

The Graph of Triangulations of a Point Configuration with $d + 4$ Vertices Is 3-Connected*

M. Azaola and F. Santos

Departamento de Matemáticas, Estadística y Computación,
Universidad de Cantabria, Av. de los Castros s/n,
39071 Santander, Spain
{azaola,santos}@matesco.unican.es

Abstract. We study the graph of bistellar flips between triangulations of a vector configuration \mathcal{A} with $d + 4$ elements in rank $d + 1$ (i.e. with *corank* 3), as a step in the Baues problem. We prove that the graph is connected in general and 3-connected for *acyclic* vector configurations, which include all point configurations of dimension d with $d + 4$ elements. Hence, every pair of triangulations can be joined by a finite sequence of bistellar flips and, in the acyclic case, every triangulation has at least three geometric bistellar neighbours. In corank 4, connectivity is not known and having at least four flips is false. In corank 2, the results are trivial since the graph is a cycle.

Our methods are based on a dualization of the concept of triangulation of a point or vector configuration \mathcal{A} to that of a *virtual chamber* of its Gale transform \mathcal{B} , introduced by de Loera et al. in 1996. As an additional result we prove a *topological representation theorem* for virtual chambers, stating that every virtual chamber of a rank 3 vector configuration \mathcal{B} can be realized as a cell in some *pseudo-chamber complex* of \mathcal{B} in the same way that regular triangulations appear as cells in the usual chamber complex.

All the results in this paper generalize to triangulations of corank 3 oriented matroids and virtual chambers of rank 3 oriented matroids, realizable or not. The details for this generalization are given in the Appendix.

Introduction

A *point configuration* \mathcal{A} in \mathbb{R}^d is a finite spanning set of points in the affine space \mathbb{R}^d . A *triangulation* of \mathcal{A} is a geometric simplicial complex which covers the convex hull of \mathcal{A}

* This research was partially funded by CajaCantabria and by Grant PB97-0358 of the Spanish Dirección General de Investigación Científica y Técnica.

and whose vertices are elements of \mathcal{A} . A *bistellar flip* (or *flip*, for short) is an elementary local transformation in a triangulation of \mathcal{A} which gives rise to another triangulation of \mathcal{A} . Triangulations, flips and other necessary notions are defined in Section 1. The *graph* $G(\mathcal{A})$ of triangulations of \mathcal{A} is the graph whose vertices are all the triangulations of \mathcal{A} and whose edges represent flips between them.

The notions of triangulation and flip can be naturally defined also for a *vector configuration* \mathcal{A} , which is a finite spanning set of vectors in a real finite-dimensional vector space (say \mathbb{R}^{d+1}). A triangulation of \mathcal{A} is a *simplicial fan* (a polyhedral fan whose cones are all spanned by independent sets of vectors) that covers the positive span of \mathcal{A} and whose 1-cones are positively spanned by elements of \mathcal{A} . A point configuration \mathcal{A} in \mathbb{R}^d can be regarded as a particular case of vector configuration in \mathbb{R}^{d+1} by embedding \mathbb{R}^d in \mathbb{R}^{d+1} as an affine hyperplane not passing through the origin. A vector configuration obtained this way is called *acyclic* or *pointed*.

It has been an open question for about a decade whether the graph of triangulations of every point or vector configuration is connected. Santos [17] has found a disconnected example, with dimension 6 and corank 317. Other previous results include:

- For point configurations in the plane the graph is connected [11] and every triangulation has at least $n - 3$ flips [6]. The graph is not known to be $(n - 3)$ -connected. For point configurations in convex position in dimension 3 every triangulation has at least $n - 4$ flips, but the graph is not known to be connected [6].
- For point or vector configurations with $n \leq d + 3$ all the triangulations are *regular* [12], [5] and, hence, the graph of triangulations is isomorphic to the 1-skeleton of a polytope of dimension $n - d - 1$ [1], [8] (the so-called *secondary polytope* of \mathcal{A}).
- For any pair of parameters n and d with $n - 5 \geq d \geq 3$ there are triangulations of point configurations with n elements in dimension d which have less than $n - d - 1$ flips. In particular, the graph is not $(n - d - 1)$ -connected. The following is an example of flip deficiency for the minimal case (with $n = 8$ and $d = 3$), based on a construction from [6].

Example 1. Let $p_1 = (0, 0, 0)$, $p_2 = (1, 1, 0)$, $p_3 = (6, 0, 0)$, $p_4 = (4, 1, 0)$, $p_5 = (0, 6, 0)$, $p_6 = (1, 4, 0)$, $q = (2, 2, 4)$ and $r = (2, 2, 6)$. Let

$$\mathcal{A} = \{p_1, p_2, p_3, p_4, p_5, p_6, q, r\}$$

and let \mathcal{T} be the triangulation

$$\{\{p_1, p_2, p_3, q\}, \{p_2, p_3, p_4, q\}, \{p_3, p_4, p_5, q\}, \{p_4, p_5, p_6, q\}, \{p_5, p_6, p_1, q\}, \\ \{p_6, p_1, p_2, q\}, \{p_2, p_4, p_6, q\}, \{p_1, p_3, q, r\}, \{p_1, p_5, q, r\}, \{p_3, p_5, q, r\}\}.$$

\mathcal{T} has only three flips, supported on the three circuits $(\{p_i, p_{j+1}\}, \{p_j, p_{i+1}\})$, for $i, j \in \{1, 3, 5\}$.

These results show that point configurations with four points more than their dimension are a border case between good and bad behaviour: with less points all triangulations are regular and with more points there are triangulations with “flip-deficiency” (less than $n - d - 1$ flips). Our main result in this paper is that in this border case all triangulations

have the expected number of flips and the graph has the expected connectivity number (Corollary 3.10):

Theorem 2. *For any vector configuration \mathcal{A} with $d + 4$ elements in \mathbb{R}^{d+1} the graph $G(\mathcal{A})$ of triangulations of \mathcal{A} is connected. If \mathcal{A} is acyclic (or if \mathcal{A} is a point configuration in \mathbb{R}^d), then $G(\mathcal{A})$ is 3-connected. In particular, every triangulation of \mathcal{A} has at least three geometric bistellar neighbours.*

Our techniques are based on the duality between a vector configuration and its Gale transform, which we now explain briefly. Any vector configuration \mathcal{A} with n elements and rank $r = d + 1$ has a *Gale transform* \mathcal{B} with n elements and rank $n - r = n - d - 1$, which is dual to \mathcal{A} in the sense of oriented matroid theory (see [4] for details on oriented matroids, or the beginning of Section 2 for some properties of them). In particular, there is a canonical bijection between the bases of \mathcal{A} and the bases of \mathcal{B} . Since a triangulation T of \mathcal{A} is just a collection of bases, it has associated a certain collection \mathcal{C} of bases of \mathcal{B} . The collections of simplices of \mathcal{B} corresponding in this way to triangulations of \mathcal{A} were called *virtual chambers* of \mathcal{B} in [5].

If we want to study the graph of triangulations of a configuration \mathcal{A} of corank 3, we can do it more simply by studying the graph of virtual chambers of its dual \mathcal{B} , which has rank 3 and can be thought of as a point configuration in the 2-sphere S^2 . This is what we do.

The name *virtual chamber* comes from the following fact: the *chamber complex* of \mathcal{B} is the polyhedral decomposition of the positive span of \mathcal{B} which results as the common refinement of all the triangulations of \mathcal{B} . Its cells of maximal dimension are called *chambers*. To each chamber \mathcal{C} of \mathcal{B} we associate the collection of bases whose positive span contains \mathcal{C} . The fundamental result in the theory of secondary polytopes (see [1] and [8]) is that the collections of bases of \mathcal{B} which arise in this way are precisely the duals to the *regular* triangulations of \mathcal{A} . (These are the triangulations which can be obtained as the projection of the lower envelope of a $(d + 1)$ -dimensional polytope.) Adjacency between chambers corresponds in this picture to bistellar flips between regular triangulations. A consequence of this duality between (geometric) chambers and regular triangulations is that the subgraph of triangulations induced by regular triangulations of a point configuration \mathcal{A} with n points and dimension d is the 1-skeleton of a polytope of dimension $n - d - 1$: the *secondary polytope* of \mathcal{A} whose normal fan is the chamber complex of \mathcal{B} .

In summary, virtual chambers are combinatorial objects which have similar properties to chambers, except that they do not exist geometrically. To illustrate this, in Example 1.5 we show the classical non-regular triangulation of the vertices of two nested triangles in the plane. In its Gale transform, the virtual chamber corresponding to this triangulation collapses (see Fig. 1).

The structure of the paper is as follows: In Section 1 we introduce the necessary definitions and notation as well as some background on Gale duality, virtual chambers and other tools that we use. Most of this section applies to arbitrary rank or number of points. The only new result here is a description of “flips between virtual chambers” of a configuration, i.e. a dualization of the concept of flip between triangulations, both in arbitrary rank (Theorem 1.8) and in rank 3 (Definition 1.10 and Corollary 1.11).

Sections 2 and 3 are the central part of the paper, leading to the proof of Theorem 2. Section 2 begins with an account of some basics of oriented matroid theory that we use frequently and then shows some geometric properties of rank 3 vector configurations. Section 3 contains the proof of Theorem 2 (Corollary 3.10). It is interesting to observe that the hardest part of this proof is showing the existence of at least one flip (Theorem 3.3, which is essential for Corollary 3.5).

A rank 3 vector configuration \mathcal{B} can be regarded as a point configuration in the sphere S^2 . Recall that the chamber complex of \mathcal{B} is the common refinement of all its triangulations, i.e. the cell decomposition of S^2 obtained drawing all possible geodesic segments between pairs of points of \mathcal{B} . In Section 4 we define pseudo-chamber complexes of \mathcal{B} by allowing non-geodesic arcs to serve as “pseudo-segments” but requiring them to reproduce the combinatorial situation given with the geodesic ones (essentially, requiring them to be consistent with the oriented matroid of \mathcal{B}). One easily proves that the full-dimensional cells of a pseudo-chamber complex represent some virtual chambers of \mathcal{B} , and cells of co-dimension 1 represent flips between them. Our task in Section 4 is to prove that every virtual chamber of \mathcal{B} realizes as a pseudo-chamber of some pseudo-chamber complex of \mathcal{B} . For instance, Fig. 3 shows the non-geometric virtual chamber of Example 1.5 as a pseudo-chamber. The main result of this section is (Theorem 4.13):

Theorem 3. *Any virtual chamber \mathcal{C} of a rank 3 vector configuration \mathcal{B} realizes as a pseudo-chamber of some pseudo-chamber complex of \mathcal{B} .*

In the Appendix we show that all the results of this paper also hold for non-realizable oriented matroids. This is motivated by the fact that the collection of triangulations of a vector or point configuration depends only on the underlying oriented matroid (this is well known and follows, for example, from the results in [5]). The concepts of triangulation and flip have been generalized to non-realizable oriented matroids in Section 9.6 of [4] and [16]. Since some of our proofs in Sections 2 and 3 are done in the language of oriented matroid theory, they are valid without change for non-realizable oriented matroids. In the rest of proofs our methods are mainly topological, so our starting point for the generalization is to have a topological picture of a non-realizable oriented matroid of rank 3. This is provided by the fact that every rank-3 oriented matroid has an *adjoint* and, hence, can be pseudo-realized as a *pseudo-configuration of points* in the sphere S^2 (see Sections 5.3 and 6.3 of [4]). Summarizing, the results in the Appendix say that:

Theorem 4. *Let \mathcal{M} be a corank 3 oriented matroid. Let \mathcal{M}^* be its dual oriented matroid and let \mathcal{B} be a pseudo-realization of \mathcal{M}^* . Then:*

1. *The graph of triangulations $G(\mathcal{M})$ of \mathcal{M} is connected and, if \mathcal{M} is acyclic, 3-connected.*
2. *Every virtual chamber of \mathcal{M}^* can be realized as a pseudo-chamber of some pseudo-chamber complex of \mathcal{B} .*

Our results are related to the so-called *Baues problem* in the following way. The poset of all polyhedral subdivisions of a point configuration \mathcal{A} is usually called the *Baues poset* $\omega(\mathcal{A})$ of \mathcal{A} . The Baues problem is to decide whether this poset is homotopy equivalent to

a sphere whose dimension is the same as that of the boundary of the secondary polytope (as usual, when referring to the topology of a poset we mean the topology of its order complex, see [3]). No non-spherical example is known but sphericity (and connectivity) has been proved only up to dimension or corank 2: In corank 2 the poset is the proper part of the face lattice of the secondary polytope; for dimension 2 see [7]. Whenever the graph $G(\mathcal{A})$ is connected the poset $\omega(\mathcal{A})$ is connected too (see [14] for a proof). This, together with our results on pseudo-realizability of virtual chambers, leads to the following:

Corollary 5. *For any point or vector configuration \mathcal{A} of corank 3 the Baues complex is connected. If \mathcal{A} is acyclic, then for every subdivision S of \mathcal{A} there is a subcomplex of $\omega(\mathcal{A})$ containing S and homeomorphic to a 2-sphere.*

It is natural to ask whether our results can hold in higher corank. An obstacle for this are examples with flip-deficiency in corank 4 (see Example 1) which imply that the graph of triangulations is not 4-connected and that a pseudo-realizability result of virtual chambers such as our Theorem 3 is not possible. For non-realizable oriented matroids things are even worse, since there exist oriented matroids of corank 4 whose Baues poset and whose graph of triangulations contain isolated elements (i.e. there are triangulations with no flips at all). See Section 4 of [16].

An optimistic possibility is that the graph of triangulations might be connected at least for oriented matroids which have an adjoint (which include realizable ones) of corank 4 (or of arbitrary corank). This conjecture is based on the fact that having an adjoint is crucial for our results in the Appendix and that the disconnected examples in corank 4 are obtained with non-Euclidean oriented matroids, which do not have adjoints. Our methods indicate that a crucial step towards knowing whether this is true in corank 4 is deciding whether every triangulation of a corank 4 vector configuration has at least one flip.

1. Triangulations, Flips and Virtual Chambers

1.1. Triangulations

We call a finite spanning set of vectors \mathcal{A} in a finite-dimensional real vector space $\mathbf{V} \cong \mathbb{R}^{d+1}$ a *vector configuration of rank $d + 1$* . For any subset $\sigma \subset \mathcal{A}$ the *positive span of σ* is the set $\text{conv}(\sigma) \subset \mathbf{V}$ of all non-negative linear combinations of σ and we call the set $\text{relconv}(\sigma)$ of strictly positive combinations the *relative interior* of σ (observe that, formally speaking, we call the relative interior of σ the *relative interior of the convex hull* of σ ; hence the notation relconv).

We call the (signed) circuits and cocircuits of the oriented matroid $\mathcal{M}(\mathcal{A})$ realized by \mathcal{A} the *circuits* and *cocircuits* of \mathcal{A} . In other words, a *circuit* is a pair (C^+, C^-) of subsets of \mathcal{A} such that $C^+ \cup C^-$ is a minimal dependent set and C^+ and C^- are the subsets of elements with positive and negative coefficient respectively in a dependence relation in $C^+ \cup C^-$. $C^+ \cup C^-$ is called the *support* of the circuit (so that (C^+, C^-) and (C^-, C^+) are the only circuits with support in $C^+ \cup C^-$). A *cocircuit* is a pair (C^+, C^-) where C^+

and C^- are the intersections with \mathcal{A} of the two open half-spaces defined by a hyperplane spanned by elements of \mathcal{A} . Again, the opposite of a cocircuit is a cocircuit.

We call the independent subsets of \mathcal{A} its *simplices*. A simplex is *maximal* or *full-dimensional* if it has $d + 1$ elements. We denote by $\Delta(\mathcal{A})$ the collection of all the full-dimensional simplices of \mathcal{A} (i.e. its *bases*). A *triangulation* of \mathcal{A} is any collection $\mathcal{T} \subset \Delta(\mathcal{A})$ of full-dimensional simplices of \mathcal{A} which:

1. *intersect properly*, i.e. for every pair of simplices $\sigma, \tau \in \mathcal{T}$ one has $\text{conv}(\sigma \cap \tau) = \text{conv}(\sigma) \cap \text{conv}(\tau)$;
2. *cover* \mathcal{A} , meaning that $\bigcup_{\sigma \in \mathcal{T}} \text{conv}(\sigma) = \text{conv}(\mathcal{A})$.

For a generic vector $v \in \text{conv}(\mathcal{A})$ we will always have that if $v \in \text{conv}(\sigma)$ for a simplex $\sigma \subset \mathcal{A}$, then σ is full-dimensional. In these case we call the collection of bases of \mathcal{A} containing v in their positive span the *chamber* of v in \mathcal{A} :

$$\mathcal{C}_{v,\mathcal{A}} := \{\sigma \in \Delta(\mathcal{A}) : v \in \text{conv}(\sigma)\}.$$

A collection of bases of \mathcal{A} is called a *chamber of \mathcal{A}* if it is a chamber of v for some generic vector $v \in \text{conv}(\mathcal{A})$. Any triangulation and any chamber of \mathcal{A} have a unique simplex in common.

1.2. Virtual Chambers

Every spanning subset ρ of \mathcal{A} with $d + 2$ vectors contains the support of a unique circuit $Z = (Z^+, Z^-)$ of \mathcal{A} . If both Z^+ and Z^- are non-empty (we say that Z is *acyclic*), then ρ can be triangulated in exactly two ways:

$$\mathcal{T}^+(\rho) := \{\rho \setminus a : a \in Z^+\}, \quad \mathcal{T}^-(\rho) := \{\rho \setminus a : a \in Z^-\}.$$

We say that $\mathcal{T}^+(\rho)$ and $\mathcal{T}^-(\rho)$ are a pair of opposite *triangulated circuits* of \mathcal{A} .

Definition 1.1 [5]. Let $\mathcal{C} \subset \Delta(\mathcal{A})$ be a collection of full-dimensional simplices of \mathcal{A} . We say that \mathcal{C} is a *virtual chamber* of \mathcal{A} if the following two conditions are satisfied:

1. \mathcal{C} has exactly one element in common with any triangulation of \mathcal{A} .
2. For every pair of opposite triangulated circuits $\mathcal{T}^+(\rho)$ and $\mathcal{T}^-(\rho)$ of \mathcal{A} , $\mathcal{C} \cap \mathcal{T}^+(\rho)$ is non-empty if and only if $\mathcal{C} \cap \mathcal{T}^-(\rho)$ is non-empty.

Remark 1.2. 1. Chambers of \mathcal{A} are virtual chambers as well, as can be easily checked.

2. Every triangulated circuit is contained in some triangulation of \mathcal{A} . Thus, part 1 of the definition implies that in part 2, if $\mathcal{C} \cap \mathcal{T}^+(\rho)$ is non-empty, then it contains a unique simplex (same for $\mathcal{C} \cap \mathcal{T}^-(\rho)$).

3. The results in [5] imply that:

- (i) If \mathcal{A} is in general position (i.e. if every subset of \mathcal{A} with no more than $d + 1$ elements is independent), then condition 1 in Definition 1.1 implies condition 2. This is not the case in general.

- (ii) Provided that condition 2 holds, saying “for any triangulation” in condition 1 is equivalent to saying “there is a triangulation”.

A *Gale transform* of a vector configuration $\mathcal{A} = \{a_1, \dots, a_n\}$ with n vectors in $\mathbf{V} \cong \mathbb{R}^{d+1}$ is a vector configuration $\mathcal{B} = \{b_1, \dots, b_n\}$ in $\mathbf{W} \cong \mathbb{R}^{n-d-1}$ such that the kernels of the two natural linear maps $\mathbb{R}^n \rightarrow \mathbf{V}(e_i \mapsto a_i)$ and $\mathbb{R}^n \rightarrow \mathbf{W}(e_i \mapsto b_i)$ are orthogonal complements in \mathbb{R}^n .

The oriented matroids of \mathcal{A} and \mathcal{B} are dual to each other; i.e. circuits of \mathcal{A} are cocircuits of \mathcal{B} and vice versa. A subset $\{a_{i_1}, \dots, a_{i_r}\}$ is spanning in \mathcal{A} if and only if the complement subset $\mathcal{B} \setminus \{b_{i_1}, \dots, b_{i_r}\}$ is independent in \mathcal{B} . In particular, $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$ are canonically identified under complementation of indices. The theory of secondary polyhedra [1], [2] implies that chambers of \mathcal{B} correspond to regular triangulations of \mathcal{A} . The definition of virtual chambers given above extends this correspondence to all the triangulations of \mathcal{A} :

Theorem 1.3 [5]. *Let \mathcal{A} and \mathcal{B} be vector configurations which are Gale transforms of each other. Then, under the natural identification of $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$, virtual chambers of \mathcal{B} correspond exactly to triangulations of \mathcal{A} and virtual chambers of \mathcal{A} to triangulations of \mathcal{B} .*

The following property of virtual chambers is not difficult but a little bit tedious to prove without Theorem 1.3 (using induction on the cardinality of $\sigma \setminus \tau$ and condition 2 of Definition 1.1). The use of Theorem 1.3 makes the proof much shorter.

Lemma 1.4. *Let \mathcal{C} be a virtual chamber of \mathcal{A} . Then, for any pair of simplices σ, τ in \mathcal{C} , the relative interiors $\text{relconv}(\sigma)$ and $\text{relconv}(\tau)$ intersect.*

Proof. Let \mathcal{B} be the Gale transform of \mathcal{A} . Let \mathcal{T} be the triangulation of \mathcal{B} corresponding to the virtual chamber \mathcal{C} . In oriented matroid terms, the relative interiors of σ and τ intersect if and only if there is no cocircuit $Z = (Z^+, Z^-)$ of \mathcal{A} with $\tau \cap Z^+ = \emptyset = \sigma \cap Z^-$ (intuitively, if no hyperplane weakly separates σ and τ). Translated into \mathcal{B} , what we need to prove is that for no pair of simplices $\sigma^c, \tau^c \in \mathcal{T}$ is there a circuit $Z = (Z^+, Z^-)$ of \mathcal{B} with $Z^+ \subset \tau^c$ and $Z^- \subset \sigma^c$.

