

PLANAR GRAPHS AS MINIMAL RESOLUTIONS OF
TRIVARIATE MONOMIAL IDEALS

EZRA MILLER

Received: May 4, 2001

Revised: March 11, 2002

Communicated by Günter M. Ziegler

ABSTRACT. We introduce the notion of rigid embedding in a grid surface, a new kind of plane drawing for simple triconnected planar graphs. Rigid embeddings provide methods to (1) find well-structured (cellular, here) minimal free resolutions for arbitrary monomial ideals in three variables; (2) strengthen the Brightwell–Trotter bound on the order dimension of triconnected planar maps by giving a geometric reformulation; and (3) generalize Schnyder’s angle coloring of planar triangulations to arbitrary triconnected planar maps via geometry. The notion of rigid embedding is stable under duality for planar maps, and has certain uniqueness properties.

2000 Mathematics Subject Classification: 05C10, 13D02, 06A07, 13F55, 68R10, 52Cxx

Keywords and Phrases: planar graph, monomial ideal, free resolution, order dimension, (rigid) geodesic embedding

CONTENTS

INTRODUCTION AND SUMMARY	44
I GEODESIC EMBEDDING IN GRID SURFACES	49
1 PLANAR MAPS	49
2 GRID SURFACES	51
3 GLUING GEODESIC EMBEDDINGS	55

44	EZRA MILLER	
4	CONTRACTING RIGID GEODESICS	57
5	TRICONNECTIVITY AND RIGID EMBEDDING	66
II	MONOMIAL IDEALS	68
6	BETTI NUMBERS	68
7	CELLULAR RESOLUTIONS	69
8	GRAPHS TO MINIMAL RESOLUTIONS	70
9	UNIQUENESS VS. NONPLANARITY	72
10	DEFORMATION AND GENERICITY	73
11	IDEALS TO GRAPHS: ALGORITHM	75
12	IDEALS TO GRAPHS: PROOF	77
III	PLANAR MAPS REVISITED	80
13	ORTHOGONAL COLORING	80
14	ORTHOGONAL FLOWS	81
15	DUALITY FOR GEODESIC EMBEDDINGS	84
16	OPEN PROBLEMS	86

INTRODUCTION

Simple triconnected planar graphs admit numerous characterizations. Two famous examples include Steinitz' theorem on the edge graphs of 3-polytopes, and the Koebe–Andreev–Thurston circle packing theorem (see [Zie95] for both). These results produce “correct” planar (or spherical) drawings of the graphs in question, from which a great deal of geometric and combinatorial information flows readily.

This paper introduces a new kind of plane drawing for simple triconnected planar graphs, from which a great deal of *algebraic* and combinatorial information flows readily. These **geodesic embeddings** inside **grid surfaces** provide methods to

- solve the problem of finding well-structured (cellular, in this case) minimal free resolutions for arbitrary monomial ideals in three variables;

- strengthen the Brightwell–Trotter bound on the order dimension of tri-connected planar maps [BT93] by giving a geometric reformulation; and
- generalize Schnyder’s angle coloring for planar triangulations [Tro92, Chapter 6] to arbitrary triconnected planar maps via geometry.

We note that Felsner’s generalization of Schnyder’s angle coloring [Fel01] coincides with the **orthogonal colorings** independently discovered here as consequences of geometric considerations. In parallel with circle-packed and polyhedral graph drawings, additional evidence for the naturality of geodesic embeddings comes from their stability under duality, and the uniqueness properties enjoyed by “correct” geodesic embeddings—called **rigid embeddings** in what follows—for a given planar map.

The plan of the paper is as follows. Immediately following this Introduction is a section containing two theorems summarizing the equivalences and constructions forming main results of the paper. After that, the paper is divided into three Parts.

Part I lays the groundwork for geodesic and rigid embeddings in grid surfaces, and is geared almost entirely toward proving Theorem 5.1: the rigid embedding theorem. Terminology for the rest of the paper is set in Section 1, which also states a standard criterion for triconnectivity under edge contraction that serves as an inductive tool in the proof of Theorem 5.1. Then Section 2 presents the definition of grid surfaces, as well as the vertex and edge axioms for geodesic and rigid embeddings. Their consequences, the region and rigid region axioms, appear in Propositions 2.3 and 2.4. The first connection with order dimension comes in Corollary 2.5.

Sections 3 and 4 consist of stepping stones to the rigid embedding theorem. The basic inductive step for abstract planar maps is Lemma 3.1, which motivates the preliminary grid surface construction of Lemma 3.2. Induction for grid surfaces occupies the three Propositions in Section 4. They have been worded so that their rather technical proofs (particularly that of Proposition 4.2) may be skipped the first time through; instead, the Figures should provide ample intuition.

Section 5 completes the induction with a few more arguments about abstract planar maps. Corollary 5.2 recovers the Brightwell–Trotter bound on order dimension from rigid embedding.

The focus shifts in Part II to the algebra of monomial ideals in three variables, specifically their minimal free resolutions. A review of the standard tools occupies Section 6, while Section 7 recaps the more recent theory of cellular resolutions, along with a triconnectivity result (Proposition 7.2) suited to the applications here. Theorem 8.4 says how geodesic embeddings become minimal free resolutions. Corollary 8.5 then characterizes triconnectivity as the condition guaranteeing that a planar map supports a minimal free resolution of some artinian monomial ideal.

Section 9 displays even more reasons why rigid embeddings are better than arbitrary geodesic embeddings: they have a strong uniqueness property (Corol-

lary 9.1), which implies in particular that every minimal cellular resolution of the corresponding monomial ideal is planar. Surprisingly, there can exist *non-planar* cell complexes supporting minimal free resolutions of trivariate artinian monomial ideals that are sufficiently nonrigid; Example 9.2 illustrates one.

Sections 10–12 are devoted to producing minimal cellular free resolutions of arbitrary monomial ideals in three variables (Theorem 11.1). The deformations reviewed in Section 10 serve as part of the algorithmic solution pseudocoded in Algorithm 11.2. The proof of correctness for the algorithm and the theorem, which occupy Section 12, are rather technical and delicate. As with Section 4, the pictures may give a better feeling for the methods than the proofs themselves, at least upon first reading.

Part III continues where Part I left off, with more combinatorial theory for planar maps. Section 13 introduces orthogonal coloring, which generalizes Schnyder’s angle coloring and abstracts the notion of geodesic embedding (Proposition 13.1). Then, Section 14 shows how orthogonal coloring encodes the abstract versions of the orthogonal flows that played crucial roles in Section 2. As a consequence, Proposition 14.2 shows that orthogonal flows are examples of—but somewhat better than—normal families of paths, connecting once again with the work of Brightwell and Trotter on order dimension. Section 15 demonstrates how Alexander duality for grid surfaces (or monomial ideals) manifests itself as duality for planar maps geodesically embedded in grid surfaces.

Finally, Section 16 presents some open problems related to the notions developed in earlier sections, including a conjecture on orthogonal colorings and some problems on classifying cell complexes supporting minimal resolutions. Further questions concern applications of the present results to broader combinatorial algebraic problems, notably how to describe the “moduli space” of all minimal free (or injective) resolutions of ideals generated by a fixed number of monomials.

After completing an earlier version of this paper, the author was informed that Stefan Felsner had independently discovered the theory in Sections 13 and 14 [Fel01, Sections 1 and 2]. In addition, Felsner proved Conjecture 16.3 in [Fel02] after reading the preliminary version of this paper. See Section 16.3 for details and consequences.

Part III is almost logically independent of Part II, the only exceptions being Lemmas 8.2 and 8.3. Thus, the reader interested primarily in the combinatorics of planar graphs (as opposed to resolutions of monomial ideals) can read Parts I and III, safely skipping everything in Part II except for these two lemmas. The reader interested primarily in resolutions of monomial ideals should skip everything in Sections 3–5 except for the statement of Theorem 5.1.

ACKNOWLEDGEMENTS

This paper grew out of conjectures developed with Bernd Sturmfels during a memorable train ride through the Alps, and subsequent discussions resulting in the expository paper [MS99], where some of the results were announced without

proof. The Alfred P. Sloan Foundation and the National Science Foundation funded various stages of this project.

