University of California
Los Angeles

Embeddings of Polytopes and Polyhedral Complexes

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mathematics

by

Stedman McDee Wilson

2012

Abstract of the Dissertation

Embeddings of Polytopes and Polyhedral Complexes

by

Stedman McDee Wilson

Doctor of Philosophy in Mathematics

University of California, Los Angeles, 2012

Professor Igor Pak, Chair

When does a topological polyhedral complex (embedded in $\mathbb{R}^d$) admit a geometric realization (a rectilinear embedding in $\mathbb{R}^d$)? What are the constraints on the possible realizations? Two classic results concerning such questions are Fáry's theorem, which states that every planar graph can be drawn in the plane such that each edge is a straight line segment, and Tutte's theorem, which provides necessary and sufficient conditions for embedding a planar graph such that all faces are convex. The present work is motivated largely by the question of whether these types of results generalize to higher dimensions.

We begin by constructing an irrational polytopal complex consisting of 1278 convex 3-polytopes in $\mathbb{R}^3$. The methods of this construction may also be used to produce small topological complexes with no geometric realization, as well as geometric complexes which encode arbitrary point and line configurations. This allows us to prove universality theorems for 3-dimensional polytopal complexes and arrangements. We also investigate geometric realizations of plane triangulations, and describe an explicit algorithm that embeds a plane triangulation with n vertices in a $4n^3 \times 8n^5$ integer grid, in such a way that at each step of the algorithm, the resulting region of the plane is convex. This embedding, by nature of its sequential convexity, may be lifted vertically to a 3-polytope. This process gives a new proof of Steinitz's Theorem for triangulations, and provides a new upper bound on the size of the integer grid necessary to embed the vertices of simplicial 3-polytopes. For certain classes of triangulations, this grid size is subexponential.

The dissertation of Stedman McDee Wilson is approved.

Amit Sahai

Benjamin Sudakov

Bruce Rothschild

Igor Pak, Committee Chair

University of California, Los Angeles

2012

To Mom, Dad, and Hadley.

Table of Contents

List of Figures

Acknowledgments

It is with great pleasure that I thank those people who have been a positive influence in my life, and who helped make this thesis possible. First I would like to thank my research advisor, Igor Pak. As our years working together come to a conclusion, he is no longer just my advisor—he is my friend. He gave me the guidance, inspiration, and encouragement that drove me to be successful in my graduate career. His stern, no-nonsense attitude towards mathematics and research was complemented by his cheerful, easy-going disposition.

I would also like to thank several other faculty members at UCLA, who taught me, wrote reference letters for me, and encouraged me. Among these faculty are Don Blasius, Robert Greene, Thomas Liggett, Bruce Rothschild, Amit Sahai, and Benny Sudakov. Of the faculty outside of UCLA, I would like to thank Matthias Beck, Jesús De Loera, Jeremy Martin, and Russ Woodroofe, for advice and helpful conversations. Special thanks goes to Stanford professor Rick Schoen, who gave me invaluable advice during my graduate school application process. I am also grateful to the staff in the UCLA mathematics department, for always being friendly and helpful. My particular thanks goes to Maggie Albert, Martha Contreras, and Leticia Dominguez.

My fellow graduate student, Jed Yang, deserves my special thanks. We shared uncountably many conversations about mathematics, the universe, and life (in that order). Always willing to listen to my ideas and proofs before I presented them to my advisor, Jed is a true friend. I would also like to acknowledge Pablo Bert, who became my best friend during graduate school, and his parents, for their frequent hospitality. The fun and camaraderie we shared made the good times better and the bad times bearable.

Finally, I would like to express my sincere and total gratitude to my family. My younger brother, Hadley, has always been my biggest advocate, and over the years we have brought out the best in one another. Most importantly, I would like to thank my parents, Clarence and Anna Wilson. They are exceptional parents and role models, and to this day their support has been unbounded.

Vita

1984	Born, Wilmington, North Carolina, USA.
2002, Summer	Cutco Cutlery salesman, Wilmington, NC.
2002-2003	Rower, Stanford University men's crew team.
2003, Summer	Programmer for GE/Global Nuclear Fuels, Wilmington, NC.
2003–2005	Keyboardist, vocalist, and composer for jazz-fusion band *The Institute for the Advancement of Funk and Soul*, Stanford, CA.
2004, Summer	Programmer for project SAUND II, experimental neutrino physics group, Stanford University.
2005, Summer	Programmer for star collapse visualization project, Kavli Institute for Astrophysics and Cosmology, Stanford University.
2006	B.S. Mathematics and Physics (double major), Stanford University.
2007–2012	Teaching Assistant, Mathematics Department, UCLA. Taught courses in calculus, discrete mathematics, computer programming, linear algebra, game theory, analysis, mathematical finance, and probability theory.
2008	M.A. Mathematics, University of California Los Angeles.

Presentations

January 5, 2012. *Strengthening Fary's Theorem: Convex and Star-Shaped Realizations of*

Polyhedral Complexes. Joint Mathematics Meetings, Boston, MA.

January 19, 2012. *Geometric Realizations of Polyhedral Complexes.* UCLA Combinatorics Seminar.

CHAPTER 1

Introduction

The notion of realizing topological configurations geometrically is illustrated elegantly in Fáry's theorem [F], which states that every planar graph can be drawn in the plane such that each edge is a straight line segment (See Figure 1.1). In two influential papers [T1, T2], Tutte first showed necessary and sufficient conditions for realizing 2-connected planar graphs, with all faces (non-strictly) convex. He then showed that for 3-connected planar graphs one can make all faces *strictly convex*. Much of this work was originally motivated by the possibility of extending Tutte's results to 3 and higher dimensions.

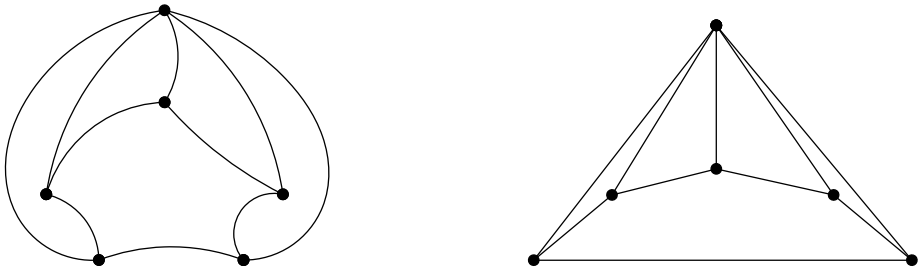

Figure 1.1: A topological plane graph (left) and a geometric embedding of that graph (right).

The starting point for our investigations is *Steinitz's Theorem* (see e.g. [G, P, R, Z1]). It states that every 3-connected planar graph is the graph of a 3-polytope. As a consequence of the proof, all polytopes in $\mathbb{R}^3$ can be realized over $\mathbb{Q}$ (i.e. realized with rational vertex coordinates) by applying small perturbations of the vertices which preserve combinatorial structure (the faces of the polytope). There are several directions into which this result has been shown to have negative analogues:

(1) In $\mathbb{R}^d$, $d \geq 4$, there exist irrational convex polytopes (see [R, RZ]),

(2) There exists an irrational 2-dimensional polyhedral surface in $\mathbb{R}^3$ (see [Z3]),

(3) There exists a 3-dimensional topological simplicial complex that is not geometrically realizable (see [Ca, HZ, K]).

In light of Steinitz's theorem, (1) and (2), it is natural to ask whether every 3-dimensional polytopal complex can be realized over $\mathbb{Q}$. A 3-dimensional polytopal complex is a natural generalization of a Schlegel diagram of a 4-polytope, so this question occupies an intermediate position between Steinitz's theorem and (1). Our first result answers this question in the negative (see below for definitions and notation).

Theorem 1.0.1 *There exists an irrational* 3*-dimensional polytopal complex in* $\mathbb{R}^3$*, consisting of 1278 convex* 3*-polytopes.*

In other words, we show that there exists a configuration of finitely many convex 3-polytopes, attached face-to-face in $\mathbb{R}^3$, which cannot be realized over $\mathbb{Q}$. In particular, in contrast with a single polytope, one cannot perturb (in unison) the vertices of the polytopes to make them all rational. Our construction also shows that Brehm's construction (2) presented in [Z3] can be replaced with a complex of convex 3-polytopes in $\mathbb{R}^3$.

The proof of Theorem 1.0.1 follows the same general approach as (1), going back to Perles's first original construction of an irrational polytope in $\mathbb{R}^8$ (see [G]). We start with an irrational point and line configuration in the plane, and then use polyhedral gadgets to constrain the realization space emulating the configuration. At the end, we explicitly construct an irrational complex consisting of 1277 triangular prisms and one pentagonal pyramid.

Our second result is a variation on (3). There are of course various topological obstructions to embedding an abstract simplicial complex into $\mathbb{R}^d$. Furthermore, a *geometric* embedding is even harder to obtain, even if we assume that we start with a *topological complex*—a complex that is already embedding into $\mathbb{R}^d$. The results in [HZ, K] (see also [AB, Ca, Wi]) rely on topological triangulations whose 1-skeletons contain a nontrivial knot with 5 or fewer

edges (this creates an obstruction to a rectilinear embedding). By a much simplified variation on a construction from the proof of Theorem 1.0.1, we show that if one replaces "simplicial" with "polyhedral", one obtains very small examples of complexes which are not geometrically realizable. These examples are much smaller than those found in [HZ, Wi].

Theorem 1.0.2 *There exists a topological* 3*-dimensional polyhedral complex* X *in* $\mathbb{R}^3$ *with* 8 *vertices and* 3 *polyhedra, that is not geometrically realizable.*

In fact, we may extend this polyhedral complex to a polyhedral subdivision of a ball:

Theorem 1.0.3 *There exists a topological* 3*-dimensional polyhedral complex* X' *in* $\mathbb{R}^3$ *consisting of* 9 *vertices and* 9 *polyhedra, such that* X' *is homeomorphic to a ball, and the complex* X *of Theorem 1.0.2 is a subcomplex of* X'*. In particular,* X' *is not geometrically realizable.*

Heuristically, both (3) and Theorem 1.0.2 say that one cannot possibly extend the Fáry and Tutte theorems to $\mathbb{R}^3$ in full generality. To put both our results into one scheme, we have:

$$\text{topological polyhedral complex} \quad \not\Rightarrow_{\text{thm 1.0.2}} \quad \text{geometric polyhedral complex},$$

$$\text{geometric polyhedral complex} \quad \not\Rightarrow_{\text{thm 1.0.1}} \quad \text{rational polyhedral complex}.$$

Our final result in a positive result complementing Theorems 1.0.1 and 1.0.2. Our actual result in full generality is somewhat involved (see Section 5.2), so we state it here only for simplicial complexes.

We restrict ourselves to simplicial complexes which are homeomorphic to a ball, and which are *vertex decomposable* (see e.g. [BP, Wo]). This is a topological property that implies *shellability*. A simplicial d-ball X is vertex decomposable if either it is a single simplex, or recursively, it has a boundary vertex $v \in \partial X$ such that the the deletion $X \setminus v$ is also a vertex decomposable d-ball. We say that X is *strongly* vertex decomposable if in addition, this vertex v is adjacent to exactly d boundary edges.

Theorem 1.0.4 *Let X be a topological d-dimensional simplicial complex in $\mathbb{R}^d$ that is homeomorphic to a d-ball and strongly vertex decomposable. Then there is a geometric simplicial complex Y in $\mathbb{R}^d$ such that Y is a realization of X.*

This result may seem restrictive, but for $d = 2$ it is equivalent Fáry's theorem. To see this, note first that it suffices to prove Fáry's theorem for triangulations (added edges can be removed later). But in the plane, every triangulation X is vertex decomposable [BP] (see also [FPP] for a short proof). But then, by definition, X is also strongly vertex decomposable, and thus Theorem 1.0.4 is just Fáry's theorem.

We have already mentioned that the original proof of Stenitz's Theorem allows us to realize any 3-polytope with rational vertex coordinates. By clearing denominators, the vertex coordinates can be made integers. A natural question then arises: How large must these integers be? Following the original paper by Onn and Sturmfels [OS] which gave the first nontrivial upper bound on the integer grid size, there have been a series of improvements [R, Ro], leading to the currently best $\exp O(n)$ bound in [BS, RRS].

The idea behind the bounds in [BS, R, RRS, Ro] is as follows. Start with the *Tutte spring embedding* of G with unit weights [T2], and lift it up to a convex surface according to the *Maxwell-Cremona theorem* (see [L, R]). Since Tutte's embedding and the lifting are given by rational equations, this embedding can be expanded to an integer embedding. However, there is only so much room for this method to work, and since the determinants are given by the number of spanning trees in G, the bounds cannot be made subexponential in the case of triangulations. The only other class of graphs for which there is a subexponential bound, is the class of triangulations corresponding to stacked polytopes [DS], which can be embedded into a polynomially-sized grid.

Although there are several interesting proofs of the Steinitz theorem [Z1, Z2], neither seem to simplify in the case of triangulations. We present a proof that follows a similar idea, but in place of the Tutte spring embedding we use an inductive construction. In essence, we construct a *sequentially convex embedding* of 2-connected triangulations, giving a simple

proof of the main result in [T2]. We make our construction quantitative, by doing this on a $4n^3 \times 8n^5$ grid, thus reproving a weak version of the main result in [BR].

We then lift the resulting triangulation directly to a convex surface. The inductive argument allows us to obtain a new type of quantitative bound $\zeta(n) = (500n^8)^{\tau(G)}$ on the height of the lifting. Unfortunately, the parameter $\tau(G)$ here and used in the induction may be linear in n for the case of "paths" or "stacked" triangulations. It is bounded from below by both the diameter of G and the diameter of the dual graph of G. However, in a number of special cases this parameter is sublinear.

Rather than discuss the main theorem (Theorem 7.1.2), we present here our main application. A *grid triangulation* of $[p \times q] = \{1, \ldots, p\} \times \{1, \ldots, q\}$ is a triangulation with all grid points as the set of vertices. These triangulations have a curious structure, and have been studied and enumerated in a number of papers (see [A, KZ, Wo] and references therein). For example, according to [Du], there are $1.999 \cdot 10^{21}$ grid triangulations of $[6 \times 6]$.

Theorem 1.0.5 *Let G be a grid triangulation of $[p \times q]$ such that every triangle fits in an $\ell \times \ell$ subgrid. Then G can be realized as the graph of a convex surface embedded in a grid of size $4(pq)^3 \times 8(pq)^5 \times (500(pq)^8)^{6\ell(p+q)}$.*

For example, for $p = q = \sqrt{n}$, and $\ell = O(1)$, we have a subexponential grid size: $O(n^3) \times O(n^5) \times \exp O(\sqrt{n}\log(n))$.

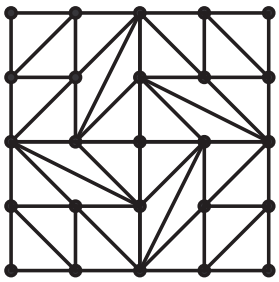

Figure 1.2: A non-regular grid triangulation of $[5 \times 5]$, with $\ell = 3$ (see [KZ]).

We should mention that not every grid triangulation is *regular*. An example found by

Santos (quoted in [KZ]), is shown in Figure 1.2. This means that one cannot embed this triangulation by a direct lifting—another plane embedding of the triangulation is necessary for that.

CHAPTER 2

Basic Definitions and Notation

Given a set $A \subseteq \mathbb{R}^n$, let $\text{conv}(A)$ and $\text{aff}(A)$ denote the convex and affine hulls of A in $\mathbb{R}^n$, respectively. We write $\text{int}(A)$ for the topological interior of A. If A is a manifold then we write ∂A for the manifold boundary of A. Let $B^d = \{x \in \mathbb{R}^d \mid \|x\| = 1\}$ denote the unit d-ball. If two topological spaces A, B are homeomorphic, we shall write $A \sim B$.

A *polytope* P is the convex hull of finitely many points $x_1, \ldots, x_k \in \mathbb{R}^n$. A polytope P is a *d-polytope* if $\text{aff}(P)$ is a d-dimensional affine subspace of $\mathbb{R}^n$. We call a poset X a *geometric d-polyhedron in $\mathbb{R}^n$* if X is the face poset of a d-polytope. By abuse of notation we will refer to a geometric polyhedron as a polytope.

For a geometric d-polyhedron X, an element $F \in X$ is a *face* of X. By a k-face of X we mean a face whose affine hull has dimension k as an affine subspace of $\mathbb{R}^n$. We shall call a 0-face of X a *vertex*, a 1-face an *edge*, a $(d-1)$-face a *facet*, and the d-face the *cell*.

A poset X, ordered by set inclusion, is a *topological d-polyhedron (in $\mathbb{R}^n$)* if there exists a geometric d-polyhedron Y in $\mathbb{R}^n$ and a poset isomorphism $\varphi : X \to Y$ such that for each $F \in X$, $F \subseteq \mathbb{R}^n$ and $F \sim \varphi(F)$. See Figure 2.1. Note that every geometric polyhedron is a topological polyhedron.

A *topological d-polyhedral complex (in $\mathbb{R}^n$)* is a set $X = \bigcup_{i=1}^n X_i$, where each X_i is a topological d-polyhedron in $\mathbb{R}^n$, and such that if $A \in X_i$ and $B \in X_j$ then $A \cap B \in X_i \cap X_j$. We call X a *geometric d-polyhedral complex (in $\mathbb{R}^n$)* if each X_i is a geometric d-polyhedron. We will also refer to a geometric polyhedral complex as a *polytopal complex*. For brevity, we will write *polyhedron* instead of *topological polyhedron* and *polyhedral complex* instead of *topological polyhedral complex*.

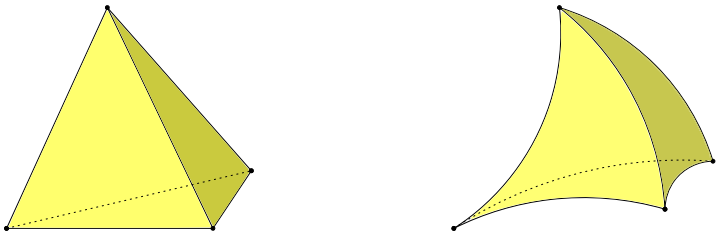

Figure 2.1: Geometric (left) and topological (right) 3-simplices.

We should mention that in Chapter 7, we will use the term *polyhedron* to mean something different—namely a region of $\mathbb{R}^n$ that is the intersection of finitely many closed half-spaces. In this usage, the difference between a polyhedron and a polytope is that a polytope is bounded, whereas a polyhedron may be unbounded. The meaning of *polyhedron* will be clear from context.

We will use the same terminology for the faces $F \in X$ of a polyhedral complex X (vertices, cells, etc.) as for the faces of a polyhedron. We say that a polyhedral complex is a *simplicial* complex if each of its faces is a simplex (of the appropriate dimension). It is worth pointing out that, in the terminology used in more general settings, we are considering here only *pure* polyhedral complexes. That is to say, for a d-polyhedral complex X as defined above, there is a fixed number d such that every face of X is a subset of some d-face of X.

Note that every polyhedron is a polyhedral complex. If $X = \bigcup_{i=1}^{k} X_i$ is a polyhedral complex in $\mathbb{R}^n$, we shall write $\mathcal{P}_X = \{X_1, \ldots, X_k\}$ and $|X| = \bigcup_{F \in X} F$. Note that $|X| \subseteq \mathbb{R}^n$ is the region of $\mathbb{R}^n$ determined by the faces of X. If $F \in X$, we define $\mathcal{P}_X(F) = \{P \in \mathcal{P}_X \mid F \in P\}$. Let $\partial X = \{F \in X \mid F \subseteq \partial X\}$. Note that ∂X is also a polyhedral complex. Two polyhedral complexes are *isomorphic*, written $X \simeq Y$, if they are isomorphic as posets under inclusion.

A *(geometric) realization* of a polyhedral complex X is a geometric polyhedral complex Y such that $X \simeq Y$. We say that Y *(geometrically) realizes* X. If a polyhedral complex has a geometric realization, then we say that it is *(geometrically) realizable*.

A *plane graph* is a planar graph together with a fixed embedding of this graph in $\mathbb{R}^2$. Let $G = (V, E)$ denote a plane graph. By abuse of notation we will identify G with the subset of

$\mathbb{R}^2$ consisting of its vertices and edges. We write $V(G)$ for the vertices of G and $E(G)$ for the edges of G. When G is 2-connected we let $\mathcal{F}(G) = \{F_1, \ldots, F_m\}$ denote the set of bounded faces of G. We define $\mathbf{F}(G) = \bigcup_i F_i$, the region of $\mathbb{R}^2$ determined by G. For a subgraph H of G, we write $H \subseteq G$ (or $H \subset G$, if the inclusion is strict).

When G is 2-connected, a vertex v is called a *boundary vertex* if $v \in \partial\mathbf{F}(G)$, and an *interior* vertex otherwise. Similarly, an edge e is called a *boundary edge* if e is completely contained in the boundary of $\mathbf{F}(G)$, and an *interior edge* otherwise. A *diagonal* of G is an interior edge whose endpoints are boundary vertices of G. For a plane graph G with vertex v, let $G - \{v\}$ denote the plane graph obtained by removing v and all edges adjacent to v.

A *geometric* plane graph is a plane graph for which each edge is a straight line segment. Two plane graphs G, G' are *isotopic*, written $G \sim G'$, if they are homotopic via a homotopy $F : \mathbb{R}^2 \times [0,1] \to \mathbb{R}^2$, and $F(x,t)$ is a homeomorphism for each fixed $t \in [0,1]$. In particular $G \sim G'$ implies that G and G' correspond to the same abstract graph. When $G \sim G'$ and v is a vertex of G, we will write v' for the corresponding vertex of G'. A *geometric embedding* of a plane graph G *in the set* $S \subseteq \mathbb{R}^2$ is a geometric plane graph G' such that $G \sim G'$ and every vertex of G' is a point of S. For a point $u = (a,b) \in \mathbb{R}^2$, we will write $x(u) = a$ and $y(u) = b$ for the standard projections.

CHAPTER 3

An Irrational 3-Polytopal Complex

In this chapter we will work in real projective space $\mathbb{RP}^d$, and we extend the definitions of convex hull and affine hull appropriately. We regard $\mathbb{R}^d \subseteq \mathbb{RP}^d$ under the standard inclusion $(p_1, \ldots, p_d) \mapsto [p_1 : \ldots : p_d : 1]$. We say that distinct points $\mathbf{p}_1, \ldots, \mathbf{p}_n \in \mathbb{RP}^d$ are *collinear* if they are contained in the same line. We say that distinct projective lines $\ell_1, \ldots, \ell_n$ are *concurrent* if $\ell_1 \cap \cdots \cap \ell_n$ is non-empty. If $e_1, \ldots, e_n$ are edges of a polytope and the edge supporting lines $\text{aff}(e_1), \ldots, \text{aff}(e_n)$ are concurrent, we say that the edges $e_1, \ldots, e_n$ are *concurrent*.

An *(abstract) point and line configuration* $\mathcal{L} = ([n], E)$ consists of a finite set $[n] = \{1, \ldots, n\}$, together with a set of *(abstract) lines* $E = \{e_1, \ldots, e_k\}$, where each $e_i \subseteq [n]$. We require that each point is contained in at least 2 lines, and each line contains at least 3 points. A *realization* of $\mathcal{L}$ is a set of points $\Lambda = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\} \subseteq \mathbb{RP}^d$ such that each collection $\{\mathbf{p}_{i_1}, \ldots, \mathbf{p}_{i_k}\}$ of 3 or more points is collinear if and only if $\{i_1, \ldots i_k\} \subseteq e$ for some $e \in E$. A line $\ell \subset \mathbb{RP}^d$ is a *line of* Λ if $\ell \cap \Lambda = \{\mathbf{p}_{i_1}, \ldots, \mathbf{p}_{i_k}\}$ and $\{i_1, \ldots, i_k\} \in E$. A point and line configuration $\mathcal{L}$ is said to be *realizable* over a field F if there is a realization Λ of $\mathcal{L}$ such that each point of Λ has coordinates in F.

A realizable point and line configuration $\mathcal{L}$ is said to be *irrational* if it is not realizable over $\mathbb{Q}$. That is, for every realization Λ of $\mathcal{L}$ there is some point $\mathbf{p} \in \Lambda$ such that $\mathbf{p}$ has an irrational coordinate. The following 9-point configuration due to Perles is irrational.

Lemma 3.0.6 (Perles, [G]) *The point and line configuration depicted in Figure 3.1 is irrational.*

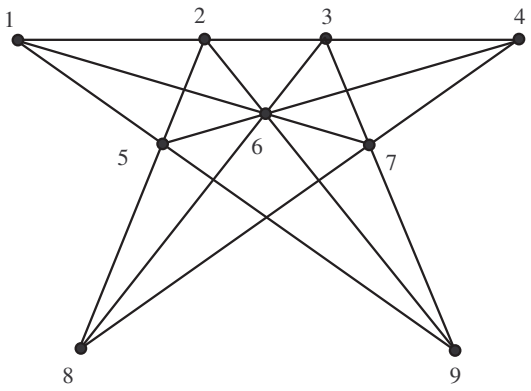

Figure 3.1: The 9-point Perles configuration

We say that a geometric polyhedral complex X *generates* a realization Λ of a point and line configuration $\mathcal{L}$, if each point of Λ is the intersection of affine hulls of faces of X. We say that a polyhedral complex X is *realizable* over a field F if there is a geometric realization X' of X such that each vertex of X' has coordinates in F. Note that *realizable over* $\mathbb{R}$ is equivalent to *realizable*. A geometric polyhedral complex X is called *irrational* if it is not realizable over $\mathbb{Q}$. In this section we will construct a geometric polyhedral complex X such that every realization of X generates the Perles configuration. This implies that X is irrational.

In what follows we let T denote any geometric realization of a triangular prism in $\mathbb{R}^3$. The edges of T not contained in the triangular facets of T are the *lateral* edges of T, and the facets containing these edges (i.e. the tetragonal facets) are the *lateral* facets.

Lemma 3.0.7 *In every geometric realization T of a triangular prism, the lateral edges of T are concurrent (see Figure 3.2).*

Proof. Let ℓ_1, ℓ_2, ℓ_3 denote the supporting lines of the lateral edges of T. These lines are pairwise coplanar. Indeed, $P_{1,2} = \text{aff}(\ell_1 \cup \ell_2)$, $P_{1,3} = \text{aff}(\ell_1 \cup \ell_3)$ and $P_{2,3} = \text{aff}(\ell_2 \cup \ell_3)$ are the supporting planes of the lateral facets of T. Since ℓ_1, ℓ_2 are coplanar projective lines, they intersect in a point $\mathbf{p}$. Thus $\{\mathbf{p}\} = \ell_1 \cap \ell_2 \subseteq P_{1,3} \cap P_{2,3} = \ell_3$. Therefore $\{\mathbf{p}\} = \ell_1 \cap \ell_2 \cap \ell_3$. $\square$

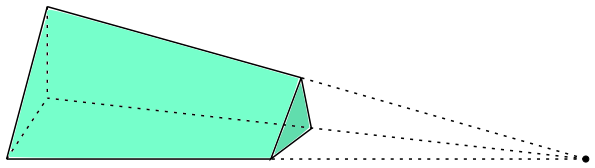

Figure 3.2: A triangular prism, and the point of concurrency of its lateral edges.

A *belt* is a polyhedral complex B consisting of triangular prisms $T_1, \ldots, T_m$ attached consecutively along their lateral facets (see Figure 3.3). We introduce notation that will be used in Lemma 3.0.8 below. Let $B = \bigcup_{i=1}^{m} T_i$ be a belt. For $i = 1, \ldots, m$ and $j = 1, 2, 3$ let $\mathbf{z}_{(i,j)}$ and $\mathbf{z}_{(i,j')}$ denote the adjacent vertices of T_i contained in opposite triangular facets. For $i = 1, \ldots, m$ and $j = 1, 2$ let $F_{(i,j)}$ denote the facet containing the vertices $\mathbf{z}_{(i,1)}, \mathbf{z}_{(i,1')}, \mathbf{z}_{(i,2)}, \mathbf{z}_{(i,2')}$. Then each $F_{(i,k)}$ is a lateral facet of T_i. We assume that the prisms are attached so that $F_{(i,2)} = F_{(i+1,1)}$ for all $i = 1, \ldots, m-1$.

