Configurations, Triangulations, Subdivisions, and Flips

2

The first goal of this chapter is to introduce the necessary mathematical language to work with triangulations. This language includes the geometry of polyhedra and cones [339] and the combinatorics of point and vector configurations, as described by their oriented matroids [55]. These two books are recommended for more details. The second goal is to provide the reader with formal definitions and notation that are strong but flexible enough to cover all kinds of point configurations, and even the more general case of vector configurations. These definitions should include degeneracies, such as collinearities and repetition of points. We will do this slowly, intending to help the reader to see why a naive definition may lead to problems.

2.1 The official languages in the land of triangulations

When one looks at their purely geometric aspects, triangulations are made of polyhedra and thus they are described by convex geometry concepts like polytopes, cones, hyperplanes, etc. But when one cares about data structures, or if one is interested in combinatorial aspects, then combinatorics comes into play, and one is concerned about such things as labels, sign vectors, simplicial complexes, posets, etc. In this chapter we develop the "bilingual" setting that we use throughout the book.

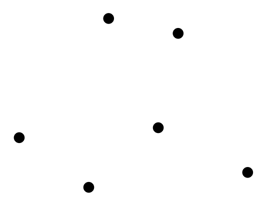

Figure 2.1: Points forming a set $\mathbf{X}$....

2.1.1 Polyhedra and cones

A *convex combination* of a finite set of points $\mathbf{x}_1, \ldots, \mathbf{x}_k$ in $\mathbb{R}^m$ is any point $\mathbf{x}$ that can be expressed as

$$\sum_{i=1}^{k} \lambda_i \mathbf{x}_i$$

with the $\lambda_i \in \mathbb{R}$ non-negative and summing to one. If non-negativity is dropped, then $\mathbf{x}$ is an *affine combination*.

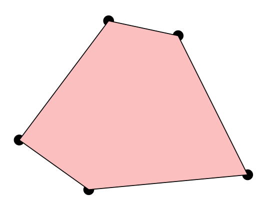

Figure 2.2: ... and the convex hull of $\mathbf{X}$.

The *convex hull* of a set $\mathbf{X} \subset \mathbb{R}^m$, denoted $\operatorname{conv}(\mathbf{X})$, is the intersection of all convex sets containing $\mathbf{X}$. Equivalently, a point $\mathbf{x}$ is in $\operatorname{conv}(\mathbf{X})$ if it is a convex combination of some finite subset of points in $\mathbf{X}$. Similarly, the set of all affine combinations of points in $\operatorname{conv}(\mathbf{X})$ is called the *affine hull* or *affine span* of $\mathbf{X}$. Affine spans are always *affine subspaces*, also called *flats* of $\mathbb{R}^m$. That is, every affine span is a translated copy of a linear subspace (conversely, linear subspaces are affine subspaces passing through the origin).

The convex hull of finitely many points is called a *convex polytope*, but we will usually just call it a *polytope* since we will very rarely be concerned with non-convex ones (which we do not even define here). A *face* of a

Jesús A. De Loera et al., *Triangulations*, Algorithms and Computation in Mathematics,
DOI 10.1007/978-3-642-12971-1_2, © Springer-Verlag Berlin Heidelberg 2010

polytope $\mathbf{P}$ is the locus $\{\mathbf{x} \in \mathbf{P} : \psi(\mathbf{x}) \geq \psi(\mathbf{y}),\ \forall \mathbf{y} \in \mathbf{P}\}$ where a certain linear form $\psi : \mathbb{R}^m \to \mathbb{R}$ is maximized (or minimized). The whole $\mathbf{P}$ is a face (obtained with $\psi = 0$) and, for convenience, the empty set is also accepted as a face. It is easy to prove that faces of a polytope are themselves polytopes and that faces of a face of $\mathbf{P}$ are faces of $\mathbf{P}$ as well. The *dimension* of a polytope, or of a face, is the dimension of its affine span. Faces of dimension zero, one, and $\dim(\mathbf{P}) - 1$ are called *vertices*, *edges*, and *facets* of $\mathbf{P}$, respectively.

Remark 2.1.1. Of course, if a polytope $\mathbf{P} \subseteq \mathbb{R}^m$ has dimension d, we can always find a projection $\mathbb{R}^m \to \mathbb{R}^d$ that "embeds" $\mathbf{P}$ in dimension d keeping all its properties of interest (face lattice, triangulations, etc). But sometimes it is easier to define, or study, a d-dimensional polytope as lying in a proper flat of a higher dimensional space. For example, the standard d-dimensional simplex is better represented in $\mathbb{R}^{d+1}$ as the intersection of the positive orthant with the hyperplane of coordinate sum equal to one. This representation preserves its symmetries. Also, one of the most important polytopes in this book, the *secondary polytope* of a point configuration, comes naturally embedded in a dimension much higher than its intrinsic dimension.

For this and other reasons (see Section 2.1.2, for example) we will make a distinction between the intrinsic dimension of $\mathbf{P}$, typically denoted d, and the ambient dimension, for which we usually reserve the letter m.

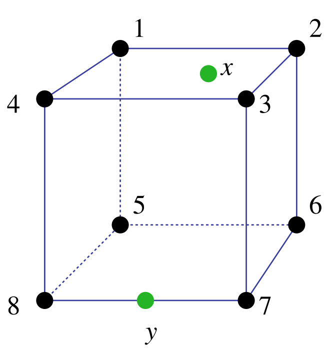

Figure 2.3: This polytope has 27 faces: 1 of dimension three, 6 of dimension 2 (facets), 12 of dimension 1 (edges), and 8 of dimension zero (vertices). Oh, and the empty face, which is usually said to have dimension -1. The carrier of $\mathbf{x}$ is the square $\mathrm{conv}(1234)$. The carrier of $\mathbf{y}$ is the edge $\mathrm{conv}(78)$.

We write $\mathbf{F} \leq \mathbf{P}$ to mean that $\mathbf{F}$ is a face of $\mathbf{P}$ and $\mathbf{F} < \mathbf{P}$ to mean that $\mathbf{F}$ is a proper face. If $\mathbf{F}$ is a non-empty proper face and ψ is a functional defining it, the hyperplane $\{\mathbf{y} \in \mathbb{R}^m : \psi(\mathbf{x}) = \psi(\mathbf{y}), \forall \mathbf{x} \in \mathbf{F}\}$ is called a *supporting hyperplane*. Equivalently, a supporting hyperplane is a hyperplane whose intersection with $\mathbf{P}$ is a non-empty face.

It is a well-known and fundamental theorem (see Chapters one and two of [339]) that a polytope can be described as either the convex hull of its set of vertices, or as the intersection of the half-spaces given by its facets. Here, a *half-space* is a set of the form $\{\mathbf{x} \in \mathbb{R}^m : \psi(\mathbf{x}) \geq c\}$, where ψ is a linear functional and $c \in \mathbb{R}$ is a constant. Any finite intersection of half-spaces is called a *polyhedron*. By Weyl-Minkowski's theorem *polytopes* are the same as bounded polyhedra.

The boundary of $\mathbf{P}$ is the union of its proper faces (equivalently, the union of its facets). The rest is the *relative interior* of $\mathbf{P}$. Every polytope is the disjoint union of the relative interiors of all its faces. Observe that the relative interior of a vertex is the vertex itself. The *carrier* in $\mathbf{P}$ of a point $\mathbf{x} \in \mathbf{P}$ or of a subset $\mathbf{X} \subseteq \mathbf{P}$ is the minimal face of $\mathbf{P}$ containing $\mathbf{x}$ or $\mathbf{X}$, respectively. Equivalently, the carrier of a point is the unique face having $\mathbf{x}$ in its relative interior.

A set of points is *affinely independent* (or *independent* for short) if none of them is an affine combination of the rest. It is called *dependent* otherwise. Equivalently, k points are independent if their convex hull has dimension $k-1$. An affinely independent set is also called a *basis* of its affine span. The convex hull of an affinely independent set of $k+1$ points is a *k-simplex*. Equivalently, a *k-simplex* is any polytope of dimension k with $k+1$ vertices. Every face of a k-simplex is a simplex and it has 2^{k+1}

possible faces. They are spanned by each of the possible subsets of points (the span of the empty set is the empty face). All k-simplices are affinely equivalent to one another. That is, if $\mathbf{P}$ and $\mathbf{Q}$ are simplices of the same dimension then there is an affine-linear bijection between $\operatorname{aff}(\mathbf{P})$ and $\operatorname{aff}(\mathbf{Q})$ sending $\mathbf{P}$ to $\mathbf{Q}$.

So far we have been speaking about "affine" objects, but we will often need to consider the analogous objects in linear algebra. We now consider the elements of $\mathbb{R}^m$ as "vectors" rather than "points".

The *positive hull*, or *positive span*, of a finite set $\mathbf{V}$ of vectors in $\mathbb{R}^m$ is the set of vectors that can be obtained as non-negative linear combinations of our points.

$$\operatorname{cone}(\mathbf{V}) := \left\{ \sum_{\mathbf{v}\in\mathbf{V}} \lambda_{\mathbf{v}}\mathbf{v} \,:\, \lambda_{\mathbf{v}\geq 0} \quad \forall \mathbf{v}\in\mathbf{V} \right\}.$$

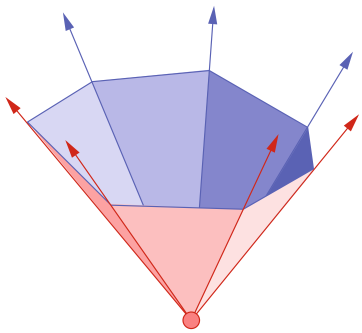

Figure 2.4: The positive hull of a finite set of vectors.

The sets of the form $\operatorname{cone}(\mathbf{V})$ for a finite set of vectors $\mathbf{V}$ is a *convex polyhedral cone*. Again, since we never deal with non-convex non-polyhedral cones, we simply call them *cones*. The *dimension* of a cone is the dimension of its *linear hull*, i.e., the linear subspace spanned by it. The *lineality space* of a cone is the largest linear subspace contained in it. A polyhedral cone is *pointed* if its lineality space is the zero subspace, or equivalently, if it does not contain any line.

As with polytopes, a *face* of a cone $\mathbf{P}$ is the subset where a linear functional is maximized or minimized. The difference now is that when this happens the minimum, or maximum, must be zero. A *supporting hyperplane* for a face $\mathbf{F}$ is a hyperplane that intersects $\mathbf{P}$ exactly on $\mathbf{F}$, and the *relative interior* of a cone is the cone minus the union of its proper faces. We now collect some basic facts about cones. The first one is essentially Theorem 1.3 in [339].

Proposition 2.1.2. *A subset $\mathbf{C}$ of $\mathbb{R}^m$ is a polyhedral cone if and only if it is the intersection of finitely many linear halfspaces, i.e., there exists a finite index set I and $\psi_i \in (\mathbb{R}^m)^*$ for $i \in I$ such that*

$$\mathbf{C} = \left\{ \mathbf{x}\in\mathbb{R}^m \,:\, \psi_i(\mathbf{x}) \geq 0 \textit{ for all } i \in I \right\}. \quad \square$$

Corollary 2.1.3. *Every linear hyperplane*

$$\left\{ \mathbf{x}\in\mathbb{R}^m \,:\, \psi(\mathbf{x}) = 0 \right\} = \left\{ \mathbf{x}\in\mathbb{R}^m \,:\, \psi(\mathbf{x}) \geq 0,\ -\psi(\mathbf{x}) \geq 0 \right\}$$

is a polyhedral cone. $\square$

Corollary 2.1.4. *The intersection of polyhedral cones is again a polyhedral cone.* $\square$

Recall that a (convex) polyhedron is any finite intersection of half-spaces. In particular, both polytopes and cones are polyhedra. Now we define polyhedral complexes:

Definition 2.1.5 (Polyhedral Complex)**.** A set $\mathscr{K}$ of polyhedra is a *polyhedral complex* if

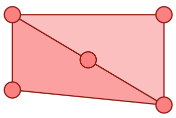

Figure 2.5: Is this a polyhedral complex?

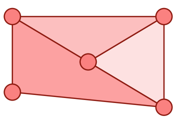

Figure 2.6: Not a polyhedral complex.

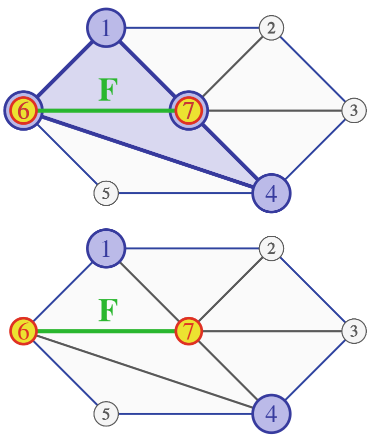

Figure 2.7: The star of the simplex $\mathbf{F} = 67$ consists of the two triangles 167 and 467; the link consists of the two vertices 1 and 4.

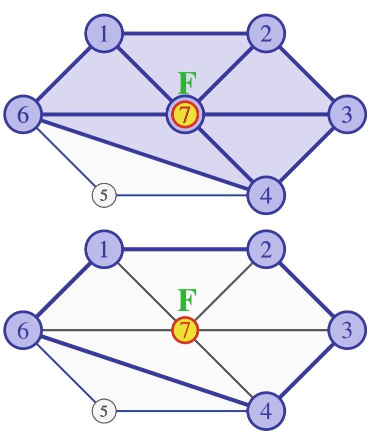

Figure 2.8: The star of the simplex $\mathbf{F} = 7$ consists of the five triangles containing it; the link consists of their edges not containing 7.

(i) $\mathbf{P} \in \mathscr{K}$ and $\mathbf{F} \leq \mathbf{P}$ implies that $\mathbf{F} \in \mathscr{K}$.

(ii) $\mathbf{P} \cap \mathbf{Q} \leq \mathbf{P}$ and $\mathbf{P} \cap \mathbf{Q} \leq \mathbf{Q}$ for all $\mathbf{P}, \mathbf{Q} \in \mathscr{K}$.

The individual polyhedra in $\mathscr{K}$ are called *cells* of the polyhedral complex. They are also sometimes called *faces* of the subdivision.

As an example, the set of all faces of a polytope, or of a polyhedron, is a polyhedral complex. The set of proper faces is another example called the *boundary complex*. Polyhedral subdivisions and triangulations of point sets, the central topic of this book, are also polyhedral complexes.

The dimension of a polyhedral complex is the highest among the dimensions of its cells. All cells of a polyhedral complex can be ordered by containment. A *maximal cell* is one that has no other cell containing it. A polyhedral complex is *pure* if all its maximal cells are of the same dimension. A *(geometric) simplicial complex* is one whose cells are all simplices. We will see more about (geometric and abstract) simplicial complexes in Section 2.6.1, but here are some important definitions.

Definition 2.1.6.

- For any $\mathbf{F} \in \mathscr{K}$, the *star* of $\mathbf{F}$ in $\mathscr{K}$, $\mathrm{st}_{\mathscr{K}}(\mathbf{F})$, is the subcomplex of $\mathscr{K}$ made up of all polyhedra of $\mathscr{K}$ having $\mathbf{F}$ as a face plus all their faces.

- The *link* of $\mathbf{F}$ in $\mathscr{K}$ is the polyhedral complex $\mathrm{link}_{\mathscr{K}}(\mathbf{F}) = \{\mathbf{C} \in \mathrm{st}_{\mathscr{K}}(\mathbf{F}) | \mathbf{F} \cap \mathbf{C} = \emptyset\}$. Note that the link of a $(k-1)$-dimensional face $\mathbf{F}$, $\mathrm{link}(\mathbf{F}, \mathscr{K})$, is a complex of dimension $d-k-1$.

- The *anti-star* of $\mathbf{F}$ in $\mathscr{K}$ is the polyhedral complex $\mathrm{ast}_{\mathscr{K}}(\mathbf{F}) = \{\mathbf{C} \in \mathscr{K} | \mathbf{F} \cap \mathbf{C} = \emptyset\}$. Alternately:
$$\mathrm{ast}_{\mathscr{K}}(\mathbf{F}) := (\mathscr{K} \setminus \mathrm{st}_{\mathscr{K}}(\mathbf{F})) \cup \mathrm{link}_{\mathscr{K}}(\mathbf{F})$$

- The *boundary* $\partial\mathscr{K}$ of a pure d-dimensional polyhedral complex $\mathscr{K}$ is the polyhedral complex whose maximal faces consist of the $d-1$ faces contained in only one d-face. For example, the boundary of a simplicial ball is a simplicial sphere, and the boundary of a simplicial sphere is empty.

- We say that a polyhedral complex $\mathscr{L}$ is a *subcomplex* of $\mathscr{K}$ if $\mathscr{L} \subset \mathscr{K}$. The boundary of $\mathscr{K}$ is an example of a subcomplex of $\mathscr{K}$.

Polyhedral complexes all of whose faces are cones are called fans:

Definition 2.1.7 (Fan). A *polyhedral fan* (or *fan* for short) in $\mathbb{R}^m$ is a polyhedral complex consisting of polyhedral cones. A fan is *pointed* if all of its cones are pointed. A fan is *complete* if the union of all its cones is $\mathbb{R}^m$. $\square$

Observe that the face consisting of the origin must be the same in all the cones of a fan, and it coincides with the lineality space of them. So if a single cone in a fan is pointed, then all of the cones are.

Definition 2.1.8 (Normal Fan)**.** Let $\mathbf{P}$ be a polyhedron in $\mathbb{R}^m$. For a point $\mathbf{x} \in \mathbf{P}$ we define the *outer normal cone of* $\mathbf{p}$ in $\mathbf{P}$ in $\mathbb{R}^m$ as

$$N_{\mathbf{P}}(\mathbf{x}) := \{\ \psi \in (\mathbb{R}^m)^* : \langle \psi, \mathbf{x} \rangle \geq \langle \psi, \mathbf{y} \rangle\ \forall \mathbf{y} \in \mathbf{P}\ \} .$$

Put differently, $N_{\mathbf{P}}(\mathbf{x})$ consists of all the linear functionals whose maximum on $\mathbf{P}$ is achieved at $\mathbf{x}$.

Similarly, for a face of $\mathbf{P}$, $\mathbf{F} \leq \mathbf{P}$ we say $N_{\mathbf{P}}(\mathbf{F})$ equals the cone $N_{\mathbf{P}}(\mathbf{x})$ for any $\mathbf{x}$ in the relative interior of $\mathbf{F}$ (this is clearly independent of the point chosen). $N_{\mathbf{P}}(\mathbf{F})$ is a polyhedral cone called the *outer normal cone* of $\mathbf{F}$ in $\mathbf{P}$ in $\mathbb{R}^m$.

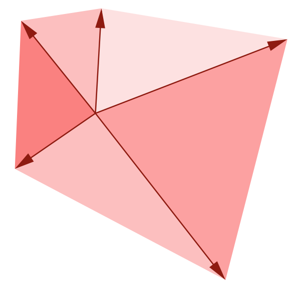

Figure 2.9: A complete pointed fan in dimension two.

The corresponding *inner normal cones* are defined as the negatives of the outer normal cones. The set

$$\mathcal{N}_{\mathbf{P}} := \{\ N_{\mathbf{P}}(\mathbf{F}) : \mathbf{F} \leq \mathbf{P}\ \} = \{\ N_{\mathbf{P}}(\mathbf{p}) : \mathbf{p} \in \mathbf{P}\ \}$$

is the *outer normal fan* of $\mathbf{P}$ in $\mathbb{R}^m$.

Two polytopes $\mathbf{P}$ and $\mathbf{P}'$ are *normally equivalent* if their outer normal fans are the same:

$$\mathbf{P} \sim \mathbf{P}' \iff \mathcal{N}_{\mathbf{P}} = \mathcal{N}_{\mathbf{P}'} . \quad \square$$

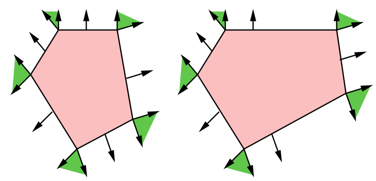

Figure 2.10: Two normally equivalent 5-gons.

Proposition 2.1.9. *Let* $\mathbf{P}$ *be a polytope in* $\mathbb{R}^m$*, not necessarily full-dimensional. Then* $N_{\mathbf{P}}(\mathbf{x})$ *is full-dimensional (i.e., of dimension m), if and only if* $\mathbf{x}$ *is a vertex of* $\mathbf{P}$. $\square$

2.1.2 Point configurations

By a *point configuration* we mean a finite set of labeled points in real affine space $\mathbb{R}^m$, but we allow our set to have repeated points which receive different labels. To see why this may be useful, suppose for a moment that you project a 3-dimensional cube, as shown in Figure 2.11, in the direction of the diagonal line joining antipodal vertices a, b. Those two points are projected on top of each other. If you want to recall the top and bottom views of the cube (with respect to this direction) you get two similar but different two-dimensional pictures (see bottom of Figure 2.11). The issue is that the interior points used in the perspectives are different. It is then a good idea to remember that the plane of projection has, so to speak, two copies of the same point.

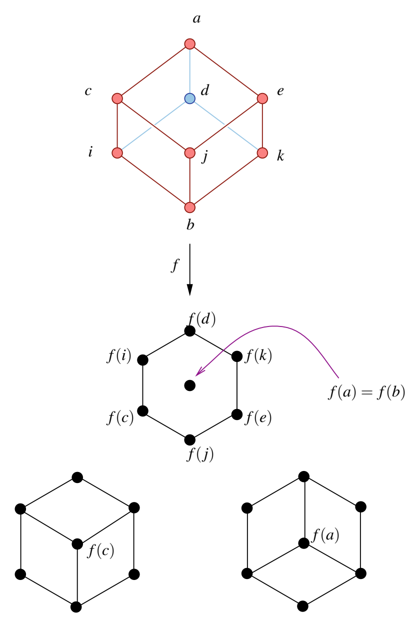

Figure 2.11: A situation when repeated points occur.

There are other fundamental operations that can be performed on point configurations which make repeated points natural and interesting. Among them are *Gale transforms* (see Section 4.1.3), *the contraction at a point* (see Section 4.2), and the *Minkowski sum* (see Section 9.2).

The best way to deal with repeated points is via labels. Every element of a point configuration will have a label, and all labels are assumed to be different. A repeated point will have several labels attached to it. Typically, but not necessarily, labels will be the first positive integer numbers. That is:

Definition 2.1.10. A *point configuration* in $\mathbb{R}^m$ is a finite set of (perhaps repeated) points with (non-repeated) labels.

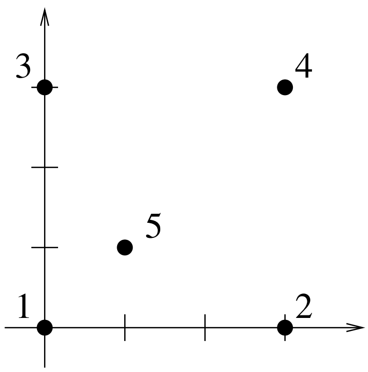

Figure 2.12: Five points in the plane.

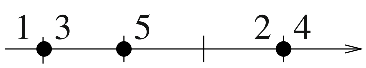

Figure 2.13: Five points along a line.

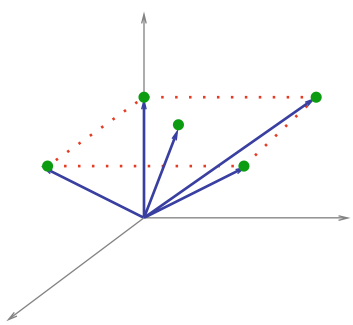

Figure 2.14: Homogenization of a five point planar configuration.