That this holds is a well-known property of triangulations, since a circuit with $Z^+ \subset \tau^c$ and $Z^- \subset \sigma^c$ would imply that σ^c and τ^c do not intersect properly (see, for example, Proposition 2.2 of [13]). □

Example 1.5 (A Non-Geometric Virtual Chamber). Let $p_1 = (4, 0, 0)$, $p_2 = (0, 4, 0)$, $p_3 = (0, 0, 4)$, $p_4 = (2, 1, 1)$, $p_5 = (1, 2, 1)$ and $p_6 = (1, 1, 2)$. Let

$$\mathcal{A} := \{p_1, p_2, p_3, p_4, p_5, p_6\}.$$

\mathcal{A} can be regarded as a point configuration in the plane, depicted in Fig. 1(a), which also shows a triangulation \mathcal{T} of \mathcal{A} that is not regular. (This is the most classical example of a non-regular triangulation. See, for instance, [19] for a proof of non-regularity.) The




Fig. 1. A non-regular triangulation of the configuration \mathcal{A} of Example 1.5 (a) and an affine Gale diagram of \mathcal{A} (b).

maximal simplices (i.e. the triangles) of \mathcal{T} are $\{p_4, p_5, p_6\}$, $\{p_1, p_2, p_4\}$, $\{p_2, p_3, p_5\}$, $\{p_1, p_3, p_6\}$, $\{p_2, p_4, p_5\}$, $\{p_3, p_5, p_6\}$ and $\{p_1, p_4, p_6\}$.

One Gale transform \mathcal{B} of \mathcal{A} is defined by the vectors $(2, 1, 1)$, $(1, 2, 1)$, $(1, 1, 2)$, $(-4, 0, 0)$, $(0, -4, 0)$ and $(0, 0, -4)$.

In Fig. 1(b) we show an affine Gale diagram of \mathcal{A} , that is, a central projection of its Gale transform \mathcal{B} to a generic affine hyperplane, in which a projected point is solid when it is a positive multiple of the vector from which it comes, and open if it is a negative multiple. By a generic hyperplane we mean one which is not parallel to any of the vectors of \mathcal{B} . We have taken the hyperplane defined by the equation $x + y + z = 4$. Figure 1(b) also shows the chamber complex of \mathcal{B} ; lines joining open and solid points have arrows since they must “pass through infinity” in this representation, and every line through an open point has been broken. For more details on affine Gale diagrams, see [19]. Each point has been labelled using the index of the corresponding element of \mathcal{A} . The affine Gale diagram in Fig. 1(b) shows that the virtual chamber \mathcal{C} of \mathcal{B} which corresponds to \mathcal{T} is not a geometric chamber. The triangles of \mathcal{C} are $\{1, 2, 3\}$, $\{3, 5, 6\}$, $\{1, 4, 6\}$, $\{2, 4, 5\}$, $\{1, 3, 6\}$, $\{1, 2, 4\}$ and $\{2, 3, 5\}$, whose relative interiors intersect in the empty set as can be seen in the affine Gale diagram.

1.3. Flips

Flips are a notion of a *minimal* or *elementary* change between triangulations (see Chapter 7 of [8] or [6] and [18]). We intend to dualize the standard definition of flip and give a definition of *flip between virtual chambers* of \mathcal{B} . For this we recall the matroidal operations of contraction and deletion and some other preliminaries.

Let $\mathcal{A} \subset \mathbf{V} \cong \mathbb{R}^{d+1}$ be a vector configuration and let $\tau \subset \mathcal{A}$. Given a linear injective map $i: \mathbf{W} \rightarrow \mathbf{V}$ whose image contains and is spanned by $\mathcal{A} \setminus \tau$, we call the vector configuration $i^{-1}(\mathcal{A} \setminus \tau)$ in the vector space \mathbf{W} the *deletion of τ* in \mathcal{A} . The deletion of τ in \mathcal{A} is unique up to linear isomorphism, so we can assume \mathbf{W} to be the subspace spanned by $\mathcal{A} \setminus \tau$ and i to be the inclusion map. For this reason the deletion of τ in \mathcal{A} is denoted by $\mathcal{A} \setminus \tau$.

Given a linear projection map $\pi: \mathbf{V} \rightarrow \mathbf{W}$ whose kernel contains and is spanned by τ , we call the vector configuration $\mathcal{A}/\tau := \pi(\mathcal{A} \setminus \tau)$ in the vector space \mathbf{W} the *contraction of τ in \mathcal{A}* . The contraction of τ in \mathcal{A} is unique modulo linear isomorphism, so the map π can be assumed to be a projection of \mathbf{V} onto a linear subspace \mathbf{W} complementary to the linear span of τ .

Contraction and deletion are dual operations: if \mathcal{A} and \mathcal{B} are Gale transforms of each other and τ is a subset of \mathcal{A} , then \mathcal{A}/τ and $\mathcal{B} \setminus \tau$ are again Gale transforms of each other (here and in what follows we identify the elements of the Gale transform \mathcal{B} with the elements of \mathcal{A} in the natural way, so that τ is considered a subset of \mathcal{B}).

We denote by $S * T := \{\sigma \cup \tau: \sigma \in S, \tau \in T\}$ the *join* of two simplicial complexes S and T [9]. If $\tau \subset \mathcal{A}$ is contained in some simplex of a triangulation \mathcal{T} of \mathcal{A} , the *link of τ in \mathcal{T}* is defined to be the following collection of subsets of \mathcal{A} :

$$\text{link}_{\mathcal{T}}(\tau) := \{\sigma \setminus \tau: \tau \subset \sigma \in \mathcal{T}\}.$$

It is clear that all the elements of $\text{link}_{\mathcal{T}}(\tau)$ are full-dimensional simplices in \mathcal{A}/τ . Even more, it can be easily proved that $\text{link}_{\mathcal{T}}(\tau)$ is a triangulation of \mathcal{A}/τ . Also, it is trivially verified that $\text{link}_{\mathcal{T}}(\tau) * \{\tau\} \subset \mathcal{T}$.

Finally, recall that the support $\overline{Z} = Z^+ \cup Z^-$ of an acyclic circuit $Z = (Z^+, Z^-)$ can be triangulated in exactly two ways:

$$\mathcal{T}^+(Z) := \{\overline{Z} \setminus \{p\} \mid p \in Z^+\}, \quad \mathcal{T}^-(Z) := \{\overline{Z} \setminus \{p\} \mid p \in Z^-\}.$$

Definition 1.6. Let \mathcal{T} be a triangulation of \mathcal{A} and let $Z = (Z^+, Z^-) \subset \mathcal{A}$ be an acyclic circuit of \mathcal{A} . Suppose that the following conditions are satisfied:

1. The triangulation $\mathcal{T}^+(Z)$ is a subcomplex of \mathcal{T} .
2. All the simplices $\overline{Z} \setminus \{p\} \in \mathcal{T}^+(Z)$ ($p \in Z^+$) have the same link L in \mathcal{T} . In particular $\mathcal{T}^+(Z) * L \subset \mathcal{T}$.

In these conditions we can obtain a new triangulation \mathcal{T}' of \mathcal{A} by replacing the subcomplex $\mathcal{T}^+(Z) * L$ of \mathcal{T} with the complex $\mathcal{T}^-(Z) * L$. This operation of changing the triangulation is called a *geometric bistellar flip* (or a *flip*, for short) supported on the circuit (Z^+, Z^-) . We say that \mathcal{T} and \mathcal{T}' are *geometric bistellar neighbours*.

Proposition 1.7. Let \mathcal{T} be a triangulation of \mathcal{A} and let $Z = (Z^+, Z^-)$ be a circuit. Then \mathcal{T} has a flip supported on Z if and only if there is a triangulation L of the contraction \mathcal{A}/\overline{Z} such that $\mathcal{T}^+(Z) * L \subset \mathcal{T}$.

Proof. The support of a circuit Z is minimal linearly dependent, so that \overline{Z} spans the same subspace as $\overline{Z} \setminus \{p\}$, for every $p \in \overline{Z}$. In particular, $\mathcal{A}/\overline{Z} = \mathcal{A}/\tau$ for every maximal simplex $\tau = \overline{Z} \setminus \{p\}$ in $\mathcal{T}^+(Z)$.

With this, the “only if” part is trivial, since the link in \mathcal{T} of any $\sigma \in \mathcal{T}^+(Z)$ will be a triangulation of the contraction \mathcal{A}/\overline{Z} . For the “if” part, let L be a triangulation of \mathcal{A}/\overline{Z} such that $\mathcal{T}^+(Z) * L \subset \mathcal{T}$. Then $L \subset \text{link}_{\mathcal{T}}(\sigma)$ for every $\sigma \in \mathcal{T}^+(Z)$. Since both L and $\text{link}_{\mathcal{T}}(\sigma)$ are triangulations of \mathcal{A}/\overline{Z} and no triangulation is properly contained in another one, $L = \text{link}_{\mathcal{T}}(\sigma)$, $\forall \sigma \in \mathcal{T}^+(Z)$. □

Theorem 1.8. *Let \mathcal{C} be a virtual chamber of a vector configuration \mathcal{B} and let $Z = (Z^+, Z^-)$ be a cocircuit of \mathcal{B} . Let \mathcal{A} be the Gale transform of \mathcal{B} and let \mathcal{T} be the triangulation of \mathcal{A} corresponding to \mathcal{C} . Then \mathcal{T} has a flip supported on the circuit $Z = (Z^+, Z^-)$ of \mathcal{A} if and only if there is a virtual chamber \mathcal{C}_Z in the deletion $\mathcal{B} \setminus \overline{Z}$ such that $\tau \cup \{p\} \in \mathcal{C}$ for every $\tau \in \mathcal{C}_Z$ and $p \in Z^+$. Moreover, in these conditions $\mathcal{C} \setminus \{\tau \cup \{p\} : \tau \in \mathcal{C}_Z, p \in Z^+\} \cup \{\tau \cup \{p\} : \tau \in \mathcal{C}_Z, p \in Z^-\}$ is the virtual chamber of \mathcal{B} corresponding to the triangulation obtained by the flip of \mathcal{T} supported on the circuit Z .*

Proof. The statement is just a dualization of Proposition 1.7 taking into account Theorem 1.3. The idea is that triangulations of \mathcal{A}/\overline{Z} and virtual chambers of $\mathcal{B} \setminus \overline{Z}$ are in bijection, since \mathcal{A}/\overline{Z} and $\mathcal{B} \setminus \overline{Z}$ are Gale transforms of each other. Calling $Z^0 = \mathcal{A} \setminus (\overline{Z})$ we have that the complement of an element $\sigma \cup (\overline{Z} \setminus \{p\})$ of $L * \mathcal{T}^+(Z)$ (where $p \in Z^+$) equals $(Z^0 \setminus \sigma) \cup \{p\}$ and vice versa. □

1.4. Virtual Chambers and Flips in Rank 3

The simplices of one, two, or three elements of a vector configuration \mathcal{B} are called *vertices*, *edges* and *triangles*, respectively. We say that a simplex τ is *empty* if $\text{conv}(\tau) \cap \mathcal{B} = \tau$.

There is a natural correspondence between cocircuits $Z = (Z^+, Z^-)$ of a vector configuration \mathcal{B} and open half-spaces H^+ whose boundary hyperplane H^0 is spanned by elements of \mathcal{B} . Indeed, given a cocircuit Z , the complement of its support Z^0 spans a hyperplane which partitions $\mathcal{B} \setminus Z^0$ as $Z^+ \cup Z^-$. Reciprocally, a half-space H^+ with these conditions provides a cocircuit $(\mathcal{B} \cap H^+, \mathcal{B} \cap H^-)$. If \mathcal{B} has rank 3, a cocircuit can thus be specified by choosing one of the two sides of an edge $\{p, q\} \subset \mathcal{B}$ and calling it “positive.” An edge of \mathcal{B} together with such a choice is called an *oriented edge*. Given an oriented edge m and the corresponding cocircuit $Z = (Z^+, Z^-)$ we denote $m^i := Z^i$ for $i \in \{+, -, 0\}$ for simplicity.

On the other hand, the deletion of the support \overline{Z} in a vector configuration \mathcal{B} is the subconfiguration $Z^0 = \mathcal{B} \setminus \overline{Z}$ of \mathcal{B} (considered as a vector configuration in the vector subspace it spans). When \mathcal{B} has rank 3, Z^0 is a rank 2 vector configuration and its virtual chambers are in bijection with its empty edges, as the following result from [5] shows.

In the following statement and in what follows, a vector configuration is called *simple* if it has no zero vectors and no pair of vectors which are positive multiples of each other. This is a slightly more general definition than the standard literature, where *simple oriented matroids* are not allowed to have negative multiples either (see [4]).

Lemma 1.9 [5]. *Let \mathcal{B} be a simple rank 2 vector configuration. Then every virtual chamber has a unique empty edge and every empty edge $\{p, q\}$ is in a unique virtual chamber, which consists of those edges whose positive span contains p and q .*

Proof. Every virtual chamber contains a unique empty edge by Theorem 1.3 since the collection of empty edges is a triangulation of \mathcal{B} .

Reciprocally, for any empty edge $\{p, q\}$, the collection of edges whose positive spans contain $\text{conv}(p, q)$ is a chamber and, in particular, a virtual chamber. It suffices to show that all virtual chambers arise in this form. For a proof of this see Proposition 5.7 of [5]. \square

The previous lemma suggests the following definition, with which Theorem 1.8 translates into Corollary 1.11 below.

Definition 1.10. Let \mathcal{C} be a virtual chamber of a simple rank 3 vector configuration \mathcal{B} and let $m = \{p, q\}$ be an empty edge of \mathcal{B} (which we consider oriented). We say that m supports a flip of \mathcal{C} if $m^+ \neq \emptyset \neq m^-$ and for every $s \in m^+$ the triangle $\{p, q, s\}$ is in \mathcal{C} .

Corollary 1.11. Let \mathcal{B} be a simple rank 3 vector configuration and let \mathcal{C} be a virtual chamber of \mathcal{B} . Let \mathcal{T} be the triangulation of \mathcal{A} corresponding to \mathcal{C} and let $Z = (Z^+, Z^-)$ be a cocircuit of \mathcal{B} . Then \mathcal{T} has a flip supported on the circuit $Z = (Z^+, Z^-)$ of \mathcal{A} if and only if there is an empty edge $m \subset Z^0$ of \mathcal{B} such that m supports a flip of \mathcal{C} . Moreover, in such conditions $(\mathcal{C} \setminus \{\{a, b, p\}: p \in m^+, \text{conv}(m) \subset \text{conv}(a, b)\}) \cup \{\{a, b, p\}: p \in m^-, \text{conv}(m) \subset \text{conv}(a, b)\}$ is the virtual chamber of \mathcal{B} corresponding to the triangulation of \mathcal{A} obtained by the flip of \mathcal{T} supported on the circuit Z .

We need the vector configuration \mathcal{B} to be simple in the previous statements since, for example, a vector configuration in which every vector has a positive multiple has no empty edges at all. However, it implies no real loss of generality for our purposes since:

Lemma 1.12. Any vector configuration \mathcal{B} has the same virtual chambers and flips as the simple vector configuration \mathcal{B}_0 obtained by removing from \mathcal{B} the zero vector and all but one of the vectors in any half-line.

Proof. The zero vector clearly does not affect the collection of triangulations or flips. For the case of positive multiples, this follows easily from the fact that if $v, w \in \mathcal{B}$ are positive multiples of each other, then $(\{v\}, \{w\})$ is a circuit. Hence, for every simplex σ containing v , every virtual chamber of \mathcal{B} either contains both σ and $\sigma \cup \{w\} \setminus \{v\}$ or none of them. In other words, the simplices containing v and containing w are equivalent with respect to virtual chambers (and flips). \square

The previous lemma appears in dual form (i.e. for triangulations of the Gale transforms of \mathcal{B} and \mathcal{B}_0) in [6] and generalized to oriented matroids in Section 4.4 of [16]. We return to it in the Appendix, Lemma A.4.

1.5. Extensions

An extension of an oriented matroid \mathcal{M} on a set E is any oriented matroid \mathcal{M}' on a set $E' \supset E$ such that every circuit of \mathcal{M} is a circuit of \mathcal{M}' as well (i.e. such that

$\mathcal{M}' \setminus (E \setminus E') = \mathcal{M}$). It is a *one-element* extension if $E' \setminus E$ has exactly one element p (see [4] or [16] for details). In this case we denote the extension as $\mathcal{M} \cup \{p\}$.

Following [16] we say that a single-element extension $\mathcal{M} \cup \{p\}$ is *interior* if there is a circuit $(\{p\}, A)$ for some $A \subset B$ and that it is in *general position* if any circuit containing p is spanning (equivalently, if \mathcal{M}/p is uniform). A key property for our purposes is that if \mathcal{M} is the oriented matroid realized by a vector configuration \mathcal{B} and $\mathcal{M} \cup \{p\}$ is an interior one-element extension in general position, then the collection $\{A \subset \mathcal{B}: (\{p\}, A) \text{ is a circuit}\}$ is a virtual chamber of \mathcal{B} [16].

We are specially interested in some one-element extensions called *lexicographic extensions*, introduced by Las Vergnas [10]. The definition we use is less general than the standard one and is adapted to rank 3.

Definition 1.13 [4], [10]. Let \mathcal{M} be a rank 3 oriented matroid (or a rank 3 vector configuration) and let $\{a_1, a_2, a_3\}$ be a triangle (i.e. a basis) of \mathcal{M} .

The *lexicographic extension of \mathcal{M} at the (ordered) basis $[a_1, a_2, a_3]$* is the unique one-element extension $\mathcal{M} \cup \{p\}$ of \mathcal{M} in which every cocircuit C of \mathcal{M} is extended to the cocircuit (which we still denote C) defined by $C(p) = C(a_i)$ for the minimal i with $C(a_i) \neq 0$.

For the existence and uniqueness of the lexicographic extension in a basis see Section 7.2 of [4]. If \mathcal{M} is an oriented matroid realized by a vector configuration \mathcal{B} , then the lexicographic extension at the basis $[a_1, a_2, a_3]$ can be realized by adding to \mathcal{B} the vector $a_1 + \varepsilon a_2 + \varepsilon^2 a_3$ for any sufficiently small positive scalar ε . Hence, the lexicographic extension has an associated chamber in the chamber complex of \mathcal{B} , which is incident to a_1 and to the edge $[a_1, a_2]$ on the side on which a_3 is. We call that chamber a *flag* chamber. The corresponding triangulation of the Gale transform is called a *pushing* triangulation [12], [4] and it is regular.

2. Triangles and Edges

In the rest of the paper (except for the Appendix) \mathcal{B} denotes a simple rank 3 vector configuration. Recall that we call the *simplices* of \mathcal{B} its independent subsets and we call a simplex a *point*, *edge* or *triangle* if it has one, two or three elements, respectively. We say that a simplex τ is *empty* if $\text{conv}(\tau) \cap \mathcal{B} = \tau$. \mathcal{B} being simple means that every point is empty.

Without loss of generality, we suppose that every vector in \mathcal{B} has unit length and we think of \mathcal{B} as a point configuration in the sphere. In this setting $\text{conv}(l)$ and $\text{conv}(\tau)$ for an edge l and a triangle τ are a geodesic segment and a geodesic triangle, respectively.

Definition 2.1. Let τ be an empty triangle of \mathcal{B} . Let l, l_1 and l_2 be edges of \mathcal{B} . We say that:

1. l_1 and l_2 *cross each other* (or l_1 *crosses* l_2) if $\text{relconv}(l_1) \cap \text{relconv}(l_2)$ is a single point. Equivalently, if (l_1, l_2) is a circuit of \mathcal{B} .
2. l *crosses* τ (or l *bisects* τ) if l crosses some edge of τ .

The following proofs are all done using either topological arguments or the language of oriented matroids, but avoiding geometric arguments. This allows us to show that all the results of this paper translate to non-realizable oriented matroids, which is done in the Appendix. Also, in some of the proofs the use of oriented matroids makes it evident that the case study involved is complete, which might not be obvious in a more geometric proof. Some basic concepts and facts of oriented matroid theory used are the following (see [4] or Chapter 6 of [19]):

- The compositions of circuits are called *vectors* and the compositions of cocircuits are called *covectors*, where the composition of (C^+, C^-) and (D^+, D^-) is by definition $(C^+ \cup (D^+ \setminus C^-), C^- \cup (D^- \setminus C^+))$.
- (C^+, C^-) is a vector if and only if C^+ and C^- are disjoint and the relative interiors $\text{relconv}(C^+)$ and $\text{relconv}(C^-)$ have a common point. Equivalently, if C^+ and C^- are the elements of \mathcal{B} with positive and negative coefficient respectively in some linear dependence among the elements of \mathcal{B} . (\emptyset, \emptyset) is a vector by convention.
- (C^+, C^-) is a covector if C^+ and C^- are the intersections with \mathcal{B} of the open half-spaces defined by some hyperplane. (\emptyset, \emptyset) is a covector by convention.
- Vectors and covectors are orthogonal to each other, where (C^+, C^-) and (D^+, D^-) are called orthogonal if $(C^+ \cap D^+) \cup (C^- \cap D^-)$ and $(C^+ \cap D^-) \cup (C^- \cap D^+)$ are either both empty or both non-empty.
- Even more, vectors are exactly the signed subsets (C^+, C^-) orthogonal to every cocircuit, and covectors are those orthogonal to every circuit.
- Given two vectors (C^+, C^-) and (D^+, D^-) (resp. two covectors) and an element $a \in C^+ \cap D^-$ there is a vector (resp. a covector) (E^+, E^-) with $C^+ \cap D^+ \subset E^+ \subset C^+ \cup D^+ \setminus \{a\}$ and $C^- \cap D^- \subset E^- \subset C^- \cup D^- \setminus \{a\}$. This is called the *elimination of a* in (C^+, C^-) and (D^+, D^-) .