When attached to other polytopes, a belt forces concurrency of the edges along which it is attached. This is ensured by the following lemma. It will be a crucial ingredient in the proofs that follow.

Lemma 3.0.8 *For every belt $B = \bigcup_{i=1}^{m} T_m$, the set of all lateral edges of all prisms T_i is a set of concurrent lines.*

Proof. Let $\ell_{(i,j)} = \mathrm{aff}(\mathbf{z}_{(i,j)}, \mathbf{z}_{(i,j')})$ denote the supporting line of the j^{th} lateral edge of T_i. For each $i = 1, \ldots, m$, by Lemma 3.0.7 there is some $\mathbf{p}_i \in \mathbb{R}^3$ such that $\{\mathbf{p}_i\} = \ell_{(i,1)} \cap \ell_{(i,2)} \cap \ell_{(i,3)}$. Let $i \in \{1, \ldots, m-1\}$. Then $\{\mathbf{p}_i\} = \ell_{(i,1)} \cap \ell_{(i,2)} \cap \ell_{(i,3)} = \ell_{(i,1)} \cap \ell_{(i,2)} = \ell_{(i+1,1)} \cap \ell_{(i+1,2)} = \ell_{(i+1,1)} \cap \ell_{(i+1,2)} \cap \ell_{(i+1,3)} = \{\mathbf{p}_{i+1}\}$. Thus

$$\{\mathbf{p}_1\} = \cdots = \{\mathbf{p}_m\} = \bigcap_{1 \le i \le m,\, 1 \le j \le 3} \ell_{(i,j)}.$$

□

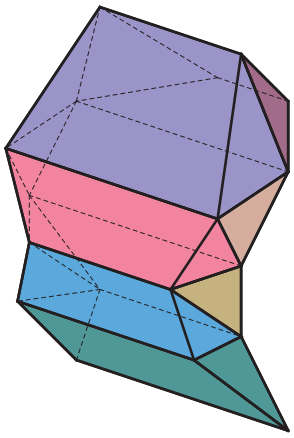

Figure 3.3: A belt of 7 prisms.

Proof of Theorem 1.0.1. For $k = 1, \ldots, 5$, let

$$\mathbf{x}_k = \left(\cos\left(\frac{2\pi k}{5}\right), \sin\left(\frac{2\pi k}{5}\right), 0\right) \quad \text{and} \quad \mathbf{x}'_k = \left(\cos\left(\frac{2\pi k}{5}\right), \sin\left(\frac{2\pi k}{5}\right), 2\right).$$

Note that the convex hull of the $\mathbf{x}_k$ and $\mathbf{x}'_k$ is a regular pentagonal prism, call it R. Let

$$\mathbf{a} = \left(0, 0, 1 - \frac{1}{\sqrt{5}}\right) \quad \text{and} \quad \mathbf{a}' = \left(0, 0, 1 + \frac{1}{\sqrt{5}}\right).$$

We subdivide R into 7 polytopes as follows. Let M denote the pentagonal pyramid with base vertices $\mathbf{x}_k$ and apex $\mathbf{a}$, and M' the pentagonal pyramid with base vertices $\mathbf{x}'_k$ and apex $\mathbf{a}'$. For $i = 1, \ldots, 5$, let T_i denote the triangular prism whose two triangular faces F_i and F'_i consist of the vertices $\mathbf{a}, \mathbf{x}_i, \mathbf{x}_{i+1}$ and $\mathbf{a}', \mathbf{x}'_i, \mathbf{x}'_{i+1}$, respectively, where addition is modulo 5. Note that the prisms T_i all share the lateral edge $\text{conv}(\mathbf{a}, \mathbf{a}')$. Now remove the "top" pentagonal pyramid M'. The result is a geometric polyhedral complex K consisting of 6 geometric polyhedra. We call K the *core* (see Fig. 3.4). It forms the centerpiece of our construction.

For $i = 1, \ldots, 5$, let $e_i = \text{conv}(\mathbf{x}_i, \mathbf{x}_{i+1})$ denote the edges of the base of the pentagonal prism M, let $e'_i = \text{conv}(\mathbf{x}'_i, \mathbf{a}')$ denote the edges on the top of K containing $\mathbf{a}'$, and let $\ell_i = \text{aff}(e_i)$ and $\ell'_i = \text{aff}(e'_i)$. Let $P_i = \text{aff}(T_i \cap T_{i-1})$, and let $P_b = \text{aff}(\mathbf{x}_1, \ldots, \mathbf{x}_5)$ denote the supporting plane of the base of the pentagonal pyramid M.

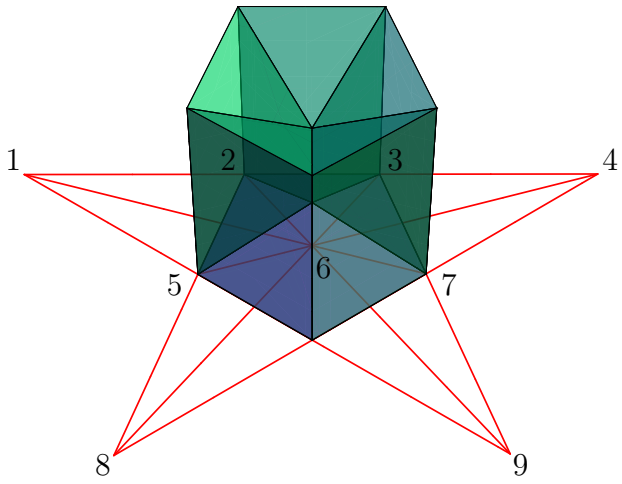

Figure 3.4: The core K, together with the resulting Perles configuration (shown in red). Compare with Figure 3.1.

We define 9 points $\mathbf{p}_1, \dots \mathbf{p}_9$ in P_b as follows. Let $\mathbf{p}_2 = \mathbf{x}_1$, $\mathbf{p}_3 = \mathbf{x}_5$, $\mathbf{p}_5 = \mathbf{x}_2$, $\mathbf{p}_7 = \mathbf{x}_4$. Define $\mathbf{p}_1, \mathbf{p}_4, \mathbf{p}_8, \mathbf{p}_9$ by

$$\{\mathbf{p}_1\} = \ell_2 \cap \ell_5, \quad \{\mathbf{p}_4\} = \ell_3 \cap \ell_5, \quad \{\mathbf{p}_8\} = \ell_1 \cap \ell_3, \quad \{\mathbf{p}_9\} = \ell_2 \cap \ell_4.$$

Finally, let $\ell_c = \text{aff}(\mathbf{a}, \mathbf{a}')$, and define $\mathbf{p}_6$ by $\{\mathbf{p}_6\} = \ell_c \cap P_b$.

One can directly check that the points $\mathbf{p}_1, \dots, \mathbf{p}_9$ constitute a realization Λ of the Perles configuration (in fact the points are labeled so that $\mathbf{p}_i$ corresponds to vertex i in Figure 3.1). Indeed, all collinearities are satisfied (see Fig. 3.4). For example, $\mathbf{p}_1, \dots, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4$ are all contained in the line ℓ_5 by definition, and $\mathbf{p}_3, \mathbf{p}_6, \mathbf{p}_8$ are collinear by a direct calculation. Furthermore, it is clear from the definitions that in any geometric realization of K, the collinearities $\{2, 5, 8\}$, $\{1, 5, 9\}$, $\{4, 7, 8\}$, $\{3, 7, 9\}$, $\{1, 2, 3, 4\}$ are satisfied, since they correspond to the lines $\ell_1, \ell_2, \ell_3, \ell_4, \ell_5$, respectively. However, the collinearities $\{2, 6, 9\}$, $\{4, 5, 6\}$, $\{1, 6, 7\}$, $\{3, 6, 8\}$ may fail. To obtain a geometric polyhedral complex X such that these last four collinearities hold in any realization of X, we attach four belts to K, one for each collinearity.

In order for X to be a polyhedral complex, when we attach four belts to K we must ensure

that they do not intersect. To achieve this, the belts we use are long thin arcs consisting of hundreds of triangular prisms. We have produced an explicit construction on the computer, of which we give an overview here. The complete code necessary to generate the complex explicitly is available in a MATHEMATICA notebook, which is included with this document as a supplemental file.

We force the collinearity $\{2, 6, 9\}$ to hold in any realization by attaching a belt B_1 to K as follows. Draw two simple arcs φ_1, φ_1', which intersect K in at most their endpoints, and such that $\varphi_1(0) = \varphi_1'(0) = \mathbf{x}_3$, $\varphi_1(1) = \mathbf{a}'$, and $\varphi_1'(1)$ is a point lying in the plane P_1, above the edge e_1' and close to but not directly above $\mathbf{a}'$. Place a large number N of points roughly equidistantly along each arc, and label these points $\mathbf{z}_1, \ldots, \mathbf{z}_N$ and $\mathbf{z}_1', \ldots, \mathbf{z}_N'$, respectively. We demand that $\mathbf{z}_1 = \mathbf{z}_1' = \mathbf{x}_3$, $\mathbf{z}_N = \mathbf{a}'$, and $\mathbf{z}_N' = \varphi_1'(1)$. These points determine a collection of triangles $\Delta_i = \mathrm{conv}(\mathbf{z}_i, \mathbf{z}_{i+1}, \mathbf{z}_{i+1}')$, for $i = 1, \ldots, N-1$ and $\Delta_i' = \mathrm{conv}(\mathbf{z}_i, \mathbf{z}_i', \mathbf{z}_{i+1}')$ for $i = 2, \ldots, N-1$.

For each triangle Δ_i, we obtain a triangular prism with triangular facet Δ_i by letting the lateral edges be segments of the lines passing through $\mathbf{p}_9$ and one of $\mathbf{z}_i, \mathbf{z}_{i+1}, \mathbf{z}_{i+1}'$. Similarly for the triangles Δ_i'. We demand in particular that the lateral edge containing $\mathbf{z}_1 = \mathbf{x}_3$ is the edge e_2, and that the lateral edge containing $\mathbf{z}_N$ is the edge e_1'. We are free to choose the length of the remaining lateral edges. By choosing this length to be very short for all lateral edges save for those containing points close to $\mathbf{z}_1$ and $\mathbf{z}_N$, we make it possible to attach other belts while avoiding intersections.

Concretely, let $m \in \mathbb{Z}_+$ and let $f_m : [0,1] \to [0,1]$ be defined by

$$f_m(t) = (1-t)^m + t^m.$$

Let $L(e)$ denote the length of the edge e of K, and let $g(t) = (1-t)L(e_2) + tL(e_1')$. Define a function h_m by $h_m(t) = f_m(t)g(t)$. Let E_i denote the lateral edge containing $\mathbf{z}_i$ and E_i' the lateral edge containing $\mathbf{z}_i'$. Then we choose the length of E_i to be $h_m(\frac{i-1}{N-1})$ and the length of E_i' to be $h_m(\frac{i-1}{N})$. Taking m large, we may ensure that the lengths of the lateral edges of each prism are very short except near the edges e_2 and e_1'. In our explicit construction we

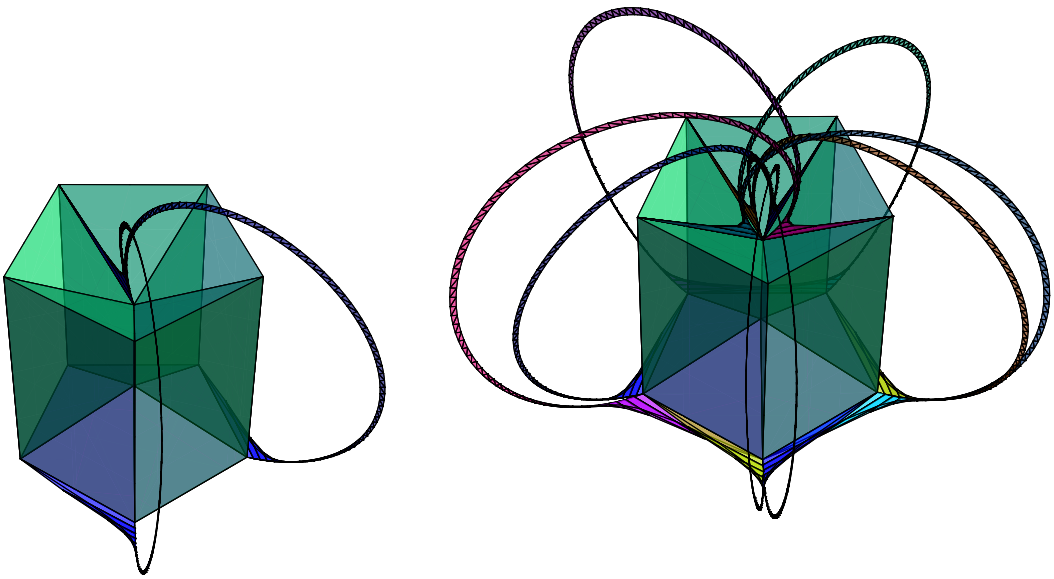

Figure 3.5: Left: Attaching one belt in the construction of Theorem 1.0.1. Right: The complete irrational complex with all four belts attached.

take $m = 80$.

The collection of the resulting triangular prisms forms a "half-belt" B_1'. Reflect B_1' across the plane P_1, and call the result B_1''. Then $B_1 = B_1' \cup B_1''$ is the desired belt. It intersects K in the three edges e_2, e_4, and e_1' (see Figure 3.5). In our explicit construction we take $N = 80$, so that B_1' and B_1'' each consist of $2(80-1)+1 = 159$ prisms, for a total of $2(159) = 318$ prisms in the belt B_1.

We now show that in any realization of the geometric polyhedral complex $K \cup B_1$, the collinearity $\{2, 6, 9\}$ is satisfied. Let Z denote any geometric realization of $K \cup B$. Let $\overline{\mathbf{p}_i}$, $\overline{\ell_i}$, $\overline{\ell_i'}$ and $\overline{P_1}, \overline{P_b}$ denote the points, lines, and planes of Z corresponding to $\mathbf{p}_i$, ℓ_i, ℓ_i', and P_1, P_b. Since the belt B is attached to K along the three edges e_2, e_4, and e_1', the lines $\overline{\ell_2}$, $\overline{\ell_4}$, and $\overline{\ell_1'}$ must be concurrent by Lemma 3.0.8, and their point of intersection is $\{\overline{\mathbf{p}_9}\} = \overline{\ell_2} \cap \overline{\ell_4}$. So in particular, $\overline{\mathbf{p}_9} \in \overline{\ell_1'}$. Clearly $\overline{\mathbf{p}_9} \in \overline{P_b}$ and $\overline{\ell_1'} \subseteq \overline{P_1}$, so we have $\overline{\mathbf{p}_9} \in \overline{\ell_1'} \cap \overline{P_b} \subseteq \overline{P_1} \cap \overline{P_b}$. But

$\overline{P_1} \cap \overline{P_b}$ is a line containing $\overline{\mathbf{p}_2}$ and $\overline{\mathbf{p}_6}$. Thus $\overline{\mathbf{p}_2}$, $\overline{\mathbf{p}_6}$, $\overline{\mathbf{p}_9}$ are collinear.

We may force the remaining three collinearities to hold in any realization by attaching three more belts to K. The construction of these remaining belts is analogous to the construction of B_1. In particular, attaching a belt B_2 to the edges e_3, e_5, and e_2' forces the collinearity $\{4, 5, 6\}$, attaching a belt B_3 to the edges e_5, e_2, and e_4' forces the collinearity $\{1, 6, 7\}$, and attaching a belt B_4 to the edges e_1, e_3, and e_5' forces the collinearity $\{3, 6, 8\}$. Let X denote the resulting geometric polyhedral complex with all 4 belts attached (see Figure 3.5). Choosing the arcs which define these belts to curve in the appropriate way, we may ensure that the belts do not intersect. In our explicit construction, each of these four belts consists of 318 triangular prisms. There are 5 triangular prisms in K, for a total of $4(318) + 5 = 1277$ triangular prisms in X. Together with the pentagonal pyramid M, we have a grand total of 1278 polytopes comprising X. For more pictures and further details see Chapter 9.

Thus in every geometric realization Y of X, the points corresponding to $\mathbf{p}_1, \ldots, \mathbf{p}_9$ form a realization Λ_P of $\mathcal{L}_P$, which is irrational by Lemma 1. Thus Y must have an irrational vertex coordinate. Otherwise, the supporting planes of Y would be the defined by equations with rational coefficients, whence the points of Λ_P would be the solutions of systems of linear equations with rational coefficients, hence rational, a contradiction. $\square$

CHAPTER 4

The Universality of 3-Polytopal Complexes and Arrangements

Let us begin by noting that the first irrational polytope result of Perles was later extended by Mnëv to a general universality theorem [Mn], and then further extended to all 4-polytopes [R]. Similarly, Brehm's result [Z3] gives a universality theorem for polyhedral 2-surfaces in $\mathbb{R}^3$.

In what follows, we extend Theorem 1.0.1 to a similar universality result. Using belts, we can in fact mimic the constructions of Theorem 1.0.1 for *any* point and line configuration. In particular, for a point and line configuration $\mathcal{L}$, we can construct a geometric 3-polyhedral complex $X(\mathcal{L})$ such that every realization of $X(\mathcal{L})$ generates a realization of $\mathcal{L}$. The universality theorem for point and line configurations (see e.g. [P]) then implies, in particular, that for any proper subfield K of the algebraic closure of $\mathbb{Q}$, there is a geometric 3-polyhedral complex that cannot be realized with all vertex coordinates in K.

Technically, the universality theorem for point and line configurations assumes that the realizations of a configuration are restricted to the projective plane $\mathbb{RP}^2$. If we allow realizations in $\mathbb{RP}^d$ for $d > 2$, some realizations may not lie entirely in a single plane. A point and line configuration $\mathcal{L}$ is said to be *planar* if every realization of $\mathcal{L}$ in $\mathbb{RP}^d$ lies in a (projective) 2-plane. If a point and line configuration is not planar, there is a straightforward way to extend it to a planar configuration, which we describe in the following lemma.

Lemma 4.0.9 *Let $\mathcal{L} = ([n], E)$ be a point and line configuration. There is a point and line configuration $\overline{\mathcal{L}} = ([3n+1], \overline{E})$ such that $E \subseteq \overline{E}$, and every realization of $\overline{\mathcal{L}}$ in $\mathbb{RP}^d$ is a*

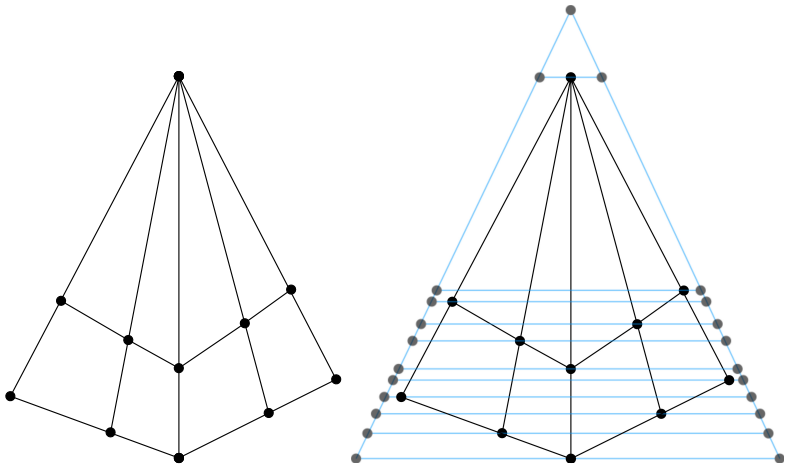

Figure 4.1: Left: A point and line configuration $\mathcal{L}$ that may have non-planar realizations in $\mathbb{RP}^3$. Right: The resulting planar configuration $\overline{\mathcal{L}}$ as constructed in Lemma 4.0.9.

subset of a 2-plane. Furthermore, every realization of $\overline{\mathcal{L}}$ contains a realization of $\mathcal{L}$.

Proof. Let $\mathcal{L} = ([n], E)$. For each point $i \in [n]$, we add two points $w_i = n+2i-1, w_i' = n+2i$ and the line $e_i = \{i, w_i, w_i'\}$. Then we add a new point $a = 3n+1$, together with two lines $l = \{a, w_1, \ldots, w_n\}$ and $l' = \{a, w_1', \ldots, w_n'\}$. See Figure 4.1. Let $\overline{\mathcal{L}} = ([3n+1], \overline{E})$ denote the resulting point and line configuration. Clearly, in every realization of $\overline{\mathcal{L}}$ in $\mathbb{RP}^d$, each point $\mathbf{p}$ must lie in the plane determined by the two intersecting lines L and L' corresponding to l and l'.

Finally, let $\overline{\Lambda}$ denote a realization of $\overline{\mathcal{L}}$. Since $X \subseteq \overline{X}$ and $E \subseteq \overline{E}$, there is a subset $\Lambda \subseteq \overline{\Lambda}$ and a map $f : X \to \Lambda$ such that f is a bijection, and a collection of points $\{f(x_i)\}_{i \in I}$ is collinear if $I \subseteq E$. To see that a collection of points $\{f(x_i)\}_{i \in I}$ is collinear only if $I \subseteq E$, note that each abstract line $e_i = \{i, w_i, w_i'\}$ contains only one point of $\mathcal{L}$, namely i. Furthermore, the two lines L and L' contain no points of $\mathcal{L}$. So the lines added to E to form $\overline{E}$ do not enforce any new concurrencies among the points of $\mathcal{L}$. □

We call the configuration $\overline{\mathcal{L}}$, constructed in Lemma 4.0.9, the *planar extension* of $\mathcal{L}$.

Theorem 4.0.10 (Weak Universality Theorem) *Let $\mathcal{L}$ be a point and line configuration, and*

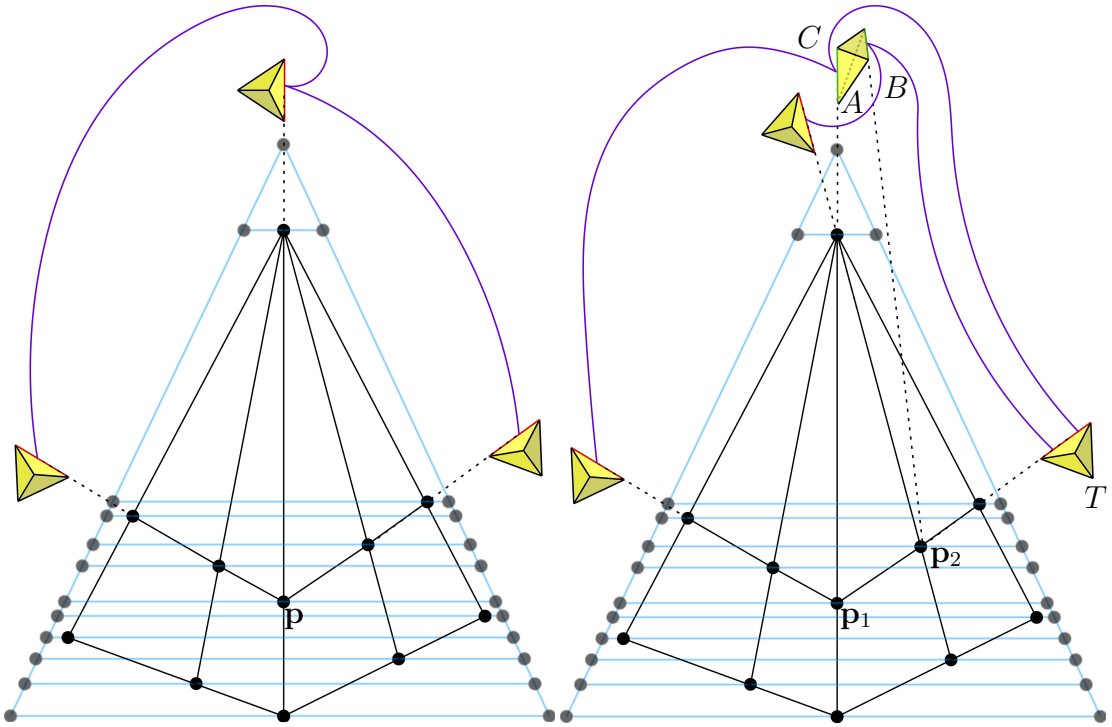

Figure 4.2: (Belts are shown schematically in purple) Left: A belt $B_{\mathbf{p}}$ which ensures that the three lines spanned by the red edges are concurrent in any realization of X. Right: Two belts B and C, attached to the tetrahedron A along the green edges, which ensure that the three lines spanned by the red edges are not concurrent in any realization of X (so that $\mathbf{p}_1 \neq \mathbf{p}_2$). Note that B and C share the red edge of the tetrahedron T, but they do not share a 2-face.

let K be a proper subfield of the algebraic closure of $\mathbb{Q}$. There exists a geometric 3-polyhedral complex $X(\mathcal{L})$ such that if $X(\mathcal{L})$ has a realization over K then $\mathcal{L}$ has a planar realization over K. Moreover, the complex $X(\mathcal{L})$ may be constructed using only triangular prisms.

Proof. We provide a sketch of the construction. By Lemma 4.0.9, we may assume that $\mathcal{L}$ is a planar configuration (if $\mathcal{L}$ is not planar, replace it with its planar extension). Let Λ denote a realization of $\mathcal{L}$ in $\mathbb{RP}^3$. For each line ℓ_i of Λ, place a tetrahedron with a marked edge e_i, such that $\ell_i = \text{aff}(e_i)$.

For each point $\mathbf{p} \in \Lambda$, let $L_{\mathbf{p}}$ denote the set of lines of Λ containing $\mathbf{p}$. For each such set $L_{\mathbf{p}}$, add a belt $B_{\mathbf{p}}$ such that for each line $\ell_i \in L_{\mathbf{p}}$, the edge e_i is identified with a lateral edge of $B_{\mathbf{p}}$. See Figure 4.2 (left).

Finally, for each collection of 3 lines ℓ_i, ℓ_j, ℓ_k which are not concurrent in the realization Λ, place a tetrahedron A_{ijk} with vertices labeled $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$. Add a belt B_{ijk} such that each of the 3 edges e_i, e_j, and $\mathbf{x}_1\mathbf{x}_2$ is identified with a lateral edge of B_{ijk}. Add another belt C_{ijk} such that each of the 3 edges e_j, e_k, and $\mathbf{x}_3\mathbf{x}_4$ is attached along a lateral edge of C_{ijk}. See Figure 4.2 (right). Call the resulting geometric 3-polyhedral complex $X(\mathcal{L})$.

Let X' be a geometric realization of $X(\mathcal{L})$. The belts $B_{\mathbf{p}}$ ensure that the concurrencies present among the edges e_i of $X(\mathcal{L})$ hold among the corresponding edges e_i' of X'. The points of intersection of the edges e_i' of X' constitute a set of points Λ'. To show that Λ' is a realization of $\mathcal{L}$, it remains to show that no further concurrencies hold among the edges e_i'. That is, we wish to show that Λ' is not a *degenerate* realization of $\mathcal{L}$. But this is ensured by the tetrahedra A_{ijk} and the corresponding pairs of belts B_{ijk} and C_{ijk}.

To see this, let ℓ_i, ℓ_j, ℓ_k be distinct lines of Λ which are not concurrent, and let e_i, e_j, e_k denote the corresponding edges of $X(\mathcal{L})$. Then $X(\mathcal{L})$ contains the tetrahedron A_{ijk} and belts B_{ijk}, C_{ijk} as described above. Let $A'_{ijk}, B'_{ijk}, C'_{ijk}$ denote the corresponding prisms and belts of X', where the vertices of A'_{ijk} are labeled $\mathbf{x}'_1, \mathbf{x}'_2, \mathbf{x}'_3, \mathbf{x}'_4$. The belt B'_{ijk} ensures that the edges e'_i, e'_j and $\mathbf{x}'_1\mathbf{x}'_2$ are concurrent, and the belt C'_{ijk} ensures that the edges e'_j, e'_k and $\mathbf{x}'_3\mathbf{x}'_4$ are concurrent. Suppose that the edges e'_i, e'_j, e'_k of X' are concurrent. From the concurrencies forced by the belts, this implies that the edges $\mathbf{x}'_1\mathbf{x}'_2$ and $\mathbf{x}'_3\mathbf{x}'_4$ are concurrent. Thus the vertices $\mathbf{x}'_1, \mathbf{x}'_2, \mathbf{x}'_2, \mathbf{x}'_3$ are coplanar, so they do not determine a tetrahedron, a contradiction.

