be the set of support configurations of all arrangements realizing Ω . We first define the set of labeled configurations $\mathcal{V}(\rho, \sigma)$ corresponding to a given swap pair (ρ, σ) . Recall that $\overline{\rho\sigma}$ records the ordered pairs of indices transposed by σ acting sequentially on ρ . Observe that if (ρ, σ) is the swap pair of a system, then $\overline{\rho\sigma}_i$ for $i \in [N]$ is the labeling of the i -th crossing of the system. Recall also that we order \mathbb{S}^1 by the parametrization by $(0, 2\pi]$. Let

$$\mathcal{V}(\rho, \sigma) := \left\{ \{(\overline{\rho\sigma}_i, \theta_i) : i \in [N]\} : \begin{array}{l} \theta_i \in \mathbb{S}^1, \theta_i \leq \theta_{i+1}, \\ \theta_i = \theta_{i+1} \Rightarrow |H(\sigma_i) - H(\sigma_{i+1})| > 1 \end{array} \right\},$$

$$\mathcal{V}(\Omega) := \bigcup_{(\rho, \sigma) \in \Omega} \mathcal{V}(\rho, \sigma).$$

Note that a vector $\theta \in \mathbb{S}^1$ might appear multiple times in $\mathcal{V}(\Omega)$ with different labels provided the pairs of indices in the labels are disjoint.

We define a metric on $\mathcal{V}(\Omega)$ as follows. For a given support configuration X and a given ordered pair of indices $(i, j) \in \binom{[N]}{2}$, let $X_{(i,j)} := \{\theta \in \mathbb{S}^1 : (i, j, \theta) \in X\}$. For two support configurations, $X, Y \subset \mathcal{L}^2 \times \mathbb{S}^1$,

$$d(X, Y) = \max_{(i,j) \in \binom{\mathcal{L}}{2}} d_H(X_{(i,j)}, Y_{(i,j)}),$$

where the distance between two direction vectors is given by their angle and d_H is the corresponding Hausdorff metric on sets.

Lemma 4.2 *For any combinatorial type Ω , the full realization space $\mathcal{R}(\Omega)$ is non-empty and equivariantly homotopic to the support configuration space $\mathcal{V}(\Omega)$.*

Proof For $\mathcal{A} \in \mathcal{R}(\Omega)$ with swap pair (ρ, σ) , we have $\text{sc}(\mathcal{A}) \in \mathcal{V}(\rho, \sigma) \subset \mathcal{V}(\Omega)$, so assigning each arrangement to its support configuration defines a map $\text{sc}: \mathcal{R}(\Omega) \rightarrow \mathcal{V}(\Omega)$, which will be one direction of the homotopy equivalence.

For the other direction, we define an embedding $\alpha: \mathcal{V}(\Omega) \rightarrow \mathcal{R}(\Omega)$. For each labeled configuration $V \in \mathcal{V}(\Omega)$, we construct a system of curves $\alpha^*(V) = \{A_i^*: i \in \mathcal{L}\}$ where $A_i^* = f_i(\mathbb{S}^1)$, $f_i: \mathbb{S}^1 \rightarrow \mathbb{R}^1$, and show that $\alpha^*(V)$ is the dual support system of an arrangement $\alpha(V)$ that has support configuration V . The system $\alpha^*(V)$ that we construct may be regarded as a smooth analog of Goodman's wiring diagram [8].

Fix $V \in \mathcal{V}(\Omega)$, let $V_i \subset \mathbb{S}^1$ denote the vectors of V with labels involving i , and let δ be the minimum angular distance between any two vectors of V with non-disjoint labels. For $v = (i, j, \theta) \in V$ define the open arc $\Theta(v) := (\theta - \delta/2, \theta + \delta/2) \subset \mathbb{S}^1$. Now define f_i to be constant on the complement of the arcs $\Theta(V_i)$, and to smoothly increase or decrease by ± 1 symmetrically about θ in each arc $\Theta(v)$ according to the label on $v \in V_i$; that is, f_i increases on $\Theta(v)$ if $(j, i, \theta) \in V$ and decreases if $(i, j, \theta) \in V$ for some j .² We claim that each f_i is well defined up to an additive constant, and these constants can be chosen so that each pair f_i, f_j coincides on $V_i \cap V_j$ and nowhere else.

Since $\sigma_N \cdots \sigma_1$ is the identity permutation, each f_i increases the same number of times as it decreases on arcs in $\Theta(V_i)$, so traversing once around \mathbb{S}^1 results in no net change in the value of f_i , which is therefore well defined up to an additive constant.

Let c be some constant to be determined later. For $\xi \in \mathbb{S}^1$ ordered by $(0, 2\pi]$, let $j(\xi)$ be the largest integer for which $\theta_{j(\xi)} < \xi$ where θ_j is the angle of the j 'th vector of V . We will show that letting $f_i(\xi) = c + \sigma_{j(\xi)} \cdots \sigma_1 \rho(i)$ on the complement of $\Theta(V_i)$ satisfies our claim. For consecutive angles $\theta_a, \theta_b \in V_i$, the labels on the vectors of V in the interval $[\theta_a + \delta/2, \theta_b - \delta/2]$ do not involving i , so i remains fixed by the corresponding adjacent transpositions of the swap sequence $\sigma_{j(\theta_a + \delta/2) + 1}, \dots, \sigma_{j(\theta_b - \delta/2)}$, which implies f_i is constant on this interval. Since the f_i are defined to increase or decrease on $\Theta(V_i)$ according to the action of the swap sequence, this completely determines a smooth function f_i . Since the f_i increase or decrease symmetrically in each arc of $\Theta(V_i)$, each pair f_i, f_j coincide exactly on $V_i \cap V_j$.

To fix the constant c , let

$$\min_{(i,\theta) \in \mathcal{L} \times \mathbb{S}^1} (f_i(\theta) + f_i''(\theta)) = 1.$$

² The definition of f_i on $\Theta(v)$ is irrelevant as long as f_i is C^2 -smooth, monotonic, symmetric about θ , and varies continuously with respect to V . A cubic spline would suffice for this.




Fig. 6 *Left:* a system and its associated tableau of local sequences Λ . *Right:* the tableau $\text{bump}(\Lambda, \{a, b\})$ and a corresponding system

By Proposition 1.1, the system $\alpha^*(V)$ defined by the functions f_i is the dual support system of an arrangement $\alpha(V) \in \mathcal{R}(\Omega)$ that is uniquely and continuously determined by $V \in \mathcal{V}(\Omega)$, and $\text{sc}(\alpha(V)) = V$. This gives us a subspace $\alpha(\Omega) := \{\alpha(V) : V \in \mathcal{V}(\Omega)\} \subset \mathcal{R}(\Omega)$ that is homeomorphic to $\mathcal{V}(\Omega) = \text{sc}(\mathcal{R}(\Omega))$. For $\mathcal{A} \in \mathcal{R}(\Omega)$ define $\mathcal{A}_t := t\alpha(\text{sc}(\mathcal{A})) + (1 - t)\mathcal{A}$ for $0 \leq t \leq 1$ by Minkowski addition on each body of the arrangement. Since $\text{sc}(\mathcal{A}) = \text{sc}(\alpha(\text{sc}(\mathcal{A})))$ and, as we linearly interpolate between two systems with the same crossings, the crossings remain fixed by Remark 4.1, $\text{sc}(\mathcal{A}_t)$ is constant for all $t \in [0, 1]$. Thus, $\alpha(\Omega)$ is an equivariant deformation retract of $\mathcal{R}(\Omega)$. □

4.2 Local Sequences and Associated Tableau

We now introduce an encoding of the combinatorial type extending the local sequences of a point set. This encoding will be used in several proofs, including in later sections of the paper.