Finally, we sometimes use the following notation for a vector, covector, circuit or cocircuit (C^+, C^-) when we are interested in a particular subset $\tau = \{a_1, \dots, a_n\} \subset \mathcal{B}$. We write a string $(\varepsilon_1, \dots, \varepsilon_n)$ of signs $\varepsilon_i \in \{+, 0, -\}$ meaning that $C^+ \cap \tau = \{a_i: \varepsilon_i = +\}$ and $C^- \cap \tau = \{a_i: \varepsilon_i = -\}$.

Lemma 2.2. *Let $\sigma = \{s, t, u\}$ be a triangle of \mathcal{B} and let $\{p, q\}$ be an edge of \mathcal{B} which crosses the edges $\{s, t\}$ and $\{t, u\}$. Suppose that p and u lie on opposite sides of $\{s, t\}$. Then q and s lie on opposite sides of $\{t, u\}$.*

Proof. In $\{p, q, s, t, u\}$ we have the circuits

$$(+, +, -, -, 0), \tag{1}$$

$$(-, -, 0, +, +) \tag{2}$$

and the cocircuit

$$(-, +, 0, 0, +). \tag{3}$$

Elimination of p between (1) and (2) gives the circuit $(0, *, -, *, +)$. Orthogonality with (3) implies that this circuit is $(0, -, -, *, +)$, which imply that q and s lie on opposite sides of $\{t, u\}$. □

Proposition 2.3. *Let $\tau = \{p, q, r\}$ be an empty triangle of \mathcal{B} and let $l = \{s, t\}$ be an empty edge of \mathcal{B} . The following conditions are equivalent:*

1. l crosses τ .
2. $\text{relconv}(\tau) \cap \text{relconv}(l) \neq \emptyset$.

Moreover, if l crosses τ , then, for every two vertices $a, b \in \tau$ which lie on opposite sides of l , l crosses $\{a, b\}$.

Proof. The equivalence $1 \Leftrightarrow 2$ is an obvious topological fact: if l_1 is an edge of τ which crosses l , since $\text{relconv}(l_1)$ is in the closure of $\text{relconv}(\tau)$ we have that $\text{relconv}(\tau) \cap \text{relconv}(l) \neq \emptyset$. Reciprocally, if $\text{relconv}(l)$ intersects the interior of the geodesic triangle $\text{relconv}(\tau)$ and since τ is empty, $\text{conv}(l)$ must intersect the boundary of the geodesic triangle $\text{conv}(\tau)$ in exactly two points and the two of them cannot be on the same edge. In particular, one of them must be in $\text{relconv}(l_1)$ for some edge l_1 of τ . This intersection point cannot be an end-point of l , since then τ would not be empty. Thus, $\text{relconv}(l) \cap \text{relconv}(l_1) \neq \emptyset$.

Let us see the “moreover.” Without loss of generality we suppose that l crosses $\{p, q\}$. Suppose that $a, b \in \tau$ lie on opposite sides of l . Without loss of generality, assume $a = q$ and $b = r$. Since l crosses τ , it must be $\text{relconv}(\tau) \cap \text{relconv}(l) \neq \emptyset$, and since s and t do not lie in $\text{relconv}(\tau)$, the boundary of $\text{conv}(\tau)$ must have another point x in common with $\text{conv}(l)$, apart from the one in $\text{relconv}(l) \cap \text{relconv}(\{p, q\})$. Since r is on one side of l , we have $x \neq r$ and thus x lies on the relative interior of either $\{p, r\}$ or $\{q, r\}$. Therefore, $x \in \text{relconv}(l)$ by emptiness of τ . If $x \in \text{relconv}(\{p, r\})$, then l crosses $\{p, r\}$ and we have that q and r are on the opposite side of l to which p is. This implies that $a = q$ and $b = r$ are on the same side of l , which contradicts the hypotheses. Thus, $x \in \text{relconv}(\{q, r\})$ and l crosses $\{a, b\}$. □

Proposition 2.4. *Let $\tau = \{p, q, r\}$ be an empty triangle of \mathcal{B} and let $l = \{s, t\}$ be any empty bisector of τ . Then either*

- (i) $l \cap \tau = \emptyset$ and l crosses exactly two edges of τ , or
- (ii) $l \cap \tau \neq \emptyset$ and l crosses exactly one edge of τ .

Proof. By Definition 2.1 we know that l crosses at least one edge of τ . We first show that l cannot cross the three edges of \mathcal{C} (this is obvious geometrically, but we include an oriented matroid proof for use in the Appendix). If it does, then the following are circuits supported on $\{p, q, r, s, t\}$:

$$(+, +, 0, -, -), \tag{4}$$

$$(0, +, +, -, -), \tag{5}$$

$$(+, 0, +, -, -). \tag{6}$$

Using (4) and (5) we eliminate q to obtain

$$(+, 0, -, *, *). \tag{7}$$

Using (6) and (7) we eliminate r to obtain $(+, 0, 0, *, *)$ and conclude that p, s and t are collinear. In the same way we can conclude that $q, r \in \text{span}\{s, t\}$ which implies that $\tau \subset \text{span}\{s, t\}$. This is obviously impossible. We conclude that l cannot cross the three edges of τ .

(ii) Suppose that $l \cap \tau \neq \emptyset$, say $t = r$. Then $l = \{r, s\}$ implies it cannot cross neither $\{p, r\}$ nor $\{q, r\}$, so l crosses $\{p, q\}$ and only $\{p, q\}$.

(i) Suppose $l \cap \tau = \emptyset$. Then the points p, q, r, s and t are distinct. We suppose that l crosses $\{p, q\}$ and no other edge of τ . The set $\{q, r, s, t\}$ contains the support of a circuit Z that satisfies $r \in \text{supp}(Z)$ (otherwise l would not cross $\{p, q\}$ since $(\{p, q\}, \{s, t\})$ would not be a circuit). Say $r \in Z^+$, then either $q \notin Z^+$ or l crosses $\{q, r\}$ (by Proposition 2.3) and we are done, so we assume $q \notin Z^+$. If l does not cross $\{p, r\}$ we can do the same reasoning with $\{p, r, s, t\}$ to get the following circuits (which we write as sign vectors on $\{p, q, r, s, t\}$):

$$(+, +, 0, -, -), \tag{8}$$

$$(0, a, +, *, *), \tag{9}$$

$$(b, 0, +, *, *), \tag{10}$$

where a and b are either 0 or $-$. Eliminating r with (9) and (10) we get $(-b, a, 0, *, *)$ which together with (8) leads us to a contradiction, except for the case in which $a = b = 0$ which we study separately.

If $a = b = 0$, then Z is of the form $(0, 0, +, *, *)$. Using simplicity of \mathcal{B} and emptiness of l the possibilities (up to exchange of the roles of s and t) are

$$(0, 0, +, +, 0), \tag{11}$$

$$(0, 0, +, +, -), \tag{12}$$

$$(0, 0, +, +, +). \tag{13}$$

By elimination of s between (11) (resp. (12)) and (8) we obtain $(+, +, +, 0, -)$, which is impossible by emptiness of τ . By elimination of s between (13) and (8) we obtain $(+, +, +, 0, *)$, where obviously $* \neq 0$. If $* = -$ we are in the previous case and if $* = +$ we eliminate t between $(+, +, +, 0, +)$ and (8) to obtain $(+, +, +, -, 0)$, which again is not possible by emptiness of τ . □

Lemma 2.5. *Let $l = \{p, q\}$, $l_1 = \{p_1, q_1\}$ and $l_2 = \{p_2, q_2\}$ be three empty edges of \mathcal{B} such that l_1 and l_2 cross l but l_1 and l_2 do not cross each other. Suppose that p_1 and p_2 are on the same side of l . Then at least one of the edges $\{p_1, q_2\}$ and $\{p_2, q_1\}$ crosses l .*

Proof. The cases $p_1 = p_2$ or $q_1 = q_2$ are trivial, so we suppose that the four points p_1, p_2, q_1 and q_2 are distinct. We consider the six edges defined by the four points $\{p_1, p_2, q_1, q_2\}$. We say that two edges overlap if their convex hulls intersect in more than one point.

Suppose first that two of the six edges overlap. This implies that for three of the points, say p_1, p_2 and q_1 , one of the three is in the relative interior of the edge formed by the other two. It is impossible that $q_1 \in \text{relconv}(p_1, p_2)$, since p_1 and p_2 lie on one

side of l and q_1 on the other. Any of the other two possibilities, $p_1 \in \text{relconv}(q_1, p_2)$ or $p_2 \in \text{relconv}(q_1, p_1)$ clearly implies that $\text{relconv}(q_1, p_2)$ intersects $\text{relconv}(l)$.

If no pair of edges overlap, then the only edges which can cross each other are the ones with disjoint end-points. However, $l_1 = \{p_1, q_1\}$ and $l_2 = \{p_2, q_2\}$ do not cross each other by hypothesis and $\{p_1, p_2\}$ and $\{q_1, q_2\}$ do not cross each other because they lie on opposite sides of l . So, only $\{p_1, q_2\}$ and $\{q_1, p_2\}$ can cross each other. We consider the two possibilities:

- If $\{p_1, q_2\}$ and $\{q_1, p_2\}$ cross each other, then we have the following three circuits among the points $\{p, q, p_1, q_1, p_2, q_2\}$:

$$(0, 0, +, -, -, +), \tag{14}$$

$$(+, +, 0, 0, -, -), \tag{15}$$

$$(+, +, -, -, 0, 0). \tag{16}$$

Eliminating q_2 between (14) and (15) we get the vector $(+, +, +, -, -, 0)$. Eliminating p_1 between this and (16) we get the vector $(+, +, 0, -, -, 0)$. Since $\{p, q, q_1, p_2\}$ has rank 3, this vector is a circuit. Thus, $\{p, q\}$ and $\{q_1, p_2\}$ cross each other.

- If $\{p_1, q_2\}$ and $\{q_1, p_2\}$ do not cross each other, then the six edges among the points $\{p_1, q_1, p_2, q_2\}$ form an embedded complete graph K_4 in the sphere, and K_4 has a unique embedding modulo topological equivalence. It is topologically obvious that since the geodesic segment $\text{conv}(l)$ crosses the edges $\text{conv}(p_1, q_1)$ and $\text{conv}(p_2, q_2)$ it must cross at least one of the remaining four edges. By hypothesis it cannot cross neither $\text{conv}(p_1, p_2)$ nor $\text{conv}(q_1, q_2)$ so we have finished. \square

Definition 2.6. Let $l = \{p, q\}$ be an empty edge of \mathcal{B} . Let $\Omega(l) = \{r: r \text{ is an edge of } \mathcal{B} \text{ which crosses } l\}$. We define the following partial ordering in $\Omega(l)$:

- For two edges $r = \{s, t\}$, $r' = \{t, u\}$ in $\Omega(l)$ with a common vertex t , we say that r' is closer to p than r and write $r <_p r'$ if u and p are on the same side of r . Equivalently, if the intersection point $\text{relconv}(r') \cap \text{relconv}(l)$ is closer to p (along $\text{conv}(l)$) than the point $\text{relconv}(r) \cap \text{relconv}(l)$.
- For arbitrary edges in $\Omega(l)$, we say that $r <_p r'$ if there is a chain $r = r_1, r_2, \dots, r_m = r'$ of edges in $\Omega(l)$ with $r_1 <_p r_2 <_p \dots <_p r_m$, where r_i and r_{i+1} share a vertex for every $i \in \{1, \dots, m - 1\}$.

Recall that an oriented edge l^+ of \mathcal{B} denotes an edge $l = \{p, q\}$ together with the choice of one of the two half-spaces (or hemispheres) defined by it.

Corollary 2.7. Let $\tau = \{p, q, r\}$ be an empty triangle of \mathcal{B} and let m and n be two bisectors of τ crossing $l = \{p, q\}$. We give m and n an orientation such that $p \in m^+ \cap n^+$. Suppose that $m <_p n$ in $\Omega(l)$ and that $r \notin m^+$. Then $r \notin n^+$.

Proof. We can assume that m and n are empty, because otherwise their convex hulls contain empty edges with the conditions required in the corollary. It is sufficient to prove

the result in the case that m and n have a common vertex and then the general case holds recursively. Say $m = \{s, t\}$, $n = \{t, u\}$. Then $p, u \in m^+$. Suppose r was in n^+ . Since n crosses $\{p, q\}$ and $p \in n^+$, we have $q \in n^-$. By Proposition 2.3, n crosses $\{q, r\}$. Then:

- If $r \notin m^-$, then $r \notin m^+$ and $r \notin m^-$, so m crosses only the edge l of τ . By Proposition 2.4, $r \in m$. It must be $r = s$, but then $s \in n^+$ and we have $n <_p m <_p n$ which is impossible.
- If $r \in m^-$, remember that $n = \{t, u\}$ crosses $\{p, q\}$ and $\{q, r\}$. Hence we have the (restricted) cocircuit $(+, -, +, *, 0, 0)$ on $\{p, q, r, s, t, u\}$. On the other hand, in this case m crosses $\{p, r\}$, so we have the circuit $(+, 0, +, -, -, 0)$. The circuit and the cocircuit must be orthogonal, so $* = +$ and we have the cocircuit

$$(+, -, +, +, 0, 0).$$

This means that r and s are on the same side of n , that is, $r, s \in n^+$. In particular, $s \in n^+$ implies that $n <_p m <_p n$ which is impossible. □

3. Main Results

Throughout this section \mathcal{C} is a virtual chamber of a simple rank 3 vector configuration \mathcal{B} .

3.1. Every Triangulation has a Flip

Definition 3.1. Let l be an edge of \mathcal{B} and consider it oriented.

- (i) We say that l has \mathcal{C} on its positive side l^+ (or that \mathcal{C} lies on l^+) if there exists $\tau \in \mathcal{C}$ such that $\tau \subset l^+ \cup l^0$ (same for l^-).
- (ii) We say that the orientation of l is \mathcal{C} -coherent if \mathcal{C} lies on l^+ .

By Lemma 1.4, an edge cannot have \mathcal{C} on both sides. However, not every edge of \mathcal{B} has \mathcal{C} on one side. For example, let $\mathcal{B} = \{p, q, r, s\}$ with $\{p, q, r\}$ being a triangle and $s \in \text{relconv}(p, q, r)$. Then $\mathcal{C} = \{\{p, q, r\}, \{p, q, s\}\}$ is a virtual chamber and $\{r, s\}$ does not have \mathcal{C} on any side.

Note that for any triangle $\tau = \{p, q, r\}$ of a virtual chamber \mathcal{C} and for any edge l (say $l = \{p, q\}$) of τ , l has \mathcal{C} on the side on which r is. This implies that there exists the \mathcal{C} -coherent orientation for l and that $r \in l^+$ for this orientation. The following proposition says that every edge which crosses an empty triangle of \mathcal{C} can also be given a \mathcal{C} -coherent orientation.

Proposition 3.2. Let \mathcal{C} be a virtual chamber of \mathcal{B} , let $\tau = \{p, q, r\} \in \mathcal{C}$ be an empty triangle and let $l = \{s, t\}$ be a bisector of τ . Then l has \mathcal{C} on one of its sides.

Proof. Let $\{p, q\}$ be an edge of τ that is crossed by l (note that $\tau \cap l$ could be non-empty, $r \in \{s, t\}$ could occur). Then p and q are on opposite sides of l , so there is at least one

vertex of τ on each side of l . Moreover, at least one side of l contains exactly one vertex (say p) of τ , because τ has only three vertices. We consider the triangle $\sigma = \{p, s, t\}$ of \mathcal{B} . We can extend σ to a triangulation \mathcal{T}_0 of $\{p, q, r, s, t\}$. We claim that for every triangle of \mathcal{T}_0 there is one side of l containing none of its vertices.

Observe that if a triangle ρ of \mathcal{T}_0 does not verify this property, then it must contain p and exactly one element of $\{q, r\}$ (ρ cannot be $\{p, q, r\} = \tau$ because τ and σ do not intersect properly, since an edge l of σ crosses an edge $\{p, q\}$ of τ). The condition “ l crosses $\{p, q\}$ ” implies that p and q cannot both be in ρ if we want ρ and σ to intersect properly, so we can assume (without loss of generality) $\rho = \{p, r, s\}$. Now, $\{p, r, s, t\}$ contains the support of a unique circuit Z and, by hypothesis, p and r lie on opposite sides of $l = \{s, t\}$, so $\{p, r\} \subset Z^+$ (up to sign reversal in Z). By Proposition 2.3 we have that l crosses $\{p, r\}$, and hence ρ and σ do not intersect properly. This means that such a triangle ρ cannot exist.

Now we extend \mathcal{T}_0 to a triangulation \mathcal{T} of \mathcal{B} and use condition 1 of Definition 1.1 to conclude that there exists exactly one triangle $\rho \in \mathcal{T} \cap \mathcal{C}$. By Lemma 1.4 and the fact that \mathcal{T}_0 covers $\text{conv}(\tau)$, we have $\rho \in \mathcal{T}_0$. □

The following result is crucial for what follows. In it, l is an edge of an empty triangle τ of the virtual chamber \mathcal{C} and $\Omega_{\mathcal{C}}(l, p)$ is the subset of $\Omega(l)$ (see Definition 2.6) consisting of the edges which cross l and have p and \mathcal{C} on the same side. Clearly, $\Omega_{\mathcal{C}}(l, p)$ inherits the partial order $<_p$ from $\Omega(l)$. For the proof that every virtual chamber has an empty triangle see Corollary 3.5.

Theorem 3.3. *Let $\tau = \{p, q, r\}$ be an empty triangle of the virtual chamber \mathcal{C} . Let $l = \{p, q\}$ and let $\{a, b\}$ be a maximal element in the poset $(\Omega_{\mathcal{C}}(l, p), <_p)$. If $\{p, a, b\} \in \mathcal{C}$, then $\{a, b\}$ supports a flip of \mathcal{C} .*

Proof. Let $m = \{a, b\}$ and consider it oriented \mathcal{C} -coherently by Proposition 3.2. We have to prove that the two conditions on m of Definition 1.10 are satisfied for the edge m . The first condition is obviously satisfied since $p \in m^+$ and $q \in m^-$. In order to prove the second condition let $s \in m^+$. We need to prove that $\{a, b, s\} \in \mathcal{C}$.

Without loss of generality we can assume that:

- $\{a, b, p, q, s\}$ are five distinct points: The first four are distinct since $\{a, b\}$ crosses $\{p, q\}$. The point s is trivially not equal to a nor b since $s \in m^+ = \{a, b\}^+$. Also, $s \neq q$ since $s \in m^+$ and $q \in m^-$. Finally, if $s = p$, then the claim $\{a, b, s\} \in \mathcal{C}$ is the hypothesis $\{a, b, p\} \in \mathcal{C}$.
- $\{a, s\}$ and $\{b, s\}$ are empty edges: if, for example, $\{a, s\}$ is not empty, then let $s' \in \text{relconv}(\{a, s\})$ such that $\{a, s'\}$ is empty. Clearly, $s \in \{a, b\}^+$ implies $s' \in \{a, b\}^+$ and then if we prove $\{a, b, s'\} \in \mathcal{C}$ we will have $\{a, b, s\} \in \mathcal{C}$ by condition 2 of Definition 1.1 applied to the circuit $(\{s'\}, \{a, s\})$.

The sets $\{a, b, p\}$ and $\{a, b, s\}$ are independent and we have $p, s \in m^+$. This implies that $\rho = \{a, b, s, p\}$ contains the support of a circuit Z in which p and s have opposite and non-zero signs. We suppose $p \in Z^+$ and $s \in Z^-$. The possibilities for Z (written

on $\{a, b, s, p, q\}$) are the following, up to exchange of the roles of a and b :

$$(0, -, -, +, 0), \tag{17}$$

$$(-, -, -, +, 0), \tag{18}$$

$$(0, +, -, +, 0), \tag{19}$$

$$(-, +, -, +, 0), \tag{20}$$

$$(+, +, -, +, 0). \tag{21}$$

Since $\sigma := \{p, a, b\}$ is in $\mathcal{T}^-(\rho) \cap \mathcal{C}$, condition 2 of Definition 1.1 implies there is a (unique) triangle $\sigma' \in \mathcal{T}^+(\rho) \cap \mathcal{C}$. We claim that $\sigma' = \{a, b, s\}$ and this finishes the proof. We study cases (17)–(21) separately.

Cases (17) and (18) are trivial because the unique point with positive sign is p , so $\mathcal{T}^+(\rho) = \{\{a, b, s\}\}$. For the remaining cases, remember that m crosses l , so we have the circuit

$$(-, -, 0, +, +). \tag{22}$$

In case (19) (resp. case (20)) we eliminate b between (19) and (22) (resp. between (20) and (22)) and obtain the vector

$$(-, 0, -, +, +). \tag{23}$$

We use the following fact repeatedly: if a vector has support contained in four points, three of which are independent, then it is a circuit. This is so because otherwise it is a composition of at least two different circuits, and there is only one circuit with support contained in a spanning set with four points. In particular, (23) is a circuit and $\{a, s\}$ crosses $l = \{p, q\}$. To prove that $\sigma' = \{a, b, s\}$ we have to show that the other triangle of $\mathcal{T}^+(\rho)$, $\{a, p, s\}$, is not in \mathcal{C} . If it was, then $p \in \{a, s\}^+$. Since $\{a, s\}$ crosses l we would have that $\{a, b\} <_p \{a, s\}$ which contradicts the hypothesis of maximality of $m = \{a, b\}$. Thus $\sigma' = \{a, b, s\}$.