Now let K be a proper subfield of the algebraic closure of $\mathbb{Q}$, and suppose that $X(\mathcal{L})$ is realizable over K. Let X' be a realization of $X(\mathcal{L})$ having all vertex coordinates in K, and let Λ be the realization of $\mathcal{L}$ generated by X'. Then the affine hulls of the faces of X' are defined by linear equations with coefficients in K. The points of Λ are the intersection of these affine hulls, hence they are solutions of a system of linear equations with coefficients in K. Thus the points of Λ have all coordinates in K. That is, $\mathcal{L}$ is realizable over K.

To prove the last claim of the theorem, note that each tetrahedron used in the above construction may be replaced in with a triangular prism. In fact, the tetrahedra whose marked edges generate the lines of Λ may be removed, as the edges of the attached belts suffice to define these lines. The tetrahedra A_{ijk} may be replaced with triangular prisms in the obvious way, by attaching the corresponding belts along two skew edges of the triangular prism, in the same way in which they were attached along two skew edges of the tetrahedron.

□

Note that the object produced by Theorem 4.0.10 is a polytopal complex, which by definition is not allowed to intersect itself (except that two polytopes may intersect in at most a common face of both). That is, a polyhedral complex is *embedded* in $\mathbb{R}^3$. This property, while geometrically appealing, makes Theorem 4.0.10 a *weak* universality theorem, in the sense that it does not imply that the *realization spaces* of $X(\mathcal{L})$ and $\mathcal{L}$ are *stably equivalent*. We would now like to investigate whether it is possible to obtain this latter type of result. We will find that by modifying our previous definitions slightly, we can in fact obtain a "true" universality theorem. To this end, we adopt the definition of stable equivalence given in [R], and we define the realization spaces of $X(\mathcal{L})$ and $\mathcal{L}$ as follows.

For a point and line configuration $\mathcal{L} = ([n], E)$, we define the *(Euclidean) realization space* of $\mathcal{L}$ (in $\mathbb{R}^3$) to be the set

$$\mathcal{R}(\mathcal{L}) = \{(\mathbf{p}_1, \ldots, \mathbf{p}_n) \in \mathbb{R}^{3n} \mid \Lambda = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\} \text{ is a realization of } \mathcal{L}\}.$$

In particular, we only allow realizations Λ in $\mathbb{R}^3$, rather than $\mathbb{RP}^3$ (that is, we do not allow points at infinity). This will be important for our final universality result. Notice that the coordinates of the realization space come with a particular order, induced by the natural order on $[n]$. That is, for each $i \in [n]$, if $(\mathbf{p}_1, \ldots, \mathbf{p}_n) \in \mathcal{R}(\mathcal{L})$ then $\mathbf{p}_i$ must be the point corresponding to i. For a geometric 3-polyhedral complex X with N vertices, the *realization space* of X is the set

$$\mathcal{R}(X) = \{(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_N) \in \mathbb{R}^{3N} \mid \boldsymbol{v}_1, \ldots, \boldsymbol{v}_N \text{ are the vertices of a realization of } X\}.$$

Consider the natural map $f : \mathcal{R}(X(\mathcal{L})) \to \mathcal{R}(\mathcal{L})$ that assigns to each realization X' of $X(\mathcal{L})$ the realization of $\mathcal{L}$ generated by X'. The following informal argument shows that f will not be a stable equivalence in general. Suppose that $X(\mathcal{L})$ is constructed so that its belts consist of a very large number of prisms. Let Λ be a realization of $\mathcal{L}$, and consider the fiber $\mathcal{A} = f^{-1}(\Lambda)$. Since the belts of $X(\mathcal{L})$ consist of a large number of prisms, for a given pair of belts B_1 and B_2 of $X(\mathcal{L})$, we may construct a realization $X' \in \mathcal{A}$ in which the corresponding belts are knotted, and a realization $Y' \in \mathcal{A}$ in which they are not knotted. Since we forbid the possibility that the belts B_1 and B_2 may intersect one another arbitrarily, the knot in X' is non-trivial. That is, there is no continuous path from X' to Y' in $\mathcal{R}(X(\mathcal{L})) \subset \mathbb{R}^{3N}$. Therefore $\mathcal{A}$ is not path-connected. But all fibers of a stable equivalence must be path-connected.

By modifying our definition of polyhedral complex slightly, we can eliminate the problem encountered in the previous paragraph. The idea is to preserve the face identifications in the complex, but allow polytopes to self-intersect (so in particular we will allow belts to intersect one another arbitrarily).

We define a *d-polytopal arrangement* $\mathcal{X} = (X, A)$ to consist of a set $X = \bigcup_{i=1}^{n} X_i$, where each X_i is a (face lattice of a) polytope, together with a set $A \subset X$ of *distinguished common faces*, such that any face $F \in A$ belongs to at least two polytopes, and any two polytopes have at most one common face that is contained in A. That is, if $F \in A$ and $F \in X_i \cap X_j$, then $G \notin A$ for all other $G \in X_i \cap X_j$.

A geometric polytopal arrangement $\mathcal{Y} = (Y, B)$ is a *realization* of $\mathcal{X} = (X, A)$ if there is a bijection $f : X \to Y$ such that $X_i \simeq f(X_i)$, and the face poset isomorphisms $g_i : X_i \to f(X_i)$ satisfy

$$F \in A \text{ if and only if } g_i(F) \in B.$$

The realization space $\mathcal{R}(\mathcal{X})$ of a polytopal arrangement $\mathcal{X}$ is defined in the obvious way.

Note that any two polytopes in $\mathcal{X}$ may *intersect* in more than a common face of both, but they can only have one common face *distinguished* by membership in A. That is, only the common face $F \in A$ is required to be a common face of both polytopes in every realization,

although the intersection of the polytopes may consist of much more.

Given a geometric d-polyhedral complex $X = \bigcup_{i=1}^{n} X_i$, we may construct a corresponding d-polytopal arrangement $\mathcal{X}$ by taking

$$A = \{F \in X \mid F \in X_i \cap X_j \text{ for some } i \neq j\}$$

and $\mathcal{X} = (X, A)$. The only difference between $\mathcal{X}$ and X is that in realizations of $\mathcal{X}$, we allow the polytopes to self-intersect arbitrarily. However, the intersections corresponding to the faces in A are required to hold in all realizations of $\mathcal{X}$. For this reason, Theorem 4.0.10 holds if we replace "polyhedral complex" with "polytopal arrangement", and the proof is identical. With this understanding, we may prove the desired universality results.

Theorem 4.0.11 *Let $\mathcal{L}$ be a point and line configuration. Then there is a 3-polytopal arrangement $\mathcal{X}(\mathcal{L})$ such that $\mathcal{R}(\mathcal{X}(\mathcal{L}))$ is stably equivalent to $\mathcal{R}(\mathcal{L})$.*

Proof. Let $\mathcal{L}$ be a point and line configuration, with realization $\Lambda \subset \mathbb{R}^3$. Let $Z(\mathcal{L})$ denote the corresponding geometric 3-polyhedral complex constructed from Λ as in Theorem 4.0.10. Let $\mathcal{Z}(\mathcal{L})$ be the polytopal arrangement corresponding to $Z(\mathcal{L})$. We begin by constructing a polytopal arrangement $\mathcal{X} = \mathcal{X}(\mathcal{L})$ by adding additional polytopes to $\mathcal{Z}(\mathcal{L})$. The purpose of adding these polytopes is simply to force the realizations of $\mathcal{L}$ generated by realizations of $\mathcal{X}$ to lie in $\mathbb{R}^3$ (rather than $\mathbb{RP}^3$), hence in $\mathcal{R}(\mathcal{L})$.

Let $\mathbf{p}_i \in \Lambda$, and place a tetrahedron (or triangular prism) T_i such that $\mathbf{p}_i$ is a vertex of T_i. Let e_1 and e_2 denote two of the edges of T_i containing $\mathbf{p}_i$. Let B denote a belt of $\mathcal{Z}(\mathcal{L})$ whose lateral edges are concurrent at $\mathbf{p}_i$. Let e_3 and e_4 denote two of the lateral edges of B. Now construct a new belt C_i that is attached along four of its lateral edges to the four edges e_1, e_2, e_3, e_4. See Figure 4.3. Adding T_i and C_i to $\mathcal{Z}(\mathcal{L})$ for each $i \in [n]$ yields the polytopal arrangement $\mathcal{X} = \mathcal{X}(\mathcal{L})$.

Consider a realization $\mathcal{X}'$ of $\mathcal{X}$, and let Λ' be the realization of Λ generated by $\mathcal{X}'$. For a point $\mathbf{p}' \in \Lambda'$, let e_1', e_2', e_3', e_4' be the edges corresponding to e_1, e_2, e_3, e_4 in the above construction. Since they are the lateral edges of a belt, these edges will be concurrent, at $\mathbf{p}'$.

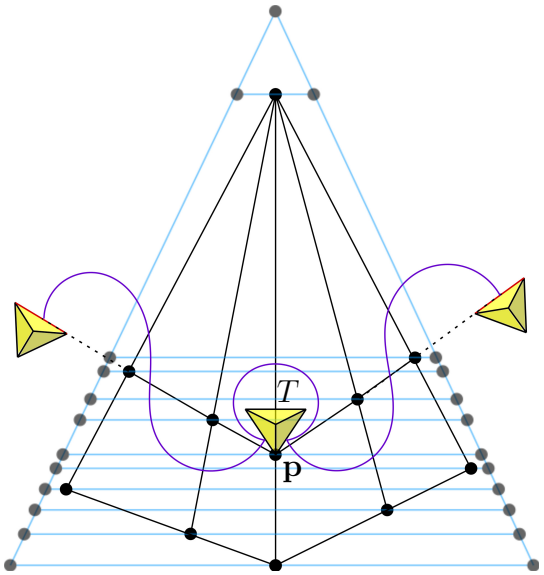

Figure 4.3: A belt, shown schematically in purple, attached to the four indicated edges of the tetrahedra. This belt forces the point $\mathbf{p}$ to be a vertex of the tetrahedron T in every realization of $\mathcal{X}$.

Since the point of concurrency of e_1' and e_2' is a vertex of $\mathcal{X}'$, we have that $\mathbf{p}'$ is a vertex of $\mathcal{X}'$. Thus, for any realization $\mathcal{X}'$ of $\mathcal{X}$, the realization Λ' of $\mathcal{L}$ generated by $\mathcal{X}'$ consists entirely of vertices of $\mathcal{X}'$. In particular, Λ' does not contain points at infinity.

We will now define a map $F : \mathcal{R}(\mathcal{X}) \to \mathcal{R}(\mathcal{L})$ such that F assigns to each realization $\mathcal{X}'$ of $\mathcal{X}$ the corresponding realization of $\mathcal{L}$ generated by $\mathcal{X}'$. We show that F is a stable equivalence, by showing that it is the composition of stable projections and rational homeomorphisms.

Let n denote the number of points of $\mathcal{L}$ and N the number of vertices of $\mathcal{X}$. Then we have $\mathcal{R}(\mathcal{L}) \subset \mathbb{R}^{3n}$, and $\mathcal{R}(\mathcal{X}) \subset \mathbb{R}^{3N}$. From the construction of $X(\mathcal{L})$ in Theorem 4.0.10, we see that each point i of $\mathcal{L}$ corresponds to a belt B_i of $\mathcal{X}$. In particular, we choose B_i to be one of the belts such that the lateral edges of B_i are concurrent at the point $\mathbf{p}_i \in \Lambda$.

For each belt B_i, fix an ordering of the its vertices $\boldsymbol{v}_{i1}, \boldsymbol{v}_{i2}, \ldots, \boldsymbol{v}_{is_i}$, such that the first four vertices are the (cyclically ordered) vertices of a lateral facet f_i of B_i, and such that $\boldsymbol{v}_{i1}\boldsymbol{v}_{i2}$ and $\boldsymbol{v}_{i3}\boldsymbol{v}_{i4}$ are the two lateral edges of f_i. Let $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_r$ denote the remaining vertices

of $\mathcal{X}$. Let $f_1 : \mathbb{R}^{3N} \to \mathbb{R}^{3N}$ be a map that permutes the coordinates, in such a way that

$$f_1(\mathbf{x}) = (\mathbf{x}_{11}, \mathbf{x}_{12}, \mathbf{x}_{13}, \mathbf{x}_{14}, \mathbf{x}_{21}, \mathbf{x}_{22}, \mathbf{x}_{23}, \mathbf{x}_{24}, \ldots, \mathbf{x}_{n1}, \mathbf{x}_{n2}, \mathbf{x}_{n3}, \mathbf{x}_{n4} \mid \mathbf{x}_{15}, \ldots, \mathbf{x}_{ns_n}, \mathbf{y}_1, \ldots, \mathbf{y}_r),$$

where each $\mathbf{x}_{ij}, \mathbf{y}_i \in \mathbb{R}^3$, and $\mathcal{X}$ is obtained by letting $\mathbf{x}_{ij}$ and $\mathbf{y}_i$ take on the values $\boldsymbol{v}_{ij}$ and $\boldsymbol{w}_i$, respectively. That is, each quadruple $\mathbf{x}_{i1}, \mathbf{x}_{i2}, \mathbf{x}_{i3}, \mathbf{x}_{i4}$ appears first in the ordering of the coordinates determined by f_1. We may choose any one of the many such maps f_1, since we do not care how the remaining coordinates (to the right of the bar) are permuted. Let $f_2 : \mathbb{R}^{3N} \to \mathbb{R}^{3(4n)}$ denote the standard projection onto the first $4n$ triples of coordinates.

Consider a realization $\mathcal{X}'$ of $\mathcal{X}$, and let $\Lambda = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\} \subseteq \mathbb{R}^3$ be the realization of $\mathcal{L}$ generated by $\mathcal{X}'$. Note that each point $\mathbf{p}_i$ is the point of intersection of the lateral edges of the belt B_i' of $\mathcal{X}'$ corresponding to the belt B_i of $\mathcal{X}$. Let $\boldsymbol{v}_{ij}'$ and $\boldsymbol{w}_i'$ denote the corresponding vertices of $\mathcal{X}'$

Note that $\boldsymbol{v}_{i2}'\boldsymbol{v}_{i3}'$ and $\boldsymbol{v}_{i1}'\boldsymbol{v}_{i4}'$ are the non-lateral edges of f_i. Let $\ell_1 = \text{aff}(\boldsymbol{v}_{i1}', \boldsymbol{v}_{i2}')$, $\ell_2 = \text{aff}(\boldsymbol{v}_{i3}', \boldsymbol{v}_{i4}')$, $\ell_3 = \text{aff}(\boldsymbol{v}_{i2}', \boldsymbol{v}_{i3}')$, $\ell_4 = \text{aff}(\boldsymbol{v}_{i1}', \boldsymbol{v}_{i4}')$ denote the lines spanned by these edges. By definition of B_i, the lines ℓ_1 and ℓ_2 intersect at $\mathbf{p}_i$. Thus the coordinates of $\mathbf{p}_i$ can be solved for in terms of those of the $\boldsymbol{v}_{ij}$. That is, $\mathbf{p}_i = g_i(\boldsymbol{v}_{i1}', \boldsymbol{v}_{i2}', \boldsymbol{v}_{i3}', \boldsymbol{v}_{i4}')$ for some rational function g_i with coefficients in $\mathbb{Q}$. Similarly, the lines ℓ_3 and ℓ_4 intersect in a point $\boldsymbol{q}_i$, so $\boldsymbol{q}_i = h_i(\boldsymbol{v}_{i1}', \boldsymbol{v}_{i2}', \boldsymbol{v}_{i3}', \boldsymbol{v}_{i4}')$ for some rational function h_i with coefficients in $\mathbb{Q}$.

Let $A \subset \mathbb{R}^{3(4)}$ denote the space of coplanar 4-tuples $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4)$, $\mathbf{x}_j \in \mathbb{R}^3$, such that no three are collinear. For each i, we define a map $s_i : A \to A$ by

$$s_i(\mathbf{x}_{i1}, \mathbf{x}_{i2}, \mathbf{x}_{i3}, \mathbf{x}_{i4}) = (g_i(\mathbf{x}_{i1}, \mathbf{x}_{i2}, \mathbf{x}_{i3}, \mathbf{x}_{i4}), \mathbf{x}_{i2}, h_i(\mathbf{x}_{i1}, \mathbf{x}_{i2}, \mathbf{x}_{i3}, \mathbf{x}_4), \mathbf{x}_{i4}).$$

From the definitions of g_i and h_i, we see that given a point in $\boldsymbol{b} = (\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3, \boldsymbol{b}_4) \in A$, we may reconstruct the unique point $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4) \in A$ for which $s_i(\mathbf{a}) = \boldsymbol{b}$. Namely, $\mathbf{a}_2 = \boldsymbol{b}_2$, $\mathbf{a}_4 = \boldsymbol{b}_4$, $\mathbf{a}_1$ is the intersection of the lines $\text{aff}(\boldsymbol{b}_1, \boldsymbol{b}_2)$ and $\text{aff}(\boldsymbol{b}_3, \boldsymbol{b}_4)$, and $\mathbf{a}_3$ is the intersection of the lines $\text{aff}(\boldsymbol{b}_1, \boldsymbol{b}_4)$ and $\text{aff}(\boldsymbol{b}_2, \boldsymbol{b}_3)$. See Figure 4.4. Thus s_i is a bijection, and clearly s_i and s_i^{-1} are continuous. Hence s_i is a homeomorphism. Since g_i and h_i are rational functions with coefficients in $\mathbb{Q}$, by definition s_i is a rational homeomorphism.

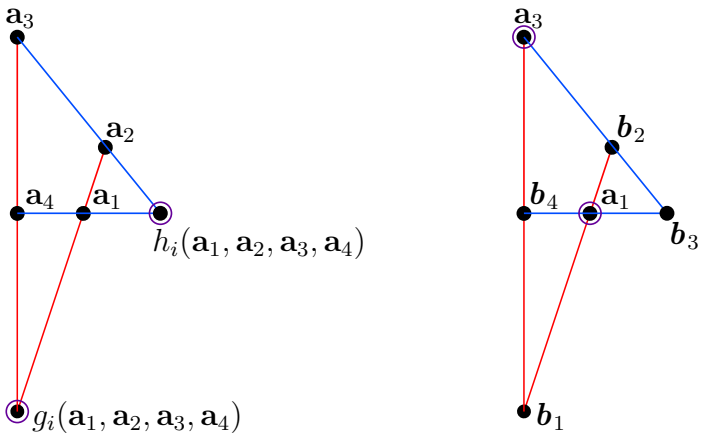

Figure 4.4: The map s_i (left) and its inverse (right). Encircled points indicate points obtained as the intersection of two lines.

It follows that the map $f_3 : A^n \to A^n$ defined by

$$f_3(\mathbf{x}_{11}, \mathbf{x}_{12}, \mathbf{x}_{13}, \mathbf{x}_{14}, \ldots, \mathbf{x}_{n1}, \mathbf{x}_{n2}, \mathbf{x}_{n3}, \mathbf{x}_{n4}) = (s_1(\mathbf{x}_{11}, \mathbf{x}_{12}, \mathbf{x}_{13}, \mathbf{x}_{14}), \ldots, s_n(\mathbf{x}_{n1}, \mathbf{x}_{n2}, \mathbf{x}_{n3}, \mathbf{x}_{n4}))$$

is a rational homeomorphism. Let $f_4 : \mathbb{R}^{3(4n)} \to \mathbb{R}^{3(4n)}$ be the coordinate permutation defined by

$$\begin{aligned} &f_4(\mathbf{x}_{11}, \mathbf{x}_{12}, \mathbf{x}_{13}, \mathbf{x}_{14}, \ldots, \mathbf{x}_{n1}, \mathbf{x}_{n2}, \mathbf{x}_{n3}, \mathbf{x}_{n4}) \\ &= (\mathbf{x}_{11}, \ldots, \mathbf{x}_{n1}, \mathbf{x}_{12}, \ldots, \mathbf{x}_{n2}, \mathbf{x}_{13}, \ldots, \mathbf{x}_{n3}, \mathbf{x}_{14}, \ldots, \mathbf{x}_{n4}). \end{aligned}$$

Let $f_5 : \mathbb{R}^{3(4n)} \to \mathbb{R}^{3n}$ denote the standard projection onto the first n triples of coordinates. Then note that

$$\begin{aligned} &f_5 \circ f_4 \circ f_3(\mathbf{x}_{11}, \mathbf{x}_{12}, \mathbf{x}_{13}, \mathbf{x}_{14}, \ldots, \mathbf{x}_{n1}, \mathbf{x}_{n2}, \mathbf{x}_{n3}, \mathbf{x}_{n4}) \\ &= (g_1(\mathbf{x}_{11}, \mathbf{x}_{12}, \mathbf{x}_{13}, \mathbf{x}_{14}), \ldots, g_n(\mathbf{x}_{n1}, \mathbf{x}_{n2}, \mathbf{x}_{n3}, \mathbf{x}_{n4})), \end{aligned}$$

so in particular

$$f_5 \circ f_4 \circ f_3(\boldsymbol{v}'_{11}, \boldsymbol{v}'_{12}, \boldsymbol{v}'_{13}, \boldsymbol{v}'_{14}, \ldots, \boldsymbol{v}'_{n1}, \boldsymbol{v}'_{n2}, \boldsymbol{v}'_{n3}, \boldsymbol{v}'_{n4}) = (\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n).$$

Let f'_1 denote the restriction of f_1 to $\mathcal{R}(\mathcal{X})$, and recursively, for $i > 1$ let f'_i denote the restriction of f_i to the range of f_{i-1}. We define $F = f'_5 \circ f'_4 \circ f'_3 \circ f'_2 \circ f'_1$. Then from the

definition of the functions f_i' we see that $F : \mathcal{R}(\mathcal{X}) \to \mathcal{R}(\mathcal{L})$ and F is a surjection. We demonstrated above that f_3 is a rational homeomorphism, and the functions f_1 and f_4 are simply permutations of the coordinates, hence rational homeomorphisms. It follows that f_1', f_3', f_4' are rational homeomorphisms.

Furthermore, note that the functions f_2 and f_5 are standard coordinate projections. We may factor f_2' and f_5' into a composition of projections, each of which deletes one coordinate at a time. Since realizations of $\mathcal{X}$ allow for self-intersection of polytopes, the only constraints on $\mathcal{R}(\mathcal{X})$ are those which require each vertex to be contained in the appropriate polytopes. That is, there are no constraints forbidding the intersection of polytopes. Thus every factor map has the property that each of its fibers is the intersection of half-spaces (in particular, no unions are taken over half-spaces). That is, these factors are stable projections, so f_2' and f_5' are the composition of stable projections. Therefore F is a stable equivalence. □

While the above result shows that the realization space of a polytopal arrangement is stably equivalent to the underlying point and line configuration, we would like a stronger result, which states that realization spaces of polytopal arrangements can be stably equivalent to arbitrary *semialgebraic* sets. For this, we will need to add an additional structure to our point and line configurations, to obtain an *oriented matroid*. We provide an equivalent definition of an oriented matroid, which will be convenient for our purposes.

A line $L \subset \mathbb{R}^d$ has two possible orientations, each of which induces a linear order on the points $\mathbf{x} \in L$. Let $\mathcal{L} = (X, E)$ be a *planar* point and line configuration, such that every two abstract lines $e_1, e_2 \in E$ share a point of X (that is, $e_1 \cap e_2 \neq \emptyset$). Let $\Lambda \subset \mathbb{R}^d$ be a realization of $\mathcal{L}$. For each $e \in E$, let L be the line of Λ realizing e, and choose an orientation γ of L. Write $e = \{i_1, \ldots, i_k\}$, where $(\mathbf{p}_{i_1}, \mathbf{p}_{i_2}, \cdots \mathbf{p}_{i_k})$ is the order of the $\mathbf{p}_j \in L$ induced by γ. We define the *oriented line* e' to be the ordered tuple $(i_1, \ldots, i_k)$. Let E' denote the resulting set of oriented lines e'. Then $\mathcal{M} = (X, E')$ is an *oriented matroid*. One can readily check that this definition is equivalent to those given elsewhere (see e.g. [R]). In particular, the fact that we require every two lines of $\mathcal{M}$ to intersect in a point of $\mathcal{M}$ means that all

realizations of $\mathcal{M}$ (which will agree on the order of the points) will agree on the set of half planes in which a given point lies.

If $\mathcal{M} = (X, E')$ is an oriented matroid, let E be the set obtained by replacing each tuple $(i_1, \ldots, i_k) \in E'$ with the set $\{i_1, \ldots, i_k\}$. Then we say that $\mathcal{L}(\mathcal{M}) = (X, E)$ is the point and line configuration corresponding to $\mathcal{M}$. A *realization* of $\mathcal{M}$ is a set $\Lambda = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\} \subset \mathbb{R}^d$ such that Λ is a realization of $\mathcal{L}(\mathcal{M})$, and such that if L is a line of Λ corresponding to $e' = (i_1, \ldots, i_k)$, then there is an orientation γ of L such that $(\mathbf{p}_{i_1}, \mathbf{p}_{i_2}, \cdots \mathbf{p}_{i_k})$ is the order of the $\mathbf{p}_j \in L$ induced by γ. In other words, a realization of $\mathcal{M}$ is a realization of the underlying point and line configuration in which the points of each line have a prescribed order, up to reversing the orientation of the line.

Theorem 4.0.12 *Let $\mathcal{M}$ be an oriented matroid. Then there is a 3-polytopal arrangement $\mathcal{Y}(\mathcal{M})$ such that $\mathcal{R}(\mathcal{Y}(\mathcal{M}))$ is stably equivalent to $\mathcal{R}(\mathcal{M})$.*

Proof. Let $\mathcal{M}$ be an oriented matroid, and let $\Lambda \subset \mathbb{R}^3$ be a realization of $\mathcal{L}(\mathcal{M})$. Let $\mathcal{X}(\mathcal{L}(\mathcal{M}))$ be the polytopal arrangement constructed from Λ as in Theorem 4.0.11. We write $\mathcal{X} = \mathcal{X}(\mathcal{L}(\mathcal{M}))$. We introduce a new polyhedral gadget which, when added to $\mathcal{X}$, will yield a polytopal arrangement $\mathcal{Y} = \mathcal{Y}(\mathcal{M})$ such that every realization of $\mathcal{Y}$ generates a realization of $\mathcal{M}$. That is, every realization of $\mathcal{Y}$ generates a realization of $\mathcal{L}(\mathcal{M})$ in which the points occur on each line in the order prescribed by $\mathcal{M}$. The proof of stable equivalence is then identical to the proof of Theorem 4.0.11.

Let $\mathbf{p}_i, \mathbf{p}_j, \mathbf{p}_k$ denote three points of Λ which appear consecutively on a line L of Λ. We construct a belt G_{ijk} consisting of two triangular prisms, call them T_1 and T_2, as follows. Let $\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_4$ denote the vertices of a lateral facet f of T_1 labeled cyclically, where $e_1 = \boldsymbol{v}_1\boldsymbol{v}_4$ and $e_2 = \boldsymbol{v}_2\boldsymbol{v}_3$ are the lateral edges of f. We may choose the vertices $\mathbf{x}_i$ so that if $\ell_1 = \text{aff}(\boldsymbol{v}_1, \boldsymbol{v}_2)$, $\ell_2 = \text{aff}(\boldsymbol{v}_1, \boldsymbol{v}_3)$, and $\ell_3 = \text{aff}(\boldsymbol{v}_1, \boldsymbol{v}_4)$, then $\mathbf{p}_i \in \ell_1$, $\mathbf{p}_j \in \ell_2$, and $\mathbf{p}_k \in \ell_3$. Note that ℓ_2 is the line spanned by the diagonal $\boldsymbol{v}_1\boldsymbol{v}_3$ of the facet f. Let $\mathbf{z}$ denote the point of intersection of the diagonals $\boldsymbol{v}_1\boldsymbol{v}_3$ and $\boldsymbol{v}_2\boldsymbol{v}_4$. We attach T_2 to T_1 along the unique lateral facet g of T_1 that contains the edge e_2 and does not contain e_1. Let e_3 denote the non-lateral

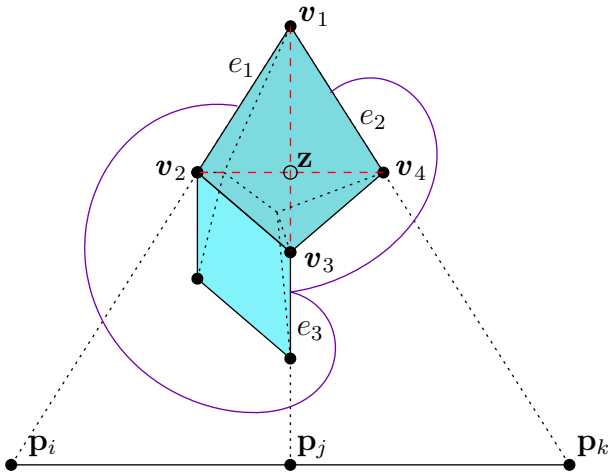

Figure 4.5: The arrangement I_{ijk}, with the two prisms T_1 and T_2 shown in blue and the belt H_{ijk} shown schematically in purple. The diagonals of the facet f of T_1 are shown in red.

edge of T_2 containing the vertex $\boldsymbol{v}_3$, and not contained in g. We may choose the vertices of T_2 such that $e_3 \subset \ell_2$.