SUMMARY THEOREMS

For the sake of perspective and completeness, we collect the main ideas of the paper into a pair of precisely stated summary theorems. Their proofs are included, in the sense that the appropriate results from later on are cited. All of the notions appearing in Theorems A and B will be introduced formally in due time; until then, brief descriptions along with Figure 1 should suffice.

Let M be a connected simple planar map—that is, a graph embedded in a surface S homeomorphic to the plane \mathbb{R}^2 . All graphs in this paper have finitely many vertices and edges. Fix a point $\infty \in S$ far from M , and define the exterior region of M to be the connected component of $S \setminus M$ containing ∞ . Given three vertices $\dot{x}, \dot{y}, \dot{z} \in M$ bordering the exterior region, form the extended map $M_\infty(\dot{x}, \dot{y}, \dot{z})$ by connecting $\dot{x}, \dot{y}, \dot{z}$ to ∞ . Call a graph triconnected either if it is a triangle, or if it has at least four vertices of which deleting any pair along with their incident edges leaves a connected graph. A set of paths leaving a fixed vertex $\nu \in M$ is said to be independent if their pairwise intersection is $\{\nu\}$.

Let $k[x, y, z]$ be the polynomial ring in three variables over a field k , and let $I \subset k[x, y, z]$ be an ideal generated by monomials. The grid surface \mathcal{S} corresponding to I is the boundary of the staircase diagram of I , which is drawn (as usual) as the stack of cubes corresponding to monomials not in I . Rigid embedding of a planar map M in \mathcal{S} involves identifying the edges of M as certain piecewise linear geodesics in \mathcal{S} , and constitutes an inclusion of the vertex-edge-face poset of M into \mathbb{N}^3 . Orthogonal coloring M involves coloring the angles in M with three colors according to certain rules. Since it would take too long to do real justice to the definitions of ‘rigid embedding’ and ‘orthogonal coloring’ here, Figure 1 will have to do for now. The outer corners in the orthogonal coloring and the vectors on the axes in \mathcal{S} are called axial vertices. The grid surface \mathcal{S} is called axial when I is artinian.

Suppose M is a cell complex (finite CW complex) whose faces are labeled by vectors in \mathbb{N}^3 , in such a way that the union $M_{\leq \alpha}$ of faces whose labels precede $\alpha \in \mathbb{N}^3$ is a subcomplex of M for every α . Roughly speaking, M supports a cellular free resolution of I if the boundary complex of $M_{\leq \alpha}$ with coefficients in k is the \mathbb{N}^3 -degree α piece of a free resolution of I , for every $\alpha \in \mathbb{N}^3$.

THEOREM A *Let M be a planar map. The following are equivalent.*

1. M has three vertices $\dot{x}, \dot{y}, \dot{z}$ bordering its exterior region for which $M_\infty(\dot{x}, \dot{y}, \dot{z})$ is triconnected.
2. M has three vertices $\dot{x}, \dot{y}, \dot{z}$ bordering its exterior region to which every vertex of M has independent paths.
3. M has an orthogonal coloring with axial vertices $\dot{x}, \dot{y}, \dot{z}$.
4. M can be rigidly embedded in an axial grid surface.




Figure 1: Orthogonal coloring and rigid embedding of an extended map

5. M supports a cellular minimal free resolution of some artinian monomial ideal in $k[x, y, z]$.

Every artinian monomial ideal in $k[x, y, z]$ has a minimal cellular resolution supported on a cell complex M satisfying these conditions; in fact, Algorithm 11.2 produces such an M automatically.

Proof. $4 \Rightarrow 3$ follows from Proposition 13.1.

$3 \Rightarrow 2$ follows from Proposition 14.2.

$2 \Rightarrow 1$ follows easily from the definitions.

$1 \Rightarrow 4$ is Theorem 5.1.

$4 \Rightarrow 5$ follows from Theorem 8.4.

$5 \Rightarrow 1$ is Proposition 7.2.

The final statement comes from Theorem 11.1 and Proposition 12.4. \square

Similar—but weaker—statements apply to minimal cellular free resolutions of arbitrary (not necessarily artinian) monomial ideals in $k[x, y, z]$.

THEOREM B *Let N be a planar map. The following two conditions are equivalent.*

1. N can be rigidly embedded in some grid surface.
2. N can be obtained by deleting $\dot{x}, \dot{y}, \dot{z}$ and all edges incident to them from some planar map M satisfying the equivalent conditions in Theorem A.

These conditions imply that

3. N supports a minimal free resolution of some monomial ideal in $k[x, y, z]$.

Every monomial ideal in $k[x, y, z]$ has a minimal free resolution supported on a planar map N satisfying conditions 1 and 2; such an N can be produced algorithmically.

Proof. $1 \Rightarrow 2$ follows from Theorem 8.4 and Lemma 8.2.

$2 \Rightarrow 1$ follows from Theorem 5.1 and Lemma 8.2.

$1 \Rightarrow 3$ follows from Theorem 8.4.

The first half of the final statement is Theorem 11.1 along with the first paragraph of its proof on p. 80; add in Proposition 12.4 for the algorithmic part. \square

In reality, the more detailed versions later on are considerably more precise, demonstrating how some of the equivalent descriptions naturally give rise to others.

PART I

GEODESIC EMBEDDING IN GRID SURFACES

1 PLANAR MAPS

Let $\mathcal{V} = \{\nu_1, \dots, \nu_r\}$ be a finite set. A **graph** G with vertex set \mathcal{V} is uniquely determined by a collection $\mathcal{E} \subseteq \binom{\mathcal{V}}{2}$ of edges, each consisting of a pair of vertices. Except for one paragraph at the beginning of Section 15, we consider only **simple graphs**—that is, without loops or multiple edges—so G is an abstract simplicial complex of dimension 1 having vertex set \mathcal{V} . Thus G can be regarded as a topological space, via any geometric realization.

Let S be a surface homeomorphic to the Euclidean plane \mathbb{R}^2 . A **plane drawing** of G in S is a continuous morphism $G \hookrightarrow S$ of topological spaces that is a homeomorphism onto its image. If G is connected, the image M is called a **planar map**. Deleting the images of the vertices and edges of G from S leaves several connected components whose closures are the **regions** of M . The unique unbounded region is called the **exterior region** of M . Two planar maps

are **isomorphic** if they result from plane drawings of the same graph G , their regions have the same boundaries in G , and the boundaries of their exterior regions correspond. We often blur the distinction between a planar map and the underlying graph, by not distinguishing a vertex (resp. edge) of G from the corresponding point (resp. arc) of M in the surface S .

A graph G is **k -connected** either if G is the complete graph on k vertices, or if G has at least $k + 1$ vertices, and given any $k - 1$ vertices ν_1, \dots, ν_{k-1} of G , the **deletion** $\text{del}(G; \nu_1, \dots, \nu_{k-1})$ is connected. Here, the deletion is obtained by removing ν_1, \dots, ν_{k-1} as well as all edges containing them from G . In case $k = 2$ or 3 , the graph G is called **biconnected** or **triconnected**, respectively.

Suppose that e is an edge of a planar map M , and that none of the (one or two) regions containing e is a triangle. The **contraction** M/e of M along e is obtained by removing the edge e and identifying the two vertices of e . The underlying graph of M/e is the topological quotient G/e ; it is still simple because e is the only edge connecting its vertices in G (so G/e has no loops) and no triangles contain e in G (so G/e has no multiple edges). Some plane drawing of M/e is obtained by literally contracting the edge e in M (technically: there is a homotopy $G \times [0, 1] \rightarrow S$ such that $G \times t \rightarrow S$ is a plane drawing of G for $t < 1$, while $G \times 1 \rightarrow S$ is a composition $G \rightarrow G/e \rightarrow S$ with the second map being a plane drawing). Contraction will be a crucial inductive tool, via a well-known criterion for triconnectivity under contraction:

PROPOSITION 1.1 *Let M be a triconnected planar map with at least four vertices, and let e be an edge. If there exist two regions F, F' of M such that*

1. $e \cap F$ and $e \cap F'$ are the two vertices of e , and
2. $F \cap F'$ is nonempty,

then either e borders a triangle or the contraction M/e fails to be triconnected. Conversely, if e borders no triangles and M/e is triconnected, then no such F, F' exist.