Only case (21) remains. We can eliminate b between (21) and (22) to get a circuit $(*, 0, -, +, +)$ where “*” cannot be zero by emptiness of $\tau = \{p, q, r\}$. Without loss of generality we can assume that “*” is a minus sign, because if it is a plus sign we eliminate a between $(+, 0, -, +, +)$ and (22) to get $(0, -, -, +, +)$ and we repeat the following argument exchanging the roles of a and b . Then the circuit becomes

$$(-, 0, -, +, +) \tag{24}$$

and the same argument as in case (19) shows that $\{a, p, s\} \notin \mathcal{C}$. The rest of the proof is devoted to showing that $\{b, p, s\} \notin \mathcal{C}$. For this, we suppose that $\{b, p, s\} \in \mathcal{C}$ and get a contradiction.

We eliminate a between (21) and (24) and we get the circuit

$$(0, +, -, +, +), \tag{25}$$

which implies that s and b are on the same side of l . If $s, b \in l^-$ (where we consider l oriented \mathcal{C} -coherently, so that $r \in l^+$) we have $\{b, p, s\} \cap l^+ = \emptyset$, hence $\{b, p, s\} \notin \mathcal{C}$. So we suppose that $s, b \in l^+$; then $a \in l^-$, because l crosses $\{a, b\}$.

Without loss of generality we can assume that $r \notin \{a, b, p, q, s\}$: the point r cannot be the same as a, p or q because $l = \{p, q\}$, $r \in l^+$ and $a \in l^-$. Circuit (25) implies $r \neq b$, or $\{p, q, r\}$ would not be empty. The same circuit implies that if $s = r$, then our claim that $\{b, p, s\} \notin \mathcal{C}$ is obvious, since $\{p, q, s\} = \{p, q, r\}$ is in \mathcal{C} and b and q lie on opposite sides of $\{p, s\}$.

So, in the following we write all the circuits on $\{a, b, s, p, q, r\}$. Circuits (21), (22), (24) and (25) become respectively

$$(+, +, -, +, 0, 0), \quad (26)$$

$$(-, -, 0, +, +, 0), \quad (27)$$

$$(-, 0, -, +, +, 0), \quad (28)$$

$$(0, +, -, +, +, 0) \quad (29)$$

when written on $\{a, b, s, p, q, r\}$. Recall also that we know $r, b, s \in \{p, q\}^+$ and $a \in \{p, q\}^-$. In other words, we have a cocircuit which, restricted to $\{a, b, s, p, q, r\}$, is

$$(-, +, +, 0, 0, +). \quad (30)$$

We now look at the cocircuit $(*, +, 0, 0, *, *)$ vanishing on $\{p, s\}$ and oriented \mathcal{C} -coherently. Orthogonality with (26) and (28) implies it is $(-, +, 0, 0, -, *)$, so it only remains to know the sign of r on that cocircuit. We know that $\text{relconv}(b, p, s) \cap \text{relconv}(p, q, r) \neq \emptyset$ by Lemma 1.4, since both simplices are in the virtual chamber \mathcal{C} . This implies (since $(-, +, 0, 0, -, *)$ implies $q \in \{p, s\}^-$) that $r \in \{p, s\}^+$. The cocircuit is then

$$(-, +, 0, 0, -, +). \quad (31)$$

By (28), together with the assumption that $\{a, s\}$ is empty, $\{a, s\}$ is an empty edge which crosses the triangle $\{p, q, r\}$. Proposition 2.4 implies $\{a, s\}$ crosses either $\{p, r\}$ or $\{q, r\}$. Cocircuit (31) implies it does not cross $\{p, r\}$, hence it crosses $\{q, r\}$ (in particular, p and r are on the same side of $\{a, s\}$) and we have the circuit

$$(+, 0, +, 0, -, -), \quad (32)$$

which together with (28) gives us a circuit $(0, 0, *, +, *, -)$. Orthogonality with (30) and (31) implies that this circuit is

$$(0, 0, +, +, -, -), \quad (33)$$

which implies that s and p lie on opposite sides of $\{q, r\}$, and, since $\{p, q, r\} = \tau \in \mathcal{C}$, we have $s \in \{q, r\}^-$.

Now we observe that p and r are on the same side of $\{a, s\}$, since both $\{p, q\}$ and $\{q, r\}$ cross $\{a, s\}$. Eliminating q and p respectively from (29) and (33) we derive the circuit

$$(0, +, x, +, 0, -) \quad (34)$$

and another circuit $(0, +, -, 0, +, +)$. The last one implies that b and s are on the same side of $\{r, q\}$. We know that $s \in \{r, q\}^-$, so $b, s \in \{r, q\}^-$. Hence $\{b, s, r\} \notin \mathcal{C}$.

Since $\{b, p, s\}$ is in the negative triangulation of (34) and we assume $\{b, p, s\}$ to be in \mathcal{C} , there must be exactly one triangle of the positive triangulation of (34) in \mathcal{C} . The triangles in the positive triangulation of (34) are $\{b, s, r\}$, $\{p, s, r\}$ and (only if $x = +$) $\{b, p, r\}$. We have already discarded $\{b, s, r\} \in \mathcal{C}$. Let us see how to discard the other two:

- $\{b, p, r\} \in \mathcal{C}$ and $x = +$ is impossible because (34) with $x = +$ implies that b and s are on opposite sides of $\{p, r\}$, $\{p, q, r\} \in \mathcal{C}$ and $\{b, p, r\} \in \mathcal{C}$ implies b and q are on the same side of $\{p, r\}$, while (33) implies that s and q are on the same side of $\{p, r\}$.
- Suppose finally that $\{p, s, r\} \in \mathcal{C}$. Since p and r are on the same side of $\{a, s\}$, $\{a, s\}$ has \mathcal{C} on one side, namely that in which p and r are. However, this means that $\{a, s\} \in \Omega_{\mathcal{C}}(l, p)$. This contradicts the maximality of $m = \{a, b\}$ in $\Omega_{\mathcal{C}}(l, p)$ since $m = \{a, b\} <_p \{a, s\}$ (s and p are on the same side of $\{a, b\}$, by circuit (26)). □

Proposition 3.4. *Let \mathcal{C} be a virtual chamber of \mathcal{B} which is not a (geometric) chamber. Let τ be an empty triangle of \mathcal{C} . Then there exist an edge l of τ , a vertex p of l and a maximal element $m = \{a, b\}$ of $\Omega_{\mathcal{C}}(l, p)$ such that $\{a, b, p\} \in \mathcal{C}$.*

Proof. Say $\tau = \{p, q, r\}$. There must be another triangle $\sigma = \{s, t, u\} \in \mathcal{C}$ such that τ is not contained in $\text{conv}(\sigma)$ because otherwise \mathcal{C} would be a chamber (just take v generic in $\text{relconv}(\tau)$ and consider $\mathcal{C}_{v, \mathcal{B}}$), and we can assume σ to be empty as well using condition 2 of Definition 1.1 iteratively. Since $\text{relconv}(\tau) \cap \text{relconv}(\sigma) \neq \emptyset$ and τ and σ are both empty, some edge of τ (say $\{p, q\}$) crosses some edge of σ (say $\{s, t\}$).

We now claim that there is an empty edge of \mathcal{B} that crosses τ , one of whose vertices is in τ . Suppose $\{s, t\} \cap \tau = \emptyset$. Then $\{s, t\}$ crosses another edge of τ (say $\{p, r\}$) by Proposition 2.4. We have the following circuits on $\{p, q, r, s, t\}$:

$$(+, +, 0, -, -), \tag{35}$$

$$(+, 0, +, -, -). \tag{36}$$

Since $\{s, t\}$ crosses $\{p, q\}$ we can assume without loss of generality that $s \in \{p, q\}^-$, so there is a circuit of the form

$$(*, *, -, -, 0). \tag{37}$$

Eliminating t between (35) and (36) we get $(*, +, -, *, 0)$, which combined with (37) implies

$$(x, +, -, -, 0). \tag{38}$$

By Lemma 2.2 $t \in \{p, r\}^-$, so we can repeat this argument to obtain

$$(y, -, +, 0, -). \tag{39}$$

Now suppose $x \neq + \neq y$. Then, eliminating q between (38) and (39), we get $(z, 0, *, -, -)$ with $z \neq +$. This contradicts (36), thus we can assume $x = +$, but then (38) becomes

$(+, +, -, -, 0)$. This implies that $\{r, s\}$ crosses $\{p, q\}$ as we claimed. Hence, we can assume that $t = r$.

Now we take a maximal element $m = \{a, b\}$ of $\Omega_{\mathcal{C}}(\{p, q\}, p)$ with $\{r, s\} \leq_p m$. Since $r \notin \{r, s\}^+$, by Corollary 2.7, we have that $r \notin m^+$. Since $q \in m^-$, m cannot cross $\{q, r\}$, so, by Proposition 2.4, m crosses $\{p, r\}$ if and only if $r \notin \{a, b\}$. We conclude that m crosses $\{p, x\}$ for every $x \in \tau \setminus \{a, b, p\} \subset \{q, r\}$. Now we extend the triangle $\{a, b, p\}$ to a triangulation \mathcal{T} of $\text{conv}(a, b, p, q, r)$ and \mathcal{T} to a triangulation \mathcal{T}' of \mathcal{B} . By definition of a virtual chamber there is exactly one triangle σ' of \mathcal{T}' in \mathcal{C} , but $\text{relconv}(\sigma') \cap \text{relconv}(\tau)$ must be non-empty by Lemma 1.4, so $\sigma' \in \mathcal{T}$. On the other hand, σ' must have a vertex in m^+ , so $p \in \sigma'$. Let x be another vertex of σ' . If $x \in \tau \setminus \{a, b, p\}$, then $\{p, x\}$ crosses m , which is an edge of $\{a, b, p\} \in \mathcal{T}$ and hence σ' and τ do not intersect properly. We conclude that $x \notin \tau \setminus \{a, b, p\}$, so $x \in \{a, b\}$. This implies that $\sigma' = \{a, b, p\}$. \square

Corollary 3.5. *Unless \mathcal{B} has only three elements, for every virtual chamber \mathcal{C} of \mathcal{B} there is an empty edge supporting a flip of \mathcal{C} .*

Proof. If \mathcal{C} is a chamber, the edges supporting its boundary support flips of \mathcal{C} , unless they are in the boundary of $\text{conv}(\mathcal{B})$. However, the only case of a chamber with all its edges on the boundary of $\text{conv}(\mathcal{B})$ is that of the unique chamber of \mathcal{B} if \mathcal{B} has three elements. For virtual chambers which are not chambers, the corollary follows immediately from Theorem 3.3 and Proposition 3.4, taking into account that an empty triangle τ in \mathcal{C} can always be obtained as the unique triangle of \mathcal{C} in a triangulation of \mathcal{B} by empty triangles (i.e. which uses all the elements of \mathcal{B}). \square

3.2. 3-Connectivity of the Graph of Virtual Chambers

Lemma 3.6. *Let $m = \{p, q\}$ be an empty edge of \mathcal{B} supporting a flip of a virtual chamber \mathcal{C} and let σ be a triangle of \mathcal{C} . Then either $m \subset \text{conv}(\sigma)$ or some edge of σ crosses m .*

Proof. Suppose m crosses no edge of σ . Consider an empty triangle $\tau \in \mathcal{C}$ of the form $\tau = \{p, q, r\}$ with $r \in m^+$ (which always exists). Then $\text{relconv}(\tau) \cap \text{relconv}(\sigma) \neq \emptyset$ by Lemma 1.4. Since τ is empty, this implies that either $\tau \subset \text{conv}(\sigma)$ (and we are done) or some edge of σ crosses τ , so we suppose that, for instance, $\{s, t\}$ crosses some edge of τ (which must be different from m). Say $\{s, t\}$ crosses $\{p, r\}$. If $\{s, t\}$ is not empty, then $\text{conv}(\{s, t\})$ contains an empty edge $\{s', t'\}$ of \mathcal{B} which crosses $\{p, r\}$. By Proposition 2.4, either $\{s', t'\}$ crosses $\{q, r\}$ or $q \in \{s', t'\}$. Anyway p and q lie on the same closed side of $\{s', t'\}$, that is, on the same closed side of $\{s, t\}$.

We have the circuit $(+, +, -, -)$ written on $\{p, r, s, t\}$ and, on the other hand, the cocircuit which vanishes on $\{p, q\}$ is $(0, +, *, *)$ when restricted to the same set. Orthogonality between both implies that one of the asterisks is a plus sign. Say this restricted cocircuit is $(0, +, +, *)$. Then $s \in m^+$, thus $\{p, q, s\} \in \mathcal{C}$. As a consequence, $\{p, q\}$ is not contained in the negative closed side of $\{s, t\}$, and hence it is contained in the

positive one. Therefore, the spherical region $D = \text{conv}(\tau) \cap \overline{\{s, t\}^+}$ contains m and its boundary is the union of the arcs $\text{conv}(m), [p, x] \subset \text{conv}(\{p, r\}), [q, y] \subset \text{conv}(\{q, r\})$ and $[x, y] \subset \text{conv}(\{s, t\})$, where x and y are the points $\text{conv}(\{p, r\}) \cap \text{conv}(\{s, t\})$ and $\text{conv}(\{q, r\}) \cap \text{conv}(\{s, t\})$ respectively, and the bracket notation has the standard meaning. Clearly, $\text{relconv}(\tau) \cap \text{relconv}(\sigma) \subset D \subset \text{conv}(\tau)$, and hence D is homeomorphic to a closed 2-ball.

Suppose $m \not\subset \text{conv}(\sigma)$. Then the convex hull of some other edge (say $\{t, u\}$) of σ intersects the boundary of D transversally. Moreover, the facts that τ is empty, $D \subset \text{conv}(\tau)$ and $\{t, u\}$ does not cross m imply that $\{t, u\}$ crosses the arcs $[p, x]$ and $[q, y]$ and for some of them (say $[p, x]$) the intersection is none of its end-points (otherwise the intersection points would be p and q necessarily and hence $m \subset \text{conv}(\sigma)$). Therefore, $\{t, u\}$ crosses $\{p, r\}$. With the same argument as above we can conclude that p and q are on the same closed side of $\{t, u\}$ and one of t and u is in m^+ . Say $t \in m^+$. Then $\{p, q, t\} \in \mathcal{C}$.

Since $p \in \{s, t\}^+$ and $\{s, t\}$ crosses $\{p, r\}$ we have $r \in \{s, t\}^-$. The facts that $\{s, t\}$ and $\{t, u\}$ cross $\{p, r\}$ and $r \in \{s, t\}^-$ imply $p \in \{t, u\}^-$ by Lemma 2.2. Therefore, p and q are on the negative closed side of $\{t, u\}$ and hence $\{p, q, t\} \cap \{t, u\}^+ = \emptyset$, which is absurd. \square

Lemma 3.7. *Let $m = \{a_1, a_2\}$ be an empty edge of \mathcal{B} supporting a flip of a virtual chamber \mathcal{C} and suppose that $\Omega_{\mathcal{C}}(m, a_2) = \emptyset$. Then \mathcal{C} is a flag chamber (in particular a chamber). More precisely, let $a_3 \in m^+$ (so $\tau = \{a_1, a_2, a_3\} \in \mathcal{C}$) and let p be the new vector in the lexicographic extension of \mathcal{B} at the basis $[a_1, a_2, a_3]$. Then $p \in \text{relconv}(\sigma)$ for every $\sigma \in \mathcal{C}$.*

Proof. If $\sigma = \{s, t, u\} \in \mathcal{C}$ is not an empty triangle, there are a point $q \in \mathcal{B}$ and a circuit $Z = (Z^+, Z^-)$ supported on $\rho = \{s, t, u, q\}$ such that $Z^- = \{q\}$ and $Z^+ \neq \emptyset$, so $\mathcal{T}^-(\rho) = \{\sigma\}$. Then the unique triangle in $\mathcal{T}^+(\rho) \cap \mathcal{C}$ is contained in $\text{conv}(\sigma)$. We can repeat this process until we get an empty triangle $\sigma' \in \mathcal{C}$ with $\sigma' \subset \text{conv}(\sigma)$ which implies $\text{relconv}(\sigma') \subset \text{relconv}(\sigma)$.

Thus, without loss of generality we assume that σ is empty.

We have to show that, for any edge of σ , p is on its positive side. We call C_1, C_2 and C_3 the (\mathcal{C} -coherently oriented) cocircuits which vanish in $\{s, t\}, \{t, u\}$ and $\{s, u\}$ respectively and let $j_k = \min\{i: C_k(a_i) \neq 0\}$ for $k = 1, 2, 3$. Then we have to show that $C_k(a_{j_k}) = +$ for $k = 1, 2, 3$. First assume m crosses some edge (say $\{s, t\}$) of σ . Since $\Omega_{\mathcal{C}}(m, a_2) = \emptyset$, we have $a_1 \in \{s, t\}^+$. Thus, $j_1 = 1$ and $C_1(a_1) = +$. No other edge of σ crosses m since clearly it would be in $\Omega_{\mathcal{C}}(m, a_2)$ (as a straightforward consequence of Lemma 2.2), thus $a_1 = u$ by Proposition 2.4. The cocircuit C_2 restricted to $\{a_1, a_2, s, t, u\}$ is $(0, *, +, 0)$, which must be orthogonal to the circuit $(+, +, -, -)$, and hence C_2 becomes $(0, +, +, 0)$. Thus, $j_2 = 2$ and $C_2(a_2) = +$. The same argument proves that $j_3 = 2$ and $C_3(a_2) = +$.

Now assume m crosses no edge of σ . Let l be an edge of σ and let C be the cocircuit vanishing in l , which we assume oriented \mathcal{C} -coherently. Let $j = \min\{i: C(a_i) \neq 0\}$. By Lemma 3.6, $m \subset \text{conv}(\sigma)$, so, if $a_1 \notin l$, then $j = 1$ and $C(a_1) = +$. If $a_1 \in l$ but $a_2 \notin l$, the same argument works with $j = 2$. Finally, if $m = l$, the same argument is valid with $j = 3$. \square

Lemma 3.8. *Let \mathcal{C} be a virtual chamber of \mathcal{B} and let $m = \{a_1, a_2\}$ be an empty edge supporting a flip of \mathcal{C} . If both $\Omega_{\mathcal{C}}(m, a_1)$ and $\Omega_{\mathcal{C}}(m, a_2)$ are nonempty, then there is a virtual chamber \mathcal{C}' of \mathcal{B} such that:*

- m supports a flip of \mathcal{C}' ,
- $\Omega_{\mathcal{C}'}(m, a_1)$ is strictly contained in $\Omega_{\mathcal{C}}(m, a_1)$,
- $\Omega_{\mathcal{C}'}(m, a_2)$ strictly contains $\Omega_{\mathcal{C}}(m, a_2)$.

Proof. Let $l = \{p, q\}$ be a maximal element of $\Omega_{\mathcal{C}}(m, a_1)$. We claim that $\{p, q, a_1\} \in \mathcal{C}$. Without loss of generality assume that $p \in m^+$. Since m supports a flip of \mathcal{C} , by Definition 1.10 the triangle $\{a_1, a_2, p\}$ is in \mathcal{C} . The spanning set $\rho = \{a_1, a_2, p, q\}$ is the support of the circuit $Z = (\{a_1, a_2\}, \{p, q\})$ since l crosses m , and $\{a_1, a_2, p\} \in \mathcal{C} \cap \mathcal{T}^-(\rho)$, so there is a triangle in $\mathcal{C} \cap \mathcal{T}^+(\rho)$ by Definition 1.1. The triangle $\{p, q, a_2\}$ is not in \mathcal{C} since $a_2 \in \{p, q\}^-$, so $\{p, q, a_1\} \in \mathcal{C}$. By Theorem 3.3, l supports a flip of \mathcal{C} , so let \mathcal{C}' be the virtual chamber which is obtained by this flip. By Corollary 1.11, the triangles of \mathcal{C} of the form $\{a_1, a_2, s\}$ have not been removed when passing to \mathcal{C}' , so the edge m still supports a flip of \mathcal{C}' by Definition 1.10. Since l is the unique element of $\Omega_{\mathcal{C}}(m, a_1)$ whose orientation has been changed, we conclude that $\Omega_{\mathcal{C}'}(m, a_1)$ is strictly contained in $\Omega_{\mathcal{C}}(m, a_1)$ and $\Omega_{\mathcal{C}'}(m, a_2)$ strictly contains $\Omega_{\mathcal{C}}(m, a_2)$. □

We recall that $G(\mathcal{A})$ denotes the graph of triangulations of \mathcal{A} . In the next theorem we do not distinguish between triangulations of \mathcal{A} and virtual chambers of its Gale transform \mathcal{B} .