More formally, a point configuration in $\mathbb{R}^m$ with set of labels J is a map $J \to \mathbb{R}^m$.

We will typically refer to elements of a configuration by their labels, not by their coordinates, saying, for example, that in Figure 2.12 the points 1, 5 and 4 are collinear. This takes care of ambiguities when we have repeated points, which are distinguished by their labels.

As a second advantage, we can carry labels from one configuration to another one, if the latter is obtained from the first by a geometric construction. For example, consider the projection of Figure 2.12 to Figure 2.13). This gives a different configuration $\mathbf{A}'$, but we can keep the same labels. In this way we can say things like: "2, 4, and 5 are independent in $\mathbf{A}$ but dependent in $\mathbf{A}'$".

It is convenient to represent a configuration as the columns of a matrix. For example, the point configuration of Figure 2.12, consisting of the five points $(0,0)$, $(3,0)$, $(0,3)$, $(3,3)$, and $(1,1)$, would be represented as:

$$\begin{array}{c} \begin{array}{ccccc} \mathit{1} & \mathit{2} & \mathit{3} & \mathit{4} & \mathit{5} \end{array} \\ \begin{pmatrix} 0 & 3 & 0 & 3 & 1 \\ 0 & 0 & 3 & 3 & 1 \end{pmatrix} \end{array}$$

The row above the matrix shows the labels we attach to the points. The matrix representation, among other things, makes repeated points familiar. They are just columns that happen to be equal.

But it is even more convenient to represent a point configuration in $\mathbb{R}^d$ as a $(d+1) \times n$ matrix by adding a constant row. That is, the previous point configuration could be written instead as:

$$\mathbf{A} = \begin{array}{c} \begin{array}{ccccc} \mathit{1} & \mathit{2} & \mathit{3} & \mathit{4} & \mathit{5} \end{array} \\ \begin{pmatrix} 0 & 3 & 0 & 3 & 1 \\ 0 & 0 & 3 & 3 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \end{array}$$

The addition of this extra coordinate is called *homogenization*. It is very helpful to use homogeneous coordinates in affine spaces, since it turns affine geometry into a special case of linear algebra. For example:

1. An affine dependence between the points $(0,0)$, $(3,0)$, $(0,3)$, $(3,3)$, and $(1,1)$ is any non-zero vector $\lambda = (\lambda_1, \ldots, \lambda_5)$ with $\sum \lambda_i \mathbf{p}_i = 0$ and $\sum \lambda_i = 0$. Equivalently, it is a vector λ such that $\mathbf{A} \cdot \lambda = 0$, if $\mathbf{A}$ is the homogeneous matrix of the point configuration.

2. Similarly, the fact that three points in $\mathbf{A}$ are collinear, or the value area of the triangle they span in case they are not, are seen in the matrix as the vanishing, or the value, of the corresponding 3×3 determinant in the homogeneous matrix.

In Section 2.5 we will introduce triangulations of *vector*, rather than *point*, configurations. The use of homogeneous coordinates for point configurations will make this task seamless.

To represent a point configuration $\mathbf{A}$ in homogeneous coordinates we do not really need to have a constant row. Homogeneous coordinates are based on the fact that the affine space $\mathbb{R}^m$ can be naturally identified with *any* hyperplane of $\mathbb{R}^{m+1}$ not passing through the origin. Hence, any matrix whose columns lie in such a hyperplane will do the job. We call such matrices *homogeneous*. Put differently, an $m \times n$ matrix $\mathbf{A}$ is homogeneous if there is a row vector $\phi \in \mathbb{R}^m$ such that $\phi\mathbf{A}$ is a vector with all entries equal.

This freedom sometimes permits us to produce simpler coordinates, or coordinates that highlight symmetries. For example, the point configuration of Figure 2.15, consisting of the vertices of two concentric regular triangles, can be represented by the following matrix, whose columns all lie in the hyperplane $x_1 + x_2 + x_3 = 4$:

$$\begin{array}{c} \begin{array}{cccccc} \mathit{1} & \mathit{2} & \mathit{3} & \mathit{4} & \mathit{5} & \mathit{6} \end{array} \\ \begin{pmatrix} 4 & 0 & 0 & 2 & 1 & 1 \\ 0 & 4 & 0 & 1 & 2 & 1 \\ 0 & 0 & 4 & 1 & 1 & 2 \end{pmatrix} \end{array}$$

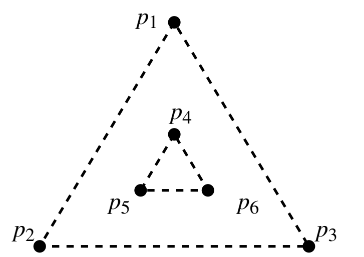

Figure 2.15: An interesting set of six points in the plane.

In this example, which we saw already in Section 1.2, each column gives the barycentric coordinates of the corresponding point with respect to the outer triangle, except that the barycentric coordinates have been normalized to add up to four, instead of one, for the convenience of getting integer numbers. Note that the coordinates can be changed via a rotation to have the points lying in the plane with coordinate $x_3 = 1$ instead of the plane $x_1 + x_2 + x_3 = 4$.

Summarizing what we have discussed:

> A d-dimensional point configuration with n elements is represented as an $m \times n$ homogeneous matrix of rank $d+1$. Usually, but not always, $m = d+1$. Columns of the matrix are the elements of the configuration and will usually be identified by their labels.

Definition 2.1.11. The rank $d+1$ of the matrix defining a point configuration is called the *rank* of the point configuration. The number $n-d+1$, where n is the number of elements, is called the *corank* of the configuration. The *dimension* of a point configuration is the dimension of its convex hull. The dimension plus one is equal to the rank.

At this point let us make explicit some typesetting and notational conventions that we have been using so far and will keep using throughout the book:

- Subsets of $\mathbb{R}^m$ (that is, polytopes, hyperplanes, cones, etc.) are denoted by upright boldface letters. This includes point configurations, as well as subconfigurations (submatrices made out from a subset of columns) $(\mathbf{A}, \mathbf{B}, \ldots)$, even if strictly speaking they are labeled subsets.

- Sets of labels are denoted by italic letters (J, I, Z). Our preference will be to use $\mathbf{A}$ for the configuration and J for the label set.

- If C is a subset of the label set, then $\mathbf{A}|_C$ denotes the corresponding subconfiguration (the submatrix consisting of the columns labeled by C).

- Similarly, boldface lowercase letters $(\mathbf{p}, \mathbf{q}, \mathbf{z})$ will denote points or vectors in $\mathbb{R}^m$ (e.g., the individual columns in a configuration) while italic lowercase letters represent individual labels (i, j, z).

- We will sometimes write matrices in abbreviated form. That is, $\mathbf{A} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ if $\{1, \dots, n\}$ is the set of labels or $\mathbf{A} = (\mathbf{p}_j)_{j \in J}$ for a general label set J.

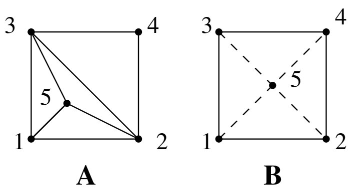

Figure 2.16: $\{125, 135, 325, 234\}$ is a triangulation of $\mathbf{A}$, but not of $\mathbf{B}$.

2.1.3 Geometry of point configurations

Let us look more closely at the consequences of our decision of referring to points by their labels. The left part of Figure 2.16 displays a triangulation of the point configuration $\mathbf{A}$ of the previous section. Triangles can be then referred to as triplets of labels, such as $\{1,2,5\}$, $\{1,3,5\}$, etc. To simplify notation, we sometimes abbreviate them as 125, 135, etc. As with point configurations, one advantage of this is that then we can also say things like "$\mathscr{T}$ is a triangulation of $\mathbf{A}$, but not of $\mathbf{B}$", where $\mathbf{B}$ is the configuration on the right of the same figure, obtained from $\mathbf{A}$ by a perturbation.

If we accept this, then sentences like "$\{1,2\}$ is a facet of $\{1,2,5\}$" have to be allowed in our language. In this setup we can speak as if the sets of labels themselves were geometric objects, which have a convex hull, faces, etc. Thus when dealing with point configurations the convex set structure takes second place to the combinatorics.

Let us repeat most of the definitions of 2.1.1 in this new "labeled" setting.

Definition 2.1.12 (Convex hull, relative interior)**.** Let $\mathbf{A} = (\mathbf{p}_j)_{j \in J}$ be a point configuration in $\mathbb{R}^m$, with set of labels J. For a subset C of J we define the *convex hull of* C *in* $\mathbf{A}$ to be the following closed convex set:

$$\operatorname{conv}_{\mathbf{A}}(C) := \left\{ \sum_{j \in C} \lambda_j \mathbf{p}_j : \lambda_j \geq 0 \text{ for all } j \in C, \text{ and } \sum_{j \in C} \lambda_j = 1 \right\}. \quad (2.1)$$

The *dimension* and *relative interior* of C are defined to be the dimension and relative interior of its convex hull. We recall that the latter is the following relatively open (i.e., open in its affine hull) convex set:

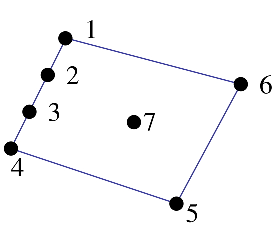

Figure 2.17: The relative interior of $\mathbf{A}$ is the rectangle without its border.

$$\operatorname{relint}_{\mathbf{A}}(C) := \left\{ \sum_{j \in C} \lambda_j \mathbf{p}_j : \lambda_j > 0 \text{ for all } j \in C, \text{ and } \sum_{j \in C} \lambda_j = 1 \right\}. \quad (2.2)$$

Remark 2.1.13. Every (nonempty) point configuration has a non-empty relative interior. For example, if C has only one element, that point is its relative interior (more precisely, $\operatorname{relint}_{\mathbf{A}}(\{b\}) = \operatorname{conv}_{\mathbf{A}}(\{b\}) = \{\mathbf{p}_b\}$).

Remark 2.1.14. Note that if $\mathbf{A}$ is a point configuration with label set J and $R \subset C \subset J$, then $\operatorname{conv}_{\mathbf{A}|_C}(C) = \operatorname{conv}_{\mathbf{A}}(C)$. Thus, whenever the configuration to which the labels refer is clear, we will simply write $\operatorname{conv}(R)$ and $\operatorname{relint}(R)$. This applies to most of the notation introduced in this section.

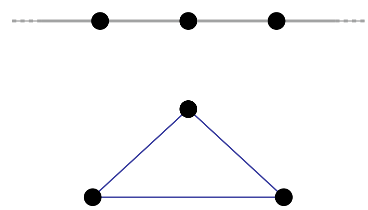

Figure 2.18: Three dependent points (top) and three independent points (bottom).

Definition 2.1.15 (dependent, spanning, general and convex position)**.** Let $\mathbf{A}$ be a point configuration of dimension d with label set J. Let C be a subset of the label set J.

(i) We say that C is *dependent in* $\mathbf{A}$ if there is a non-zero vector $(\lambda_j)_{j\in C}$ such that

$$\sum_{j\in C} \lambda_j \mathbf{p}_j = 0 \text{ and } \sum_{j\in C} \lambda_j = 0.$$

It is called *independent* otherwise.

(ii) We say that C is *spanning* if $\operatorname{conv}(C)$ has the same dimension as $\operatorname{conv}(J)$.

(iii) C is *in general position* if each subset with at most $d+1$ elements is independent. Equivalently, if every dependent set is spanning in $\mathbf{A}|_C$. If a set is not in general position we say it is *in special position*.

(iv) An element $j \in J$ is *extremal* in $\mathbf{A}$ if the corresponding point $\mathbf{p}_j$ is not repeated in $\mathbf{A}$ and is a vertex of $\operatorname{conv}_{\mathbf{A}}(J)$. A point configuration is *in convex position* if all its elements are extremal.

Remark 2.1.16. Observe that if C has a repeated point (that is, two labels i and j pointing to the same point $\mathbf{p}_i = \mathbf{p}_j$), then C is necessarily dependent, since $\mathbf{p}_i - \mathbf{p}_j = 0$ is a dependence.

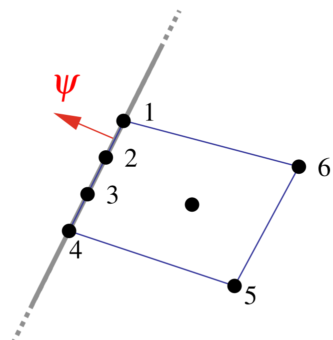

Figure 2.19: Face of C in the direction ψ is 1234.

Definition 2.1.17 (Face)**.** Let $\mathbf{A} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ be a point configuration in $\mathbb{R}^m$, with set of labels J. Let $C \subset J$. For a linear functional $\psi \in (\mathbb{R}^m)^*$, the *face of C in direction ψ* is the following subset of C:

$$\operatorname{face}_{\mathbf{A}}(C,\psi) := \left\{ j \in C \,:\, \psi(\mathbf{p}_j) = \max_{b\in C}(\psi(\mathbf{p}_b)) \right\}. \tag{2.3}$$

The affine hyperplane $\{\, \mathbf{x} \in \mathbb{R}^m \,:\, \psi(\mathbf{x}) = \psi(\mathbf{p}_j), j \in F \,\} \subset \mathbb{R}^m$ is called a *supporting hyperplane* of the face F. The empty subset is considered a face and C is always a face of C, obtained with the zero functional. If F is a face of C , we write $F \leq_{\mathbf{A}} C$. If, moreover, $F \neq C$ then we write $F <_{\mathbf{A}} C$, and we say that F is a *proper face* of C. A *facet* of C is a face of dimension one less than the dimension of C, that is, it is a maximal proper face.

Remark 2.1.18. A face of a face of C is also a face of C, as is any intersection of faces.

Remark 2.1.19. $\operatorname{relint}_{\mathbf{A}}(C) = \operatorname{conv}_{\mathbf{A}}(C) \setminus \bigcup_{F<C} \operatorname{conv}_{\mathbf{A}}(F)$. Thus, $\operatorname{conv}_{\mathbf{A}}(C)$ equals $\dot{\bigcup}_{F\leq C} \operatorname{relint}_{\mathbf{A}}(F)$, where $\dot{\bigcup}$ denotes "disjoint union".

Remark 2.1.20. We have the following connections to the corresponding concepts of convex geometry:

(i) $\mathrm{conv}_{\mathbf{A}}(\mathrm{face}_{\mathbf{A}}(C,\psi))$ equals the face of the polytope $\mathrm{conv}_{\mathbf{A}}(C)$ in the direction of ψ.

(ii) If $F <_{\mathbf{A}} C$, then $\mathrm{conv}_{\mathbf{A}}(F) < \mathrm{conv}_{\mathbf{A}}(C)$, but in general the reverse implication is not true.

(iii) If B is in general position, then $\mathrm{conv}_{\mathbf{A}}(F) < \mathrm{conv}_{\mathbf{A}}(C)$ implies $F <_{\mathbf{A}} C$ too. This is because all proper faces are spanned by independent subconfigurations.

(iv) $\bigcup_{F<C} \mathrm{conv}_{\mathbf{A}}(F) = \partial(\mathrm{conv}_{\mathbf{A}}(C))$.

Remark 2.1.21. A single element $\{j\}$ is a face of C if and only if $\mathbf{p}_j$ is not repeated and is a vertex of $\mathrm{conv}_{\mathbf{A}}(C)$ in the usual sense. This is exactly what we defined as an *extremal point* of C.

Definition 2.1.22 (Carrier)**.** Let $\mathbf{A}$ be a configuration in $\mathbb{R}^m$. Let C be a set of labels and let $\mathbf{X} \subseteq \mathrm{conv}_{\mathbf{A}}(C)$ and $R \subset C$. The *carrier* of $\mathbf{X}$ in C is the smallest face of C whose convex hull contains $\mathbf{X}$. That is,

$$\mathrm{carrier}_{\mathbf{A}}(\mathbf{X},C) := \bigcap_{\substack{\mathbf{X} \subseteq \mathrm{conv}_{\mathbf{A}}(F) \\ F \leq C}} F. \tag{2.4}$$

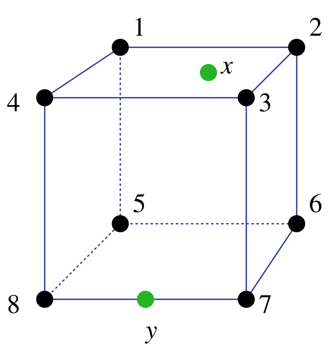

Figure 2.20: The carrier of x is $\{1,2,3,4\}$. The carrier of y is $\{7,8\}$.

The carrier of R in C is $\mathrm{carrier}_{\mathbf{A}}(R,C) := \bigcap_{R \subseteq F \leq C} F$. Clearly, $\mathrm{carrier}_{\mathbf{A}}(R,C)$ equals $\mathrm{carrier}_{\mathbf{A}|_C}(R,C)$.

Sometimes the easiest way to prove that something is a face is to check that it coincides with its own carrier. For this trick, which will be applied several times in Section 4.5.1, the following lemma is useful.

Lemma 2.1.23. *For every pair of subsets R,C of labels of a point configuration $\mathbf{A}$ that satisfy $R \subseteq C$ we have*

$$\mathrm{relint}_{\mathbf{A}}(R) \subseteq \mathrm{relint}_{\mathbf{A}}(\mathrm{carrier}_{\mathbf{A}}(R,C)). \tag{2.5}$$

Remark 2.1.24. In fact, something more general is true: since different faces have disjoint relative interiors, the carrier of C is the unique face F with $\mathrm{relint}_{\mathbf{A}}(C) \subseteq \mathrm{relint}_{\mathbf{A}}(F)$. In particular, as in the geometric case, the carrier of a point $\mathbf{x}$ in C is the unique (by Remark 2.1.19) face F of C with $\mathbf{x} \in \mathrm{relint}_{\mathbf{A}}(F)$.

Proof. We have to show that $\mathrm{relint}_{\mathbf{A}}(R) \cap \partial\, \mathrm{conv}_{\mathbf{A}}(\mathrm{carrier}_{\mathbf{A}}(R,C)) = \emptyset$.

Assume, for the sake of contradiction, that there exists a point $\mathbf{x}$ in $\mathrm{relint}_{\mathbf{A}}(R) \cap \partial\, \mathrm{conv}_{\mathbf{A}}(\mathrm{carrier}_{\mathbf{A}}(R,C))$. Since $\mathbf{x}$ is not in the boundary of $\mathrm{conv}_{\mathbf{A}}(R)$, each hyperplane $\mathbf{H} \subset \mathbb{R}^m$ through $\mathbf{x}$ must either separate the points in $\mathrm{conv}_{\mathbf{A}}(R)$ (i.e., there are points in $\mathrm{conv}_{\mathbf{A}}(R)$ on both sides of $\mathbf{H}$) or $\mathbf{H}$ must contain all of $\mathrm{conv}(R)$. In particular, it must either separate the points in $\mathbf{A}|_R$ or contain all points in $\mathbf{A}_R$. Since $R \subseteq \mathrm{carrier}_{\mathbf{A}}(R,C)$ by definition of the carrier, each hyperplane through $\mathbf{x}$ must either separate $\mathrm{carrier}_{\mathbf{A}}(R,C)$, or contain all points of $\mathbf{A}|_R$. Since $\mathbf{x}$ is in the boundary of $\operatorname{carrier}_{\mathbf{A}}(R,C)$, there is a hyperplane $\mathbf{H}$ through $\mathbf{x}$ neither separating $\operatorname{carrier}_{\mathbf{A}}(R,C)$ nor containing all of $\operatorname{carrier}_{\mathbf{A}}(R,C)$. Therefore, $\mathbf{H}$ must contain $\mathbf{A}|_R$ completely. This, however, implies that $\mathbf{A}|_R$ is contained in $\operatorname{carrier}_{\mathbf{A}}(R,C) \cap \mathbf{H} \neq \operatorname{carrier}_{\mathbf{A}}(R,C)$, a contradiction with the minimality of the carrier. □

Lemma 2.1.25. *(i) For subsets $R \subseteq F \leq C$ of labels of a point configuration $\mathbf{A}$, we have*

$$\operatorname{carrier}_{\mathbf{A}}(R,F) = \operatorname{carrier}_{\mathbf{A}}(R,C). \tag{2.6}$$

(ii) For subsets $F \leq C$ of labels of a point configuration $\mathbf{A}$ and a point $\mathbf{x} \in \operatorname{conv}_{\mathbf{A}}(F)$, we have

$$\operatorname{carrier}_{\mathbf{A}}(\mathbf{x},F) = \operatorname{carrier}_{\mathbf{A}}(\mathbf{x},C). \tag{2.7}$$

(iii) For subsets $F \leq C$ of labels of a point configuration $\mathbf{A}$ and a subset $\mathbf{X} \subseteq \operatorname{conv}_{\mathbf{A}}(F)$ we have

$$\operatorname{carrier}_{\mathbf{A}}(\mathbf{X},F) = \operatorname{carrier}_{\mathbf{A}}(\mathbf{X},C). \tag{2.8}$$

Proof. We only prove the first assertion; the remaining items are analogous and left to the reader. Let $R \subseteq F \leq C$ be subsets of labels of a point configuration $\mathbf{A}$. Then:

$$\begin{aligned}
\operatorname{carrier}_{\mathbf{A}}(R,F) &= \bigcap_{\substack{G \leq F \\ R \subseteq G}} G \\
&= F \cap \bigcap_{\substack{G \leq C \\ R \subseteq G}} G \\
&= \bigcap_{\substack{G \leq C \\ R \subseteq G}} G \\
&= \operatorname{carrier}_{\mathbf{A}}(R,C).
\end{aligned}$$

The second to last equality is true because $F \leq C$. □

In words, this means that the carrier does not change if you (properly) enlarge the object in which the carrier is taken.

2.2 A closer look at the definition of triangulation

Here will demonstrate that the language developed in Section 2.1.3 is now adequate for our purposes. Nevertheless, let us first work with the probably more familiar language of Section 2.1.1 and pinpoint some of its limitations. We will discover that things actually become simpler by the use of a seemingly more abstract language.

We take as starting point the definition of triangulations given at the beginning of the book.

Definition 2.2.1. A *triangulation* of a point configuration $\mathbf{A}$ is a collection $\mathscr{T}$ of simplices, with vertices in $\mathbf{A}$, that satisfies the following properties:

1. All faces of simplices of $\mathscr{T}$ are in $\mathscr{T}$. (*Closure Property*)
2. The intersection of any two simplices of $\mathscr{T}$ is a (possibly empty) face of both. (*Intersection Property*.)
3. The union of all these simplices equals $\operatorname{conv}(\mathbf{A})$. (*Union Property*)

Note that the first two properties are the definition of a (geometric) *simplicial complex*. In other words: a triangulation of $\mathbf{A}$ is *a simplicial complex with vertex set contained in* $\mathbf{A}$ *and which covers* $\operatorname{conv}(\mathbf{A})$.

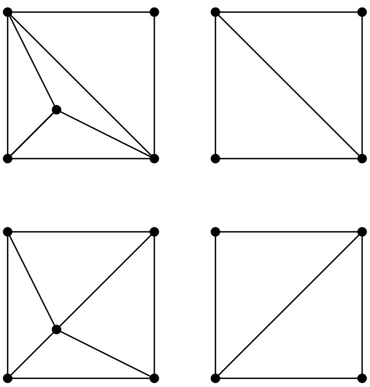

Figure 2.21: The four triangulations of the point configuration of Figure 2.12.