For a system \mathcal{S} , the *local sequence* $\lambda_i = (\lambda_{i,1}, \dots, \lambda_{i,n_i})$ of $i \in \mathcal{L}$ lists the indices of the curves that S_i crosses ordered by $(0, 2\pi]$; see Fig. 6. Analogously for a pair $(\rho, \sigma) \in \Omega$, the local sequence λ_i lists the indices $\lambda_{i,j}$ appearing together with i as part of a pair $(\lambda_{i,j}, i)$ or $(i, \lambda_{i,j})$ in the incidence sequence $\overline{\rho\sigma}$. The *associated tableau* Λ of (ρ, σ) is the tableau whose rows are the local sequences for (ρ, σ) ordered by ρ from bottom to top. We say Λ is a *tableau representation* of the combinatorial type Ω . The local sequences and associated tableau of a system or arrangement are that of its swap pair.

Like in the definition of combinatorial type, we define an equivalence relation on associated tableau. We say a pair $j, k \in \mathcal{L}$ are adjacent in a tableau Λ with rows λ_i when

$$\lambda_j = (k, \lambda_{j,2}, \dots, \lambda_{j,n_j}) \quad \text{and} \quad \lambda_k = (j, \lambda_{k,2}, \dots, \lambda_{k,n_k}).$$

In this case, *bumping* $\{j, k\}$ sends Λ to the tableau $\Lambda' = \text{bump}(\Lambda, \{j, k\})$ with rows

$$\lambda'_i = \begin{cases} (\lambda_{j,2}, \dots, \lambda_{j,n_j}, k), & i = j, \\ (\lambda_{k,2}, \dots, \lambda_{k,n_k}, j), & i = k, \\ \lambda_i, & i \neq j, k \end{cases}$$

and the order of rows j, k in the tableau is reversed.

Remark 4.3 A pair of tableaux represent the same combinatorial type if and only if one can be obtained from the other by a sequence of bumps, since cyclic shifts correspond to bumps and independent transpositions do not change the associated tableau.

For a pair of tableau Λ^1, Λ^2 with respective rows λ_i^1, λ_i^2 appearing in the same order ρ , let $\Lambda^1 \cdot \Lambda^2$ denote row-wise concatenation. That is, $\Lambda^1 \cdot \Lambda^2$ has rows

$$\lambda_i = (\lambda_{i,1}^1, \dots, \lambda_{i,n_{1,i}}^1, \lambda_{i,1}^2, \dots, \lambda_{i,n_{2,i}}^2).$$

The *periodicity* p of a tableau Λ is the largest p such that Λ is a p -fold concatenation of a tableau Λ' ,

$$\Lambda = \underbrace{\Lambda' \cdots \Lambda'}_{\times p}.$$

We call the tableau Λ' , the *period* of Λ . We say Λ is *non-periodic* when $p = 1$. The periodicity of a combinatorial type Ω is that of any tableau representing Ω , which is well defined by the following lemma.

Lemma 4.4 *The period of a tableau is unique, and all tableau representing a fixed combinatorial type have the same periodicity.*

Proof For uniqueness, just observe that row λ'_i of Λ' consists of the first n_i/p entries of row λ_i of Λ where n_i is the length of λ_i . Now consider a pair of tableau Λ_1 and $\Lambda_2 = \text{bump}(\Lambda_1, \{i, j\})$ representing the same Ω with periodicities p_1, p_2 respectively, and let Λ'_1 be the period of Λ_1 . Then Λ_2 is the p_1 -fold concatenation of $\Lambda'_2 = \text{bump}(\Lambda'_1, \{i, j\})$. Therefore, $p_2 \geq p_1$. Since any tableau representing Ω can be obtained from any other tableau representing Ω by a sequence of bumps, all tableaux have the same periodicity. \square

4.3 Standard Configurations (Non-periodic)

For non-periodic non-layered combinatorial types Ω , we will construct a standard set of labeled vector configurations $\mathcal{W}(\Omega) \subset \mathcal{V}(\Omega)$ parametrized by \mathbb{S}^1 . Here we arbitrarily choose a standard configuration similar to the “compressed form” given in [20, p. 31]. For layered combinatorial types, we can apply this construction independently to each layer. The periodic case will be dealt with in the next subsection.

We first construct a labeled vector configuration $W(\Omega, \theta, \delta)$ for $\theta \in \mathbb{S}^1$ and $\delta > 0$ sufficiently small as follows. Let Λ_{\min} be the lexicographically minimal tableau

representing Ω for which there exists exactly one adjacent pair. By Lemma 4.5 below, such a tableau always exists, provided Ω is non-layered. We will define a sequence of configurations W_t recursively starting from $t = 0$. To start, set $\Lambda_0 = \Lambda_{\min}$, $\xi_0 = \theta$, $W_0 = \emptyset$. Let $\{(i_{t,1}, j_{t,1}), \dots, (i_{t,m_t}, j_{t,m_t})\}$ be the set of all adjacent pairs in Λ_t ordered according to the rows of Λ_t . Let

$$W_{t+1} = W_t \cup \{(i_{t,1}, j_{t,1}, \xi_t), \dots, (i_{t,m_t}, j_{t,m_t}, \xi_t)\},$$

$\xi_{t+1} = \xi_t + \delta$, and let Λ_{t+1} be the tableau obtained from Λ_t by interchanging the corresponding pairs of rows and deleting the first entry from each of these rows. Note that this is similar to a bump except that initial entries are deleted instead of being moved to the end of their respective rows. Eventually, $\Lambda_T = \emptyset$ for some minimal T . Let $W(\Omega, \theta, \delta) = W_T$. Finally, let

$$\mathcal{W}(\Omega) = \{W(\Omega, \theta, 2\pi/N) : \theta \in \mathbb{S}^1\}.$$

Lemma 4.5 *For any non-layered combinatorial type Ω , $\mathcal{V}(\Omega)$ is equivariantly homotopic to $\mathcal{W}(\Omega) \simeq \mathbb{S}^1$.*

Proof (Ω non-periodic) For each $V \in \mathcal{V}(\Omega)$, we will inductively define a sequence of partitioned vector configurations $W_k \cup R_k \in \mathcal{V}(\Omega)$ starting from $k = 1$, and maps $\psi_k : \mathcal{V}(\Omega) \times [0, 1] \rightarrow \mathcal{V}(\Omega)$ that perform each step of the induction continuously. Together, these maps will define an equivariant deformation retraction from $\mathcal{V}(\Omega)$ to $\mathcal{W}(\Omega)$.