Thinking of the surface $S \cong \mathbb{R}^2$ as the 2-sphere minus ∞ , many of the planar maps M in this paper result by embedding some graph G_∞ in the sphere with ∞ as a vertex, and then considering the induced plane drawing M of $\text{del}(G_\infty; \infty)$. When this is the case, we frequently need to consider the subset $M_\infty \subset S$ obtained by omitting the point ∞ from the plane drawing of G_∞ in the sphere; thus some of the vertices in M connect to the missing point ∞ by unbounded arcs in S . More generally, define an **extended map** $M_\infty \subset S$ to be the union of a planar map M and a set of infinite nonintersecting arcs connecting some of its vertices to ∞ . The closure \overline{M}_∞ of M_∞ in the sphere need not be a simple graph because it can have doubled edges: some vertex in M could have two or more unbounded arcs in M_∞ containing it.

Suppose the edges contained in the exterior region of M form a simple closed curve, called the **exterior cycle**. This occurs, for instance, when M is triconnected. Three vertices $\hat{x}, \hat{y}, \hat{z} \in M$ are called **axial** if they are encountered

(in order) proceeding counterclockwise around the exterior cycle. Having chosen axial vertices, define $M_\infty(\dot{x}, \dot{y}, \dot{z}) \subset S$ to be the union of M and three unbounded arcs, called the x , y , and z -**axes**, connecting $\dot{x}, \dot{y}, \dot{z}$ to ∞ . We sometimes blur the distinction between $M_\infty(\dot{x}, \dot{y}, \dot{z})$ and its closure $\overline{M}_\infty(\dot{x}, \dot{y}, \dot{z})$ in the sphere. For instance, we say that $M_\infty(\dot{x}, \dot{y}, \dot{z})$ is triconnected if the graph underlying $\overline{M}_\infty(\dot{x}, \dot{y}, \dot{z})$ is.

2 GRID SURFACES

Let \mathbb{R} denote the real numbers. Write vectors in \mathbb{R}^3 as $\alpha = (\alpha_x, \alpha_y, \alpha_z)$, and partially order \mathbb{R}^3 by setting $\alpha \preceq \beta$ (read ‘ α **precedes** β ’) whenever $\alpha_u \leq \beta_u$ for all $u \in \{x, y, z\}$. Say that $\alpha \in \mathbb{R}^3$ **strongly precedes** $\beta \in \mathbb{R}^3$ when $\alpha_u < \beta_u$ for all $u = x, y, z$; this is stronger than saying $\alpha \prec \beta$. (Throughout this paper, the letter u denotes any one of x, y, z , in the same way that x_i denotes one of x_1, \dots, x_n .) Use $\alpha \vee \beta$ and $\alpha \wedge \beta$ to denote the **join** (componentwise maximum) and **meet** (componentwise minimum) of $\alpha, \beta \in \mathbb{R}^3$.

Let $\mathcal{V} \subset \mathbb{N}^3 \subset \mathbb{R}^3$ be a set of pairwise incomparable elements, where \mathbb{N} denotes the set of nonnegative integers. The order filter

$$\langle \mathcal{V} \rangle = \{ \alpha \in \mathbb{R}^3 \mid \alpha \succeq \nu \text{ for some } \nu \in \mathcal{V} \}$$

generated by \mathcal{V} is a closed subset of the topological space \mathbb{R}^3 . Its boundary $\mathcal{S}_\mathcal{V}$ is called a **grid surface** or **staircase**. Orthogonal projection onto the plane $x + y + z = 0$ restricts to a homeomorphism $\mathcal{S}_\mathcal{V} \cong \mathbb{R}^2$. (This homeomorphism gives the correspondence between rhombic tilings of the orthogonal projection of the $|\dot{x}| \times |\dot{y}| \times |\dot{z}|$ parallelepiped and plane partitions of the $|\dot{x}| \times |\dot{y}|$ grid with parts at most $|\dot{z}|$. The grid surface in Figure 1 clearly demonstrates the homeomorphism: the diagram is, after all, drawn faithfully on the two-dimensional page.)

One of the basic properties of grid surfaces is that $\alpha \in \mathcal{S}_\mathcal{V}$ whenever $\rho, \sigma \in \mathcal{S}_\mathcal{V}$ and $\rho \preceq \alpha \preceq \sigma$. Therefore, if $\rho, \sigma \in \mathcal{S}_\mathcal{V}$ and $\rho \preceq \sigma$, then $\mathcal{S}_\mathcal{V}$ contains the line segment in \mathbb{R}^3 connecting ρ to σ . In particular, if $\nu, \omega \in \mathcal{V}$ satisfy $\nu \vee \omega \in \mathcal{S}_\mathcal{V}$, then $\mathcal{S}_\mathcal{V}$ contains the union $[\nu, \omega]$ of the two line segments joining ν and ω to $\nu \vee \omega$; we refer to such arcs as **elbow geodesics**¹ in $\mathcal{S}_\mathcal{V}$. When ν and ω are the *only* vectors in \mathcal{V} preceding $\nu \vee \omega$, the arc $[\nu, \omega]$ is called a **rigid geodesic**.

Denote the nonnegative rays of the coordinate axes in \mathbb{R}^3 by X, Y, Z , and use the letter U to refer to any of X, Y, Z . The ray $\nu + U$ intersects $\mathcal{S}_\mathcal{V}$ in an oriented line segment U_ν called the **orthogonal ray** leaving ν in the direction of U . Thus every point in \mathcal{V} has precisely three orthogonal rays, one parallel to each coordinate axis and all contained in $\mathcal{S}_\mathcal{V}$, although some orthogonal rays may be unbounded while others are bounded.

If U_ν is bounded, so it has an endpoint besides ν , then the other endpoint of U_ν can always be expressed as a join $\nu \vee \omega$ for some $\omega \in \mathcal{V}$. When there is

¹Elbow geodesics do minimize length for the metric on $\mathcal{S}_\mathcal{V}$ induced by the usual metric on \mathbb{R}^3 , but this fact has no practical application in this paper.

exactly one such point ω , so $[\nu, \omega]$ is a rigid geodesic, we say that ν or U_ν **points toward** ω .

Observe that $\nu \vee \omega$ must share two coordinates with at least one (and perhaps both) of ν and ω , so every elbow geodesic contains at least one orthogonal ray. Making compatible choices of elbow geodesics containing all orthogonal rays yields a planar map. To be precise, a plane drawing $M \hookrightarrow \mathcal{S}_\mathcal{V}$ is a **geodesic grid surface embedding**, or simply a **geodesic embedding** in $\mathcal{S}_\mathcal{V}$, if the following two axioms are satisfied:

(Vertex axiom) The vertices of M coincide with \mathcal{V} .

(Elbow geodesic axiom) Every edge of M is an elbow geodesic in $\mathcal{S}_\mathcal{V}$, and every bounded orthogonal ray in $\mathcal{S}_\mathcal{V}$ is part of an edge of M .

With the following stronger edge axiom instead, $M \hookrightarrow \mathcal{S}_\mathcal{V}$ is a **rigid embedding**, which we sometimes phrase by saying that M is **rigidly embedded** in $\mathcal{S}_\mathcal{V}$:

(Rigid geodesic axiom) The elbow geodesic axiom holds, and every edge of M is a rigid geodesic in $\mathcal{S}_\mathcal{V}$.

The rigid geodesic axiom really consists of three parts, each of which puts nontrivial restrictions on $\mathcal{S}_\mathcal{V}$ or M : every bounded orthogonal ray in $\mathcal{S}_\mathcal{V}$ is part of a rigid geodesic (a priori, this has nothing to do with M); every rigid geodesic in $\mathcal{S}_\mathcal{V}$ is an edge of M ; and every edge of M is a rigid geodesic in $\mathcal{S}_\mathcal{V}$.