In our definition we do not assume $\operatorname{conv}(\mathbf{A})$ to be full-dimensional. In particular, we may speak of triangulations of a single point (there is one!), or of triangulations of a face of $\operatorname{conv}(\mathbf{A})$, as in the following statement, proved in Lemma 2.3.4 in a more general context.

Lemma 2.2.2. *Let* $\mathscr{T}$ *be a triangulation of a point configuration* $\mathbf{A}$ *and let* $\mathbf{F}$ *be a face of* $\operatorname{conv}(\mathbf{A})$. *Then, the following is a triangulation of* $\mathbf{A} \cap \mathbf{F}$*:*

$$\mathscr{T}_{\mathbf{F}} := \{ \sigma \in \mathscr{T} \, : \, \sigma \subset \mathbf{F} \}.$$

Observe that we do not require *all* the points of $\mathbf{A}$ to be used as vertices in a triangulation. For example, the configuration of Figure 2.12 has the four triangulations shown in Figure 2.21. Two of them use the five points and have four triangles, and two use only four points and have two triangles. Similarly, the six points in Figure 2.15 have 18 triangulations, only 8 of which use all points. Of course, all vertices of $\operatorname{conv}\mathbf{A}$ are used in all triangulations.

2.2.1 There is always a triangulation

Our first goal is to show that every point configuration has at least one triangulation. The method we are going to use is conceptually the simplest way to compute triangulations of point configurations. It is surprisingly general and it is central to the structure of the set of all triangulations of $\mathbf{A}$. The process, illustrated in Figure 2.22, is as follows:

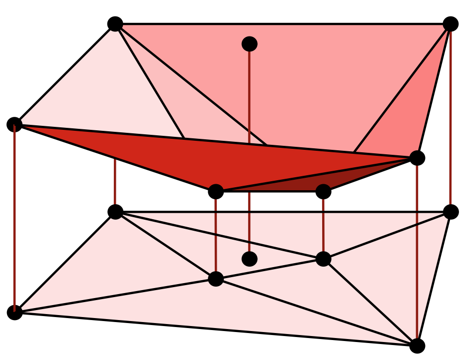

Figure 2.22: The lifting construction.

Let $\mathbf{A} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ be a point configuration in $\mathbb{R}^m$:

1. Pick a *height function* $\omega : \mathbf{A} \to \mathbb{R}$ (ω can be thought of as a vector $\omega = (\omega_1, \dots, \omega_n) \in \mathbb{R}^n$, with $\omega_i = \omega(\mathbf{p}_i)$), and consider the *lifted point configuration* in $\mathbb{R}^{n+1}$
$$\mathbf{A}^\omega := \begin{pmatrix} \mathbf{p}_1 & \dots & \mathbf{p}_n \\ \omega_1 & \dots & \omega_n \end{pmatrix}.$$

2. Compute the face structure of the polytope $\mathbf{P} = \mathrm{conv}(\mathbf{A}^\omega)$. A *lower face* of $\mathbf{P}$ is a face "visible from below", that is, a face that can be defined by minimizing a linear functional ψ which is positive on the last coordinate (equivalently, a face that has a non-vertical supporting hyperplane $\mathbf{H}$ and with $\mathbf{P}$ above $\mathbf{H}$).

3. Lower faces, since they are not vertical, project bijectively to polytopes inside $\mathrm{conv}(\mathbf{A})$. Moreover, as we will prove in a more general context in Lemma 2.3.11, the collection of projected lower faces always satisfies the three properties in the definition of triangulation. So, *if all lower faces of* $\mathbf{P}$ *are simplices*, their projections form a triangulation of $\mathbf{A}$.

Definition 2.2.3 (Regular Triangulation)**.** A triangulation of a point configuration $\mathbf{A}$ in $\mathbb{R}^m$ is called *regular* if it can be obtained by projecting the lower envelope of a lifting of $\mathbf{A}$ to $\mathbb{R}^{m+1}$.

Regular triangulations have appeared in different mathematical contexts and have actually received different names, such as *convex*, *weighted Delaunay*, *Gale*, or *coherent* triangulations.

Proposition 2.2.4. *Every point configuration has regular triangulations.*

While reading the proof, observe that we profit from the use of homogeneous coordinates for our point configuration.

Proof. We need to check that for any given $\mathbf{A}$ there are height functions ω for which $\mathbf{A}^\omega$ has the property that "all lower faces of $\mathrm{conv}(\mathbf{A}^\omega)$ are simplices". A sufficient condition for this to happen can be stated as: Every affine basis $\mathbf{B}$ contained in $\mathbf{A}$ is lifted so that the unique hyperplane containing the lifted point set $\mathbf{B}^\omega$ contains no other point $\mathbf{p}^\omega$ of $\mathbf{A}^\omega$.

For a given basis $\mathbf{B} = \{\mathbf{p}_1, \dots, \mathbf{p}_{d+1}\} \subseteq \mathbf{A}$ and extra point $\mathbf{p} \in \mathbf{A} \setminus \mathbf{B}$ this condition is equivalent to the non-vanishing of the determinant
$$\begin{vmatrix} \mathbf{p}_1 & \dots & \mathbf{p}_{d+1} & \mathbf{p} \\ \omega(\mathbf{p}_1) & \dots & \omega(\mathbf{p}_{d+1}) & \omega(\mathbf{p}) \end{vmatrix}.$$
This determinant is a linear equation on the ω's with non-zero coefficient on (at least) $\omega(\mathbf{p})$, hence it is non-zero except on a certain hyperplane in the space $\mathbb{R}^n$ of possible heights. As a conclusion, almost any choice satisfies the condition. More technically, there is an open dense subset of choices of ω in $\mathbb{R}^n$ that satisfies the condition: any ω lying in the complement of a certain union of finitely many hyperplanes. $\square$

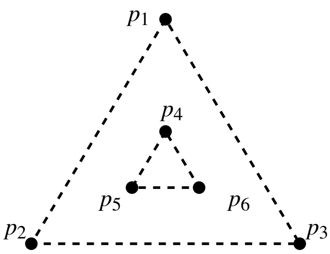

Figure 2.23: The mother of all examples and its six labeled points.

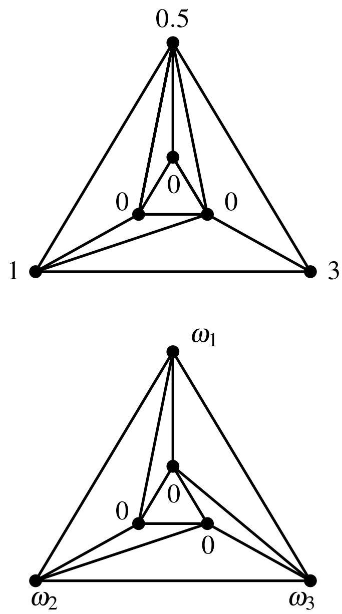

Figure 2.24: Two triangulations of the mother of all examples.

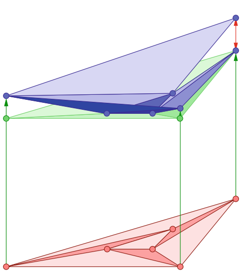

Figure 2.25: Fixing the height on the middle triangle and on the vertex in the background followed by tweaking the remaining heights (in order to "fold" along the missing diagonals) leads to a contradiction at the height of the vertex in the background.

The reader can verify without trouble that all the triangulations of (the vertex set of) a convex n-gon are regular. The same happens for the triangulations of our five-point example, displayed in Figure 2.21 (we will check this in detail in Section 2.2.3). But some point configurations have *non-regular triangulations*, which resemble Escher's famous "impossible pictures"; they look like projections of something from a higher dimension, but they really are not. Here is a *very* important example:

Example 2.2.5 (The mother of all examples)**.** Consider the two concentric triangles of Figure 2.23, whose coordinates we recall:

$$\begin{array}{c} \\ \mathbf{M} := \end{array} \begin{array}{c} \begin{matrix} \mathit{1} & \mathit{2} & \mathit{3} & \mathit{4} & \mathit{5} & \mathit{6} \end{matrix} \\ \begin{pmatrix} 4 & 0 & 0 & 2 & 1 & 1 \\ 0 & 4 & 0 & 1 & 2 & 1 \\ 0 & 0 & 4 & 1 & 1 & 2 \end{pmatrix} \end{array}$$

Figure 2.24 shows two triangulations of this configuration. The triangulation on top is produced by the heights shown in the picture. But the other one is not produced by any choice of heights, as we prove by contradiction. To simplify the argument, we first observe that there is no loss of generality in assuming that the three height values for the interior triangle are zero (see Exercise 2.1). The other three values must be strictly positive in order for the three interior points to be lower vertices of the lifted configuration.

Moreover, using the labels of Figure 2.23, the exact condition on ω_1, ω_2, and ω_3 that makes the edge 15 appear in the second triangulation is that $\omega_1 < \omega_2$. The same happens for the diagonals 26 and 34, giving rise to the following impossible sequence of inequalities.

$$\omega_1 < \omega_2 < \omega_3 < \omega_1 .$$

See Figure 2.25 for an illustration of the lifting discussed.

2.2.2 A famous example: the Delaunay triangulation

Perhaps due to their easy construction, regular triangulations are often the prime examples of triangulations. Actually, special choices of the height vector ω produce some of the most studied and nicest triangulations. Here we study one of them, and more important examples will come in Section 4.3.

The *Delaunay triangulation* is arguably the most important triangulation of a point set for applications, and there is a vast literature about it (see, for example, [45, 129]). On the one hand, it is geometrically dual to the *Voronoi diagram* of the point set. The Voronoi cell of a point $\mathbf{p}$ in a point configuration $\mathbf{A}$ is the locus of points that are at least as close to $\mathbf{p}$ as to any other point of $\mathbf{A}$. The Voronoi diagram of $\mathbf{A}$ is the polyhedral complex whose maximal cells are the Voronoi cells. Clearly, Voronoi diagrams are important tools for solving proximity questions. This means Delaunay triangulations are useful too, since they carry exactly the same combinatorial information as Voronoi diagrams. As another example of usefulness, the set of edges of a Delaunay triangulation of any point set in the plane contains the minimum Euclidean spanning tree of the point set.

On the other hand, the Delaunay triangulation of a point set is considered to be one of the most uniform triangulations of it, meaning that its simplices are, on average, as close to regular simplices as possible. This makes a Delaunay triangulation a good candidate for meshing problems. For example, in two dimensions the Delaunay triangulation of $\mathbf{A}$ has, for example, the following properties: It is the triangulation that minimizes both the maximum angle and the maximum circumradius of its triangles (it also maximizes the minimum circumradius). See Corollary 3.2.7 for the proofs of these and other properties.

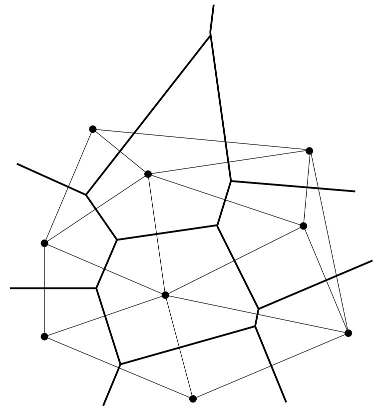

Figure 2.26: A Voronoi diagram and its dual Delaunay triangulation.

Delaunay triangulations are also important examples of regular triangulations.

Definition 2.2.6. Let $\mathbf{A} \subset \mathbb{R}^d$ be a finite point set of dimension d. If the lifting procedure applied to the heights

$$\omega(i) = ||\mathbf{p}_i||^2 = {a_1}^2 + {a_2}^2 + \cdots + {a_d}^2, \text{ for each } \mathbf{p}_i = (a_1, \ldots, a_d) \in \mathbf{A},$$

produces a triangulation, then this triangulation is called the *Delaunay triangulation* of $\mathbf{A}$.

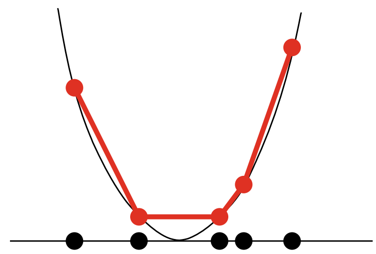

Figure 2.27: Delaunay triangulation of a 1-dimensional configuration.

Put differently, we are lifting the points onto a paraboloid (as indicated in Figure 2.27). Let us further analyze this triangulation and see when it is well-defined. Observe that we are assuming that our configuration is full-dimensional. This simplifies some of the arguments in the following results.

Lemma 2.2.7. *Let $\mathbf{C} \subset \mathbb{R}^{d+1}$ be the paraboloid given by the equation*

$$x_{d+1} = {x_1}^2 + \cdots + {x_d}^2,$$

and let $\mathbf{H} \subset \mathbb{R}^{d+1}$ be a non-vertical hyperplane, that is, one whose normal vector has non-zero last coordinate. Let $\mathbf{S}$ be the projection of $\mathbf{H} \cap \mathbf{C}$ into $\mathbb{R}^d$ obtained by dropping the last coordinate. Then $\mathbf{S}$ is either empty, a single point, or a sphere in $\mathbb{R}^d$. See Figure 2.28.

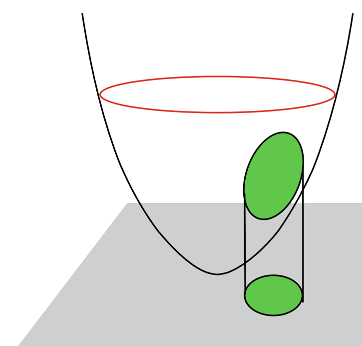

Figure 2.28: A hyperplane intersected with the paraboloid.

Proof. Since $\mathbf{H}$ is not vertical, we can isolate the variable x_{d+1} in its defining equation. That is, $\mathbf{H}$ is defined by an equation

$$x_{d+1} = \lambda_1 x_1 + \cdots + \lambda_d x_d + \lambda_0.$$

The intersection of $\mathbf{H}$ and $\mathbf{C}$ then satisfies the equation

$$x_1^2 + \cdots + x_d^2 = \lambda_1 x_1 + \cdots + \lambda_d x_d + \lambda_0,$$

which is equivalent to

$$(x_1 - \lambda_1/2)^2 + \cdots + (x_d - \lambda_d/2)^2 = (\lambda_1/2)^2 + \cdots + (\lambda_d/2)^2 + \lambda_0.$$

Since this equation does not involve the variable x_{d+1}, it is satisfied on the projection of $\mathbf{C} \cap \mathbf{H}$. Depending on whether the right-hand side is negative, zero or positive, it defines the empty set, a single point, or a sphere. $\square$

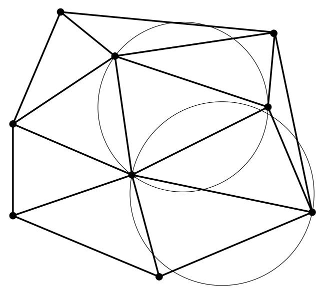

Figure 2.29: The empty-sphere property for an edge and a triangle in a Delaunay triangulation.

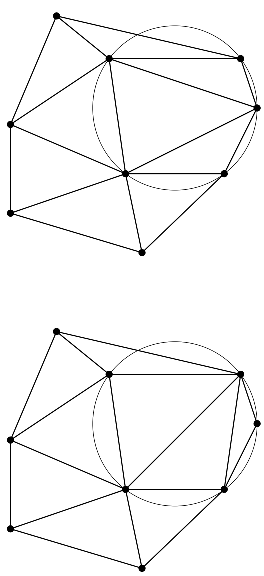

Figure 2.30: The empty sphere property of two Delaunay triangulations for a degenerate point configuration.

The following result indicates that the facets (highest dimension cells) of the Delaunay subdivision can be identified by a simple "empty sphere criterion". That this is the case beyond two dimensional space was first observed in [204].

Corollary 2.2.8. *Let* $\mathbf{A} \subset \mathbb{R}^d$ *be a finite point set and let* $\omega : \mathbf{A} \to \mathbb{R}$ *be the height function*

$$\omega(i) = ||\mathbf{p}_i||^2.$$

Let $\mathbf{B} \subset \mathbf{A}$ *be a subset such that* $\operatorname{conv}(\mathbf{B})$ *is full-dimensional. Then* $\mathbf{B}$ *corresponds to the vertex set of a lower facet of the lifted point set if and only if there is a sphere passing through all points of* $\mathbf{B}$ *and leaving all points of* $\mathbf{A} \setminus \mathbf{B}$ *outside.*

Proof. Since $\mathbf{B}$ is full-dimensional, there is at most one sphere passing through all points of $\mathbf{B}$. If there is no such sphere, then by Lemma 2.2.7, $\mathbf{B}$ is not lifted to a hyperplane, hence it does not lie on a facet of the lifted point set. If there is such a sphere $\mathbf{S}$, then there is a hyperplane $\mathbf{H_B}$ passing through the lifted point set. Points in the interior, on the surface and in the exterior of the sphere $\mathbf{S}$ are lifted respectively to points below $\mathbf{H_B}$, on $\mathbf{H_B}$ or above $\mathbf{H_B}$ (see Figure 2.28). Then $\mathbf{B}$ is lifted to the vertex set of a lower facet if and only if there is no point in the interior of the sphere and the points on the surface are precisely those of $\mathbf{B}$. ☐

In particular, simplices in the Delaunay triangulation are characterized by the "empty sphere" property: σ is a simplex in it if and only if there is a Euclidean sphere circumscribed to σ (i.e., with all vertices of σ on the surface of the sphere) with the rest of the points of $\mathbf{A}$ outside; see Figure 2.29.

The corollary also gives a simple sufficient condition to guarantee that the Delaunay triangulation is uniquely defined: that no $d+2$ points lie in the surface of any sphere. But, what if this is not the case and the lifting to the paraboloid does not produce a triangulation? In this case there are two approaches to defining a Delaunay triangulation. You can either

1. Give up uniqueness and call *Delaunay triangulations* all the triangulations that "refine" the projection of the lower envelope of the lifted point set, or

2. Give up simpliciality and call the projection that you get a *Delaunay subdivision*. This is indeed a subdivision of $\operatorname{conv}(\mathbf{A})$ into convex polytopes that intersect face to face (that is, it is a *polyhedral complex*).

In applications, the first approach is what is usually used. After all, one of the reasons to construct the Delaunay triangulation is that one wants a triangulation in the first place! In this setting, a Delaunay triangulation is any triangulation $\mathscr{T}$ whose simplices have the following "weak" empty sphere property: for any $\sigma \in \mathscr{T}$, there is a Euclidean sphere circumscribed to σ with no point of $\mathbf{A}$ inside (but perhaps with extra points on its surface).

But, conceptually, the second approach is nicer. For example, the Delaunay subdivision is still dual to the corresponding Voronoi diagram; as

Figure 2.31 shows, its cells are characterized by the same empty sphere property as in the general position case: a polytope $\sigma \subset \mathrm{conv}(\mathbf{A})$ with vertex set contained in $\mathbf{A}$ is a cell in the complex if and only if there is a sphere circumscribed to σ with the rest of the points of $\mathbf{A}$ outside.

As a curious historical remark, although the spelling "Delaunay" is the standard one in Computational Geometry, the spelling "Delone" triangulations is also used in other areas. Both spellings honor the same Russian geometer Boris Nikolaevich Delone [123], who studied these triangulations mostly for infinite periodic point configurations (lattices) [103].

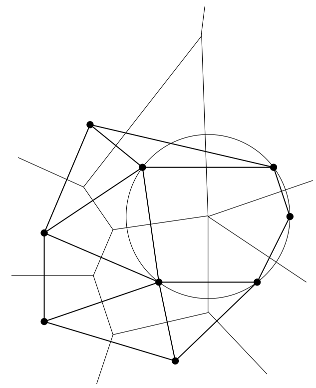

Figure 2.31: The Delaunay subdivision of a degenerate point set and its dual Voronoi diagram.

2.2.3 Regular subdivisions and their structure

Let a configuration $\mathbf{A}$ be given. In Proposition 2.2.4 we have seen that for any sufficiently generic choice of a height function $\omega : \mathbf{A} \to \mathbb{R}$, the projection of the lower envelope of $\mathbf{A}^\omega$ is a regular triangulation.

What happens if we take a height vector ω for which the lifted point set $\mathbf{A}^\omega$ has some non-simplicial lower facets? As has been hinted in Figure 2.31 of the Delaunay subdivision of a degenerate point set, instead of being a triangulation, the projection of the lower envelope is a collection of more complicated polytopes, each with vertices in $\mathbf{A}$. This collection is called the *regular (polyhedral) subdivision* of $\mathbf{A}$, produced by the height vector ω. The different polytopes in it are called *cells* of the subdivision. Clearly, they still satisfy the three properties of Definition 2.2.1.

Polyhedral subdivisions play a fundamental role in this book; triangulations are nothing but particular cases of them. Even more, sometimes you need to understand general polyhedral subdivisions even if your ultimate goal is to work only with triangulations.

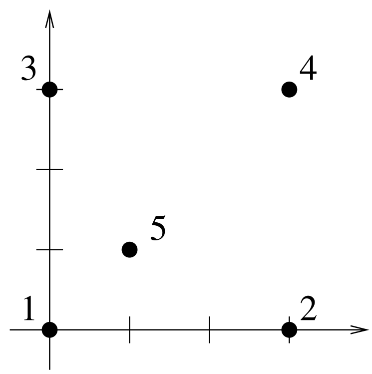

Figure 2.32: Five points in the plane.

To justify the need for a more careful definition, other than just saying that a polyhedral subdivision is a *polyhedral complex that covers* $\mathrm{conv}(\mathbf{A})$, let us look carefully at a particular example:

Example 2.2.9. Consider the configuration

$$\begin{array}{cc} & \begin{array}{ccccc} \mathit{1} & \mathit{2} & \mathit{3} & \mathit{4} & \mathit{5} \end{array} \\ \mathbf{A} = & \begin{pmatrix} 0 & 3 & 0 & 3 & 1 \\ 0 & 0 & 3 & 3 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \end{array}$$

of five points in the plane that appeared already in Figure 2.12. Figure 2.33 shows the six possible regular polyhedral subdivisions of it. Four of them are triangulations.

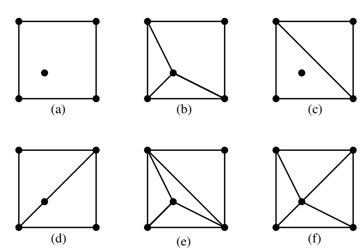

Figure 2.33: Subdivisions of the point set introduced in Figure 2.32.

To check that these are indeed regular subdivisions and that the list is complete, we look at precisely what height vectors produce each of these six subdivisions. As in the previous section, there is no loss of generality in restricting attention to height vectors with $\omega_1 = \omega_2 = \omega_3 = 0$ (see Exercise 2.1). Hence, we do a case study depending on the values of ω_4 and ω_5. The results can be represented in the plane. The exact conditions that produce the six subdivisions (a) to (f) are, respectively (check this!):

(a) $\omega_5 \geq 0$, $\omega_4 = 0$.