Let $\angle(y, x)$ denote the angular distance from x to y in the counter-clockwise direction, and as in the proof of Lemma 4.2, let δ be the minimum angular distance between vectors of V that have non-disjoint labels, so that for all $x, y \in V_i$ we have $0 < \delta \leq \angle(y, x) < 2\pi$. Fix a labeled vector $v \in V$, and set $W_1 = \{v\}$ and $R_1 = V \setminus \{v\}$. We will continuously deform V to a configuration $W(v)$. Later we will make a choice of $v = \tilde{v}$ that depends continuously on the initial configuration $V \in \mathcal{V}(\Omega)$.

At each step, rotate the sub-configuration R_k clockwise by a continuous rotation $\phi_{k,t} : [0, 1] \rightarrow SO(2)$, $\phi_{k,0} = \text{id}$, that decreases angular distances among the vectors of $W_k \cup \phi_{k,t}R_k$ with non-disjoint labels, until reaching the minimal rotation $\phi_k = \phi_{k,1}$ such that there is some $x \in R_k$ and $y \in W_k$ with non-disjoint labels and $\angle(\phi_k x, y) = \delta$. Let X be the set of all such $x \in R_k$, let $W_{k+1} = W_k \cup \phi_k X$ and $R_{k+1} = \phi_k(R_k \setminus X)$, and continue inductively until $R_K = \emptyset$. Since Ω is non-layered, such a pair x, y exists provided $R_k \neq \emptyset$, and since each step removes elements from R_k , the process terminates in a configuration $W_K = W_K(v)$ such that for every $x \in W_K \setminus \{v\}$ there exists a $y \in W_K$ such that $\angle(x, y) = \delta$ and x, y have non-disjoint labels; see Fig. 7. Let $\psi_k(V, t) = W_k \cup \phi_{k,t}R_k$ for $k \leq K$. Since vectors of $\psi_k(V, t)$ are only rotated through other vectors with disjoint labels, this process only changes the swap pair of a configuration by elementary operations, which implies $\text{ct}(\mathcal{A}_{\psi_k(V,t)}) = \text{ct}(\mathcal{A}_V)$. Finally, let $\psi_{K+1}(V, t)$ start from $\psi_{K+1}(V, 0) = W_K$ and continuously scale the angular distance of each vector from v by sending a labeled vector at u to u_t where $\angle(u, v) = r\delta$, $\angle(u_t, v) = tr\delta' + (1 - t)r\delta$, and $\delta' = 2\pi/N$. Let $W(v) = \psi_{K+1}(V, 1)$.

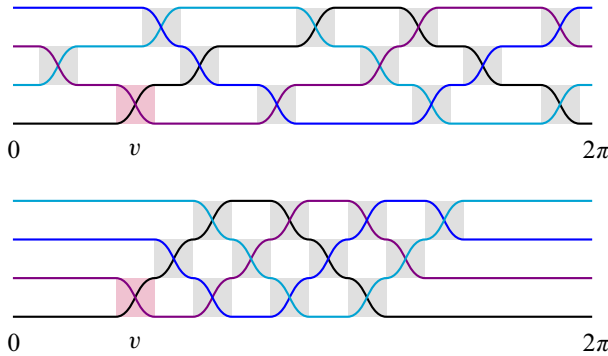


Fig. 7 Sending V (top) to its compressed form $W_K(v)$ (bottom)

We now define $\tilde{v} = \tilde{v}(V)$ in terms of the configurations $\{W(v) : v \in V\}$. If we rotate $W = W(v)$ by ϕ such that there are no elements of ϕW in the arc $(0, \phi v) \subset \mathbb{S}^1$, then the associated tableau Λ of ϕA_W will have exactly one adjacent pair, which is given by the support information of $v \in V$. Since Ω is *non-periodic*, there is a unique choice of $v = \tilde{v}$ such that $\Lambda = \Lambda_{\min}$. With this, our final vector configuration becomes $W = W(\tilde{v}) = W(\Omega, \theta, 2\pi/N)$, where $\tilde{v} = (i, j, \theta)$, and the desired equivariant deformation retraction is $\psi_{K+1} \cdots \psi_1$. \square

4.4 Standard Configurations (Periodic)

To define $\mathcal{W}(\Omega)$ in the periodic case, set Λ_0 to be the period of Λ_{\min} and otherwise proceed as in the non-periodic case to obtain a configuration W_T at the minimal T such that $\Lambda_T = \emptyset$. Let

$$W(\Omega, \theta, \delta) = \{(i, j, \xi + 2\pi k/p) : (i, j, \xi) \in W_T, k \in [p]\}$$

and define $\mathcal{W}(\Omega)$ as above. Note that although $\mathcal{W}(\Omega)$ is homeomorphic to \mathbb{S}^1 , the map $\mathbb{S}^1 \rightarrow \mathcal{W}(\Omega), \theta \mapsto W(\Omega, \theta, 2\pi/N)$ is a p -to-1 covering.

Proof of Lemma 4.5 (Ω periodic) We proceed the same way as in the non-periodic case, except instead of fixing a single vector $v = \tilde{v}$, we start by fixing all vectors $v_1, \dots, v_p = (i, j, \xi_1), \dots, (i, j, \xi_p)$ corresponding to Λ_{\min} . After continuously rotating certain vectors clockwise and then continuously rescaling the angular distance of each vector from the corresponding v_i as above, we obtain $\psi_{K+1}(V, 1) = W(\Omega', \xi_1, \delta') \cup \dots \cup W(\Omega', \xi_p, \delta')$ where Ω' is the combinatorial type represented by the period Λ'_{\min} of Λ_{\min} and $\delta' = 2\pi/N$. Note that Λ'_{\min} is the lexicographically minimal tableau representing Ω' with exactly one adjacent pair. Now, rotate each sub-configuration $W(\Omega', \xi_i, \delta')$ for which $\angle(\xi_i, \xi_{i-1}) > 2\pi/p$ clockwise continuously by $\phi_{i, K+2, t}$ until the vectors $\theta_i = \phi_{i, K+2, 1} \xi_i$ are spaced evenly around the circle, to obtain

$$\psi_{K+2}(V, 1) = W(\Omega, \theta_1, \delta') = W(\Omega', \theta_1, \delta') \cup \dots \cup W(\Omega', \theta_p, \delta') \in \mathcal{W}(\Omega). \quad \square$$

4.5 Contractibility

Proof of Theorem 2.4 In the depth 1 case, the full realization space $\mathcal{R}(\Omega)$ is homotopic to the space of support configurations $\mathcal{V}(\Omega)$ by Lemma 4.2, which is homotopic to \mathbb{S}^1 by Lemma 4.5. Since these homotopies are equivariant, by Remark 2.3, $\widetilde{\mathcal{R}}(\Omega)$ is contractible.