Theorem 3.9. *Let \mathcal{C} be a virtual chamber of a rank 3 vector configuration \mathcal{B} and let \mathcal{A} be the Gale transform of \mathcal{B} . If \mathcal{C} is not a chamber, then there are three vertex-disjoint paths in $G(\mathcal{A})$ joining \mathcal{C} to three distinct flag chambers.*

Proof. Since \mathcal{C} is not a chamber, by Corollary 3.5 there is an empty edge $m = \{a_1, a_2\}$ supporting a flip of \mathcal{C} . If $\Omega_{\mathcal{C}}(m, a_1) = \emptyset$ or $\Omega_{\mathcal{C}}(m, a_2) = \emptyset$, then \mathcal{C} would be a chamber by Lemma 3.7, so $\Omega_{\mathcal{C}}(m, a_1) \neq \emptyset \neq \Omega_{\mathcal{C}}(m, a_2)$.

Set m^+ as the side of m in which \mathcal{C} lies. We successively apply the previous lemma to obtain a path from \mathcal{C} to a virtual chamber \mathcal{C}_1 (in which every virtual chamber \mathcal{C}' involved lies on m^+ and has $\Omega_{\mathcal{C}'}(m, a_1)$ strictly contained in $\Omega_{\mathcal{C}}(m, a_1)$) with $\Omega_{\mathcal{C}_1}(m, a_1) = \emptyset$ and such that m supports a flip of \mathcal{C}_1 . By Lemma 3.7 (exchanging the roles of a_1 and a_2), \mathcal{C}_1 is a flag chamber of \mathcal{B} . Also we get a second path to a flag chamber \mathcal{C}_2 in which every virtual chamber \mathcal{C}' involved lies on m^+ and has $\Omega_{\mathcal{C}'}(m, a_1)$ strictly containing $\Omega_{\mathcal{C}}(m, a_1)$. Finally, performing the flip of \mathcal{C} supported on m we get a virtual chamber which lies on m^- . From it, the previous lemma produces a path to a flag chamber \mathcal{C}_3 of \mathcal{B} in which all intermediate virtual chambers lie on m^- .

These three paths are obviously vertex-disjoint. □

Corollary 3.10. *For any vector configuration \mathcal{A} with $d + 4$ elements and rank $d + 1$ the graph $G(\mathcal{A})$ of triangulations of \mathcal{A} is connected. If, moreover, \mathcal{A} is acyclic (or if \mathcal{A} is a spanning point configuration in \mathbb{R}^d), then $G(\mathcal{A})$ is 3-connected. In particular, every triangulation of \mathcal{A} has at least three geometric bistellar neighbours.*

Proof. Let \mathcal{B} be the Gale transform of \mathcal{A} (which we regard as a point configuration in S^2). Since a common edge of the chamber complex between two chambers is a flip, the subgraph $G_{\text{reg}}(\mathcal{A})$ of $G(\mathcal{A})$ induced by chambers of \mathcal{B} is connected and, if $\text{conv}(\mathcal{B}) = S^2$ (which is equivalent to \mathcal{A} being acyclic), 3-connected. Both results follow also from the theory of secondary polyhedra: the subgraph $G_{\text{reg}}(\mathcal{A})$ is the 1-skeleton of a three-dimensional convex polyhedron, and of a 3-polytope if \mathcal{A} is acyclic; see [1] and [2].

With this, Theorem 3.9 implies that $G(\mathcal{A})$ is connected and, using Lemma 3.11 below, 3-connected if \mathcal{A} is acyclic. \square

Lemma 3.11. *Let H be a k -connected subgraph of a graph G such that every vertex of $G \setminus H$ can be joined by k vertex-disjoint paths in G to k distinct vertices of H . Then G is k -connected.*

Proof. Suppose we remove $k - 1$ vertices of G and their incident edges. The k -connectivity of H implies that H remains connected after the removal. On the other hand, after the removal, for every vertex v in $G \setminus H$ that has not been removed there is at least one path joining it to one of the remaining vertices of H , since before the removal there were at least k vertex-disjoint paths joining v to k distinct vertices of H . \square

The following example exhibits the three different ways of going from a non-geometric virtual chamber to three geometric chambers.

Example 3.12. Let \mathcal{A} and \mathcal{B} be as in Example 1.5. Let \mathcal{C} be the virtual chamber of \mathcal{B} shown in the same example. In Fig. 2(a) we show an affine Gale diagram of \mathcal{A} in which the crucial edges of \mathcal{B} are oriented in the \mathcal{C} -coherent way.

The empty triangle $\{1, 2, 3\}$ belongs to \mathcal{C} . Since the edge $\{1, 4\}$ is maximal in $\Omega_{\mathcal{C}}(\{2, 3\}, 2)$ and the triangle $\{1, 2, 4\}$ is in \mathcal{C} , $\{1, 4\}$ supports a flip of \mathcal{C} . The edge $\{3, 6\}$ is maximal in $\Omega_{\mathcal{C}}(\{1, 4\}, 1)$. We first flip on $\{3, 6\}$ which makes $\{2, 3\}$ maximal. Finally we flip on $\{2, 3\}$ to get the geometric chamber (in fact a flag one) \mathcal{C}_1 which is represented in Fig. 2(b) by the shaded region (which is connected in the 2-sphere). On the other hand, $\{2, 5\}$ is maximal in $\Omega_{\mathcal{C}}(\{1, 4\}, 4)$. By flipping on $\{2, 5\}$ we obtain the geometric chamber \mathcal{C}_2 depicted in Fig. 2(c). Finally, if we first flip on $\{1, 4\}$ and then on $\{2, 5\}$, we obtain the geometric chamber \mathcal{C}_3 shown in Fig. 2(d).

4. Virtual Chambers and Pseudo-Chambers

4.1. Definition and Properties of Pseudo-Chamber Complexes

In this section we prove that every virtual chamber of a rank 3 vector configuration \mathcal{B} , which is a combinatorial object, can be realized in a certain sense as a geometric (or, rather, topological) object: as a cell of a cell complex in the sphere S^2 very similar to the chamber complex of \mathcal{B} .

The cell complexes which appear are called *pseudo-chamber complexes of \mathcal{B}* and they have the same good properties as the chamber complex: their full-dimensional cells




Fig. 2. Illustration of Example 3.12. (a) The \mathcal{C} -coherent orientation in the Gale diagram of \mathcal{A} . (b)–(d) The chambers \mathcal{C}_1 , \mathcal{C}_2 and \mathcal{C}_3 , respectively.

represent triangulations of the Gale transform \mathcal{A} and adjacent full-dimensional cells correspond to triangulations of \mathcal{A} which differ by a flip.

Definition 4.1. Let \mathcal{B} be a rank 3 vector configuration, regarded as a point configuration in S^2 . Schönflies theorem, see [15], implies that all the constructions below can be considered PL-topological without loss of generality, although we drop the prefix PL. For example, if $C \subset S^2$ is an embedded S^1 and C^+ and C^- are the two connected components of $S^2 \setminus C$, then $C^+ \cup C$ (same for C^-) is a topological disk with boundary C and interior C^+ .

- A *pseudo-edge* c of \mathcal{B} is the image of a topological embedding $\varphi: [0, 1] \rightarrow S^2$ such that $p = \varphi(0)$ and $q = \varphi(1)$ are non-antipodal distinct elements of \mathcal{B} . We say that c *joins* p and q .
- Let $\{p, q, r\}$ be a triangle of \mathcal{B} . Let c_1, c_2 and c_3 be pseudo-edges joining respectively $\{p, q\}, \{q, r\}$ and $\{p, r\}$ and such that they intersect only in their end-points. Then their union is homeomorphic to S^1 and divides $S^2 \setminus (c_1 \cup c_2 \cup c_3)$ into two connected components D_1 and D_2 . We say that D_1 is a *pseudo-triangle* of \mathcal{B} if D_1 and $\text{conv}(p, q, r)$ define, together with the ordered triple (p, q, r) , the same orientation for S^2 .

Definition 4.2. Let \mathcal{B} be a rank 3 vector configuration.

- A *pseudo-chamber complex* Γ of \mathcal{B} is the cellular decomposition of $\text{conv}(\mathcal{B})$ induced by a collection of pseudo-edges of \mathcal{B} (called the pseudo-edges of Γ) satisfying:
 - (i) For every edge $\{p, q\}$ of \mathcal{B} , there is exactly one pseudo-edge of Γ joining p and q .
 - (ii) For every two edges $\{p_1, q_1\}$ and $\{p_2, q_2\}$ of \mathcal{B} and the corresponding pseudo-edges c_1 and c_2 of Γ one has:
 - $c_1 \subset c_2$ if and only if $\text{conv}(p_1, q_1) \subset \text{conv}(p_2, q_2)$.
 - If $c_1 \not\subset c_2$, then c_1 intersects c_2 if and only if $\text{conv}(p_1, q_1)$ intersects $\text{conv}(p_2, q_2)$. In this case, c_1 and c_2 intersect in exactly one point.
 - (iii) Every pseudo-triangle of \mathcal{B} defined by pseudo-edges of Γ contains exactly the same points of \mathcal{B} as the convex hull of the corresponding triangle.
- We write $\text{conv}_\Gamma(p, q) = c$ for an edge $\{p, q\}$ of \mathcal{B} , where c is the pseudo-edge joining p and q . Also, we call $c \setminus \{p, q\}$ the *relative interior* of c and write $\text{relconv}_\Gamma(p, q)$.
- The pseudo-triangles of \mathcal{B} defined by pseudo-edges of Γ are called *pseudo-triangles of Γ* . As we did for pseudo-edges, for a triangle $\{p, q, r\}$ of \mathcal{B} we denote $\text{conv}_\Gamma(p, q, r)$ as the corresponding pseudo-triangle and define its *relative interior* by

$$\text{relconv}_\Gamma(p, q, r) := \text{conv}_\Gamma(p, q, r) \setminus (\text{conv}_\Gamma(p, q) \cup \text{conv}_\Gamma(p, r) \cup \text{conv}_\Gamma(q, r)).$$
- A pseudo-edge or a pseudo-triangle is called *empty* if the only points of \mathcal{B} it contains are its two or three vertices, respectively.
- We say that a pseudo-chamber complex is *generic* if there are no three empty pseudo-edges whose relative interiors intersect.
- A *pseudo-triangulation* of Γ is a topological triangulation of $\text{conv}(\mathcal{B})$ by pseudo-triangles of Γ .
- The full-dimensional cells of Γ are called *pseudo-chambers* of Γ .

Example 4.3. Figure 3 shows the virtual chamber \mathcal{C} considered in Examples 1.5 and 3.12 realized as a pseudo-chamber. It is easy to check that the pseudo-triangles defined by the triangles of \mathcal{C} (listed in Example 1.5) intersect in the shaded region.

Remark 4.4. The “*sides*” of a pseudo-edge of \mathcal{B} can be defined locally in a topological sense, although they cannot have the global meaning that sides of an edge of \mathcal{B} have since pseudo-edges do not define hemispheres of S^2 . The following are some properties of any pseudo-chamber complex Γ which can be easily proved for this local definition of sides.

1. For any two triangles σ and τ of Γ , the corresponding pseudo-triangles overlap (meaning that their relative interiors intersect) if and only if $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$.
2. Two pseudo-triangles of Γ are incident to opposite sides of a pseudo-edge if and only if the corresponding triangles of \mathcal{B} are incident to opposite sides of the corresponding edge.




Fig. 3. A realization of \mathcal{C} as a pseudo-chamber (Example 4.3).

3. If two pseudo-edges c and c' intersect in a point r of their relative interiors, then the two components of $c \setminus \{r\}$ (resp. $c' \setminus \{r\}$) are incident to opposite sides of c' (resp. c). In other words, the intersection of c and c' is topologically transversal.
4. In the conditions of the previous point, there exists a neighbourhood D of r homeomorphic to a closed disk whose boundary intersects alternately c and c' (when going along the boundary of D in any direction) and exactly once per segment (of those defined by r).

Lemma 4.5. Any rank 3 vector configuration \mathcal{B} has a generic pseudo-chamber complex.

Proof. Start with an arbitrary pseudo-chamber complex Γ_0 of \mathcal{B} (for example, the chamber complex having as pseudo-edges the geodesic arcs) and then, whenever three empty pseudo-edges intersect in a point, perturb one of them slightly but keep the property that it intersects the others transversally and in a unique point. That this can be done is obvious in the PL category. □

Proposition 4.6. For any pseudo-chamber complex Γ of \mathcal{B} , there is a natural bijection between triangulations of \mathcal{B} and pseudo-triangulations of Γ .

Proof. The result follows straightforwardly from condition (ii) of Definition 4.2 and the properties in Remark 4.4. □

Proposition 4.7. Let Γ be a pseudo-chamber complex of \mathcal{B} .

1. For every pseudo-chamber \mathcal{C} of Γ , the collection of pseudo-triangles of Γ which contain \mathcal{C} correspond to a virtual chamber of \mathcal{B} .
2. If two pseudo-chambers \mathcal{C}_1 and \mathcal{C}_2 of Γ are adjacent, then the corresponding virtual chambers of \mathcal{B} (which we also denote by \mathcal{C}_1 and \mathcal{C}_2) differ by a flip.

Proof. Part 1 is trivial taking into account Proposition 4.6.

For part 2, let \mathcal{C}_1 and \mathcal{C}_2 be two adjacent pseudo-chambers of Γ . That is to say, the closures of \mathcal{C}_1 and \mathcal{C}_2 share a subsegment of a pseudo-edge $c = \text{conv}_\Gamma(p, q)$, which we can assume to be empty. Then every pseudo-triangle which contains \mathcal{C}_1 either contains \mathcal{C}_2 as well or is incident to c in the opposite side as \mathcal{C}_2 , and vice versa.

In particular, the pseudo-triangles of \mathcal{C}_1 incident to c must all be incident to the same side of c since they contain \mathcal{C}_1 (which is incident to c), and the pseudo-triangles of \mathcal{C}_2 incident to c must all be incident to the opposite side, since \mathcal{C}_1 and \mathcal{C}_2 are incident to the opposite sides of c .

We consider the edge $l = \{p, q\}$ corresponding to the pseudo-edge c . For any r in \mathcal{B} such that $\{p, q, r\}$ is a triangle the pseudo-triangle $\text{conv}_\Gamma(p, q, r)$ is incident to c and, thus, contains either \mathcal{C}_1 or \mathcal{C}_2 . Moreover, whether it contains \mathcal{C}_1 or \mathcal{C}_2 depends only on whether $r \in l^+$ or $r \in l^-$ (for a suitable orientation of l), thanks to the role of orientation in Definition 4.1. Thus, l is in the conditions of Definition 1.10 and, hence, it supports a flip of \mathcal{C}_1 and \mathcal{C}_2 . That this flip exchanges between \mathcal{C}_1 and \mathcal{C}_2 follows from Corollary 1.11. \square

4.2. Every Virtual Chamber Realizes as a Pseudo-Chamber

We have shown that pseudo-chamber complexes have the good properties we announced. Now we want to prove that every virtual chamber of \mathcal{B} realizes as a pseudo-chamber of some pseudo-chamber complex of \mathcal{B} .

First observe that for a virtual chamber \mathcal{C} of \mathcal{B} and for an empty edge $l = \{p, q\}$ supporting a flip of \mathcal{C} , if $l' = \{p', q'\}$ is an edge crossing l , then either $\{p', q', p\}$ or $\{p', q', q\}$ is in \mathcal{C} , since either $\{p, q, p'\}$ or $\{p, q, q'\}$ is in \mathcal{C} (we are using condition 2 of Definition 1.1 applied to $\rho = \{p, q, p', q'\}$ which supports the circuit $(\{p, q\}, \{p', q'\})$). Since the sides of l' correspond to (local) sides of the pseudo-edge $\text{conv}_\Gamma(l')$, it makes sense to say that $\text{conv}_\Gamma(l')$ has \mathcal{C} on a certain side.

Lemma 4.8. *Let Γ be a generic pseudo-chamber complex of \mathcal{B} and let $l = \{p, q\}$, $l_1 = \{r, s\}$ and $l_2 = \{r, t\}$ be empty edges of \mathcal{B} such that l_1 and l_2 cross l . Then, when going (along $\text{conv}(l)$) from p to q , we cross $\text{conv}(l_1)$ and $\text{conv}(l_2)$ in the same order as we cross $\text{conv}_\Gamma(l_1)$ and $\text{conv}_\Gamma(l_2)$ (when going from p to q along $\text{conv}_\Gamma(l)$).*

Proof. Assume we cross $\text{conv}(l_1)$ first (along $\text{conv}(l)$). This implies $\text{conv}(l_2) \cap \text{conv}(\{p, r, s\}) = \{r\}$. If we cross $\text{conv}_\Gamma(l_2)$ first (along $\text{conv}_\Gamma(l)$), then either $\text{conv}_\Gamma(l_2) \cap \text{conv}_\Gamma(\{p, r, s\}) \neq \{r\}$ (which implies that $\text{conv}_\Gamma(l_2)$ crosses $\text{conv}_\Gamma(\{p, s\})$, and this violates condition (ii) of Definition 4.2) or the open interval of $\text{conv}_\Gamma(l)$ between p and $\text{conv}_\Gamma(l_1)$ is, by condition (ii) of Definition 4.2, contained in $S^2 \setminus \text{conv}_\Gamma(\{p, r, s\})$ (but then point 3 of Remark 4.4 forces $q \in \text{conv}_\Gamma(\{p, r, s\})$, which violates condition (iii) of Definition 4.2). \square

Definition 4.9. Let \mathcal{C} be a virtual chamber of \mathcal{B} and let $l = \{p, q\}$ be an empty edge supporting a flip of \mathcal{C} . Let Γ be a generic pseudo-chamber complex of \mathcal{B} and let $c = \text{conv}_\Gamma(l)$ be the pseudo-edge of Γ which corresponds to l . Let c_1 and c_2 be two pseudo-

edges of Γ which intersect the relative interior of c . We say that:

- c_1 and c_2 are \mathcal{C} -incoherent along c if, when we go along c (in any direction), we cross the first one from the side of \mathcal{C} and the second one from the side opposite to \mathcal{C} . In any other case we say that c_1 and c_2 are \mathcal{C} -coherent along c .
- c_1 and c_2 are c -neighbours if there is no pseudo-edge of Γ which intersects c in the open arc between c_1 and c_2 .

Proposition 4.10. *Let \mathcal{C} be a virtual chamber of \mathcal{B} and let $l = \{p, q\}$ be an empty edge supporting a flip of \mathcal{C} . Let Γ be a generic pseudo-chamber complex of \mathcal{B} and let $c = \text{conv}_\Gamma(l)$. Let c_1 and c_2 be two pseudo-edges of Γ which cross c . If c_1 and c_2 do not cross each other, then they are \mathcal{C} -coherent along c .*

Proof. Clearly, if the statement is true when c_1 and c_2 are empty pseudo-edges, then it is true for non-empty ones as well. So, we assume that c_1 and c_2 are empty.

Suppose first that c_1 and c_2 share a vertex. Say $c_1 = \text{conv}_\Gamma(r, p_1)$ and $c_2 = \text{conv}_\Gamma(r, p_2)$ and the five points $\{p, q, p_1, p_2, r\} \subset \mathcal{B}$ are distinct.

Since l supports a flip of \mathcal{C} , one of the triangles $\{p, r, p_1\}$ and $\{q, r, p_1\}$ and one of the triangles $\{p, r, p_2\}$ and $\{q, r, p_2\}$ are in \mathcal{C} . Say $\tau = \{p, r, p_1\} \in \mathcal{C}$. If $\{p, r, p_2\} \in \mathcal{C}$, then c_1 and c_2 are \mathcal{C} -coherent. So, suppose that $\sigma = \{q, r, p_2\}$ is in \mathcal{C} . What we want to prove is that when going along c from p to q we cross c_2 first and c_1 afterwards (if this happens c_1 and c_2 are \mathcal{C} -coherent).

Set $l_1 = \{p_1, r\}$ and $l_2 = \{p_2, r\}$. By Lemma 4.8 we have to show that when going along $\text{conv}(l)$ from p to q we cross $\text{conv}(l_2)$ first and $\text{conv}(l_1)$ afterwards. However, if this was not the case we would have $\sigma \cap l_1^+ = \emptyset$, which is a contradiction, so we are done.

Now assume that c_1 and c_2 do not share a vertex. Say $c_1 = \text{conv}_\Gamma(p_1, q_1)$ and $c_2 = \text{conv}_\Gamma(p_2, q_2)$. Without loss of generality assume that p_1 and p_2 are on the same side of l . By Lemma 2.5, either $\{p_1, q_2\}$ crosses l or $\{p_2, q_1\}$ crosses l . Say that $\{p_1, q_2\}$ crosses l and let $c_3 = \text{conv}_\Gamma(p_1, q_2)$. Then the fact that c_3 is \mathcal{C} -coherent with both c_1 and c_2 and that it intersects c between c_1 and c_2 (which follows straightforward from Lemma 4.8) implies that c_1 and c_2 are \mathcal{C} -coherent as well. \square

Lemma 4.11. *Let Γ be a generic pseudo-chamber complex of \mathcal{B} and let c_1, c_2 and c_3 be three pseudo-edges of Γ such that every two of them cross each other. Let γ be the closed simple curve defined by the segments of the three pseudo-edges between the intersection points. Then one of the two connected components of $S^2 \setminus \gamma$ contains all six vertices of c_1, c_2 and c_3 .*

Proof. For $1 \leq i < j \leq 3$ let r_{ij} be the intersection point between c_i and c_j , and for $1 \leq k \leq 3$ let c'_k be the arc of c_k between the intersection points. We consider an arbitrary intersection point, for instance, $r_{1,2}$. This point $r_{1,2}$ divides c_1 (resp. c_2) into two arcs, one of them containing c'_1 (resp. c'_2). Let the other be c''_1 (resp. c''_2). This latter arc has as end-points $r_{1,2}$ and a vertex of c_1 (resp. c_2) and contains no other intersection point.