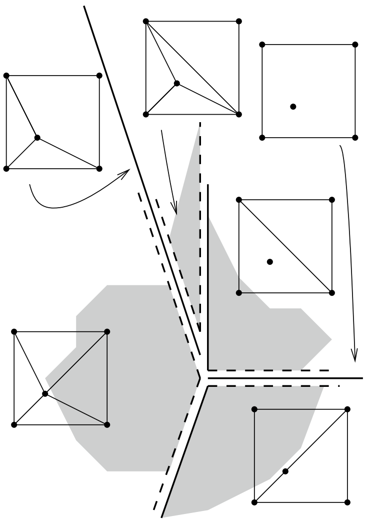

Figure 2.34: Stratification in the space of lifts.

(b) $\omega_5 < 0$, $\omega_4 + 3\omega_5 = 0$.

(c) $\omega_5 \geq 0$, $\omega_4 > 0$.

(d) $\omega_4 < 0$, $\omega_4 \leq 3\omega_5$.

(e) $\omega_5 < 0$, $\omega_4 + 3\omega_5 > 0$.

(f) $\omega_5 < 0$, $\omega_4 + 3\omega_5 < 0$, $\omega_4 > 3\omega_5$.

Figure 2.34 shows these conditions pictorially. Height ω_5 is the horizontal coordinate and ω_4 is the vertical one.

One thing we observe is that the regions that produce triangulations are precisely the ones of full-dimension, in agreement with our assertion that sufficiently random heights produce triangulations.

But Figure 2.34 also shows that the stratification we obtain in the two-dimensional space of heights has defects. Some cones are neither open nor closed, but contain only one of their two boundary rays. This, and other unsatisfactory features come from the fact that our temporary definition of regular subdivision does not distinguish between the case where the interior point $\mathbf{p}_5$ is lifted to lie right *on* the lower envelope of the convex hull of the rest of the points (which should be considered a degenerate situation), and the case where it is lifted *above* the lower envelope of the rest.

To make this distinction, the language of Section 2.1.3 comes to the rescue. Instead of defining the regular subdivision by projecting the lower faces of the *polytope* $\operatorname{conv}(\mathbf{A}^\omega)$, we project the lower faces of the *point configuration* $\mathbf{A}^\omega$:

Definition 2.2.10. Let $\mathbf{A} \subset \mathbb{R}^m$ be a point configuration with n elements and let $\omega : J \to \mathbb{R}$ be a "height vector". We use indistinctly $\omega(j)$ and ω_j to refer to the height given to j, although the latter is preferred.

Let $\mathbf{A}^\omega$ be the *lifted point configuration*, which has the same labels as $\mathbf{A}$ and a point $\mathbf{p}_j^\omega := (\mathbf{p}_j, \omega_j) \in \mathbb{R}^{d+1}$ for each $j \in J$. A *lower face* of $\mathbf{A}^\omega$ is any face $F = \operatorname{face}_{\mathbf{A}}(J, \psi)$ of $\mathbf{A}$ in the direction of some functional ψ with last coordinate positive (put simply, a face that is visible from below).

The *regular subdivision* of $\mathbf{A}$ produced by ω is the set of lower faces of the point configuration $\mathbf{A}^\omega$. This regular subdivision will be denoted by $\mathscr{S}(\mathbf{A}, \omega)$.

Observe in this definition the convenience of using the same labels in $\mathbf{A}^\omega$ and in $\mathbf{A}$. This allows us to say that the lower faces of the former are cells in a subdivision of the latter. Also, observe that here (and in the rest of the book) our height function is defined on the labels rather than on the points. Aside from slightly simplifying notation (we write $(\mathbf{p}_i, \omega_i)$ instead of $(\mathbf{p}_i, \omega(\mathbf{p}_i))$ for the lifted points), this takes care of the possibility that $\mathbf{A}$ may have repeated copies of a point and that we may want to lift them to different heights.

We now look at regular subdivisions of the point set of Example 2.2.9 in this new language. To give a subdivision we only list the full-dimensional cells, since the others are simply the faces of them (put differently, lower faces of $\mathbf{A}^\omega$ are all the faces of lower facets). Also, as an abbreviation, we write 123, instead of $\{1,2,3\}$.

Example 2.2.11 (Example 2.2.9 continued)**.** The triangulation (c) in the above list can be represented by the following list of *subsets* of $\mathbf{A}$:

$$\{123,234\}.$$

But there is another regular subdivision with the same set of geometric cells, i.e., the same convex hulls. It is not a triangulation, since it contains cells with dependences. Its list of full-dimensional cells is:

$$\{1235,234\}.$$

This subdivision is produced by heights lying in the positive vertical axis in the representation of Figure 2.34.

Let us work out the list of regular subdivisions of the five point example again, with this new definition, and with the same convention that $\omega_1 = \omega_2 = \omega_3 = 0$. Our point set now has exactly nine regular polyhedral subdivisions, namely:

(a1) $\{12345\}$, obtained whenever $\omega_5 = 0$, $\omega_4 = 0$.

(a2) $\{1234\}$, obtained whenever $\omega_5 > 0$, $\omega_4 = 0$.

(b) $\{135,125,2345\}$, obtained whenever $\omega_5 < 0$, $\omega_4 + 3\omega_5 = 0$.

(c1) $\{1235,234\}$, obtained whenever $\omega_5 = 0$, $\omega_4 > 0$.

(c2) $\{123,234\}$, obtained whenever $\omega_5 > 0$, $\omega_4 > 0$.

(d1) $\{1345,1245\}$, obtained whenever $\omega_4 < 0$, $\omega_4 = 3\omega_5$.

(d2) $\{134,124\}$, obtained whenever $\omega_4 < 0$, $\omega_4 < 3\omega_5$.

(e) $\{125,135,235,234\}$, obtained whenever $\omega_5 < 0$, $\omega_4 + 3\omega_5 > 0$.

(f) $\{125,135,245,345\}$, obtained whenever $\omega_5 < 0$, $\omega_4 + 3\omega_5 < 0$, $\omega_4 > 3\omega_5$.

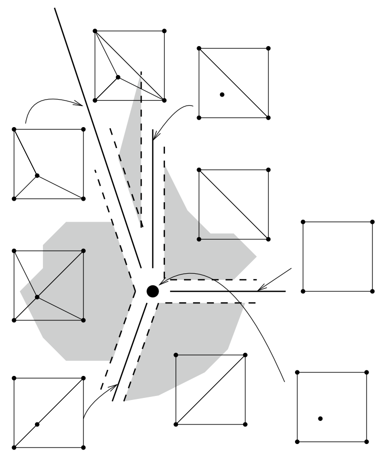

Figure 2.35: Stratification in the space of lifts, with the modified definition.

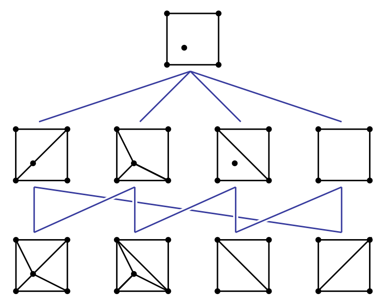

Figure 2.36: Poset of the subdivisions.

Figure 2.35 illustrates this catalogue. The difference between Figure 2.35 and Figure 2.34 is that each of the subdivisions (a), (c), and (d) of our first computation splits into two subdivisions, depending on whether the interior point $\mathbf{p}_5$ is lifted to lie in the lower envelope of $\mathrm{conv}(\mathbf{A}^\omega)$ or above it. In the figures, we distinguish these two cases by drawing the point or not. Observe that there are two different subdivisions into a single cell, (a1) and (a2): (a1) is called the trivial subdivision because the cell is the whole of $\mathbf{A}$ (Example 2.3.5).

To finish convincing the reader that this combinatorial framework gives a nicer set of regular subdivisions, let us look at the *refinement poset of regular subdivisions* of a point configuration. Given two regular subdivisions $\mathscr{S}_1$ and $\mathscr{S}_2$ of $\mathbf{A}$, we say that $\mathscr{S}_1$ *refines* $\mathscr{S}_2$ if every element of $\mathscr{S}_1$ is contained in some element of $\mathscr{S}_2$. This relation induces a partial order; hence, the set of regular subdivisions is a partially ordered set, the *regular refinement poset* for short. Intuitively, if ω is a height vector and $\omega' = \omega + \varepsilon$ is a sufficiently small perturbation of it, then the regular subdivision $\mathscr{S}(\mathbf{A}, \omega')$ will be a refinement of $\mathscr{S}(\mathbf{A}, \omega)$. This is proved in Lemma 2.3.15.

With our temporary approach to regular polyhedral subdivisions, the refinement poset has three levels. On top, there is the trivial subdivision (a). Below this are (b), (c), and (d), which are incomparable to one another. Finally, (f) is the common refinement of (b) and (d), and (e) is the common refinement of (b) and (c). With the combinatorial definition, the poset has a different and a nicer structure in which, for example, the four triangulations are exactly the elements in the lower level (the minimal elements in the poset). Actually, the poset becomes isomorphic to the poset of non-empty faces of a quadrilateral. Furthermore Figure 2.35 shows the relative interiors of the normal cones of the quadrilateral of Figure 2.37. That the same happens for regular subdivisions of any point configuration is one of the fundamental results in this field.

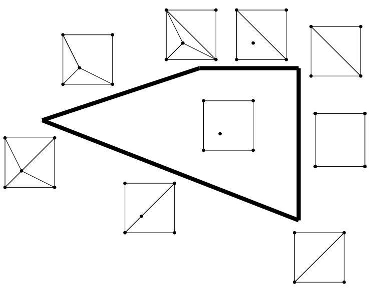

Figure 2.37: The poset of regular subdivisions, as the faces of a quadrilateral.

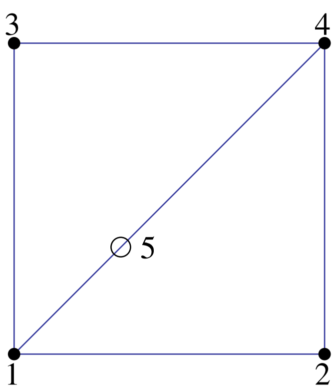

Figure 2.38: $\{1235, 234\}$ is not a subdivision while $\{1235, 2345\}$ is and $\{123, 234\}$ is another.

2.3 A bullet-proof definition of polyhedral subdivisions

In this section we elaborate more on the idea of representing cells in triangulations and subdivisions of $\mathbf{A}$ as subconfigurations of $\mathbf{A}$, rather than as convex hulls of subsets, which led to an unsatisfactory structure. We continue using geometric language to refer to those subsets of indices since we introduced the combinatorial framework in which, for example, faces of a configuration are subconfigurations. Faces, carriers, etc. were defined for configurations already.

2.3.1 Polyhedral subdivisions

We are finally ready to present our *master definition* that considers polyhedral subdivisions and triangulations of a point configuration $\mathbf{A}$ with label set J as collections of subsets of J with certain properties, mimicking the ones in our "intuitive" definition of triangulations provided in Definition 2.2.1.

Definition 2.3.1 (Polyhedral subdivision)**.** Let $\mathbf{A}$ be a point configuration in $\mathbb{R}^m$, with a set of labels J. A collection $\mathscr{S}$ of subsets of J is a *polyhedral subdivision* of $\mathbf{A}$ if it satisfies the following conditions:

(CP) If $C \in \mathscr{S}$ and $F \leq C$, then $F \in \mathscr{S}$ as well. (Closure Property)

(UP) $\bigcup_{C \in \mathscr{S}} \operatorname{conv}_{\mathbf{A}}(C) \supseteq \operatorname{conv}_{\mathbf{A}}(J)$. (Union Property)

(IP) If $C \neq C'$ are two cells in $\mathscr{S}$, then $\operatorname{relint}_{\mathbf{A}}(C) \cap \operatorname{relint}_{\mathbf{A}}(C') = \emptyset$. (Intersection Property)

The elements of a polyhedral subdivision $\mathscr{S}$ are called *cells*. Cells of dimension k are called k-cells. Cells of the same dimension as $\mathbf{A}$ are *full-dimensional* or *maximal*. Cells of dimension zero are called *vertices* of $\mathscr{S}$. Independent cells are called *simplices*. A *triangulation* of $\mathbf{A}$ is a polyhedral subdivision all of whose cells are simplices. The set of polyhedral subdivisions of a point configuration $\mathbf{A}$ will be denoted $\operatorname{Subdivs}(\mathbf{A})$.

Sometimes we want to state (IP) for a cell and all of its faces. For this, we introduce a handy term.

Definition 2.3.2. Two subsets C and C' of J *intersect properly* if all their faces satisfy Property (IP) of Definition 2.3.1; they intersect *improperly* otherwise.

Remark 2.3.3.

- Property (IP) can be reformulated as: All pairs of maximal cells intersect properly.
- The union property could have been written with an equality, since containment in the opposite direction is obvious.
- Similarly, Property (UP) could have been written with "relint" instead of "conv" in one or both sides, because Property (CP) implies that once a relative interior is covered, the convex hull is covered too.
- In property (UP), the restriction of the union to maximal cells is equivalent.

We now collect some facts in order to get used to the definition:

Lemma 2.3.4. *Let $\mathscr{S}$ be a polyhedral subdivision of a point configuration $\mathbf{A} \subset \mathbb{R}^m$ with a label set J. Then:*

(i) All maximal cells in $\mathscr{S}$ are full-dimensional, i.e., their dimension equals the dimension of $\mathbf{A}$.

(ii) For all $C, C' \in \mathscr{S}$, if $C \subseteq C'$ then $C \leq C'$.

(iii) For all $C, C' \in \mathscr{S}$, if $C \subseteq C'$ then either $C = C'$ or $\dim C < \dim C'$.

(iv) For any face F of $\mathbf{A}$, the set $\mathscr{S}|_F \subseteq \mathscr{S}$ of all cells C in $\mathscr{S}$ with $C \subseteq F$ is a polyhedral subdivision of $\mathbf{A}|_F$.

(v) If C and C' are cells of $\mathscr{S}$, then $C \cap C'$ is a face of both (and hence a cell in $\mathscr{S}$). Moreover,

$$\operatorname{conv}_{\mathbf{A}}(C \cap C') = \operatorname{conv}_{\mathbf{A}}(C) \cap \operatorname{conv}_{\mathbf{A}}(C').$$

Proof. In order to prove Part (i), assume there is a maximal k-cell F with $k < \dim \mathbf{A} = d$. Pick an arbitrary point $\mathbf{x} \in \operatorname{relint}(F)$, and consider a point $\mathbf{y} \in \operatorname{conv}(\mathbf{A})$ very close to $\mathbf{x}$ in general position, i.e., no point on the half-open segment $(\mathbf{x}, \mathbf{y}]$ from $\mathbf{y}$ to $\mathbf{x}$ is contained in any cell of dimension less than d. Then, by (UP), $\mathbf{y}$ must be contained in the relative interior of a full-dimensional cell $C_{\mathbf{y}}$. We claim that F is a face of $C_{\mathbf{y}}$. First note that, because of (IP), $\mathbf{x}$ cannot be in the relative interior of $C_{\mathbf{y}}$. However, all points on $(\mathbf{x}, \mathbf{y}]$ must be in $C_{\mathbf{y}}$ as well because $(\mathbf{x}, \mathbf{y}]$ does not intersect any lower-dimensional face. Therefore, $\mathbf{x}$ is in the boundary of $C_{\mathbf{y}}$. Consider the carrier of point $\mathbf{x}$ in $C_{\mathbf{y}}$, $\operatorname{carrier}_{\mathbf{A}}(\mathbf{x}, C_{\mathbf{y}})$. This is a face of $C_{\mathbf{y}}$, thus it is in $\mathscr{S}$, and it contains $\mathbf{x}$ in its relative interior. Since F contains $\mathbf{x}$ in its relative interior as well, F must be equal to $\operatorname{carrier}_{\mathbf{A}}(\mathbf{x}, C_{\mathbf{y}})$ by Property (IP). Thus F is a face of $C_{\mathbf{y}}$, and we are done since we have reached a contradiction.

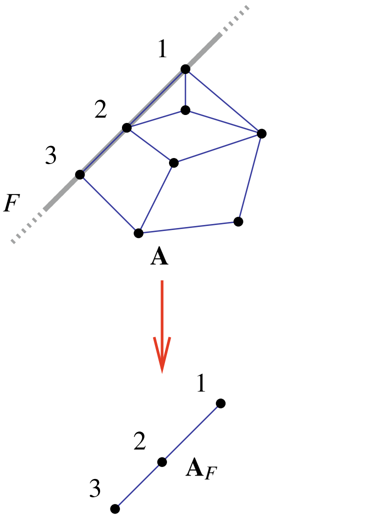

Figure 2.39: A polyhedral subdivision of $\mathbf{A}$ and $\mathbf{A}_F$.

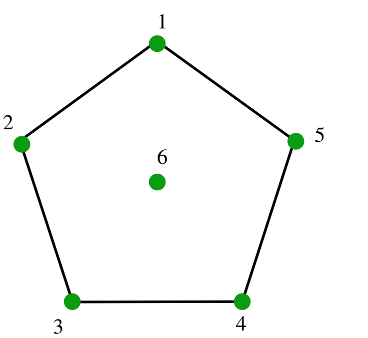

Figure 2.40: $S = \{123456\}$ is the trivial subdivision.

Part (ii) can be proved as follows: Let $C, C' \in \mathscr{S}$ with $C \neq C'$ and $C \subset C'$. Property (CP) implies that all faces of C' are in $\mathscr{S}$ as well. In particular, the carrier of C in C' is in $\mathscr{S}$. Since the relative interior of C is contained in the relative interior of $\mathrm{carrier}(C, C')$ (in particular, the relative interiors of these subconfigurations have a non-empty intersection), Property (IP) implies that $C = \mathrm{carrier}(C, C')$, which means $C \leq C'$.

Part (iii) is a direct consequence of Part (ii): The only face C of a point configuration C' with $\dim(C) = \dim(C')$ is C' itself. All other faces have strictly smaller dimensions.

In order to see Part (iv) we first note that (IP) is trivially fulfilled for all subsets of a polyhedral subdivision, in particular for $\mathscr{S}|_F$. Moreover, (CP) is satisfied because Property (CP) holds for $\mathscr{S}$, and whenever $C \subseteq F$, then $C' \subseteq F$ holds also for all faces C' of C.

It remains to show that $\mathrm{conv}_{\mathbf{A}}(F)$ is covered by all cells contained in it (without loss of generality, F is a proper face). Since Property (UP) holds for $\mathscr{S}$, each point $\mathbf{x}$ in $\mathrm{conv}_{\mathbf{A}}(F)$ is contained in $\mathrm{conv}_{\mathbf{A}}(C'_{\mathbf{x}})$ for some cell $C'_{\mathbf{x}} \in \mathscr{S}$. Let $\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}}) \in \mathscr{S}$ be the carrier of $\mathbf{x}$ in $C'_{\mathbf{x}}$. Then, by Lemma 2.1.23, $\mathbf{x} \in \mathrm{relint}(\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}}))$.

We claim that $\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})$ is in $\mathscr{S}|_F$. For this we simply need to show that $\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}}) \subseteq F$. Since $\mathbf{A}|_F = \mathbf{A} \cap \mathbf{H}_F$ for a supporting hyperplane $\mathbf{H}_F$ of F, we know that $\mathbf{A}|_{\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})} \cap \mathbf{H}_F \subseteq \mathbf{A} \cap \mathbf{H}_F = \mathbf{A}|_F$.

Since $\mathbf{A}|_{\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})} \cap \mathbf{H}_F$ is a face of $\mathbf{A}|_{\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})}$ that contains $\mathbf{x}$, we have, by the minimality of the carrier, that $\mathbf{A}|_{\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})} \subseteq \mathbf{H}_F$. Thus,

$$\mathbf{A}|_{\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})} \subseteq \mathbf{A}|_{\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})} \cap \mathbf{H}_F \subseteq \mathbf{A} \cap \mathbf{H}_F = \mathbf{A}|_F, \tag{2.9}$$

and $\mathrm{carrier}_{\mathbf{A}}(\mathbf{x}, C'_{\mathbf{x}})$ is in $\mathscr{S}|_F$, as desired.

Part (v) is trivially true if $C \cap C'$ is empty. If it is not empty, let R and R' be the carriers of $C \cap C'$ in C and C', respectively. R and R' are faces of C and C', hence they are cells in $\mathscr{S}$, and have $C \cap C' \neq \emptyset$ in their relative interiors. Hence, by (IP), $R = R'$.

In the equation about convex hulls, the containment $\subseteq$ is true for the convex hulls of arbitrary sets. So, let $\mathbf{x} \in \mathrm{conv}(C) \cap \mathrm{conv}(C')$ and let us see that $\mathbf{x}$ is also in $\mathrm{conv}(C \cap C')$. Let R and R' be the carriers of $\mathbf{x}$ in C and C', respectively. We again conclude that $R = R'$, hence $R \subseteq C \cap C'$ and $\mathbf{x} \in \mathrm{conv}(R) \subset \mathrm{conv}(C \cap C')$. $\square$

Example 2.3.5 (The trivial subdivision)**.** Every point configuration $\mathbf{A}$ has the *trivial subdivision*, consisting of all the faces of $\mathbf{A}$ (including the non-proper one, that is, the full set J of labels). That this is indeed a subdivision according to Definition 2.3.1 is easy to show: (CP) follows from the fact that "a face of a face is a face". (UP) is obvious since J is one of our cells, and (IP) follows from the fact that different faces of $\mathbf{A}$ have disjoint relative interiors.

If $\mathbf{A}$ is independent, then the trivial subdivision is the only polyhedral subdivision, and it is a triangulation.

Example 2.3.6 (0-dimensional point configurations)**.** A configuration of dimension zero consists of several copies (say k) of a single point. By property (IP), a subdivision cannot have more than one (non-empty) cell. Conversely, any non-empty subconfiguration of $\mathbf{A}$ (together with the empty set, which is its unique proper face) is a polyhedral subdivision. That is, $\mathbf{A}$ has exactly $2^k - 1$ subdivisions in bijection with the non-empty subconfigurations of $\mathbf{A}$, labeled by subsets of $\mathbf{A}$'s label set J. Triangulations are the ones corresponding to one-element subconfigurations, and there are k of them.

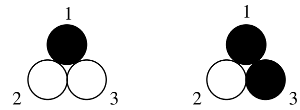

Figure 2.41: Two subdivisions of a zero-dimensional configuration with three copies of the same point. The one on the left, $\{1\}$, is a triangulation.The one in the right, $\{13\}$, is not.

Example 2.3.7 (1-dimensional point configurations, no repeated points)**.** For a 1-dimensional point configuration, being in general position is equivalent to not having repeated points. The convex hull is a segment, and any subdivision $\mathscr{S}$ is geometrically a chain of sub-segments going from one end to the other with the endpoints of sub-segments elements of $\mathbf{A}$. A sub-segment may also contain some points of $\mathbf{A}$ in its interior. To get a polyhedral subdivision in the sense of Definition 2.3.1, we need to specify whether these interior points are part of the cell or not.

Put differently, if $\mathbf{A} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ is our point set (and the labels denote its ordering along the line $\mathbb{R}$), in a polyhedral subdivision $\mathscr{S}$ of $\mathbf{A}$ each of the $n-2$ points $\mathbf{p}_2$ through $\mathbf{p}_{n-1}$ will have one of the following three statuses: It may be an endpoint of two consecutive sub-segments of $\mathscr{S}$, it may belong to a cell of $\mathscr{S}$ without being one of its endpoints, or it may not belong to any cell of $\mathscr{S}$. If we represent these three possibilities as the three signs "$-$", "0", and "$+$" respectively, we get:

> Each polyhedral subdivision of $\mathbf{A}$ has an associated string of length $n-2$ on the alphabet $\{+, 0, -\}$.