In the depth $d > 1$ case, partition Ω into layers $\Omega = \Omega_1 \cup \dots \cup \Omega_d$. If we restrict a support configurations of Ω to vectors with labels in a layer Ω_i , then we obtain a support configuration of Ω_i . Hence, $\mathcal{V}(\Omega) \subset \mathcal{V}(\Omega_1) \times \dots \times \mathcal{V}(\Omega_d)$. In the other direction, if we are given support configurations $V_i \in \mathcal{V}(\Omega_i)$, then $\bigcup_{i \in [d]} V_i \in \mathcal{V}(\Omega)$. Hence $\mathcal{V}(\Omega) = \mathcal{V}(\Omega_1) \times \dots \times \mathcal{V}(\Omega_d)$, and therefore by Lemmas 4.2 and 4.5, $\mathcal{R}(\Omega)$ is homotopic to a product of d circles, and again by Remark 2.3, $\widetilde{\mathcal{R}}(\Omega)$ is homotopic to a product of $d-1$ circles. \square

5 Topological Invariance

In this section, we prove that combinatorial type is a complete topological invariant of systems of curves on the cylinder. Specifically, we show that the associated tableaux of two systems are related by a sequence of bumps if and only if the systems are related by a self-homeomorphism of the cylinder that preserves orientation and $+\infty$. Theorem 1.2 then follows from Remark 4.3.


For $0 \leq \theta < 2\pi$, let $\zeta_\theta = \{\theta\} \times \mathbb{R}^1$. We call a curve $\gamma : \mathbb{R} \rightarrow \mathbb{S}^1 \times \mathbb{R}^1$ a *cut-path* of a system \mathcal{S} when γ diverges to $\pm\infty$, is oriented from $-\infty$ to $+\infty$, intersects each curve of \mathcal{S} exactly once, and intersects each curve one at a time away from any crossings of \mathcal{S} and away from ζ_0 . We associate to each component γ' of $\gamma \setminus \zeta_0$, the *region* of the cylinder to the left of γ' and bounded by ζ_0 . We assume every component of $\gamma \setminus \zeta_0$ intersects some curve of \mathcal{S} ; otherwise we can perform an isotopy of the cylinder that preserves \mathcal{S} and removes any components that do not intersect \mathcal{S} . Let $M = M(\mathcal{S}, \gamma)$ denote the sum of the number of curves intersecting each region plus the number of crossings in each region. Note that $M \geq n := |\mathcal{C}|$, since a cut-path intersects all curves and defines at least one region.

We define two classes of isotopies of the cylinder, called moves, sending one system and cut-path to another, while preserving the combinatorial type of the system; see Fig. 8.

Remark 5.1 Move (i) removes an crossing from one or more regions and changes the associated tableau of the system by bumping the pair of indices of the curves that cross at that crossing. While move (ii) decreases the number of curves intersecting some region, possibly deleting that region, the move does not change the associated tableau. Thus, these moves decrease the value of M and preserve combinatorial type.

Proof of Theorem 1.2 By Theorem 2.4 for any pair of systems \mathcal{S}, \mathcal{T} of the same combinatorial type, there is a path in $\mathcal{R}(\Omega)$ from \mathcal{S} to \mathcal{T} , so there is an isotopy sending \mathcal{S} to \mathcal{T} .

Fig. 8 The possible moves on a system and cut-path: (i) If a pair of curves cross each other after ζ_0 before crossing any other curve or the cut-path, then deform the curves to send this crossing through ζ_0 in the clockwise direction. (ii) If a curve crosses the cut-path before intersecting any other curve and the cut-path intersects ζ_0 either (a) immediately after or (b) immediately before this crossing, then deform the cut-path by sending this crossing through ζ_0 in the clockwise direction



For the other direction, suppose we are given a pair of systems \mathcal{S}, \mathcal{T} on \mathcal{L} and an orientation preserving homeomorphism $\varphi : \mathbb{S}^1 \times \mathbb{R}^1 \rightarrow \mathbb{S}^1 \times \mathbb{R}^1$ that preserves $+\infty$. We will find a sequence of bumps sending the associated tableau of \mathcal{T} to that of \mathcal{S} . Choose some $\varepsilon > 0$ sufficiently small so that $M(\mathcal{S}, \zeta_\varepsilon) = n$. That is, ζ_ε is a cut-path for \mathcal{S} with a single region that does not contain any crossings. And, choose ε generically so that $\eta := \varphi(\zeta_0)$ is a cut-path of \mathcal{T} . We will perform a sequence of moves starting from (\mathcal{T}, η) .

Since each move decreases the value of M and $M \geq n$, we can perform a sequence of moves on (\mathcal{T}, η) until no more moves are possible. If some region contains a crossing then we can perform move (i), and if there are multiple regions then we can perform move (ii). Therefore, we obtain a system \mathcal{U} with associated tableau Λ and cut-path γ with a single region that contains no crossings. The local sequence of a curve $U_i \in \mathcal{U}$ is given by the order in which U_i crosses the other curves after crossing γ , which is the same as that of $S_i \in \mathcal{S}$. The rows of Λ are the local sequences of \mathcal{U} in the order γ crosses the curves of \mathcal{U} , which is the same as order as that of \mathcal{S} . Hence \mathcal{U} and \mathcal{S} have the same the associated tableau Λ . Furthermore, Λ is obtained from the associated tableau of \mathcal{T} by preforming the sequence of bumps corresponding to

the above sequence of moves. Therefore \mathcal{U} , \mathcal{T} , and \mathcal{S} all have the same combinatorial type. □

6 Universality

In this section, we prove a universality theorem for arrangements of k -gons. We actually prove the following slightly more specific result.

Lemma 6.1 *For any k order types χ_1, \dots, χ_k on $[n]$, where at least two are distinct, there is a combinatorial type Ω on $[n]$ such that its k -gon realization space $\tilde{\mathcal{R}}_k(\Omega)$ is homotopy equivalent to $\tilde{\mathcal{R}}_1(\chi_1) \times \dots \times \tilde{\mathcal{R}}_1(\chi_k)$.*

Theorem 2.5 follows immediately.

Proof of Theorem 2.5 Fix a basic primary semialgebraic set \mathcal{Z} and $k > 1$. Let χ_1 be the order type of the Mnëv point set with point realization space homotopic to \mathcal{Z} . Let χ_2, \dots, χ_k all be the order type of n points in convex position. Note that the point realization space of n points in convex position is contractible. With this, the k -gon realization space of Ω from Lemma 6.1 is also homotopic to \mathcal{Z} . □

To show Lemma 6.1, we construct a combinatorial type Ω such that for every realization \mathcal{A} of Ω by k -gons, the vertices of each k -gon can be labeled. That is, each vertex can be uniquely identified using only information encoded in the combinatorial type. Note that this is not possible in general, as combinatorial type does not provide information about individual vertices directly. Furthermore, we construct Ω so that the order type of the vertices of \mathcal{A} is the same in every realization and each χ_i appears as a subset of the vertices.

