Polytopal Digraphs

Topics in Enumeration and Reconstruction

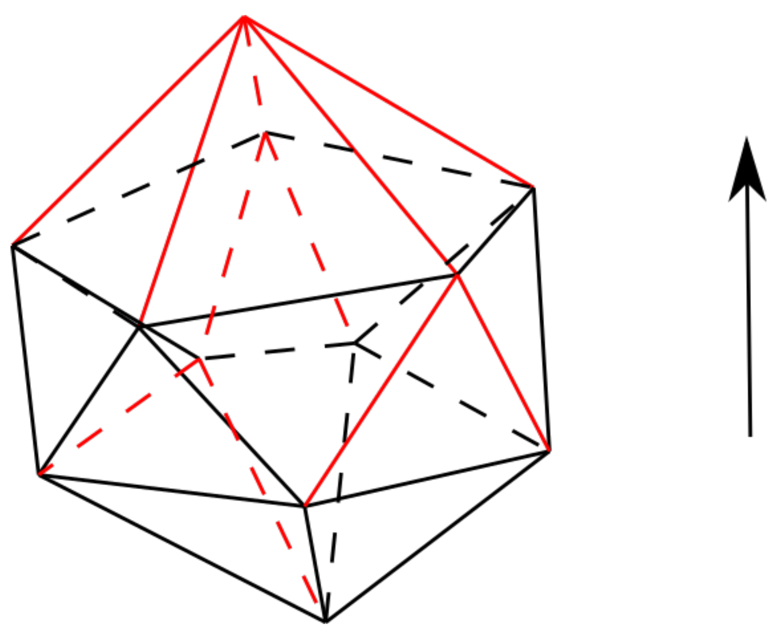

ΜΕΤΑΠΤΥΧΙΑΚΗ ΕΡΓΑΣΙΑ - ΤΜΗΜΑ ΜΑΘΗΜΑΤΙΚΩΝ, ΕΘΝΙΚΟ ΚΑΙ ΚΑΠΟΔΙΣΤΡΙΑΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

Γεώργιος Κοντογεωργίου

Αθήνα 2021

Επιβλέπων Καθηγητής: Χρήστος Αθανασιάδης

Τριμελής Επιτροπή
Χρήστος Αθανασιάδης
Δημήτριος Θηλυκός
Γιάννης Μούρτος

Contents

Overview and Preliminaries

This thesis was born out of the author's recent fascination with directed polytopes and polytopal digraphs, so it is only appropriate to start it by defining these objects:

Definition 1. *A directed polytope is a pair* (P, f) *where* P *is a polytope in* $\mathbb{R}^d$ *and* $f : \mathbb{R}^d \to \mathbb{R}$ *is a linear functional which is generic over the vertices of* P*.*

Definition 2. *The digraph* $\omega(P, f)$ *of a directed polytope* (P, f) *is obtained by applying the orientation induced by* f *on the graph of* P*, i.e. every edge of* $\omega(P, f)$ *is directed from the end with the lower* f*-value to the end with the higher* f*-value. If a digraph is isomorphic to the digraph of some directed polytope, then it is called polytopal.*

These pretty things have resided in the collective consciousness of mathematicians for three quarters of a century, emerging once in a while to provide crucial insight about undirected polytopes, as in the strengthening of Balinski's Theorem by Holt and Klee [11], or in Kalai's cunning proof of the Blind-Mani-Levitska Theorem [13]. And yet, although directed polytopes and polytopal digraphs have occasionally proven to be useful tools, they have rarely been the object of study. Rather, they have mostly been treated as a setting for the simplex algorithm.

The purpose of this thesis is triple.

- Firstly, to lay some foundations for this topic. This is a twofold matter.
  - For one, we need a language for directed polytopes. For example, what does it mean for a directed polytope to be "unique" in satisfying a property? We can only make such statements if we have a notion of equivalence over directed polytopes. We will provide such notions, along with a wider vocabulary of definitions.
  - For another, we wish to gather all the low-hanging fruit concerning directed polytopes and their digraphs.
- Secondly, to collect some of the mathematical work which has been done on directed polytopes in one document.

- Thirdly, to compare results concerning polytopes and their graphs to results concerning directed polytopes and their digraphs.

Two topics appear to be particularly relevant to directed polytopes, in the sense that some work can be done (and has been done) on them. One is reconstruction; determining polytopal digraphs, deducing whether there is a unique directed polytope associated with each of them, and providing a (preferably polynomial-time) algorithm for reconstructing said polytope. The other is enumeration; counting various structures which are encountered on directed polytopes. This thesis is split between these two topics.

This work is organised as follows.

The rest of the introduction sets the necessary foundations for working with directed polytopes. In 0.1, we settle our definitions and provide some relevant theorems and observations. One particularly important theorem is the one bearing Balinski's name; in 0.2, we present the proof of Balinski's Theorem provided by Holt and Klee. For reasons that we will explain, it makes for a great introductory result.

In the first half of the thesis, we will be preoccupied with reconstruction results: propositions which identify certain polytopal (di)graphs, answer whether their corresponding (directed) polytopes are unique, and provide an effective way to compute such (directed) polytopes.

In Chapter 1, we will focus on 3-polytopes, of both the directed and undirected varieties. Steinitz's Theorem is presented in 1.1, whereas its directed analogue, the Mihalisin-Klee Theorem, follows in 1.2.

Chapter 2 is devoted to simple polytopes. In 2.1, Kalai's Theorem easily resolves both the directed and undirected cases, proves the uniqueness of the reconstruction, and yields an exponential algorithm for performing it. In 2.2, we show how the reconstruction can be done in polynomial time instead, using Friedman's algorithm.

In Chapter 3, we conclude the reconstruction part by mentioning a few additional classes of reconstructible directed polytopes.

The second half of the thesis is about enumeration results: counting certain structures which appear in directed polytopes.

In Chapter 4, we review results from a recent paper ([3]) of Athanasiadis, De Loera, and Zhang, which contains several enumeration theorems related to arborescences and monotone paths on directed polytopes.

In Chapter 5, we tackle Conjecture 4.6 appearing in [3], which states that the maximum number of monotone paths in a simple directed 3-polytope on $2n$ vertices is $F_{n+2}+1$, where F_n is the n^{th} Fibonacci number. We prove that this is true provided that no directed 3-polytope (P, f) on $2n$ vertices with maximum number of monotone paths has a non-extremal triangle, i.e. a triangular face which contains neither the top nor the bottom vertex of (P, f).

Finally, in Chapter 6, we pose and partially answer the problem of constructing directed polytopes with a given set of different outdegrees. We show that, for every $d \geq 3$ and every finite sequence $0 < 1 < D_1 < ... < D_k$ of natural numbers with each D_i being large enough, there exists a directed d-polytope (P, f) such that the different outdegrees of the vertices of the digraph $\omega(P, f)$ are exactly the numbers of the given sequence. In particular, in every dimension $d \geq 3$ there exists a d-polytope with only three different outdegrees. We also discuss other relevant results.

0.1 Definitions and Known Results

We begin with the basic definitions. Note that, in this thesis, all graphs are simple and all polytopes are convex.

Graphs:

Given a (di)graph G, a *vertex deletion* consists of removing a vertex from $V(G)$ and all of its incident edges from $E(G)$. An *edge deletion* consists of removing a single edge from $E(G)$. A *vertex smoothing* consists of deleting a vertex of degree 2 and adding its neighbourhood to $E(G)$ (so that the outdegrees of the remaining vertices do not change).

A graph H is called a *subgraph* of a graph G if it can be obtained from G by applying vertex and edge deletions. In particular, it is called an *induced subgraph* of G if it can be obtained through vertex deletions alone, and a *spanning subgraph* of G if it can be obtained through edge deletions alone. Moreover, H is called a *topological minor* of G if it can be obtained from G by applying vertex and edge deletions and vertex smoothings.

A *graph isomorphism* between two graphs G and H is a function which maps the vertices of G bijectively to the vertices of H, and the edges of G bijectively to the edges of H, so that the ends of an edge are always mapped to the ends of its image. A *digraph isomorphism* between two digraphs G and H is a graph isomorphism between the underlying graphs which also preserves orientation. If there exists an isomorphism between to (di)graphs, then we call them *isomorphic*.

The *topological realisation* of a graph is the 1-complex whose 0-cells are bijectively labelled by the vertices of the graph, its 1-cells are bijectively labelled by the edges of the graph, and the two 0-cells which form the boundary of an 1-cell are labelled by the ends of the edge which labels the 1-cell. By an *embedding* of a graph H into a graph G we mean a topological embedding of the topological realisation of H into the topological realisation of G. A graph H is embeddable in a graph G if and only if it is a topological minor of G.

A (di)graph is *k-connected* if any deletion of $k-1$ or fewer vertices yields a connected sub(di)graph.

A digraph is *weakly connected* if there exists a directed path between any two of its vertices, and it is *strongly connected* if, for every directed pair of its vertices, there exists a directed path from the first vertex to the second vertex.

Polytopes:

A *(convex) d-polytope* is the convex hull of a finite set of points in $\mathbb{R}^d$.

An intersection of a d-polytope with any of its supporting hyperplanes is called a *face*. In particular, a face whose affine hull is a k-dimensional affine space is a k-face ($\emptyset$ is conventionally considered to have dimension -1).

The set of faces of a polytope P ordered by inclusion is called the *face lattice* of P (because it is a lattice), and is denoted $\mathcal{F}(P)$. A maximal chain of $\mathcal{F}(P)$ is called a *flag*. The number of k-faces of P is denoted f_k, and the vector $(f_{-1}, ..., f_d)$ is called the *f-vector* of P. We also define, for $0 \leq k \leq d$, $h_k := \Sigma_{i=0}^{k}(-1)^{k-i}\binom{d-i}{k-i}f_{i-1}$. The vector $(h_0, ..., h_d)$ is called the *h-vector* of P, and its relevance to our particular topic will be made clear in chapter 5.

The $(d-1)$-faces of a polytope are commonly called *facets*. The affine hull of a facet F divides $\mathbb{R}^d$ into two half-spaces. The points which lie in the half-space which does not intersect P are said to be *beyond* F, whereas the points which lie in the half-space which does intersect P are said to be *beneath* F.

The 0-faces of a polytope are commonly called *vertices*, and its 1-faces are commonly called *edges*. The *graph* $G(P)$ of a d-polytope P is the graph whose vertices are the vertices of the polytope and whose edges are the edges of the polytope. A d-polytope is called *simple* if its graph is d-regular.

A linear functional $f : \mathbb{R}^d \to \mathbb{R}$ is *generic* over a d-polytope P if, for every pair of vertices u, v of P, $f(u) \neq f(v)$.

A *directed polytope* is an ordered pair (P, f), where P is a polytope and f is a generic linear functional over P. Each directed polytope (P, f) determines a digraph $\omega(P, f)$ as follows: the underlying graph of $\omega(P, f)$ is the graph of P, and each edge $\{u, v\}$ with $f(u) < f(v)$ is oriented from u to v. If an orientation over the graph of a polytope can be obtained in this manner, it is called *LP-admissible*.

A *unique sink orientation* over a plane graph is an orientation in which every face has a unique sink, and the graph as a whole also has a unique sink.

It is a well-known folklore theorem (e.g. see []) that every LP-admissible orientation is an acyclic unique sink orientation. The sink of $\omega(P, f)$ is the *top vertex*. The sink of $\omega(P, -f)$, i.e. the source of $\omega(P, f)$, is the *bottom vertex*.

In general, it is convenient to use the notation v_k to denote the vertex of P with the k^{th} lowest f-value, and the notation d_k to mean the out-degree of v_k in $\omega(P, f)$. For example, if (P, f) has n vertices, we will write v_1 to denote its bottom vertex, and v_n to denote its top vertex.

Occasionally, we will refer to the vertices of (P, f) using terminology from order theory. By default, such terms should be understood to concern not the linear order of the vertices by f-value, but rather the order induced by $\omega(P, f)$. For instance, when we say that a vertex v is *less* than a vertex u, and write $v <_{(P,f)} u$, we do not simply mean that $f(v) < f(u)$, but also that there is a directed path from v to u in $\omega(P, f)$. When we only wish to communicate that $f(v) < f(u)$, we will say that v is *lower* than u.

A (di)graph is *d-polytopal* if it is isomorphic to the (di)graph of a (directed) d-polytope.

Reconstruction:

Now, we will define the notion of reconstruction, which will be our topic throughout the first part.

We say that a class of (directed) polytopes C is a *reconstruction class* up to an equivalence relation R, or, equivalently, that C is R-reconstructible, if every two (directed) polytopes in C with the same (di)graph are R-equivalent. Informally, C is a reconstruction class if, when given the graph G of a polytope P and the knowledge that $P \in C$, we can always determine P up to R. Here is an example:

Example 1. *The class of all polytopes is a reconstruction class up to quasi-isometry.*

This is not a particularly interesting example. With the possible exception of geometric group theorists, few mathematicians will gaze upon the multitude of polytopes and declare that they are all the same simply because they are bounded. Intuitively, an R-reconstructible class is interesting if R is some "natural" equivalence between (directed) polytopes. For instance, it comes very naturally to a combinatorialist to equate polytopes with the same combinatorial structure. The corresponding equivalence relation is well known:

Definition 3. *Two polytopes P, P' are combinatorially equivalent if their face lattices are order isomorphic.*

We will not make extensive use of any other equivalence relation between undirected polytopes: combinatorial equivalence is well-established as the most interesting one. Our challenge is rather to invent an equally natural equivalence relation for directed polytopes. Unfortunately, we can come up with two.

Definition 4. *Two directed polytopes (P, f), (P', f') are orientation equivalent if there exists an order isomorphism $\phi : \mathcal{F}(P) \to \mathcal{F}(P')$ between their face lattices which, when restricted to the digraph of (P, f), yields a digraph isomorphism.*

Definition 5. *Two directed polytopes (P, f), (P', f') are order equivalent if there exists an order isomorphism $\phi : \mathcal{F}(P) \to \mathcal{F}(P')$ between their face lattices such that, for any two vertices u, v of P, $f(u) < f(v) \Leftrightarrow f'(\phi(u)) < f'(\phi(v))$.*

We will work with both of these relations. The reader may have noticed that order equivalence implies orientation equivalence, since fixing the linear order of vertices by f-value also fixes the directions of the edges; Figure 1 shows that the reverse implication is false.

That said, the following theorem holds:

Theorem 1. *If a class of directed polytopes is a reconstruction class up to orientation equivalence and the digraphs of its directed polytopes are weakly connected, then it is also a reconstruction class up to order equivalence.*

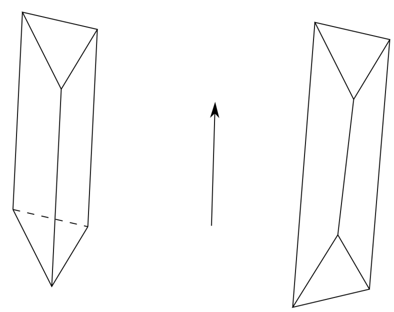

Figure 1: The depicted triangular prisms are orientation equivalent, but not order equivalent.