We leave it to the reader to prove the following:

1. This correspondence is a bijection between the set of subdivisions of $\mathbf{A}$ and the 3^{n-2} element set $\{+, 0, -\}^{n-2}$.
2. Triangulations correspond to the strings without zeroes, hence there are 2^{n-2} of them.

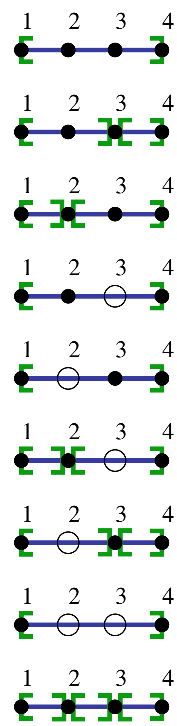

Figure 2.42: All subdivisions of four collinear points.

The most prominent structure on the set $\mathrm{Subdivs}(\mathbf{A})$ of all polyhedral subdivisions of a point configuration $\mathbf{A}$ is the *refinement poset*. It is often called the *Baues poset* of $\mathbf{A}$. Refinement of a polyhedral subdivision means roughly that some pieces of that subdivision are subdivided further. We have seen in Section 2.2.3 a restricted version using only the regular triangulations. Now we present a bigger more inclusive definition.

Definition 2.3.8 (Refinement)**.** Let $\mathscr{S}$ and $\mathscr{S}'$ be two polyhedral subdivisions of a point configuration $\mathbf{A}$. Then, $\mathscr{S}$ is a *refinement* of $\mathscr{S}'$—in brief notation: $\mathscr{S} \preceq \mathscr{S}'$—if for each $C \in \mathscr{S}$, there is a $C' \in \mathscr{S}'$ with $C \subseteq C'$.

With the refinement relation, $\mathrm{Subdivs}(\mathbf{A})$ is a *partially ordered set* (*poset* for short). This means the following (by definition of a poset):

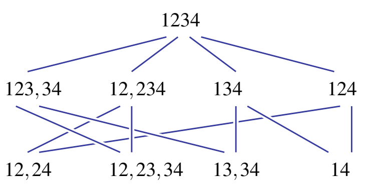

Figure 2.43: Refinement poset of four colinear points.

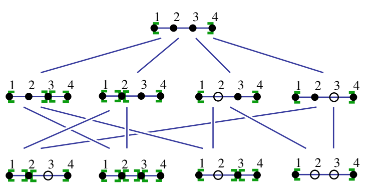

Figure 2.44: Refinement poset of four colinear points.

Lemma 2.3.9. *The refinement relation induces a partial order on the set of all polyhedral subdivisions of* $\mathbf{A}$*. That is, for every triple* $\mathscr{S}, \mathscr{S}', \mathscr{S}''$ *of polyhedral subdivisions of* $\mathbf{A}$*, one has:*

(i) $\mathscr{S} \preceq \mathscr{S}$

(ii) $\mathscr{S} \preceq \mathscr{S}'$ *and* $\mathscr{S}' \preceq \mathscr{S}$ *imply* $\mathscr{S} = \mathscr{S}'$

(iii) $\mathscr{S} \preceq \mathscr{S}'$ *and* $\mathscr{S}' \preceq \mathscr{S}''$ *imply* $\mathscr{S} \preceq \mathscr{S}''$

Using the refinement we can say a subdivision is *coarsest* if it refines the trivial subdivision and no other. Triangulations are the *finest subdivisions*.

Proof. Part (i) is clear by the definition of "$\preceq$". To prove Part (ii), let $\mathscr{S} \preceq \mathscr{S}'$ and $\mathscr{S}' \preceq \mathscr{S}$. Let C be an arbitrary maximal cell in $\mathscr{S}$. Since $\mathscr{S} \preceq \mathscr{S}'$, there is an $C' \in \mathscr{S}'$ with $C \subset C'$. Since $\mathscr{S}' \preceq \mathscr{S}$, there is a $C'' \in \mathscr{S}$ with $C' \subseteq C''$. Thus, $C \subseteq C''$, and, by Lemma 2.3.4(iii), either $C = C''$ or $\dim(C) < \dim(C'')$. Since C was chosen to be a maximal cell, only $C = C''$ is possible. This implies that in the chain $C \subseteq C' \subseteq C''$ we have equality everywhere. The remaining Part (iii) is again straightforward by definition of "$\preceq$". □

Recall that a maximal or minimal element in a poset is an element which is not strictly smaller or strictly greater, respectively, than any other one.

Lemma 2.3.10. *Let* $\mathbf{A}$ *be a point configuration. Then:*

(i) $\mathrm{Subdivs}(\mathbf{A})$ *has a unique maximal element, the trivial subdivision.*

(ii) A subdivision is a minimal element of $\mathrm{Subdivs}(\mathbf{A})$ *if and only if it is a triangulation.*

Proof. For Part (i), observe simply that every cell of every subdivision is contained in the full-dimensional cell J of the trivial subdivision. Hence, every subdivision refines the trivial one.

For Part (ii), here we only prove that triangulations are minimal elements in the poset. The converse is Corollary 2.3.18 in the next section, and will be proved using regular triangulations.

Let $\mathscr{S}$ be a triangulation and let $\mathscr{S}'$ be a subdivision that refines $\mathscr{S}$. Every cell F' of $\mathscr{S}'$ is contained in a cell F of $\mathscr{S}$ and, since F is independent, F' is a face of it. Hence, by property (CP) F' is a cell of $\mathscr{S}$. Conversely, if F is a face of $\mathscr{S}$, let $\mathbf{x}$ be a point in $\mathrm{relint}(F')$. Then $\mathbf{x}$ must also be in the relative interior of some face F' of $\mathscr{S}'$, by property (UP). By the previous argument, F' is a face of $\mathscr{S}$ as well. Since F and F' intersect in their relative interiors, $F = F'$ by property (IP). □

Let us close this section with a reassuring remark. The reader may wonder what parts of Section 2.2 need to be re-worked to match the combinatorial setting introduced in this one. The answer is "almost none", for the following reason. In that section we were primarily interested in triangulations. With the new definition, cells in a triangulation are affinely independent subsets and, hence, they are the vertex set of their convex hull. In particular, there is a one-to-one correspondence between the triangulations allowed by Definition 2.2.1 and those allowed by Definition 2.3.1, except in the following case: If the point configuration has repeated points, there will be several combinatorial representatives of each "geometric" triangulation. Indeed, each combinatorial triangulation will select a particular copy of each repeated point, since a cell with repeated points is not independent and two combinatorial simplices using different copies of a repeated point violate the intersection property (IP).

2.3.2 Regular subdivisions, again

We now look again at polyhedral subdivisions. The first thing to check is that they are indeed polyhedral subdivisions, according to our "bullet-proof" definition.

Lemma 2.3.11. *$\mathscr{S}(\mathbf{A},\omega)$ is a polyhedral subdivision of $\mathbf{A}$, for every ω.*

Proof. Observe that every lower face F of $\mathbf{A}^{\omega}$, by definition, lies in a non-vertical hyperplane (the one in which the functional ψ is constant). Hence, the projection $\pi : \mathbf{A}^{\omega} \to \mathbf{A}$ that forgets the last coordinate is an affine isomorphism between $\mathbf{A}_F$ and $\mathbf{A}^{\omega}_F$. In particular, the face structure of F is the same in $\mathbf{A}$ and $\mathbf{A}^{\omega}$.

The closure property (CP) then follows from the fact that if $F' < F \le J$ are faces of $\mathbf{A}^{\omega}$ then every functional for the face F can be slightly perturbed to a functional of the face F' (for people familiar with polyhedral geometry, what we are saying is that the normal cone of F is a face of the normal cone of F'). In particular, if F is a lower face, then any $F' < F$ is a lower face as well. The intersection property for the projected faces follows from the intersection property of the faces of $\mathbf{A}^{\omega}$ [339].

For the union property, let $\mathbf{x}$ be a point in the relative interior of $\mathbf{A}$. The intersection of $\mathbf{x} \times \mathbb{R}$ with $\mathrm{conv}(\mathbf{A}^{\omega})$ is a vertical segment from a bottom point $\mathbf{x}_1$ to a top point $\mathbf{x}_2$. Let F be any proper face of $\mathbf{A}^{\omega}$ with $\mathbf{x}_1 \in \mathrm{conv}_{\mathbf{A}^{\omega}} F$, which exists since $\mathbf{x}_1$ is in the boundary of $\mathrm{conv}(\mathbf{A}^{\omega})$. Let ψ be any linear functional selecting F as a face of $\mathbf{A}^{\omega}$. We have that $\psi(\mathbf{x}_2) > \psi(\mathbf{x}_1)$, otherwise ψ would be constant in the whole segment from $\mathbf{x}_1$ to $\mathbf{x}_2$, hence constant on $\mathrm{conv}(\mathbf{A}^{\omega})$, because this segment crosses its relative interior, contradicting the fact that F is a proper face.

But $\psi(\mathbf{x}_2) > \psi(\mathbf{x}_1)$ implies that ψ is positive on the last coordinate, so F must be a lower face. Its projection is a cell in $\mathscr{S}(\mathbf{A},\omega)$ covering the point $\mathbf{x}$. ☐

Example 2.3.12 (Example 2.3.5 continued)**.** The trivial subdivision is the regular subdivision obtained with the zero height vector or, more generally, with any height vector that is the restriction to $\mathbf{A}$ of an affine function $\mathbb{R}^m \to \mathbb{R}$. In this case, $\mathbf{A}^{\omega}$ has the same face structure as $\mathbf{A}$, since they are affinely equivalent configurations, and every face is a lower face.

Example 2.3.13 (Example 2.3.6 continued)**.** Every subdivision of a zero-dimensional point configuration is regular. Remember that every subdivision

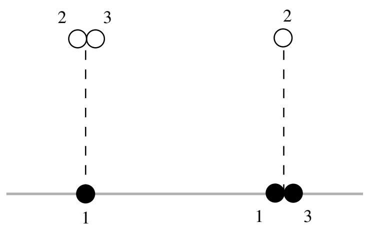

Figure 2.45: Heights that produce the regular subdivisions shown in Figure 2.41.

has a unique (non-empty) cell, say F. A height function that produces this subdivision is one that gives the same height to all points in F and greater height to the others.

Example 2.3.14 (Example 2.3.7 continued)**.** In this case, all subdivisions are again regular. The proof is left to the reader as Exercise 5.1.

Lemma 2.3.15. *Let $\mathbf{A}$ be a point configuration in $\mathbb{R}^m$ and let $\omega : J \to \mathbb{R}$ be a height function. Let $\mathscr{S} = \mathscr{S}(\mathbf{A}, \omega)$ be the regular subdivision of $\mathbf{A}$ produced by ω. Then,*

1. *If ω is sufficiently generic, then $\mathscr{S}$ is a triangulation.*
2. *If F is a face of $\mathbf{A}$, then the restriction of $\mathscr{S}$ to F equals the regular subdivision of $\mathbf{A}_F$ obtained with ω. Symbolically:*
$$\mathscr{S}(\mathbf{A}, \omega)|_F = \mathscr{S}(\mathbf{A}_F, \omega|_F).$$
3. *There is an $\varepsilon > 0$ such that, for every height function $\omega' : J \to \mathbb{R}$ that is ε-close to ω, namely $|\omega(j) - \omega'(j)| < \varepsilon$ for all j, we have that $\mathscr{S}(\mathbf{A}, \omega') \preceq \mathscr{S}(\mathbf{A}, \omega)$.*

Proof. (1) Assume, without loss of generality, that $\mathbf{A}$ is represented by a matrix of full rank. That is, that $\mathbf{A}$ has dimension $m-1$.

For every affine basis $C \subset J$ and every point $j \in J \setminus C$, there is a linear equation on the ω's that expresses the fact that j is lifted to lie in the hyperplane containing the lift of C, namely the determinant of the lifted points. Since C is an affine basis, this equation has non-zero coefficient on, at least, $\omega(j)$ and, in particular it is not identically zero. Hence, it holds only on a hyperplane in $\mathbb{R}^{m+1}$. If ω does not lie in any of the hyperplanes obtained for the different choices of C and j, then ω must produce a triangulation: no lower facet, hence no lower face, projects to a dependent set.

(2) This follows from the fact that the lower faces of the lifted face $(\mathbf{A}_F)^{\omega|_F} = (\mathbf{A}^\omega)_F$ are just the lower faces of $\mathbf{A}^\omega$ contained in F.

(3) Recall that the determinant of $m+2$ points in $\mathbb{R}^{m+1}$ is zero if the points lie in a hyperplane and non-zero if they are independent. In this case the sign gives what is the relative position of each point with respect to the hyperplane determined by the other $m+1$. Also, recall that the determinant is continuous with respect to the coordinates of the points.

Let $\mathbf{A}^\omega$ be the lifted point configuration for the original height function ω. By continuity of the determinant, for each independent subconfiguration C of $m+2$ points in $\mathbf{A}^\omega$ there is an $\varepsilon_C > 0$ such that the determinant does not change sign under an ε_C-close perturbation of C. We take ε to be the minimum of all the ε_C's for the different choices of C.

Let now ω' be an ε-close perturbation of ω and let F be a maximal cell in $\mathscr{S}(\mathbf{A}, \omega')$. Maximality means that $\mathbf{A}_F$ is the projection of a facet $\mathbf{A}_F^{\omega'}$ of $\mathbf{A}^{\omega'}$, because every lower face is a face of a lower facet. In particular, the hyperplane $\mathbf{H}'$ containing $\mathbf{A}_F^{\omega'}$ is not vertical and leaves the rest of $\mathbf{A}^{\omega'}$ above it. Let now C be a maximal independent subset in $\mathbf{A}_F^{\omega'}$, with $m+1$ points.

Consider the corresponding set C^ω in $\mathbf{A}$ and the hyperplane $\mathbf{H}$ spanned by it. Our choice of ε implies that no point of $\mathbf{A}^\omega$ is below $\mathbf{H}$ and that the points that are above $\mathbf{H}$ have their corresponding points in $\mathbf{A}^{\omega'}$ above $\mathbf{H}'$. That is, $\mathbf{H}$ is the supporting hyperplane of a lower facet of $\mathbf{A}^\omega$ and the projection of this facet contains (as a subconfiguration, as well as geometrically) the projection of F'. $\square$

Using regular subdivisions we can easily prove that a (regular or not) subdivision that cannot be refined further must be a triangulation. We need the following construction:

Lemma 2.3.16. *Let $\mathscr{S}$ be a polyhedral subdivision of $\mathbf{A}$. Let $\omega : J \to \mathbb{R}$ be a height vector. Then the following is a polyhedral subdivision of $\mathbf{A}$ that refines $\mathscr{S}$:*

$$\mathscr{S}_\omega := \cup_{C \in \mathscr{S}} \mathscr{S}(\mathbf{A}|_C, \omega|_C).$$

Moreover, if $\mathscr{S}$ is regular, i.e., $\mathscr{S}(\mathbf{A}, \omega_0)$ for some ω_0, then $\mathscr{S}_\omega$ is also regular and equals $\mathscr{S}(\mathbf{A}, \omega_0 + \varepsilon\omega)$ for any sufficiently small positive ε.

Recall that $\mathbf{A}|_C$ denotes C considered as a subconfiguration of $\mathbf{A}$. That is, $\mathscr{S}_\omega$ is obtained by refining each cell of $\mathscr{S}$ in the regular way given by the height vector ω.

Proof. Our first goal is to prove that $\mathscr{S}_\omega$ is a subdivision of $\mathbf{A}$. That it refines $\mathscr{S}$ is obvious.

Clearly, $\mathscr{S}_\omega$ satisfies (UP) since the cells in each $\mathscr{S}(\mathbf{A}|_C, \omega|_C)$ cover the convex hull of C, by the union property of $\mathscr{S}(\mathbf{A}|_C, \omega|_C)$. It also satisfies (CP), since each $\mathscr{S}(\mathbf{A}|_C, \omega|_C)$ does. So, we only need to prove (IP). To get a contradiction, let R_1 and R_2 be two different cells in $\mathscr{S}_\omega$ with $\mathbf{x} \in \operatorname{relint}(R_1) \cap \operatorname{relint}(R_2)$. By the intersection property of each $\mathscr{S}(\mathbf{A}|_C, \omega|_C)$, R_1 and R_2 come from two different regular subdivisions $\mathscr{S}(\mathbf{A}|_{C_1}, \omega|_{C_1})$ and $\mathscr{S}(\mathbf{A}|_{C_2}, \omega|_{C_2})$ of cells of $\mathscr{S}$. Then, $\mathbf{x} \in \operatorname{conv}(C_1) \cap \operatorname{conv}(C_2)$. Let F_1 and F_2 be the carriers of $\mathbf{x}$ in C_1 and C_2, so that $\mathbf{x} \in \operatorname{relint}(F_1) \cap \operatorname{relint}(F_2)$. By (CP) of $\mathscr{S}$, F_1 and F_2 are cells in $\mathscr{S}$. Then, by (IP), F_1 and F_2 are *the same* cell of $\mathscr{S}$. Let us call it simply F. Now, by Part (iv) of Lemma 2.3.4, the subdivision of F obtained via $\mathscr{S}(\mathbf{A}|_{C_2}, \omega|_{C_2})$ and $\mathscr{S}(\mathbf{A}|_{C_1}, \omega|_{C_1})$ are the same, namely $\mathscr{S}(\mathbf{A}|_C, \omega|_C)$. That is, R_1 and R_2 are both cells of $\mathscr{S}(\mathbf{A}|_C, \omega|_C)$, which contradicts the fact that they are different cells with non-empty common relative interiors.

We now prove regularity. By Lemma 2.3.15, $\mathscr{S}(\mathbf{A}, \omega_0 + \varepsilon\omega)$ refines $\mathscr{S}$ if ε is sufficiently small. On the other hand, since ω_0 is a linear height vector on each cell $C \in \mathscr{S}$, the height vectors $\omega_0 + \varepsilon\omega$, $\varepsilon\omega$ and ω produce the same regular subdivision of that cell. Hence $\mathscr{S}_\omega = \mathscr{S}(\mathbf{A}, \omega_0 + \varepsilon\omega)$, as stated. $\square$

Definition 2.3.17. The polyhedral subdivision $\mathscr{S}_\omega$ of the previous lemma is called the *regular refinement* of $\mathscr{S}$ for the height vector ω.

Observe that the regular refinement of a non-regular subdivision may be a non-regular subdivision itself. See three examples of regular refinements in Figure 2.46.

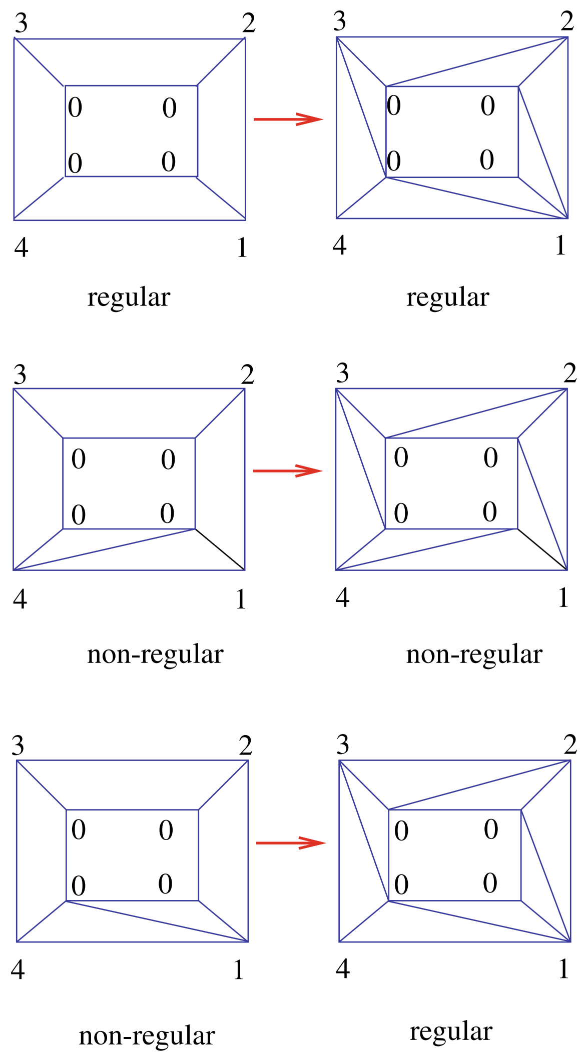

Figure 2.46: Subdivisions of a point configuration with heights and their regular refinements.

Corollary 2.3.18. *Every polyhedral subdivision of* $\mathbf{A}$ *can be refined to a triangulation. Moreover, every regular subdivision of* $\mathbf{A}$ *can be refined to a regular triangulation.*

Proof. For the first sentence, observe that if ω is sufficiently generic, then $\mathscr{S}_\omega$ is a triangulation since each $\mathscr{S}(\mathbf{A}|_C,\omega)$ is a triangulation. The second sentence follows from the last part of Lemma 2.3.16. $\square$

We now address the following question: if we are given a polyhedral subdivision $\mathscr{S}$ of $\mathbf{A}$ and a height vector ω, what is an easy way to check if $\mathscr{S} = \mathscr{S}(\mathbf{A},\omega)$?

From the definition of regular subdivision we can certainly derive an algorithm to answer this question. Namely: $\mathscr{S} = \mathscr{S}(\mathbf{A}, \omega)$ if and only if every cell $C \in \mathscr{S}$ is lifted to lie in a hyperplane $\mathbf{H}_C$ and every $j \notin C$ is lifted above that hyperplane. The decision question can be then phrased as the solvability of a system of linear inequalities (more about this later). What we want to show here is that this second part needs only be checked for some of the points $j \notin C$ determined by the cell adjacency (thus, giving a smaller system of inequalities).

Definition 2.3.19. A *wall* in a polyhedral subdivision of a point configuration $\mathbf{A}$ is a cell C of codimension one that is a face of two maximal cells (equivalently, that does not lie in a facet of $\mathbf{A}$). We say that C *separates* those two cells.

Theorem 2.3.20. *Let $\mathscr{S}$ be a polyhedral subdivision of a point configuration $\mathbf{A} \in \mathbb{R}^m$, and let $\omega : J \to \mathbb{R}$ be a height function. Then, one has $\mathscr{S} = \mathscr{S}(\mathbf{A}, \omega)$ if and only if*

(i) For every full-dimensional cell $C \in \mathscr{S}$, the lifted subconfiguration $\mathbf{A}^{\omega}|_C$ lies in a hyperplane in $\mathbb{R}^{m+1}$ (the coplanarity condition*).*

(ii) For every wall $C_0 \in \mathscr{S}$, with incident full-dimensional cells C_1 and C_2, all points in $\mathbf{A}^{\omega}_{C_1 \setminus C_2}$ lie above the hyperplane containing $\mathbf{A}^{\omega}|_{C_2}$ and vice versa (the local folding condition*).*

The coplanarity condition and the folding condition will play a crucial role in Chapter 5. Before going into the proof, observe that the "vice versa" in the second part is not an additional condition to be checked, but is equivalent to the stated condition. Also note that to verify the condition it is enough to check it for a single element of $C_1 \setminus C_2$. In fact, checking condition (ii) is equivalent to the computation of a certain $(d+2) \times (d+2)$ determinant. Condition (i), in turn, is equivalent to the vector $\omega|_C$ lying in the row span of the matrix $\mathbf{A}|_C$.