LEMMA 2.1 *Let $M \hookrightarrow \mathcal{S}_\mathcal{V}$ be a geodesic or rigid embedding. Suppose \mathcal{V} is in order-preserving bijection with another set $\tilde{\mathcal{V}}$ of vertices via $\nu \leftrightarrow \tilde{\nu}$, so that $\nu_u \leq \omega_u \Leftrightarrow \tilde{\nu}_u \leq \tilde{\omega}_u$ for all $\nu, \omega \in \mathcal{V}$ and $u \in \{x, y, z\}$. Then the elbow or rigid geodesics in $\mathcal{S}_{\tilde{\mathcal{V}}}$ constitute another geodesic or rigid embedding of M . In particular, linearly scaling one or more coordinate axes by integer factors preserves geodesic or rigid embeddings.*

Proof. Purely order-theoretic properties of \mathcal{V} determine whether ν and ω are the endpoints of an elbow geodesic, or whether ν points toward ω . \square

Any geodesic embedding $M \hookrightarrow \mathcal{S}_\mathcal{V}$ determines an extended map

$$M_\infty = M \cup (\text{unbounded orthogonal rays}).$$

A special case occurs when $\mathcal{S}_\mathcal{V}$ is **axial**, having axial vectors

$$\hat{x} = (|\hat{x}|, 0, 0), \quad \hat{y} = (0, |\hat{y}|, 0), \quad \text{and} \quad \hat{z} = (0, 0, |\hat{z}|)$$

in \mathcal{V} for nonzero $|\hat{x}|, |\hat{y}|, |\hat{z}| \in \mathbb{N}$. Thus, if M is geodesically embedded in an axial grid surface $\mathcal{S}_\mathcal{V}$, we can define the axial vertices of M to be the axial vectors in \mathcal{V} , and set $M_\infty(\hat{x}, \hat{y}, \hat{z}) = M \cup X_{\hat{x}} \cup Y_{\hat{y}} \cup Z_{\hat{z}}$. (Precisely two bounded orthogonal rays leave each axial vertex, while all three orthogonal rays leaving any other vertex are bounded). Conversely, if M comes equipped with axial

vertices $\dot{x}, \dot{y}, \dot{z}$, then we require any geodesic embedding $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ to send these axial vertices to axial vectors in \mathcal{V} .

Suppose M is geodesically embedded in the axial grid surface $\mathcal{S}_{\mathcal{V}}$. The edge of M leaving any vertex $\nu \neq \dot{z}$ along the vertical orthogonal ray Z_{ν} connects ν to another vertex ω with strictly larger z -coordinate, but weakly smaller x and y -coordinates. Continuing in this manner constructs an **orthogonal flow** $[\nu, \dot{z}]$ from ν to \dot{z} that is increasing in z , but weakly decreasing in x and y . It follows that $[\nu, \dot{z}]$ and the similarly constructed paths $[\nu, \dot{x}]$ and $[\nu, \dot{y}]$ are **independent**, meaning that they intersect pairwise only at ν itself. Since $[\dot{x}, \dot{y}]$, $[\dot{y}, \dot{z}]$, and $[\dot{z}, \dot{x}]$ partition the exterior cycle of M into three arcs, the contractible sets bounded by

$$[\dot{x}, \nu, \dot{y}] := [\nu, \dot{x}] \cup [\nu, \dot{y}] \cup [\dot{x}, \dot{y}]$$

and its cyclically permuted analogues partition the regions of M .

LEMMA 2.2 *Suppose $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ is an axial geodesic embedding, and $\nu \in \mathcal{V}$ borders a region contained in $[\dot{x}, \omega, \dot{y}]$. Then $\nu_z \leq \omega_z$, with strict inequality if $\nu \notin [\omega, \dot{x}] \cup [\omega, \dot{y}]$. A similar statement holds for arbitrary permutations of x, y, z .*

Proof. The orthogonal flow $[\nu, \dot{z}]$ must cross $[\omega, \dot{x}]$ or $[\omega, \dot{y}]$, at $\nu' \in [\omega, \dot{x}]$, say. Concatenating the part of $[\omega, \dot{x}]$ from ω to ν' with the part of $[\nu, \dot{z}]$ from ν' to ν yields a path from ω to ν that is weakly decreasing in z . This path is strictly decreasing if $\nu \notin [\omega, \dot{x}] \cup [\omega, \dot{y}]$, for then it traverses (downwards) the vertical orthogonal ray Z_{ν} . \square

PROPOSITION 2.3 (REGION AXIOM) *Let $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ be an axial geodesic embedding, and F a bounded region of M . If α_F is the join of the vertices of F , then $\alpha_F \in \mathcal{S}_{\mathcal{V}}$, and every vertex $\nu \in F$ shares precisely one coordinate with α_F .*

Proof. If $\omega \in \mathcal{V}$, then Lemma 2.2 implies there is some $u \in \{x, y, z\}$ such that $\nu_u \leq \omega_u$ for all $\nu \in F$. This shows ω cannot strongly precede α_F , so $\alpha_F \in \mathcal{S}_{\mathcal{V}}$; every vertex $\nu \in F$ therefore shares at least one coordinate with α_F . Suppose by symmetry that $\nu_z = (\alpha_F)_z$. The two edges of F containing ν cannot increase in z , so they exit ν counterclockwise of X_{ν} and clockwise of Y_{ν} . At least one of these edges strictly increases in x , and another strictly increases in y , completing the proof. \square

The next proposition says that regions in rigid axial embeddings are 2-dimensional analogues of rigid geodesics. In addition to its applications throughout Sections 3 and 4, this fact will play a crucial role in Corollary 5.2, by way of Corollary 2.5, below. Refer to Figure 2 for an illustration of the rigid region axiom as well as its failure for nonrigid embeddings; ν and σ are vertices of an elbow geodesic, but $\rho \preceq \nu \vee \sigma$ in the latter case. The diagram is labeled as in the coming proof of Proposition 2.4.




Figure 2: The rigid region axiom and its failure for nonrigid embeddings

PROPOSITION 2.4 (RIGID REGION AXIOM) *Let $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ be a rigid axial embedding, and F a region of $M_{\infty}(\dot{x}, \dot{y}, \dot{z})$. If α_F is the join of the vertices of F and $\omega \in \mathcal{S}_{\mathcal{V}}$, then $\omega \in F \Leftrightarrow \omega \preceq \alpha_F$.*

Proof. The claim is obvious if F is one of the three unbounded regions, so assume F is bounded. Since $\omega \in F \Rightarrow \omega \preceq \alpha_F$ by definition, let $\omega \notin F$, and assume F is contained in $[\dot{x}, \omega, \dot{y}]$ by symmetry. Either $\omega_z > \nu_z$ for some vertex $\nu \in F$ with maximal z -coordinate $\nu_z = (\alpha_F)_z$, in which case the proof is trivial, or all vertices of F with z -coordinate $(\alpha_F)_z$ lie on $[\omega, \dot{x}] \cup [\omega, \dot{y}]$, by Lemma 2.2. Assume there is one on $[\omega, \dot{y}]$ by transposing x and y if necessary, and let $\nu \in F \cap [\omega, \dot{y}]$ be closest to ω .

The vertex $\rho \in [\omega, \dot{y}]$ pointing toward ν has the same z -coordinate as ν (because $\omega_z \geq \rho_z \geq \nu_z$ and $\omega_z = \nu_z$), so the rigid geodesic $[\rho, \nu]$ consists of the orthogonal rays Y_{ρ} and X_{ν} . Of ν 's two neighbors in F , let σ have smaller y -coordinate. Since $\rho \neq \sigma$ and $\sigma \notin [\omega, \dot{y}]$, the edge connecting ν to σ exits ν strictly counterclockwise of X_{ν} and strictly clockwise of Y_{ρ} . Thus $\sigma_z < \nu_z$, whence $\sigma_x = (\alpha_F)_x$ by the region axiom (Proposition 2.3). But $\rho \not\preceq \nu \vee \sigma$ by the rigid geodesic axiom, while $\rho_z = \nu_z = (\nu \vee \sigma)_z$ and $\rho_y < \sigma_y = (\nu \vee \sigma)_y$ by construction. Therefore $\rho_x > (\nu \vee \sigma)_x = \sigma_x = (\alpha_F)_x$. The proof is complete because ω decreases in x along $[\omega, \dot{y}]$ to ρ . \square

COROLLARY 2.5 *Let $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ be an axial rigid embedding. If P is a vertex, edge, or bounded region of M , let α_P denote join of the vertices in P . The map sending $P \mapsto \alpha_P$ constitutes an embedding in \mathbb{N}^3 of the vertex-edge-face poset of M .*

Proof. Immediate from the vertex, rigid geodesic, and rigid region axioms. \square

Recall that the **order dimension** of a partially ordered set (poset) \mathcal{P} is the smallest $d \in \mathbb{N}$ such that \mathcal{P} includes into the poset \mathbb{R}^d . The previous corollary says that the order dimension of an axial rigidly embedded planar map is no greater than 3. See [Tro92] for more on order dimension.