Proof. Let (P, f), (Q, g) be two orientation equivalent directed polytopes whose (isomorphic) digraphs are weakly connected. Since $\omega(P, f)$ is weakly connected, it induces a linear order in its vertices, with $u <_{(P,f)} v \Leftrightarrow f(u) < f(v)$. The same is true for $\omega(Q, g)$. Hence, an orientation equivalence between (P, f) and (Q, g) is also an order equivalence. □

The following question remains:

Question 1. *Are there digraphs which are not weakly connected yet determine a directed polytope up to order equivalence?*

It would be interesting to answer this question within various contexts, such as for 3-polytopes or simple polytopes.

Enumeration:

Finally, we define the directed polytopes which we shall use and the structures which we shall enumerate in the second part of the thesis.

A polytope is called *k-neighbourly* if every k of its vertices form a face. In particular, a polytope is called 2-*neighbourly* if its graph is complete. A $[\frac{d}{2}]$-neighbourly polytope is also simply called *neighbourly*. When discussing directed neighbourly polytopes, we will not provide specifications about their direction, since, evidently, this does not affect their directed graph, which is the unique directed, acyclic, complete graph on the required number of vertices.

A polytope is called *stacked* if it can be obtained from a simplex by repeatedly gluing other simplices of the same dimension along common facets, preserving convexity at each step. A useful class of directed stacked 3-polytopes are those produced by gluing, for each $k \in \{1, ..., n-4\}$, the k^{th} tetrahedron on the face of the $(k-1)^{th}$ tetrahedron which is not incident to its lowest vertex, in such a manner that the unique new vertex is the top vertex of the k^{th} tetrahedron. A polytope on n vertices thus constructed will be denoted $X(n)$. An example is shown in Figure 2.

A *k-fold inverted pyramid* over a directed d-polytope (P, f) is a directed $(d+k)$-polytope (Q, g) which is constructed by adding to (P, f) k apices $a_1, ..., a_k$,

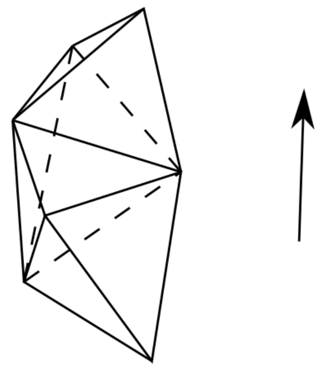

Figure 2: This is X(7). For a different, artistically competent example, see [3].

such that $a_i \notin aff(P \cup \{a_1, ..., a_{i-1}\})$ $\forall i \in \{1, ..., k\}$, and such that $g|_{aff(P)} = f$ and each apex is lower than all the vertices of P and all the apices with smaller indices.

A *prism* is a 3-polytope which is the convex hull of two polygons (its *bases*) which lie in parallel planes. A prism is *orthogonal* if all of its quadrilateral facets, except perhaps for its bases, are rectangles. The digraph of a directed prism (P, f) does not depend upon the particular choice of f, which will therefore be omitted.

A *wedge over a vertex* is a 3-polytope which is produced by contracting exactly one non-base edge of a prism. A wedge is othogonal if it originated from an orthogonal prism.

A *spindle* is a polytope which has a pair of vertices $\{u, v\}$ such that every facet contains exactly one of the vertices u, v. An alternative definition can be obtained from [2], in which a simplex Σ is called *special* for a polytope P if its vertices are vertices of P and each facet of P contains all the vertices of Σ but one. In that sense, a spindle is simply a polytope which admits a special 1-simplex.

A *partial monotone* path of (P, f) is a directed path in $\omega(P, f)$ which begins at v_1. The number of partial monotone paths in (P, f) which end at v_k is denoted $\mu_k(P, f)$, or $\mu^{v_k}(P, f)$. A partial monotone path in $(P, -f)$ is also called a *partial antitone path* in (P, f).

A *monotone path* of (P, f) is a partial monotone path which ends at the top vertex. The number of monotone paths in (P, f) is denoted $\mu(P, f)$.

An *arborescence* of (P, f) is a directed tree which contains exactly one outgoing edge for each vertex of $\omega(P, f)$.

Theorems:

Having finished the presentation of the basic definitions, we proceed to present some useful theorems.

The first theorem that we list is a directed version of Menger's Theorem ([8], Theorem 11.6):

Theorem 2. *Let x and y be two vertices of a digraph D, such that x is not joined to y. Then the number of pairwise vertex-disjoint directed (x, y)-paths in D is equal to the minimum number of vertices (other than x and y) whose deletion destroys all directed (x, y)-paths in D.*

We continue with Whitney's Theorem, a classic:

Theorem 3. *Every 3-connected graph has a unique embedding in S^2 up to composition with homeomorphisms of S^2.*

We will require the Inductive Construction Theorem:

Theorem 4. *Given a polytope P and a point v, the faces of $conv(P \cup \{v\})$ are exactly the following:*

- *every face F of P incident to a facet F' of P such that v is beyond F',*
- *every set $conv(F \cup \{v\})$, where F is incident to two facets F', F'' such that v is beyond F' and beneath F''.*

And we finish with Balinski's Theorem, which we will discuss in detail immediately:

Theorem 5. *Every d-polytopal graph is d-connected.*

0.2 Balinski's Theorem

In 1961, Balinski proved in his doctoral thesis that every d-polytope has a d-connected graph. His proof famously employs the simplex algorithm. Thirty-seven years later, Holt and Klee reproduced this result, albeit with a single-page proof. In fact, Holt and Klee proved the following stronger result:

Theorem 6. *There are d vertex-disjoint monotone paths joining the top and bottom vertices of any directed d-polytope.*

Let us see why Theorem 6 implies Balinski's Theorem. Given any d-polytope P and its vertices x, y, there exists a directed d-polytope (Q, f) such that Q is combinatorially equivalent to P, x is the bottom vertex, and y is the top vertex; simply choose two hyperplanes H_x, H_y such that $H_x \cap P = \{x\}$ and $H_y \cap P = \{y\}$, then perform a perturbation on P as in [5] to make H_x and H_y parallel, and finally choose f to be perpendicular to the two hyperplanes. Consequently, from the Holt-Klee Theorem we obtain Balinski's Theorem.

In brief, Holt and Klee used directed polytopes to obtain a result about directed polytopes, which by the way implied a famous result about undirected polytopes, the latter having been previously obtained in a much more arduous manner. This is excellent motivation for us to examine the proof of the Holt-Klee Theorem, and in fact we shall do so immediately as a warm-up.

The proof is done by induction: in dimension 2, the assertion is obvious. Suppose that it also holds for all dimensions less than d, and let us examine what happens for dimension d. Let (P, f) be a directed d-polytope, and let x and y be its bottom and top vertices, respectively. Without loss of generality, we can identify x with the origin. Let Φ denote the projection to the $d-1$ coordinates perpendicular to f, so that any point $q \in \mathbb{R}^d$ can be written as $q = \Phi(q) + f(q)e_1$. Again, without loss of generality, we can assume that $\Phi(y) = x$.

Next, we recall the directed version of Menger's Theorem; to prove the Holt-Klee Theorem, it is sufficient to show that for every vertex set S of cardinality $d-1$ there exists a monotone path in (P, f) which does not intersect S.

Let H be the hyperplane at which $f(q) = 0$ (i.e. the first coordinate is zero). Let J be the hyperplane at which the second coordinate is zero, J^+ the half-space in which it is positive, and J^- the half-space in which it is negative. The projection $\Phi(S)$ of the vertex set S is contained in a $(d-2)$-dimensional affine space G within H. With the aid of an appropriate isometry (a rotation around the e_1 axis and perhaps a reflection along J), we may assume that $S \subset J \cup J^-$. We now distinguish cases.

- If the e_1 axis does not skewer P, then we have $S \subset J^-$ and $P \cap J^+ = \emptyset$. Then $P \cap J$ is a face of P incident to both x and y and does not intersect S. By the induction hypothesis, there is a monotone path in $(P \cap J, f)$ from x to y missing S.

- If the e_1 axis skewers P, then we have $S \subset J \cup J^-$ and $P \cap J^+ \neq \emptyset$. Let Π be the projection to the first two coordinates. Then $(\Pi(P), f)$ is a directed polygon which has a monotone path from x to y whose internal vertices lie in J^+. This path can be lifted to a monotone path in (P, f). Indeed, each edge of $\Pi(P)$ is the projection of a face of P, and its vertices are the projections of the top and bottom vertices of the face. Hence, each edge of the monotone path from x to y in $(\Pi(P) \cap J^+, f)$ can be lifted to a monotone path on the boundary of a face of (P, f). Concatenating these paths yields a monotone path from x to y in (P, f) which misses S.

By the directed Menger's Theorem, the proof is now complete.

Part I: Reconstruction

Chapter 1

3-polytopes

The theory of reconstruction of polytopes from their graphs has been most successful in dimension three; indeed, the matter is mostly settled for both undirected and directed 3-polytopal graphs. As for most great advances on 3-polytopes, this was greatly facilitated by the use of planar graphs, common spatial intuition, and humankind's cultural familiarity with the topic. As we will see in later chapters, relatively little has been achieved in higher dimensions, where we do not have the aforementioned advantages and, most importantly, interesting classes are rarely reconstructible.

The first theorem that we present in this chapter is arguably the most important theorem in 3-polytopes: Steinitz's Theorem ([10]).

Theorem 7. *A graph is* 3*-polytopal if and only if it is planar and 3-connected.*

The counterpart of Steinitz's Theorem in content and importance for digraphs is the Mihalisin-Klee Theorem [15]:

Theorem 8. *A digraph is* 3*-polytopal if and only if it is acyclic, has a unique source and a unique sink, has three vertex-independent directed paths connecting its source to its sink, and its underlying graph is* 3*-polytopal.*

Both of these theorems are indeed constructive, in that they provide a method to produce a (directed) polytope from a polytopal (di)graph. That the obtained (directed) polytope is unique is a consequence of Whitney's Theorem. Indeed, an embedding $f : G(P) \to S^2$ of a polytopal graph determines P up to combinatorial equivalence, as it determines all the incidences of edges with faces. Homeomorphisms of S^2 preserve these incidences. Hence, all the planar embeddings of $G(P)$, being of the form $h \circ f$, $h \in Hom(S^2)$, determine the same polytope P up to combinatorial equivalence. For digraphs in particular, any isomorphism between the digraphs of two directed 3-polytopes can be extended in this manner to a combinatorial equivalence. These observations yield the following corollaries:

Corollary 1. 3*-polytopes are a reconstruction class up to combinatorial equivalence.*

Corollary 2. *Directed 3-polytopes are a reconstruction class up to orientation equivalence.*

Corollary 3. *Directed 3-polytopes with weakly connected digraphs are a reconstruction class up to order equivalence.*

Finally, Figure 1 establishes that directed 3-polytopes are not a reconstruction class up to order equivalence. There is only one question remaining:

Question 2. *Are there digraphs which are not weakly connected but determine a 3-polytope up to order equivalence?*

1.1 Steinitz's Theorem

Steinitz's Theorem, being Steinitz's Theorem, has been proven in a variety of ways since its conception. Some of the most famous proofs are those found in [4], [16] and of course [10], which we presently follow.

That the graph of a 3-polytope is planar is obvious. That it is 3-connected is a special case of Balinski's Theorem. Hence, we will focus on proving the inverse implication, namely that every 3-connected planar graph G is polytopal.

The proof comes in four distinct steps:

- **Step 1:** We consider G as a plane graph. By Whitney's Theorem, this can be achieved in a unique way.
- **Step 2:** We prove that G either has a vertex of degree 3, or it has a triangular face.
- **Step 3:** We show that G can be reduced by a finite sequence of $Y-\Delta$ and $\Delta-Y$ transforms to a plane graph G' with fewer edges such that, if we have a 3-polytope P' with $G(P')\cong G'$, then we can obtain a 3-polytope P with $G(P)\cong G$.
- **Step 4:** We apply induction on the number of edges of G.

So, consider G as a plane graph, and let n_k be the number of vertices with degree k and p_k be the number of k-sided faces. Then $\Sigma_{k\geq 3}kn_k = 2f_1 = \Sigma_{k\geq 3}kp_k$. Hence, by Euler's formula, $\Sigma_{k\geq 3}k(n_k+p_k) = 4f_0+4f_2-8 \Rightarrow n_3+p_3 = 8+\Sigma_{k\geq 5}(k-4)(n_k+p_k)\geq 8$, which certainly implies that $f(G)$ has either a vertex of degree 3 or a triangular face.

A $Y-\Delta$ transform over a plane graph consists of deleting a vertex of degree 3 and connecting with edges its non-adjacent neighbours. A $\Delta-Y$ transform consists of introducing a new vertex inside a triangular face, connecting it with the vertices of the face, deleting the edges of the face, and smoothing out any vertices of degree 2. Examples are seen in Figure 1.1 and Figure 1.2.

Applying a $Y-\Delta$ or $\Delta-Y$ transform over G produces a plane graph G'. Observe that G' is also 3-connected, as seen by applying Menger's Theorem; for any three vertex-independent paths between two vertices of G, only one can

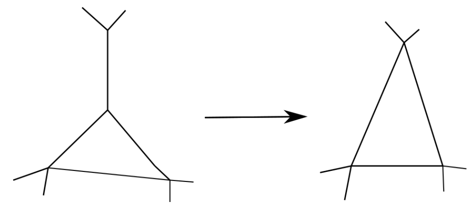

Figure 1.1: A $Y - \Delta$ transform.

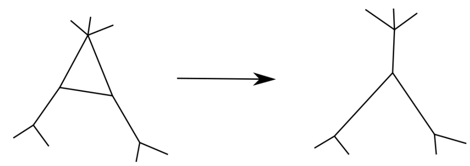

Figure 1.2: A $\Delta - Y$ transform.

pass from any vertex of degree 3 or triangular face, and it can be transformed to fit the obtained graph.

Now we proceed to step 3; let P' be a polytope with $G(P') \cong G'$. Suppose that G' has been obtained from G through a $\Delta - Y$ transform producing a vertex v. Consider the plane which contains the neighbours of v and take the intersection of P' with the half-space of this plane which does not contain v. For the resulting polytope P we have $G(P) \cong G$. On the other hand, if G' has been obtained from G through a $Y - \Delta$ transform producing a triangular face F from a vertex u, then we modify P' as follows: if 1, 2, or 3 pairs of neighbours of u are connected in G, then we select a point p in P' which rests in the affine hulls of 1, 2 or 3 of the faces surrounding F, respectively. We also take care to pick p so that it is beyond the affine hull of F. This is straightforward in the first two cases, but in the third case p is fully determined as the unique point of the intersection of the affine hulls of the three faces surrounding F; hence, we might need to perform an appropriate projective transformation on P' to ensure that this condition is fulfilled (such a transformation exists as long as P' is not the tetrahedron). Having selected p, we take $P := conv(P' \cup \{p\})$, and again we have that $G(P) \cong G$. Thus, in all cases, if G' is 3-polytopal, then G is 3-polytopal.