Attach a belt H_{ijk} to G_{ijk} along the edges e_1, e_2, e_3 of G_{ijk}. Call the resulting arrangement I_{ijk} (see Figure 4.5). Then in every realization of I_{ijk}, the edges e_1, e_2, e_3 will be concurrent. Thus the vertex $\boldsymbol{v}_1$ will be contained in $\text{aff}(e_3)$. Since $\boldsymbol{v}_3 \in e_3$ by construction, this implies that $\ell_2 = \text{aff}(e_3)$. That is, the diagonal $\boldsymbol{v}_1\boldsymbol{v}_3$ will be collinear with the line spanned by e_3. If B_a denotes a belt of $\mathcal{X}$ with all lateral edges concurrent at $\mathbf{p}_a$, then attach a lateral edge of the belt B_i to the edge e_1 of I_{ijk}, a lateral edge of B_k to e_2, and a lateral edge of B_j to e_3. Let $\mathcal{Y}$ denote the arrangement obtained by adding I_{ijk} to $\mathcal{X}$ for each such consecutive collinear triple of points $\mathbf{p}_i, \mathbf{p}_j, \mathbf{p}_k$, and performing the belt attachments just described.

Now suppose $\mathbf{p}_i, \mathbf{p}_j, \mathbf{p}_k$ are three consecutive collinear points of Λ, contained in a line L of Λ. Let $\mathcal{Y}'$ be a realization of $\mathcal{Y}$, and for each point, edge, line, or facet a of $\mathcal{Y}$, let a' denote the corresponding object of $\mathcal{Y}'$. Then $\mathbf{z}'$ must lie between $\boldsymbol{v}_2'$ and $\boldsymbol{v}_4'$ on the line $\text{aff}(\boldsymbol{v}_2', \boldsymbol{v}_4')$, since $\mathbf{z}'$ is the intersection of the diagonals $\boldsymbol{v}_1'\boldsymbol{v}_3'$ and $\boldsymbol{v}_2'\boldsymbol{v}_4'$ of f', and f' is convex. From the attachment of the belts B_i, B_j, B_k to I_{ijk}, we see that $\mathbf{p}_i \in \ell_1'$, $\mathbf{p}_j \in \text{aff}(e_3') = \ell_2'$, and $\mathbf{p}_k \in \ell_3'$. Thus the point $\mathbf{p}_j'$ must lie between the points $\mathbf{p}_i'$ and $\mathbf{p}_k'$ on the line L' containing them.

It follows that in every realization $\mathcal{Y}'$ of $\mathcal{Y}$, the realization Λ' of $\mathcal{L}(\mathcal{M})$ generated by $\mathcal{Y}'$ has the property that all points occur along each line in the order prescribed by $\mathcal{M}$. That is, Λ' is a realization of $\mathcal{M}$. □

From Theorem 4.0.12, together with the universality theorem for oriented matroids (see [R]), we obtain the following universality theorem for polytopal arrangements.

Corollary 4.0.13 (Strong Universality Theorem) *Let $V \subseteq \mathbb{R}^m$ be a basic primary semialgebraic set, defined over $\mathbb{Z}$. Then there is a 3-polytopal arrangement $\mathcal{X}$ such that $\mathcal{R}(\mathcal{X})$ is stably equivalent to V.*

CHAPTER 5

Geometric Realizations of Polyhedral Complexes

5.1 Unrealizable Complexes

Using the techniques of the construction in the proof of Theorem 1.0.1 (namely belts), it is easy to construct topological polyhedral complexes that have no geometric realization at all. We will prove Theorems 1.0.2 and 1.0.3 by constructing two such complexes. The second of these complexes is homeomorphic to a 3-ball.

Proof of Theorem 1.0.2. Let A be a 3-simplex in $\mathbb{R}^3$, with vertices labeled $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$. Let $e_{i,j} = \text{conv}(\mathbf{x}_i, \mathbf{x}_j)$, $i \neq j$, denote the edges of A. Attach a topological belt B consisting of 2 triangular prisms to the two edges $\mathbf{x}_1\mathbf{x}_2$ and $\mathbf{x}_3\mathbf{x}_4$ as shown in Figure 5.1. Call the resulting topological polyhedral complex X.

Suppose that X has a geometric realization Y, with vertices $\mathbf{y}_i$ corresponding to $\mathbf{x}_i$. Since the edges $e_{1,2}$ and $e_{3,4}$ both belong to the belt B, the corresponding edges of Y must be concurrent by Lemma 3.0.8. But then the vertices $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \mathbf{y}_4$ are all coplanar. Hence they do not determine a 3-simplex, a contradiction. □

Proof of Theorem 1.0.3. Let X denote the polyhedral complex of Theorem 1.0.2. As shown in Figure 5.1, the vertices $\mathbf{x}_1, \mathbf{x}_3, \mathbf{a}_1$ and the edges between them determine a topological 2-simplex, call it Δ_1. Similarly, the vertices $\mathbf{x}_2, \mathbf{x}_4, \mathbf{a}_2$ and the edges between them determine a topological 2-simplex Δ_2. That is, Δ_1 and Δ_2 are topological 2-polyhedral complexes. We define $S = X \cup \{\Delta_1, \Delta_2\}$.

Note that S is not a 3-polyhedral complex by our definition, because it contains facets

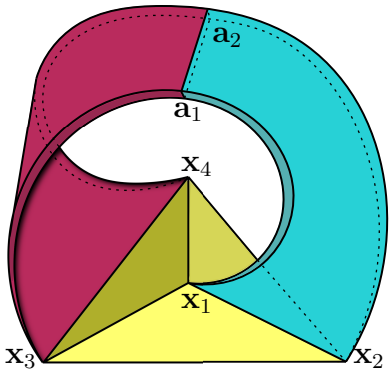

Figure 5.1: The unrealizable topological polyhedral complex of Theorem 1.0.2.

Δ_1 and Δ_2 which are not contained in any cell of S. However, we may create a polyhedral complex from S as follows. Let R denote the bounded component of the complement of S (i.e. R is the region surrounded by S). Let $c \in R$, and for each facet F in the boundary of (the closure of) R, add to S the topological cone with apex c and base F. In other words, cone from the point c. Let X' denote the resulting 3-polyhedral complex.

Then clearly X is a subcomplex of X', and $|X'| \sim B^3$. Furthermore, X' has exactly 9 vertices, 24 edges, 25 facets, and 9 cells. The cells of X' are comprised of 5 tetrahedra, 2 triangular prisms, and 2 tetragonal pyramids. □

It is worth noting that the unrealizable complex of Theorem 1.0.2 is minimal, in the sense that any topological 3-polyhedral complex consisting of two polyhedra is geometrically realizable. To see this, let X be a topological 3-polyhedral complex consisting of two 3-polyhedra Q_1 and Q_2. If Q_1 and Q_2 share less than a 2-face, the result is immediate. So suppose that Q_1 and Q_2 share a 2-face F. Let P_1 and P_2 be polytopes isomorphic to Q_1 and Q_2, respectively. Let F_1 and F_2 denote the facets of P_1 and P_2, respectively, that correspond to F. Barnette and Grünbaum [BG] proved that the shape of one facet of a 3-polytope may be arbitrarily prescribed. Therefore we may choose P_1 and P_2 so that F_1 and F_2 are congruent. Apply an affine transformation that identifies F_1 with F_2. The result is a geometric 3-polyhedral complex isomorphic to X.

5.2 Sufficient Conditions for Realization

Having described complexes which are not realizable, we now undertake the task of devising reasonable sufficient conditions for a polyhedral complex to be geometrically realizable. The first is Theorem 1.0.4, which tells us that we may geometrically realize a certain class of simplicial complexes in arbitrary dimension. The second is an analogous result for general polyhedral complexes, and holds only for $d \leq 3$. Before we state and prove these theorems we will need to introduce some relevant definitions. Some of these definitions (such as *star*, *link*, and vertex decomposable) are standard, while others (such as *strongly vertex decomposable* and *vertex truncatable*) are not.

Let X be a polyhedral complex and let F be a face of X. The *star* of F in X is the set $\mathrm{st}_X(F) = \{A \in X \mid F \subseteq A\}$. Let $\mathrm{cst}_X(F) = \{A \in X \mid A \subseteq B \in \mathrm{st}_X(F)\}$ denote the *closed star* of F. Note that $\mathrm{cst}_X(F)$ is a polyhedral complex, while $\mathrm{st}_X(F)$ may not be (since it may not be closed under taking subfaces). For $F \in X$, let $X \smallsetminus F$ be a shorthand for $X \smallsetminus \mathrm{st}_X(F)$, which is called the *deletion* of F from X. The deletion is clearly a polyhedral complex. The *link* of F in X is the polyhedral complex $\mathrm{lk}_X(F) = \mathrm{cst}_X(F) \smallsetminus F$.

Let X be a topological d-polyhedral complex such that $|X| \sim B^d$, and let v be a boundary vertex of X. We say that v is *boundary minimal* if v is contained in exactly d boundary facets, each of which is a $(d-1)$-simplex. We call v a *shedding vertex* of X if $|X \smallsetminus v| \sim B^d$, and a *strong shedding vertex* if in addition v is boundary minimal. We say that X is *strongly vertex decomposable* if either X is a single polyhedron, or recursively, X has a strong shedding vertex v such that $X \smallsetminus v$ is strongly vertex decomposable. A boundary vertex w of X is a *solitary vertex* if there is exactly one polyhedron $P_w \in \mathcal{P}_X$ such that $w \in P_w$, and a *strong solitary vertex* if in addition w it is boundary minimal.

If we restrict ourselves to *simplicial* complexes, then we find that in any dimension d, a strongly vertex decomposable simplicial d-ball is always geometrically realizable. This is the statement of Theorem 1.0.4, and the proof is straightforward.

Proof of Theorem 1.0.4. Let X denote a d-simplicial complex, such that $|X| \sim B^d$ and X

is strongly vertex decomposable. We proceed by induction on n, the number of vertices of X. If $n = d + 1$, then X consists of a single d-simplex, which is obviously realizable. If $n > d+1$, let v be a strong shedding vertex of X, and let $X' = X \setminus v$. Then by definition X' is a strongly vertex decomposable d-ball. Since X' has one fewer vertex than X, by the induction hypothesis X' has a geometric realization Y'. Since v is a strong shedding vertex, it is adjacent to exactly d boundary vertices of X, call them $w_1, \dots w_d$. Then $\mathrm{lk}_{\partial X}(v)$ (the link of v in the complex ∂X) has exactly d faces of maximal dimension, call them $F_1, \dots, F_d$, each of which is a $(d-2)$-face of X. Each face F_i is contained in exactly one facet τ_i of X'. Let $H_i = \mathrm{aff}(\tau_i)$ and let $H' = \mathrm{aff}(w_1, \dots, w_d)$.

Since d hyperplanes in $\mathbb{RP}^d$ always intersect in a point, we may let x denote the point of intersection of the hyperplanes H_i. If x is a point at infinity, or x lies on the same side of the hyperplane H' as $|Y'|$, then apply a projective transformation so that x is a finite point and x and $|Y'|$ lie on different sides of H'. Now let u be a point contained in $\mathrm{conv}(Y' \cup v)$ such that u is very close to v. Add straight line segments between u and all vertices of $|Y'|$ that correspond to neighbors of v in X. By taking u arbitrarily close to v, we may ensure that these added line segments intersect $|Y'|$ only in the desired vertices. These line segments, together with u and its neighbors, determine a collection of k-simplices for $2 \leq k \leq d$. Adding these simplices to Y' yields a geometric d-simplicial complex Y such that $X \simeq Y$. □

Now we consider the general case that X is a d-polyhedral complex. We will need to develop some more involved definitions in order to prove a result analogous to Theorem 1.0.4.

Suppose X_1 and X_2 are two d-polyhedra that share exactly one facet Q. Then let $X_1 \#_Q X_2$ denote the polyhedron obtained from $X_1 \cup X_2$ by replacing the two cells $|X_1|$ and $|X_2|$ by the single cell $|X_1| \cup |X_2|$ and removing the facet Q.

We now define a construction called *subdivision by facet*. Let X be a polyhedral complex, let v be a boundary vertex of X, and let $P \in \mathcal{P}_X(v)$. Let $\tau(v, P) \subseteq |P|$ denote a (topological) $(d-1)$-ball such that $\tau(v, P) \cap |\mathrm{lk}_{\partial P}(v)| = \partial\tau(v, P)$. In particular, the vertices of X contained

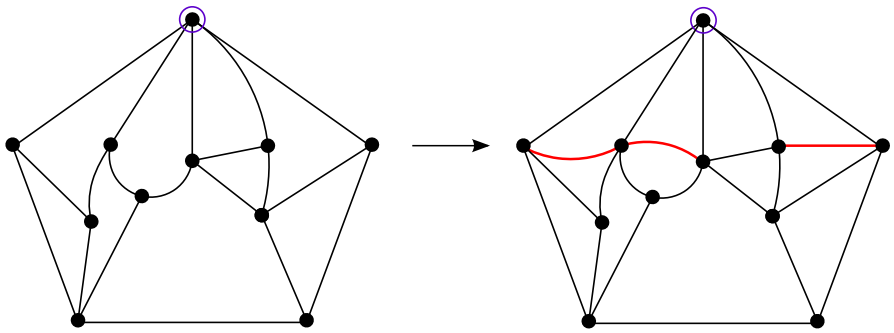

Figure 5.2: A 2-polyhedral complex with vertex v encircled, and the resulting complex $X^*(v)$ with the added facets $\tau(v, P)$ shown in red.

in $\tau(v, P)$ are exactly the neighbors of v in ∂X. We may construct a new polyhedral complex $X \oplus \tau(v, P)$, called the *subdivision of X at v and P*, as follows. If X already contains a facet $F \subseteq P$ such that $F \cap |\mathrm{lk}_{\partial P}(v)| = \partial F$, then take $X \oplus \tau(v, P) = X$. If X contains no such facet, then $X \oplus \tau(v, P)$ is obtained by adding the facet $\tau(v, P)$ and replacing $P \in X$ with two new cells σ_1, σ_2 such that $\sigma_1 \cup \sigma_2 = P$ and $\sigma_1 \cap \sigma_2 = \tau(v, P)$. Note that subdivision by facet has no effect on simplicial complexes.

For a boundary vertex v of X we define the *full subdivision of X at v* by

$$X^*(v) = X \bigoplus_{P \in \mathcal{P}_X(v)} \tau(v, P).$$

That is, $X^*(v)$ is obtained from X by repeatedly doing subdivision by facet, in effect subdividing X at v and P for each $P \in \mathcal{P}_X(v)$ (see Figure 5.2).

We say that X is *vertex truncatable* if $|X| \sim B^d$ and at least one of the following holds:

(a) X is a d-simplex.

(b) X has a strong shedding vertex v such that $X^*(v) \setminus v$ is vertex truncatable.

(c) X has a strong solitary vertex v such that $X^*(v) \setminus v$ is vertex truncatable.

Note that if $|X| \sim B^d$ and v is a strong shedding vertex of X, then v is clearly a strong shedding vertex of $X^*(v)$. Thus $|X^*(v) \setminus v| \sim B^d$. It follows that strongly vertex decomposable implies vertex truncatable.

Two k-faces $F_1, F_2 \in X$ are said to be *strongly adjacent* if $\dim(F_1 \cap F_2) = k-1$. For X_1 a polytope and $F_1, F_2 \in X_1$ strongly adjacent facets, we define $\alpha(F_1, F_2)$ to be the interior angle of X_1 formed between F_1 and F_2. More precisely, $\alpha(F_1, F_2)$ is the angle between the normal vectors to the hyperplanes $\text{aff}(F_1)$ and $\text{aff}(F_2)$. Technically, there are two such angles—we take the angle less than π so that it is interior to X_1. If $d = 2$ then $\alpha(F_1, F_2)$ is a vertex angle, and if $d = 3$ then $\alpha(F_1, F_2)$ is a dihedral angle.

Given a d-polytope X in $\mathbb{R}^d$ and a facet $F \in X$, the *Schlegel set* of X with respect to F, which we denote by $\xi(X, F)$, is the closed convex set bounded by the hyperplanes $\text{aff}(F)$ and $\text{aff}(G_i)$ for each facet G_i strongly adjacent to F. That is, if $G_1, \ldots, G_k \in X$ denote the facets strongly adjacent to F, the set H_{G_i} is the closed halfspace bounded by $\text{aff}(G_i)$ and meeting the interior of X, and the set H_F is the closed halfspace bounded by $\text{aff}(F)$ and *not* meeting the interior of X, then $\xi(X, F) = H_F \cap \bigcap_{i=1}^{k} H_{G_i}$. A *Schlegel point* is a point $y \in \text{int}(\xi(X, F))$.

If the facet F is a $(d-1)$-simplex, the hyperplanes $\text{aff}(G_i)$ intersect in a (possibly infinite) point, which we call the *apex* of $\xi(X, F)$. If the apex x is a finite point, then $x \in \xi(X, F)$ if and only if x and X lie on opposite sides of the hyperplane $\text{aff}(F)$. If the apex x is a finite point and $x \in \xi(X, F)$, then $\xi(X, F)$ is just the cone with base F and apex x.

Theorem 5.2.1 *Let $d \leq 3$ and let X be a topological d-polyhedral complex in $\mathbb{R}^d$ such that $|X| \sim B^d$ and all interior facets F (that is, $F \in X \smallsetminus \partial X$) are simplices. If X is vertex truncatable, then there is a geometric polyhedral complex Y in $\mathbb{R}^d$ such that $X \simeq Y$. Furthermore, we may choose Y to have all vertices rational.*

Proof. To prove the theorem, we strengthen it, which in turn strengthens our induction hypothesis. Namely, we claim that if such a Y exists, then furthermore:

(i) For any polytope Z such that $\partial Z \simeq \partial X$, we may choose Y such that $|Y| = |Z|$.

(ii) Let $\epsilon > 0$. Let $X_1, \ldots X_k \in \mathcal{P}_X$ such that $X_i \cap X_j \subseteq \partial X$ and $X_i \cap \partial X$ is a facet of X, call it F_i. Let $Y_i \in \mathcal{P}_Y$ be the polyhedron corresponding to X_i and let $G_i \in \partial Y$ be

the face corresponding to $F_i \in \partial X$. Then we may choose Y such that $\alpha(H_j, G_i) < \epsilon$ for all $i = 1, \ldots, k$ and every face $H_j \in Y_i$ strongly adjacent to G_i.

If $d \leq 1$ the theorem is trivial, so assume $d \in \{2, 3\}$. Suppose that X is a vertex collapsible d-polyhedral complex with simplicial interior facets. We proceed by induction on n, the number of vertices of X. In the base case X has $d+1$ vertices, and we may take Y to be any geometric d-simplex in $\mathbb{R}^d$. Now suppose $n > d+1$.

Since X is vertex truncatable, choose a boundary vertex w_0 satisfying either condition (b) or (c) of the definition. Let $X' = X^*(w_0) \setminus w_0$. By definition, X' is vertex truncatable, so in particular $|X'| \sim B^d$. Furthermore, all interior facets of X' are clearly simplices.

Let Z be a polytope such that $\partial Z \simeq \partial X$. If $d = 3$, such a polytope Z exists by Steinitz's Theorem. If $d = 2$ then Z may be any strictly convex polygon with the same number of boundary vertices as X. Since w_0 is boundary minimal, it has d neighbors $w_1, \ldots, w_d$ in ∂X. Let v_0 denote the vertex of Z corresponding to w_0. For each $i = 1, \ldots, d$, the vertex w_i corresponds to a vertex v_i of Z, where v_i is a neighbor of v_0. The vertices v_i lie on a hyperplane $H = \text{aff}(v_1, \ldots, v_d)$. The hyperplane H splits Z into two polytopes Q_1, Q_2, each defined by taking all faces of Z lying on a given side of H, which includes the facet $T_Z = \text{conv}(v_1, \ldots, v_d)$ in both cases. One of these two polytopes, say Q_1, contains v_0. Then note that Q_1 is a d-simplex.

Clearly $|\text{lk}_{X^*(w_0)}(w_0)|$ is homeomorphic to a $(d-1)$-ball. So if $d = 3$ then $\text{lk}_{X^*(w_0)}(w_0)$ is isomorphic to the Schlegel diagram of some 3-polytope by Steinitz's theorem. Call this polytope A. If $d = 2$ then we may simply take A to be a convex polygon with one more edge than $\text{lk}_{X^*(w_0)}(w_0)$. Let $u_1, \ldots, u_d$ be the vertices of A corresponding to the vertices $v_1, \ldots, v_d$ of $\text{lk}_{X^*(w_0)}(w_0)$, and let $T_A = \text{conv}(u_1, \ldots, u_d)$. Note that T_A is a $(d-1)$-simplex, and a facet of A. Apply an affine transformation to $|A|$ so that $T_A = T_Z$. If A is not a simplex, apply a projective transformation to $|A|$ that fixes T_A and takes the apex x of $\xi(A, T_A)$ to a finite point on the side of the hyperplane $\text{aff}(T_A)$ not containing $|A|$.

For a set $S \subseteq \mathbb{R}^d$, let $(S, 1) = \{(x, 1) \in \mathbb{R}^d \times \mathbb{R} \mid x \in S\}$. Form the $(d+1)$-cone C with

base $(|A|, 1)$ and apex $a \in \mathbb{R}^{d+1}$, $a_{d+1} \neq 1$. Let F_C denote the (clearly simplicial) d-face of C containing a and $(T_A, 1)$. Let W denote the Schegel projection of C onto its facet F_C, with respect to a Schlegel point y. Then W is a geometric polyhedral complex. Clearly W contains a subcomplex B such that $B \simeq A$. Since $(T_A, 1) \subseteq F_C$, the $(d-1)$-face $(T_A, 1)$ is fixed by the Schlegel projection, so in fact $(T_A, 1) \in B$ is the face of B corresponding to $T_A \in A$. Apply an affine transformation to $|W|$ that maps $(T_A, 1)$ to $T_A = T_Z$ and maps a to v_0. This transforms $|B|$ accordingly. In particular, we now have $|W| = |Q_1|$. Thus $Z' = B \#_{T_Z} Q_2$ is a polytope such that $\partial Z' \simeq \partial X'$.

By the induction hypothesis and (i), since X' contains one fewer vertex than X, there is a geometric polyhedral complex Y' such that $X' \simeq Y'$ and $|Y'| = |Z'|$. Let $Y^* = Y' \cup W$. Then clearly $X^*(w_0) \simeq Y^*$. Removing the facets of Y^* corresponding to the facets $\tau(w_0, P)$ of $X^*(w_0)$, we obtain a polyhedral complex Y such that $X \simeq Y$ and $|Y| = |Z|$. This establishes (i), provided that each cell of Y is convex. We now show that this is the case.

Let $X'_1, \ldots, X'_k \in \mathcal{P}_{X'}$ denote the d-polyhedra of X' such that $|X'_i| \subseteq |X|$ for some $X \in \mathcal{P}_X(w_0)$. We will let X_i denote the unique d-polyhedron of X such that $|X'_i| \subseteq |X_i|$. For all $i, j = 1, \ldots, k$, since $w_0 \in X_i \cap X_j$ and X is a polyhedral complex, we must have $X'_i \cap X'_j \subseteq \partial X'$, for otherwise X_i and X_j would intersect in more than a unique common face. If w_0 is a strong solitary vertex of X, then $k = 1$. Thus the single polyhedron $Y_1 \in \mathcal{P}_Y$ corresponding to X_1 is clearly convex from the above construction.

Now suppose that w_0 is a strong shedding vertex of X. Note that any shedding vertex w_0 of X must satisfy $\mathrm{cst}_X(w_0) \cap \partial X = \mathrm{cst}_{\partial X}(w_0)$, for otherwise $|X \smallsetminus w_0|$ would not be homeomorphic to B^d. Thus we must have $X'_i \cap \partial X' = \tau(w_0, X_i)$. So $X'_1, \ldots X'_k$ satisfy the hypotheses of (ii). Let $Y'_i \in \mathcal{P}_{Y'}$ be the polyhedron corresponding to X'_i and $G_i \in Y'$ the facet corresponding to $\tau(w_0, X_i)$. Let $Y_i \in \mathcal{P}_Y$ be the unique polyhedron such that $|Y'_i| \subseteq |Y_i|$. Since our induction hypothesis is enhanced by (ii), we may assume that the angles $\alpha(H_j, G_i)$ are arbitrarily small for each i and each face $H_j \in Y'_i$ strongly adjacent to G_i. Since the cells $|Y'_i|$ are convex by induction and form arbitrarily small angles with G_i, we may ensure that on removing G_i the resulting cell $|Y_i| \in Y$ is convex. Finally, if $\sigma \in Y$ is a cell not having

any of the $|Y_i'|$ as a subset, then either σ is a d-simplex (hence convex) or $\sigma \in Y'$. In the latter case σ is convex because Y' is a geometric polyhedral complex.

Now we must show that (ii) holds. Let $X_1, \ldots, X_\ell \in \mathcal{P}_X$ be a collection of polytopes satisfying the hypotheses of (ii), and let $Y_1, \ldots Y_\ell \in \mathcal{P}_Y$ denote the corresponding polytopes of Y. If $Y_i \in \mathcal{P}_{Y'}$ then we obtain the conclusion of (ii) from the induction hypothesis. Now suppose $Y_i \notin \mathcal{P}_{Y'}$. Then $Y_i \in \mathcal{P}_W$. If the A in the above construction is a simplex, then $\ell = 1$. Let v be the vertex of A not contained in T_A. Then we may clearly choose v arbitrarily close to the facet $Y_i \cap \partial Y$. If A is not a simplex, then because the point x is finite and lies on the side of $\text{aff}(T_A)$ opposite to that of A, $x \in \xi(C, F_C)$. In particular $\xi(C, F_C)$ is a cone with base F_C and apex $x \in \mathbb{R}^{d+1}$. By choosing the Schlegel point y in the above construction arbitrarily close to x, we may ensure that all such angles $\alpha(F_i, F_j)$ in the projection W are arbitrarily small.

Finally, that Y may be chosen to be rational is an immediate consequence of the methods of the proof. Broadly speaking, we obtained Y by first using Steinitz's theorem to produce polytopes Z and A, and then manipulating these polytopes using projective transformations. But from the proof of Steinitz's theorem, we may take both Z and A to be rational. By then using only rational affine and projective transformations in the above constructions, we may in fact obtain a *rational* geometric polyhedral complex Y such that $X \simeq Y$. □

It is straightforward to show that all 2-polyhedral complexes are vertex truncatable. Thus Theorem 5.2.1 implies that *all* 2-polyhedral complexes are geometrically realizable. Note that every 2-polyhedral complex is a 2-connected plane graph. However, the converse is not true. In fact, if a 2-connected plane graph is not a polyhedral complex (i.e. if the intersection of two faces is more than a unique common edge or vertex of both), then it clearly does not admit an embedding such that all faces are strictly convex.

Therefore we obtain necessary and sufficient conditions for a 2-connected plane graph G to admit an isotopic embedding G' such that all faces of G' are strictly convex. Specifically, a 2-connected plane graph G has a strictly convex isotopic embedding if and only if G is a

2-polyhedral complex. From this we recover Tutte's theorem [T1].