The arcs c''_1 and c''_2 cannot intersect γ by condition (ii) of Definition 4.2, so each one of them is completely contained in one of the two regions defined by γ . By point 4 of

Remark 4.4 there is a neighbourhood D of $r_{1,2}$ homeomorphic to a disk such that c'_1, c'_2, c''_1 and c''_2 intersect the boundary of D once each and in this order for a certain choice of orientation of D .

Then the sector of D between the arcs c'_1 and c'_2 is precisely the intersection between D and one of the components of $S^2 \setminus \gamma$, so the other region contains the other three sectors. This forces c'_1 and c'_2 to be contained in the same component of $S^2 \setminus \gamma$. In particular, the vertices of c_1 and c_2 incident to c'_1 and c'_2 are in the same region. So we have proved that vertices of c_i and c_j which are “adjacent” (in the obvious sense) to the same intersection point r_{ij} between pseudo-edges are in the same component of $S^2 \setminus \gamma$.

Now suppose that both components of $S^2 \setminus \gamma$ contain some of the six vertices. Then there are four vertices in one component and two vertices in the other one and there is at least one pseudo-edge (say c_1) whose two vertices are in the same component.

We denote by p_i and q_i the vertices of c_i for $i = 1, 2, 3$ and assume that p_2 is in the opposite region to that containing p_1 and q_1 . Since the edges $\{p_1, q_1\}$ and $\{p_2, q_2\}$ cross each other, the points p_1, q_1 and p_2 are independent, so they define a pseudo-triangle $\text{conv}_\Gamma(\tau)$ of Γ .

The pseudo-edge $\text{conv}_\Gamma(p_2, p_1)$ cannot intersect c'_1 nor c'_2 by condition (ii) of Definition 4.2, but it must intersect γ , so it intersects c'_3 . Similarly, $\text{conv}_\Gamma(p_2, q_1)$ intersects c'_3 . We conclude that the three pseudo-edges of $\text{conv}_\Gamma(p_2, p_1)$ cross c_3 , which is impossible by condition (ii) of Definition 4.2 and Proposition 2.4. \square

Proposition 4.12. *Let $l = \{p, q\}$ be an empty edge of \mathcal{B} supporting a flip of a virtual chamber \mathcal{C} . Then there exists a generic pseudo-chamber complex Γ such that every two pseudo-edges of Γ which intersect the pseudo-edge $c = \text{conv}_\Gamma(l)$ are \mathcal{C} -coherent along c .*

Proof. We proceed by induction on the number of pairs of c -neighbours which are \mathcal{C} -incoherent along c .

First take any generic pseudo-chamber complex Γ_1 (we can do so by Lemma 4.5). If there is no pair of non-coherent pseudo-edges of Γ_1 we are done. If there are non-coherent pairs, it is clear that we can find at least one pair of non-coherent pseudo-edges c_1 and c_2 which are c -neighbours. Let $c_1 = \text{conv}_\Gamma(p_1, q_1)$ and $c_2 = \text{conv}_\Gamma(p_2, q_2)$.

By Proposition 4.10, c_1 and c_2 intersect in a point r which is not in \mathcal{B} . Without loss of generality we can assume that c_1 and c_2 are empty pseudo-edges, so that c, c_1 and c_2 intersect pairwise in three points. Let c', c'_1 and c'_2 denote the closed arcs of c, c_1 and c_2 between the intersection points and let $\gamma = c' \cup c'_1 \cup c'_2$. By Lemma 4.11 one of the connected components of $S^2 \setminus \gamma$ contains the six points $\{p, q, p_1, q_1, p_2, q_2\}$. We call the other one D .

We claim that the closure of D does not contain any point of \mathcal{B} : Since c_1, c_2 and c_3 are empty pseudo-edges, it suffices to show that $D \cap \mathcal{B} = \emptyset$. Suppose there exists $x \in D \cap \mathcal{B}$. The point x is antipodal to at most one of the points p_1, q_1, p_2 and q_2 , so we assume without loss of generality that it is not antipodal to neither p_1 nor q_1 . Then we consider the pseudo-edges c_3 and c_4 joining x to p_1 and q_1 , which must intersect γ . Neither c_3 nor c_4 can intersect the relative interior of c' , since c_1 and c_2 are c -neighbours. If one of them (say c_3) intersects $c'_1 \subset \text{relconv}_\Gamma(p_1, q_1)$, then either $c_3 \subset c_1$ (impossible, since $x \notin c_1$) or $c_1 \subset c_3$ (impossible, since then c_1 should intersect $\gamma \setminus c_1$). Thus, both c_3 and c_4 intersect c_2 . Then the three pseudo-edges $c_1 = \text{conv}_\Gamma(p_1, q_1)$,

$c_3 = \text{conv}_\Gamma(p_1, x)$ and $c_4 = \text{conv}_\Gamma(q_1, x)$ intersect c_2 . We have the following circuits supported on $\{p_2, q_2, p_1, q_1, x\}$: $(+, +, -, -, 0)$, $(+, +, -, 0, -)$ and $(+, +, 0, -, -)$. Elimination of x in the last two gives $(*, *, -, +, 0)$ which contradicts the fact that the first one is a circuit.

Once we know that the closure of D does not contain any points of \mathcal{B} we observe the following fact: a pseudo-edge c'' of Γ_1 intersects c'_1 if and only if it intersects c'_2 : Suppose c'' intersects c'_1 . The intersection point is in the relative interior of c'_1 (since Γ_1 is generic) and c'' must intersect γ in a second point, since both its end-points are outside D . Since it cannot intersect c' because c_1 and c_2 are c -neighbours, it intersects c'_2 .

With all this information we are going to perturb the pseudo-edge c_1 of Γ_1 to obtain a new pseudo-chamber complex Γ_2 in which c has less \mathcal{C} -incoherent pairs than in Γ_1 : we consider an open arc u_1 of c_1 containing c'_1 and with no intersections with any pseudo-edges of Γ_1 apart from those in c'_1 . We remove u_1 from c_1 and insert instead an open arc v_1 with the same extremal points as u_1 but which intersects c and c_2 in the opposite order as u_1 does. We can do so by drawing v_1 in two pieces, one “parallel” to c' , very close to it and outside D and the other “parallel” to c'_2 , very close to it and outside D .

The fact that c'_1 and c'_2 intersect exactly the same pseudo-edges of Γ_1 implies that we can do this in such a way that u_1 and v_1 intersect exactly the same pseudo-edges of Γ_1 . Thus, conditions (i) and (ii) of Definition 4.2 are preserved. Using the fact that D does not contain points of \mathcal{B} it is not hard to see that condition (iii) is preserved as well. In other words, we have constructed a new generic pseudo-chamber complex Γ_2 in which the number of c -incoherent pairs of c -neighbours has been decreased by one with respect to Γ_1 . □

Theorem 4.13. *Any virtual chamber \mathcal{C} of a rank 3 vector configuration \mathcal{B} realizes as a pseudo-chamber of some pseudo-chamber complex of \mathcal{B} .*

Proof. We take an empty edge l of \mathcal{B} which supports a flip of \mathcal{C} (we can do so by Corollary 3.5). By Proposition 4.12, there exists a generic pseudo-chamber complex Γ such that every two pseudo-edges of Γ that intersect the pseudo-edge $c = \text{conv}_\Gamma(l)$ are \mathcal{C} -coherent along c .

We travel along c from one vertex to the other. If there is a pseudo-edge c' of Γ which we cross from the side of c' in which \mathcal{C} is, then the same must occur with any other pseudo-edge we meet after c' , by the \mathcal{C} -coherence assumption. Similarly if we cross c' from the side opposite to \mathcal{C} , then the same has occurred for any pseudo-edge we crossed before c' .

In other words, the pseudo-edges we cross are divided in two subsets E_1 and E_2 (perhaps empty) such that we cross first all the pseudo-edges of E_1 from the side opposite to \mathcal{C} and then all the pseudo-edges of E_2 from the side of \mathcal{C} . Between these two groups there must be an open arc I of c which is on the same side as \mathcal{C} (along c) of every pseudo-edge that intersects c . Let \mathcal{C}_Γ be the collection of triangles τ of \mathcal{B} such that $\text{relconv}_\Gamma(\tau)$ contains the pseudo-chamber incident to I on the side of \mathcal{C} . We claim that $\mathcal{C} = \mathcal{C}_\Gamma$, which finishes the proof.

Let $\tau \in \mathcal{C}$. We first see that $I \subset \text{conv}_\Gamma(\tau)$.

- If l is not an edge of τ , then either $l \subset \text{conv}(\tau)$ (in which case c is contained in $\text{conv}_\Gamma(\tau)$ by condition (iii) of Definition 4.2) or some edge (or edges) of τ crosses l . In the last case, the corresponding pseudo-edge (resp. pseudo-edges) of $\text{conv}_\Gamma(\tau)$ intersect c , so it (they) has I in the same side as \mathcal{C} . This again implies that I is contained in $\text{conv}_\Gamma(\tau)$.
- If l is an edge of τ , then $\text{conv}_\Gamma(\tau)$ contains I trivially.

So I is contained in all the pseudo-triangles of \mathcal{C} . Moreover, since I is an open arc of c which does not intersect any other pseudo-edge of Γ , for every $\tau \in \mathcal{C}$ either I is fully contained in $\text{relconv}_\Gamma(\tau)$ or fully contained in a pseudo-edge of $\text{conv}_\Gamma(\tau)$. In the last case, $\text{relconv}_\Gamma(\tau)$ is incident to c on the side of the virtual chamber \mathcal{C} (this is the definition of “ \mathcal{C} is in one side of c ”). Thus, the pseudo-chamber incident to I on the side of \mathcal{C} intersects (and, hence, is fully contained in) the relative interior of every pseudo-triangle of \mathcal{C} .

Hence, $\mathcal{C} \subset \mathcal{C}_\Gamma$. By Proposition 4.7, \mathcal{C}_Γ is a virtual chamber of \mathcal{B} and this implies $\mathcal{C} = \mathcal{C}_\Gamma$ since otherwise we would have two different triangulations of the Gale transform \mathcal{A} , one contained in the other, which is impossible. □

Remark 4.14. Taking into account condition (ii) of Definition 4.2, each pseudo-chamber of a pseudo-chamber complex of \mathcal{B} must be incident to at least three pseudo-edges. On the other hand, all the flag chambers appear in any pseudo-chamber complex, since they correspond to pseudo-chambers which are incident to both a pseudo-edge and one of its vertices. With these observations it is easy to obtain Theorem 3.9 as a corollary of Theorem 4.13.

Example 4.15. The following example shows that there exist virtual chambers of rank 3 configurations which are not geometric chambers in any realization of the oriented matroid. This kind of virtual chambers was called *truly virtual chambers* in [5]. In addition, we show a corank 3 point configuration \mathcal{A}' whose Gale transform \mathcal{B}' has no virtual chambers but the oriented matroid of \mathcal{B}' being a mere reorientation of that of \mathcal{B} . This shows that having non-geometric virtual chambers or truly virtual chambers does depend on such subtle things as reorientations of the oriented matroid.

Let \mathcal{A} be the (corank 3) point configuration in \mathbb{R}^3 consisting of the points $p_1 = (2, 0, 0)$, $p_2 = (0, 2, 0)$, $p_3 = (0, 0, 2)$, $p_4 = (1, 0, 0)$, $p_5 = (0, 1, 0)$, $p_6 = (0, 0, 1)$ and $p_7 = (1, 1, 1)$. \mathcal{A} is the set of vertices of a truncated tetrahedron together with an extra point beyond the untouched facet (defined by p_1, p_2 and p_3) of the tetrahedron. We consider in \mathcal{A} the triangulation \mathcal{T} defined by the tetrahedra $\{p_4, p_5, p_6, p_7\}$, $\{p_1, p_2, p_4, p_7\}$, $\{p_2, p_3, p_5, p_7\}$, $\{p_1, p_3, p_6, p_7\}$, $\{p_2, p_4, p_5, p_7\}$, $\{p_3, p_5, p_6, p_7\}$ and $\{p_1, p_4, p_6, p_7\}$.

Figure 4(a) shows an affine Gale diagram of \mathcal{A} . Point 7 is contained in the linear spans of the edges $\{1, 4\}$, $\{2, 5\}$ and $\{3, 6\}$, but in none of their positive spans. Hence, in any realization of the oriented matroid of \mathcal{B} , the relative interiors of the edges $\{1, 4\}$, $\{2, 5\}$ and $\{3, 6\}$ must intersect in a common point (since the three of them must contain the opposite of 7 in their relative interiors), and therefore \mathcal{C} , the virtual chamber of \mathcal{B} which corresponds to \mathcal{T} , must remain unrealized as a geometric chamber. Equivalently,




Fig. 4. Affine Gale diagram of the configuration \mathcal{A} in Example 4.15 (a) and realization of the virtual chamber \mathcal{C} as a pseudo-chamber (b).

for any realization of the oriented matroid of \mathcal{A} , the triangulation \mathcal{T} remains non-regular. Note that \mathcal{C} is defined by the same triangles as in Example 1.5. This happens because the triangulation of Example 1.5 is the link of the point p_7 in \mathcal{T} .

Nevertheless, our virtual chamber \mathcal{C} must be realizable as a pseudo-chamber of some pseudo-chamber of \mathcal{B} (as shown in Fig. 4(b)), by Theorem 4.13.

Now let \mathcal{A}' be the point configuration in \mathbb{R}^3 consisting of the points $p_1 = (2, 0, 0)$, $p_2 = (0, 2, 0)$, $p_3 = (0, 0, 2)$, $p_4 = (1, 0, 0)$, $p_5 = (0, 1, 0)$, $p_6 = (0, 0, 1)$ and $p_7 = (-1, -1, -1)$. \mathcal{A}' is the set of vertices of the same truncated tetrahedron together with an extra point beyond the apex (the point $(0, 0, 0)$). Figure 5 depicts an affine Gale diagram of \mathcal{A}' , which coincides with that of \mathcal{A} except for a reorientation of point 7. It is easy to see that no additional pseudo-chamber of the Gale transform \mathcal{B}' of \mathcal{A}' can be created. Thus, by Theorem 4.13, every virtual chamber of \mathcal{B} defines a geometric chamber. Equivalently, every triangulation of \mathcal{A}' is regular. Observe that \mathcal{A}' is just a reorientation of \mathcal{A} .




Fig. 5. Affine Gale diagram of the configuration \mathcal{A}' in Example 4.15.

Appendix. The Graph of Triangulations of Corank 3 Oriented Matroids

The results presented in this paper have been proved using the language of vector configurations. Since the collection of triangulations and flips of a vector configuration depends only on the oriented matroid defined by the dependences among its elements (see, for example, [5]) what we have done so far is deal with the graph of triangulations of realizable corank 3 oriented matroids (via the graph of virtual chambers of its rank 3 dual).

The purpose of this Appendix is to show that all the results of the previous sections hold also for non-realizable oriented matroids of corank/rank 3. The starting point is the definition of triangulations of an oriented matroid introduced in Section 9.6 of [4] for acyclic oriented matroids and generalized in [16] for non-acyclic ones. This definition agrees with the geometric definition if the oriented matroid is realizable [4, Proposition 9.6.2] and has the property that triangulations of an oriented matroid \mathcal{M} are dual by complementarity of simplices to virtual chambers of the dual oriented matroid \mathcal{M}^* , where the virtual chambers of \mathcal{M}^* are defined exactly using our Definition 1.1 (this is proved in Theorem 3.8 of [16]).

Thus, we are led to study virtual chambers of rank 3 oriented matroids. The crucial property of rank 3 oriented matroids is that they admit a “Type II” realization, i.e. a topological representation as a pseudo-configuration of points in the sphere S^2 . First we see what this means.

A.1. Pseudo-Configurations of Points

Definitions A.1 and A.2 below are taken from [4], except that we give them for the case of rank 3.

Definition A.1 [4, Definitions 5.1.2 and 5.1.3]. Let S^2 denote the standard sphere of dimension 2.

- A *pseudo-circle* S in S^2 is a (topological) subspace of S^2 which is PL-homeomorphic to S^1 . The two connected components of $S^2 \setminus S$ are called *sides of S* or *open hemispheres* and are denoted by S^+ and S^- . Their closures $S \cup S^+$ and $S \cup S^-$ are called *closed sides of S* or *closed hemispheres* and are denoted $\overline{S^+}$ and $\overline{S^-}$.
- A *pseudo-circle arrangement* (or a *pseudo-sphere arrangement in S^2*) is a finite set Λ of pseudo-circles such that:
 1. The intersection of any subset of at least two spheres in Λ is either empty or a pair of points.
 2. For any $S, S' \in \Lambda$ with $S \neq S'$, $S \cap S'$ is a pair of points and the two connected components of $S' \setminus S$ coincide with $S^+ \cap S'$ and $S^- \cap S'$. If $S'' \in \Lambda$ is another pseudo-sphere and does not contain the two points $S \cap S'$, then one of the points lies in S''^+ and the other in S''^- .
 3. The intersection of an arbitrary collection of closed sides of pseudo-circles of Λ is either a (topological) sphere or a ball.

Axiom 3 in the above definition of pseudo-circle arrangement is implied by 1 and 2, as shown by Edmonds and Mandel. Nevertheless, we keep axiom 3 in the definition because it will be useful later. The intersection points of different pseudo-circles in a pseudo-circle arrangement are called *vertices* of the arrangement.

Definition A.2 [4, Definition 5.3.1]. A pseudo-configuration of points in S^2 is a pair (Λ, \mathcal{B}) where Λ is a pseudo-circle arrangement and \mathcal{B} is a collection of vertices of Λ such that:

1. Any pair of points in \mathcal{B} are contained in some pseudo-circle in Λ .
2. For every pseudo-circle $S \in \Lambda$ there is a subset $\mathcal{B}' \subset \mathcal{B}$ which is contained in S and in no other pseudo-circle of Λ .

Any configuration of points \mathcal{B} in S^2 defines a pseudo-configuration of points whose pseudo-circles are the great circles passing through every pair of non-antipodal points of \mathcal{B} .

For simplicity, when referring to a pseudo-configuration of points (Λ, \mathcal{B}) we denote it just \mathcal{B} and assume that the pseudo-circles passing through the points in \mathcal{B} are given implicitly. Any pseudo-configuration of points \mathcal{B} has an associated oriented matroid $\mathcal{M}_{\mathcal{B}}$ whose set of cocircuits C^* is as follows:

$$C^* := \{\sigma(S): S \in \Lambda\} \sqcup \{-\sigma(S): S \in \Lambda\},$$

where $\sigma(S)$ is the function $\mathcal{B} \rightarrow \{+, 0, -\}$ (i.e. a signed subset of \mathcal{B}) defined as

$$\sigma(S)_p := \begin{cases} + & \text{if } p \in S^+, \\ - & \text{if } p \in S^-, \\ 0 & \text{if } p \in S. \end{cases}$$

We say that \mathcal{B} is a *pseudo-realization* or a *type II realization* of $\mathcal{M}_{\mathcal{B}}$. If Λ is an essential arrangement (i.e. if the intersection of all the pseudo-circles is empty), then $\mathcal{M}_{\mathcal{B}}$ has rank 3. We always assume this to be the case.

The cocircuits of $\mathcal{M}_{\mathcal{B}}$ are defined exactly in the same way as cocircuits for a point configuration \mathcal{B} in the sphere, with Λ being the collection of great circles passing through every pair of points in \mathcal{B} . So, every realization of a rank 3 oriented matroid \mathcal{M} as a point configuration in the sphere is, in particular, a pseudo-realization of \mathcal{M} .

We say that an oriented matroid is *simple* if it has no loops or positively parallel elements, i.e. if every circuit has at least two elements and those with two elements are of type $(\{p, q\}, \emptyset)$ (we say in this case that p and q are *antiparallel* or *opposite*). Our definition is slightly more general than the standard one (see [4]) in which a simple oriented matroid is not allowed to have antiparallel elements. The following pseudo-realizability result for simple rank 3 oriented matroids is not true in higher rank, since not every oriented matroid has an *adjoint*.

Theorem A.3. *Every simple rank 3 oriented matroid can be pseudo-realized as a pseudo-configuration of points in S^2 .*

Proof. According to Proposition 6.3.6 (Goodman–Cordovil–Pollack) of [4], every rank 3 oriented matroid has an adjoint. A simple oriented matroid has an adjoint if and only if it has a pseudo-realization [4, Theorem 5.3.6]. \square

Let us recapitulate. We start with a corank 3 oriented matroid \mathcal{M}^* whose graph of triangulations and flips we want to study. Theorem 3.8 in [16] tells us that this is equivalent to studying the graph of virtual chambers and flips of the dual oriented matroid \mathcal{M} , which has rank 3. If \mathcal{M} is simple, we consider \mathcal{M} pseudo-realized in the sphere S^2 by a pseudo-configuration of points \mathcal{B} and we show that the results of the previous sections hold for \mathcal{B} . Before continuing let us see that the assumption of \mathcal{M} being simple is not a loss of generality, because of the following fact which we already mentioned for vector configurations in Section 1 (see Lemma 1.12).