We define Ω in two ways: in the primal we construct an arrangement of k -gons, then in the dual we construct a system of curves. We show in Lemma 6.4 that these two constructions define the same combinatorial type. We will use both the primal and the dual construction in the proof of Lemma 6.1.

Before defining Ω , index the order types χ_i so that the cyclic ordering $\chi_1, \chi_2, \dots, \chi_k, \chi_1, \dots$ is not periodic with period smaller than k . This is possible by the assumption that there are at least two distinct order types.

6.1 The Primal Construction

The primal construction \mathcal{A} depends on certain arbitrary choices that will not affect the combinatorial type. Assume for the primal construction that each χ_i is realizable; the non-realizable case is defined by the dual construction only.

Let \mathcal{A}_0 be an arrangement of $2k$ points in convex position denoted by $a_1^1, a_1^n, a_2^1, a_2^n, \dots, a_k^1, a_k^n$ in counter-clockwise order, such that the lines ℓ_i spanning a_i^n and a_{i+1}^1 bound a convex k -gon B .³ Observe that $B \setminus \text{conv}(\mathcal{A}_0)$ consists of k triangular regions. We construct \mathcal{A} by placing a point set realizing one of the χ_i in each of these triangular

³ Here subscripts are indices over \mathbb{Z}_k , so in particular ℓ_k is the line spanning a_k^n and a_1^1 .




Fig. 9 The point set \mathcal{P}_1 on the left is mapped to points on the right by the projective transformation determined by $p_1^1 \mapsto a_1^1, q_1^1 \mapsto a_2^1, p_1^6 \mapsto a_1^6, q_1^6 \mapsto a_2^6$

regions, then we define the k -gons A^s to have vertices consisting of one point from each region; see Fig. 9 for an example with $n = 6, k = 4$.

Let χ_i be defined on the index set $\left\{ \binom{1}{i}, \dots, \binom{n}{i} \right\}$, and let $\mathcal{P}_i = \{p_i^1, \dots, p_i^n\}$ be a realization of χ_i . Furthermore, let χ_i be indexed so that p_i^1 and p_i^2 appear on the boundary of the convex hull of \mathcal{P}_i and the local sequence of p_i^1 is $p_i^2, p_i^3, \dots, p_i^n$. That is, the angles θ_i^s at p_i^1 from the semiline through p_i^2 to the semiline through p_i^s are increasing in the counter-clockwise direction, $0 = \theta_i^2 < \theta_i^3 < \dots < \theta_i^n < \pi$. Note that this implies p_i^n is also on the boundary of the convex hull of \mathcal{P}_i , which we will call the convex boundary for short. Now augment \mathcal{P}_i by two points as follows. Let $\mathcal{Q}_i = \mathcal{P}_i \cup \{q_i^1, q_i^n\}$ such that $p_i^n, q_i^1, q_i^n, p_i^1$ appear consecutively in counter-clockwise order on the convex boundary of \mathcal{Q}_i and no line through any two points of \mathcal{P}_i separates the points q_i^1, q_i^n, p_i^1 . Note that this uniquely determines the order type of \mathcal{Q}_i ; see Fig. 9 (left).

Now define projective transformations ϕ_i such that

$$\phi_i(q_i^n) = a_{i-1}^n, \quad \phi_i(p_i^1) = a_i^1, \quad \phi_i(p_i^n) = a_i^n, \quad \phi_i(q_i^1) = a_{i+1}^1,$$

and let

$$\mathcal{P} = \{a_1^1, a_1^2, \dots, a_2^1, \dots, a_n^n\}$$

where $a_i^s = \phi_i(p_i^s)$. Note that ϕ_i sends \mathcal{Q}_i into B and preserves the orientation of q_i^1, q_i^n, p_i^1 , so $\mathcal{Q}_i \xrightarrow{\text{proj}} \{a_i^1, \dots, a_i^n, a_{i-1}^n, a_{i+1}^1\}$ with appropriate relabeling. Let $\mathcal{A} = \{A^1, \dots, A^n\}$ where $A^s = \text{conv}(\{a_1^s, a_2^s, \dots, a_k^s\})$.

Finally, let Ω denote the combinatorial type of \mathcal{A} , let χ denote the order type of \mathcal{P} , and let ψ_i denote the order type of \mathcal{Q}_i .

6.2 Path Systems

We call the graph of some indexed family of functions defined over an interval a *path system*. We say two path systems are *equivalent* when they are related by an orientation preserving self-homeomorphism of the plane that also preserves indices and the orientations of the paths. We will always assume that the end-points of a path system are all distinct, and that the paths satisfy the same genericity conditions given in Sect. 2.1 for systems of curves. For path systems $\mathcal{L}_1, \mathcal{L}_2$ over intervals $I_1, I_2 \subset \mathbb{R}$ with the same number of paths, the concatenation $\mathcal{L}_1 \cdot \mathcal{L}_2$ is the path system obtained by identifying the right edge of $I_1 \times \mathbb{R}$ with the left edge of $I_2 \times \mathbb{R}$ by a homeomorphism sending the right end-points of \mathcal{L}_1 to the left end-points of \mathcal{L}_2 . Here indices must be dealt with appropriately. If the left end-points of \mathcal{L}_1 appear in the same order as the right end-points, then we may form a system of curves $\circlearrowleft \mathcal{L}_1$ by identifying the left and right edges of $I_1 \times \mathbb{R}$ by a homeomorphism that identifies the left and right end-points of each path in \mathcal{L}_1 . Let $\updownarrow \mathcal{L}_1$ denote the path system obtained by flipping \mathcal{L}_1 vertically by the map $(x, y) \mapsto (x, -y)$. Given an order type χ , we say a path system \mathcal{L} is a *pseudoline representation* of χ when $\mathcal{S} = \circlearrowleft(\mathcal{L} \cdot \updownarrow \mathcal{L})$ is an orientable system with order type χ as in Theorem 1.3. Note that classes of equivalent path pseudoline representations of χ bijectively correspond to tableau representations of χ . We say an element i is on the *convex boundary* of χ when the corresponding curve S_i appears on the upper envelope of a corresponding system \mathcal{S} .

Remark 6.2 For each element i on the convex boundary of an order type χ , there is a unique class of equivalent pseudoline representations \mathcal{L} where L_i starts as the top most path and crosses all other paths, thereby going to the bottom, before any other crossings occur.