Finally, we note that the condition $d \leq 3$ in the statement of Theorem 5.2.1 plays a crucial role. Namely, it allows us to invoke Steinitz's theorem (for $d = 3$). For example, by Steinitz's theorem a simplicial 2-sphere is always (combinatorially) isomorphic to the boundary of a 3-polytope. In higher dimensions the analogous statement is not true (see [GS]).

CHAPTER 6

A Sequentially Convex Embedding Algorithm for Plane Triangulations

For a plane graph G with n vertices and an ordering of the vertices $\mathbf{a} = (a_1, \ldots, a_n)$, we define a sequence of plane graphs $G_0(\mathbf{a}), \ldots, G_n(\mathbf{a})$ recursively by $G_n(\mathbf{a}) = G$ and $G_{i-1}(\mathbf{a}) = G_i(\mathbf{a}) - \{a_i\}$. We will write G_i for $G_i(\mathbf{a})$ when $\mathbf{a}$ is understood. If v is a vertex of G_i then we let $d_i(v)$ denote the degree of v *in the graph* G_i.

A *plane triangulation* is a 2-connected plane graph G such that each bounded face of G has exactly 3 vertices. Note in particular that if G is a plane triangulation then $\mathbf{F}(G)$ is homeomorphic to a 2-ball. A boundary vertex v of a plane triangulation G is a *shedding vertex* of G if $G - \{v\}$ is a plane triangulation. Let G be a plane triangulation with n vertices. A vertex sequence $\mathbf{a} = (a_1, \ldots, a_n)$ is called a *shedding sequence* for G if a_i is a shedding vertex of $G_i(\mathbf{a})$ for all $i = 4, \ldots, n$. We have the following useful lemma.

Lemma 6.0.2 ([FPP], section 2) *Let G be a plane triangulation. Then for any boundary edge uv of G, there is a shedding sequence $\boldsymbol{a} = (a_1, \ldots, a_n)$ for G such that $u = a_1$ and $v = a_2$.*

We say that a *strictly* convex polygon $P \subset \mathbb{R}^2$ with edge e is *projectively convex* with respect to e if P is contained in a triangle having e as an edge. A shedding sequence $\mathbf{a} = (a_1, \ldots, a_n)$ for a plane triangulation G is a *convex shedding sequence* if the region $\mathbf{F}(G_i(\mathbf{a}))$ is a projectively convex polygon with respect to the edge a_1a_2 for all $i = 3, \ldots, n$. A geometric embedding G' of G is *sequentially convex* if G' has a convex shedding sequence. In this section we describe an algorithm for producing a sequentially convex embedding of any

given plane triangulation, such that the vertices of the embedding lie in a polynomially-sized integer grid.

6.1 A Rational Embedding

First we address a much easier question: How does one obtain a sequentially convex embedding of G in $\mathbb{Q}^2$ (that is, with vertex coordinates *rational*)? We describe a simple construction that produces such an embedding. The method used to accomplish this easier task will provide part of the motivation and intuition behind the more involved method we will use to obtain a polynomially-sized embedding in $\mathbb{Z}^2$.

Theorem 6.1.1 *Let G be a plane triangulation with n vertices and boundary edge uv, and let $\boldsymbol{a} = (a_1, \ldots, a_n)$ be a shedding sequence for G with $u = a_1$, $v = a_2$. Then G has a geometric embedding G' in $\mathbb{Q}^2$, such that the corresponding sequence $\boldsymbol{a'} = (a'_1, \ldots, a'_n)$ is a convex shedding sequence for G'.*

Proof. We proceed by induction on n. If $n = 3$ then any triangle with rational coordinates is a sequentially convex embedding of G. If $n > 3$, then by the inductive hypothesis there is an embedding G'_{n-1} of G_{n-1} in $\mathbb{Q}^2$ such that $(a'_1, \ldots, a'_{n-1})$ is a convex shedding sequence for G'_{n-1}. We may assume that $u'v'$ lies along the x-axis, and that u' lies to the left of v', and that G'_{n-1} lies in the upper half-plane. Let $w_1, \ldots, w_k$ denote the neighbors of a_n in G, and let $w'_1, \ldots, w'_k$ denote the corresponding vertices of G'_{n-1}, ordered from left to right. If $w'_1 \neq u'$, then let z'_1 denote the left boundary neighbor of w'_1. Similarly, if $w'_k \neq v'$, let z'_2 denote the right boundary neighbor of w'_k.

For adjacent vertices u and v, we will denote the slope of the edge uv by $s(uv)$. Similarly, we will denote the slope of a line ℓ by $s(\ell)$. Consider the lines $\ell_1, \ell_2, \ell_3, \ell_4$ spanned by the edges $z'_1w'_1$, $w'_1w'_2$, $w'_{k-1}w'_k$, and $w'_kz'_2$, respectively. If $w'_1 = u'$, we may take ℓ_1 to be any non-vertical line passing through u', with slope satisfying $s(\ell_1) > s(\ell_2)$. Similarly, if $w'_k = v'$, we may take ℓ_4 to be any non-vertical line passing through v', with slope satisfying $s(\ell_3) > s(\ell_4)$.

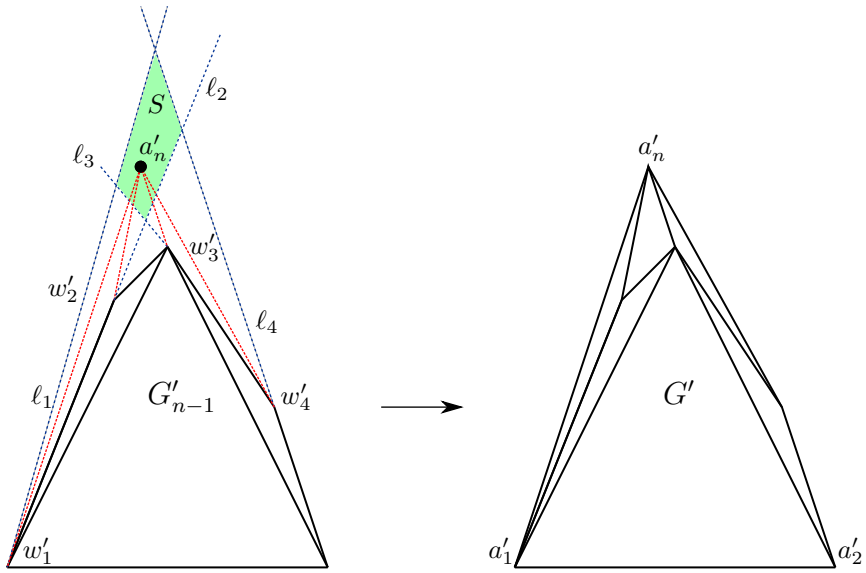

Figure 6.1: The new vertex a'_n, chosen as a rational point of the set S.

Let $A_1, A_4 \subset \mathbb{R}^2$ denote the open half-planes below the lines ℓ_1 and ℓ_4, respectively, and let A_2 and A_3 denote the open half-planes above the lines ℓ_2 and ℓ_3, respectively. Since $\mathbf{F}(G'_{n-1})$ is projectively convex with respect to $u'v'$, the slopes of the lines ℓ_i must satisfy $s(\ell_1) > s(\ell_2) > s(\ell_3) > s(\ell_4)$. Thus the region $S = A_1 \cap A_2 \cap A_3 \cap A_4$ is non-empty. See Figure 6.1. Since each set A_i is open, the set S is open, so we may choose a rational point in S, call it a'_n. For each $j = 1, \ldots, k$ add a straight line segment e_j between a'_n and the vertex w'_j. Since a'_n lies in the region above the lies ℓ_2 and ℓ_3, each line segment e_j will intersect G'_{n-1} only in the vertex w'_j.

Let G' denote the plane graph obtained from G'_{n-1} by adding the vertex a'_n and the edges e_j. Then G' is clearly a geometric embedding of $G_n = G$, such that each vertex a_i corresponds to a'_i, for $i = 1, \ldots, n$. Furthermore, since a'_n lies in the region below the lines ℓ_1 and ℓ_4, the region $\mathbf{F}(G')$ is projectively convex with respect to $u'v'$. From this, together with the fact that $(a'_1, \ldots, a'_{n-1})$ is a convex shedding sequence for G'_{n-1}, we have that $\mathbf{a}' = (a'_1, \ldots, a'_n)$ is a convex shedding sequence for G'. □

6.2 The Shedding Tree of a Plane Triangulation

Now we turn to the problem of embedding the triangulation G on an integer grid. The idea behind our construction is roughly as follows. We start with a triangular base whose the horizontal width is very large. We then show that, because this horizontal width is large enough, we may add each vertex in a manner similar to that used in the proof of Theorem 6.1.1, and we will always have enough room to find an acceptable integer coordinate. The crucial part of the construction is the careful method in which we add each new vertex. In particular, there are two distinct methods for adding the new vertex a_i, depending on whether $d_i(a_i) = 2$ or $d_i(a_i) > 2$. To facilitate the proper placement of the vertices a_i with $d_i(a_i) = 2$, we will appeal to a certain tree structure determined by the shedding sequence $\mathbf{a}$. We introduce the following definitions.

Let G be a plane triangulation with shedding sequence $\mathbf{a}$. We may assume that G is embedded geometrically as in Theorem 6.1.1. Proceeding recursively, we define a binary tree $T = T(G, \mathbf{a})$, such that the nodes of T are edges of G, and the edges of T correspond to faces of G.

Let ν_2 denote the edge of G containing vertices a_1, a_2, and let T_2 be the tree consisting of the single node ν_2. Now let $3 \leq i \leq n$, and let ν_i, ν_i' denote the boundary edges of $G_i(\mathbf{a})$ immediately to the left and right of a_i, respectively (this is well defined because G is embedded as in Theorem 6.1.1). Assume that we have already constructed T_{i-1}, and that all boundary edges of $G_{i-1}(\mathbf{a})$ are nodes of T_{i-1}. Let ξ, ξ' be the boundary edges of $G_{i-1}(\mathbf{a})$ such that ξ shares a face with ν_i, and ξ' shares a face with ν_i'.

Define $T_i = T_i(G, \mathbf{a})$ to be the tree obtained from T_{i-1} by adding ν_i and ν_i' as nodes, and adding the edges (ξ, ν_i) and (ξ', ν_i'), designated *left* and *right*, respectively. Then clearly all boundary edges of $G_i(\mathbf{a})$ are vertices of T_i. Thus we have a recursively defined sequence of trees $(T_2, T_3, \ldots, T_n)$, and nodes $(\nu_2, \nu_3, \nu_3', \ldots, \nu_n, \nu_n')$, such that T_i has nodes $\nu_2, \nu_3, \nu_3', \ldots, \nu_i, \nu_i'$. We define $T = T_n$ (see Figure 6.2). Note that for all $i = 2, \ldots, n$, we have $T_i(G, \mathbf{a}) = T(G_i(\mathbf{a}), (a_1, \ldots, a_i))$.

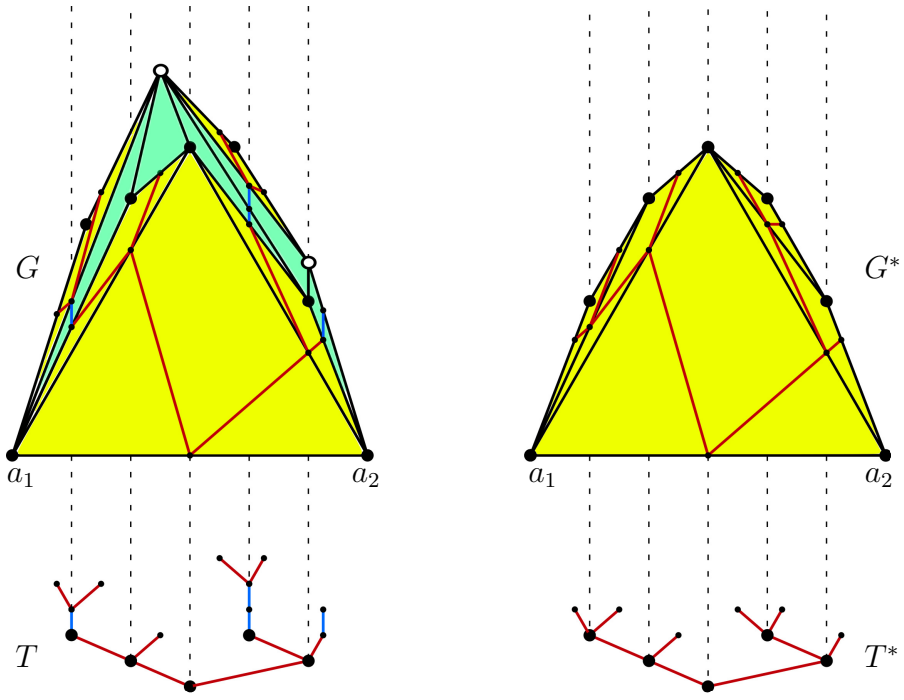

Figure 6.2: The triangulations G and G^*, together with corresponding trees T and T^*. Each node of T corresponds to an edge of G, and similarly for T^* and G^*. The tree T^* is obtained from T by contracting the blue edges. The large nodes of T are the internal nodes of T^*, and correspond to the vertices of G^* other than a_1 and a_2.

Let $\mathcal{R} = \{i \in \{1, \ldots, n\} \mid d_i(a_i) \leq 2\}$, and define a set of edges

$$E_{\mathcal{R}} = \{(\xi, \sigma) \in E(T) \mid \sigma = \nu_j \text{ or } \sigma = \nu'_j \text{ for some } j \notin \mathcal{R}\}.$$

Let $T_i^* = T_i^*(G, \mathbf{a})$ be the tree obtained from T_i by contracting all edges in $E_{\mathcal{R}}$ (shown in blue in Figure 6.2). Note that each T_i^* is a *full* binary tree. We define $T^* = T_n^*$, and we call the T_i's the *reduced trees* of G. If $d_i(a_i) = 2$ for all $i \geq 3$, then $T^* = T$.

A fundamental idea behind our integer grid embedding is that the reduced tree T^* contains the crucial information needed for carrying out the embedding of G. For example, for each $i \in \mathcal{R} - \{1, 2\}$, the vertex a_i corresponds to an internal node of T^* (shown as large dots in Figure 6.2). Thus the structure of T^* tells us how to horizontally space the vertices a_i with $d_i(a_i) = 2$, and how to choose the slopes of the boundary edges adjacent to them.

On the other hand, when adding vertices a_i with $d_i(a_i) > 2$, our construction will have the property that the boundary slopes will be perturbed only slightly, and the horizontal distances between vertices will only increase. Furthermore, throughout our entire construction the total horizontal width of the embedding will remain fixed.

If a plane triangulation G has the reduced trees $T_2^*, \ldots, T_n^*$, then there is a natural construction which produces a triangulation G^* and shedding sequence $\mathbf{a}^*$, such that T_i^* is isomorphic to $T_i(G^*, \mathbf{a}^*)$, for all $i = 2, \ldots, n$. We call the resulting triangulation G^* the *reduced triangulation* of G (see Figure 6.2). Furthermore, the construction of G^* may be carried out on a grid with size polynomial in the number of vertices of G^*. This is the content of the next theorem. It can be though of as a special instance of our main result, in the case that $d_i(a_i) = 2$ for all $i \geq 3$.

Lemma 6.2.1 *Let $(t_2, \ldots, t_n)$ be a sequence of full binary trees, such that t_{i-1} is a subtree of t_i, and t_i has $1 + 2(i-2)$ nodes, for all $i = 2, \ldots, n$. Then there is a sequentially convex plane triangulation H with n vertices, embedded in a $2(n-2) \times \binom{n-1}{2}$ integer grid, and a convex shedding sequence $\boldsymbol{a}$ for H, such that t_i is isomorphic to $T_i(H, \boldsymbol{a})$ for all $i = 2, \ldots, n$.*

Proof. Let m and m' denote the number of internal nodes of t_n to the left and right of the root node, respectively. Note that $m + m' + 3 = n$. Without loss of generality we may assume $m \leq m'$. For $-m \leq k \leq m' + 1$, we define

$$x_k = k, \qquad y_k = \binom{m'+2}{2} - \binom{|k|+1}{2}.$$

Additionally, define

$$x_{-m-1} = -x_{m'+1}, \qquad y_{-m-1} = 0.$$

We have now defined n pairs (x_k, y_k), which will serve as the (x, y)-coordinates of the vertices of the triangulation H.

Note that for all $i = 3, \ldots, n$, since t_i is full and contains two more nodes than t_{i-1}, it follows that t_i contains exactly one more internal node than t_{i-1}. Let $(\xi_3, \ldots \xi_n)$ denote

the sequence of internal nodes so obtained. Note that ξ_3 is the root node of all the trees t_i, and t_2 consists of the single node ξ_3. The nodes of t_n may be linearly ordered by a depth-first search on t_n, such that left nodes are visited before right nodes. Call this order Ω. This restricts to a linear order on the internal nodes of t_n, which induces a permutation $\omega : \{3, \ldots, n\} \to \{3, \ldots, n\}$. That is, node ξ_i has position $\omega(i)$ in the order Ω. Then we define a sequence of points $\mathbf{a} = (a_1, a_2, \ldots, a_n)$ by

$$\begin{aligned} a_1 &= (x_{-m-1}, y_{-m-1}), \\ a_2 &= (x_{m'+1}, y_{m'+1}), \\ a_i &= (x_{\sigma(i)-m-3}, y_{\sigma(i)-m-3}) \quad \text{for } 3 \leq i \leq n. \end{aligned}$$

For $i = 3, \ldots, n$, let $\omega_i : \{3, \ldots, i\} \to \{3, \ldots, i\}$ denote the permutation induced by restricting the order Ω to the internal nodes $(\xi_3, \ldots, \xi_i)$ of t_i. Note that $\omega = \omega_n$. For each $i = 3, \ldots, n$, we would like to determine the internal nodes of t_i that immediately precede and succeed ξ_i in the order ω_i. For this purpose, we define functions $f, g : \{3, \ldots, n\} \to \{1, \ldots, n\}$ by

$$f(i) = \begin{cases} \omega_i^{-1}(\omega_i(i) - 1) & \text{if } \omega_i(i) > 3, \\ 1 & \text{otherwise} \end{cases}$$

$$g(i) = \begin{cases} \omega_i^{-1}(\omega_i(i) + 1) & \text{if } \omega_i(i) < i, \\ 2 & \text{otherwise.} \end{cases}$$

We may now define a sequence of plane triangulations $H_1, \ldots, H_n$ recursively. Let H_1 consist of the single vertex a_1, and let H_2 consist of the vertices a_1, a_2 and the line segment a_1a_2. Now let $3 \leq i \leq n$, and suppose we have constructed H_{i-1}. We obtain H_i by adding the vertex a_i and the line segments $a_ia_{f(i)}$ and $a_ia_{g(i)}$ to H_{i-1}. This completes the construction of the graphs $H_2, \ldots, H_n$. We write $H = H_n$.

We now check that $\mathbf{a} = (a_1, \ldots, a_n)$ is a convex shedding sequence for H. By construction, we have immediately that $H_{i-1} = H_i - \{a_i\}$ for all $i = 2, \ldots, n$, and furthermore $d_i(a_i) = 2$ for all $i \geq 3$. For $k \geq -m$, the slope of the edge between adjacent vertices (x_k, y_k)

and (x_{k+1}, y_{k+1}) of H is

$$\frac{y_{k+1}-y_k}{x_{k+1}-x_k} = \binom{|k|+1}{2} - \binom{|k+1|+1}{2} = \begin{cases} -(k+1) & \text{if } k \geq 0, \\ -k & \text{if } k < 0. \end{cases}$$

Additionally, the slope of the edge between (x_{-m-1}, y_{-m-1}) and (x_{-m}, y_{-m}) is

$$\begin{aligned} &\frac{y_{-m}-y_{-m-1}}{x_{-m}-x_{-m-1}} = \frac{y_{-m}}{-m+(m'+1)} = \frac{1}{m'-m+1}\left(\binom{m'+2}{2} - \binom{m+1}{2}\right) \\ &= \frac{(m'+2)(m'+1)-(m+1)m}{2(m'-m+1)} = \frac{(m'-m+1)(m'+m+2)}{2(m'-m+1)} \\ &= \frac{m'+m+2}{2} \geq \frac{m+m+2}{2} = m+1. \end{aligned}$$

Thus the boundary edge slopes of H are strictly decreasing from left to right. Since $d_i(a_i) = 2$ for all $i \geq 3$, this implies that the boundary edge slopes of each H_i are strictly decreasing from left to right. It follows that H_i is plane triangulation and $\mathbf{F}(H_i)$ is projectively convex, for all $i \geq 3$. Hence $\mathbf{a}$ is a convex shedding sequence for H.

To see that t_i is isomorphic to $T_i(H, \mathbf{a})$ for all $i = 2, \ldots, n$, we construct an explicit isomorphism. We define a map $\psi_2 : t_2 \to T_2(H, \mathbf{a})$ by $\psi_2(\xi_3) = a_{f(3)}a_{g(3)} = a_1a_2$, which is trivially an isomorphism. For $i \geq 3$, and $j = 3, \ldots, i$, let ξ_j^- and ξ_j^+ denote the left and right child, respectively, of the internal node ξ_j of t_i. We define a map $\psi_i : t_i \to T_i(H, \mathbf{a})$ by

$$\begin{aligned} \psi_i(\xi_j) &= a_{f(j)}a_{g(j)}, \\ \psi_i(\xi_j^-) &= a_j a_{f(j)}, \\ \psi_i(\xi_j^+) &= a_j a_{g(j)} \qquad \text{for } j = 3, \ldots, i. \end{aligned}$$

From the definition of the order Ω and the resulting functions f and g, it is straightforward to check that ψ_i is well-defined and bijective. Since the triangle $(a_j a_{f(j)} a_{g(j)})$ is a face of H_j for all $j = 3, \ldots, i$, the pairs $(a_{f(j)}a_{g(j)}, a_j a_{f(j)})$ and $(a_{f(j)}a_{g(j)}, a_j a_{g(j)})$ are edges of $T_i(H, \mathbf{a})$. Thus ψ_i is clearly a tree isomorphism. We may think of ψ_n as providing a correspondence between the internal node ξ_i and the vertex a_i (whose neighbors in H_i are $a_{f(i)}$ and $a_{g(i)}$), for all $i = 3, \ldots, n$ (See Figure 6.2).

Finally, the width of the grid is

$$x_{m'+1} - x_{-m-1} = 2x_{m'+1} = 2(m'+1) \leq 2(n-2),$$

and the height of the grid is

$$y_0 = \binom{m'+2}{2} \leq \binom{n-1}{2}.$$

Therefore H is embedded in an integer grid of size $2(n-2) \times \binom{n-1}{2}$. □

6.3 The Integer Grid Embedding

Given a plane triangulation G with n vertices and shedding sequence $\mathbf{a}$, let $t_i = T_i^*(G, \mathbf{a})$ denote the reduced trees of G. Let ρ denote the unique increasing bijection from $\mathcal{R}$ to $\{1, \ldots, R\}$, where $\mathcal{R}$ is the subset of $\{1, \ldots, n\}$ defined above, and $R = |\mathcal{R}|$. Note that $1, 2, 3 \in \mathcal{R}$, so $\rho(i) = i$ for $i \leq 3$. Define a map $h : \{1, \ldots, n\} \to \{1, \ldots, R\}$ by taking $h(i)$ to be the unique index for which

$$\rho^{-1}(h(i)) \leq i < \rho^{-1}(h(i)+1).$$

In particular, if $i \in \mathcal{R}$ then $h(i) = \rho(i)$.

The sequence of distinct trees $t_{\rho^{-1}(1)}, \ldots, t_{\rho^{-1}(R)}$ satisfies the hypotheses of Theorem 6.2.1. Therefore we let G_i^* denote the sequentially convex triangulation constructed from $t_{\rho^{-1}(i)}$ as in Theorem 6.2.1. We call the G_i^* the *reduced triangulations* of G, and write $G^* = G_R^*$ (See Figure 6.2). Let $\mathbf{a}^* = (a_1^*, \ldots, a_R^*)$ denote the corresponding convex shedding sequence of G^* produced by Theorem 6.2.1. Note that each vertex a_i^* has degree 2 in G_i^*. So we may think of G^* as being obtained from G by "throwing away" all vertices a_i for which $d_i(a_i) > 2$. It was this property that originally motivated our definition of G^*.

As we will see, the triangulation G^* will tell us exactly how to add vertices of degree 2, in our construction of a sequentially convex embedding of G. A particular property of the reduced triangulations makes this possible. Namely, for any boundary edge e of G_i, there is a *corresponding boundary edge* e^* of $G_{h(i)}^*$, which we define as follows. First note that from the definitions, and the construction of Theorem 6.2.1, the trees $T_i(G^*, \mathbf{a}^*)$ and $T_i^*(G, \mathbf{a})$ are isomorphic. Thus we may think of an edge of G_i^* (which is a node of $T_i(G^*, \mathbf{a}^*)$) as a node of $T_i^*(G, \mathbf{a})$. For a boundary edge e of G_i, we denote by e^* the unique boundary edge of

$G^*_{h(i)}$, thought of as a node of $T^*_i(G, \mathbf{a})$, that is identified with the node e of $T_i(G, \mathbf{a})$ upon contracting the edges in the set $E_{\mathcal{R}}$.

Theorem 6.3.1 *Let G be a plane triangulation with n vertices and boundary edge uv, and let $\boldsymbol{a} = (a_1, \ldots, a_n)$ be a shedding sequence for G with $u = a_1$, $v = a_2$. Then G has a geometric embedding G' in a $4n^3 \times 8n^5$ integer grid, such that the corresponding sequence $\boldsymbol{a}' = (a'_1, \ldots, a'_n)$ is a convex shedding sequence for G'.*

Proof. We recursively construct a sequence of graphs $G'_1, \ldots, G'_n$, and a sequence of vertices $a'_1, \ldots, a'_n$, such that each G'_i is a geometric embedding of G_i with convex shedding sequence $\mathbf{a}' = (a'_1, \ldots, a'_i)$, where a'_i is the vertex of G' corresponding to a_i. Let G^*_i denote the reduced triangulations of G, and let $\mathbf{a}^* = (a^*_1, \ldots, a^*_{|R|})$ denote the corresponding shedding sequence for G^*. Let m' denote the number of vertices of G^* lying between a^*_3 and a^*_2, and m the number of vertices lying between a^*_1 and a^*_3. Then $m' + m + 3 = n$. Without loss of generality we may assume that $m \leq m'$.

For points $v_1, v_2 \in \mathbb{R}^2$ and $e = v_1v_2$ the line segment between them, we write

$$x(e) = |x(v_1) - x(v_2)| \qquad \text{and} \qquad y(e) = |y(v_1) - y(v_2)|.$$

We first scale G^* to obtain a much larger triangulation, which we will use to construct the triangulations G'_i. We choose the scaling factors large enough so that we will have "enough room" to carry out our constructions. Specifically, let $\alpha = 2n^2 + n + 1$ and $\beta = 2n\alpha$. We define Z_i to be the result of scaling G^*_i by a factor of α in the x dimension and β in the y dimension. That is, for each $i = 1, \ldots, R$, we define $z_i = (\alpha x(a^*_i), \beta y(a^*_i))$. Then $\mathbf{z} = (z_1, \ldots, z_R)$ is the shedding sequence for Z_R corresponding to $\mathbf{a}^*$. We write $Z = Z_R$.