Proof. The "only-if"-direction is clear because we have taken a subset of the regularity conditions, and our result is implied by the definition of regularity.

Assume now that $\mathscr{S}$ is some polyhedral subdivision of $\mathbf{A}$ and $\omega : J \to \mathbb{R}$ is a height function satisfying the assumptions (i) and (ii).

We have to show that the remaining conditions are implied for ω: for every full-dimensional cell $C \in \mathscr{S}$ and every element $j \in J \setminus C$, point j is lifted above the hyperplane containing the lifted cell C^{ω}.

Let $\mathbf{p}_j$ be an arbitrary point and let $C \in \mathscr{S}$ be an arbitrary cell not containing j. Choose a point $\mathbf{x}$ in general position in $\operatorname{conv}(C)$ and consider the straight line segment ℓ from $\mathbf{p}_j$ to $\mathbf{x}$. This line segment intersects a sequence of full-dimensional cells $C_0, C_1, \dots, C_k = C$ with $C_i \in \mathscr{S}$ for $i = 0, 1, \dots, k$. Let $\mathbf{H}_i$ be the hyperplane spanned by the lifted points in C_i for $i = 0, 1, \dots, k$.

If $k = 0$ then $j \in C$, contradicting the choice of C. Therefore, $k > 0$ and $j \in C_0 \setminus C_1$. We will now prove that for all $i = 1, \dots, k$ $\mathbf{p}_j$ is lifted above $\mathbf{H}_i$.

For $i = 1$ the claim is literally Assumption (ii) in the Theorem. Let therefore $i > 1$, and let F_i be the wall between C_i and C_{i-1}. This wall is, by

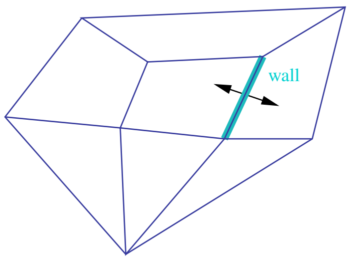

Figure 2.47: A wall in a two-dimensional point configuration.

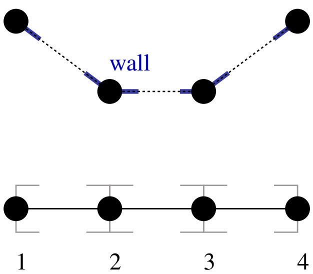

Figure 2.48: A wall in a one-dimensional point configuration.

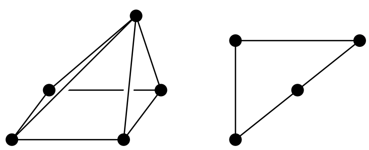

Figure 2.49: Two corank-one configurations.

construction, pierced by the segment ℓ. Therefore $\mathbf{p}_j$ lies on the same side of the vertical hyperplane $\mathbf{H}_{F_i}$ spanned by F_i as the points in $C_{i-1} \setminus C_i$.

Assumption (ii) for C_i and C_{i-1} says that the points in $C_{i-1} \setminus C_i$ are lifted above the hyperplane $\mathbf{H}_i$. In particular, the whole part of the hyperplane $\mathbf{H}_{i-1}$ that is on the same side of $\mathbf{H}_F$ as $C_{i-1} \setminus C_i$ is above $\mathbf{H}_i$. Moreover, by induction, $\mathbf{p}_j$ is lifted above $\mathbf{H}_{i-1}$.

Putting everything together yields: $\mathbf{p}_j$ is lifted above $\mathbf{H}_{i-1}$ by induction, which in turn is lifted higher than $\mathbf{H}_i$ on the side of $\mathbf{p}_j$ by Assumption (ii), which proves that $\mathbf{p}_j$ is lifted above $\mathbf{H}_i$ as well.

Setting $i = k$ completes the proof, since $C_k = C$. □

2.4 Flips and the graph of triangulations

Flips between triangulations are a central topic in this book. Flips are local changes that transform one triangulation into another. Reasons for introducing and understanding them come both from applications (they are a computationally simple way of searching for particular triangulations, or enumerating them) and theory (they highlight the rich structure to the set of all triangulations of a configuration). In Chapter 4, we will see how to technically deal with flips. Here, we simply introduce the concept. The core objects are configurations of corank one.

2.4.1 Corank-one configurations and circuits

Recall that the *corank* of a d-dimensional point configuration with n points is the number $n-d-1$. Independent configurations are exactly the configurations of corank zero. Let $\mathbf{A}$ be a configuration of corank one. Clearly, all non-trivial subdivisions of $\mathbf{A}$ are triangulations, because every full-dimensional proper subconfiguration has corank zero and so is independent. Here we prove that there are exactly two such triangulations, and characterize them.

A configuration has *corank one* if and only if it has a unique affine dependence relation $\sum_{j\in J} \lambda_j \mathbf{p}_j = 0$, with $\sum_{j\in J} \lambda_j = 0$ (uniqueness is, of course, up to multiplication of all λ's by the same constant). This affine dependence divides J into three subsets

$$J_+ := \{\, j \in J \,:\, \lambda_j > 0 \,\}, J_0 := \{\, j \in J \,:\, \lambda_j = 0 \,\}, J_- := \{\, j \in J \,:\, \lambda_j < 0 \,\}.$$

J_+ and J_- are the only disjoint subsets of J with the property that their relative interiors intersect. They intersect at the point

$$\sum_{j\in J_+} \lambda_j \mathbf{p}_j = \sum_{j\in J_-} |\lambda_j| \mathbf{p}_j,$$

where the λ's are assumed to be normalized so that

$$\sum_{j\in J_+} \lambda_j = \sum_{j\in J_-} |\lambda_j| = 1.$$

The set $J_+ \cup J_-$ containing the "relevant" part of the unique affine dependence is called a *circuit* in J. The pair (J_+, J_-) is classically called the

Radon partition of **A**, or the *oriented circuit* of **A**. Here is the formal definition, but much more will be said about circuits in Chapter 4:

Definition 2.4.1. Let **A** be a point configuration, with index set J. A subset of J is called a *circuit*, Z, if it is a minimal dependent set (that is, it is dependent but every proper subset is independent).

The partition (Z_+, Z_-) of Z into two parts such that $\operatorname{conv}(Z_+) \cap \operatorname{conv}(Z_-)$ is non-empty (which exists and is unique except for the swap of Z_+ and Z_-) is called an *oriented circuit*, or *signed circuit*. We say that the circuit is *of type* $(|J_+|, |J_-|)$.

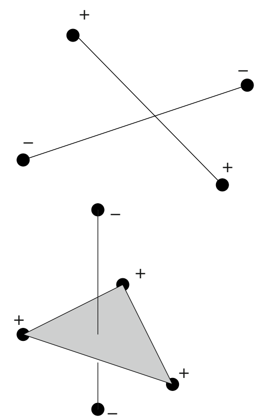

Figure 2.50: Two circuits represented as Radon partitions.

Since oriented circuits play a more prominent role than unoriented ones in this book, we will typically abuse language and drop the word "oriented" when referring to them. The underlying unoriented circuit Z will be called the *support* of (Z_+, Z_-). Circuits are studied in a broader context in Section 4.1.

Candidates for full-dimensional simplices in a triangulation of **A** are of the form $J \setminus \{j\}$ with $j \in J_+ \cup J_-$, since $j \in J_0$ means that $J \setminus \{j\}$ still has corank one.

Lemma 2.4.2. *Let* **A** *be a configuration of corank one and let* $J = J_+ \cup J_0 \cup J_-$ *be its label set, partitioned by the unique Radon partition of* **A**. *Then the following are the only two triangulations of* **A**:

$$\mathscr{T}_+ = \{C \subset J : J_+ \not\subseteq C\}, \quad \textit{and} \quad \mathscr{T}_- = \{C \subset J : J_- \not\subseteq C\}.$$

The two triangulations are regular.

Remark 2.4.3. The formulas given in this statement for $\mathscr{T}_+$ and $\mathscr{T}_-$ are equivalent to saying that their sets of maximal simplices are, respectively,

$$\{J \setminus \{j\} : j \in J_+\}, \quad \text{and} \quad \{J \setminus \{j\} : j \in J_-\}.$$

These are actually the formulas we prove.

Proof. Clearly, no triangulation can simultaneously contain J_+ and J_- since their relative interiors intersect. That is, every triangulation is either contained in $\mathscr{T}_+$ or in $\mathscr{T}_-$. It then suffices to show that $\mathscr{T}_+$ and $\mathscr{T}_-$ are indeed regular triangulations. For this, let $\omega : J \to \mathbb{R}$ be a height function and consider the quantity

$$\sum_{j \in J} \lambda_j \omega(j), \tag{2.10}$$

where the λ_j are the real coefficients of the unique (because of corank-one) dependence. If (2.10) equals zero, then the lifted point set lies in a hyperplane and $\mathscr{S}(\mathbf{A}, \omega)$ is the trivial subdivision. If (2.10) is not zero, then the lifted configuration is independent, because every affine dependence in the lifted configuration should be a dependence in **A** too. Hence, $\mathscr{S}(\mathbf{A}, \omega)$ is a triangulation. Which triangulation we obtain is governed only by the sign of $\sum_{j \in J} \lambda_j \omega(j)$. More precisely, if it is positive, then each point in J_+

is lifted above the hyperplane spanned by the lift of the other points, which implies that the set of lower facets of the lifted point set is

$$\{J \setminus \{j\} : j \in J_+\}.$$

Similarly, if it is negative then the lower facets are

$$\{J \setminus \{j\} : j \in J_-\}.$$

□

Observe that all we have said is valid for the case where the corank one configuration J has a repeated point. In this case, $\mathbf{A}$ is an independent set together with a second copy of one of its points, say $\mathbf{p}_j$. Then the affine dependence is $\mathbf{p}_j - \mathbf{p}_j = 0$ and J_+ and J_- have a single element each, the labels j and j' of the two copies of $\mathbf{p}_j$. In accordance with Lemma 2.4.2, $\mathbf{A}$ has two triangulations, each with a unique maximal simplex $J \setminus \{j\}$ or $J \setminus \{j'\}$.

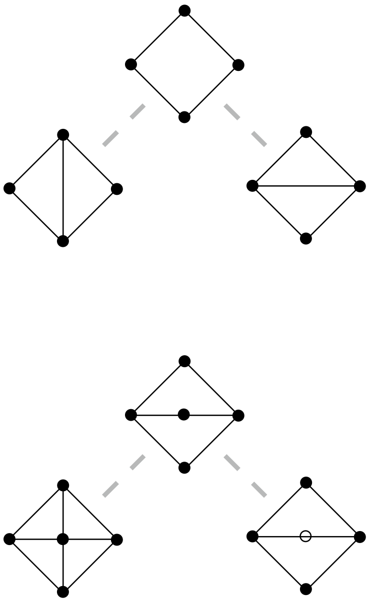

Figure 2.51: An almost-triangulation only has two triangulations that refine it.

2.4.2 Almost-triangulations and flips

The poset of subdivisions gives a way to quantify how close to being a triangulation a subdivision is: for $\mathscr{S} \in \mathrm{Subdivs}(\mathbf{A})$, let us call *height* of $\mathscr{S}$ the maximal length of chains of proper refinements of $\mathscr{S}$. Triangulations are the subdivisions at height zero.

Definition 2.4.4. A subdivision $\mathscr{S}$ is *almost a triangulation* (or an *almost-triangulation*) if it is at height one in the poset of subdivisions. That is, it is not a triangulation but all its proper refinements are triangulations.

As a first example, see Figure 2.51. We saw in Section 2.4.1 that the trivial subdivision of a corank-one configuration is an almost-triangulation: it has only two proper refinements, which are both triangulations. In what follows we show that this example is essentially unique. Every almost-triangulation is a triangulation except for a small part of it, where it is the almost-triangulation of a corank-one subconfiguration.

For this, observe that if $\mathbf{A}$ has corank one and Z is its unique circuit, then for every $j \in J \setminus Z$ we have that $J \setminus \{j\}$ still has corank one, hence it must have dimension one less than $\mathbf{A}$, hence it is a face of it (the latter because there is a single point $\mathbf{p}_j$ out of the hyperplane containing $J \setminus \{j\}$). In particular, the two triangulations of $\mathbf{A}$ are the pyramids, with apex $\mathbf{p}_j$, over the two triangulations of $\mathbf{A} \setminus \{\mathbf{p}_j\}$ (a pyramid means that the face is constructed by the taking the convex hull of each of the triangulation of $\mathbf{A} \setminus \{\mathbf{p}_j\}$ and the point $\mathbf{p}_j$). This is the only possibility to build full-dimensional cells. This fact follows also from the description of these triangulations given in Lemma 2.4.2. Together with the existence and properties of regular refinements, it is at the heart of the following characterization of almost-triangulations:

Lemma 2.4.5. *Let $\mathscr{S}$ be a polyhedral subdivision that is not a triangulation. $\mathscr{S}$ is almost a triangulation if and only if:*

(i) All its cells have corank at most one, and

(ii) All its cells of corank one contain the same circuit.

Proof. We start by proving necessity of the conditions: Condition (i) is easy with the ideas in the proof of Lemma 2.3.16. Suppose a cell $C \in \mathscr{S}$ has corank $k \geq 2$ and let $j \in C$ be such that $C \setminus \{j\}$ has at least corank one. Let $\omega : J \to \mathbb{R}$ be the height function that is everywhere zero except at j, where $\omega(j) > 0$. Then, the refinement $\mathscr{S}_\omega$ of $\mathscr{S}$ is proper and is still not a triangulation, since it contains the cell $C \setminus \{j\}$. This means $\mathscr{S}$ is a not an almost-triangulation.

Similarly, if Condition (ii) fails this means there are two different minimal dependent cells C_1 and C_2, both of corank one. Let $j \in C_1 \setminus C_2$. Minimality of C_1 implies that $C_1 \setminus \{j\}$ is independent thus of corank zero, thus of the same dimension as C_1. In particular, it is not a facet of C_1, and hence $\mathscr{S}_\omega$, with ω exactly as before, again produces a proper refinement. Since $j \notin C_2$, this finer subdivision still has C_2 as a dependent cell, hence, it is not a triangulation.

Conversely, if Conditions (i) and (ii) hold, then $\mathscr{S}$ has a unique minimal dependent cell Z, which is the circuit contained in all the dependent cells. Every proper refinement of $\mathscr{S}$ will, in particular, refine Z to one of its two triangulations. Moreover, every proper refinement refines every dependent cell as a cone over this triangulation. In particular, every proper refinement is a triangulation, and is completely characterized by which of the two triangulations of Z it contains. $\square$

Corollary 2.4.6. *Every almost-triangulation has exactly two proper refinements, which are both triangulations.*

Proof. This follows from the "sufficiency" part of the proof of the previous lemma. Every proper refinement refines the unique circuit common to all dependent cells in one of the two possible ways, and the refinement of this circuit determines the refinement of every cell. $\square$

Corollary 2.4.6 suggests that we call the change from a triangulation $\mathscr{T}_1$ to another one $\mathscr{T}_2$ a *flip* if they are the two proper refinements of the same almost-triangulation $\mathscr{S}$. This is exactly what we do:

Definition 2.4.7 (Flip)**.** Two triangulations of the same point configuration are *connected by a flip supported on the almost triangulation $\mathscr{S}$* if they are the only two triangulations refining $\mathscr{S}$.

In the literature, flips are defined in a more constructive way. The equivalence of our definition to that one will be the content of Theorem 4.4.1 in Chapter 4. For the time being we show some flips in action.

Figure 2.52 illustrates some examples of flips in the plane. The first is the traditional *diagonal-edge flip* common in computational geometry. The second is the insertion or deletion of a point in the interior of a triangle. There are other flips, not shown in the picture: the insertion or deletion of a point in the interior of an edge and the exchange between two copies of a repeated point.

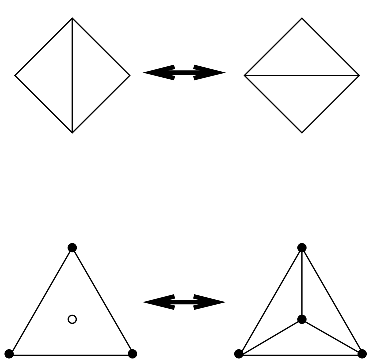

Figure 2.52: Some next-to-minimal subdivisions in the plane, each together with the two triangulations refining it.

The flip relation between triangulations of the same point configuration can be interpreted as an adjacency relation. The result of this is the graph of triangulations of a point configuration.

Definition 2.4.8. We call the *graph of triangulations* of $\mathbf{A}$ the graph whose nodes are all the triangulations of $\mathbf{A}$ and whose edges are flips between them. We denote it $\mathscr{G}_{\mathrm{tri}}\,\mathbf{A}$.

Example 2.4.9 (Example 2.3.6 continued)**.** The poset of subdivisions here is (isomorphic to) the Boolean lattice of all subsets of the label set J. The refinement poset is then the face lattice of a simplex with $|J|$ vertices. The graph of flips is the graph of this simplex, a complete graph.

Example 2.4.10 (Example 2.3.7 continued)**.** By the description in Example 2.3.7, the poset of subdivisions is isomorphic to the face lattice of a cube. The graph of triangulations is the graph of this cube.

Our final result in this section extends Corollary 2.3.18. There we saw that triangulations are the minimal elements in the refinement poset. Here we see that almost triangulations are the "next-to-minimal" elements.

Proposition 2.4.11. *Every polyhedral subdivision $\mathscr{S}$ other than a triangulation can be refined to an almost-triangulation. Moreover, if $\mathscr{S}$ is regular, then it can be refined to a regular almost-triangulation.*

Proof. We use the regular refinements of Lemma 2.3.16. For this, let B be a cell of $\mathscr{S}$ that is not a simplex (this exists because $\mathscr{S}$ is not a triangulation). Let C be a circuit contained in B. Consider a height vector ω that is sufficiently generic except at C; for example, let ω be zero on C and random in the rest of the elements of $\mathbf{A}$.

Then, $\mathscr{S}_\omega$ is an almost-triangulation since each $\mathscr{S}(\mathbf{A}|_D, \omega)$ is either a triangulation or an almost-triangulation with C as the unique circuit contained in some cell.

The sentence about regularity follows from the last part of Lemma 2.3.16. □

This result has a nice interpretation in the language of posets. A graph, as any simplicial complex, can be considered a poset. It has two levels, the bottom one consisting of the nodes of the graph and the top one consisting of the edges. Proposition 2.4.11 says that:

Corollary 2.4.12. *The graph of triangulations $\mathscr{G}_{\mathrm{tri}}\,\mathbf{A}$, as a poset, consists of the lowest two levels of the refinement poset* $\mathrm{Subdivs}(\mathbf{A})$.

Proof. This follows from the fact that edges in $\mathscr{G}_{\mathrm{tri}}\,\mathbf{A}$ correspond to flips between triangulations, and a flip between two triangulations corresponds to an almost-triangulation. □

2.5 Vector configurations and their triangulations

The reader should now be prepared for a final generalization in our framework: the study of *vector configurations* instead of *point configurations* as well as their triangulations and subdivisions. There are several immediate justifications for this: on the one hand, it is a harmless generalization; no extra complications are introduced, except perhaps for a slight need of "mental adjustment". On the other hand, it is a case that arises naturally in several contexts: for example, the normal fan of a polytope is a polyhedral subdivision of the vector configuration consisting of facet normals. Also, we need vector configurations if we want to apply the contraction operation (to be introduced in Section 4.2.4) to non-extremal points in a point configuration.

But the most profound justification for the study of vector configurations is the Gale duality that exists between a point configuration and its *Gale transform*, which is a vector configuration. The study of this duality will be started in Chapter 4, but it will be taken much further in Chapters 5 and 8. Figure 2.53 shows two examples of vector configurations one of them a point configuration as well.

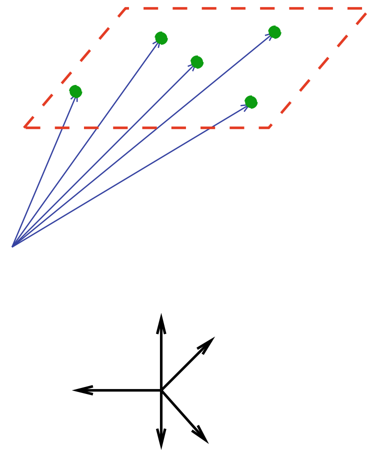

Figure 2.53: Miscellaneous vector configurations.

2.5.1 Vector configurations

A vector configuration is a finite collection of vectors in $\mathbb{R}^m$. We apply to vector configurations the same conventions that we used for point configurations: Repeated vectors are allowed, and distinguished by their labels, and we normally refer to vectors by their labels.

In fact, the homogeneous representation of a point configuration that we have been using so far is an example of a vector configuration. What we have called rank of a point configuration (its dimension plus one) is actually the rank of the linear space generated by the vectors. But there is no reason why we should not allow general, non-homogeneous vector configurations. One way to do this is to interpret what the concepts for a point configuration mean in a homogeneous vector configuration, and then apply them to non-homogeneous ones. For example:

Definition 2.5.1 (Vector configuration)**.** A *vector configuration* in $\mathbb{R}^m$ is a finite set $\mathbf{A} = (\mathbf{p}_j : j \in J)$ of labeled vectors $\mathbf{p}_j \in \mathbb{R}^m$. Its *rank* is its rank as a set of vectors. Its *corank* is $n - r$, where n is its number of elements and r is its rank. A *subconfiguration* is any (labeled) subset of it.

A vector (sub)configuration is *independent* if it does not have repeated vectors and its vectors are linearly independent. It is *dependent* otherwise. A vector configuration of rank r is *in general position* if each r-element subconfiguration is independent. Otherwise, it is *in special position*.

The *positive span*, or *conical hull*, of a subset $C \subseteq J$ of a vector configuration $\mathbf{A}$ with label set J is the following closed polyhedral cone.

$$\operatorname{cone}_{\mathbf{A}}(C) := \left\{ \sum_{j \in C} \lambda_j \mathbf{p}_j \, : \, \lambda_j \geq 0 \text{ for all } j \in C \right\}. \tag{2.11}$$

Its *relative interior* is

$$\operatorname{relint}_{\mathbf{A}}(C) := \left\{ \sum_{j \in C} \lambda_j \mathbf{p}_j \, : \, \lambda_j > 0 \text{ for all } j \in C \right\}. \tag{2.12}$$

For a linear functional $\psi \in (\mathbb{R}^m)^*$ with the property that $\psi(\mathbf{p}_j) \geq 0$ for every $j \in C$, the *face of C in direction ψ* is the following set of labels:

$$\operatorname{face}_{\mathbf{A}}(C, \psi) := \{\, j \in C \,:\, \psi(\mathbf{p}_j) = 0 \,\} \tag{2.13}$$

If F is a face of C, we write $F \leq C$. Moreover, if $F \neq C$ then we write $F < C$, and we say that F is a *proper face* of C. Observe that C is always a face of C, obtained when ψ is the zero functional. However, the empty set is *not always* a face in a vector configuration. See Remark 2.5.6 below. A *facet* of C is a face of rank one less than the rank of C. That is, it is a maximal proper face. The linear hyperplane $\{\, \mathbf{x} \in \mathbb{R}^m \,:\, \psi(\mathbf{x}) = 0 \,\} \subset \mathbb{R}^m$ is a *supporting hyperplane* of the face $\operatorname{face}_{\mathbf{A}}(C, \psi)$. An element $j \in C$ is *extremal* if $\{j\}$ is a face. A configuration is *in convex position* if all its elements are extremal.