6.3 The Dual Construction

Let χ_i be an order type on elements $\left\{ \binom{1}{i}, \dots, \binom{n}{i} \right\}$ indexed as in the primal construction, and let \mathcal{L}_i be a pseudoline representation of χ_i with paths L_i^1, \dots, L_i^n such that L_i^1 starts at the top and crosses all other paths first as in Remark 6.2. Let $\mathcal{C} = \{C_1, \dots, C_k\}$ be the dual system of k points in convex position indexed in counter-clockwise order, and observe that each curve C_i appears exactly once on the upper envelope and once on the lower envelope of \mathcal{C} . Let \mathcal{S} be a system of curves where each curve $C_i \in \mathcal{C}$ is replaced by n curves $\{S_i^1, \dots, S_i^n\}$ in a small tubular neighborhood about C_i crossing to form a copy of \mathcal{L}_i above all other curves of \mathcal{S} and a copy of $\updownarrow \mathcal{L}_i$ below all other curves of \mathcal{S} . Let T^s be the upper envelope of the curves S_1^s, \dots, S_k^s , and let $\mathcal{T} = \{T^1, \dots, T^n\}$. Equivalently, let \mathcal{U} be the path system of size n where each path from bottom to top crosses all paths below itself (beginning with the bottom path crossing no other paths and ending with the top path crossing all other paths), and let $\mathcal{T} = \circlearrowleft(\mathcal{L}_1 \cdot \mathcal{U} \cdot \mathcal{L}_2 \cdot \mathcal{U} \cdots \mathcal{L}_k \cdot \mathcal{U})$. See Fig. 10 for an example with $n = 6, k = 4$.

Finally, let Ω denote the combinatorial type of \mathcal{T} and let χ denote the order type of \mathcal{S} .




Fig. 10 Top left: the pseudoline representation \mathcal{L}_1 of χ_1 . Top right: the system \mathcal{C} . Center: the system \mathcal{S} . Bottom: the system \mathcal{T} of combinatorial type Ω

Lemma 6.3 \mathcal{P}^* and \mathcal{S} are orientable and have the same well defined combinatorial type χ .

Proof Consider the orientations of a triple $a_u^r, a_v^s, a_w^t \in \mathcal{P}$, and the corresponding curves $S_u^r, S_v^s, S_w^t \in \mathcal{S}$. If u, v, w are distinct, then these points have the same orientation as a_u^1, a_v^1, a_w^1 , since each a_i^x among these points is between a_i^1 and a_i^n in the local sequence of a_{i-1}^n among $\phi_i(Q_i)$, which implies a_i^x is in convex position together with A_0 between a_i^1 and a_i^n in counter-clockwise order. Furthermore, the curves have the same orientation as C_u, C_v, C_w , which is the same as that of a_u^1, a_v^1, a_w^1 , since each S_i^x is in a small tubular neighborhood about C_i . If $u = v = w$ then the orientation of a_u^r, a_u^s, a_u^t as well as that of S_u^r, S_u^s, S_u^t is determined by χ_u . If $u = v$ and w is distinct, then both the points and the curves have the same orientation as p_u^r, p_u^s, p_u^1 . In any case, each triple of \mathcal{S} is orientable and its orientation is fixed and is the same as that of \mathcal{P}^* , and since order types are completely determined by the orientation of each triple, this determines the order type of \mathcal{P} and \mathcal{S} . \square

Lemma 6.4 \mathcal{A}^* and \mathcal{T} have the same well defined combinatorial type Ω .

Proof This follows from Lemma 6.3 and the fact that the support function of a polygon is the upper envelope of the support functions of its vertices. \square

6.4 Proof of Lemma 6.1

First observe that $\mathcal{R}_1(\chi_i)$ is homotopic to $\mathcal{R}_1(\psi_i)$, since ψ_i is reducible by q_i^1 and $\omega_i := \psi_i \setminus q_i^1$ is reducible by q_i^n [2, Lem. 8.2.1]. That is, for any realization \mathcal{P}_i of χ_i , there is a non-empty convex region where \mathcal{P}_i can be augmented by a point q_i^n to obtain a realization of ω_i , and the fibers of the deletion map $\delta: \mathcal{R}_1(\omega_i) \rightarrow \mathcal{R}_1(\chi_i)$ defined by deleting the point q_i^n are given by this convex region, which implies $\mathcal{R}_1(\omega_i)$ and $\mathcal{R}_1(\chi_i)$ are homotopic. Likewise, $\mathcal{R}_1(\psi_i)$ and $\mathcal{R}_1(\omega_i)$ are homotopic.

Let ω denote the non-simple order type of the point set $\mathcal{A}_0 \cup \{b_1, \dots, b_k\}$ where b_i is the meet of lines ℓ_{i-1}, ℓ_i . We claim that $\mathcal{R}_1(\chi)$ is homeomorphic to $\mathcal{R}_1(\omega) \times \widetilde{\mathcal{R}}_1(\psi_1) \times \dots \times \widetilde{\mathcal{R}}_1(\psi_k)$. To see this, let

$$\pi_0: \mathcal{R}_1(\chi) \rightarrow \mathcal{R}_1(\omega) \quad \text{and} \quad \pi_i: \mathcal{R}_1(\chi) \rightarrow \widetilde{\mathcal{R}}_1(\psi_i)$$

where π_0 is defined by restricting to the points $a_1^1, a_1^n, b_1, \dots, a_n^1, a_n^n, b_n$, and π_i is defined by restricting to the points $a_i^1, \dots, a_i^n, a_{i-1}^n, a_{i+1}^1$ respectively relabeled as $p_i^1, \dots, p_i^n, q_i^n, q_i^1$ and identifying realizations that are projectively equivalent. Let $\pi = \pi_0 \times \dots \times \pi_k$. The construction of \mathcal{P} in the primal definition of Ω shows that π is surjective, and π is injective since the points of Ω_i are uniquely determined by the projective equivalence class to which they belong and the positions of the points $q_i^n, p_i^1, p_i^n, q_i^1$, which are fixed by a realization of ω . Hence π is a bijection, and since it and its inverse are continuous, the claim holds.

Next, we claim $\mathcal{R}_k(\Omega)$ is homeomorphic to $\mathcal{R}_1(\chi)$. Let $\varphi: \mathcal{R}_1(\chi) \rightarrow \mathcal{R}_k(\Omega)$ by taking convex hulls as in the primal definition of Ω . Assume $\mathcal{R}_k(\Omega)$ is non-empty and consider $\mathcal{A} = \{A^1, \dots, A^n\} \in \mathcal{R}_k(\Omega)$; the existence of φ implies that if $\mathcal{R}_k(\Omega)$ were empty then $\mathcal{R}_1(\chi)$ would also be empty. We will show that there is a unique way of indexing the vertices of each body A^s by $\{a_1^s, \dots, a_n^s\}$ so that together they realize χ , which implies there is the unique point set $\mathcal{V} \in \mathcal{R}_1(\chi)$ for which $\varphi(\mathcal{V}) = \mathcal{A}$.