3 GLUING GEODESIC EMBEDDINGS

Let M be a planar map with axial vertices $\dot{x}, \dot{y}, \dot{z}$ and extended map $M_\infty = M_\infty(\dot{x}, \dot{y}, \dot{z})$. Suppose C is a simple cycle in M having three counterclockwise ordered vertices $\ddot{x}, \ddot{y}, \ddot{z}$, and furthermore that C bounds a closed disk $R \subset S$ that is a union of bounded regions in M . Following Brightwell and Trotter (cf. [Tro92, Chapter 6], although our definition differs slightly), we call C a **ring** if every edge of M not contained in R intersects R in a (possibly empty) subset of $\{\ddot{x}, \ddot{y}, \ddot{z}\}$. The double-dotted vertices play the roles of axial vertices for a smaller map $N = M \cap R$ “glued into” M by external edges emanating from N at $\ddot{x}, \ddot{y}, \ddot{z}$. Although we allow \ddot{u} for $u \in \{x, y, z\}$ to equal the original axial vertex $\dot{u} \in M$, we exclude the case where C is the exterior cycle of M by referring to a **proper ring**.

Assume for each $u = x, y, z$ that the vertex \ddot{u} meets at least one edge in M_∞ not contained in R (this occurs when M_∞ is triconnected). If there are at least two such edges then set $\underline{\ddot{u}} = \ddot{u}$. Otherwise, call the unique edge $e_{\ddot{u}}$ and name its other endpoint $\underline{\ddot{u}}$. Here, $\underline{\ddot{u}} = \infty$ is allowed because $\ddot{u} = \dot{u}$ is; but if C is a *proper* ring, then at most one of $\underline{\ddot{x}}, \underline{\ddot{y}}, \underline{\ddot{z}}$ can equal ∞ , because there are no proper rings containing two axial vertices $\dot{u} \in \{\dot{x}, \dot{y}, \dot{z}\}$ such that R contains all of their edges in M . Indeed, it would be impossible to choose the third vertex \ddot{u} from the pair of last points on the exterior cycle of M going from the two axial vertices on C toward the third.

The closure in the 2-sphere of the subset $\overline{M}_\infty \setminus R$ is a planar map whose intersection with R equals $\{\ddot{x}, \ddot{y}, \ddot{z}\}$. Construct the **contraction** \overline{M}_∞/R by leaving off the edges $\{e_{\ddot{u}} \mid \underline{\ddot{u}} \neq \ddot{u}\}$ as well as their endpoints \ddot{u} on C , and then connecting $\underline{\ddot{x}}, \underline{\ddot{y}}, \underline{\ddot{z}}$ to a new vertex τ inside R . View $M_\infty/R := (\overline{M}_\infty/R) \setminus \infty$ as being the extension of a map $M/R = \text{del}(\overline{M}_\infty/R; \infty)$. Thus $M_\infty/R = (M/R)_\infty$ has τ as a vertex, and still has axes drawn to ∞ , although τ might replace one of $\dot{x}, \dot{y}, \dot{z}$ as an axial vertex. When τ replaces \dot{u} , however, we are free to choose $\tau = \dot{u}$, so we still write $(M/R)_\infty(\dot{x}, \dot{y}, \dot{z})$.

LEMMA 3.1 *Let $M_\infty(\dot{x}, \dot{y}, \dot{z})$ be triconnected and M contain a proper ring C as above. Then both $M \cap R$ and M/R are planar maps, with axial vertices, whose extended maps $(M \cap R)_\infty(\ddot{x}, \ddot{y}, \ddot{z})$ and $(M/R)_\infty(\dot{x}, \dot{y}, \dot{z})$ are triconnected. Each of $M \cap R$ and M/R contains fewer edges and strictly fewer regions than M .*

Proof. Deleting from M_∞ any pair of vertices in M leaves every remaining vertex $\nu \in M \cap R$ connected to $\{\ddot{x}, \ddot{y}, \ddot{z}\}$, because every path connecting ν to ∞ in M_∞ passes through $\{\ddot{x}, \ddot{y}, \ddot{z}\}$. By the same argument, every vertex in M_∞ that remains after deleting any pair of vertices in M is connected to R —and hence to τ —in the deletion. (The removal of the edges $e_{\ddot{u}}$ ensures that M/R has no bivalent vertices on the way to τ .) The fact that $\dot{x}, \dot{y}, \dot{z}$ and $\ddot{x}, \ddot{y}, \ddot{z}$ can be chosen as axial vertices follows from the triconnectivity of the extended maps of M/R and $M \cap R$.

Now M has at least one region inside (resp. outside) R , because C is a simple cycle (resp. a proper ring). Thus M/R (resp. $M \cap R$) has strictly fewer regions

than M . The edge number inequality is obvious for $M \cap R$. For M/R , the number of edges is at most $E + 3$, counting the edges to τ , where E is the number of edges in $M \setminus R$. But the number of edges in M is at least $E + 3$, because R contains the cycle C . \square

Let M and N be planar maps with axial vertices $(\dot{x}, \dot{y}, \dot{z})$ and $(\ddot{x}, \ddot{y}, \ddot{z})$, respectively, such that $M_\infty(\dot{x}, \dot{y}, \dot{z})$ and $N_\infty(\ddot{x}, \ddot{y}, \ddot{z})$ are both triconnected. We now show how to **glue N into M** at a vertex $\tau \in M$ that is trivalent in M_∞ . Let the counterclockwise ordered neighbors of τ be α, β, γ (one of which might be ∞) in M_∞ . (Think of M and N as M/R and $M \cap R$ from Lemma 3.1, respectively.) Start by replacing τ with a small triangle in M_∞ (a ‘ $Y-\Delta$ ’ transformation), adding three new vertices in the process. This action requires working in M_∞ rather than M if $\tau \in \{\dot{x}, \dot{y}, \dot{z}\}$. Next, replace the new triangle and its interior with N , in such a way that $\alpha, \beta, \gamma \in M$ connect to the axial vertices $\ddot{x}, \ddot{y}, \ddot{z} \in N$ via edges $e_{\ddot{x}}, e_{\ddot{y}}, e_{\ddot{z}}$, called **tethers** in M_∞ . The result is an extended map for $M \cup_\tau N$, the **tethered gluing** of N into M at τ . Contracting some or all of the tethers yields a **gluing** of N into M , provided the resulting map is simple and triconnected.

The construction of the tethered gluing works at the level of grid surfaces. For instance, the hypotheses in the next lemma can easily be attained by scaling M . This is a key observation, making the induction in the proof of Theorem 5.1 possible. The left columns of Figures 5 and 4 illustrate examples of $\mathcal{S}_{\mathcal{V}_M}, \mathcal{S}_{\mathcal{V}_N}, \tau$, and $M \cup_\tau N \hookrightarrow \mathcal{S}_{\mathcal{V}}$.

LEMMA 3.2 *Let $M \hookrightarrow \mathcal{S}_{\mathcal{V}_M}$ and $N \hookrightarrow \mathcal{S}_{\mathcal{V}_N}$ be rigid embeddings with respective axial vertices $\dot{x}, \dot{y}, \dot{z}$ and $\ddot{x}, \ddot{y}, \ddot{z}$, and suppose $\tau \in M_\infty(\dot{x}, \dot{y}, \dot{z})$ is trivalent. If U_τ has length at least $m + 1$ for $U = X, Y, Z$, then τ is the unique vector in \mathcal{V}_M preceding $\tau + m\mathbf{1}$, where $\mathbf{1} = (1, 1, 1)$. If, in addition, $|\ddot{u}| \leq m$ for $u = x, y, z$ and*

$$\mathcal{V} = (\mathcal{V}_M \setminus \tau) \cup (\tau + \mathcal{V}_N)$$

then the rigid geodesics in $\mathcal{S}_{\mathcal{V}}$ provide a rigid embedding of the map $M \cup_\tau N$.