Conceptually speaking, the proof is near completion. We would now like to continue with step 4, saying something akin to: "every $Y - \Delta$ and $\Delta - Y$ transform reduces the total number of edges, so, if we apply enough of those to G, we will obtain a sequence $G \to G' \to G'' \to ... \to G^{(k-1)} \to G^{(k)} = K_4$, since K_4 is the planar 3-connected graph with the minimum number of edges. However, K_4 is 3-polytopal, being the graph of the tetrahedron, so the same is true for $G^{(k-1)}$, $G^{(k-2)}$, etc. up to G, completing the proof".

However, the beginning of the above argument is not quite accurate. A $Y - \Delta$ transform on a vertex with pairwise non-adjacent neighbours does not

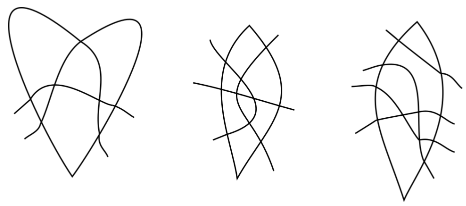

Figure 1.3: A non-lens, decomposable lens, and indecomposable lens, in that order.

decrease the number of edges. A $\Delta - Y$ transform on a face without vertices of degree 3 does not decrease the number of edges, either. That is, in order for a planar 3-connected graph G to have a $Y - \Delta$ or $\Delta - Y$ transform which produces a graph G' with fewer edges, G must have a triangular face incident to a vertex of degree 3. We will now show that, even if this is not true for G, it is true for some graph H which can be obtained from G through $Y - \Delta$ and $\Delta - Y$ transforms. This will conclude the proof. Firstly, we need to establish some useful properties of 4-regular graphs.

Let $\mathcal{J}$ be a 4-regular, 3-connected plane graph. An edge $\{x, y\}$ of $\mathcal{J}$ has a *direct extension* $\{y, z\}$ if the path xyz separates the other two edges of y. If every edge in a path (resp. cycle) is a direct extension of the edges of the path (resp. cycle) which are adjacent to it, then the path (resp. cycle) is called a *geodesic* (resp. *closed geodesic*) of $\mathcal{J}$ (this is unrelated to the metric geodesic).

A subgraph $\mathcal{L}$ of $\mathcal{J}$ is called a *lens* if:

- $\mathcal{L}$ consists of a cycle C composed of two geodesics A and B, and of all the vertices and edges contained in one of the connected components of $\mathbb{S}^2 \setminus C$, called the *inner* vertices and edges, and
- no inner edge of $\mathcal{L}$ is incident to either of the two points of C where the geodesics which make up C intersect (these points are called the *poles* of $\mathcal{L}$)

A lens $\mathcal{L}$ of $\mathcal{J}$ is *indecomposable* if no lens of $\mathcal{J}$ is a proper subgraph of $\mathcal{L}$. At least one indecomposable lens is contained in $\mathcal{J}$; if we consider the set S of subgraphs of $\mathcal{J}$ which comprise of a cycle C consisting of at most two geodesics and all the vertices and edges in one of the connected components of $\mathbb{S}^2 \setminus C$, then $S \neq \emptyset$, as we can simply pick a geodesic starting from a point and follow it until it intersects itself. Indecomposable lenses are exactly the minimal elements of S ordered by inclusion. For examples and non-examples of lenses and indecomposable lenses, view Figure 1.3.

The following statements hold for an indecomposable lens $\mathcal{L}$:

- Each point a in the boundary geodesic A that is not a pole is the end of exactly one maximal geodesic contained in the interior of $\mathcal{L}$. We call this geodesic the *cut* of a and denote it $C(a)$.

- For every $a \in A$, $C(a) \cap A = \emptyset$, since otherwise $C(a)$ and a part of A would bound a lens properly contained in $\mathcal{L}$, a contradiction. Hence, the other end of $C(a)$ is some $b \in B$. In particular, the two geodesics which bound an indecomposable lens have the same length.
- Any two cuts intersect at one point at most, otherwise the intervals between the intersection points would again bound a lens properly contained in $\mathcal{L}$.
- Every inner edge of $\mathcal{L}$ belongs to exactly one cut. Every inner vertex belongs to exactly two cuts.

Now we prove the following useful lemma:

Lemma 1. *Every indecomposable lens $\mathcal{L}$ contains a triangular face incident to the boundary C of $\mathcal{L}$.*

Proof. If $\mathcal{L}$ has no inner vertices, then its faces which are incident to the poles are triangular. Otherwise, let $v_1, ..., v_n$ be the inner vertices which are adjacent to vertices $a_1, ..., a_n$ on the boundary geodesic A. Let $h(v_i)$ be the number of faces contained in the triangular region determined by the cuts $C(a_i), C(a^i)$ which meet at v_i and the interval of A between a_i and a^i. Then the inner vertex v_i with the smallest $h(v_i)$ defines a triangular face $a_i v_i a^i$. □

Let us now apply these observations to the proof of Steinitz's Theorem. For our 3-connected plane graph G, we define the graph $I(G)$ such that the vertices of $I(G)$ are the edges of G, and two vertices are connected in $I(G)$ if the corresponding edges of G are incident and also belong to the same face. Note that $I(G)$ is a 4-regular plane graph. Additionally, since G is 3-connected, it is also 3-edge-connected, which implies that $I(G)$ is 3-connected. Hence, all of our previous notes concerning lenses are true for $I(G)$.

Each face of $I(G)$ corresponds to exactly one vertex or face of G, so $f_2(I(G)) = f_2(G) + f_0(G)$. In particular, the number of sides of a face of $I(G)$ equals the degree of the corresponding vertex. Thus, in order to show that G contains a vertex of degree 3 incident to a triangular face, it suffices to show that $I(G)$ contains adjacent triangular faces, i.e. $K_4^- \subseteq I(G)$.

Let $g(G)$ be the minimum number of faces in any indecomposable lens in $I(G)$. If $g(G) = 2$, then observe that the only indecomposable lens with two faces is K_4^-, so we are done. If $g(G) > 2$, then let $\mathcal{L}$ be the indecomposable lens with $g(G)$ faces, and let T be a triangular face of $\mathcal{L}$ incident to the boundary of $\mathcal{L}$; such exists from Lemma 1.

Figure 1.4 illustrates how to perform an appropriate $Y-\Delta$ or $\Delta-Y$ transform on G to obtain a graph G' with $g(G') < g(G)$: if T as a face of $I(G)$ corresponds to a face of G, then the $\Delta-Y$ transform (a) diminishes the number of faces in $\mathcal{L}$, whereas for T corresponding to a vertex of G, the $Y-\Delta$ transform (b) has the same effect.

Consequently, a finite number of $Y-\Delta$ and $\Delta-Y$ transforms suffices to transform G into a graph H which has a vertex of degree 3 incident to a triangular face. The proof is complete.

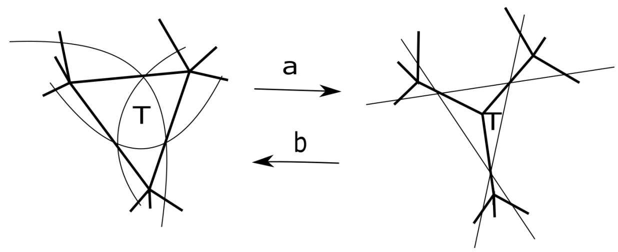

Figure 1.4: Heavy edges denote G, light edges denote $I(G)$.

Time Complexity: We seek to calculate the number of steps ($Y - \Delta$ or $\Delta - Y$ transforms) required to reach the polytope determined by G starting from the tetrahedron. That is, we want to find the length of the sequence $G \to G' \to G'' \to ... \to G^{(k-1)} \to G^{(k)} = K_4$.

Denote $|V(G)|$ by n. Since G is a plane graph, $|E(G)| = O(n)$, hence $I(G)$ has $O(n)$ vertices. Since $I(G)$ is a plane graph, it also has $O(n)$ faces. An indecomposable lens of $I(G)$ has at most as many faces as $I(G)$ itself, i.e. $O(n)$. Consequently, applying $Y - \Delta$ or $\Delta - Y$ transforms to reach a graph $G^{(i)}$ with $g(G^{(i)}) = 2$ starting from G requires $O(n)$ steps. At that point, we can perform an additional $Y - \Delta$ or $\Delta - Y$ transform to obtain a graph $G^{(i+1)}$ with fewer edges than G. Since $E(G) = O(n)$, to reach the tetrahedron, we must repeat this process $O(n)$ times.

Note that all the graphs $I(G^{(i)})$ also have $O(n)$ faces, and so $g(G^{(i)}) = O(n)$ $\forall i \in [n-1]$; indeed, the graphs $G^{(i)}$ have at most $E(G)$ edges, so the graphs $I(G^{(i)})$ have at most $E(G)$ vertices, hence at most $E(G) + 1 = O(n)$ vertices, since they are connected, hence $O(n)$ faces, since they are plane graphs.

We deduce that the number of steps required to transition from the tetrahedron to the polytope determined by G is $O(n) \cdot O(n) = O(n^2)$.

This is the best currently known upper bound for the time complexity of reconstructing a 3-polytope from its graph. In a recent paper ([9]), Chang, Cossarini, and Erickson proved that $\Omega(n^{\frac{3}{2}})$ steps are necessary in certain cases.

1.2 The Mihalisin-Klee Theorem

Both this result and its proof have many parallels to Steinitz's Theorem. Again, one direction is straightforward: for a 3-polytopal digraph $\omega(P, f)$, it is evident that the underlying graph of G is 3-polytopal, G is acyclic, and G has a unique source and a unique sink. Finally, by the Holt-Klee Theorem, G must also contain three vertex-independent directed paths connecting its source to its sink.

For the converse, let G be a digraph which is acyclic, has a unique source and a unique sink, has three vertex-independent directed paths connecting its source to its sink, and has a 3-polytopal underlying graph. For brevity's sake, we will call G and other digraphs with all of the above properties *3-monotone*.

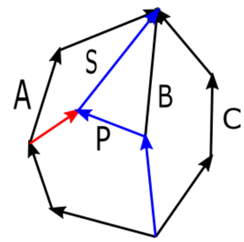

Figure 1.5: B' is noted in blue and $P \setminus P'$ is noted in red.

Additionally, as we did in Steinitz's Theorem, we consider G as a plane graph in the unique possible way.

Our first observation is that G contains the digraph T of the tetrahedron as a topological minor, and indeed T can be obtained from G without deleting either the source x or the sink y of G.

To see this, let $\mathcal{Q}$ be the set of quadruples (A, B, C, P), where:

- A, B, C are directed paths in G connecting x to y,
- P is a path in G (not necessarily directed) connecting an interior point of A to an interior point of B,
- A, B, C, P are vertex-independent.

Since G is 3-monotone, it contains three directed paths A, B, C connecting x and y. Since G is simple, two of these, say A and B, must contain interior vertices, say u and v. Since G is 3-connected, it contains a path P which connects u and v and avoids x and y. Let $u' \in A$, $v' \in B$ be two consecutive points of $P \cap (A \cup B)$ in P, and let R be the subpath of P between u' and v'. Then $(A, B, C, R) \in \mathcal{Q}$. This proves that $\mathcal{Q} \neq \emptyset$.

Now let $(A, B, C, P) \in \mathcal{Q}$ such that P is of minimum length. Suppose that P connects $u \in A$ to $b \in B$. Let p be the other endpoint of the longest directed subpath of P containing b, say P'.

If P' is directed from b to p, then let S be a path leading from p to y. S meets A or B before meeting C. Observe that S meets A before y. Were it otherwise, (A, B, C, P) would not be a quadruple with its last element being of minimum length in $\mathcal{Q}$, since $(A, B', C, P \setminus P') \in \mathcal{Q}$, where B' coincides with B from x to b, and then coincides with P' and S (see Figure 1.5). Let S' be the initial segment of S up to the vertex where it intersects A for the first time. Then (A, B, C, S') is a subdivision of the digraph T of the tetrahedron, as required.

The process is quite similar if P leads from p to b; in this case we consider a path S leading from x to p, and proceed as above.

The next step of the proof is to notice that there is a sequence $(J_0 := T, f_0), ..., (J_{k-1}, f_{k-1}), (J_k := G, f_k)$ of 3-monotone topological minors of G and their embeddings in G such that:

- each J_i is obtained from J_{i+1} by deleting an edge and smoothing out its end vertices,
- $f_0(J_0)$ contains the source x and sink y of G,
- $f_i|_{V(J_{i-1})} = f_{i-1}|_{V(J_{i-1})}$ $\forall i \in \{1, ..., k\}$,
- each f_i, when possible, maps the edges of J_i to paths of length ≥ 2 in G.

Indeed, we have already shown that T is a 3-monotone topological minor of G, and that it can be embedded in G so that the source and the sink of G are contained in the image of the embedding. Additionally, if $J_i \neq G$, then there exists a monotone path in G which is not intersected internally by $f_i(J_i)$; for instance, one can concatenate a directed path leading from x to a vertex $v \in G \setminus f_i(J_i)$ and another leading from v to y, then take a subpath of the resulting path which is not internally intersected by $f_i(J_i)$. If such a vertex v does not exist, then the required path is an edge of G.

Let us denote such a path P_i. We construct H_{P_i} by adding an appropriately oriented edge e to J_i connecting the points of $f_i^{-1}(f_i(J_i) \cap P_i)$, and we define the embedding h_{P_i} which satisfies $h_{P_i}|_{J_i} := f_i$ and $h_{P_i}(e) := P_i$. Let us call this pair (H_{P_i}, h_{P_i}) the *canonical extension* of (J_i, f_i) through P_i. A cursory look suggests that we can simply define (J_{i+1}, f_{i+1}) to be a canonical extension of (J_i, f_i). After all, it is trivial to see that any H_{P_i} is indeed 3-monotone. However, there is a reasonable concern that H_{P_i} might not be simple. We must take care to counter this possibility. We distinguish two cases:

- Suppose that $f_i(E(J_i)) \subset E(G)$. Then the ends of P_i are non-adjacent vertices of J_i, so H_{P_i} is simple.
- Suppose that some edge e of J_i is mapped to a path of length ≥ 2 in G. Let P be a path connecting a vertex $u \in f_i(e)$ and a vertex $v \in f_i(J_i) \setminus f_i(e)$ and avoiding the ends of $f_i(e)$. We connect each relative source or sink s_j of P to a vertex in $f_i(J_i)$ through directed paths p_j which are not internally intersected by $f_i(J_i)$; in particular, if $s_l \in f_i(J_i)$, we set p_l to be the trivial path. Let s_r be the relative source or sink of P nearest to u such that the other end of p_r, say b, is not in $f_i(e)$, but rather in the image $f_i(e')$ of some other edge of J_i. If the other end, say a, of p_{r-1} is not in $f_i(e')$, then let M be the subgraph of G which is obtained by concatenating p_{r-1}, the path between s_{r-1} and s_r, and p_r, and set P_i to be a path from a to b in M (see Figure 1.6).