Lemma A.4. *Let \mathcal{M}^* be an oriented matroid and suppose that \mathcal{M} is not simple. Let \mathcal{M}_0 denote the oriented matroid obtained by deleting from \mathcal{M} all the loops and all but one copy of all parallel classes of elements. Let \mathcal{M}_0^* be the dual of \mathcal{M}_0 . Then \mathcal{M}^* and \mathcal{M}_0^* have the same graph of triangulations, and \mathcal{M} and \mathcal{M}_0 the same graph of virtual chambers.*

Proof. Deleting loops from \mathcal{M} is the same as deleting coloops from \mathcal{M}^* , which does not affect the collection of triangulations. On the other hand, if \mathcal{M} is obtained from \mathcal{M}_0 by adding parallel elements, then \mathcal{M}^* is obtained from \mathcal{M}_0^* by a *reoriented Lawrence construction* in the sense of Section 4.4 of [16]. Theorem 4.18 in that paper proves that \mathcal{M}^* and \mathcal{M}_0^* have the same collection of triangulations. The proof for flips follows the same lines. \square

A property of pseudo-configurations of points that will be useful later is that any point in $S^2 \setminus \mathcal{B}$ induces an extension of the pseudo-configuration (perhaps not only one):

Proposition A.5. *Let \mathcal{B} be a pseudo-configuration of points in S^2 and let $p \in S^2 \setminus \mathcal{B}$. Then there is a pseudo-configuration $\mathcal{B}' = \mathcal{B} \cup \{p\}$ which extends \mathcal{B} in p . By this we mean that for each $q \in \mathcal{B}$ there exists a pseudo-circle S_q containing p and q such that $\Lambda' := \Lambda \cup \{S_q : q \in \mathcal{B}\}$ is an arrangement of pseudo-circles and (\mathcal{B}', Λ') is a pseudo-configuration of points (where Λ is the collection of pseudo-circles of the pseudo-configuration \mathcal{B}).*

Observe that in these conditions the oriented matroid $\mathcal{M}_{\mathcal{B}'}$ is a single element extension of the oriented matroid $\mathcal{M}_{\mathcal{B}}$.

Proof. This is a consequence of Levi’s Enlargement Lemma [4, Proposition 6.3.4]. This lemma asserts that for any arrangement of pseudo-circles Λ in S^2 and any pair of points $x, y \in S^2$ there exists a pseudo-circle S in S^2 which contains x and y and extends Λ to an arrangement $\Lambda \cup \{S\}$ (unless x and y already lie in a pseudo-circle of Λ , in which case we will not need to add S). Doing this iteratively with $x = p$ and $y \in \mathcal{B}$ gives the result. \square

A.2. Basic Properties of Pseudo-Configurations

Definition A.6. Let \mathcal{B} be a pseudo-configuration of points in S^2 .

- We say that two points $p, q \in S^2$ (not necessarily in \mathcal{B}) are *antipodal* if $\{p, q\}$ is the intersection of two pseudo-circles of the arrangement.
- For any two different points $p, q \in \mathcal{B}$ which are non-antipodal we say that $\{p, q\}$ is an *edge* of \mathcal{B} .
- For three points $p, q, r \in \mathcal{B}$, we say that $\{p, q, r\}$ is a *triangle* of \mathcal{B} if no pseudo-circle contains the three of them.

Remark A.7. 1. Single points, edges and triangles of \mathcal{B} are the independent sets of rank 1, 2 and 3 respectively of the oriented matroid $\mathcal{M}_{\mathcal{B}}$ pseudo-realized by \mathcal{B} . We call them *simplices* of \mathcal{B} and call the rank of any subset $\sigma \subset \mathcal{B}$ the rank of the maximal simplices it contains. The rank of \emptyset is 0. The rank of any other set is 3, 2 or 1 according to whether there is no pseudo-circle, exactly one pseudo-circle or more than one pseudo-circle containing it.

2. If p and q are two antipodal points, then for every pseudo-circle S of the arrangement either $\{p, q\} \subset S$ or p and q lie in opposite sides of S (this is a consequence of axiom 2 of Definition A.1). In other words, p and q are opposite elements in the oriented matroid $\mathcal{M}_{\mathcal{B}}$.

3. If $\{p, q\} \subset \mathcal{B}$ is a pair of non-antipodal points, then there is a unique pseudo-circle containing them (existence is axiom 1 of Definition A.2, uniqueness is the definition of non-antipodal).

4. If $\{p, q, r\}$ is a triangle no pair of them can be antipodal (a pseudo-circle containing p and q will not contain r , which proves that r is not antipodal with p nor q). Actually, $\{p, q, r\}$ is a triangle if and only if no pair of them is antipodal and the three pseudo-circles passing through the edges $\{p, q\}$, $\{p, r\}$ and $\{q, r\}$ are distinct (equivalently, two of them are distinct).

Definition A.8. Let \mathcal{B} be a pseudo-configuration of points in S^2 and let Λ be the corresponding pseudo-circle arrangement. For any subset $\tau \subset \mathcal{B}$ we define:

- The *convex hull* of τ (denoted $\text{conv}(\tau)$) as the intersection of all the closed sides of pseudo-circles of Λ containing τ .
- The *relative interior* of τ (denoted $\text{relconv}(\tau)$) as the intersection $(\bigcap_{\tau \subset S} S) \cap (\bigcap_{\tau \subset S^+, \tau \not\subset S} S^+)$, where S ranges over all the pseudo-circles in Λ .

From the definition it follows that $\text{relconv}(\tau)$ is the topological interior of $\text{conv}(\tau)$ in the intersection of all pseudo-circles containing τ and that $\text{conv}(\tau)$ is the closure of $\text{relconv}(\tau)$ both in S^2 and in that intersection of pseudo-circles. By axiom 3 of Definition A.1, $\text{conv}(\tau)$ is either a sphere or a ball of some dimension ≤ 2 . In particular, it is a topological manifold perhaps with boundary. Its interior, in the manifold sense, equals $\text{relconv}(\tau)$.

The following property of the convex hull is immediate from the definition: if $\sigma, \tau \subset \mathcal{B}$ are such that $\sigma \subset \text{conv}(\tau)$, then $\text{conv}(\sigma) \subset \text{conv}(\tau)$.

We are specially interested in $\text{conv}(\tau)$ and $\text{relconv}(\tau)$ when τ is a *simplex* of \mathcal{B} , i.e. a point, segment or triangle. For a single point we trivially have $\{p\} = \text{relconv}(p) = \text{conv}(p)$.

Lemma A.9. *The following properties hold for convex hulls and relative interiors in a pseudo-configuration of points \mathcal{B} :*

1. *Let $\{p, q\}$ be an edge and let $S \in \Lambda$ be the unique pseudo-circle containing the edge $\{p, q\}$. If $S' \in \Lambda \setminus \{S\}$ is such that p and q are on the same (closed) side of S' , then exactly one of the two connected components of $S \setminus \{p, q\}$ is completely contained in one side of S' .*
2. *This component is independent of the choice of S' and, actually, it equals $\text{relconv}(\{p, q\})$. Also, $\text{conv}(\{p, q\}) = \text{relconv}(\{p, q\}) \cup \{p, q\}$.*
3. *For a triangle $\{p, q, r\}$ let S_r, S_q and S_p be the pseudo-circles containing $\{p, q\}, \{p, r\}$ and $\{q, r\}$, respectively, and assume that $r \in S_r^+, q \in S_q^+, p \in S_p^+$. Then $\text{relconv}(\{p, q, r\}) = S_e^+ \cap S_f^+ \cap S_g^+$ and $\text{conv}(p, q, r) = \overline{S_e^+} \cap \overline{S_f^+} \cap \overline{S_g^+}$ (i.e. the closure of $\text{relconv}(\{p, q, r\})$). They are respectively an open and a closed 2-ball.*

Proof. 1. Assume without loss of generality that $p, q \in \overline{S'^+}$. $S \cap S'$ consists of two points p' and q' . By axiom 2 of Definition A.1 (with S and S' exchanged) the two components of $S \setminus \{p', q'\}$ coincide with $S'^+ \cap S$ and $S'^- \cap S$. Let $s^+ = S'^+ \cap S$. Then $s^+ \cup \{p', q'\} = \overline{S'^+}$ and hence $p, q \in s^+ \cup \{p', q'\}$. Thus, one of the connected components of $S \setminus \{p, q\}$ is contained in $s^+ \subset S'^+$. If the other one was contained in S'^- , then we would have $\{p, q\} = \{p', q'\}$ which cannot be the case since p and q are not antipodal.

2. Suppose that for two different pseudo-circles S' and S'' we had one component of $S \setminus \{p, q\}$ contained in one side of S' and the other component in one side of S'' . Then the intersection of S with the corresponding closed sides is disconnected and, by axiom 3 of Definition A.1, it equals the two points p, q . However, this implies p and q are antipodal.

This implies that if S' is a pseudo-circle other than S with p and q on the same closed side, the component l of $S \setminus \{p, q\}$ contained in some closed side of S' is contained in $\text{relconv}(\{p, q\})$. For the converse, let S' and S'' be pseudo-circles containing respectively p but not q and q but not p .

3. It is clear that $\text{conv}(p, q, r) \subset \overline{S_e^+} \cap \overline{S_f^+} \cap \overline{S_g^+}$. Also, since any arrangement of three pseudo circles in S^2 is homeomorphic to the arrangement of three great circles with no common point, $\overline{S_e^+} \cap \overline{S_f^+} \cap \overline{S_g^+}$ is a 2-ball whose interior is $S_e^+ \cap S_f^+ \cap S_g^+$. It suffices to show that $\overline{S_e^+} \cap \overline{S_f^+} \cap \overline{S_g^+} \subset \text{conv}(\{p, q, r\})$. In other words, that if $\overline{S^+}$ is a closed hemisphere containing $\{p, q, r\}$, then $\overline{S_e^+} \cap \overline{S_f^+} \cap \overline{S_g^+} \subset \overline{S^+}$.

By part 1, the three closed arcs $\text{conv}(p, q)$, $\text{conv}(p, r)$ and $\text{conv}(q, r)$ are contained in $\overline{S^+}$. These three arcs are the boundary of the 2-ball $S_e^+ \cap S_f^+ \cap S_g^+$. Since $\overline{S_e^+} \cap \overline{S_f^+} \cap \overline{S_g^+} \cap \overline{S^+}$ has to be a ball or a sphere and contains the boundary of the 2-ball $\overline{S_e^+} \cap \overline{S_f^+} \cap \overline{S_g^+}$ either it equals the whole 2-ball (and we have finished) or it equals its boundary. However, the latter is only possible if S contains p, q and r , which is not the case (the intersection of two closed 2-balls in S^2 is a circle only if this circle is the boundary of the two 2-balls). □

Remark A.10. Parts 1 and 2 of the previous lemma have the following easy consequence: Given an edge $l = \{p, q\} \subset S$ of \mathcal{B} and two pseudo-circles S_p, S_q such that $p \in S_p \cap S_q^+$ and $q \in S_q \cap S_p^+$. Then $\text{relconv}(l) = S \cap S_p^+ \cap S_q^+$ and $\text{conv}(l) = S \cap \overline{S_p^+} \cap \overline{S_q^+}$.

Also, for any pseudo-circle $S' \in \Lambda$ other than S , $S' \cap \text{conv}(\{p, q\})$ has at most one point. Indeed, if p and q lie on the same closed side of S' , then the previous lemma implies that $S' \cap \text{conv}(\{p, q\})$ contains at most p and q , and it cannot contain both since p and q are not antipodal. If p and q lie in opposite open sides of S' , then let x and y be the two antipodal intersection points of S and S' . Since p and q lie in different components of $S \setminus S'$, x and y lie in different components of $S \setminus \{p, q\}$ and hence only one of them can be in $\text{conv}(\{p, q\})$, again by the previous lemma.

The following statement is probably true for pseudo-configurations of points of arbitrary rank.

Lemma A.11. *Let σ and τ be two disjoint simplices of a pseudo-configuration of points \mathcal{B} in S^2 . The following conditions are equivalent:*

1. $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$.
2. (σ, τ) is a vector (in the oriented matroid $\mathcal{M}_{\mathcal{B}}$ of \mathcal{B}).

Proof. (1 \Rightarrow 2) First we suppose $\sigma = \{p\}$ is a single point. Then τ is an edge or a triangle. If $\tau = \{q, r\}$ is an edge, then let S, S_q and S_r be pseudo-circles with $\{q, r\} \subset S$, $q \in S_q \setminus S_r$ and $r \in S_r \setminus S_q$, so that $\text{relconv}(\tau) = S \cap S_q^+ \cap S_r^+$ for the appropriate choice of the sign for the sides of S_q and S_r . Then $p \in \text{relconv}(\tau) \subset S$. Hence $\{p, q, r\}$ has rank 2 and contains the support of a circuit. The only way for this circuit to be orthogonal to the cocircuits defined by S_q and S_r is $(\{p\}, \{q, r\})$. In the same way, if τ is a triangle, then $p \cup \tau$ contains the support of a circuit and the only possibility for this circuit to be orthogonal to the cocircuits defined by the three sides of τ is $(\{p\}, \tau)$.

Finally, for general σ and τ we consider any extension \mathcal{B}' of \mathcal{B} by a point $p \in \text{relconv}(\tau) \cap \text{relconv}(\sigma)$, as in Proposition A.5. Then (p, σ) and (p, τ) are vectors in the extended oriented matroid $\mathcal{M}_{\mathcal{B}'}$ and, hence, (σ, τ) is a vector in $\mathcal{M}_{\mathcal{B}'}$ and in $\mathcal{M}_{\mathcal{B}}$.

(2 \Rightarrow 1) Since every vector is a composition of circuits, there is one circuit (C^+, C^-) with $C^+ \subset \sigma$ and $C^- \subset \tau$. Since σ and τ are simplices, $C^+ \neq \emptyset \neq C^-$. We consider separately the following cases:

In case σ (or τ) is a vertex $\{p\}$, the orthogonality between the vector $(\{p\}, \tau)$ and the cocircuits which define $\text{relconv}(\tau)$ implies trivially that $p \in \text{relconv}(\tau)$.

We assume now that $\sigma = \{p, q\}$ and $\tau = \{r, s\}$ are both edges. Let S_σ and S_τ be the unique pseudo-circles containing σ and τ , respectively. If $S_\tau = S_\sigma$, then $\sigma \cup \tau$ has rank 2 and every circuit with support contained in $\sigma \cup \tau$ has at most three elements. Thus, one of C^+ and C^- has only one element (say $C^+ = \{p\}$) and the other has two (i.e. $C^- = \tau$). However, this implies that p is in the open arc $\text{relconv}(\tau)$ and, hence, the two open arcs $\text{relconv}(\sigma)$ and $\text{relconv}(\tau)$ in $S_\sigma = S_\tau$ intersect.

If $S_\tau \neq S_\sigma$, then orthogonality between the vector (σ, τ) and the cocircuits defined by S_τ and S_σ implies that τ has exactly one point on each side of S_σ , σ one point on each side of τ and thus $C^+ = \sigma$ and $C^- = \tau$. In particular, $S_\sigma \cap \text{relconv}(\tau) \neq \emptyset$ and $S_\tau \cap \text{relconv}(\sigma) \neq \emptyset$.

Let x be the point (not necessarily in \mathcal{B}) in $S_\tau \cap \text{relconv}(\sigma) \subset S_\tau \cap S_\sigma$. We claim that $x \in \text{relconv}(\tau)$. Otherwise, $y \in \text{relconv}(\tau)$ for the only point y other than x in $S_\tau \cap S_\sigma$. However, in this case in the extension of \mathcal{B} by the points x and y we have the following circuits supported on $\{p, q, r, s, x, y\}$:

$$\begin{aligned} & (+, +, 0, 0, -, 0), \\ & (0, 0, +, +, 0, -), \\ & (0, 0, 0, 0, +, +). \end{aligned}$$

By eliminating x and y we get the vector $(+, +, +, +, 0, 0)$. Then $\sigma \cup \tau$ is the support of two different and not opposite vectors, and eliminating in them we would get a circuit supported on three of the four points of $\sigma \cup \tau$. This contradicts the fact that $(C^+, C^-) = (\sigma, \tau)$ is a circuit.

In the cases not considered yet, $\sigma \cup \tau$ has at least five elements. We have the following possibilities:

- C^+ is a point p and C^- is a triangle $\{q, r, s\}$. Then $\text{relconv}(\tau)$ is an open subset of S^2 . By the first case studied, $p \in \text{relconv}(\tau)$, which implies that $\text{conv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$ and since $\text{conv}(\sigma)$ is the topological closure of $\text{relconv}(\sigma)$, $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$.
- Both C^+ and C^- are edges. By the second case studied, $\text{relconv}(C^+)$ and $\text{relconv}(C^-)$ are open edges which cross each other transversally. No matter whether σ and τ coincide with C^+ and C^- or have one more element, it is clear that $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$.
- $C^+ = \{p\}$ is a point and $C^- = \{q, r\}$ is an edge. Let $S_{\{q,r\}}$ be the unique pseudo-circle containing C^- . Again, by the first case studied, $p \in \text{relconv}(\{q, r\})$. First suppose $\tau = \{q, r\}$. The cases in which σ is a point or an edge have already been discussed, so assume it is a triangle. Orthogonality of (σ, τ) with the circuit defined by $S_{\{q,r\}}$ implies that σ has one point on each side of $S_{\{q,r\}}$ which, together with the fact that p is in the open segment from q to r , makes it clear that $\text{conv}(\tau) \cap \text{relconv}(\sigma) \neq \emptyset$. Since $\text{relconv}(\sigma)$ is open and $\text{conv}(\tau)$ is the closure of $\text{relconv}(\tau)$, we have $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$.

Finally suppose $\tau = \{q, r, s\}$ for some s . The argument is almost identical to the previous one. In this case, orthogonality of (σ, τ) with the cocircuit defined by $S_{\{q,r\}}$ implies that σ has some point in the open side of $S_{\{q,r\}}$ containing s . Since p is a point in the open segment from q to r , it is clear then that $\text{conv}(\sigma)$ intersects the open triangle $\text{relconv}(\tau)$. Since $\text{relconv}(\tau)$ is open and $\text{conv}(\sigma)$ is the closure of $\text{relconv}(\sigma)$, $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$. □

A.3. How to Adapt Sections 1–4 to Non-Realizable Oriented Matroids

In what follows we show that all the results in the previous sections hold for non-realizable oriented matroids as well, with only some changes in the language to be used.

Section 1

We start by giving the definition of triangulation of an oriented matroid.

Definition A.12 [4, Definition 9.6.1]. Let \mathcal{M} be an oriented matroid of rank k . Let \mathcal{T} be a collection of bases of \mathcal{M} . \mathcal{T} is a triangulation of \mathcal{M} if the following properties are satisfied:

1. Any pair σ, τ of elements of \mathcal{T} *intersect properly* meaning by this that for any single element extension $\mathcal{M} \cup p$ of \mathcal{M} , if there are $\sigma_0 \subset \sigma$ and $\tau_0 \subset \tau$ such that $(\{p\}, \sigma_0)$ and $(\{p\}, \tau_0)$ are circuits of $\mathcal{M} \cup p$, then there is also a $\rho \subset \sigma \cap \tau$ such that $(\{p\}, \rho)$ is a circuit as well.
2. For every independent subset τ of size $k - 1$ such that there is a $\sigma \in \mathcal{T}$ with $\tau \subset \sigma$, either τ is in a facet of \mathcal{M} (i.e. there is a nonnegative cocircuit vanishing on τ) or there are at least two bases in \mathcal{T} containing τ .

In Theorem 2.4 of [16] it is proved that:

Lemma A.13. *Let \mathcal{M} be an oriented matroid of rank k . Let \mathcal{T} be a collection of bases of \mathcal{M} . Then \mathcal{T} is a triangulation of \mathcal{M} if and only if:*

1. *For every independent subset τ of size $k - 1$ such that there is a $\sigma \in \mathcal{T}$ with $\tau \subset \sigma$, either*
 - *τ is contained in a facet of \mathcal{M} (i.e. there is a positive cocircuit vanishing on τ),*
 - *or*
 - *there is exactly another $\sigma' \in \mathcal{T}$ with $\tau \subset \sigma'$ and σ and σ' are in opposite sides of τ (i.e. the unique cocircuit vanishing on τ has opposite sign at the elements $\sigma \setminus \tau$ and $\sigma' \setminus \tau$).*
2. *There is a single element extension $\mathcal{M} \cup p$ of \mathcal{M} such that exactly one element σ of \mathcal{T} has $p \in \text{conv}_{\mathcal{M} \cup \{p\}}(\sigma)$ (meaning that $(\{p\}, \sigma)$ is a circuit of $\mathcal{M} \cup p$).*

Moreover, any triangulation \mathcal{T} of \mathcal{M} covers \mathcal{M} meaning by this that if $\mathcal{M} \cup p$ is a single element extension in which there is a circuit of the form $(\{p\}, \tau)$, then there exists another circuit $(\{p\}, \sigma)$ of $\mathcal{M} \cup p$ where σ is a subset of an element of \mathcal{T} .

It would be interesting for us to find a characterization of triangulations for a rank 3 oriented matroid pseudo-realized as a pseudo-configuration of points \mathcal{B} .

Lemma A.14. *Let \mathcal{M} be a simple rank 3 oriented matroid and let \mathcal{B} be a pseudo-configuration of points in the sphere S^2 which pseudo-realizes \mathcal{B} . A collection \mathcal{T} of triangles of \mathcal{M} (or of \mathcal{B}) is a triangulation of \mathcal{M} if and only if:*

1. *\mathcal{T} realizes geometrically as a simplicial complex in the sphere, i.e. for any pair of triangles $\sigma, \tau \in \mathcal{T}$ one has*

$$\text{conv}(\sigma) \cap \text{conv}(\tau) = \text{conv}(\sigma \cap \tau).$$

2. *\mathcal{T} covers the “convex hull” of the pseudo-configuration \mathcal{B} , i.e.*

$$\bigcup_{\sigma \in \mathcal{T}} \text{conv}(\sigma) = \bigcup_{\sigma \text{ is a triangle of } \mathcal{B}} \text{conv}(\sigma).$$

Proof. We first see that if \mathcal{T} is a triangulation of \mathcal{M} , then it satisfies 1 and 2. In fact, 1 follows from axiom 1 in Definition A.12: if there is a point $p \in \text{conv}(\sigma) \cap \text{conv}(\tau) \setminus \text{conv}(\sigma \cap \tau)$, then this point provides an extension (via Proposition A.5) which violates the axiom. In the same way, property 2 follows from the final part of Lemma A.13.