Proof. The orthogonal rays X_τ, Y_τ, Z_τ point toward α, β, γ (one of these may be ∞) in $\mathcal{S}_{\mathcal{V}_M}$ because τ is trivalent. Each vertex $\nu \in \mathcal{V}_M$ with $\nu \neq \tau$ has $\nu_x \geq \alpha_x$, $\nu_y \geq \beta_y$, or $\nu_z \geq \gamma_z$. Indeed, if ν lies in $[\dot{x}, \tau, \dot{y}]$ (say), then considering where the orthogonal flow $[\nu, \dot{z}]$ intersects $[\tau, \dot{x}] \cup [\tau, \dot{y}]$ shows that either $\nu_x \geq \alpha_x$ or $\nu_y \geq \beta_y$. Thus $\tau \preceq \tau + m\mathbf{1}$ is unique in $\mathcal{S}_{\mathcal{V}_M}$; the vertex axiom for \mathcal{V} is immediate.

The part of $\mathcal{S}_{\mathcal{V}}$ preceding $\tau + m\mathbf{1}$ equals $\tau +$ (the part of $\mathcal{S}_{\mathcal{V}_N}$ preceding $m\mathbf{1}$), by the uniqueness in M of $\tau \preceq \tau + m\mathbf{1}$. Thus every vertex, rigid geodesic, or bounded orthogonal ray in $N \hookrightarrow \mathcal{S}_{\mathcal{V}_N}$ gets translated by τ to the corresponding feature in $\mathcal{S}_{\mathcal{V}}$. Similarly, the parts of $\mathcal{S}_{\mathcal{V}}$ and $\mathcal{S}_{\mathcal{V}_M}$ not preceded by τ agree, so any vertex, rigid geodesic, or bounded orthogonal ray in $M \hookrightarrow \mathcal{S}_{\mathcal{V}_M}$ survives in $\mathcal{S}_{\mathcal{V}}$, as long as it is contained in a (perhaps unbounded) region of M_∞ not containing τ , by the rigid region axiom.

The only orthogonal rays unaccounted for as yet for the rigid geodesic axiom are those leaving $\tau + \ddot{u}$ and α, β, γ . Observe that $\tau + \ddot{u}$ is the unique element of \mathcal{V} on the orthogonal ray $U_\tau \subset \mathcal{S}_{\mathcal{V}_M}$. Thus an orthogonal ray leaving $\tau + \ddot{u}$ either points toward the corresponding one of α, β, γ whenever the latter is not ∞ , or it points away from \mathcal{V}_M . An orthogonal ray leaving α, β, γ either points back toward \ddot{u} along a rigid geodesic $e_{\ddot{u}}$, or it points away from $\tau + \mathcal{V}_N$. We conclude that the rigid geodesics in $\mathcal{S}_{\mathcal{V}}$ form a planar map isomorphic to $M \cup_\tau N$. \square

4 CONTRACTING RIGID GEODESICS

PROPOSITION 4.1 *Let M be a planar map with axial vertices $\dot{x}, \dot{y}, \dot{z}$. Suppose e is the edge in the exterior cycle of M leaving \dot{x} toward \dot{y} , and that e borders no triangles in $M_\infty(\dot{x}, \dot{y}, \dot{z})$. If M/e can be rigidly embedded in some grid surface, then so can M .*

Proof. Letting $\nu \in M$ be the other endpoint of e , we have $\nu \neq \dot{y}$ because the unbounded region of M_∞ containing e is not a triangle. The edge in M leaving ν clockwise from e determines an edge f in any rigid embedding $N \hookrightarrow \mathcal{S}_{\mathcal{V}}$ isomorphic to M/e . Note that $f \in N$ does not contain the orthogonal ray $Y_{\dot{x}}$ because $f \neq e$; and $f \not\supset Z_{\dot{x}}$ because ν sits between \dot{x} and \dot{y} . Therefore the orthogonal ray X_ω at the other endpoint ω of f in N points toward \dot{x} .

Assume all coordinates of vectors in \mathcal{V} are even, by scaling. The claim is that the rigid geodesics in $\mathcal{S}_{\mathcal{V} \cup \nu}$ constitute a rigid embedding isomorphic to M , where the coordinates of ν are defined by

$$\nu = (\nu_x, \nu_y, \nu_z) = (|\dot{x}| - 1, \omega_y, 0).$$

The addition of ν to \mathcal{V} affects at most the rigid geodesics in N containing one of the following: an orthogonal ray X_σ for some vertex $\sigma \in \mathcal{V}$ pointing toward \dot{x} ; an orthogonal ray at ν ; or $Y_{\dot{x}}$. All other rigid geodesics lie behind the plane $x = |\dot{x}| - 1$.

If $\sigma_y < \omega_y$, then X_σ is unaffected by ν , while if $\sigma_y \geq \omega_y$, then \dot{x} and ν are the only elements of \mathcal{V} preceding $\sigma \vee \dot{x} = \sigma \vee \nu + (1, 0, 0)$. Thus X_σ points toward ν if $\sigma_y \geq \omega_y$, because $\dot{x} \not\leq \sigma \vee \nu$. The three orthogonal rays X_ν, Z_ν , and Y_ν leaving ν point respectively toward \dot{x}, ω , and the vertex to which $Y_{\dot{x}} \subset \mathcal{S}_{\mathcal{V}}$ points in N . Finally, $Y_{\dot{x}} \subset \mathcal{S}_{\mathcal{V} \cup \nu}$ points toward ν . (Figure 3 illustrates the transition $M/e \rightsquigarrow M$.) \square

In the situation of Lemma 3.1, gluing $M \cap R$ into M/R may involve contracting some of the tethers in $(M/R) \cup_\tau (M \cap R)$. For the rest of this section, let $M \hookrightarrow \mathcal{S}_{\mathcal{V}_M}$ and $N \hookrightarrow \mathcal{S}_{\mathcal{V}_N}$ be rigid embeddings having respective axial vertices $\dot{x}, \dot{y}, \dot{z}$ and $\ddot{x}, \ddot{y}, \ddot{z}$, with $\tau \in M_\infty(\dot{x}, \dot{y}, \dot{z})$ a trivalent vertex having neighbors α, β, γ . Let B be the region of $M \cup_\tau N$ containing $e_{\dot{y}}$ and $e_{\ddot{z}}$.

PROPOSITION 4.2 *Assume that $\tau \neq \dot{y}$, that $Y_{\dot{x}}$ points toward τ , and that no edge in M has vertices $\{\dot{x}, \gamma\}$. Contracting $e_{\ddot{x}}$ along with neither, either, or*

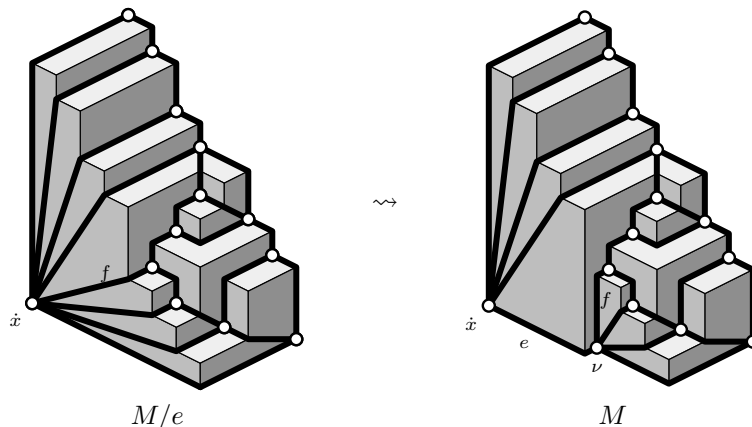


Figure 3: Uncontracting the lower-left edge

(if B has at least five vertices) both of e_z and e_{ij} in $M \cup_\tau N$ yields a planar map possessing a geodesic grid surface embedding.