We will first see that such an indexing exists, which implies φ is surjective. Notice that A^{1*} and A^{t*} cross $2k$ times for each $t \neq 1$, so A^1 and A^t each appear k times on the boundary of $\text{conv}(A^1 \cup A^t)$, which implies each A^s must be a k -gon and the vertices of A^1 and A^t are in convex position. First choose some indexing of the vertices of A^1 by v_1^1, \dots, v_k^1 in counter-clockwise order, and then index the vertex of each A^t between v_i^1 and v_{i+1}^1 by v_i^t , and let $\mathcal{V}_i = \{v_i^1, \dots, v_i^n\}$. Now let μ_i^t be the outward normal direction of the common supporting tangent of A^1 and A^t through v_i^1 and v_i^t , and let ξ_i^t be the outward normal direction of the common supporting tangent through v_i^t and v_{i+1}^1 . The support curve of the point v_i^t coincides with A^{t*} on the half open interval $[\mu_i^t, \xi_i^t) \subset \mathbb{S}^1$, and by Lemma 6.4, has a crossing with A^{s*} in $[\mu_i^s, \xi_i^s)$ corresponding to a common supporting tangent through a_i^t, a_i^s for t, s distinct. This fixes a pseudoline representation for each \mathcal{V}_i , which is equivalent to \mathcal{P}_{i-h}^* for an appropriate cyclic shift of indices by $h \in \mathbb{Z}_k$, so $\mathcal{P}_i := \mathcal{V}_{i+h} \in \mathcal{R}_1(\chi_i)$. If we perform the primal construction by choosing the point sets \mathcal{P}_i as realizations of χ_i , and choosing $q_i^1, q_i^n, \mathcal{A}_0$ among the vertices of A^1 and A^n so that each map ϕ_i is the identity, then we obtain the same arrangement \mathcal{A} that we started with. Thus, by Lemma 6.3 the vertices of \mathcal{A} labeled by $a_i^s = v_{i+h}^s$ have order type χ .




Fig. 11 Construction for $k = 2, 3, 4$ from the proof of universality for non-crossing combinatorial types

We will now see this indexing is unique, which implies φ is injective. Since the vertices of A^1 and A^i appear in an alternating order around the convex boundary of their union, the indexing of the vertices is fixed up to a cyclic shift, and since the cyclic ordering of χ_1, \dots, χ_k does not have periodicity smaller than k , h is the unique cyclic shift of indices for which the vertices have order type χ . Hence φ is a bijection, and since it and its inverse are continuous, the claim holds.

Finally, $\mathcal{R}_k(\Omega)$ is homeomorphic to $\mathcal{R}_1(\chi)$, which is homeomorphic to $\mathcal{R}_1(\omega) \times \tilde{\mathcal{R}}_1(\psi_1) \times \dots \times \mathcal{R}_1(\psi_k)$, so by identifying projectively equivalent realizations, $\mathcal{R}_k(\Omega)$ is homeomorphic to $\mathcal{R}_1(\omega) \times \tilde{\mathcal{R}}_1(\psi_1) \times \dots \times \mathcal{R}_1(\psi_k)$, and since $\tilde{\mathcal{R}}_1(\omega)$ is contractible and $\mathcal{R}_1(\psi_i)$ is homotopic to $\tilde{\mathcal{R}}_1(\chi_i)$, $\mathcal{R}_k(\Omega)$ is homotopic to $\tilde{\mathcal{R}}_1(\chi_1) \times \dots \times \tilde{\mathcal{R}}_1(\chi_k)$. \square

7 Open Problems and Concluding Remarks

What is the smallest integer k_n such that any order type on n elements can be realized by convex k_n -gons? Theorem 1.4 gives asymptotic bounds, but there remains a wide gap.

Does universality hold for realizations of *order types* by k -gons? The authors were able to establish the weaker result, that universality holds for *non-crossing* arrangements of k -gons. That is, arrangements for which every *pair* of bodies has the same combinatorial type as a pair of distinct points, instead of every triple. The proof is similar to that of Theorem 2.4, but depends on a construction where certain pairs bodies of the arrangement are disjoint in every realization by k -gons; see Fig. 11.

Is the realization space of k -gons with the combinatorial type of n points in convex position contractible? While this may be the simplest order type, difficulties arise when the k -gons intersect; see Fig. 12.

Can every *order type* be realized by arrangements of *pairwise disjoint* bodies? Ziegler has given a construction of $2^{O(n^2)}$ distinct order types, all of which can be realized by pairwise disjoint bodies [2, Thm. 7.4.2]. This fails for combinatorial types

Fig. 12 Four triangles having the same combinatorial type as vertices of a square





Fig. 13 An arrangement with combinatorial type that cannot be realized by disjoint bodies



in general, as there exists an arrangement of four *non-crossing* bodies in which some pair must always intersect; see Fig. 13.

Our results mostly focus on the cases where k is a constant or k is infinite. It would be interesting to understand how the realization space depends on k as a function on n . For instance, Theorem 2.4 states that $\tilde{\mathcal{R}}(\Omega)$ is contractible for any non-layered combinatorial type Ω on $[n]$, but is there a function $p(n)$ which guarantees that $\tilde{\mathcal{R}}_{p(n)}(\Omega)$ is contractible? A natural guess would be $p(n) \in O(n^2)$. For combinatorial types Ω on $[n]$, are there upper bounds on the Betti numbers of $\tilde{\mathcal{R}}_k(\Omega)$ in terms of k and n ?

Unlike the order type, the definition of combinatorial type is non-local, in the sense that order type depends only on information about triples, whereas combinatorial type depends on global information. In fact, there are pairs of arrangements where for each triple, the corresponding combinatorial type of that triple is the same for both arrangements, but the combinatorial type of the two arrangements overall is not the same; see Fig. 14. Can combinatorial type be defined from local information when the local complexity is bounded? More specifically, is there an integer m_t , such that the combinatorial type of an arrangement is uniquely determined by the combinatorial type of each sub-arrangement of m_t bodies, provided that each pair of bodies has at most t common supporting tangents?

We have bounded the complexity of the bodies by working in the space of k -gons. Are there other measures of complexity that yield a universality theorem? Any continuous map from $\mathbb{R}^{c(n)}$ to the space of arrangements of convex bodies defines a finite dimensional subspace. In the case of k -gons in the plane, $c(n) = 2kn$. Does univer-

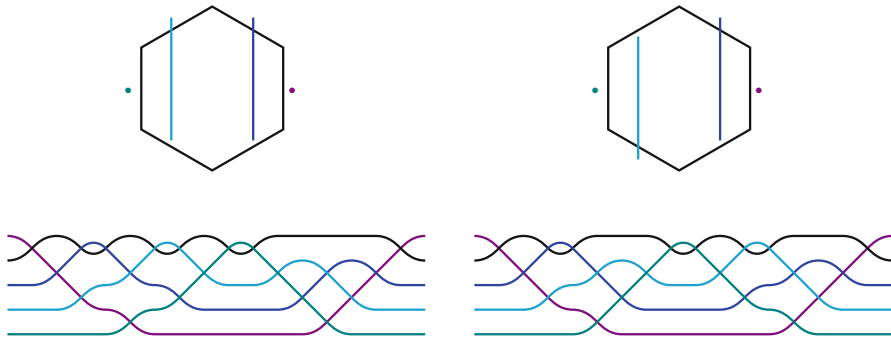


Fig. 14 Top two arrangements that are related by a bijection preserving the combinatorial type of triples, but that do not have the same combinatorial type. Bottom their dual support system

ality hold for any other such map? Consider for instance subspaces having bounded VC-dimension.