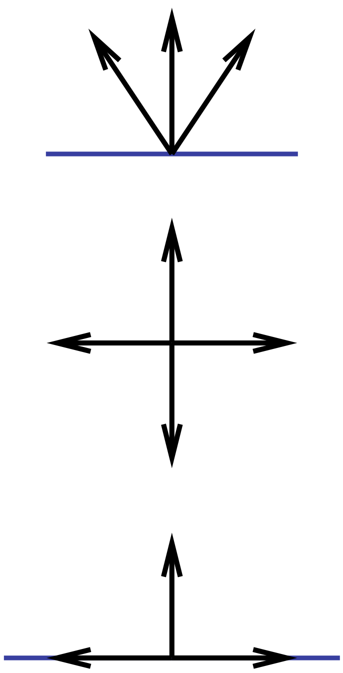

Figure 2.54: Three vector configurations: one acyclic, one totally cyclic, and one uninteresting (with three, four, and three vectors respectively).

Remark 2.5.2. $\operatorname{cone}\mathbf{A} = \dot\bigcup_{F \leq J} \operatorname{relint} F$, where $\dot\cup$ denotes "disjoint union".

Definition 2.5.3 (Carrier)**.** For a subconfiguration $S \subseteq \mathbf{A}$, the *carrier* of F in $\mathbf{A}$ is the smallest face of $\mathbf{A}$ containing S, i.e.,

$$\operatorname{carrier}_{\mathbf{A}}(S) := \bigcap_{S \subseteq F \leq J} F \tag{2.14}$$

Remark 2.5.4 (Lineality space). Let $\mathbf{L}$ be the maximal linear subspace contained in $\operatorname{cone}(J)$, for a configuration $\mathbf{A}$. $\mathbf{L}$ is usually called the *lineality space* of $\operatorname{cone}(J)$. The lineality space of a cone is contained in every face, and since it is maximal, it equals the intersection of all faces of $\operatorname{cone}(J)$. That is, it is the unique minimal face. For a vector configuration, its unique minimal face is the set of elements lying in the lineality space of its conical hull.

Here are two important, but more abstract, definitions:

Definition 2.5.5. A vector configuration $\mathbf{A}$ with index set J is *acyclic* if there is a linear functional that is positive in all the elements of the configuration. It is *totally cyclic* if $\operatorname{cone}_{\mathbf{A}}(J)$ is equal to the vector space spanned by $\mathbf{A}$.

In Section 2.1 we said we would *represent* point configurations as homogeneous matrices. In the same way, vector configurations will be represented by arbitrary matrices. In particular, a homogeneous matrix can be read both as a "homogeneous vector configuration" and as a point configuration. This is not a source of ambiguity; is simply that these two things are the same, for all the purposes of this book.

Furthermore, scaling the vectors of a configuration by positive scalars does not affect the face structure of the cone they span, or the set of subdivisions and triangulations of it (see next section). If a configuration is acyclic, by positive scaling we can make it homogeneous and thus equivalent to a point configuration. The bottom line is that acyclic vector configurations behave exactly as point configurations. Non-acyclic ones have a couple of strange new features. The first one is that they do not have an empty face:

Remark 2.5.6 (The empty face). In a point configuration, the empty set is considered a face because there is an affine functional that is positive at all points (for example, the constant functional $\psi(\mathbf{x}) = 1$). In vector configurations we only consider *linear* functionals and the existence of an always positive one is our definition of acyclic. That is, the empty set is a face of a vector configuration $\mathbf{A}$ only if $\mathbf{A}$ is acyclic. This agrees with the fact that the trivial subspace $\{0\}$ is a face of a cone only if the cone is *pointed* (that is, the positive span of an acyclic configuration). Observe that homogeneous vector configurations (also known as point configurations) are always acyclic.

Observe also that the zero vector may be an element in a vector configuration. In this case, the configuration cannot be independent, nor acyclic, nor in general position.

2.5.2 Polyhedral subdivisions of vector configurations

The definition of subdivision for a vector configuration is the same as for point configurations, except the role of the convex hull is now played by the positive hull (also called the positive span or the conic hull):

Definition 2.5.7 (Subdivision of a vector configuration)**.** A collection $\mathscr{S}$ of subconfigurations of a vector configuration $\mathbf{A}$ in $\mathbb{R}^m$ is a *polyhedral subdivision* of $\mathbf{A}$ if it satisfies the following conditions:

(CP) If $C \in \mathscr{S}$ and $F \leq C$ then $F \in \mathscr{S}$ as well. (Closure Property)

(UP) $\bigcup_{C\in\mathscr{S}} \operatorname{cone} C \supseteq \operatorname{cone} \mathbf{A}$. (Union Property)

(IP) $\operatorname{relint} C \cap \operatorname{relint} C' \neq \emptyset$ for $C, C' \in \mathscr{T}$ implies $C = C'$. (Intersection Property)

The elements of a polyhedral subdivision $\mathscr{S}$ are called *cells*. Cells of the same rank as $\mathbf{A}$ are *full-dimensional* or *maximal*. Cells of rank 1 are usually called *rays* of $\mathscr{T}$, but we will sometimes call them *vertices*, as if we were dealing with a point configuration. Independent cells are called *simplicial cells* or *simplices*.

Two subconfigurations *intersect properly* if they and all of their faces satisfy (IP); they intersect *improperly* otherwise.

A *triangulation* of $\mathbf{A}$ is a polyhedral subdivision all of whose cells are simplices.

A subdivision $\mathscr{S}$ *refines* another one $\mathscr{S}'$—symbolically: $\mathscr{S} \preceq \mathscr{S}'$—if for each $C \in \mathscr{S}$ there is a $C' \in \mathscr{S}'$ with $C \subseteq C'$. The poset of polyhedral subdivisions of a vector configuration $\mathbf{A}$ will be denoted $\operatorname{Subdivs}(\mathbf{A})$.

Multiplying a vector (or more) of a vector configuration by a *positive* constant does not change its combinatorics (faces, positive spans, relative interiors, etc). In particular, it preserves the collection of subdivisions and triangulations of it. Since every acyclic vector configuration can be homogenized by normalizing each vector by the value that a certain positive functional ψ takes on it, acyclic vector configurations do not introduce extra complications into our picture. But non-acyclic ones may. Let us see two examples:

Example 2.5.8 (Four vectors in general position in rank 2)**.** Consider the following four vectors in the plane

$$\begin{array}{cc} & \begin{array}{cccc} \mathit{1} & \mathit{2} & \mathit{3} & \mathit{4} \end{array} \\ \mathbf{A} = & \begin{pmatrix} 1 & -2 & 1 & 0 \\ 0 & 1 & -2 & 1 \end{pmatrix}. \end{array}$$

We encourage the reader to check that $\mathbf{A}$ has the eight subdivisions presented in Figure 2.55. We arranged them in a suggestive way to indicate the structure of the poset (without drawing the edges of a Hasse diagram) and, as usual, we only include the list of maximal cells in each. The top picture represents the trivial subdivision.

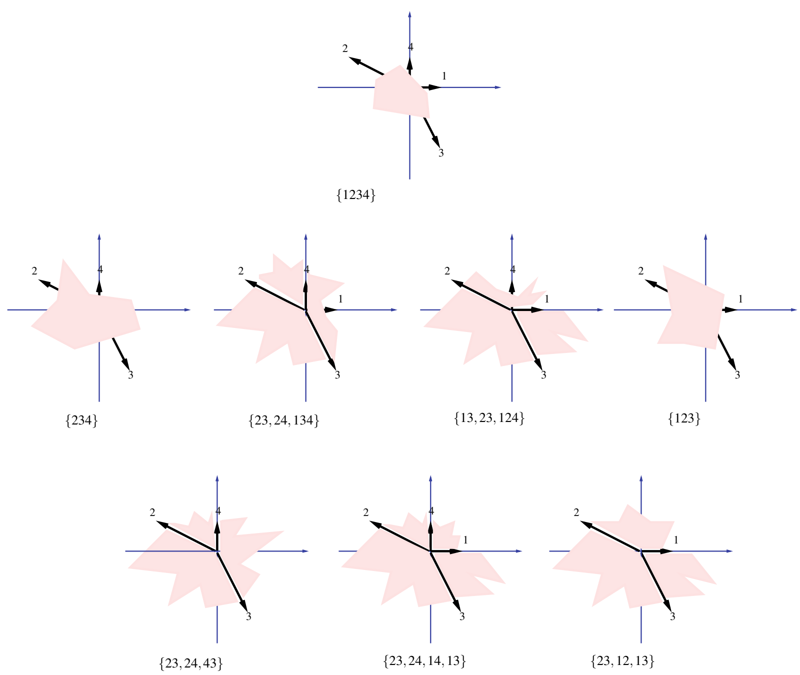

Figure 2.55: All subdivisions of four vectors in general (and non-acyclic) position in the plane.

Remark 2.5.9 (The empty set as a cell). We have said that the empty set is a face of a configuration only if the configuration is acyclic. Another feature of the empty set is that, by convention, its relative interior and hull form the zero cone $\{0\} \subseteq \mathbb{R}^m$ (a justification for this convention is that the result of an empty sum is zero).

In particular, if a subdivision contains a (non-empty, and necessarily not acyclic) cell with 0 in its relative interior, then the empty set is not a cell, or the intersection property would be violated. Conversely, if the empty set is not a cell, then the subdivision must contain a non-acyclic cell, then a unique minimal such one, in order for (UP) to hold: 0 is not in the relative interior of any acyclic set.

Summing up, a non-acyclic vector configuration has two types of subdivisions: those whose cells are all acyclic, which contain $\emptyset$ as a cell, and those whose cells are all not acyclic, which contain a unique minimal non-empty cell. In the previous example there are three and five respective subdivisions of each type.

Example 2.5.10 (Five vectors in general position in the plane)**.** We now let $\mathbf{A} \subset \mathbb{R}^2$ consist of five vectors of the same length, spaced equally (that is, pointing in the directions of the vertices of a regular pentagon). Each maximal simplex is a set of two vectors $\{i, j\}$, and a triangulation will just be a circular sequence of such pairs covering the circle of directions once. There are three possibilities, modulo the symmetries of the configuration:

1. The triangulation consisting of five cones $\{12, 23, 34, 45, 15\}$.
2. Five triangulations consisting of four cones, such as $\{13, 34, 45, 15\}$ and the ones obtained from it by symmetry.
3. Five triangulations consisting of three cones, such as $\{13, 35, 15\}$ and the ones obtained from it by symmetry.

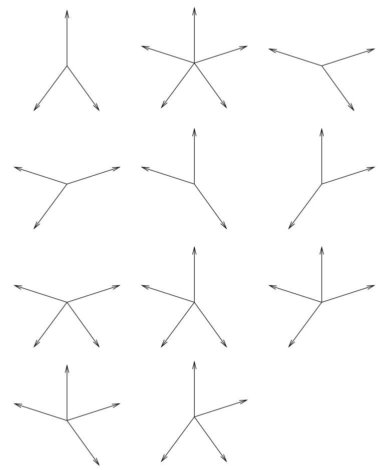

Figure 2.56: Triangulations of five vectors in general position in the plane.

There are six types of non-simplicial subdivisions. We encourage the reader to verify the definition for each, and check that the list is complete:

1. The trivial subdivision $\{12345\}$.
2. The non-trivial subdivisions with only one maximal cell, which come in two types: $\{1345\}$ and $\{135\}$ (observe that $\{123\}$, for example, is not a subdivision because it does not cover the whole positive span of $\mathbf{A}$, which is the whole plane).
3. The subdivisions whose cells are acyclic but still are not triangulations, which come in three types again: $\{123, 345, 15\}$, $\{13, 345, 15\}$ and $\{123, 34, 45, 15\}$.

2.5.3 Regular subdivisions of vector configurations

Again, for vector configurations we can take word for word the definition of regularity used for point configurations (see Definition 2.2.10). As an example, Figure 2.57 shows how the nine types of subdivisions of our last example arise as regular subdivisions. The lifting of vectors happens in $\mathbb{R}^3$.

As expected, the zero height vector (first picture in the figure) produces the trivial subdivision. Triangulations (bottom row) are produced by sufficiently generic height vectors.

There is, however, one substantial difference with the case of point configurations. What would have happened if we took negative height vectors? In the case of point configurations this causes no trouble, since changing all entries of a height vector by a constant does not change the combinatorics. In particular, it does not change what the lower faces of the lifted point configuration are. The lifted point configuration lies below the hyperplane at height zero but the lower faces remain unchanged.

In the case of a non-acyclic vector configuration, however, if the heights are negative then the lifted vector configuration *has no lower faces*! Indeed, a lower face is one for which there is a functional ψ with last coordinate positive and minimized at that face. In our "downwards lifted" vector configuration, however, such functionals do not achieve a minimum. For the time being let us be satisfied with just saying that *height vectors for vector configurations should better be taken with non-negative entries*, since this is sufficient (and essentially necessary) to obtain a regular subdivision from the height.

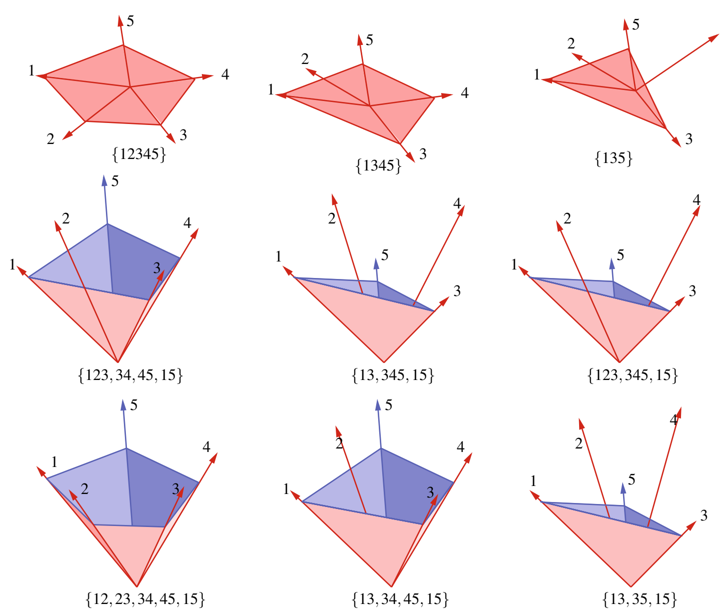

Figure 2.57: The subdivisions of Example 2.5.10, obtained as regular subdivisions. The three triangulations are in the bottom, the three subdivisions with a single cell are on top and the rest are in the middle.

Lemma 2.5.11. *Let* $\mathbf{A}$ *be a vector configuration and let* $\omega : J \to \mathbb{R}$*. Then the set* $\mathscr{S}(\mathbf{A},\omega)$ *of lower faces of* $\mathbf{A}^\omega$ *is a subdivision of* $\mathbf{A}$ *if and only if* ω *differs from a nonnegative height only by a linear function.*

Proof. Let ω be a height function that differs from a nonnegative height $\omega' \leq 0$ only by a linear function. Then the set of lower faces of $\mathbf{A}^\omega$ and $\mathbf{A}^{\omega'}$ coincide. We prove that the set of lower faces of $\mathbf{A}^{\omega'}$ is a subdivision of $\mathbf{P}$ by proving (CP), (IP), and (UP) for it. Conditions (CP) and (IP) again follow from elementary polyhedral geometry [339]; there is nothing new here. For nonnegative height functions, the proof of (UP) is completely analogous to the proof of (UP) for Theorem 2.3.11: For general heights ω and some $\mathbf{x} \in \operatorname{cone}(\mathbf{A})$, the intersection of $\mathbf{x} \times \mathbb{R}$ with $\operatorname{cone}(\mathbf{A}^\omega)$ (the *fiber over* $\mathbf{x}$) might not have a negative last coordinate, and such an $\mathbf{x}$ would then not be contained in any lower face of $\operatorname{cone}(\mathbf{A}^\omega)$. This may happen if $\mathbf{A}$ is totally cyclic, as discussed before. However, if the heights are all nonnegative, then all fibers have lowest elements, and the proof of (UP) can proceed as before.

It remains to show that every regular subdivision is induced by a nonnegative height function. For an arbitrary cell σ in such a regular subdivision, consider a supporting hyperplane of the corresponding lower face of cone($\mathbf{A}^{\omega}$). Moreover, let ω_σ be the unique linear height function that lifts $\mathbf{A}$ to lie in that hyperplane. Then, $\omega - \omega_\sigma$ is a nonnegative height function with the desired properties. ☐

In Section 4.1.3, Theorem 4.1.39, we will see why the restriction to nonnegative heights is no loss of generality as far as the set of all regular subdivisions is concerned. A much more detailed look at the structure of the *space of height vectors* can be found in Section 5.4.1.

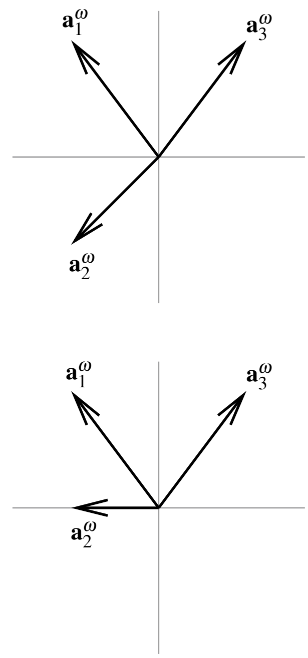

Figure 2.58: The one-dimensional vector configuration $(-1,-1,1)$ and two of its liftings. When $\omega = (2,-1,2)$ it has negative entries but it still defines a regular subdivision because $\omega + (1,1,-1) = (3,0,1)$.

2.6 Triangulations as simplicial complexes

Triangulations of a point or vector configuration are, among other things, simplicial complexes. In this section, we review several basic notions regarding simplicial complexes and how they behave in triangulations of configurations. Most specifically, we are interested in the numbers of faces of various dimensions. A good reference for the contents of this section is [339].

2.6.1 Simplicial complexes

An *abstract simplicial complex* $\mathscr{K}$ is a family of finite subsets of a label set J, such that, for every $F \in \mathscr{K}$, all subsets of F are also in $\mathscr{K}$. We will always assume that J is finite so that $\mathscr{K}$ is finite too. In parallel to the geometric objects we studied earlier, the elements of a simplicial complex $\mathscr{K}$ will be called *cells* or *simplices*. The dimension of a simplex is its cardinality minus one, and simplices of dimensions zero and one are called *vertices* and *edges* respectively. The maximal dimension among all cells is called the dimension of the complex, and the complex is called *pure* if all its maximal cells have the same dimension. In a pure complex of dimension d, the cells of dimensions d and $d-1$ are called, respectively, facets and ridges. For example, the boundary of a simplicial polytope forms a simplicial complex whose facets and ridges correspond to those of the polytope.

There are good topological and combinatorial reasons to study abstract simplicial complexes on their own. We already defined a *geometric simplicial complex* as a polyhedral complex, all of whose faces are simplices. It is worth remarking that any abstract simplicial complex, which is just a set of sets, can in fact be effectively achieved as a geometric simplicial complex. For this we need a map $J \to \mathbb{R}^m$ that sends the set of labels into vectors of Euclidean space with the following properties: the image of every $F \in \mathscr{K}$ is affinely independent (that is, it is the vertex set of a simplex); and different (geometric) simplices of the complex only intersect in common faces. Indeed, any map $J \to \mathbb{R}^{|J-1|}$ with an affinely independent image satisfies these. In particular, we can study every simplicial complex as if it is embedded geometrically. It is not true that every simplicial complex of a certain dimension d can be embedded in that same $\mathbb{R}^m$. The theory of embeddability of complexes is quite rich (see [227]).

Let $\mathscr{K}$ be a simplicial complex. For any $F \in \mathscr{K}$, the *star* of F in $\mathscr{K}$, $st_{\mathscr{K}}(F)$, is the subcomplex of $\mathscr{K}$ induced by all simplices of $\mathscr{K}$ that contain F as a face, plus all their faces (a word of caution: In the literature sometimes this is called the *closed star*). The *link* of F in $\mathscr{K}$ is the simplicial complex $\operatorname{link}_{\mathscr{K}}(F) = \{C \in st_{\mathscr{K}}(F) | F \cap C = \emptyset\}$. If $\mathscr{K}, \mathscr{L}$ are simplicial complexes, their *join* is $K * L = \{F \cup G | F \in K, G \in L\}$. For a geometric simplicial complex $\mathscr{K}$ we denote by $|\mathscr{K}|$ the underlying topological space, that is, the union of the geometric simplices. If $|\mathscr{K}|$ is homeomorphic to a ball or sphere we say that $\mathscr{K}$ is a *simplicial ball* or *simplicial sphere*, respectively. Every triangulation of a point configuration is a simplicial ball, while triangulations of a vector configuration can be simplicial balls or simplicial spheres, the latter if and only if the configuration is totally cyclic.

2.6.2 The f-vector of a simplicial complexes

In what follows the actual geometry is irrelevant. The notion of an *f-vector* of a simplicial complex has been central to the development of the combinatorial theory of polytopes [339]. The f-vector of an abstract simplicial complex $\mathscr{K}$ of dimension d is the vector

$$f(\mathscr{K}) = \big(f_{-1}(\mathscr{K}), f_0(\mathscr{K}), f_1(\mathscr{K}), f_2(\mathscr{K}), \dots, f_d(\mathscr{K})\big),$$

where $f_i(\mathscr{K})$ denotes the number of simplices of dimension i in $\mathscr{K}$. As a convention, in this definition $f_{-1}(\mathscr{K}) = 1$ for the empty set. The top entry in the f-vector, that is, $f_d(\mathscr{K})$, is sometimes called the *size* of $\mathscr{K}$.

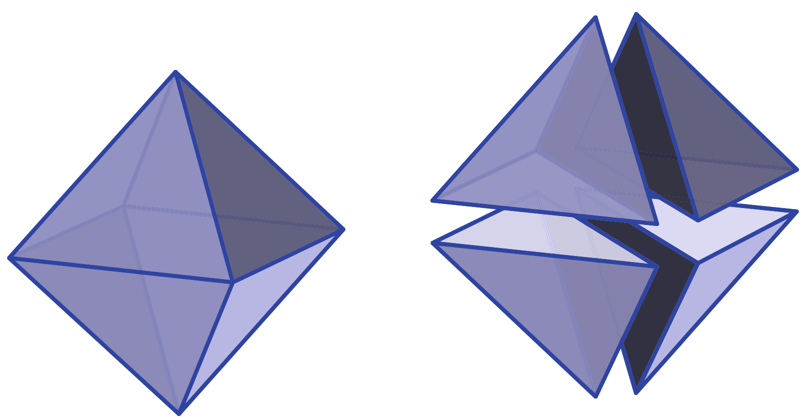

Figure 2.59: The f-vector of the boundary of the octahedron is $(1,6,12,8)$; the f-vector of the triangulated regular octahedron is $(1,6,13,12,4)$ because there are four tetrahedra adding four interior triangles and one interior edge; the number of vertices remains unchanged.

Two of the most important questions about f-vectors of simplicial complexes are what constraints different entries of a same f-vector must satisfy and how big can the individual entries of f-vectors be for a given dimension and number of vertices. We will be interested in these questions only for simplicial balls and spheres. We first look at the smallest possible size.