Conversely, let \mathcal{T} be in the conditions of the statement and let us see that \mathcal{T} satisfies 1 and 2 of Lemma A.13. Statement 2 is easy: any point $p \in S^2$ in the relative interior of a triangle of \mathcal{T} provides an extension of the oriented matroid in the required conditions.

To prove 1, let $\tau = \{p, q\}$ be an edge of a triangle $\sigma = \{p, q, r\} \in \mathcal{T}$, and suppose that τ is not in a facet of \mathcal{M} . This means that there is a point $s \in \mathcal{B}$ such that r and s lie in opposite sides of the pseudo-circle containing τ . Let $a \in S^2$ be a point “very close” to the relative interior of τ and on the side on which s is. This point is in $\text{relconv}(\{p, q, s\})$ and, hence, by condition 2 in the statement, there is a $\sigma' \in \mathcal{T}$ such that $a \in \text{conv}(\sigma')$. The only way in which σ and σ' can intersect properly in the sense of condition 1 in the statement is that $\tau \subset \sigma'$. Thus, there is a $\sigma' \in \mathcal{T}$ with $\tau \subset \sigma'$ and with σ and σ' in opposite sides of τ . Finally, there cannot be any other $\sigma'' \in \mathcal{T}$ with $\tau \subset \sigma''$ because then condition 1 would not be fulfilled either for σ and σ'' or for σ' and σ'' . \square

Remark A.15. In fact $\bigcup_{\sigma \text{ is a triangle of } \mathcal{B}} \text{conv}(\sigma) = \text{conv}(\mathcal{B})$, but we do not make use of this assertion.

It is still true that an acyclic circuit of an oriented matroid can be triangulated in exactly two ways and thus we can define virtual chambers of \mathcal{M} using Definition 1.1 with the only substitution of “full-dimensional simplices of \mathcal{A} ” by “bases of \mathcal{M} .” With this, Theorem 1.3 is proved in Theorem 3.8 of [16] for non-realizable oriented matroids.

Lemma 1.4 says that the relative interiors of any two simplices of a virtual chamber intersect. This statement makes sense for an oriented matroid if it is pseudo-realized as a pseudo-configuration of points \mathcal{B} and the relative interior of a simplex is defined as in Definition A.8. However, the proof given in Section 1 is not valid for non-realizable oriented matroids (this is related to Remark 2.5(v) of [16] where it is said not to be known whether a circuit can have its positive and negative parts respectively contained in two simplices of a triangulation). Thus, we provide a new proof for it:

Lemma A.16 (Lemma 1.4 for Oriented Matroids). *Let \mathcal{M} be a rank 3 oriented matroid pseudo-realized by a point configuration \mathcal{B} in S^2 . Let \mathcal{C} be a virtual chamber of \mathcal{M} . Then, for any pair of simplices σ, τ in \mathcal{C} , the relative interiors $\text{relconv}(\sigma)$ and $\text{relconv}(\tau)$ intersect.*

Proof. Let \mathcal{T} be the triangulation which corresponds to \mathcal{C} in the dual oriented matroid \mathcal{M}^* . The link $\mathcal{T}' := \text{link}_{\mathcal{T}}[\mathcal{B} \setminus (\sigma \cup \tau)]$ of $\mathcal{B} \setminus (\sigma \cup \tau)$ in \mathcal{T} is a triangulation of the contraction $\mathcal{M}^*/[\mathcal{B} \setminus (\sigma \cup \tau)]$ [16, Theorem 2.4(e)]. Thus, the set \mathcal{C}' of triangles of \mathcal{B} complementary to simplices of \mathcal{T}' is a virtual chamber of the oriented matroid \mathcal{M} restricted to $\sigma \cup \tau$. It is clear that \mathcal{C}' consists of those triangles of \mathcal{C} contained in $\sigma \cup \tau$. Since the restriction of \mathcal{M} is a realizable oriented matroid (it has at most six elements), Lemma 1.4 holds for some realization of it and in this realization $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$ is equivalent to (σ, τ) being a vector in the oriented matroid. However, then (σ, τ) is a vector in the orig-

inal oriented matroid \mathcal{M} too, and Lemma A.11 implies that $\text{relconv}(\sigma) \cap \text{relconv}(\tau) \neq \emptyset$ in the pseudo-realization \mathcal{B} . □

Again, we can use Definition 1.6 of flips in a triangulation as it is for oriented matroids. This is equivalent to the definition given in Definition 3.11 of [16] as shown in Definition 3.13 of [16]. Proposition 1.7 and Theorem 1.8 hold in this setting since their proofs only use oriented matroid duality and the fact that any link in a triangulation is a triangulation of the contraction, which we have already used and is proved in Theorem 2.4(d) of [16]. Lemma 1.9 holds in the oriented matroid setting trivially, since every rank 2 oriented matroid is realizable.

As in Section 1, a simplex σ in a pseudo-configuration of points \mathcal{B} is said to be empty if $\text{conv}(\sigma) \cap \mathcal{B} = \sigma$. With this, Definition 1.10 of a flip of a rank 3 virtual chamber on an empty edge, and Corollary 1.11 hold without change. Lemma 1.12 has already been proved as Lemma A.4, and Definition 1.13 comes from general oriented matroid theory.

Summing up, we have defined virtual chambers of an oriented matroid so that they are in (flip-preserving) bijection with triangulations of the dual oriented matroid. In particular, for a corank 3 oriented matroid \mathcal{M} , triangulations of \mathcal{M} are in one-to-one (flip-preserving) correspondence with virtual chambers of the dual oriented matroid \mathcal{M}^* , which realizes as a pseudo-configuration of points in S^2 .

Section 2

As in Section 2, given two edges l_1 and l_2 of a pseudo-configuration of points and an empty triangle τ , we say that l_1 and l_2 cross each other if (l_1, l_2) is a circuit and that l_1 crosses τ if it crosses some edge of τ .

Corollary A.17. *Let \mathcal{B} be a pseudo-configuration of points in S^2 . Let σ be a simplex and let l_1 and l_2 be two edges. Then:*

1. σ is empty if and only if there is no circuit Z such that $Z^+ \subset \sigma$ and Z^- is a single point.
2. (l_1, l_2) is a circuit (i.e. l_1 and l_2 cross each other) if and only if $\text{relconv}(l_1)$ and $\text{relconv}(l_2)$ intersect in a single point.
3. If l_1 and l_2 are empty edges and $\text{relconv}(l_1) \cap \text{relconv}(l_2) \neq \emptyset$, then (l_1, l_2) is a circuit.

Proof. 1. If $(Z^+, \{p\})$ is a circuit, then $p \in \text{relconv}(Z^+) \subset \text{conv}(\sigma)$. On the other hand, $p \notin \sigma$ since $Z^+ \cup \{p\} \subset \sigma \cup \{p\}$ contains the support of a circuit. Reciprocally, if σ is not empty then let $p \notin \sigma$ be in $\text{conv}(\sigma)$. From the geometric description of the convex hull of a triangle and an edge given in Lemma A.9 it follows that $p \in \text{relconv}(\tau)$ for some $\tau \subset \sigma$. Then, by Lemma A.11, we have that $(\tau, \{p\})$ is a vector and, since τ is independent, it is a circuit.

2. If (l_1, l_2) is a circuit, then it is a vector and the pseudo-circles containing l_1 and l_2 are distinct. Then $\text{relconv}(l_1) \cap \text{relconv}(l_2) \neq \emptyset$ by Lemma A.11 and it has only one point by Remark A.10. Conversely, if $\text{relconv}(l_1) \cap \text{relconv}(l_2) \neq \emptyset$, then (l_1, l_2) is a vector by Lemma A.11. It is a circuit unless $l_1 \cup l_2$ has rank 2, in which case l_1 and l_2 are

contained in a pseudo-circle S . If two arcs in a pseudo-circle intersect in their relative interiors, then they intersect in more than one point.

3. As in part 2, $\text{relconv}(l_1) \cap \text{relconv}(l_2) \neq \emptyset$ implies that (l_1, l_2) is a circuit unless l_1 and l_2 are contained in a pseudo-circle S and overlap in their relative interiors. However, this implies that one of them is not empty. \square

In the statement of Lemma 2.2 (and in what follows) the sentence “two vertices p and q lie on opposite sides of an edge l ” has to be understood as “the unique (up to sign reversal) cocircuit vanishing on l has opposite signs at p and q ” (which is what is shown in the proof of Lemma 2.2) or equivalently as “ p and q lie on opposite sides of the unique pseudo-circle containing the edge l .” With this in mind, the proofs of Lemma 2.2 and Propositions 2.3 and 2.4 are valid for pseudo-configurations without change since they only use topological or matroidal tools.

For Lemma 2.5 it suffices to observe that it involves at most six points, so the restricted oriented matroid can be realized. We just apply Lemma 2.5 to a realization and observe that the resulting circuit must be a circuit of the whole pseudo-configuration too.

For Definition 2.6 to make sense we need to prove the following:

Lemma A.18. *Let $l = \{p, q\}$ be an empty edge of a pseudo-configuration of points \mathcal{B} in S^2 and let s, t, u be three points such that the edges $r = \{s, t\}$ and $r' = \{t, u\}$ cross l . Then the following are equivalent:*

1. u and p are on the same side of r .
2. The intersection point $\text{relconv}(r') \cap \text{relconv}(l)$ is closer to p (along l) than the point $\text{relconv}(r) \cap \text{relconv}(l)$.

Proof. Let $p' = \text{relconv}(r') \cap \text{relconv}(l)$ and consider an extension of \mathcal{B} by a point at p' as given by Proposition A.5.

By axiom 2 of Definition A.1, p' is closer to p (along l) than the point $\text{relconv}(r) \cap \text{relconv}(l)$ if and only if p' and p are on the same side of the pseudo-sphere containing r , i.e. in the same side of r . On the other hand, p' and u are clearly on the same side of r , so that p' and p are on the same side of r if and only if u and p are on the same side of r . \square

With this, the proof of Corollary 2.7 needs no changes.

Section 3

Definition 3.1 adapts without problem to pseudo-configurations of points: a virtual chamber \mathcal{C} is said to lie on a certain side of an edge l if there is a triangle in \mathcal{C} contained in one (closed) side of the pseudo-circle containing l .

The proof of Proposition 3.2 is based in the fact that if \mathcal{B}_0 is a subconfiguration of a configuration \mathcal{B} and we have a triangulation \mathcal{T}_0 of \mathcal{B}_0 , then there is a triangulation \mathcal{T} of \mathcal{B} with $\mathcal{T}_0 \subset \mathcal{T}$. This holds in the oriented matroid setting substituting \mathcal{B}_0 and \mathcal{B} by a restriction \mathcal{M}_0 of an oriented matroid \mathcal{M} , as shown in Corollary 2.11 of [16].

The proof of Theorem 3.3 is completely based on dealing with circuits and cocircuits, and thus it is valid for oriented matroids without change.

For the last results of Section 3 we need to find an analogy in the oriented matroid setting of a geometric chamber. The natural one is the following:

Definition A.19. Let \mathcal{B} be a pseudo-configuration of points in the sphere S^2 . A point $p \in S^2$ is said to be *interior and generic* if it lies in the convex hull of some triangle of \mathcal{B} but not in the convex hull of any edge of \mathcal{B} . The *chamber* \mathcal{C}_p associated to an interior generic point p is the collection of triangles of \mathcal{B} which have p in their convex hull.

Geometrically, the chamber \mathcal{C} corresponding to a point p can be thought of as the region of $\text{conv}(\mathcal{B})$ obtained by the intersection of the convex hulls of all the triangles of \mathcal{C} . As in the realizable case, we can therefore define the *chamber complex* of \mathcal{B} as the coarsest common refinement of all the triangulations of \mathcal{B} or, equivalently, as the decomposition of $\text{conv}(\mathcal{B})$ defined by the chambers of \mathcal{B} . That this decomposition is a cellular decomposition follows from the following properties of a chamber \mathcal{C} of \mathcal{B} (considered as a closed region in S^2):

- By definition \mathcal{C} is the intersection of certain closed hemispheres. By axiom 3 of Definition A.1, \mathcal{C} is either a sphere or a closed ball. Since there exists a generic interior point p which is in the relative interior of every triangle of \mathcal{C} and since there are finitely many triangles in \mathcal{C} , we have that \mathcal{C} has a non-empty interior and hence it has dimension 2.
- Since \mathcal{C} is contained in the convex hull of any of its triangles, \mathcal{C} cannot contain a pair of antipodal points and in particular cannot contain an entire pseudo-circle. Thus, \mathcal{C} is a 2-dimensional ball.
- Let S_1, \dots, S_k be a family of pseudo-circles of \mathcal{B} oriented so that the chamber \mathcal{C} equals $\bigcap_{i=1}^k \overline{S_i^+}$ in an irredundant way. Then the boundary of the chamber \mathcal{C} (which is homeomorphic to S^1 by the previous point) is a union of k closed arcs each contained in one of the S_i 's.
- Let \mathcal{C} and \mathcal{C}' be two chambers. If their interiors have a common point p , then this point is clearly interior and generic and hence every triangle in \mathcal{C} or in \mathcal{C}' is also in \mathcal{C}_p . Since a chamber cannot properly contain another one, we conclude that $\mathcal{C} = \mathcal{C}'$.
- Let \mathcal{C} and \mathcal{C}' be two different chambers and suppose that their boundaries intersect in a non-empty set I . I is an intersection of closed sides and, hence, it is a ball or a sphere of dimension at most 1. If it has dimension 1, then let \mathcal{C} be expressed as an irredundant intersection of closed hemispheres, as above. One of the pseudo-circles S_i in this expression intersects I in a one-dimensional arc, and \mathcal{C} and \mathcal{C}' have to be contained respectively in the two opposite closed hemispheres defined by S_i . In particular, $I \subset S_i$ and I is a closed arc.

If the intersection has dimension 0, then we see that it has only one point. Let x be one of the intersection points and let S_1^+, \dots, S_k^+ be the set of all the open hemispheres such that $x \in S_i$ and one of the chambers \mathcal{C}_1 or \mathcal{C}_2 is contained in $\overline{S_i^+}$ for every $i = 1, \dots, k$. Let $I' = \bigcap_{i=1}^k \overline{S_i^+}$. I' contains I , I' contains the antipodal point y to x and is an intersection of closed hemispheres. Since I' is a sphere or a ball, either $I' = \{x, y\}$, in which case $I = \{x\}$ because I cannot contain antipodal pairs, or I' contains an arc joining x to y . However, in the second case, since I is the

intersection of I' with some closed pseudo-spheres all of which contain x in their interior, I contains a subarc of this arc and, in particular, I is not zero-dimensional.

For the same reasons as in the realizable case, every chamber is in particular a virtual chamber and edges in the chamber complex correspond to flips (using the characterization in Lemma A.14 for triangulations of a pseudo-configuration of points \mathcal{B}). Moreover, flag chambers of the oriented matroid \mathcal{M} appear as chambers in any pseudo-realization of it. With this, the proofs of Proposition 3.4, Corollary 3.5, Lemmas 3.6–3.8 and Theorem 3.9 are valid in the context of non-realizable oriented matroids.

For Corollary 3.10, we state and prove the new version of it as follows:

Corollary A.20. *Let \mathcal{M}^* be a corank 3 oriented matroid. Then:*

1. *The graph of triangulations of \mathcal{M}^* is connected.*
2. *If \mathcal{M}^* is acyclic, then the graph is 3-connected.*

Proof. By Lemma A.4, without loss of generality we can assume that the dual oriented matroid \mathcal{M} of \mathcal{M}^* is simple. Let \mathcal{B} be a pseudo-realization of \mathcal{M}^* in S^2 . By Theorem 3.9, any virtual chamber of \mathcal{B} which is not a chamber can be joined through three vertex-disjoint paths to three distinct chambers. On the other hand, any two chambers can be connected by a sequence of flips since $\text{conv}(\mathcal{B})$ is either S^2 or homeomorphic to a disk and since two chambers differ by a flip if and only if they are adjacent. Moreover, if \mathcal{M}^* is acyclic, then $\text{conv}(\mathcal{B}) = S^2$ and we prove below that the subgraph of $G(\mathcal{A})$ induced by triangulations which appear as chambers of the chamber complex is 3-connected (this graph is the adjacency graph of the chamber complex). As in the realizable case, this implies that the graph $G(\mathcal{A})$ is 3-connected.

We now prove that the adjacency graph of the chamber complex is 3-connected, i.e. that it has at least four vertices and that it remains connected when we remove any two chambers. Since \mathcal{B} is totally cyclic, any triangulation \mathcal{T} of \mathcal{B} has at least four triangles and no two of them belong to the same chamber (by Definition 1.1). Thus, there are at least four chambers. Since the intersection of any two closed chambers is empty, a point or a 1-ball, removing them from the chamber complex leaves something connected in the sphere, homeomorphic to an open 2-ball or an open annulus. In particular, the adjacency graph of the chamber complex remains connected when removing two of its vertices. □

Section 4

Our next (and final) goal is to show that, for a pseudo-configuration of points \mathcal{B} in S^2 , pseudo-chamber complexes can be defined and every virtual chamber realizes as a pseudo-chamber of some pseudo-chamber complex of \mathcal{B} . Definitions 4.1 and 4.2 can be naturally translated into pseudo-configurations of points. Now we observe that the collection of arcs $\{\text{conv}(\{p, q\}) : \{p, q\} \text{ is an edge of } \mathcal{B}\}$ defines a pseudo-chamber complex of \mathcal{B} . Thus, throughout Section 4 we can assume that \mathcal{B} is a pseudo-configuration of points in S^2 given with an initial pseudo-chamber complex and this initial pseudo-chamber complex will take the role of the chamber complex of a vector configuration.

Taking into account the properties of pseudo-configurations of points enumerated in Lemma A.9, Lemma A.11 and Corollary A.17, all the notions and proofs provided in Section 4 are valid for pseudo-configurations of points in S^2 . For example, in the proof of Proposition 4.6 (triangulations of an oriented matroid are the same as triangulations of any of its pseudo-chamber complexes) the result for the initial pseudo-chamber complex follows from Lemma A.14.

References

1. L. J. Billera, P. Filliman and B. Sturmfels, Constructions and complexity of secondary polytopes, *Adv. Math.* **83** (1990), 155–179.
2. L. J. Billera, I. M. Gel'fand and B. Sturmfels, Duality and minors of secondary polyhedra, *J. Combin. Theory Ser. B* **57** (1993), 258–268.
3. A. Björner, Topological methods, in *Handbook of Combinatorics* (R. L. Graham, M. Grötschel and L. Lovász, eds.), Elsevier, Amsterdam, 1995, pp. 1819–1872.
4. A. Björner, M. Las Vergnas, B. Sturmfels, N. White and G. M. Ziegler, *Oriented Matroids*, Cambridge University Press, Cambridge, 1992.
5. J. de Loera, S. Hoşten, F. Santos and B. Sturmfels, The polytope of all triangulations of a point configuration, *Doc. Math. J. DMV* **1** (1996), 103–119.
6. J. de Loera, F. Santos and J. Urrutia, The number of geometric bistellar neighbours of a triangulation, *Discrete Comput. Geom.* **21** (1999), 131–142.
7. P. H. Edelman and V. Reiner, Visibility complexes and the Baues problem for triangulations in the plane, *Discrete Comput. Geom.* **20** (1998), 35–59.
8. I. M. Gel'fand, M. M. Kapranov and A. V. Zelevinsky, *Multidimensional Determinants, Discriminants and Resultants*, Birkhäuser, Boston, 1994.
9. L. Glaser, *Geometrical Combinatorial Topology*, Van Nostrand Reinhold, New York, 1970.
10. M. Las Vergnas, Extensions ponctuelles d'une géométrie combinatoire orientée, in *Problèmes Combinatoires et théorie des graphes* (Actes Coll. Orsay, 1976), Colloques Internationaux, C.N.R.S., No. 260, pp. 265–270.
11. C. L. Lawson, Software for C^1 -interpolation, in *Mathematical Software III* (John Rice, ed.), Academic Press, New York, 1977, pp. 161–194.
12. C. W. Lee, Regular triangulations of convex polytopes, in *Applied Geometry and Discrete Mathematics—The Victor Klee Festschrift* (P. Gritzmann and B. Sturmfels, eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 4, American Mathematical Society, Providence, RI, 1991, pp. 443–456.
13. J. Rambau, Triangulations of cyclic polytopes and higher Bruhat orders, *Mathematika* **44** (1997), 162–194.
14. V. Reiner, The generalized Baues problem, in *New Perspectives in Algebraic Combinatorics* (L. J. Billera, A. Björner, C. Greene, R. E. Simion and R. P. Stanley, eds.), MSRI Publications, No. 38, Cambridge University Press, Cambridge, 1999.
15. T. B. Rushing, *Topological Embeddings*, Academic Press, New York, 1973.
16. F. Santos, Triangulations of oriented matroids, *Mem. Amer. Math. Soc.*, accepted.
17. F. Santos, A point set whose space of triangulations is disconnected, Preprint, 1999.
18. F. Santos, Triangulations with very few geometric bistellar neighbors, *Discrete Comput. Geom.*, to appear.
19. G. M. Ziegler, *Lectures on Polytopes*, Springer-Verlag, New York, 1994.

Received March 1, 1999, and in revised form September 7, 1999.