  Then H_{P_i} is simple. On the other hand, if $a \in f_i(e')$, then $a \in f_i(e) \cap f_i(e')$, so a is the image of a vertex of J_i which is an end of both e and e'. In that case, we do not take (J_{i+1}, f_{i+1}) to be the canonical extension of (J_i, f_i); instead, J_{i+1} is constructed and embedded as in Figure 1.7. It is straightforward to see that J_{i+1} is 3-monotone and simple.

Now that we have constructed the sequence $(J_i, f_i)_{i=0}^n$, we enter the final stage of the proof of the Mihalisin-Klee Theorem. By Steinitz's Theorem, all the

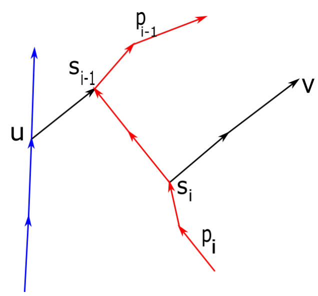

Figure 1.6: The local change in the embedding in the case $a \notin f(e')$. The edge-embedding $f(e)$ is noted in blue, whereas the new edge-embedding corresponding to the path M is noted in red.

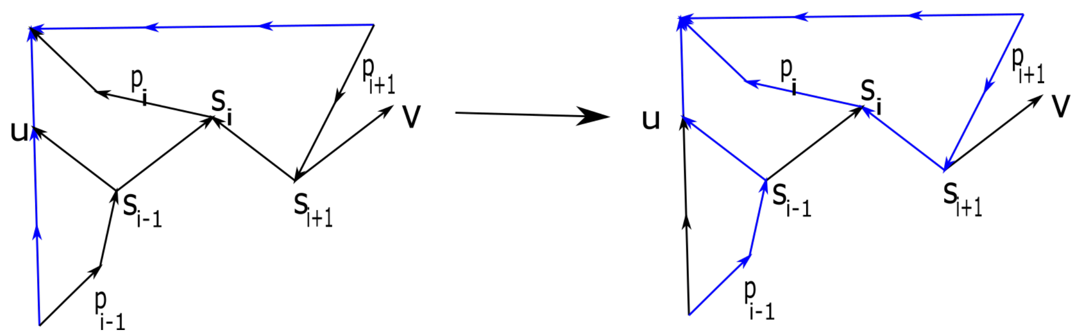

Figure 1.7: The local change in the embedding (noted in blue) in the case $a \in f(e) \cap f(e')$.

underlying graphs of the J_i are 3-polytopal; specifically, each defines a unique 3-polytope up to combinatorial equivalence. Consider a linear extension L of the poset induced by G, and embed the tetrahedron in $\mathbb{R}^3$ so that the z-coordinate of each vertex v of its 1-skeleton J_0 equals the rank of $f_0(v)$ in L. For each subsequent J_i, we produce a 3-polytope which has J_i as its digraph by slightly "bending" a face of the 3-polytope corresponding to J_{i-1} to obtain the extra edge; this is always geometrically feasible with a slight perturbation of the vertices, as proven in [**?**] At each such step, we take care that the z-coordinate of any newly introduced vertex is determined by its corresponding rank in L. In this manner we obtain a directed 3-polytope J_n which has G as its digraph. The proof is complete.

Time Complexity: Denote $V(G)$ by n. Each pair (J_i, f_i) requires $O(n^2)$ steps to be constructed from (J_{i-1}, f_{i-1}), and there are $E(G) - 6 = O(n)$ such pairs, yielding a total execution time of $O(n^3)$.

Chapter 2

Simple Polytopes

During the 1980's, Perles conjectured in various occasions, most notably in two separate Oberwolfach meetings on convex bodies ('84, '86), that simple polytopes can be determined by their graphs. This was eventually proven in 1987 by Blind and Mani-Levitska ([7]).

Theorem 9. *Simple polytopes are a reconstruction class up to combinatorial equivalence.*

Soon after, Gil Kalai provided a proof in a mere page ([13]); it is this proof that we present here. Apart from brevity, Kalai's proof has two other major advantages. For one, it uses directed polytopes (generally, "good" orientations), and in that manner showcases directed polytopes as a natural and useful mathematical object. For another, it can be manipulated to produce an analogue of the Blind-Mani-Levitska Theorem for directed polytopes:

Corollary 4. *Directed simple polytopes are a reconstruction class up to orientation equivalence.*

Corollary 5. *Directed simple polytopes with weakly connected digraphs are a reconstruction class up to order equivalence.*

Once again, Figure 1 gives a counterexample to an analogue for order reconstruction of simple polytopes in general. Additionally, let us ask:

Question 3. *Are there simple digraphs which are not weakly connected but determine a simple polytope up to order equivalence?*

Despite its merits, Kalai's work does not provide a good algorithm for reconstructing a (directed) polytope from its (di)graph; considering all the possible orientations of (the underlying graph) G, determining the good ones, and applying Kalai's criterion to find the faces of the reconstructed polytope requires exponential time. Devising an algorithm which achieves combinatorial reconstruction for simple polytopes in polynomial time is far from trivial.

Of course, this problem stirred a lot of interest among combinatorialists. In 2002, Joswig, Kaibel, and Körner provided polynomial certificates, and soon after Friedman published a polynomial algorithm for reconstructing a simple polytope from its graph. We review Friedman's algorithm in this chapter.

2.1 Kalai's Theorem

Let G be the graph of a simple polytope P. To begin, note that G must be d-regular, and that $dimP = d$.

For an acyclic orientation O of G we say that O is *good* if every non-empty face of P is incident to exactly one sink of G^O. There exist, of course, good orientations of G (e.g. LP-admissible orientations). Let h_k^O be the number of vertices with in-degree k of O. Define $f^O := \Sigma_{i=0}^d 2^i h_i^O$.

If x is a vertex of P with in-degree k, then it is the sink of 2^k different faces (since P is simple, each set of edges incident to x defines a face that contains it). Hence, if O is a good orientation, $f^O = f$, since f^O enumerates pairs of faces and their sinks. Otherwise, $f^O > f$.

Now, the faces of P can be determined uniquely; they are defined exactly by the induced, connected, regular subgraphs of G the vertex sets of which are ideals of posets induced by good orientations of G. Indeed, every face F of P matches this description - simply embed P in $\mathbb{R}^d$ and derive O from a linear functional which obtains its lowest values at the vertices of F.

For the converse, let H be an induced, connected, regular subgraph of G and let O be a good orientation such that $V(H)$ is an ideal of the poset induced by O. Let x be the sink of H with respect to O. There are k edges incident to and oriented towards x, therefore x is incident to a k-face F which contains these k edges. Since O is a good orientation, x is the unique sink of G incident to F, therefore the vertices of F lie in $\downarrow x$ in the poset induced by O. This implies that $V(F) \subseteq V(H)$. Since both H and $G(F)$ are k-regular and connected, $V(H) = V(F)$, thus $H \cong G(F)$. The proof is complete.

Remark 1. *If we label the vertices of $G(P)$, we observe that Kalai's criterion uniquely determines the faces of P, i.e. any automorphism of $G(P)$ can be extended to a combinatorial equivalence $\phi : P \to P$. Hence, any two directed simple polytopes with the same digraph are orientation equivalent, proving Corollary 5.*

2.2 Friedman's Algorithm

Let G be the graph of a simple polytope P, and let O be an acyclic orientation of G with source a specific vertex $v \in V(G)$. A *2-frame* is a pair of incident edges, and its *centre* is the vertex at the intersection of the edges. A 2-frame is called a *sink* if both of its edges are oriented toward their common vertex. Let $h(O)$ be the number of 2-frames which are sinks under O. A *2-system* is a set of cycles of G such that every 2-frame is contained in a unique such cycle. Let S be a 2-system of G. The following lemma can be found in [12]:

Lemma 2. *$|S| \leq f_2 \leq h(O)$, with $f_2 = h(O)$ if and only if O is good.*

Proof. Since O is acyclic, every cycle must contain a sink, including those which bound 2-faces; hence, $f_2 \leq h(O)$, with equality if and only if O is good. Since S is a 2-system, $|S| \leq h(O)$. Hence, for a good orientation O' we have $|S| \leq h(O') = f_2$. Putting together these inequalities proves the required result. □

Moreover, we use the above lemma to define a *pseudopolytopal multigraph* as the 1-skeleton of a 2-complex for which Lemma 2 holds. An important observation is that performing a 2-face contraction on a simple polytope transforms its graph to a pseudopolytopal multigraph.

We now present the main part of the proof. Let W be the set of cycles of G and let T be the set of 2-frames of G. We introduce the following optimisation problems:

- IP-S: $\max \Sigma_{w \in W} x_w$ s.t. $\forall t \in T\ \Sigma_{w \supset t} x_w = 1$, $x_w \in \{0, 1\}$,
- LP-S: $\max \Sigma_{w \in W} x_w$ s.t. $\forall t \in T\ \Sigma_{w \supset t} x_w \leq 1$, $x_w \geq 1$,
- LP-SD: $\min \Sigma_{t \in T} v_t$ s.t. $\forall w \in W\ \Sigma_{t \subset w} v_t \geq 1$, $v_t \geq 0$.

By Lemma 2, a solution to IP-S is a 2-system of size f_2. Since LP-S is a relaxation of IP-S, we have for their optimal objective values that $Opt(IP-S) \leq Opt(LP-S)$. Since LP-SD is the dual problem of LP-S, $Opt(LP-S) = Opt(LP-SD)$.

Now we prove that LP-S can be solved in polynomial time using the ellipsoid algorithm. This is because the dual, LP-SD, has a polynomial number of variables and all constraints have polynomially bounded size (since all the coefficients are 0 or 1). Then, because of the equivalence of separation and optimization (see Corollary 14.1g in [17]) it remains to show is that there exists a polynomial separation algorithm for LP-SD.

We construct the directed graph H. The vertices of H are the labelled by the ordered 2-frames of G (i.e. if $t := (e, e') \in T$ then $u_{(e,e')}$ and $u_{(e',e)}$ are both vertices of H). Additionally, there is an edge from $u_{(a,a')}$ to $u_{(b,b')}$ if $a' = b$ and the centres of (a, a') and (b, b') differ. For any two 2-frames t, t' we define the cost of an edge of H from u_t to $u_{t'}$ to be $\frac{v_t + v_{t'}}{2}$. Every cycle in H corresponds in a natural way to a cycle w of G with cost $\Sigma_{t \subset w} v_t$. Hence, solving LP-SD is equivalent to finding the cycle of minimum cost in H, which can be done in polynomial time trough linear programming. We have shown that an optimal solution for LP-SD, and hence its dual LP-S, can be found in polynomial time.

Furthermore, the optimal solution is unique. Indeed, suppose that there are two distinct optimal solutions x and x' for LP-S. Then there exists some $w \in W$ such that $x_w = 1$ and $x'_w = 0$, otherwise $(1 + \epsilon)x - \epsilon x'$ would also be an optimal solution, implying that x' is not an extreme point, a contradiction. By contracting the 2-face implied by x'_w we obtain a pseudopolytopal multigraph

G'. The optimal solution of IP-S for G' must have objective value $f_2 - 1$. However, the projection of x onto G' is feasible for LP-S but has objective value f_2, which is a contradiction. Hence, $LP - S$ has a unique optimal solution.

We have shown that the 2-skeleton of P can be reconstructed in polynomial time. However, it is straightforward to reconstruct a simple polytope from its 2-skeleton (in [13], Kalai credits Perles as the first to have made this observation).

First, we reconstruct the facets. To produce a facet, start with any $(d-1)$-frame centered at u and let u' be the unique vertex adjacent to u which is not in that frame. Consider any other vertex, u'', in the frame. Clearly there exists another frame in the same facet as the first $(d-1)$-frame which is centered at u''. Now consider the vertex $\tilde{u}$ adjacent to u'' which is contained in the 2-face that contains the 2-frame (u, u', u''). It is shown in [?]hat the $(d-1)$-frame centered at u'' in that facet contains all the vertices except $\tilde{u}$. We can use this procedure to find all the facets in polynomial time.

We note that the face lattice may contain exponentially many faces, so explicitly presenting it may require exponential time. However, one can compute the face lattice in polynomial time, since the faces of P are exactly the intersections of its facets.

Chapter 3

Other Cases

There is no dearth of reconstruction classes up to our chosen equivalence relationships, but a few of them are particularly interesting. Up to this point, we have examined two important reconstruction classes of (directed) polytopes: 3-polytopal and simple. In this brief chapter, we mention a few more. Let us begin with some definitions.

The *cyclic polytope* $C(n, d)$ is the convex hull of n distinct points on the moment curve $\gamma = (t, t^2, ..., t^d)$.

The *Kleetope* P^K of a polytope P is the convex hull of P and points $x_1, ..., x_{f_{d-1}(P)}$ chosen so that each point x_i lies beyond exactly one facet of P and, for every two points x_i, x_j, the line $conv(x_i, x_j)$ intersects the interior of P.

We call a polytopal graph *dimensionally ambiguous* if it is isomorphic to the graphs of two polytopes of different dimensions. We call it *strongly* (d)-*ambiguous* if it is isomorphic to the graph of two (d)-polytopes which are not combinatorially equivalent. We call it *weakly* (d)-*ambiguous* if it is isomorphic to the graph of a (d)-polytope P through two different graph isomorphisms ϕ and ψ, and $\phi \circ \psi^{-1}$ cannot be extended to a combinatorial equivalence of P with itself.

3.1 Kleetopes of Large Cyclic Polytopes

Our first example has been sorely missing from our analysis. It is an example of an interesting class of polytopes which is reconstructible up to combinatorial equivalence, but its directed analogue is not reconstructible up to orientation equivalence.

Theorem 10. *The Kleetopes of cyclic polytopes are a reconstruction class up to combinatorial equivalence.*

Proof. The vertices of a cyclic polytope $C(n, d)$ have degrees $\geq d$, hence the only vertices of $C(n, d)^K$ which have degree d are the added ones, and every other vertex of $C(n, d)^K$ has degree strictly greater than d. Consequently, given the graph of $C(n,d)^K$, we know that d is equal to the minimum degree of the graph, and that n is equal to the number of vertices which do not have the minimum degree. Since $C(n,d)^K$ is a class of combinatorial equivalence, the result follows. □

Remark 2. *Directed Kleetopes of cyclic polytopes are not a reconstruction class up to orientation equivalence; any combinatorial equivalence $\phi : C(n,d)^K \to C(n,d)^K$ must map vertices of $C(n,d)$ to vertices of $C(n,d)$. If the second copy of $C(n,d)$ is a slight rotation of the first which transposes v_1 and v_2 and fixes every other vertex in the f-ordering, then since the graph of $C(n,d)$ is weakly connected, ϕ must respect the f-ordering of the vertices of $C(n,d)$, so it must transpose v_1 and v_2 and fix every other vertex. But then, if $n \geq d+3$, there exists a vertex in the Kleetope of the first copy which is adjacent exactly to v_1 and to the d highest vertices (because by the Gale Evenness Criterion these form a facet), but there does not exist a vertex in the Kleetope of the second copy which is adjacent exactly to $\phi(v_1) = v_2$ and to the d highest vertices (because by the Gale Evenness Criterion these do not form a facet). Hence, there does not exist an isomorphism between the digraphs of the two copies of $C(n,d)^K$, much less an orientation equivalence between them.*

3.2 Polytopes with Unambiguous Graphs

We finish our reconstruction analysis by noting how reconstruction interacts with ambiguity:

- A class of polytopes (directed polytopes) with dimensionally unambiguous graphs (underlying graphs) which contains at most one polytope of each dimension is combinatorially (orientation) reconstructible.
- The set of (d-)polytopes with strongly (d-)unambiguous graphs is combinatorially reconstructible.
- The set of directed (d-)polytopes with weakly (d-)unambiguous underlying graphs is orientation reconstructible.