Since $G'_i \sim G_i$ (as we verify below) the edges of G'_i and G_i are in correspondence. Thus every boundary edge e of G'_i corresponds to a boundary edge e^* of $G^*_{h(i)}$, as defined above. We write $Z(e)$ for the edge of $Z_{h(i)}$ corresponding to e^*. Note that if e^* has slope s, then $Z(e)$ has slope $\frac{\beta}{\alpha}s$. In particular, since $m' + 1$ is the largest magnitude of the slope of any edge of G^*, we see that $\frac{\beta}{\alpha}(m' + 1) = 2n(m' + 1)$ is the largest magnitude of the slope of any

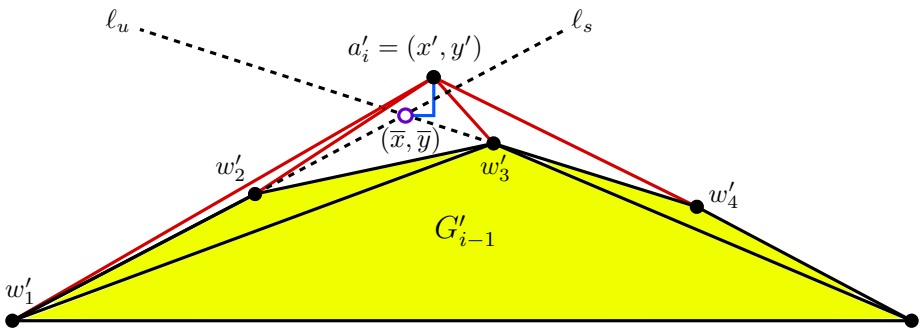

Figure 6.3: The construction of vertex a'_i when $d_i(a_i) > 2$. In this example, $d_i(a_i) = 4$.

edge of Z. Let M denote this slope, and note that $M \leq 2n^2$. Note also that the absolute difference of two boundary edge slopes of Z is at least $\frac{\beta}{\alpha} = 2n$. It follows immediately that for each $i = 1, \ldots, R$, the absolute difference of two boundary edge slopes of Z_i is at least $2n$.

Define $a'_1 = z_1$ and $a'_2 = z_2$. Take G'_1 to consist of the single vertex a'_1, and take G'_2 to consist of the vertices a'_1, a'_2, together with the line segment $a'_1 a'_2$. Now let $3 \leq i \leq n$, and suppose we have constructed G'_{i-1}. To define a'_i, we consider two cases, namely whether $d_i(a_i) = 2$ or $d_i(a_i) > 2$.

Construction of a'_i, in the case $d_i(a_i) > 2$. If $d_i(a_i) > 2$, then let $w_1, \ldots, w_k$ denote the neighbors of a_i in G_i, and let $w'_1, \ldots, w'_k$ denote the corresponding vertices of G'_{i-1}, ordered from left to right. Let s denote the slope of the edge $w'_1 w'_2$, and u the slope of the edge $w'_{k-1} w'_k$. Let ℓ_s denote the line of slope s containing the point w'_1, and ℓ_u the line of slope u containing the point w'_k. We denote by $(\overline{x}, \overline{y})$ the point of intersection of the lines ℓ_s and ℓ_u. Let $x' = \lceil \overline{x} \rceil$ and $\gamma = x' - \overline{x}$, and let $y' = \lceil \overline{y} \rceil + \lfloor \gamma s \rfloor + 1$. We now define $a'_i = (x', y')$ (see Figure 6.3). We obtain G'_i from G'_{i-1} by adding the vertex a'_i, together with all line segments between a'_i and the vertices $w'_1, \ldots, w'_k$.

Construction of a'_i, in the case $d_i(a_i) = 2$. If $d_i(a_i) = 2$, then let Δ be the triangle of Z_i containing $z_{\rho(i)}$. Let w_1, w_2 denote the boundary neighbors of a_i in G_i, and let w'_1, w'_2 denote the corresponding vertices of G'_{i-1}, so that w'_1 lies to the left of w'_2. We are going to construct a triangle Δ', such that Δ' is the image of Δ under an affine map which is the composition

of a uniform scaling and a translation. Furthermore, we will place Δ' in a specific position with respect to the triangulation G'_{i-1}. In particular, if v_1, v_2, v_3 denote the vertices of Δ', we require that $x(v_1) = x(w'_1)$, $v_2 = w'_2$, and $x(v_1) < x(v_3) < x(v_2)$ (see Figure 6.4). It is easily verified that these conditions, together with the requirement that Δ' is a scaled, translated copy of Δ, determine the vertices v_1, v_2, v_3 of Δ' uniquely.

Let $\eta = y(v_1) - y(w'_1)$. Let b_1 and b_2 denote the left and right boundary neighbors, respectively, of $z_{\rho(i)}$ in $Z_{\rho(i)}$. Notice that the edge $b_1b_2 = Z(w'_1w'_2)$ is a boundary edge of $Z_{h(i-1)}$. We will now define a ratio κ, which describes how far away $z_{\rho(i)}$ is from b_2, in the x direction. That is, we define

$$\kappa = \frac{x(b_2) - x(z_{\rho(i)})}{x(b_2) - x(b_1)}.$$

Using η and κ will allow us to appropriately define the new vertex a'_i. We begin by letting $\overline{v_3} = (x(v_3), y(v_3) - \kappa\eta)$. So $\overline{v_3}$ is obtained by pushing v_3 up or down, by an amount determined by κ.

Note that z_3 is the apex of Z_i for all $i = 1, \ldots, R$, and $x(z_3) = \alpha x(a^*_3) = \alpha \cdot 0 = 0$. So if $x(z_{\rho(i)}) < 0$ then $x(z_{\rho(i)})$ lies to the left of the apex of Z_i, and if $x(z_{\rho(i)}) > 0$ then $z_{\rho(i)}$ lies to the right of the apex. With this understanding, we set

$$v'_3 = \begin{cases} (\lfloor x(\overline{v_3}) \rfloor, \lceil y(\overline{v_3}) \rceil) & \text{if } x(z_{\rho(i)}) \leq 0, \\ (\lceil x(\overline{v_3}) \rceil, \lceil y(\overline{v_3}) \rceil) & \text{if } x(z_{\rho(i)}) > 0. \end{cases}$$

We now define $a'_i = v'_3$. We obtain G'_i from G'_{i-1} by adding the vertex a'_i and the two line segments between a'_i and the vertices w'_1, w'_2.

Verification of the construction. We have now explicitly described the construction. It remains to show that the above constructions actually produce a *convex* shedding sequence $\mathbf{a}' = (a'_1, \ldots, a'_n)$ for G', and that G' lies in the grid size indicated. Clearly the horizontal width of the grid remains constant throughout the construction. Specifically, the triangulations $G'_1, G'_2, \ldots, G'_n$ all have the same width $\alpha 2(n-2)$, which is the width of the Z_i. So to show that the construction is sequentially convex, and that the bound on the height of

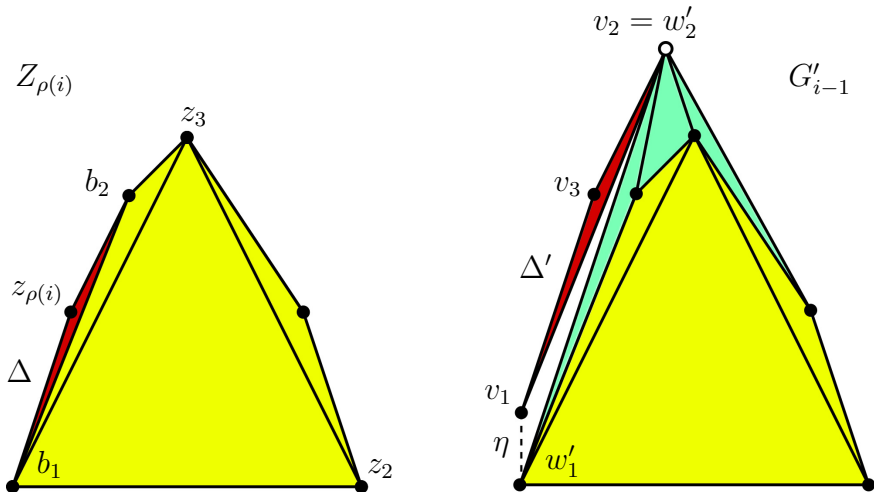

Figure 6.4: The first stage of the construction of vertex a_i' when $d_i(a_i) = 2$. The red triangles Δ and Δ' differ by a uniform scaling and a translation.

G' is correct, we will calculate how the boundary slopes are modified when we add the new vertex a_i' in the two cases $d_i(a_i) = 2$ and $d_i(a_i) > 2$.

For each $3 \le i \le n$, let $\mathbf{C}(i)$ denote the following three-part claim:

$\mathbf{C}(i, 1)$. For every boundary edge e of G_i', we have $x(e) \ge x(Z(e))$.

$\mathbf{C}(i, 2)$. For every boundary edge e of G_i', the slopes of e and $Z(e)$ differ by at most i.

$\mathbf{C}(i, 3)$. $G_i' \sim G_i$ and G_i' has convex shedding sequence $(a_1', \dots a_i')$.

To prove $\mathbf{C}(i)$ for each $i = 3, \dots, n$, we proceed by induction on i.

For $i = 3$, note that $a_1' = z_1$, $a_2' = z_2$, and the vertex a_3' is chosen so that in particular the triangle $(a_1'a_2'a_3')$ is a scaling of the triangle $(z_1z_2z_3) = \mathbf{F}(Z_3)$. This implies that $a_3' = z_3$. Thus $G_3' = Z_3$, which immediately establishes $\mathbf{C}(3)$.

Now let $i > 3$, and suppose that $\mathbf{C}(i-1)$ holds. As in the construction, we consider separately the cases $d_i(a_i) > 2$ and $d_i(a_i) = 2$.

Verification of $\mathbf{C}(i)$ in the case $d_i(a_i) > 2$. Clearly $Z(w_1'a_i') = Z(w_1'w_2')$ and $Z(a_i'w_k') = Z(w_{k-1}'w_k')$. Note that $x(w_2') \le \overline{x} \le x(w_{k-1}')$, and thus $x(w_2') \le x(a_i') \le x(w_{k-1}')$. Therefore

$$x(w_1'a_i') = x(a_i') - x(w_1') \ge x(w_2') - x(w_1') = x(w_2'w_1') \ge x(Z(w_2'w_1')) = x(Z(w_1'a_i')),$$

where the last inequality follows from $\mathbf{C}(i-1,1)$. Similarly, we have $x(a_i'w_2') \geq x(Z(a_i'w_2'))$. Thus $\mathbf{C}(i,1)$ holds.

We now show that the slopes of the boundary edges of G_i' containing a_i' differ only slightly from the slopes s and u defined in the above construction. That is, let s' denote the slope of the line passing through w_1' and $a_i' = (x', y')$, and let u' denote the slope of the line passing through a_i' and w_k'. Since

$$y' - \overline{y} = (\lceil \overline{y} \rceil - y) + \lfloor \gamma s \rfloor + 1 \geq \lfloor \gamma s \rfloor + 1 > \gamma s,$$

we clearly have $s' - s > 0$. On the other hand,

$$\begin{aligned} s' - s &= \frac{y' - y(w_1')}{\overline{x} + \gamma - x(w_1')} - s = \frac{y' - y(w_1') - (\overline{x} - x(w_1'))s - \gamma s}{\overline{x} + \gamma - x(w_1')} = \frac{(y' - \overline{y}) - \gamma s}{\overline{x} + \gamma - x(w_1')} \\ &= \frac{(\lceil \overline{y} \rceil - y) + (\lfloor \gamma s \rfloor - \gamma s) + 1}{\overline{x} + \gamma - x(w_1')} < \frac{2}{\overline{x} + \gamma - x(w_1')} \leq \frac{2}{\overline{x} - x(w_1')} \leq \frac{2}{x(w_2') - x(w_1')} \\ &\leq \frac{2}{x(Z(w_2'w_1'))} \leq \frac{2}{\alpha} \leq 1. \end{aligned}$$

In the last line we have invoked $\mathbf{C}(i-1,1)$.

Since $s > u$, we have $y' - \overline{y} > \gamma s > \gamma u$. From this, together with the fact that $x' \geq \overline{x}$, it follows that $u - u' > 0$. By $C(i-1,2)$, we have $|s| \leq M + (i-1)$. Thus

$$\begin{aligned} u - u' &= \frac{y(w_k') - \overline{y}}{x(w_k') - \overline{x}} - \frac{y(w_k') - y'}{x(w_k') - x'} \leq \frac{y(w_k') - \overline{y}}{x(w_k') - x'} - \frac{y(w_k') - y'}{x(w_k') - x'} = \frac{y' - \overline{y}}{x(w_k') - x'} \\ &= \frac{(\lceil \overline{y} \rceil - \overline{y}) + \lfloor \gamma s \rfloor + 1}{x(w_k') - x'} < \frac{|s| + 2}{x(w_k') - x'} \leq \frac{|s| + 2}{x(w_k') - x(w_{k-1}')} \leq \frac{|s| + 2}{x(Z(w_{k-1}'w_k'))} \\ &\leq \frac{|s| + 2}{\alpha} \leq \frac{M + (i-1) + 2}{\alpha} \leq \frac{2n^2 + i + 1}{\alpha} \leq \frac{2n^2 + n + 1}{\alpha} = 1. \end{aligned}$$

In the second line we have invoked $\mathbf{C}(i-1,1)$.

Let $Z(s)$ denote the slope of the edge $Z(w_1'w_2')$, and let $Z(u)$ denote the slope of $Z(w_{k-1}'w_k')$. By $\mathbf{C}(i-1,2)$, we have $|s - Z(s)| \leq i-1$ and $|u - Z(u)| \leq i-1$. Thus

$$|s' - Z(s)| \leq |s' - s| + |s - Z(s)| \leq 1 + (i-1) = i,$$

and similarly $|u' - Z(u)| \leq i$, so $\mathbf{C}(i,2)$ holds.

Since $s' - s > 0$ and $u - u' > 0$, each line segment $a'_i w'_j$ intersects G'_{i-1} only in the vertex w'_j, for all $j = 1, \ldots k$. Thus G'_i is a plane triangulation, and $G'_i \sim G_i$. It remains to show that $\mathbf{F}(G'_i)$ is projectively convex, in order to establish $\mathbf{C}(i,3)$. To do this, we will show that when the slope s is changed to s' for example, convexity is preserved at the vertex w'_1. That is, the slope s', while greater than s, is still less than the slope of the boundary edge to the left of the edge $w'_1 w'_2$.

Let $\hat{s}$ denote the slope of the boundary edge of G'_i adjacent and to the left of w'_1, if such an edge exists, and let $\hat{u}$ denote the slope of the boundary edge of G'_i adjacent and to the right of w'_k, if such an edge exists. Let $Z(\hat{s})$ and $Z(\hat{u})$ denote the boundary slopes of $Z_{h(i-1)}$ corresponding to $\hat{s}$ and $\hat{u}$, respectively. By $\mathbf{C}(i-1,2)$, we have $s - Z(s) < i - 1$ and $Z(\hat{s}) - \hat{s} < i - 1$. Thus

$$\begin{aligned}
\hat{s} - s &= (\hat{s} - Z(s)) - (s - Z(s)) \geq (\hat{s} - Z(s)) - (i-1) \\
&= (Z(\hat{s}) - Z(s)) - (Z(\hat{s}) - \hat{s}) - (i-1) \geq (Z(\hat{s}) - Z(s)) - 2(i-1) \\
&\geq 2n - 2i + 2 \geq 2.
\end{aligned}$$

We conclude that

$$\hat{s} - s' = (\hat{s} - s) - (s' - s) \geq (\hat{s} - s) - 1 \geq 2 - 1 = 1 > 0.$$

An analogous calculation shows that $u' - \hat{u} > 0$. Thus w'_1 and w'_k are convex vertices of G'_i. Because the region $\mathbf{F}(G'_{i-1})$ is projectively convex by $\mathbf{C}(i-1,3)$, we conclude that $\mathbf{F}(G'_i)$ is projectively convex. The sequence $(a'_1, \ldots, a'_{i-1})$ is a convex shedding sequence for G'_{i-1} by $\mathbf{C}(i-1,3)$, hence $(a'_1, \ldots, a'_i)$ is a convex shedding sequence for G'_i. Thus $\mathbf{C}(i,3)$ holds. We have now established $\mathbf{C}(i)$ in the case that $d_i(a_i) > 2$.

Verification of $\mathbf{C}(i)$ in the case $d_i(a_i) = 2$. First note that by $\mathbf{C}(i-1,1)$, we have

$$x(v_1 v_2) = x(w'_1 w'_2) \geq x(Z(w'_1 w'_2)) = x(b_1 b_2).$$

This implies that $x(w'_1 \overline{v_3}) = x(v_1 v_3) \geq x(b_1 z_{\rho(i)})$, because Δ' is a scaled, translated copy of Δ. Since $x(a'_i) = \lfloor x(\overline{v_3}) \rfloor$, and $x(b_1 z_{\rho(i)})$ and $x(w'_1)$ are integers, we also have

$$x(w'_1 a'_i) \geq x(b_1 z_{\rho(i)}) = x(Z(w'_1 a'_i)).$$

Similarly, we obtain $x(a_i'w_2') \geq x(z_{\rho(i)}b_2) = x(Z(a_i'w_2'))$. Thus $\mathbf{C}(i,1)$ holds.

We now consider two relevant slopes in the construction of G_i'. Let r denote the slope of the edge $w_1'w_2'$ of G_{i-1}', and let $Z(r)$ denote the slope of the corresponding edge $Z(w_1'w_2') = b_1b_2$ of $Z_{h(i-1)}$. Since the triangle Δ' is a scaled, translated copy of Δ, we see that $Z(r)$ is also the slope of the edge v_1v_2 of Δ'. Let $\varepsilon = r - Z(r)$, and note that

$$\varepsilon = \frac{y(w_2') - y(w_1')}{x(w_2') - x(w_1')} - \frac{y(v_2) - y(v_1)}{x(v_2) - x(v_1)} = \frac{y(w_2') - y(w_1')}{x(v_2) - x(v_1)} - \frac{y(w_2') - y(v_1)}{x(v_2) - x(v_1)} = \frac{\eta}{x(v_2) - x(v_1)}.$$

We consider another important pair of slopes arising in the construction of G_i', together with the corresponding pair of slopes of $Z_{\rho(i)}$. Specifically, let q_1 and q_2 denote the slopes of the line segments v_1v_3 and v_3v_2, respectively. Since Δ' is a scaled, translated copy of Δ, these slopes q_1 and q_2 are also the slopes of the boundary edges $b_1z_{\rho(i)}$ and $z_{\rho(i)}b_2$ of $Z_{\rho(i)}$, respectively. Let $\overline{q_1}$ and $\overline{q_2}$ denote the slopes of the line segments $\overline{v_3}w_1'$ and $\overline{v_3}w_2'$, respectively. We may think of $\overline{q_1}$ and $\overline{q_2}$ as *modifications* of the slopes q_1 and q_2, which arise from replacing the vertices v_1 and v_3 of Δ' with the vertices w_1' and $\overline{v_3}$, respectively. A consequence of our definitions is that $\overline{q_1} - q_1 = \overline{q_2} - q_2 = \varepsilon$. Indeed,

$$\begin{aligned}
\overline{q_1} - q_1 &= \frac{y(\overline{v_3}) - y(w_1')}{x(\overline{v_3}) - x(w_1')} - \frac{y(v_3) - y(v_1)}{x(v_3) - x(v_1)} = \frac{y(\overline{v_3}) - y(w_1')}{x(v_3) - x(v_1)} - \frac{y(v_3) - y(v_1)}{x(v_3) - x(v_1)} \\
&= \frac{(y(v_1) - y(w_1')) + (y(\overline{v_3}) - y(v_3))}{x(v_3) - x(v_1)} = \frac{\eta - \kappa\eta}{x(v_3) - x(v_1)} = (1-\kappa)\frac{\eta}{x(v_3) - x(v_1)} \\
&= \frac{x(z_{\rho(i)}) - x(b_1)}{x(b_2) - x(b_1)} \cdot \frac{\eta}{x(v_3) - x(v_1)} = \frac{x(v_3) - x(v_1)}{x(v_2) - x(v_1)} \cdot \frac{\eta}{x(v_3) - x(v_1)} \\
&= \frac{\eta}{x(v_2) - x(v_1)} = \varepsilon.
\end{aligned}$$

In the third line, we have used the fact that Δ' is a scaled, translated copy of Δ. An analogous calculation shows that $\overline{q_2} - q_2 = \varepsilon$.

The result is that if ε is the difference between a current boundary slope r of G_{i-1}' and the corresponding boundary slope $Z(r)$ of $Z_{h(i-1)}$, then this difference is propagated, but not

increased, by the addition of a_i'. That is, the slopes of the boundary edges of G_i' adjacent to a_i' will differ by ε from the corresponding boundary slopes of $Z_{\rho(i)}$.

We now investigate how the slopes $\overline{q_1}$ and $\overline{q_2}$ change when we move $\overline{v_3}$ to the integer point $v_3' = a_i'$. We may assume without loss of generality that $x(z_{\rho(i)}) < 0$, and hence that $a_i' = (\lfloor x(v_3) \rfloor, \lceil y(v_3) \rceil)$, as the other case is treated identically. We will let q_1' and q_2' denote the slopes that result from replacing $\overline{v_3}$ with $v_3' = a_i'$. That is, let q_1' be the slope of the line passing through a_i' and w_1', and let q_2' be the slope of the line passing through a_i' and w_2'.

Since $z_{\rho(i)}$ lies to the left of the apex of $Z_{\rho(i)}$, the vertex a_i' lies below and to the left of the apex of G_i'. Therefore we clearly have $q_1' - \overline{q_1} > 0$ and $\overline{q_2} - q_2' > 0$. By $\mathbf{C}(i-1, 2)$, we have $|\varepsilon| < i - 1$. Therefore we obtain

$$
\begin{aligned}
q_1' - \overline{q_1} &= \frac{\lceil y(\overline{v_3}) \rceil - y(w_1')}{\lfloor x(\overline{v_3}) \rfloor - x(w_1')} - \overline{q_1} < \frac{y(\overline{v_3}) - y(w_1') + 1}{x(\overline{v_3}) - x(w_1') - 1} - \overline{q_1} \\
&= \frac{y(\overline{v_3}) - y(w_1') + 1 - (x(\overline{v_3}) - x(w_1'))\overline{q_1} + \overline{q_1}}{x(\overline{v_3}) - x(w_1') - 1} = \frac{1 + \overline{q_1}}{x(\overline{v_3}) - x(w_1') - 1} = \frac{1 + \overline{q_1}}{x(v_3) - x(v_1) - 1} \\
&\leq \frac{1 + \overline{q_1}}{x(z_{\rho(i)}) - x(b_1) - 1} \leq \frac{1 + \overline{q_1}}{\alpha - 1} = \frac{1 + q_1 + \varepsilon}{\alpha - 1} \leq \frac{1 + M + \varepsilon}{\alpha - 1} \leq \frac{1 + M + (i-1)}{\alpha - 1} \\
&\leq \frac{2n^2 + n}{\alpha - 1} = 1.
\end{aligned}
$$

In the third line we have used the fact that $x(v_3) - x(v_1) \geq x(z_{\rho(i)}) - x(b_1)$, which we demonstrated above in order to establish $\mathbf{C}(i, 1)$. An analogous calculation shows that $\overline{q_2} - q_2' < 1$.

We may now compute

$$|q_1' - q_1| \leq |q_1' - \overline{q_1}| + |\overline{q_1} - q_1| = |q_1' - \overline{q_1}| + \varepsilon < 1 + \varepsilon \leq 1 + (i-1) = i.$$

An identical caluculation shows that $|q_2' - q_2| \leq i$. Note that q_1' is the slope of the boundary edge $w_1'a_i'$ of G_i' and q_1 is the slope of the edge $b_1 z_{\rho(i)} = Z(w_1'a_i')$, and similarly for q_2' and q_2. This establishes $\mathbf{C}(i, 2)$.

From the construction of a_i' it is clear that the line segments $w_1'a_i'$ and $a_i'w_2'$ intersect G_{i-1}' only in the vertices w_1' and w_2'. Thus G_i' is a plane triangulation, and $G_i' \sim G_i$. To show

that $\mathbf{F}(G'_i)$ is projectively convex, we carry out a calculation similar to that of the $d_i(a_i) > 2$ case.

Let $\hat{q_1}$ denote the slope of the boundary edge of G'_i adjacent and to the left of w'_1, if such an edge exists, and let $\hat{q_2}$ denote the slope of the boundary edge of G'_i adjacent and to the right of w'_2, if such an edge exists. Let $Z(\hat{q_1})$ and $Z(\hat{q_2})$ denote the boundary slopes of $Z_{h(i-1)}$ corresponding to $\hat{q_1}$ and $\hat{q_2}$, respectively. By $\mathbf{C}(i-1,2)$, we have $\overline{q_1} - q_1 = \epsilon \leq i-1$ and $Z(\hat{q_1}) - \hat{q_1} \leq i-1$. Thus

$$\begin{aligned}
&\hat{q_1} - \overline{q_1} = (\hat{q_1} - q_1) - (\overline{q_1} - q_1) \geq (\hat{q_1} - q_1) - (i-1) \\
&= (Z(\hat{q_1}) - q_1) - (Z(\hat{q_1}) - \hat{q_1}) - (i-1) \geq (Z(\hat{q_1}) - q_1) - 2(i-1) \\
&\geq 2n - 2i + 2 \geq 2.
\end{aligned}$$

We conclude that

$$\hat{q_1} - q'_1 = (\hat{q_1} - \overline{q_1}) - (q'_1 - \overline{q_1}) \geq (\hat{q_1} - \overline{q_1}) - 1 \geq 2 - 1 = 1 > 0.$$

An analogous calculation shows that $q'_2 - \hat{q_2} > 0$. Thus w'_1 and w'_2 are convex vertices of G'_i. Because the region $\mathbf{F}(G'_{i-1})$ is projectively convex by $\mathbf{C}(i-1,3)$, we conclude that $\mathbf{F}(G'_i)$ is projectively convex. By $\mathbf{C}(i-1,3)$, the sequence $(a'_1, \ldots, a'_{i-1})$ is a convex shedding sequence for G'_{i-1}, hence $(a'_1, \ldots, a'_i)$ is a convex shedding sequence for G'_i. Thus $\mathbf{C}(i,3)$ holds. We have now established $\mathbf{C}(i)$ in the case that $d_i(a_i) = 2$. This completes the induction, and we conclude that $\mathbf{C}(i)$ holds for all $3 \leq i \leq n$. Thus the triangulation $G' = G'_n$ is a sequentially convex embedding of G, with convex shedding sequence $\mathbf{a'} = (a'_1, \ldots, a'_n)$.

We have immediately that the x dimension of G' is

$$\alpha 2(n-2) = (2n^2 + n + 1)(2n - 4) = 4n^3 - 6n^2 - 2n - 4 \leq 4n^3.$$

Since $\mathbf{C}(n,2)$ holds, we conclude that the largest absolute value of a boundary slope of G' is at most $M + n \leq 2n^2 + n$. Thus the y dimension of G' is at most

$$\alpha 2(n-2)(2n^2+n) = (4n^3 - 6n^2 - 2n - 4)(2n^2 + n) = 8n^5 - 8n^4 - 10n^3 - 10n^2 - 4n \leq 8n^5.$$

Therefore G' is embedded in a $4n^3 \times 8n^5$ integer grid. □

CHAPTER 7

Lifting a Plane Triangulation to a Simplicial 3-Polytope

7.1 Lifting and the Shedding Diameter

Let $G = (V, E)$ be a plane triangulation. After obtaining a sequentially convex embedding of G using the algorithm of the previous chapter, it is a straightforward matter to construct a simplicial 3-polyhedron that projects vertically onto G. We introduce some definitions which will allow us to precisely describe how "tall" this polyhedron may need to be.

Let $\mathcal{A}_G$ denote the set of all shedding sequences for G. For $\mathbf{a} = (a_1, \ldots, a_n) \in \mathcal{A}_G$, we write $a_j \rightarrow_{\mathbf{a}} a_i$ if a_j is adjacent to a_i in $G_i(\mathbf{a})$. Then we define the *height* of each vertex a_i recursively, by

$$\tau(a_i, \mathbf{a}) = \begin{cases} i & i \leq 3 \\ 1 + \max\{\tau(a_j, \mathbf{a}) \mid a_j \rightarrow_{\mathbf{a}} a_i\} & i > 3 \end{cases}.$$

We define the *height* of the shedding sequence $\mathbf{a} \in \mathcal{A}_G$ by

$$\tau(\mathbf{a}) = \max_i \tau(a_i, \mathbf{a}),$$

and the *shedding diameter* of G by

$$\tau(G) = \min_{\mathbf{a} \in \mathcal{A}_G} \tau(\mathbf{a}).$$

Taking the transitive closure of the relation $\rightarrow_{\mathbf{a}}$, we obtain a partial order $\preceq_{\mathbf{a}}$ on the vertices of G. The height $\tau(\mathbf{a})$ of the sequence $\mathbf{a}$ is then precisely the height of $\preceq_{\mathbf{a}}$. That is, $\tau(\mathbf{a})$ is the maximal length of a chain in $\preceq_{\mathbf{a}}$.