On a simplicial ball or sphere, the *adjacency graph* or *dual graph* of $\mathscr{K}$ is the graph whose nodes correspond to the maximal simplices in $\mathscr{K}$ and whose edges correspond to simplices intersecting in a common *ridge*, i.e, a simplex of dimension one less than the maximal simplex. Not much can be said in general about dual graphs of triangulations (see exercises), but they are useful to estimate a lower bound for the size of a triangulation:

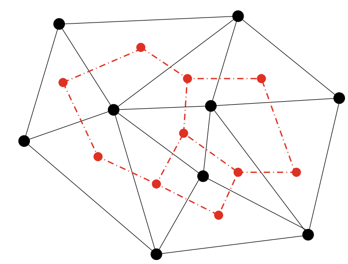

Figure 2.60: The dual graph of a triangulation appears in red.

Theorem 2.6.1 (Lower bound theorem for balls)**.** *The size of a simplicial d-ball with n vertices is at least $n-d$. Moreover, the equality is achieved precisely if the following (equivalent) conditions occur:*

1. *The dual graph of the ball is a tree.*
2. *Every $(d-2)$-cell of the ball lies in the boundary of* $\operatorname{conv}(\mathscr{K})$.

Proof. Observe that the dual graph of our ball $\mathscr{K}$ is a connected graph. In particular, it is possible to order its nodes (the full-dimensional simplices σ_1, σ_2, etc. of $\mathscr{K}$) in such a way that the first i of them form a connected subgraph, for every i.

We now imagine we are "building" $\mathscr{K}$ from scratch by adding the simplices one by one. With the first simplex, σ_1, we are inserting $d+1$ vertices at the same time, but any subsequent simplex will either use only vertices that were already there or insert a single new one (because every σ_i, $i>1$ shares at least one facet with one of the previous simplices). In particular, we need *at least* $n-d-1$ simplices other than σ_1 to insert all vertices, which proves the first sentence:

$$f_d(\mathscr{K}) \geq 1 + (n-d-1) = n-d.$$

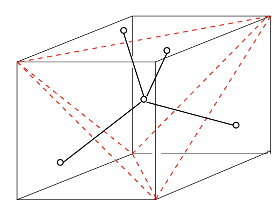

Figure 2.61: A triangulation of the 3-cube with 5 tetrahedra. Its dual graph is a tree.

Moreover, we get equality if and only if *every* simplex after σ_1 indeed inserts a new vertex, which is easily seen to be equivalent to not having cycles in the dual graph (if there is a cycle, the last simplex of the cycle that we insert has two facets in common with previous simplices, hence it introduces no new vertex).

We now prove that having a dual graph which is a tree is equivalent to excluding interior $(d-2)$-faces (interior edges in the case of three dimensions, for example). One direction is easy: if F is an interior $d-2$-cell, its link is itself a cycle, with nodes being the $(d-1)$-cells containing F and two consecutive nodes belonging to the same d-cell. This cycle is dual to a cycle in the adjacency graph of $\mathscr{K}$.

For the other direction (although the statement is valid for abstract simplicial complexes too), we will assume that our ball $\mathscr{K}$ is a triangulation of a point configuration. That is, that we have it geometrically realized in $\mathbb{R}^m$ and that $|\mathscr{K}|$ is a convex polytope. (If the reader is familiar with topological arguments, he or she can probably change the language in the proof to make it work for an abstract d-ball).

Assume that the dual graph of $\mathscr{K}$ contains a cycle. Let σ_0, $\sigma_1, \dots \sigma_k = \sigma_0$ denote the d-simplices forming the cycle, and let $\mathbf{p}_1, \dots, \mathbf{p}_k$ denote the barycenters of them. We consider the cycle geometrically embedded with the $\mathbf{p}_i$'s as nodes. Let $\mathbf{p}$ be a point that "moves" along the cycle, starting at $\mathbf{p}_1$ and ending at $\mathbf{p}_{k-1}$. We look at what happens to the segment $\mathbf{p}_0\mathbf{p}$

during this motion. At the first and last moments, the segment crosses two different $(d-1)$-faces of σ_0. Since at every moment in time $\mathbf{p}$ is outside σ_0, there must be some instant when $\mathbf{p}_0\mathbf{p}$ intersects a $(d-2)$-face of σ_0. The intersection point must be in the interior of $|\mathscr{K}|$ by convexity, since both $\mathbf{p}_0$ and $\mathbf{p}$ are in the interior. $\square$

Remark 2.6.2. What is the largest size of a triangulation? The answer is essentially given by the *Upper Bound Theorem* for spheres, that we state below. This theorem was first proved by Stanley in 1980 [306], and a detailed proof can be found in [339]. The cyclic d-polytope with n vertices, denoted by $\mathbf{C}(n,d)$, which appears in the statement, is the convex hull of n arbitrary points taken from the moment curve $\{(t,t^2,t^3,\dots,t^d) : t \in \mathbb{R}\} \subset \mathbb{R}^d$. The proper faces of this polytope form a simplicial $(d-1)$-sphere that we still denote $\mathbf{C}(n,d)$. Cyclic polytopes are among the most important constructions in polytope theory. We will study them again in detail in Section 6.1 (and they make an appearance also in Section 3.6.1).

Theorem 2.6.3 (*Upper Bound Theorem*)**.** *For any simplicial d-sphere $\mathscr{K}$ with n vertices*

$$f_i(\mathscr{K}) \le f_i(\mathbf{C}(n,d+1)), \quad 0 \le i \le d.$$

In order to adapt this statement to balls, we need to define the *deletion* of a vertex v in a simplicial complex $\mathscr{K}$. We denote it $\mathscr{K} \setminus v$ and by it we mean the subcomplex of $\mathscr{K}$ consisting of the cells that do not contain v. A formula for it is

$$\mathscr{K} \setminus v = (\mathscr{K} \setminus \mathrm{st}_{\mathscr{K}}(v)) \cup \mathrm{link}_{\mathscr{K}}(v).$$

Lemma 2.6.4.

$$f_j(\mathscr{K}) = f_j(\mathscr{K} \setminus v) + f_{j-1}(\mathrm{link}_{\mathscr{K}}(F)),\ \textit{for}\ -1 \le j \le d-1.$$

Proof. Suppose C is a cell of $\mathscr{K}$ of dimension j. If $v \in C$ then C appears as a $(j-1)$-face $C \setminus \{v\}$ in $\mathrm{link}_{\mathscr{K}}(F)$. Otherwise, C is in $\mathscr{K} \setminus v$. In either case the presence of C is counted in exactly one summand of the formula. $\square$

Now we are ready to present the upper bound on the the size of triangulations. This result is known to be tight, and the bound is achieved by the cyclic polytopes.

Corollary 2.6.5. *The size of a simplicial d-ball with n vertices is bounded above by $f_d(\mathbf{C}(n+1,d+1)) - (d+1)$.*

That is, the largest size of a triangulation is asymptotically $O(n^{\lceil (d+1)/2 \rceil})$. Similar bounds hold for the i-th entry of the f-vector of a triangulation.

Proof. What we do is we embed the triangulation in question inside a simplicial d-sphere. For this, think of your d-ball as embedded in $\mathbb{R}^d$ (here we do not really need it to be *geometrically* embedded; a topological embedding is enough, and it obviously exists since $|\mathscr{K}|$ is a d-ball).

Now, think of the "point at infinity" in $\mathbb{R}^d$ as an extra $(n+1)$-th vertex of your simplicial complex, joined to every boundary cell of $\mathscr{K}$. If you feel more comfortable being able to visualize this process, project $\mathbb{R}^d$ stereographically to a d-sphere, as shown in Figure 2.62 so that the simplices of $\mathscr{K}$ become spherical simplices in the sphere. Then use the center of the projection (the "north pole" in the figure) as the new vertex of your complex, joining it to all the boundary of $\mathscr{K}$.

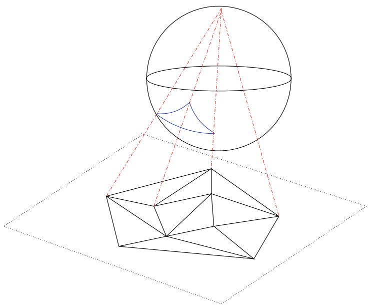

Figure 2.62: "Printing" a triangulation in the plane over the surface of the sphere.

In this way you get a simplicial d-sphere $\mathscr{K}'$ which contains $\mathscr{K}$ as a subcomplex, namely the subcomplex $\mathscr{K}' \setminus \{v\}$, where v is the new vertex.

In this situation, Lemma 2.6.4 reads

$$f_j(\mathscr{K}') = f_j(\mathscr{K}) + f_{j-1}(\partial\mathscr{K}), \qquad -1 \le j \le d.$$

Hence,

$$f_d(\mathscr{K}) = f_d(\mathscr{K}') - f_{d-1}(\partial\mathscr{K}) \le f_j(\mathbf{C}(n,d)) - (d+1),$$

as stated. In the last inequality we are using the upper bound theorem and the fact that $f_{d-1}(\partial\mathscr{K}) \ge d+1$ (since a simplicial $(d-1)$-sphere needs at least $d+1$ vertices). $\square$

Figure 2.63: Leonard Euler.

2.6.3 Linear constraints on the f-vector

We start this section by stating perhaps the most important formula involving the f-numbers, *Euler's formula*. Proofs of this result can be found in most sources in algebraic topology (see [237]). We will use Euler's formula heavily when we study the space of planar triangulations (in particular how they are connected).

Lemma 2.6.6. *For every simplicial d-ball $\mathscr{K}$:*

$$\sum_{j=0}^{d} (-1)^j f_j(\mathscr{K}) = 1.$$

For every simplicial d-sphere $\mathscr{K}$:

$$\sum_{j=0}^{d} (-1)^j f_j(\mathscr{K}) = 1 + (-1)^d.$$

The two versions of Euler's formula are easy to derive from one another. For example, if $\mathscr{K}'$ is a d-sphere, removing the interior of any particular d-simplex from it gives a a simplicial d-ball $\mathscr{K}$, with the same f-vector except for $f_d(\mathscr{K}) = f_d(\mathscr{K}') - 1$. For the converse, the reader can use the same trick (stereographic projection) as in the proof of Corollary 2.6.5.

Can one characterize the f-vectors of triangulations of convex (simplicial) polytopes or point configurations? It turns out one can indeed, but in doing so it is convenient to first translate the f-vector into a different form, called the *h-vector*. Formally speaking, the transformation of f-vectors into h-vectors is just a linear change of coordinate system in the space $\mathbb{R}^{d+2}$ where the f-vector $(f_{-1}, f_0, \ldots, f_d)$ lies. Some benefits of this change are:

(a) Complicated relations among entries of the f-vector become neat clean relations in terms of the entries of the h-vector,

(b) The f_j are nonnegative linear combination of the h_i's, therefore upper and lower bounds on the h_i imply upper and lower bounds on the f_j,

(c) The h-vector has a useful geometric meaning in terms of *shellings* of a simplicial complex [339] (we will come back to this in Section 9.5.2), and

(d) The h-vector has also an amazing algebraic meaning in terms of the Stanley-Reisner rings [306].

Definition 2.6.7. Let $\mathscr{K}$ be a $(d-1)$-dimensional pure simplicial complex with f-vector $\big(f_{-1}(\mathscr{K}), f_0(\mathscr{K}), f_1(\mathscr{K}), \ldots, f_{d-1}(\mathscr{K})\big)$. The *h-vector* of $\mathscr{K}$ is the vector

$$h(\mathscr{K}) := \big(h_{-1}(\mathscr{K}), h_0(\mathscr{K}), h_1(\mathscr{K}), \ldots, h_{d-1}(\mathscr{K})\big),$$

where:

$$h_k(\mathscr{K}) := \sum_{i=0}^{k} (-1)^{k-i} \binom{d-i}{k-i} f_{i-1}, \qquad 0 \le k \le d.$$

A compact way of representing the f and h vector is as the coefficients of two univariate polynomials of degree d, written in reverse order. Namely:

$$F_{\mathscr{K}}(t) = \sum_{i=0}^{d} f_{i-1} t^{d-i}, \qquad H_{\mathscr{K}}(t) = \sum_{i=0}^{d} h_i t^{d-i},$$

The reader can verify that in this notation the relation between the two is simply that:

$$H_{\mathscr{K}}(t) = F_{\mathscr{K}}(t-1),$$

from which

$$F_{\mathscr{K}}(t) = H_{\mathscr{K}}(t+1),$$

that is,

$$f_{k-1}(\mathscr{K}) := \sum_{i=0}^{k} \binom{d-i}{k-i} h_i, \qquad 0 \le k \le d.$$

In the example of a triangulated regular octahedron of Figure 2.59 the h-vectors are $h(\mathscr{T}) = (1,2,1,0,0)$ and $h(\partial\mathscr{T}) = (1,3,3,1)$. Observe also that $h_0 = f_{-1} = 1$. Euler's formula (Lemma 2.6.6) translates to the following simpler statement:

Corollary 2.6.8. *For every simplicial d-ball, $h_d = 0$, while for every simplicial d-sphere, $h_d = 1$.*

More generally, one has the following statements:

Lemma 2.6.9. *Suppose $\mathscr{K}$ is a $(d-1)$-simplicial complex. Then*

- $h_d(\mathcal{K}) = (-1)^d(1-\chi(\mathcal{K}))$ *where* $\chi(\mathcal{K})$ *is the Euler characteristic of the simplicial complex, i.e.,* $\chi(\mathcal{K}) = \sum_{j=0}^{d-1} f_j(\mathcal{K})$.
- $h_0(\mathcal{K}) = 1$.
- $h_1(\mathcal{K}) = f_0(\mathcal{K}) - d$.
- $f_{d-1}(\mathcal{K}) = \sum_{i=0}^{d} h_i(\mathcal{K})$.

It is easy to find other relations among the entries in the f-vector of a simplicial $(d-1)$-sphere. For example, since every $(d-1)$-cell has d facets and every $(d-2)$-cells is a facet of precisely two $(d-1)$-cells, one has that $df_{d-1} = 2f_{d-2}$. In a sense, this equation can be understood as Euler's formula summed over the links of every $(d-2)$-cell. Indeed, such links are 0-spheres, and Euler's formula for the 0-sphere is "every 0-sphere has two vertices". Similarly, the link of every $(d-2-k)$-cell is going to be a k-sphere, so we will get one Euler's formula for each dimension k. Perhaps surprisingly, when written in terms of the h-vectors, all these relations are equivalent to the following beautiful and symmetric formulas:

Lemma 2.6.10 (Dehn-Sommerville equations)**.** *The h-vector of any simplicial $(d-1)$-sphere $\mathcal{K}$ is symmetric, that is,*

$$h_i(\mathcal{K}) = h_{d-i}(\mathcal{K}) \qquad 0 \le i \le d.$$

We now show another simple relation between the h-vectors of a ball and of its boundary. A proof for shellable balls will be sketched in Section 9.5.2.

Theorem 2.6.11 (McMullen and Walkup (1971))**.** *For any simplicial $(d-1)$-ball $\mathcal{K}$ we have:*

$$h_i(\partial\mathcal{K}) = \sum_{k=0}^{i} h_k(\mathcal{K}) - h_{d-k}(\mathcal{K}), \quad 0 \le i \le d-1.$$

Proof. We use the Dehn-Sommerville equations with two spheres: the $(d-2)$-dimensional sphere $\partial\mathcal{K}$, and the $(d-1)$-sphere $\mathcal{K}'$ that we constructed for the proof of Corollary 2.6.5. The equation we settled there was

$$f_j(\mathcal{K}') = f_j(\mathcal{K}) + f_{j-1}(\partial\mathcal{K}), \qquad -1 \le j \le d.$$

The above linear equation carries over to give

$$h_i(\mathcal{K}') = h_i(\mathcal{K}) + h_{i-1}(\partial\mathcal{K}).$$

Hence, we have $h_i(\mathcal{K}) = h_i(\mathcal{K}') - h_{i-1}(\partial\mathcal{K})$, where it is understood that $h_k(\mathcal{K}) = 0$ for $k < 0$. Therefore $h_i(\mathcal{K}) = h_{d-i}(\mathcal{K}') - h_{i-1}(\partial\mathcal{K})$ by the Dehn-Sommerville relations applied to $\mathcal{K}'$. Thus, $h_i(\mathcal{K}) - h_{d-i}(\mathcal{K}) = h_i(\partial\mathcal{K}) - h_{i-1}(\partial\mathcal{K})$, applying again the previous equation and the Dehn-Sommerville equations to $\partial\mathcal{K}$. Adding enough copies of this expression yields the desired result. $\square$

Intuitively, the above theorem says that the h-vector (equivalently, the f-vector) of the boundary of a simplicial ball is determined by the f-vector of the ball. We see two curious applications below.

Corollary 2.6.12. *A simplicial $(d-1)$-ball $\mathscr{K}$ has no interior face of dimension $d-k-1$ if and only if $h_i(\mathscr{T}) = 0$ for all $i \geq k$.*

Proof. Notice that due to the containments of faces, when $\mathscr{K}$ has no interior face of dimension $d-k-1$, then $\mathscr{K}$ has no interior face of dimension less than that either. That is, $f_i(\mathscr{K})$ and $f_i(\partial\mathscr{K})$ are equal for all $0 \leq i \leq d-k$ and thus the same happens for the corresponding h-vector entries. On the other hand we saw in the proof of the McMullen-Walkup equations that $h_i(\mathscr{K}) - h_{d-i}(\mathscr{K}) = h_i(\partial\mathscr{K}) - h_{i-1}(\partial\mathscr{K})$. Hence $h_{d-i}(\mathscr{K}) = 0$ for $0 \leq i \leq d-k$. This means $h_i(\mathscr{K}) = 0$ for $k \leq i \leq d$. The converse implication is a reversal of the above arguments. $\square$

Exercises

Exercise 2.1. Let $\mathbf{A}$ be a point configuration with label set J and consider two height functions $\omega, \omega' : J \to \mathbb{R}^n$ such that $\omega' - \omega$ is the restriction to $\mathbf{A}$ of an affine map $\mathbb{R}^m \to \mathbb{R}$. Show that ω and ω' produce the same regular subdivision of $\mathbf{A}$. From this conclude that:

1. If $\mathscr{S}$ is a regular subdivision and C is an affinely independent subset of $\mathbf{A}$, then there is a height function that produces $\mathscr{S}$ and gives height zero to all points in C.

2. If ω is itself the restriction to $\mathbf{A}$ of an affine map, then it produces the trivial subdivision.

Exercise 2.2. How many distinct circuits are there in the point configuration of the vertices of a regular hexagon? What happens if the points are not on a circle?

Exercise 2.3. Show that for a 1-dimensional point configuration without repeated points all subdivisions are regular.

Exercise 2.4. Let $\mathbf{A}$ be a point configuration. Let C_1 be an affinely independent subset of points of $\mathbf{A}$. Show that there is a regular triangulation of $\mathbf{A}$ in which C_1 is used as a simplex.

Exercise 2.5. Let $\mathbf{A}$ be a point configuration. Let C_1 and C_2 be two subsets of points of $\mathbf{A}$. Suppose they are both affinely independent, and that they intersect properly. Show that there is a regular triangulation of $\mathbf{A}$ in which both are used as simplices.

Exercise 2.6. Let $\mathbf{A}$ be a point configuration. Let C_1, C_2 and C_3 be three subsets of points of $\mathbf{A}$. Suppose the three are affinely independent, and that they intersect properly to one another. Show by an example that there may not be any regular triangulation of $\mathbf{A}$ in which the three are used as simplices.

Exercise 2.7. Let $\mathbf{A}$ be a point configuration. Let C_1, C_2 and C_3 be three subsets of points of $\mathbf{A}$. Suppose the three are affinely independent, and that they intersect properly to one another. Show by an example that there may not be any triangulation, regular or not, of $\mathbf{A}$ in which the three are used as simplices.

(*Hint:* Now you need to go to dimension three.)

Exercise 2.8. Find all subdivisions of the vector configuration

$$\begin{array}{cc} & \begin{array}{cccc} \mathit{1} & \mathit{2} & \mathit{3} & \mathit{4} \end{array} \\ \mathbf{A} = & \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 1 & -1 & -1 \end{pmatrix}. \end{array}$$

Note that this example is neither totally cyclic nor acyclic.

Exercise 2.9. Let $\mathscr{K}$ be a simplicial complex and let F_1 and F_2 be two of its faces. Suppose that $F_1 \cap F_2 = \emptyset$ and that $F_1 \cup F_2$ is a face of $\mathscr{K}$. Prove that $\operatorname{link}(F_2, \operatorname{link}(F_1, \mathscr{K})) = \operatorname{link}(F_1 \cup F_2, \mathscr{K})$.

Exercise 2.10. Let $\mathscr{K}$ be a simplicial d-ball with F a boundary face of $\mathscr{K}$ such that it has dimension $k-1$ and $\operatorname{link}(F, \mathscr{K})$ is a simplicial $d-k$ ball. Then $\operatorname{link}(F, \partial\mathscr{K}) = \partial(\operatorname{link}(F, \mathscr{K}))$.

Exercise 2.11. Let $\mathscr{K}$ be the boundary complex of a simplicial d-polytope $\mathbf{P}$ and $\mathbf{v}$ a vertex of $\mathbf{P}$. Prove that $\partial(\mathscr{K} - \mathbf{v})$ equals $\operatorname{link}(\mathbf{v}, \mathscr{K})$.

Exercise 2.12. Give a formula for the f and h-vectors and polynomials associated to a d-simplex. Do the same for the boundary of a d-simplex (this is a sphere now).

Exercise 2.13. Verify a couple of the properties of Lemma 2.6.9.

Exercise 2.14. Consider a simplicial complex $\mathscr{K}$ which is the union or intersection of two other simplicial complexes. Can you write formulas for the h-vector of $\mathscr{K}$ in terms of the h-vectors of its parts?

Exercise 2.15. What are the conditions that a vector with integer coordinates must satisfy to be the f-vector of a 3-dimensional convex polytope?

Exercise 2.16. Compute the f-vectors of triangulations of a regular 3-cube. What is the dimension of the convex hull of the f-vectors? Do you notice anything interesting?

Exercise 2.17. Can you find an example of a positive integral vector that satisfies the McMullen-Walkup equations but is not the f-vector of a triangulation of a simplicial polytope?

Exercise 2.18. Let $\mathbf{A}$ be a point configuration of even dimension d, such that $\operatorname{conv}(\mathbf{A})$ is a simplicial polytope. Prove that, for every pair of triangulations of $\mathbf{A}$, $\mathscr{T}_1$ and $\mathscr{T}_2$, we have $f_d(\mathscr{T}_1) = f_d(\mathscr{T}_2)$ modulo 2. The parity of $f_d(\mathscr{T})$ is in fact determined by $\operatorname{conv}(\mathbf{A})$ (*Hint:* you can prove this using the McMullen-Walkup equations, but there is also a very elementary way of doing it. Think spheres!).

Exercise 2.19. If two triangulations $\mathscr{T}_1$ and $\mathscr{T}_2$ differ by a flip, what can you say about their f-vectors? In other words, what are the f-vectors associated to flips? Investigate especially the case when the points are in general position.

Exercise 2.20. Prove that the dimension of the linear subspace spanned by the f-vectors of all full triangulations of a three-dimensional point configuration is one.

http://www.springer.com/978-3-642-12970-4

Triangulations
Structures for Algorithms and Applications
De Loera, J.; Rambau, J.; Santos, F.
2010, XIII, 535 p., Hardcover
ISBN: 978-3-642-12970-4