Part II: Enumeration

Chapter 4

Counting

Traditionally, the only quantity that is regularly counted in directed polytopes is the length of their longest monotone path. The purpose is, of course, to calculate the worst-case time complexity of the simplex algorithm; each monotone path of a directed polytope represents an execution of the simplex algorithm starting from the minimum basic feasible solution, hence an upper bound on the greatest possible length of a monotone path of a directed d-polytope on k facets is also an upper bound on the number of steps of an execution of the simplex algorithm for k constraints on d variables.

For this reason, the problem of bounding the directed diameter of directed polytopes has been a very popular topic for both research and surveys, to the extent that there is little that the author might add in this brief reckoning. A famous survey on the matter is [1]. For some recent progress, see [6].

Here we will instead present the available results concerning the enumeration of objects related to directed polytopes, and in particular the following two: arborescences and monotone paths. All of these results were only discovered too recently, in a paper by Athanasiadis, De Loera, and Zhang [3].

4.1 Arborescences

Let (P, f) be a directed polytope and let $\tau(P, f)$ be the number of its arborescences. The basic result concerning arborescences is the following:

Theorem 11. *The number of arborescences in (P, f) is given by the formula $\tau(P, f) = \prod_{i=1}^{n-1} d_i$. In particular, if P is simple, then we obtain the formula $\tau(P, f) = \prod_1^d i^{h_i(P)}$.*

Proof. An arborescence of (P, f) is exactly the union of a set of outgoing edges, one for each vertex of $\omega(P, f)$, so the first formula is derived from the multiplication principle. For the second formula we take note of an interesting combinatorial interpretation which the h-vector affords for simple polytopes: $h_k(P)$ is exactly the number of vertices of $\omega(P, f)$ with outdegree k, regardless of the choice of f (see Sections 3.4 and 8.3 and Exercise 8.10 in [18]). This observation, combined with the first formula, yields the second one. □

Theorem 11 implies the following quick bounds for the number of arborescences in a simple directed polytope:

Corollary 6. *For $m > d \geq 4$, the maximum number of arborescences over simple directed d-polytopes with m facets is*

$$max\tau(P,f) = \prod_{i=1}^{[\frac{d}{2}]} i^{\binom{m-d+i-1}{i}} \prod_{i=0}^{[\frac{d-1}{2}]} (d-i)^{\binom{m-d+i-1}{i}},$$

and it is achieved by the duals of neighbourly polytopes.

For $m > d \geq 4$, the minimum number of arborescences over simple directed d-polytopes with m facets is

$$min\tau(P,f) = d((d-1)!)^{m-d},$$

and is achieved by the duals of stacked polytopes.

Finally, for simple 3-polytopes with m facets, $\tau(P,f) = 3 \cdot 2^{m-3}$.

Proof. The case $d \geq 4$ is due to the fact that neighbourly polytopes (respectively, stacked polytopes) maximise (respectively, minimise) the entries of the h-vector. The case $d = 3$ is due to the fact that $h_0(P) = h_3(P) = 1$ and $h_1(P) = h_2(P) = m - 3$ for every simple directed 3-polytope P. □

Now let us view the known bounds for general polytopes.

Theorem 12. *For $n > d \geq 3$, the maximum number of f-arborescences over all d-dimensional polytopes with n vertices is achieved by the stacked polytope $X(n)$ for $d = 3$ and by any 2-neighborly polytope for $d \geq 4$. This number is equal to $2 \cdot 3^{n-3}$ and $(n-1)!$ in the two cases, respectively.*

Proof. Since $d_k \leq n - k$ $\forall k \in [n]$, we get $\tau(P,f) \leq (n-1)!$. The equality holds for 2-neighbourly polytopes. Such polytopes exist for every n in dimensions $d \geq 4$ (e.g. the cyclic polytopes), but not in dimension 3. If (P,f) is a directed 3-polytope, then its graph is planar, so it has $\Sigma_{i=1}^{n-1} d_i \leq 3n - 6$ edges. Hence, $\prod_{i=1}^{n-1} d_i = \prod_{i=1}^{n-2} d_i \leq 2 \cdot 3^{n-3}$, with equality when $d_{n-2} = 2$ and $d_i = 3$ for $i \leq n - 3$. This is exactly the multiset of outdegrees of the stacked polytope $X(n)$. □

Theorem 13. *The minimum number of arborescences in directed 3-polytopes is $2(n-1)$ and is exhibited by the inverted pyramid over the $(n-1) - gon$.*

Proof. Let k be the number of outdegrees greater than 1, and denote these $d', d'', ..., d^{(k)}$. Then $k \geq 2$. Note that $2\Sigma_{i=1}^{n} d_i = \Sigma_{i=1}^{n} D_i$. Also, $d_n = 0$, $D_1 = d_1 \geq 3$ and $D_i \geq \max\{d_i + 1, 3\}$ $\forall i > 1$. These considerations yield $2\Sigma_{i=1}^{n} d_i \geq \Sigma_{i=1}^{k} d^{(i)} + k - 1 + 3(n-k) \Rightarrow \Sigma_{i=1}^{k} d^{(i)} \geq k - 1 + 3(n-k) - [2(n-k) - 2] = n + 1$. From this, it is an easy minimisation problem to obtain $\prod_{i=1}^{k} d^{(i)} \geq 2(n-1)$. It is straightforward to see that this is achieved by the inverted pyramid over an $(n-1)$-gon. □

Unfortunately, little can be said about the outdegrees of $\omega(P, f)$ when P is of dimension $d \geq 4$ (although we will answer a couple of relevant questions in a later chapter). Hence, the above proof cannot be generalised to higher dimensions. As a consequence, the following question remains open:

Question 4. *What is the minimum number of arborescences in a d-polytope on n vertices? Is it* $2(n-1)...(n-d+2)$*, achieved by the* $(d-2)$*-fold inverted pyramid over the* $(n-d+2)$*-gon?*

4.2 Monotone Paths

Let (P, f) be a directed polytope, let $\mu(P, f)$ be the number of its monotone paths, and let $\mu_k(P, f)$ be the number of its partial monotone paths ending at v_k. The basic theorem about the enumeration of monotone paths is the following:

Theorem 14. $\mu(P, f) = 1 + \Sigma_{k=1}^{n-1}(d_k - 1)\mu_k(P, f)$

Proof. Let $V_k = \{v_1, ..., v_k\}$ and let I_k be the ideal generated by V_k in the poset induced by $\omega(P, f)$. Let S_k be the set of partial monotone paths which end at successors of elements which are maximal in I_k. Then $|S_k| - |S_{k-1}| = (d_k - 1)\mu_k(P, f)$, since there are $\mu_k(P, f)$ partial monotone paths ending at elements of $S_{k-1} \setminus S_k = \{v_k\}$, and $d_k\mu_k(P, f)$ partial monotone paths ending at elements of $S_k \setminus S_{k-1}$ (i.e. the d_k successors of v_k). So $\mu(P, f) = S_{n-1} = \Sigma_{k=2}^{n-1}(d_k-1)\mu_k(P, f)+|S_1| = \Sigma_{k=2}^{n-1}(d_k-1)\mu_k(P, f)+(d_1-1)\mu_1(P, f)+\mu_1(P, f) = 1 + \Sigma_{k=1}^{n-1}(d_k - 1)\mu_k(P, f)$, which proves the formula. □

The above theorem can be used to deduce the maximum possible value of $\mu(P, f)$ over all directed 3-polytopes on n vertices. Recall that the Tribonacci sequence is defined as $T_1 = T_2 = 1$, $T_3 = 2$, and $T_n = T_{n-1} + T_{n-2} + T_{n-3}$ for $n > 2$.

Theorem 15. *The maximum number of monotone paths for a directed* 3*-polytope over n vertices is* T_n*. This is achieved by the stacked polytope* $X(n)$*.*

Proof. We proceed by induction. The claim is obviously true for $n = 4$. Suppose now that the claim holds for every natural number less than n, and let (P, f) be a simple directed polytope on n vertices. Note that every partial monotone path ending at v_k is a monotone path on the convex hull of $v_1, ..., v_k$ directed by f, so, by the induction hypothesis, $\mu_k(P, f) \leq T_k$ $\forall k \in [n-1]$. Thus, by Theorem 14, we have $\mu(P, f) \leq 1 + \Sigma_{k=1}^{n-1}(d_k - 1)T_k$.

To bound the right-hand side, first observe that, for any $k \in [n-1]$ $\Sigma_{i=1}^{k}d_{n-i} \leq 3k - 3$, since this quantity counts the number of edges of the planar subgraph of $G(P)$ with vertices $v_{n-k}, ..., v_n$. We can use a telescopic sum and this inequality, along with the convention that $T_0 := 0$, to obtain a bound as follows: $\Sigma_{k=1}^{n-1}T_k = \Sigma_{k=1}^{n-1}(\Sigma_{i=1}^{k}d_{n-i})(T_{n-k} - T_{n-k-1}) \leq T_{n-1} - T_{n-2} + (3k - 3)\Sigma_{k=2}^{n-1}(T_{n-k} - T_{n-k-1}) = T_{n-1} + 2T_{n-2} + ... + 3T_1 = \Sigma_{k=2}^{n}T_k$.

We conclude that $\mu(P, f) \leq 1 + \Sigma_{k=1}^{n-1}(d_k - 1)T_k = 1 + \Sigma_{k=1}^{n-1}d_kT_k - \Sigma_{k=1}^{n}T_k \leq T_n$, which concludes the induction.

To finish the proof, it is easy to verify that the stacked polytope $X(n)$ has exactly T_n monotone paths. $\square$

For the minimum number of monotone paths over directed 3-polytopes on n vertices, we have:

Theorem 16. *The minimum number of monotone paths on a 3-dimensional polytope (P, f) on n vertices is equal to $[\frac{n}{2}] + 2$. This is achieved by orthogonal prisms, when n is even, and by orthogonal wedges of polygons over a vertex, when n is odd.*

Proof. We have $\mu(P, f) \geq 1 + \Sigma_{k=1}^{n-1}(d_k - 1) = \Sigma_{k=1}^{n-1} d_k - n + 2$. Since $\Sigma_{k=1}^{n-1} d_k$ is the number of edges of $\omega(P, f)$, which is bounded below by $[\frac{3n}{2}]$, we obtain $\mu(P, f) \geq [\frac{n}{2}] + 2$. It is straightforward to prove that this is realised as asserted in the claim. $\square$

Matters are less simple in higher dimensions. On one hand, it is easy to deduce the maximum number of monotone paths:

Theorem 17. *The maximum number of monotone paths in a directed d-polytope (P, f) on n vertices is equal to 2^{n-1}. This is achieved by 2-neighbourly polytopes.*

Proof. Since monotone paths can be mapped injectively to sets of vertices, they are at most 2^{n-1}. The map becomes surjective for 2-neighbourly polytopes. $\square$

On the other hand, the minimum number is harder to calculate. Naturally, we can operate in the exact same manner as in Theorem and obtain a lower bound of $[\frac{[(d-2)n]}{2}] + 2$, but this bound is not expected to be optimal.

Question 5. *What is the minimum number of monotone paths on a directed d-polytope on n vertices?*

We close this chapter by kicking off the next one. Athanasiadis, De Loera and Zhang conjectured the following upper bound on $\mu(P, f)$ in terms of the Fibonacci numbers ($F_1 = F_2 = 1$):

Conjecture 1. *The least number of monotone paths encountered in a simple directed 3-polytope on $2n$ vertices is $F_{n+2} + 1$, where F_n is the n^{th} Fibonacci number.*

Chapter 5

Maximum Number of Monotone Paths in Simple 3-Polytopes

5.1 Introduction

In this chapter we prove the following intermediate result towards Conjecture 1. Let (P, f) be a simple directed 3-polytope. A *non-extremal triangle* is a triangular face which contains neither the sink nor the source of $\omega(P, f)$. Theorem 18 states that in order to prove Conjecture 1, it suffices to show that simple directed 3-polytopes with maximum number of monotone paths do not contain non-extremal triangles. We call this "the triangle condition".

Theorem 18. *For the set $S(n)$ of simple directed 3-polytopes on $2n$ vertices with the maximum number of monotone paths, exactly one of the following two statements holds:*

i) there is a member of $S(n)$ which has a non-extremal triangle
ii)the only member of $S(n)$, up to order equivalence, is the staircase wedge, and it has exactly $F_{n+2} + 1$ monotone paths.

This chapter is structured as follows. In Section 2 we provide definitions of important concepts and we mention some useful results. In Sections 3 and 4 we proceed with the proof of Theorem 1. Specifically, we assume the triangle condition for $S(n)$ and prove proposition (ii). Finally, in Section 5 we examine the event that the triangle condition does not hold for $S(n)$, i.e. that there are simple directed 3-polytopes with maximum number of monotone paths which contain non-extremal triangles. We prove results about the possible structure of these polytopes.

5.2 Basic Definitions and Preliminaries

Let (P, f) be a directed polytope.

The number of monotone paths which contain a set of edges E is denoted $\mu^E(P, f)$ (if E is a singleton, the curly bracket notation can be omitted).

Remark 3. *Observe that $\mu(P, f) = \mu(P, -f)$ and, in general, $\mu^E(P, f) = \mu^E(P, -f)$ for every set of edges E.*

Suppose that (P, f) is a simple directed 3-polytope. We say that a vertex v_i of (P, f) is of type A (denoted $v_i \in A$) if $d_i = 2$, and of type B (denoted $v_i \in B$) if $d_i = 1$. Observe that, with the exception of the top and bottom vertices, every vertex is of one of these types.

A useful corollary of Theorem 14 for simple polytopes is the following:

Corollary 7. *If (P, f) is a simple directed 3-polytope, $\mu(P, f) = 3 + \Sigma_{v_k \in A} \mu_k(P, f)$.*

Lemma 3. *Let (P, f) be a simple directed 3-polytope and let I be an ideal of $\omega(P, f)$. Then $|A \cap I| \geq |B \cap I|$.*

Proof. If I is the vertex set of $\omega(P, f)$, the proof is simply a matter of equating the sum of the in-degrees and the sum of the out-degrees of $\omega(P, f)$, so let us assume that I is a proper subset of the vertex set. Consider the subgraph H of $\omega(P, f)$ induced by I. Suppose that it contains l vertices of type A and m vertices of type B. The sum of the in-degrees of H equals $l+2m$. The sum of the out-degrees equals at most $2l+m$ (we have taken into account that v_1 contributes up to 3 edges, but also that, by Steinitz's Theorem, there are at least 3 edges between vertices of H and $\omega(P, f) \setminus H$). Hence, $l + 2m \leq 2l + m \Rightarrow l \geq m$. □

Corollary 8. *Let $V_k := \{v_1, ..., v_k\}$. Then $|A \cap V_k| \geq |B \cap V_k|$*

An *extremal triangle* is a triangular face which contains either the top or the bottom vertex of (P, f). A triangular face which is not extremal is called *non-extremal.*

The set of directed polytopes on $2n$ vertices with maximum number of monotone paths is denoted $S(n)$. We say that the *triangle condition* holds for $S(n)$ if there is no simple directed 3-polytope with $2n$ vertices with a non-extremal triangle which has the maximum number of monotone paths.