Lemma 7.1.1 *Let G be a plane triangulation with n vertices and shedding sequence $\boldsymbol{a} \in \mathcal{A}_G$, embedded as in Theorem 6.3.1, so that $\boldsymbol{a}$ is a convex shedding sequence for G. Then there is*

a convex (unbounded) 3-polyhedron P_i that projects vertically onto G_i, for each $i = 3, \ldots, n$. Furthermore, if $h(a_i)$ denotes the height of the vertex of P_i projecting to a_i, then we may choose $h(a_i)$ to be an integer such that $h(a_i) \leq 499n^8 m_i + 1$, where

$$m_i \;=\; \max\{h(a_j) \mid a_j \rightarrow_{\boldsymbol{a}} a_i\}.$$

Proof. We proceed by induction on i. Let $h(v)$ denote the height assigned to the vertex v, and let $\varphi(v) = (v, h(v)) \in \mathbb{R}^3$ denote the point of $\mathbb{R}^3$ projecting vertically to v. We set $h(a_1) = h(a_2) = h(a_3) = 0$, and let $P_3 = \{(x, y, z) \in \mathbb{R}^3 \mid (x, y) \in \text{conv}(a_1, a_2, a_3), z \geq 0\}$. That is, P_3 is the unbounded prism with triangular face $a_1a_2a_3$ and lateral edges extending in the positive vertical direction, parallel to the z-axis. If $i > 3$, then by the induction hypothesis, G_{i-1} is the projection of a convex polyhedral surface P_{i-1}. To obtain a lifting of G_i, we must choose $h(a_i)$ properly.

Let $\mathcal{S}_i$ denote the set of faces of G_{i-1} having a vertex v such that $v \rightarrow_{\mathbf{a}} a_i$. Choose the height $h(a_i)$ large enough so that $\varphi(a_i)$ is not coplanar with $\varphi(v_1), \varphi(v_2), \varphi(v_3)$, for any face $F \in \mathcal{S}_i$ with vertices v_1, v_2, v_3. Let ℓ_i denote the ray with vertex $\varphi(a_i)$ and extending in the positive vertical direction, parallel to the z-axis. Since $\mathbf{F}(G_{i-1})$ is convex, clearly $P_i = \text{conv}(P_{i-1} \cup \ell_i)$ is a convex polyhedral surface P_i that projects vertically to G_i.

We now determine an upper bound on the height $h(a_i)$ for which $\varphi(a_i)$ is coplanar with $\varphi(v_1), \varphi(v_2), \varphi(v_3)$, for any $F \in \mathcal{S}_i$ with vertices v_1, v_2, v_3. Let $(x_0, y_0, z_0) \in \mathbb{R}^3$ denote the coordinates of $\varphi(a_i)$. Fix an $F \in \mathcal{S}_i$ and let v_1, v_2, v_3 denote the vertices of F. Let $(x_j, y_j, z_j) \in \mathbb{R}^3$ denote the coordinates of $\varphi(v_j)$. So in particular $(x_j, y_j) = v_j$. Then coplanarity of $\varphi(a_i)$ with $\varphi(v_1), \varphi(v_2), \varphi(v_3)$ means that

$$z_0 = a_1 x_0 + a_2 y_0 + a_3,$$

where

$$\begin{pmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} z_1 \\ z_2 \\ z_3 \end{pmatrix}.$$

Let A denote the matrix on the left side of this equation, and write $\mathbf{x} = (x_1, x_2, x_3), \mathbf{y} = (y_1, y_2, y_3), \mathbf{z} = (z_1, z_2, z_3)$. By Cramer's rule, $a_i = \frac{\det(A_i)}{\det(A)}$, where A_i is obtained by replacing the i^{th} column of A with $\mathbf{z}$.

Since G is embedded as in Theorem 6.3.1, the vertices of G lie in a $4n^3 \times 8n^5$ integer grid. Furthermore, from the construction of Theorem 6.3.1, the point $(0, 0)$ is contained in the edge a_1a_2 of G. This implies that $|x_j| \leq 4n^3$ and $|y_j| \leq 8n^5$ for each $j = 0, 1, 2, 3$. Therefore

$$\|\mathbf{x}\| \leq \sqrt{3} \max_{1 \leq j \leq 3} |x_j| \leq 4\sqrt{3}n^3 \qquad \text{and} \qquad \|\mathbf{y}\| \leq \sqrt{3} \max_{1 \leq j \leq 3} |y_j| \leq 8\sqrt{3}n^5.$$

Note also that

$$\|\mathbf{z}\| \leq \sqrt{3} \max_{1 \leq j \leq 3} |z_i| \leq \sqrt{3}m_i.$$

Since A is an invertible integer matrix, we have $|\det(A)| \geq 1$. Thus

$$\begin{aligned}
\frac{|\det(A_1)|}{|\det(A)|} &\leq |\det(A_1)| \\
&\leq \|(1, 1, 1)\|\|\mathbf{z}\|\|\mathbf{y}\| = \sqrt{3}\|\mathbf{z}\|\|\mathbf{y}\| \\
&\leq \sqrt{3}(\sqrt{3}m_i)(8\sqrt{3}n^5) = 24\sqrt{3}n^5m_i,
\end{aligned}$$

where we have invoked Hadamard's inequality in the second line.

By a similar argument,

$$\frac{|\det(A_2)|}{|\det(A)|} \leq 12\sqrt{3}n^3m_i \qquad \text{and} \qquad \frac{|\det(A_3)|}{|\det(A)|} \leq 96\sqrt{3}n^8m_i.$$

Thus when $\varphi(a_i)$ is coplanar with $\varphi(v_1), \varphi(v_2), \varphi(v_3)$, we have

$$\begin{aligned}
z_0 &= \frac{\det(A_1)}{\det(A)}x_0 + \frac{\det(A_2)}{\det(A)}y_0 + \frac{\det(A_3)}{\det(A)} \\
&\leq \frac{|\det(A_1)|}{|\det(A)|}|x_0| + \frac{|\det(A_2)|}{|\det(A)|}|y_0| + \frac{|\det(A_3)|}{|\det(A)|} \\
&\leq 24\sqrt{3}n^5m_i(4n^3) + 12\sqrt{3}n^3m_i(8n^5) + 96\sqrt{3}n^8m_i \\
&= 288\sqrt{3}n^8m_i \leq 499n^8m_i.
\end{aligned}$$

So letting z_0 be the smallest integer greater than $499n^8m_i$ will ensure that $\varphi(a_i)$ is not coplanar with $\varphi(v_1), \varphi(v_2), \varphi(v_3)$. Thus we may take $h(a_i) = z_0 \leq 499n^8m_i + 1$, as desired.

□

Theorem 7.1.2 *Let G be a plane triangulation with n vertices. Then G is the vertical projection of a convex 3-polyhedron with vertices lying in a $4n^3 \times 8n^5 \times (500n^8)^{\tau(G)}$ integer grid.*

Proof. Choose a shedding sequence $\mathbf{a} \in \mathcal{A}_G$ such that $\tau(G) = \tau(\mathbf{a})$. By Theorem 6.3.1, we may embed G in a $4n^3 \times 8n^5$ integer grid such that $\mathbf{a} = (a_1, \ldots, a_n)$ is a convex shedding sequence for G. For each vertex a_i we may assign a height $h(a_i)$ as follows. For $i = 1, 2, 3$ we may set $z_i = 0$. For $i > 3$, by Lemma 7.1.1 we may choose $h(a_i)$ such that G_i is the projection of a polyhedral surface, and

$$h(a_i) \leq \left(499n^8 + 1\right)^{\tau(a_i, \mathbf{a})} \leq (500n^8)^{\tau(a_i, \mathbf{a})} \leq (500n^8)^{\tau(\mathbf{a})} = (500n^8)^{\tau(G)}.$$

□

Note that if G has at least 4 vertices, and the boundary G is a triangle (that is, $\partial\mathbf{F}(G)$ contains exactly 3 vertices), then the 3-polyhedron of Theorem 7.1.2 may be replaced with a (bounded) 3-polytope. Indeed, simply truncate the polyhedron with the hyperplane that is defined by the lifts of the three boundary vertices of G. Then the three boundary vertices of G lift to the vertices of a triangular face of the resulting 3-polytope.

7.2 Triangulations of a Rectangular Grid

For $p, q \in \mathbb{Z}$, $p, q \geq 2$, let $[p \times q] = \{1, \ldots, p\} \times \{1, \ldots, q\}$. We may think of the integer lattice $[p \times q]$ as the vertices of $(p-1)(q-1)$ unit squares. A geometric plane triangulation G is a *triangulation of* $[p \times q]$ if the vertices of G are exactly the vertices of $[p \times q]$, and every boundary edge of $[p \times q]$ is an edge of G. We call G a *grid triangulation*. An $\ell \times \ell$ *subgrid* of $\mathbb{Z}^2$ is an integer translation of the lattice $[\ell \times \ell] = \{1, \ldots, \ell\} \times \{1, \ldots, \ell\}$. By an $\ell \times \ell$ subgrid of $[p \times q]$ we mean an $\ell \times \ell$ subgrid of $\mathbb{Z}^2$ that is a subset of $[p \times q]$. In this section we state and prove the following result concerning the shedding diameter of grid triangulations.

Theorem 7.2.1 *Let G be a triangulation of $[p \times q]$ such that for every edge e of G, there is*

a subgrid of size $\ell \times \ell$ that contains the endpoints of e. Then $\tau(G) \leq 6\ell(p+q)$.

This gives a class of triangulations with sublinear shedding diameter, if ℓ is held constant. According to Theorem 7.1.2, such a triangulation can be drawn in the plane so that it is the vertical projection of a simplicial 3-polyhedron embedded in a subexponential grid. That is, this class of triangulations corresponds to a class of simplicial polyhedra which may be embedded in an integer grid whose size is subexponential in the number of vertices.

Let $\leq_{\mathbb{Z}^2}$ denote the linear order on $\mathbb{Z}^2$ defined by

$$(x_1, y_1) \leq_{\mathbb{Z}^2} (x_2, y_2) \text{ if and only if } y_1 < y_2 \text{ or } y_1 = y_2 \text{ and } x_1 \leq x_2.$$

That is, $\leq_{\mathbb{Z}^2}$ is a lexicographic order in which y-coordinates take precedence in determining the order. We state without proof the following lemma, which summarizes some standard properties of shedding vertices of planar triangulations. See [BP] for a proof and references.

Lemma 7.2.2 ([BP], section 3) *Let G be a plane triangulation, and let v be a boundary vertex of G. If v is not a shedding vertex of G, then v is the endpoint of a diagonal e of G. Furthermore, each of the two connected components of $\mathbf{F}(G) \setminus e$ contains a shedding vertex of G.*

The rough idea of the proof of Theorem 7.2.1 is as follows (we provide the details below). We begin by constructing a particular shedding sequence $\mathbf{a}$ for G. To do this, we first subdivide $[p \times q]$ into a grid of $\lceil \frac{pq}{\ell^2} \rceil$ subgrids, (most of) which are squares of size $\ell \times \ell$. These squares form $\lceil \frac{p}{\ell} \rceil$ columns and $\lceil \frac{q}{\ell} \rceil$ rows.

We shed G in three stages. In Stage 1, we take every fourth column $U(1), U(5), U(9), \ldots$ and shed the vertices of each of these columns from top to bottom. When shedding $U(i)$, we may need to shed vertices in the column $U(i-1)$ or $U(i+1)$, for a total of at most $3q\ell$ vertices shed in the process of shedding the column $U(i)$. Because of their spacing, the shedding vertices for each column $U(1), U(5), U(9), \ldots$ do not interact. Specifically, at each step we have a collection of shedding vertices, one for each column, which we may think of as shedding "all at once". This collection of vertices is then an antichain with respect to $\preceq_{\mathbf{a}}$.

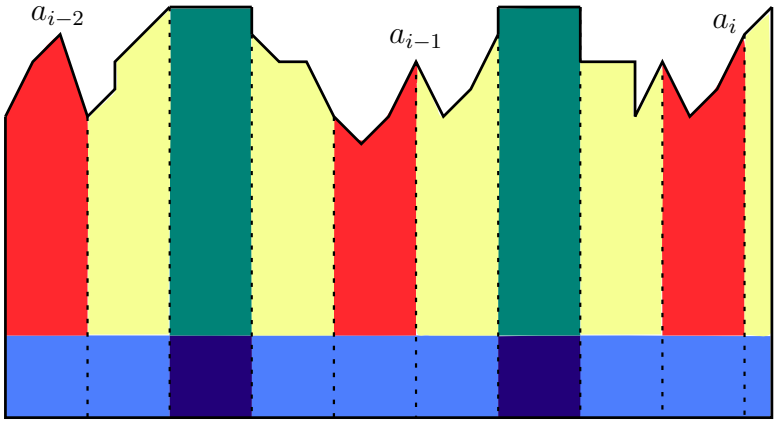

Figure 7.1: A graph G_i produced during Stage 1 of the construction of $\mathbf{a}$ in the proof of Theorem 7.2.1. Distinct columns $U(j)$ are separated by dashed lines. The columns of the form $U(1+4j)$ are shown in red, while the columns $U(3+4j)$ are shown in green. The bottom row $R(1)$ is shown in blue.

When shedding the vertices of each such column, for topological reasons we do not shed the vertices (x, y) with $y \leq \ell$. See Figure 7.1.

After Stage 1 is complete, what remains are a set of "jagged columns", each of which consists of the remaining vertices of three adjacent columns. Hence each jagged column contains at most $3q\ell$ vertices. In Stage 2, we shed these columns, but for topological reasons we do not shed vertices (x, y) with $y \leq 2\ell$. As before, these jagged columns do not interact, and at each step we have a set of shedding vertices, each of which belongs to a different jagged column. Hence this set forms an antichain. Finally, in Stage 3 we shed the remaining vertices, which are contained in the bottom two rows of G. There are at most $2p\ell$ such vertices, and we simply define a singleton set (which is trivially an antichain) for each of them. Therefore we see that G may be partitioned into at most $2p\ell + 3q\ell + 3q\ell = \ell(2p + 6q) \leq 6\ell(p + q)$ antichains of $\preceq_{\mathbf{a}}$. This implies that $\tau(\mathbf{a}) \leq 6\ell(p + q)$, since $\tau(\mathbf{a})$ is the length of some chain in $\preceq_{\mathbf{a}}$. The detailed proof follows.

Proof of Theorem 7.2.1. Let G be such a triangulation of $[p \times q]$. For each $i \in \mathbb{Z}$, let

$$U(i) = \{(x, y) \in [p \times q] \mid \ell(i-1) + 1 \leq x \leq \ell i\}, \quad \text{and}$$
$$R(i) = \{(x, y) \in [p \times q] \mid \ell(i-1) + 1 \leq y \leq \ell i\},$$

where some of these sets may be empty. We think of the sets $U(i)$ as *columns* of width ℓ and the $R(i)$ as *rows* of height ℓ. For each $i \in \mathbb{Z}$, we also define

$$T(i) = U(i-1) \cup U(i) \cup U(i+1),$$

which may be empty. We call $T(i)$ a *tricolumn* of $[p \times q]$.

We construct the shedding sequence $\mathbf{a}$ recursively. Suppose that we have a sequence of shedding vertices $a_{i+1}, a_{i+2}, \ldots, a_n$ (for the initial step of the recursion, $i = n$ and this sequence is empty), and therefore we also have graphs $G_i, G_{i+1}, \ldots, G_n = G$, where as usual $G_{j-1} = G_j - \{a_j\}$ for all $j = i+1, \ldots, n$. For each $i = 1, \ldots, n$, let $R_i(1)$ denote the subgraph of G_i induced by all vertices of G_i contained in $R(1)$. Similarly, for each $i = 1, \ldots, n$ and $j \in \mathbb{Z}$, let $U_i(j)$ denote the subgraph of G_i induced by all vertices of G_i contained in $U(j)$. We let $\mathbf{C}(i, 1)$ and $\mathbf{C}(i, 2)$ denote the following inductive claims:

$\mathbf{C}(i, 1)$. We have $U_i(3 + 4j) = U_n(3 + 4j)$ (for $j \in \mathbb{Z}$ such that $U(3 + 4j) \neq \emptyset$).

$\mathbf{C}(i, 2)$. The graph $R_i(1)$ is connected.

Note that $\mathbf{C}(n, 1)$ holds by definition. We have $G_n = G$, so the vertices of $U_n(1 + 3j)$ are exactly those of $U(1+3j)$, and similarly for $R_n(1)$ and $R(1)$. Since G is a grid triangulation, it follows that $U_n(1 + 3j)$ and $R_n(1)$ are connected. Thus in particular $\mathbf{C}(n, 2)$ holds.

To construct the next vertex a_i of the shedding sequence we break the construction into three stages, described below. As can readily be seen, Stage 1 occurs for a consecutive sequence of indices $i = k, k + 1, \ldots, n$, Stage 2 occurs for a consecutive sequence $i = j, j + 1, \ldots, k - 1$, and Stage 3 occurs for a consecutive sequence $i = 1, 2, \ldots, j - 1$. We will also see that $\mathbf{C}(i, 1)$ holds $i = k, k+1, \ldots, n$, and $\mathbf{C}(i, 2)$ holds for $i = j, j+1, \ldots, k, \ldots, n$.

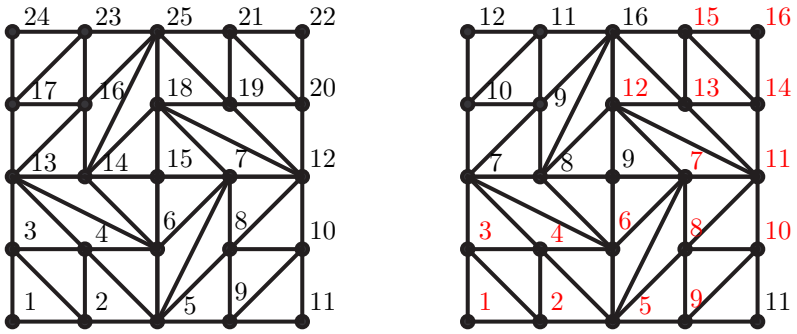

Figure 7.2: The grid triangulation of Figure 1.2, together with the indices i of the shedding sequence $\mathbf{a}$ defined in Theorem 7.2.1 (left) and the corresponding values of $\tau(a_i, \mathbf{a})$ (right). A chain of maximal length $\tau(\mathbf{a}) = 16$ is shown in red.

Stage 1. Some column of the form $U(1+4j)$ contains a vertex (x, y) of G_i with $y > \ell$. See Figure 7.1. Let $U(1+4j_1), \ldots, U(1+4j_r)$ denote all such columns, where $j_1 < \cdots < j_r$. Assume that $\mathbf{C}(i, 1)$ and $\mathbf{C}(i, 2)$ hold.

For each $k = 1, \ldots, r$, let v_k be the $\leq_{\mathbb{Z}^2}$-greatest vertex of $U(1+4j_k)$. If v_k is a shedding vertex of G_i, define $w_k = v_k$. Otherwise, by Lemma 7.2.2, the vertex v_k is the endpoint of a diagonal of G_i. Let u_k denote $\leq_{\mathbb{Z}^2}$-greatest vertex of G_i such that the edge $u_k v_k$ is a diagonal of G_i. Write $e_k = u_k v_k$.

By the Jordan curve theorem, $\mathbf{F}(G_i) \setminus e_k$ consists of two connected components, call them A_i and A_i'. Since the vertices u_k and v_k are adjacent, by assumption they are contained in an $\ell \times \ell$ subgrid of $[p \times q]$. It follows that $u_k \in T(1+4j_k)$. Thus $u_k, v_k \notin U(3+4j)$ for all j. Furthermore, since $y(v_k) > \ell$, we have $y(u_k) > 1$. By this and $\mathbf{C}(i, 1)$, one of the components of $\mathbf{F}(G_i) \setminus e_k$, say A_i, does not intersect any of the columns $U(3+4j)$. It follows that all vertices in A_i are contained in $T(1+4j_k)$. By Lemma 7.2.2, the region A_i contains a shedding vertex of G_i. We define w_k to be the $\leq_{\mathbb{Z}^2}$-greatest such shedding vertex.

We now have a collection of shedding vertices $w_1, \ldots, w_r$ of G_i. Since each of these vertices lies in a distinct column $U(1+4j_k)$, no two of these vertices are adjacent to a common vertex. Thus the vertex w_{r-1} is a shedding vertex of $G_i - \{w_r\}$, the vertex w_{r-2} is a shedding vertex

of $G_i - \{w_r, w_{r-1}\}$, etc. That is, these vertices remain shedding vertices after deleting any finite subset of them from G_i. So for each $k = 1, \ldots, r$, we may define $a_{i-r+1}, \ldots, a_i$ by $a_{i-r+k} = w_k$. Since no two of the vertices $a_{i-r+1}, \ldots a_i$, are adjacent, the set $\{a_{i-r+1}, \ldots, a_i\}$ is an antichain of $\preceq_{\mathbf{a}}$.

We now show that $\mathbf{C}(i-r+k, 1)$ and $\mathbf{C}(i-r+k, 2)$ hold for all $k = 1, \ldots, r$. By construction we have $a_{i-r+1}, \ldots, a_i \notin U(3+4j)$ for all j. Thus $U_{i-r+k}(3+4j) = U_i(3+4j)$ for all $k = 1, \ldots, r$. That is, none of the vertices of the columns $U(3+4j)$ are deleted in Stage 1. Thus $\mathbf{C}(i-r+k, 1)$ clearly holds.

On the other hand, some vertices of $R_i(1)$ may be deleted in Stage 1. But any vertex a_{i-r+k} which is a vertex of $R_i(1)$ is part of a sequence of vertices, whose deletion turns a diagonal of G_i into a boundary edge of G_{i-r+k}, for some k. It follows that $R_{i-r+k}(1)$ is path-connected for all $k = 1, \ldots, r$, so $\mathbf{C}(i-r+k, 2)$ holds.

Stage 2. No column of the form $U(1+4j)$ contains vertices (x, y) of G_i with $y > \ell$, but some tricolumn of the form $T(3+4j)$ contains vertices of (x, y) with $y > 2\ell$. Let $T(3+4j_1), \ldots, T(3+4j_r)$ denote all such tricolumns, where $j_1 < \cdots < j_r$. Assume that $\mathbf{C}(i, 2)$ holds.

For each $k = 1, \ldots, r$, let v_k be the $\leq_{\mathbb{Z}^2}$-greatest vertex of $T(3+4j_k)$. If v_k is a shedding vertex of G_i, define $w_k = v_k$. Otherwise, by Lemma 7.2.2, the vertex v_k is the endpoint of a diagonal of G_i. Let u_k denote $\leq_{\mathbb{Z}^2}$-greatest vertex of G_i such that the edge $u_k v_k$ is a diagonal of G_i. Write $e_k = u_k v_k$.

By the Jordan curve theorem, $\mathbf{F}(G_i) \setminus e_k$ consists of two connected components, call them A_i and A_i'. Since the vertices u_k and v_k are adjacent, by assumption they are contained in an $\ell \times \ell$ subgrid of $[p \times q]$. Since $y(v_k) > 2\ell$, it follows that $y(u_k) > \ell$. Thus $u_k, v_k \notin R(1)$. Then by $\mathbf{C}(i, 2)$, one of the components of $\mathbf{F}(G_i) \setminus e_k$, say A_i, does not intersect $R(1)$. By Lemma 7.2.2, the region A_i contains a shedding vertex of G_i. We define w_k to be the $\leq_{\mathbb{Z}^2}$-greatest such shedding vertex.

We now have a collection of shedding vertices $w_1, \ldots, w_r$ of G_i. Since each of these

vertices lies in a distinct tricolumn $T(3+4j_k)$, and the columns of the form $U(1+4j)$ are empty save for vertices in $R(1)$, no two of the vertices $w_1, \dots, w_r$ are adjacent to a common vertex. Thus these vertices remain shedding vertices after deleting any finite subset of them from G_i. So for each $k = 1, \dots, r$, we may define $a_{i-r+1}, \dots, a_i$ by $a_{i-r+k} = w_k$. Since no two of the vertices $a_{i-r+1}, \dots a_i$, are adjacent, the set $\{a_{i-r+1}, \dots, a_i\}$ is an antichain of $\preceq_{\mathbf{a}}$.

Finally, note that by construction we have $a_{i-r+1}, \dots, a_i \notin R(1)$ for each $k = 1, \dots, r$. That is, none of the vertices of the bottom row are deleted in Stage 2. Thus $R_{i-r+k}(1) = R_i(1)$ for all $i = 1, \dots, k$, so the claim $\mathbf{C}(i-r+k, 2)$ clearly holds.

Stage 3. All vertices (x, y) of G_i have $y \leq 2\ell$. If $i > 3$ we define a_i to be the $\leq_{\mathbb{Z}^2}$-greatest shedding vertex of G_i, which exists by Lemma 6.0.2. If $i \leq 3$ we define a_i to be the $\leq_{\mathbb{Z}^2}$-greatest vertex of the triangle G_3. The singleton set $\{a_i\}$ is trivially an antichain of $\preceq_{\mathbf{a}}$.

This completes the construction of the shedding sequence $\mathbf{a} = (a_1, \dots, a_n)$ (See Figure 7.2). It is straightforward to count the number of antichains of $\preceq_{\mathbf{a}}$ obtained from this construction. Stage 1 requires as many steps as it takes for the last column of the form $U(1+4j)$ to run out of vertices (x, y) with $y > \ell$. Since each vertex a_{i-r+k} of Stage 1 is contained in some tricolumn of the form $T(1+4j_k)$, this requires at most $|T(1+4j)| = 3q\ell$ steps, each of which produces an antichain. Similarly, Stage 2 requires as many steps as it takes for the last tricolumn of the form $T(3+4j)$ to run out of vertices (x, y) with $y > 2\ell$. This requires at most $|T(3+4j)| = 3q\ell$ steps, each of which produces an antichain. Finally, each set $\{a_i\}$ is trivially an antichain, so taking the singleton of each vertex a_i defined in Stage 3 yields at most $2p\ell$ antichains.

The set of antichains of $\preceq_{\mathbf{a}}$ produced by these three cases clearly forms a partition of $V(G) = [p \times q]$. There are at most $2p\ell + 3q\ell + 3q\ell = \ell(2p+6q)$ antichains in this partition. Thus, since $\tau(\mathbf{a})$ is the length of some chain in $\preceq_{\mathbf{a}}$, we have

$$\tau(G) \leq \tau(\mathbf{a}) \leq \ell(2p+6q) \leq 6\ell(p+q).$$

□

Theorem 1.0.5 is now an immediate corollary of Theorems 7.1.2 and 7.2.1.

CHAPTER 8

Final Remarks and Further Directions

It is perhaps not obvious why Theorem 1.0.1 does not follow from existence of irrational 4-polytopes. Indeed, one can take a Schlegel diagram Q of an irrational 4-polytope P and conjecture that this is the desired irrational polyhedral complex. The logical fallacy here is that the implications go the other way. If the Schlegel diagram of P is irrational, then indeed P must be an irrational polytope. However, the converse is not true. There is no reason why all realizations of the Schegel diagram Q must have irrational coordinates, even if P is irrational. In fact, after computing degrees of freedom one should *expect* additional realizations of Q. Similarly, it is only in $\mathbb{R}^3$ that one can have (and does have) the Maxwell-Cremona theorem [R]. In $\mathbb{R}^4$ and higher dimensions it easily fails.

There is a rather simple reason why Tutte's theorems are delicate and unlikely to allow a direct extension to higher dimensions, even ignoring the obstructions discussed in Chapter 1. Consider the two graphs as in Figure 8.1 below. The smaller of the two has a non-strict convex realization, while the bigger does not. Tutte's result produces a non-strict convex realization for the smaller one, and no realization for the second. His theorem is "if and only if", and he explains that the difference between the two graphs is combinatorial rather than geometric or topological. Of course, neither have a (strict) convex realization. This can be explained from the fact that this graph is not 3-connected and thus its *spring embedding* collapses.