Proof. Paragraph headings are included below to make parts of the proof easier to follow and cross-reference.

PLAN OF PROOF

Given the conditions of the present proposition, assume all hypotheses and notation of Lemma 3.2, as well; this is possible by Lemma 2.1. If $|\mathcal{V}_N| = n$, scale M so that all coordinates of vectors in \mathcal{V}_M are divisible by $n + 1$. Until further notice (see the special construction for (8), below), assume in addition that the orthogonal rays X_τ, Y_τ, Z_τ all have length exactly $m + 1$. Set

$$\nu' := \tau + \nu \quad \text{for } \nu \in \mathcal{V}_N.$$

The plan of the proof is to split into a number of cases, each of which demands slightly different treatment. In every case, Lemma 2.1 allows a judicious choice of coordinates for vectors in \mathcal{V}_N . Most often, omitting one or more of the vertices $\{\ddot{x}', \ddot{y}', \ddot{z}', \beta\}$ from \mathcal{V} leaves a set $\underline{\mathcal{V}}$ such that the desired contraction of $M \cup_\tau N$ geodesically embeds into $\mathcal{S}_{\underline{\mathcal{V}}}$; two of the cases require additional fiddling with the surviving vertices to get the desired grid surface.

EIGHT CASES

Let A and B be the bounded regions of $M \cup_\tau N$ containing $e_{\ddot{x}}$ and $e_{\ddot{y}}$, respectively. By the region axiom, we think of $A = (A_x, A_y, A_z)$ and $B = (B_x, B_y, B_z)$ as the vectors given by the joins of their vertices. Denote by $A(u)$ the set of

vertices of A having u -coordinate A_u . For instance, $A(x) = \{\dot{x}\} = \{\alpha\}$, and $\beta \in B(y)$.

Here is the list of constructions yielding $\mathcal{S}_{\underline{\mathcal{V}}}$. In each of (1)–(7), choose \mathcal{V}_N so that every coordinate of every vector in \mathcal{V}_N is at most n ; we treat the last case (8) separately later, since it involves somewhat different choices. Construct $\underline{\mathcal{V}}$ from $\mathcal{V} = (\mathcal{V}_M \setminus \tau) \cup (\tau + \mathcal{V}_N)$ by omitting the indicated vectors, and (in (5) and (8)) making the specified alterations.

- (1) To contract only $e_{\dot{x}}$: omit \dot{x}' .
- (2) To contract $e_{\dot{x}}$ and $e_{\dot{y}}$: omit \dot{x}', \dot{y}' .
To contract $e_{\dot{x}}$ and $e_{\dot{z}}$...
- (3) if no edge in N has endpoints $\{\dot{x}, \dot{z}\}$: omit \dot{x}', \dot{z}' .
if $\{\dot{x}, \dot{z}\}$ are the endpoints of an edge in N ...
- (4) and $A_z > \gamma_z$: omit \dot{x}', \dot{z}' .
- (5) and $A_z = \gamma_z$: omit \dot{x}', \dot{z}' ; then add 1 to ν_z for all $\nu \in A(z) \setminus \gamma$.
To contract $e_{\dot{x}}$, $e_{\dot{y}}$, and $e_{\dot{z}}$...
- (6) if no edge in N has endpoints $\{\dot{y}, \dot{z}\}$: omit \dot{y}' after (3)–(5).
if $\{\dot{y}, \dot{z}\}$ are the endpoints of an edge in N ...
- (7) and $\beta_x \geq \gamma_x$: omit β after (3)–(5).
- (8) and $\beta_x < \gamma_x$: make the special construction below.

In general, observe that the only vertices connected to \dot{x} , \dot{y} , or \dot{z} in $M \cup_{\tau} N$ are α, β, γ , and some vertices in N . Also, one of (6)–(8) must occur if B has at least 5 vertices. Representative instances of the cases (1)–(8) appear in Figures 4 and 5.

OMITTING VERTICES

In general, omitting one or more elements from \mathcal{V} always leaves a set of pairwise incomparable vectors. The vertex axiom will follow immediately in the applications below, in the sense that the surviving vertex vectors $\underline{\mathcal{V}}$ are in obvious bijection with the vertices of the desired map. To check the rigid geodesic axiom after omitting one vertex ν , we must verify that any orthogonal ray U pointing toward ν before the omission of ν points to some other uniquely determined surviving vertex afterwards. This will show that the surviving vertices $\underline{\mathcal{V}}$ define a rigid embedding $L \hookrightarrow \mathcal{S}_{\underline{\mathcal{V}}}$ for some map L .

Specifically, L is obtained from $M \cup_{\tau} N$ by first deleting ν along with all of its incident edges, and then reconnecting the neighbors of ν to other surviving vertices according to where their orthogonal rays point. It is important to remember that some edges incident to ν may fail to reappear upon reconnecting its neighbors: these non-reappearing edges are precisely those rigid geodesics that, before the omission of ν , do not contain any orthogonal ray pointing toward ν . (There are at most three such edges, because each one must contain an orthogonal ray leaving ν .)

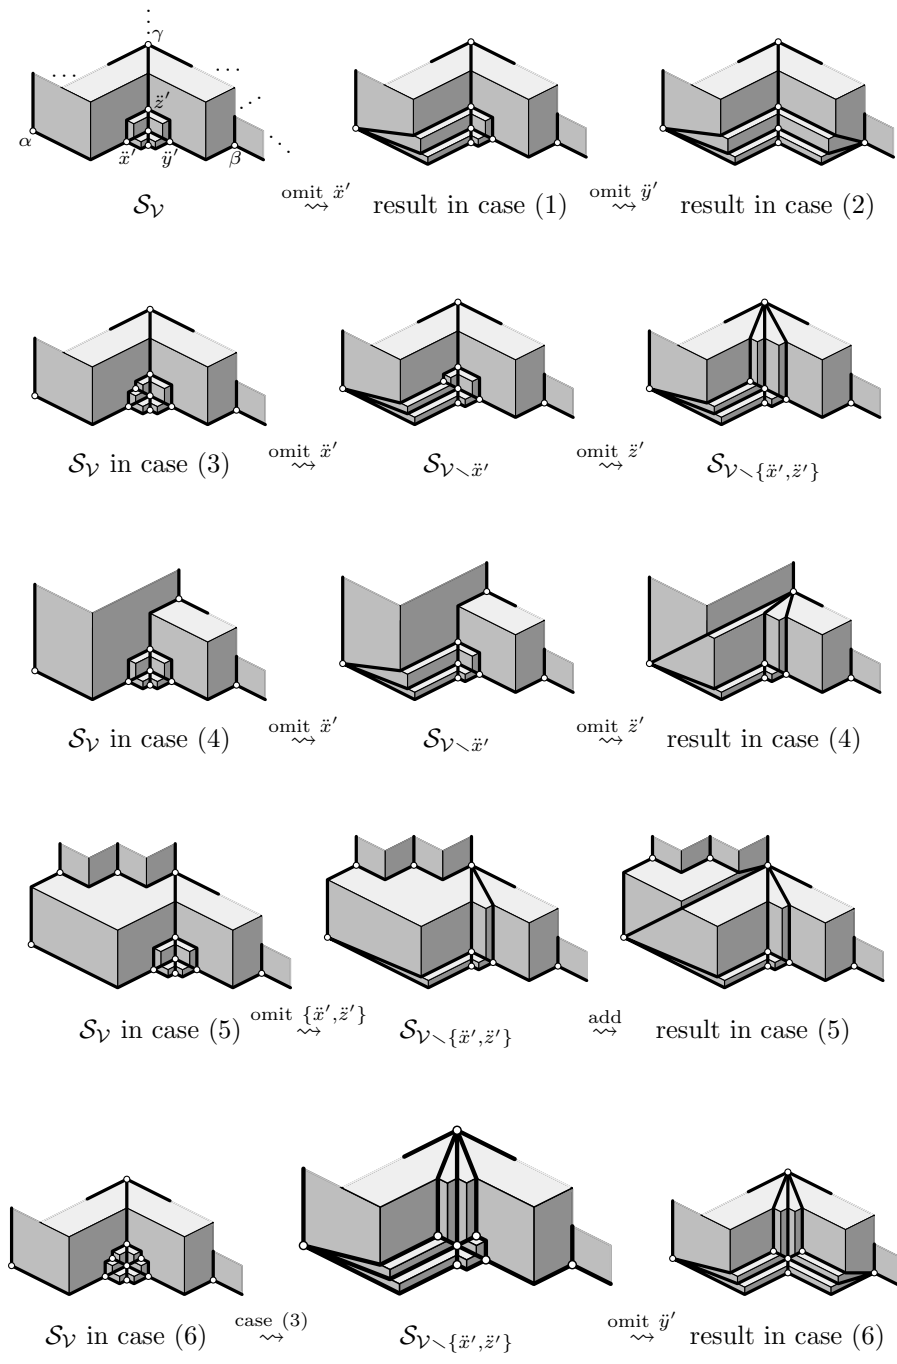


Figure 4: Gluing grid surfaces: cases (1)–(6)