Let Q be a polygon with its vertices cyclically ordered $(v_1, ..., v_n)$. The *wedge* W of Q is obtained from the prism $Q \times [0, 1]$ of Q by collapsing the face $[v_1, v_n] \times [0, 1]$ to $[v_1, v_n] \times \{0\}$. The *staircase wedge* is a directed wedge (W, f) such that:

- $\forall i \in [n]$, $j \in \{0, 1\}$, $f(v_1 \times \{0\}) \leq f(v_i \times \{j\}) \leq f(v_n \times \{0\})$,
- $f(v_i \times \{0\}) < f(v_i \times \{1\})$ for every odd i, and
- $f(v_i \times \{0\}) > f(v_i \times \{1\})$ for every even i,

whenever the above vertices are defined.

A staircase wedge is illustrated in Figure 5.1.

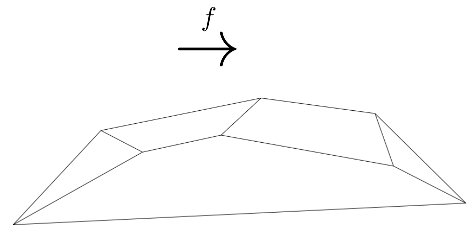

Figure 5.1: A staircase wedge on eight vertices.

5.3 Forbidden Faces

To prove Theorem 18, we will require one more lemma. In Sections 3 and 4, (P, f) is a simple directed 3-polytope on $2n$ vertices with the maximum number of monotone paths, and the triangle condition is assumed to hold for $S(n)$.

Let G be a plane digraph, $e := (u_1, u_2), e_1 := (u_1, v_1), e_2 := (v_2, u_2) \in E(G)$ distinct, $F_1, F_2 \supset e$ distinct faces of G, $e_1 \subset F_1$, $e_2 \subset F_2$. An *exchange transform* deletes e_1, e_2 and replaces them with (v_2, u_1) and (u_2, v_1). An example of an exchange transform is illustrated in Figure 5.2.

Lemma 4. *(P, f) does not have any forbidden faces, i.e. faces with at least four vertices, all of which are of type A except for one, and which do not include v_1 or v_{2n}.*

Proof. Let F be a high forbidden face in (P, f), i.e. there is no forbidden face the bottom vertex of which is higher than the bottom vertex of F. We will first prove the result in the case that the top vertex and the bottom vertex of F are not adjacent.

We can perform an exchange transform, depicted in Figure 5.2, between the top vertex u of F and a predecessor v of u, to obtain a 3-regular plane digraph G. We will show that there is a choice of v (from the two available options) such that G is a 3-polytopal digraph. Let us begin the proof by making an arbitrary choice of v, and producing the corresponding version of G.

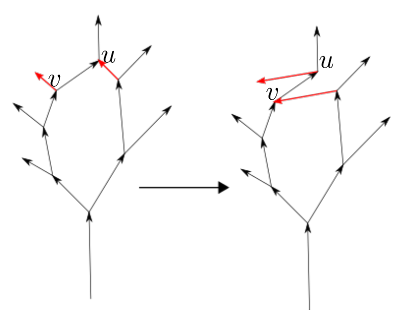

Figure 5.2: The exchange transform between u and v.

Evidently, G is planar, acyclic, and has a unique source and a unique sink.

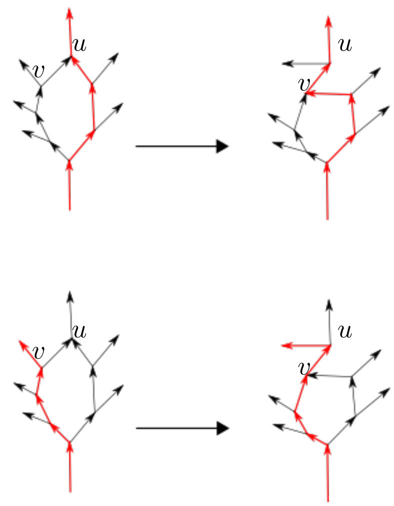

Figure 5.3: A diagram of the two changes to which the monotone path M_1 may be subject. The edges of M_1 are shown in red.

Additionally, G contains three pairwise vertex-independent monotone paths. Indeed, $\omega(P, f)$, being a 3-polytopal digraph, has three pairwise vertex-independent monotone paths, say M_1, M_2, M_3, and only one of them can intersect F, say M_1, because a path which intersects F must necessarily pass through the bottom vertex of F. Moreover, $M_1 \cap F$ is connected, because the only entry point in F for a directed path is the bottom vertex, which therefore belongs to every connected component of $M_1 \cap F$. In addition, $M_1 \cap F$ is contained in one of the two maximal directed paths of F. Hence, to obtain three pairwise vertex-independent monotone paths in G, one first needs to consider three pairwise vertex-independent monotone paths of $\omega(P, f)$, retain $M_1 \cup M_2 \cup M_3 \setminus F$, retain $M_1 \cap F$ if it contains neither u nor v or traverses (v, u), and alter it as depicted in Figure 5.3, otherwise.

Finally, G is 3-connected.

Truly, if we delete from G two vertices $x, y \notin \{u, v\}$, then $G \setminus \{x, y\}$ is isomorphic to the graph $\omega(P, f) \setminus \{x, y\}$ after an exchange transform, if neither of x, y is incident to an exchanged edge, or after an "edge slide", if exactly one of x, y is incident to an exchanged edge, or unaltered, otherwise. However, $\omega(P, f)$ is 3-polytopal, hence 3-connected, so $\omega(P, f) \setminus \{x, y\}$ is connected, and exchange transforms and edge slides preserve connectedness.

If we delete $\{u, v\}$, then every vertex in $G \setminus \{u, v\}$ still connects to either v_1 or v_{2n} through a partial monotone or partial antitone path, respectively, and v_1 connects to v_{2n} through a monotone path (since there are three vertex-independent monotone paths connecting v_1 to v_{2n} in G), hence $G \setminus \{u, v\}$ is connected.

Finally, if we delete a vertex $x \in \{u, v\}$ and a vertex $y \notin \{u, v\}$, then in

$G \setminus x$ every vertex connects to v_{2n} through a partial antitone path, and every vertex except perhaps those in $\uparrow_G x$ connects to v_1 through a partial monotone path. Hence, if x and y are incomparable in the order induced by G, or if $y <_G x$ and $y \neq v_1$, then every vertex in $G \setminus \{x, y\}$ connects to either v_1 or v_{2n} through a partial monotone path or a partial antitone path, respectively, and v_1, v_{2n} connect through a monotone path, so $G \setminus \{x, y\}$ is connected. If $y = v_1$, then every vertex in $G \setminus \{x, y\}$ connects to v_{2n} through a partial antitone path, so $G \setminus \{x, y\}$ is connected. So, suppose that $x <_G y$, and that $G \setminus \{x, y\}$ is disconnected.

If a minimal vertex w of $(x, y)_G$ is of type B, then it is adjacent to a vertex outside of $\uparrow_G x$, which means that the vertices in $\uparrow_G w$ connect to v_1 through a partial monotone path. Similarly, if a maximal vertex z of $(x, y)_G$ is of type A, then $y \neq v_{2n}$ and z is adjacent to a vertex outside $\downarrow_G y$, which implies that the vertices in $\downarrow_G z$ connect to v_{2n} through a partial antitone path. There are at most two minimal vertices and at most two maximal vertices in $(x, y)_G$, so $(x, y)_G$ has at most two connected components. If both of them have either a minimal vertex of type B or a maximal vertex of type A, then every vertex in $G \setminus \{x, y\}$ connects to either v_1 or v_{2n}, which in turn connect through a monotone path, so $G \setminus \{x, y\}$ is connected, a contradiction. So, $(x, y)_G$ contains at least one vertex of type A (in particular, one that is minimal) and at least one vertex of type B (in particular, one that is maximal).

Next, observe that $(x, y)_G$ must be connected. Indeed, suppose otherwise. Then $(x, y)_G$ has two minimal vertices, say w and w'. Let z be a vertex that is minimal among vertices of type B in $(x, y)_G$. Then z would be either in $\uparrow_G w$ or in $\uparrow_G w'$ (but not both, or else $(x, y)_G$ would be connected), and so each lower bound of the predecessors of z would be greater than or equal to either w or w'. Thus, we would obtain either a non-extremal triangle or a forbidden face in $(x, y)_G$, so in $(x, y)_P$ as well, since the exchange transform does not alter this part of the graph. This contradicts our assumption that the triangle condition holds and that F is a high forbidden face.

Next, since $(x, y)_G$ is connected, and $G \setminus \{x, y\}$ has been assumed to be disconnected, $(x, y)_G$ is a connected component of $G \setminus \{x, y\}$, as otherwise $(x, y)_G$ would connect to either v_1 or v_{2n} through a partial monotone path or a partial antitone path, respectively, and so does every other vertex in G, and v_1 and v_{2n} connect through a monotone path, so $G \setminus \{x, y\}$ would be connected, a contradiction. In particular, $(x, y)_G$ being a connected component implies that all the minimal vertices of $(x, y)_G$ are of type A, and all its maximal vertices are of type B.

Since $\{x, y\}$ is a cut set for G producing a connected component $(x, y)_G$, $\{v, u, y\}$ is a cut set for $\omega(P, f)$ producing a connected component $(v, y)_P \cup (u, y)_P$. Now we utilise the freedom to make a choice for v between the two predecessors of u. Let v' be the predecessor of u in $\omega(P, f)$ which is not v. Then there must be a vertex $y' >_P u, v'$ such that $\{v', u, y'\}$ is a cut set for $\omega(P, f)$ producing a connected component $(v', y')_P \cup (u, y')_P$ - for, if not, then we simply opt to perform an exchange transform between u and v' in $\omega(P, f)$, rather than between u and v, and the graph G' that we obtain is 3-connected and otherwise observes all the criteria of the Mihalisin-Klee Theorem, which proves our claim. However, y is comparable with every vertex of $\uparrow_P u$, which includes y'. If $y' \leq_P y$ and z' is minimal among the elements of type B in $(v', y')_P \cup (u, y')_P$, then v' connects to $z' \in (u, y')_P \subseteq (u, y)_P$ through a directed path which does not cross v, u or y, so it connects to the subgraph $(v, y)_P \cup (u, y)_P$ in $\omega(P, f) \setminus \{v, u, y\}$, despite not belonging to it, which contradicts the fact that $(v, y)_P \cup (u, y)_P$ is a connected component of $\omega(P, f) \setminus \{v, u, y\}$. Hence, $y' >_P y$. In an entirely symmetric manner, we obtain $y >_P y'$. We have reached a contradiction, and we therefore conclude that either G or G' (say, G without loss of generality) is 3-connected, hence a 3-polytopal digraph.

Moreover, $\mu(G) \geq \mu(P, f) + \mu^{v'}(P, f) > \mu(P, f)$. Hence, there is a simple directed 3-polytope which has a greater number of monotone paths than (P, f), a contradiction.

We have concluded the case that the top vertex and the bottom vertex of F are not adjacent in $\omega(P, f)$.

Now suppose instead that the top vertex and the bottom vertex of F are adjacent in $\omega(P, f)$. If the exchange transform between u and v produces a 3-connected, hence polytopal digraph, we are done. Otherwise, we compare the partial antitone paths of u and v. If $\mu^u(P, -f) \leq \mu^v(P, -f)$, then we simply reverse the orientation of the edge (v, u). The resulting graph G obeys all of the Mihalisin-Klee criteria, has at least as many monotone paths as $\omega(P, f)$, and since v is not adjacent to the bottom vertex of F (because F has at least four vertices) we revert to the first case and reach a contradiction.

If $\mu^u(P, -f) > \mu^v(P, -f)$, then we perform the operation depicted in Figure 5.4. This again produces a digraph G which obeys the requirements of the Mihalisin-Klee Theorem. Indeed, planarity, acyclicity, and uniqueness of source and sink are straightforward. If one of three vertex-independent monotone paths of $\omega(P, f)$ passes from (v, u), say M_1, we can easily change its intersection with F (Figure 5.5); there is no danger that this new monotone path will intersect one of the other two monotone paths, since those do not intersect F. Finally, G is 3-connected. Indeed, let $\{a, b\}$ be a cut set of size 2 for G, then it is also a cut set for $H := \omega(P, f) \setminus (v, u)$. Since $\omega(P, f)$ is 3-connected, v, u must be in distinct connected components of $H \setminus \{a, b\}$, or $\omega(P, f) \setminus \{a, b\}$ would also be disconnected. Since in H there are three vertex-independent paths between the successors s and t of u and v respectively, as illustrated in Figure 5.6, at least one of these successors belongs to $\{a, b\}$; without loss of generality, say that $a = s$. Then b lies in the lower path between s and the successor of v in Figure 5.6. The only choice of b which disconnects $H \setminus a$ is the bottom vertex of F in $\omega(P, f)$. However, this choice of the pair $\{a, b\}$ does not disconnect G, a contradiction.

This 3-polytopal, 3-regular digraph G has the same number of monotone paths as (P, f), but it also contains a non-extremal triangle, which contradicts the triangle condition.

We conclude that, under the triangle condition, simple directed 3-polytopes which exhibit maximum number of monotone paths do not have forbidden faces. □

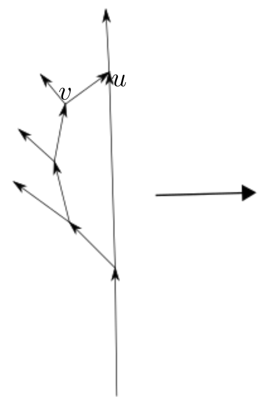

Figure 5.4: This transformation deletes the edge (v, u), smooths out v, and adds a new edge which preserves the number of monotone paths and forms a triangle.

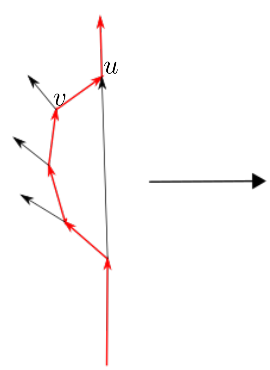

Figure 5.5: The edges of M_1 are in red.