Now in $\mathbb{R}^3$, replace the middle square with an octahedron, and the other two cells with non-convex polyhedra whose union with the octahedron forms a square prism (a natural 3-dimensional analogue of the middle graph in Figure 8.1). Then the connectivity obstacle

disappears. However, one can readily see that in this case, the intersection of all three polyhedra yields the four non-apex vertices of the octahedron, and the edges between them. Thus the intersection is topologically S^1, hence not contractible. From this and similar reasoning, we found that a more natural requirement than graph connectivity is that the topological polyhedra form a *complex*. That is, we require that the intersection of any two (or more) polyhedra is at most a single common face of both. In particular, this implies that the intersection must be contractible.

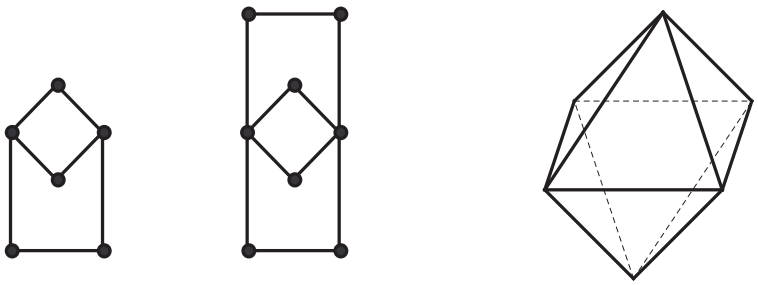

Figure 8.1: A graph with a non-strict convex realization (left), a graph with no convex realization (middle), and an octahedron (right).

Recently, two new explicit examples of simplicial balls with further properties were announced in [BL]. They have 12 and 15 vertices, respectively. These balls are not geometrically realizable, a consequence of the fact that their 1-skeletons contain knots consisting of few edges. They can be contrasted with the topological polyhedral ball X' we construct in the proof of Theorem 1.0.3, which has only 9 vertices.

It is obvious that the triangulation produced in Theorem 1.0.4 can be made rational—simply perturb all the vertices. In particular, this explains why we must use non-simplicial polytopes in the proof of Theorem 1.0.1. It is perhaps less obvious that all geometric realizations produced by Theorem 5.2.1 can be rational. Although the resulting polyhedral complex must have simplicial *interior faces*, the boundary faces can be arbitrary. Here rationality is a corollary resulting from the nature of the proof. All steps, in particular all projective transformations, can be done over $\mathbb{Q}$.

The *strong* vertex decomposability condition in Theorem 1.0.4 is somewhat restrictive already in $\mathbb{R}^3$. For example, it cannot apply to any triangulation of the icosahedron since its boundary has no vertices of degree three. Note the implications

$$\text{strongly vertex decomposable} \quad \Rightarrow \quad \text{vertex decomposable} \quad \Rightarrow \quad \text{shellable.}$$

Strong vertex decomposability is sufficient but not necessary for the existence of geometric realizations, and in fact neither vertex decomposability nor shellability are necessary. For all $d \geq 3$, there are examples of simplicial d-balls with $d+4$ vertices which are geometrically realizable, but not vertex decomposable (see [Lu]). In fact, there are simplicial 3-balls which are geometrically realizable (with boundary that of a tetrahedron), but not even shellable (see [Ru, Wz]).

Finding necessary and sufficient conditions for geometric realizability, hence improving the statements of Theorems 1.0.4 and 5.2.1 to *if and only if*, would be a remarkable accomplishment. We can attempt to place these types of results into a general framework. We have mentioned in Chapter 1 that the presence of knots in the 1-skeleton of the complex can be an obstruction to realizability. In the construction of Theorem 1.0.2, we have something slightly different. There are no knots present, and we could have used an arbitrary number of prisms in the belt. Indeed, the only reason we use two prisms is so that the construction is minimal.

In the case of non-trivial knots made of very few edges, the obstruction to realizability is a *topological one*. The edges cannot be knotted in a geometric (convex) realization, from which it follows that the complex must self-intersect in any such realization. In the case of the complex of Theorem 1.0.2, the obstruction to realizability is a *geometric one*. The issue is not self-intersection, but rather that in any geometric realization, the tetrahedron must be degenerate. That is, the belt imposes a linear constraint (namely, affine dependence) on the vertices of the tetrahedron.

If we limit ourselves to simplicial rather than general polyhedral complexes, then faces do not impose linear relations on their vertices, hence the only obstructions to realizability

are topological. The goal is then to classify the possible topological obstructions. Above we considered realizations of d-simplicial complexes in $\mathbb{R}^d$, but it may be enlightening to consider more generally realizations of d-simplicial complexes in $\mathbb{R}^k$, with $k \geq d$. Already by Fáry's theorem and Kuratowski's theorem, we have a precise description of when a graph (a 1-simplicial complex) has a geometric realization in $\mathbb{R}^2$. Namely, such a realization exists if and only if the graph G does not contain a subdivision of K_5 or $K_{3,3}$ (see also [Th]). On the other hand, a simple volume argument shows that every graph may be geometrically realized in $\mathbb{R}^3$.

The general idea is that either passing to a lower dimensional space, or requiring a geometric (convex) realization, may force certain non-trivial topological data to collapse, preventing such an embedding. There are many results (for example in [HZ]) concerning the implications of knots for properties of simplicial complexes such as constructibility, shellability, and vertex decomposability. However, the precise relationship between knots or other topological structures and geometric realizability remains unknown.

A natural question to ask is whether the geometric realizability of a simplicial complex is determined (at least in some reasonable cases) by the topological properties of a lower-dimensional subcomplex. For example, we conjecture the following.

Conjecture 8.0.3 *Let X be a topological 3-simplicial complex such that $|X| \sim B^3$. Suppose there is a topological 3-simplicial complex Z for which $X \simeq Z$, and such that the 1-skeleton of Z is ambient isotopic to a geometric graph G in $\mathbb{R}^3$. Then X has a geometric realization Y, such that G is the 1-skeleton of Y.*

We began our discussion with Steinitz's Theorem, and we conclude with it as well. An immediate consequence of Theorem 7.1.2 is that every plane triangulation may be lifted to a simplicial 3-polyhedron. In fact, if the boundary of the triangulation is a triangle, it may be lifted to a (bounded) 3-polytope. A natural question to ask is whether the proof of this theorem (in particular, the proofs of Theorems 6.1.1 and 7.1.1) can be readily extended to work for *all* 2-connected plane graphs (i.e. with general polygonal faces).

We have recently shown that it can, which gives an elementary proof of Tutte's theorem [T1] by induction on the number of vertices of the graph, and also yields a new proof of Steinitz's Theorem in the general case. In fact, since we consider general 2-connected plane graphs, we obtain something slightly stronger. Namely, our proof shows that every 2-connected planar graph is the graph of a 3-dimensional *Alexandrov cap*, which can be thought of as an unbounded polyhedron with one vertex at infinity. This brings us to the end of our discourse, but it is not the end of our story.

CHAPTER 9

Additional Figures and Notes

It follows from Theorem 1.0.1 that the vertices of the belts used in the proof of that theorem cannot be given by integer coordinates, so we cannot easily tabulate them. However, using the computer we are able to give an explicit description of the belts. This is done by specifying the arcs on which the triangular facets of the prisms lie, as well as a function describing their lateral lengths (see Figure 9.3). Note that the lateral lengths of the prisms are made small except at the boundary to ensure that the belts do not intersect. Furthermore, we must ensure that the arcs bend sufficiently to avoid each other at the top of the core.

To create the arcs we start with a family of circles, and then apply a parametrized family of rotations to stretch them. A MATHEMATICA notebook with code describing the explicit details of the construction in the proof of Theorem 1.0.1, and used to generate the irrational complex and images of it, is included with this document as a supplemental file.

In each belt, our construction uses 318 triangular prisms, exactly $2(80-1)+1 = 159$ prisms per half-belt. The core consists of 5 triangular prisms and 1 pentagonal pyramid. The complete irrational complex thus consists of a total of $4 \cdot 318 + 5 = 1277$ triangular prisms and 1 pentagonal pyramid, as in the theorem.

Since the belts come close to intersecting near the boundary of the core, some checking is necessary. In Figure 9.4 we show how the belts near-miss each other due to their shape. We conclude with additional images of the irrational 3-polytopal complex of Theorem 1.0.1.

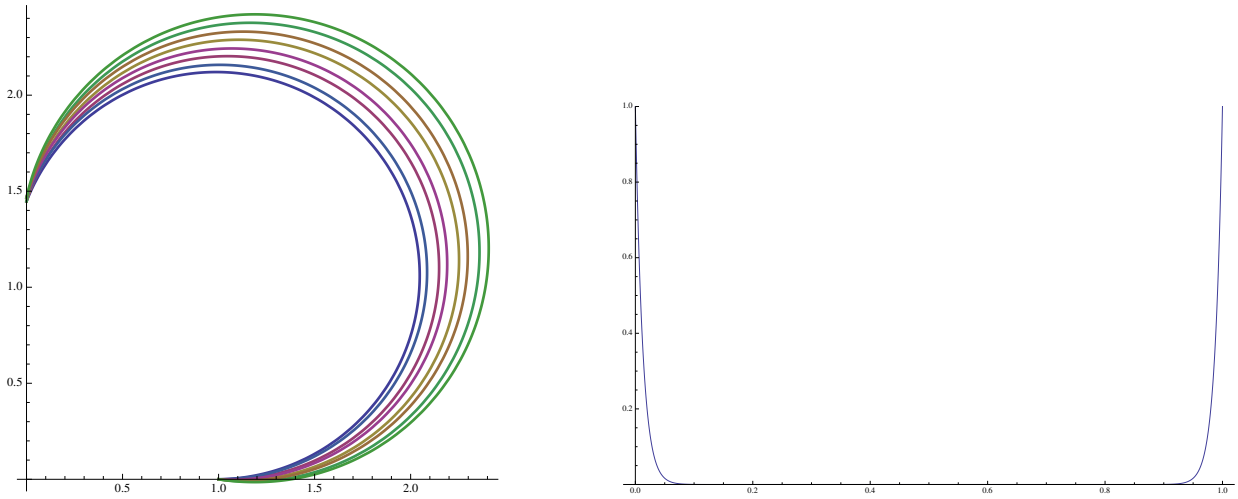

Figure 9.1: Left: Circles used in the first stage of the design of each belt. Right: The prism length function f_m for $m = 80$, described in the proof of Theorem 1.0.1.

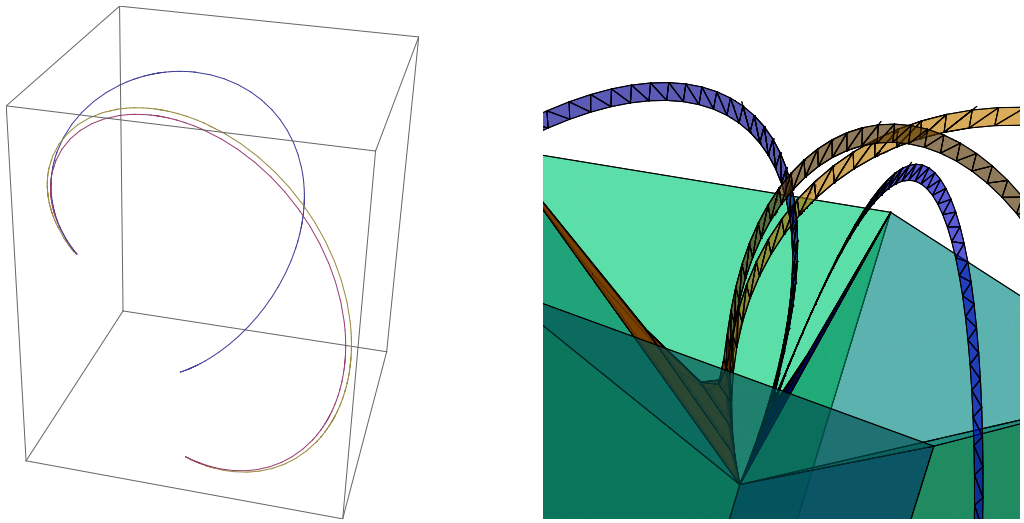

Figure 9.2: Left: The shape of one belt, obtained by stretching the circles of Fig. 9.3. Right: Belts do not intersect at the top of the core.

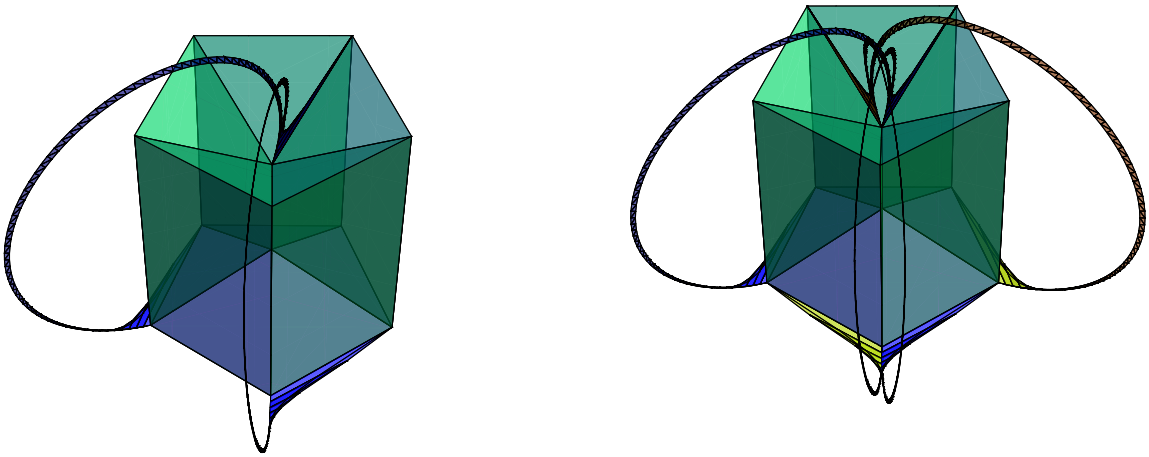

Figure 9.3: Left: One belt attached to the core. Right: Two belts attached.

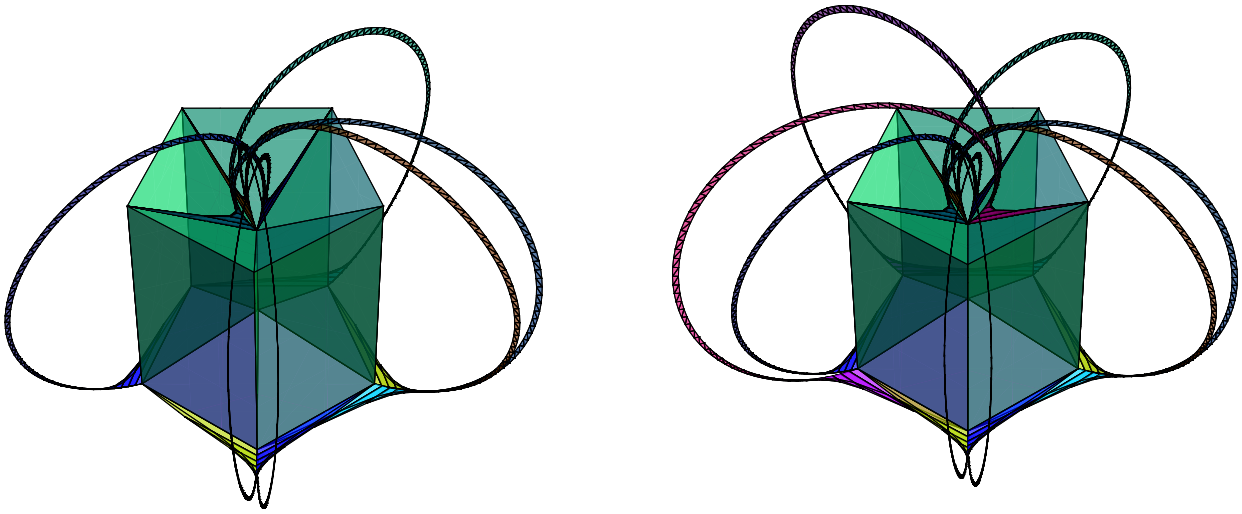

Figure 9.4: Left: Three belts attached. Right: All four belts attached.

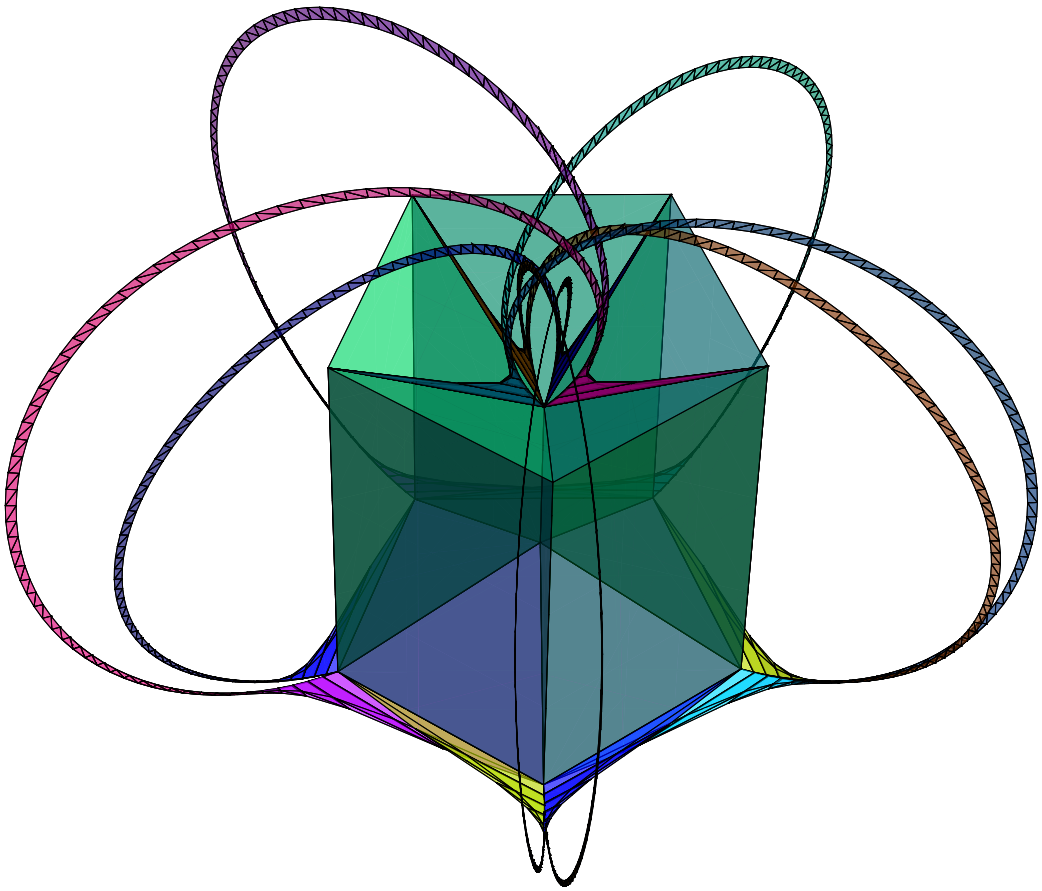

Figure 9.5: The complete irrational 3-polytopal complex of Theorem 1.0.1.

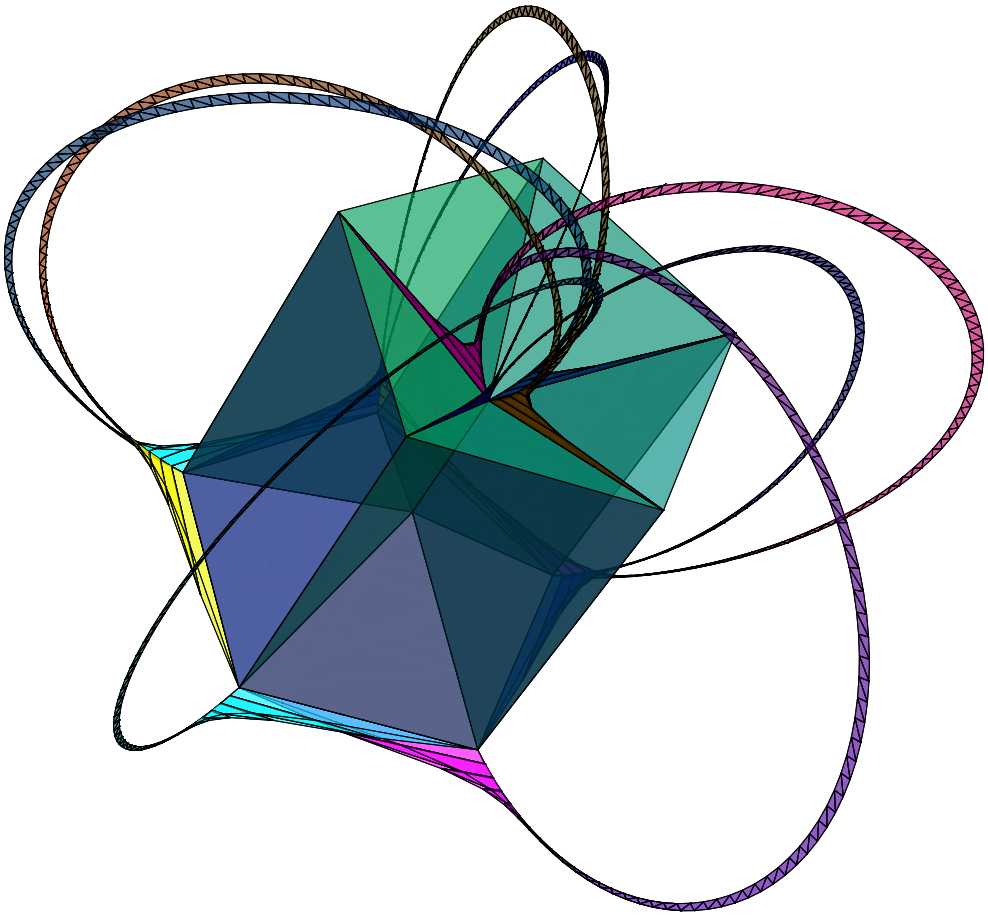

Figure 9.6: A rotated view of the irrational complex.

References

[A] E. E. Anclin, An upper bound for the number of planar lattice triangulations, *J. Combin. Theory Ser. A* **103** (2003), no. 2, 383–386.

[AB] K. Adiprasito and B. Benedetti, Metric geometry and collapsibility, available at `http://arxiv.org/abs/1107.5789`.

[BP] L. Billera and J. Provan, Decompositions of simplicial complexes related to diameters of convex polyhedra, *Math. Op. Res.* **5** (1980), 576–594.

[BG] D. W. Barnette and G. Grünbaum, Preassigning the shape of a face, *Pacific J. Math.* **32** (1970), 299–302.

[BL] B. Benedetti and F. Lutz, Non-evasiveness, collapsibility and explicit knotted triangulations, *Oberwolfach Report*, **8** (2011), 1–2.

[BR] I. Bárány and G. Rote, Strictly convex drawings of planar graphs, *Documenta Math.* **11** (2006), 369–391.

[BS] K. Buchin and A. Schulz, On the number of spanning trees a planar graph can have, in *Proc. ESA 2010*, 576–594.

[Ca] S. S. Cairns, Triangulated manifolds which are not Brouwer manifolds, *Ann. Math.* **41** (1940), 792–795.

[Ch] D. R. J. Chillingworth, Collapsing three-dimensional convex polyhedra, *Proc. Camb. Phil. Soc.* **63** (1967), 353–357; Erratum in **88** (1980), 307–310.

[DRS] J. A. De Loera, J. Rambau and F. Santos, *Triangulations: Structures and Algorithms*, Springer, 2008.

[DS] E. D. Demaine and A. Schulz, Embedding stacked polytopes on a polynomial-size grid, in *Proc. SODA 2011*, 1177–1187.

[Du] J. Dunkel, How many triangulations exist for a 6×6 grid?, preprint (2003).

[F] I. Fary, On straight line representations of planar graphs, *Acta Univ. Szeged. Sect. Sci. Math.* **11** (1948), 229–233.

[FPP] H. de Fraysseix, J. Pach and R. Pollack, Small sets supporting Fary Embeddings of Planar Graphs, in *Proc. STOC 1988*, ACM, 426–433.

[G] B. Grünbaum, *Convex polytopes*, Springer, Berlin, 1995.

[GS] B. Grünbaum and V.P. Sreedharan, An Enumeration of Simplicial 4-Polytopes with 8 vertices, *J. Comb. Theory* **2** (1967), 437–465.

[HZ] M. Hachimori and G. M. Ziegler, Decomposition of simplicial balls and spheres with knots consisting of few edges, *Math. Z.* **235** (2000), 159–171.

[K] E. R. van Kampen, Remark on the address of S. S. Cairns, in *Lectures in Topology*, 311–313, University of Michigan Press, Ann Arbor, MI, 1941.

[KZ] V. Kaibel, G. Ziegler, Counting lattice triangulations, in *Surveys in combinatorics*, 277–307, London Math. Soc. Lecture Note Ser., **307**, Cambridge Univ. Press, Cambridge, 2003.

[L] L. Lovász, *Geometric representations of graphs*, lecture notes; available electronically at `http://www.cs.elte.hu/~lovasz/geomrep.pdf`.

[Lu] F. Lutz, Vertex-Minimal Not Vertex-Decomposable Balls, *Electronic Geometry Model* No. 2003.06.001, available at `http://tinyurl.com/bmgsgyo`.

[Mn] N. Mnëv, The universality theorems..., in *Lecture Notes in Math.* **1346**, Springer, Berlin, 1988, 527–543.

[MT] B. Mojar and C. Thomassen, *Graphs on surfaces*, JHU Press, 2001.

[OS] S. Onn and B. Sturmfels, A quantitative Steinitz' theorem, *Beiträge zur Algebra und Geometrie* **35** (1994), 125–129.

[P] I. Pak, *Lectures on discrete and polyhedral geometry*, monograph draft; available electronically at `http://www.math.ucla.edu/~pak/book.htm`.

[R] J. Richter-Gebert, *Realization spaces of polytopes*, *Lecture Notes in Math.*, Springer, 1996.

[Ro] G. Rote, The number of spanning trees in a planar graph, in *Oberwolfach Reports* vol. 2, EMS, 2005, 969–973.

[Ru] M. Rudin, An unshellable triangulation of a tetrahedron, *Bull. Am. Math. Soc.* **64** (1958), 90–91.

[RRS] A. Ribó Mor, G. Rote and A. Schulz, Small grid embeddings of 3-polytopes, *Discrete Comput. Geom.* **45** (2011), 65–87.

[RZ] J. Richter-Gebert and G. M. Ziegler, Realization spaces of 4-polytopes are universal, *Bull. Amer. Math. Soc.* **32** (1995), 403–412.

[T1] W. T. Tutte, Convex representations of graphs, *Proc. Lond. Math. Soc.* **10** (1960), 304–320.

[T2] W. T. Tutte, How to draw a graph, *Proc. Lond. Math. Soc.* **13** (1963), 743–768.

[Th] C. Thomassen, Planarity and duality of finite and infinite graphs, *J. Combin. Theory Ser. B* **29** (1980), no. 2, 244–271.

[W] E. Welzl, The number of triangulations on planar point sets, in *Lecture Notes in Comput. Sci.* **4372**, Springer, Berlin, 2007, 1–4.

[Wi] N. Witte, A vertex decomposable 3-ball and 3-sphere with a knot consisting of 6 edges, *Electronic Geometry Model* No. 2001.05.003, available at `http://tinyurl.com/7bxdoxp`.

[Wo] R. Woodroofe, Vertex decomposable graphs and obstructions to shellability, *Proc. Amer. Math. Soc.* **10** (2009), 3235–3234.

[Wz] R. Wotzlaw, Rudin's non-shellable ball, *Electronic Geometry Model* No. 2004.08.001, available at `http://tinyurl.com/cumazdd`.

[Z1] G. M. Ziegler, *Lectures on polytopes*, Springer, New York, 1995.

[Z2] G. M. Ziegler, Convex polytopes: extremal constructions and f-vector shapes, in *Geometric Combinatorics*, AMS, Providence, RI, 2007, 617–691, available at `http://arxiv.org/abs/math/0411400`.

[Z3] G. M. Ziegler, Non-rational configurations, polytopes, and surfaces, *Math. Intell.* **30** (2008), 36–42.