Our notion of combinatorial type extends to higher dimensions and to systems of sections of vector bundles other than the cylinder. It would be interesting to see how our results may extend in these cases.

Acknowledgements M. G. Dobbins was supported by National Research Foundation Grant NRF-2011-0030044 (SRC-GAIA) funded by the government of South Korea. 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 and by the Advanced Grant of the European Research Council GUDHI (Geometric Understanding in Higher Dimensions).

References

1. Aichholzer, O., Miltzow, T., Pilz, A.: Extreme point and halving edge search in abstract order types. *Comput. Geom.* **46**(8), 970–978 (2013)
2. Björner, A., Las Vergnas, M., Sturmfels, B., White, N., Ziegler, G.M.: *Oriented Matroids*. Encyclopedia of Mathematics and Its Applications, vol. 46, 2nd edn. Cambridge University Press, Cambridge (1999)
3. Dhandapani, R., Goodman, J.E., Holmsen, A., Pollack, R.: Interval sequences and the combinatorial encoding of planar families of convex sets. *Rev. Roum. Math. Pures Appl.* **50**(5–6), 537–553 (2005)
4. Dobbins, M.G., Holmsen, A., Hubard, A.: The Erdős–Szekeres problem for non-crossing convex sets. *Mathematika* **60**(2), 463–484 (2014)
5. Dobbins, M.G., Holmsen, A.F., Hubard, A.: Regular systems of paths and families of convex sets in convex position. *Trans. Am. Math. Soc.* **368**(5), 3271–3303 (2016)
6. Felsner, S., Valtr, P.: Coding and counting arrangements of pseudolines. *Discrete Comput. Geom.* **46**(3), 405–416 (2011)
7. Folkman, J., Lawrence, J.: Oriented matroids. *J. Comb. Theory Ser. B* **25**(2), 199–236 (1978)
8. Goodman, J.E.: Proof of a conjecture of Burr, Grünbaum, and Sloane. *Discrete Math.* **32**(1), 27–35 (1980)
9. Goodman, J.E., Pollack, R.: On the combinatorial classification of nondegenerate configurations in the plane. *J. Comb. Theory Ser. A* **29**(2), 220–235 (1980)
10. Goodman, J.E., Pollack, R.: Semispaces of configurations, cell complexes of arrangements. *J. Comb. Theory Ser. A* **37**(3), 257–293 (1984)
11. Goodman, J.E., Pollack, R.: Upper bounds for configurations and polytopes in \mathbb{R}^d . *Discrete Comput. Geom.* **1**(3), 219–227 (1986)

12. Goodman, J.E., Pollack, R.: The combinatorial encoding of disjoint convex sets in the plane. *Combinatorica* **28**(1), 69–81 (2008)
13. Groemer, H.: Geometric Applications of Fourier Series and Spherical Harmonics. *Encyclopedia of Mathematics and Its Applications*, vol. 61. Cambridge University Press, Cambridge (1996)
14. Grünbaum, B.: Arrangements and spreads. In: *CBMS Regional Conference Series in Mathematics*, vol. 10. American Mathematical Society, Providence, RI (1972)
15. Habert, L., Pocchiola, M.: Computing pseudotriangulations via branched coverings. *Discrete Comput. Geom.* **48**(3), 518–679 (2012)
16. Habert, L., Pocchiola, M.: LR characterization of chirotopes of finite planar families of pairwise disjoint convex bodies. *Discrete Comput. Geom.* **50**(3), 552–648 (2013)
17. Hubard, A.: Erdős–Szekeres para cuerpos convexos. Bachelor’s Thesis, UNAM (2005)
18. Hubard, A., Montejano, L., Mora, E., Suk, A.: Order types of convex bodies. *Order* **28**(1), 121–130 (2011)
19. Kapovich, M., Millson, J.J.: Universality theorems for configuration spaces of planar linkages. *Topology* **41**(6), 1051–1107 (2002)
20. Knuth, D.E.: Axioms and Hulls. *Lecture Notes in Computer Science*, vol. 606. Springer, Berlin (1992)
21. Levi, F.: Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade. *Ber. Math. Phys. Kl. Sächs. Akad. Wiss.* **78**, 256–267 (1926)
22. Mnev, N.E.: Varieties of combinatorial types of projective configurations and convex polyhedra. *Dokl. Akad. Nauk SSSR* **283**(6), 1312–1314 (1985) (in Russian)
23. Mnev, N.E.: The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In: Viro, O.Ya., Vershik, A.M. (eds.) *Topology and Geometry—Rohlin Seminar. Lecture Notes in Mathematics*, vol. 1346, pp. 527–543. Springer, Berlin (1988)
24. Novick, M.: Allowable interval sequences and line transversals in the plane. *Discrete Comput. Geom.* **48**(4), 1058–1073 (2012)
25. Novick, M.: Allowable interval sequences and separating convex sets in the plane. *Discrete Comput. Geom.* **47**(2), 378–392 (2012)
26. Pach, J., Tóth, G.: Families of convex sets not representable by points. In: Bhattacharya, B.B., Surkolay, S., Nandy, S.C., Bagchi, A. (eds.) *Algorithms, Architectures and Information Systems Security. Statistical Science and Interdisciplinary Research*, vol. 3, pp. 43–53. World Scientific, Hackensack (2009)
27. Padrol, A., Theran, L.: Delaunay triangulations with disconnected realization spaces. In: *30th Annual Symposium on Computational Geometry (SoCG’14)*, pp. 163–170. ACM, New York (2014)
28. Richter-Gebert, J.: Realization Spaces of Polytopes. *Lecture Notes in Mathematics*, vol. 1643. Springer, Berlin (1996)
29. Ringel, G.: Teilungen der Ebene durch Geraden oder topologische Geraden. *Math. Z.* **64**, 79–102 (1956)
30. Shor, P.W.: Stretchability of pseudolines is NP-hard. In: Gritzmann, P., Sturmfels, B. (eds.) *Applied Geometry and Discrete Mathematics. DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, vol. 4, pp. 531–554. American Mathematical Society, Providence, RI (1991)
31. Tsukamoto, Y.: New examples of oriented matroids with disconnected realization spaces. *Discrete Comput. Geom.* **49**(2), 287–295 (2013)
32. Vakil, R.: Murphy’s law in algebraic geometry: badly-behaved deformation spaces. *Invent. Math.* **164**(3), 569–590 (2006)