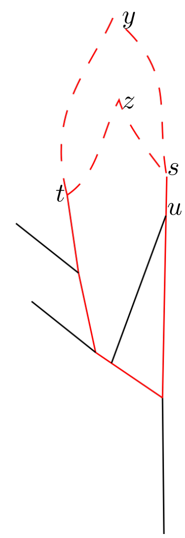

Figure 5.6: The three vertex-independent paths between s and t. The upper two paths do not intersect, otherwise $\omega(P, f)$ would have a cut set of size 2.

Remark 4. *Since $\mu(P, -f) = \mu(P, f)$, Lemma 4 implies more generally that, under the triangle condition, in a simple directed 3-polytope with maximum number of monotone paths, every face contains at least two vertices from each type.*

5.4 Main Proof

Now, to prove Theorem 18. By Corollary 7, we have that $\mu(P, f) = 3 + \Sigma_{v_k \in A} \mu_k(P, f)$. We seek to calculate this expression.

Each term $\mu_k(P, f)$ is obtained in one of two ways: it is either equal to a previous term, if $v_k \in A$, or it is the sum of two distinct previous terms, if $v_k \in B$. Thus, the finite sequence $\mu_k(P, f), 1 \leq k \leq 2n$ is composed of an initial term 1, a final term $\mu(P, f)$, and two finite sequences a_k (for type A vertices) and b_k (for type B vertices), each having $n-1$ terms. Additionally, by Corollary 8, the term a_k always comes before the term b_k.

Firstly, we prove that $b_k \leq F_{k+2}$ $\forall k \in [n-1]$. We show this by induction. For $k = 1$, we have $b_1 = 2 \leq 2 = F_3$. Suppose that the claim is true for $k \leq m$. Then it is also true for $k = m + 1$. Indeed:

i) If $b_{m+1} := b_i + b_j$, then $i < j \leq m$, so $b_{m+1} \leq F_{i+2} + F_{j+2} \leq F_{m+3}$.

ii) If $b_{m+1} := a_i + b_j$, then suppose that $j = m$. Then it is impossible to have a finite sequence $a_i := a_i' := a_i'' := ... := b_m$, because then the vertex corresponding to the term b_m would have total degree of at least 4, a contradiction. Hence, $a_i = b_r$ for $r < m$ or $a_i = 1$. In either case, we have $b_{m+1} \leq F_{m+1} + F_{m+2} = F_{m+3}$. On the other hand, if $j < m$, then again we have $b_{m+1} \leq F_{m+2} + F_{m+1} = F_{m+3}$.

iii) If $b_{m+1} := a_i + a_j$, then there do not exist finite sequences $a_i := ... := b_m$ and $a_j := ... := b_m$, since this would imply that (P, f) contains a forbidden face or a non-extremal triangle, which is impossible by Lemma 4 and the triangle condition. Hence, at least one of a_i, a_j is equal to b_r with $r < m$. But then $b_{m+1} \leq F_{r+2} + F_{m+2} \leq F_{m+3}$.

iv) If one of the summants in the definition of b_{m+1} is the initial term 1, then $b_{m+1} \leq F_{m+2} + 1 \leq F_{m+3}$.

This concludes the induction. We have shown that, if (P, f) maximizes the number of monotone paths among simple directed 3-polytopes on $2n$ vertices, then $b_k \leq F_{k+2}$ for every $k \in [n-1]$.

By Corollary 8, we have that a_k precedes b_k in the finite sequence $\mu_k(P, f)$. This implies that $a_k = b_r$ for $r < k$, hence $a_k \leq F_{r+2} \leq F_{k+1}$ $\forall k \in [n-1]$. Then $\mu(P, f) = 2 + 1 + \Sigma_{v_k \in A} \mu_k(P, f) \leq 2 + \Sigma_{k=1}^{n} F_k = 2 + F_{n+2} - 1 = F_{n+2} + 1$.

We have shown that, under the triangle condition, $F_{n+2}+1$ is the maximum number of monotone paths in a simple directed 3-polytope on $2n$ vertices. Additionally, this bound can be achieved by staircase wedges. Indeed, in staircase wedges, vertices of type A and type B alternate when ordered by f-value, and we have $a_1 := 1$, $a_k := b_{k-1}$ for $k > 1$, $b_1 := 1+a_1$, and $b_k := a_k+a_{k-1} = b_{k-1}+b_{k-2}$ for $k > 1$. From the above recursion we obtain $b_k = F_{k+2}$ $\forall k \in [n-1]$ and $\mu(W, f) := b_{n-1} + a_{n-1} + 1 = F_{n+1} + F_n + 1 = F_{n+2} + 1$.

It only remains to prove that staircase wedges are unique in this aspect. Indeed, under the triangle condition, for any simple directed 3-polytope with maximum number of monotone paths we obtain a sequence with $a_k = F_{k+1}$ $\forall k \in [n-1]$. Since each term of type A is unequal to all the previous ones, the "roots" of the defining sequences of these terms must be distinct. The first such root is the term 1, and, by Corollary 8, it is followed by a_1, which expends the term 1 as a root for a type A term. Hence, a_2 will require a new root, so a_1 must be followed by b_1. Then, by Corollary 8, a_2 follows b_1, expending it. By repeating this argument indefinitely, we deduce that the terms a_k and b_k alternate in the sequence, and that $a_k := b_{k-1}$, so $b_k = F_{k+2}$, $\forall k \in [n]$.

As far as adjacencies for terms of type B are concerned, for $k = 1$, we necessarily have $b_1 := 1 + a_1$. For greater k, there are only four two-term sums which yield the values F_{k+2} for b_k. If $b_k := b_{k-1} + b_{k-2}$, then we obtain degrees of 4 for the vertices which correspond to the terms b_{k-1} and b_{k-2}, which is absurd. The same contradiction is encountered for the sums $b_{k-1} + a_{k-1}$ and $b_{k-2} + a_k$. Hence, $b_k := a_k + a_{k-1}$. Finally, v_{2n} is adjacent to the vertices corresponding to a_{n-1}, b_{n-1}, and the initial term 1.

We have shown that, under the triangle condition, every simple directed 3-polytope (P, f) on $2n$ vertices with maximum number of monotone paths has the same digraph as the staircase wedge on $2n$ vertices. However, this polytopal digraph is weakly connected, thus, by Theorem 1, it determines a unique simple directed polytope up to order equivalence, namely the staircase wedge.

5.5 Considering Non-extremal Triangles

In order to prove Conjecture 1, it suffices to prove that simple directed 3-polytopes bearing non-extremal triangles which maximize the number of monotone paths do not exist. In this section, we mention some basic statements which hold for this (possibly null) class V of "vertebrate" polytopes. These statements have double purpose. On one hand, they provide additional peculiarities to the objects of V. If we gather enough of those, we may spot an inconsistency which will collapse V to nothingness. On the other hand, their contrapositives provide conditions which are useful for discerning directed simple 3-polytopes with non-extremal triangles which do not maximise the number of monotone paths.

A *bevelling* around a vertex v in a 3-polytope P is performed by choosing a plane H which separates v from the other vertices, then taking the intersection of P with the half-space of H which does not contain v to receive a 3-polytope Q. If P is simple, then Q is also simple. From a graph perspective, a bevelling around a vertex v of a simple 3-polytope is equivalent to subdividing two edges incident to v and connecting the new vertices with an edge. The inverse operation, *filling*, i.e. removing the affine hull H of a triangular face from the set of planes which bound a given 3-polytope Q, is only defined in the class of polytopes if the intersection of the affine hulls of the surrounding faces lies beyond Q with respect to H.

A bevelling in P around a vertex v which yields Q will be denoted $P \to_v Q$. The inverse filling will simply be denoted $Q_v \leftarrow P$.

If v is of type A, then the non-extremal triangle that we obtain from a bevelling v sprouts an upward edge from its middle vertex, whereas, if v is of type B, the resulting triangle sprouts a downward edge from its middle vertex. We call these triangles of type A and of type B, respectively. Fillings also preserve type.

Now, let $(P, f) \in V$. Then the following results hold.

Lemma 5. *Filling is well-defined over every non-extremal triangle F of (P, f).*

Proof. Let H be the affine hull of F, and let p be the intersection of the affine hulls of the three faces surrounding F. Then p exists and lies beyond H. Indeed, if p does not exist or lies beneath H, then P is contained within a right triangular prism which has F as one of its bases. One of the vertices of F is the top or bottom vertex of the triangular prism, hence of (P, f) as well. This contradicts the assumption that F is a non-extremal triangle. Since p lies beyond H, the filling over F is well-defined. □

Lemma 6. *Consider a bevelling $P \to_p Q$ of a simple directed 3-polytope P and let $e = (p, q)$ be the edge of v which is subdivided to produce the top (if $v \in A$) or the bottom (if $v \in B$) vertex of the obtained triangular face F. Then $\mu(Q, f) = \mu(P, f) + \mu^e(P, f)$.*

Proof. We prove the part for $p \in A$, then apply it to $(P, -f)$ to receive the result for $p \in B$.

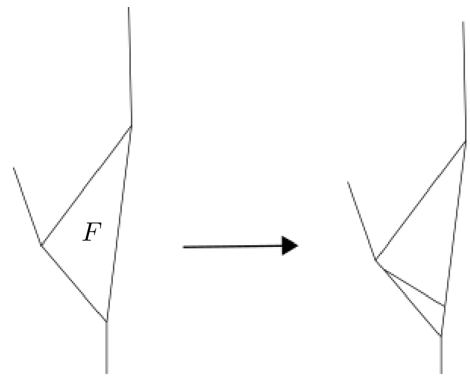

Figure 5.7: The transformation for triangles of type A. The corresponding transformation for triangles of type B can be seen by turning this figure upside-down.

The simple directed 3-polytope (Q, f) has exactly one more edge than (P, f), say e'. The only additional monotone paths in (Q, f) are those which pass from the added edge e', so $\mu(Q, f) = \mu(P, f) + \mu^{e'}(Q, f)$. However, there are $\mu^p(Q, f) = \mu^p(P, f)$ partial monotone paths reaching the lower end of e' and $\mu^q(Q, -f) = \mu^q(P, -f)$ partial antitone paths reaching its higher end, hence $\mu^{e'}(Q, f) = \mu^p(P, f) \cdot \mu^q(P, -f) = \mu^e(P, f)$. □

Theorem 19. *The non-extremal triangles of* (P, f) *all lie upon the same monotone path.*

Proof. Let F, F' be distinct forbidden triangles in (P, f), with the top vertex of F being lower than the top vertex of F'. Then the top vertex of F is less than the bottom vertex of F'. Indeed, suppose otherwise, and consider $P_p \leftarrow P'_{p'} \leftarrow P''$. Let e and e' be the edges of p and p', respectively, in (P'', f) which are subdivided during this transformation to produce the top (for a type A triangle) or the bottom (for a type B triangle) vertex of F and F', respectively. Note that $\mu(P', f) = \mu(P'', f) + \mu^{e'}(P'', f)$ and $\mu(P, f) = \mu(P'', f) + \mu^{e'}(P'', f) + \mu^e(P'', f) + \mu^{\{e,e'\}}(P'', f)$. If the top vertex of F is not smaller than the bottom vertex of F', then $\mu^{\{e,e'\}}(P'', f) = 0$, so $\mu(P, f) = \mu(P'', f) + \mu^{e'}(P'', f) + \mu^e(P'', f)$. Now, suppose that $\mu^e(P'', f) \geq \mu^{e'}(P'', f)$ and consider the bevelling $P' \rightarrow_p Q$ of Figure 5.7. We have $\mu(Q, f) > \mu(P'', f) + 2\mu^e(P'', f) \geq \mu(P, f)$, a contradiction. A similar contradiction is obtained if $\mu^e(P'', f) < \mu^{e'}(P'', f)$. Thus, the top vertex of F must be smaller than the bottom vertex of F'. Since F and F' were arbitrary, we conclude that all the non-extremal triangles of (P, f) lie upon a single monotone path. □

Theorem 20. *The ideal of the bottom vertex* p *of a non-extremal triangle includes at least two of the successors of* v_1*, and the filter of its top vertex* q *includes at least two of the predecessors of* v_{2n}*.*

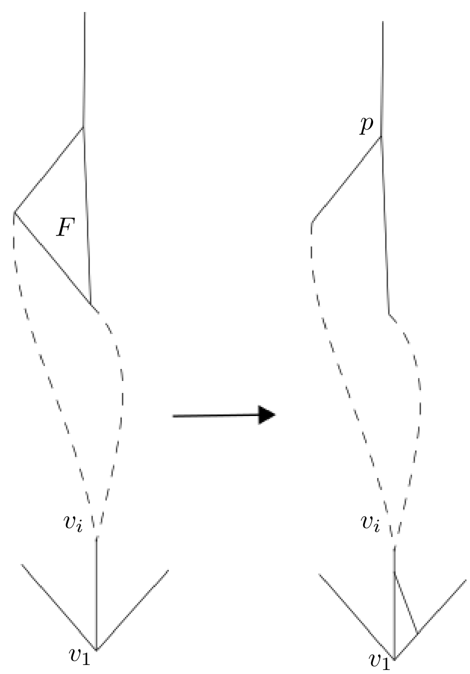

Figure 5.8: A transformation that increases the number of monotone paths.

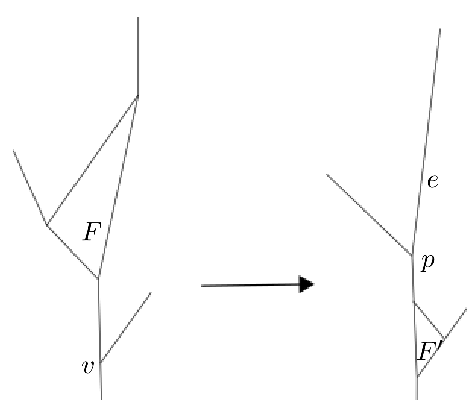

Figure 5.9: Another transformation that increases the number of monotone paths.

Proof. We prove the first part of this result; the second is obtained by applying the first on $(P, -f)$.

Suppose not. Let $v_i \in \downarrow_P p$ be the unique successor of v_1 in the ideal of p. We perform the transformation $P_p \leftarrow P' \rightarrow_{v_1} Q$ as depicted in Figure 5.8. Then $\mu(Q, f) = \mu(P', f) + \mu^{(v_1,v_i)}(P', f) = \mu(P, f) - \mu^e(P, f) + \mu^{(v_1,v_i)}(P', f) > \mu(P, f)$, a contradiction. The last inequality is strict because there is a monotone path which passes from (v_1, v_i) but not from e, namely one which passes from an edge which has exactly one end in $(v_i, p)_P$ (such an edge must exist, or $\{v_i, p\}$ would be a cut set of $\omega(P, f)$). □

Theorem 21. *The top vertex of every triangle of type A initiates at least as many partial antitone paths as the middle vertex. The bottom vertex of every triangle of type B initiates at least as many partial monotone paths as the middle vertex.*

Proof. We prove the first part of this result; the second is obtained by applying the first on $(P, -f)$.

If not, we can obtain a simple directed 3-polytope with more monotone paths by reversing the orientation of the edge connecting the middle to the top vertex through a transformation $P_p \leftarrow P' \rightarrow_p Q$. □

Theorem 22. *The vertex v directly below the bottom vertex of a non-extremal triangle F of type A is of type B. The vertex v directly above the top vertex of a non-extremal triangle F of type B is of type A.*

Proof. We prove the part for triangles of type A, then apply it to $(P, -f)$ to receive the result for triangles of type B.

If v is of type A, then we perform a transformation $P_p \leftarrow P' \rightarrow_v Q$ (Figure 5.9) and we have $\mu(Q, f) > \mu(P, f)$, a contradiction. □

Figure 5.10 illustrates some of these results. It is now apparent why we call the members of V vertebrates.

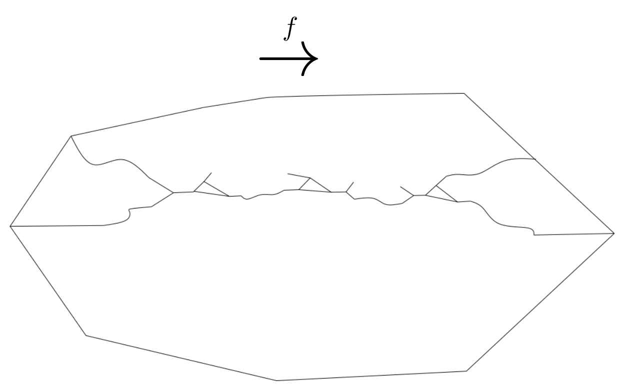

Figure 5.10: Vertebrate 3-polytopes have all of their non-extremal triangles arranged in a "spine".

Chapter 6

Feasible Sets of Outdegrees for Directed Polytopes

In a private communication, Athanasiadis asked whether it is possible for the digraph $\omega(P, f)$ of a directed d-polytope to contain no vertex of outdegree 2. Such a thing is indeed possible: simply let (P, f) be a square pyramid as pictured in Figure 6.1. Moreover, taking the $(d-3)$-fold inverted pyramid over (P, f) yields a directed d-polytope with this property for every $d \geq 4$.

In general, we observe that 0 and 1 are always the outdegrees of the top and the next-to-top vertices, so they can never be avoided. However, every other natural number can be easily avoided by the set of outdegrees of an appropriate directed d-polytope: to avoid $k \geq 3$, simply take the $(d-2)$-fold pyramid over a $(k+1)$-gon.

After establishing that every individual integer greater than or equal to 2 can be avoided as an outdegree, it is not hard to extend this result to finite sets S not containing 0 and 1. Let $m = \max S$, and let (P, f) be the stacked polytope on $m+3$ vertices of Figure 6.2; here, every vertex has an outdegree of 0, 1, $m+1$, or $m+2$. Hence, we obtain a directed 3-polytope which avoids all the numbers of S as outdegrees; moreover, taking the $(d-3)$-fold inverted pyramid over (P, f) yields a directed d-polytope avoiding S for every $d \geq 4$.

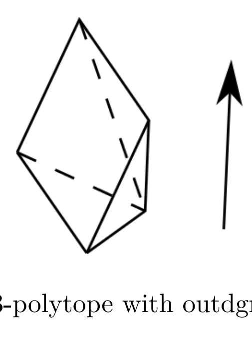

Figure 6.1: A 3-polytope with outdgrees 0, 1, and 3.

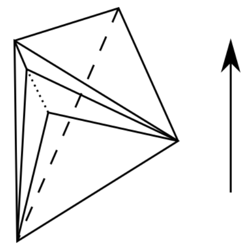

Figure 6.2: A 3-polytope with outdgrees 0, 1, m, and $m+1$.

This brief exercise in "avoidable" sets whets the appetite for the real question: which finite sets of natural numbers are sets of outdegrees for d-polytopes? For $d \geq 3$, we call a set $S \subseteq \mathbb{N}$ *d-feasible* if there exists a directed d-polytope (P, f) such that the set of outdegrees of $\omega(P, f)$ is S. As a convenient notation, we present the elements of a d-feasible set, i.e. the distinct outdegrees present in a polytope, in increasing order: $0 < 1 < D_1 < ... < D_k$. Essentially, what we showed above is simply that every cofinite set of natural numbers has a d-feasible subset.

The following theorems and their proofs appear in detail in [14] (under preparation):

Theorem 23. *Every subset of* $\mathbb{N}$ *containing* 0, 1*, and a number greater than or equal to* 3 *is* 3*-feasible.*

Proof. To construct a 3-polytope with outdegrees $0 < 1 < D_1 < ... < D_k$, start with the spindle depicted in Figure 6.3, in which $d_1 = D_k$, $d_i = 2$ for $2 \leq i \leq n$, $d_i = 1$ for $n+1 \leq i \leq 2n-1$, and $d_{2n} = 0$. Each d_i, $2 \leq i \leq n$ can be increased to 3 by lifting v_i beyond the unique face F_i which is incident to both v_i and v_n. This does not alter any other outdegree. Then, each d_i can be increased further by splitting a triangular face to which it is incident any number of times. By the Mihalisin-Klee Theorem, this yields a 3-polytope. An example for the set $\{0, 1, 2, 3, 5, 6\}$ can be seen in Figure 6.4. □

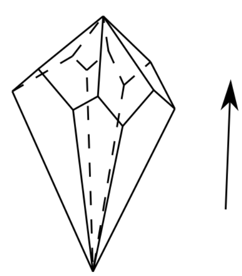

Figure 6.3: The spindle with $d_1 = 6$.

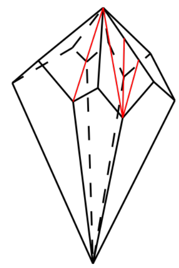

Figure 6.4: The edges added to the initial spindle are noted with red.

By taking a $(d-3)$-fold inverted pyramid over the above construction, we obtain:

Corollary 9. *Every finite set of natural numbers is a subset of a d-feasible set.*

For $m \in \mathbb{N}$, we define the m^{th} triangular number $t_m := \frac{m(m+1)}{2}$.

By performing an inductive construction over cyclic polytopes, we have:

Theorem 24. *Assume that $d > 4$ and $0 < 1 < D_1 < ... < D_k$ is a finite sequence of positive integers. If, for every $i \in \{1, ..., k\}$, $D_i \geq t_{d-2} + (i-1)(d-1) + 2$ (for even d) or $D_i \geq t_{d-2} + (i-1)(d-2) + 3$ (for odd d), then there exists a d-polytope with outdegrees $0 < 1 < D_1 < ... < D_k$.*

Proof. To begin, define $s : \mathbb{N} \to \mathbb{N}$ such that $s(j) = \min\{m \in \mathbb{N} | \frac{(m-1)m}{2} > j\}$.

Consider the directed d-simplex (P, f). It exhibits all the outdegrees from 0 through d exactly once. Let F be the facet which is not incident to the top vertex. We can introduce a new vertex v which is beyond F, within the affine hull of every facet incident to an edge e of our choice, and higher than every vertex of F except for v_d. We choose e to be the edge between v_d and v_{d-2}. It is an easy consequence of the Inductive Construction Theorem that v will connect to all the vertices of F, but the edge e will disappear. Hence, all the outdegrees of all the vertices of F will increase, except for v_{d-2}. Additionally, the d-polytope $P' := conv(P \cup \{v\})$ has a facet F' formed by the vertices $v, v_{d-1}, v_{d-2}, ..., v_1$, which have outdegrees 1, 3, 3, 5, 6, ..., $d+1$. Let us call this process "splitting".

We then split F' by introducing a new vertex v', setting $e' := \{v, v_{d-3}\}$, and producing a new d-polytope $P'' := conv(P' \cup \{v'\})$ with a facet F'' formed by the vertices $v', v_{d-1}, v_{d-2}, ..., v_1$, which have outdegrees 1, 4, 4, 5, 7, ..., $d+2$. We continue repeating this process, splitting every facet $F^{(j)}$ from F through $F^{(t_{d-2}-1)}$ by introducing a new vertex $v^{(j)}$, setting $e^{(j)} := \{v^{(j-1)}, v_{d-1-s(j)}\}$, and producing a new d-polytope $P^{(j+1)} := conv(P \cup \{v^{(j)}\})$ with a facet $F^{(j+1)}$ formed by the vertices $v^{(j)}, v_{d-1}, v_{d-2}, ..., v_1$, which have outdegrees 1, $3+j$, $3+j$, ..., $3+j$, $2+\frac{s(j)(s(j)+1)}{2}$, $4+j+s(j)$, ..., $d+j+1$. Eventually, after all of these splits, every vertex in the resulting polytope $P^{(t_{d-2})}$ has outdegree 0, 1, or $t_{d-2}+2$, and it has a simplicial facet $F^{t_{d-2}}$ which is formed by the vertices $v^{t_{d-2}}, v_{d-1}, v_{d-2}, ..., v_1$ with outdegrees 1, $t_{d-2}+2$, ..., $t_{d-2}+2$. To obtain a polytope with outdegrees 0, 1, r with $r \geq t_{d-2} + 2$, simply expand $F^{t_{d-2}}$ by stacking d-simplices on it, which will increase the outdegrees of the vertices v_1 through v_{d-1} for the required amount, leaving every other outdegree fixed. Hence, every such set $\{0, 1, r\}$ is d-feasible.

To expand this result to set with k elements, we consider the cyclic polytope $C(d+(k-1)(d-1), d)$, if d is even, or $C(d+(k-1)(d-2), d)$, if d is odd. We use the Gale Evenness Criterion to separate the vertices of the cyclic polytope into k simplicial facets $F_1, ..., F_k$, with consecutive facets F_i, F_{i+1} sharing only the vertex $v_{d+(i-1)(d-1)}$ (if d is even) or only the edge $\{v_{d-1+(i-1)(d-2)}, v_{d+(k-1)(d-2)}\}$ (if d is odd). Then we separately treat each F_i with the splitting process described above, until all of the vertices in each F_i have outdegree D_i. The proof is complete. □

Theorem 25. *Let* $0 < 1 < D_1 < ... < D_k$ *be a finite sequence of positive integers. If* $D_i \geq 4 + 3(i-1)$ *for every* $i \in \{1, ..., k\}$*, then there exists a* 4*-polytope with outdegrees* $0 < 1 < D_1 < ... < D_k$*.*

Proof. As above, with the additional trick that in this particular case we can take F to be the facet of the 4-simplex which is not incident to the bottom vertex. Then the splitting process yields a 4-polytope with eight vertices of degrees 0, 1, 1, 1, 4, 4, 4, 4. □

Corollary 10. *For every* $d \geq 3$ *and every finite set* $S \subset \mathbb{N}$*, if* S *contains* 0 *and* 1 *and every other number in* S *is greater than a function* $f(d, |S|) = \Omega(d(|S|+d))$*, then* S *is d-feasible.*

Corollary 11. *The d-feasible sets with exactly three elements for* $d = 3$ *or* $d = 4$ *are exactly those of the form* $\{0, 1, r\}$ *with* $r \geq d$*.*

As an aside, note also that, for any vertex v of a polytope P with total degree k and for every natural number $l \leq k$, there exists a linear functional f such that $deg_{out}(v) = l$ in $\omega(P, f)$. Indeed, it is a simple application of the Hyperplane Separation Theorem that $N(v)$ can be divided by a hyperplane H into two disjoint subsets with l and $k - l$ elements, respectively. Hence, taking $f :=< x, z >$ to be the appropriate $z \perp H$ proves the claim. Thus, for the right choice of f, any outdegree less than or equal to d can be present in the digraph of a directed d-polytope (P, f).

Another intriguing question is this: for a class C of polytopes, a set $S \subset \mathbb{N}$ is called *d-feasible in relation to C* if there exists a directed d-polytope $(P, f) \in C$ such that the set of outdegrees of $\omega(P, f)$ is S. It would be interesting to search for the d-feasible sets in relation to certain important classes. Here is an example:

Remark 5. *The only d-feasible set in relation to the class of simple polytopes is* $\{0, ..., d\}$

Proof. As proven in [18], the h-vector of a directed simple polytope informs us about the number of vertices which have each outdegree. But the h-vector of a simple polytope only has strictly positive elements. □

One might also ask:

Question 6. *Which are the d-feasible sets in relation to the class of simplicial polytopes?*

Bibliography

[1] I. Adler, C. Papadimitriou, and A. Rubinstein. On simplex pivoting rules and complexity theory. *Integer Programming and Combinatorial Optimization Lecture Notes in Computer Science*, page 13–24, 2014.

[2] C. A. Athanasiadis. Ehrhart polynomials, simplicial polytopes, magic squares and a conjecture of stanley. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2005(583):163–174, 2005.

[3] C. A. Athanasiadis, J. A. De Loera, and Z. Zhang. Enumerative problems for arborescences and monotone paths on polytope graphs. *Journal of Graph Theory*, 2021.

[4] D. W. Barnette. On steinitzs theorem concerning convex 3-polytopes and on some properties of planar graphs. *The Many Facets of Graph Theory Lecture Notes in Mathematics*, page 27–40, 1969.

[5] D. W. Barnette. A short proof of the d-connectedness of d-polytypes. *Discrete Mathematics*, 137(1-3):351–352, 1995.

[6] M. Blanchard, J. A. De Loera, and Q. Louveaux. On the length of monotone paths in polyhedra. *SIAM Journal on Discrete Mathematics*, 35(3):1746–1768, 2021.

[7] R. Blind and P. Mani-Levitska. Puzzles and polytope isomorphisms. *Aequationes Mathematicae*, 34(2-3):287–297, 1987.

[8] J. A. Bondy and U. S. R. Murty. *Graph theory with applications.* Macmillan, 1977.

[9] H. Chang, M. Cossarini, and J. Erickson. Lower bounds for electrical reduction on surfaces. *35th International Symposium on Computational Geometry, Leibniz International Proceedings in Informatics*, 129, 2019.

[10] B. Grunbaum, V. Kaibel, V. Klee, and Ziegler G. M. *Convex polytopes.* Springer, 2003.

[11] F. Holt and V. Klee. A proof of the strict monotone 4-step conjecture. *Contemporary Mathematics Advances in Discrete and Computational Geometry*, page 201–216, 1999.

[12] M. Joswig, V. Kaibel, and F. Korner. On the k-systems of a simple polytope. *Israel Journal of Mathematics*, 129(1):109–117, 2002.

[13] G. Kalai. A simple way to tell a simple polytope from its graph. *Journal of Combinatorial Theory, Series A*, 49(2):381–383, 1988.

[14] G. Kontogeorgiou. Outdegrees of directed polytopal graphs. *under preparation*, 2021.

[15] J. Mihalisin and V. Klee. Convex and linear orientations of polytopal graphs. *Discrete and Computational Geometry*, 24(2):421–436, 2000.

[16] A. Ribo Mor, G. Rote, and A. Schulz. Small grid embeddings of 3-polytopes. *Discrete and Computational Geometry*, 45(1):65–87, 2010.

[17] A. Schrijver. *Theory of linear and integer programming.* Wiley, 2011.

[18] G.M. Ziegler. *Lectures on Polytopes.* Springer-Verlag, 1995.