Figure 5: Gluing grid surfaces: cases (7)–(8)

SCALE PRINCIPLE

We chose the relative sizes of \mathcal{V}_N and \mathcal{V}_M so that some arguments in what follows can rely on the following principle: If ω, σ are two vectors whose coordinates are all divisible by $n + 1$, then $\omega \preceq \sigma$ if and only if $\omega \preceq \sigma + n\mathbf{1}$.

RAYs LEAVING $\tau + \mathcal{V}_N$

An orthogonal ray leaving $\nu' \in \tau + \mathcal{V}_N$ and pointing toward z' in \mathcal{S}_ν must be $Z_{\nu'} \subset \mathcal{S}_\gamma$. We claim that after omitting z' , the ray $Z_{\nu'}$ points toward γ , regardless of whether or not one or both of x', y' has already been omitted. Moreover, this statement remains valid after permuting the roles of x, y, z . To see why, suppose that $\omega \preceq \nu' \vee \gamma$ for $\omega \in \mathcal{V}$. We must show that

$\omega \in \{\nu', \tilde{z}', \gamma\}$. If $\omega \in \tau + \mathcal{V}_N$ then $\omega_z \leq \tau_z + |\tilde{z}|$ and thus $\omega \preceq \nu' \vee \tilde{z}'$, so $\omega \in \{\nu', \tilde{z}'\}$ because $[\nu', \tilde{z}']$ is a rigid geodesic (Lemma 3.2). When $\omega \in \mathcal{V}_M \setminus \tau$, apply the scale principle with $\sigma = \gamma \vee \tilde{z}'$, using the fact that $\gamma \vee \nu' \preceq \gamma \vee \tilde{z}' + n\mathbf{1}$ because $\tau \preceq \tilde{z}', \nu' \preceq \tau + n\mathbf{1}$. The argument is invariant under permutation of x, y, z .

PROOF OF (1)

The rigid geodesics containing orthogonal rays $X_{\nu'}$ pointing toward \tilde{x}' before the omission and toward α afterwards account for all of the necessary edges in the contraction. It remains only to verify that $Y_\alpha \subset \mathcal{S}_{\mathcal{V} \setminus \tilde{x}'}$ points toward the next vertex after \tilde{x}' whose z -coordinate is zero—that is, the counterclockwise next vertex after \tilde{x}' on the exterior cycle of $M \cup_\tau N$. This easy argument is left to the reader, completing the proof of (1). Until the proof of (8), assume \tilde{x}' has been omitted.

PROOF OF (2)

The orthogonal rays pointing toward \tilde{y}' in $\mathcal{S}_{\mathcal{V} \setminus \tilde{x}'}$ are $Y_{\nu'}$ for some vertices $\nu \in \mathcal{V}_N$, the ray X_β , and possibly Y_α (if \tilde{x} and \tilde{y} are the vertices of an edge in N). The arguments in ‘Rays in $\tau + \mathcal{V}_N$ ’ and ‘Proof of (1)’ apply as well to the omission of \tilde{y}' , including the fact that the ray X_β points toward the next vertex after \tilde{y}' whose z -coordinate is zero (which may be α). None of the edges incident to \tilde{y}' vanish (see ‘omitting vertices’), although X_β and the Y -orthogonal ray at height $z = 0$ pointing toward \tilde{y}' point toward each other after omitting \tilde{y}' , effectively contracting $e_{\tilde{y}}$.

CASES (3)–(5)

Every rigid geodesic incident to \tilde{z}' in $\mathcal{S}_{\mathcal{V} \setminus \tilde{x}'}$ contains an orthogonal ray pointing toward \tilde{z}' , except the geodesic connecting α to \tilde{z}' , if there is one (this occurs in (4) and (5) only, where $X_{\tilde{z}'}$ points toward α in $\mathcal{S}_{\mathcal{V} \setminus \tilde{x}'}$), and sometimes the geodesic connecting γ to \tilde{z}' .

After omitting \tilde{z}' , we will verify the rigid geodesic axiom at γ separately for each of (3)–(5), by checking *only* that an orthogonal ray leaving γ and pointing toward \tilde{z}' (if there is one) points instead to another surviving vertex after omitting \tilde{z}' . Any other orthogonal ray in $\mathcal{S}_{\mathcal{V} \setminus \tilde{x}'}$ either leaves some vector in $\tau + \mathcal{V}_N$ before omitting \tilde{z}' (these have been dealt with in a paragraph above), or points toward another vertex in $\mathcal{V}_M \setminus \tau$ both before and after omitting \tilde{z}' . Verifying the rigid geodesic axiom will complete the proof unless the edges connecting α and γ to \tilde{z}' (if these exist) *both* vanish upon omitting \tilde{z}' , for otherwise all of the edges in the contraction of $e_{\tilde{z}}$ are accounted for as in the proofs of (1) and (2). Both geodesics vanish in (5) only: in (3), there is no edge connecting α to \tilde{z}' ; while in (4), the region axiom implies $\gamma_y = A_y$, whence X_γ points toward \tilde{z}' . In (5), we show that the addition procedure reconstructs a rigid geodesic connecting α to γ .

PROOF OF (3)

If X_γ points toward \ddot{z}' in $\mathcal{S}_{\mathcal{V} \setminus \ddot{x}'}$, and ν' is the vertex to which $X_{\ddot{z}'}$ points in $\mathcal{S}_{\mathcal{V} \setminus \ddot{x}'}$, then X_γ points toward $\nu' \in \tau + \mathcal{V}_N$ after omitting \ddot{z}' by the scale principle (the hypothesis for (3) guarantees that $\nu' \neq \ddot{x}'$). A similar statement holds by switching the roles of x and y (but $\nu' \in \tau + \mathcal{V}_N$ is always guaranteed to exist, since \dot{y}' has not been omitted). This verifies the rigid geodesic axiom at γ .

PROOF OF (4)

The assumption $A_z > \gamma_z$ implies $A_y = \gamma_y$, so that X_γ points toward \ddot{z}' in $\mathcal{S}_{\mathcal{V} \setminus \ddot{x}'}$. If $\nu \in \mathcal{V} \setminus \dot{x}$ and $\nu_y < \gamma_y$, then $\nu_z > \gamma_z$, by the region axiom. The omission of \ddot{z}' therefore causes the ray $X_\gamma \subset \mathcal{S}_{\mathcal{V} \setminus \{\ddot{x}', \ddot{z}'\}}$ to point toward α .

PROOF OF (5)

The region axiom implies $\gamma_y < \tau_y$, so X_γ does not point toward \ddot{z}' . When Y_γ points toward \ddot{z}' , the rigid geodesic axiom holds for $\mathcal{S}_{\mathcal{V} \setminus \{\ddot{x}', \ddot{z}'\}}$ by the argument in the proof of (3), although no elbow geodesic connects α to γ in $\mathcal{S}_{\mathcal{V} \setminus \{\ddot{x}', \ddot{z}'\}}$. Now we verify that the addition procedure outputs a rigid geodesic embedding, and that an edge connecting α to γ is the only new rigid geodesic. We can safely ignore all orthogonal rays contained in rigid geodesics on the positive side of the plane $y = \gamma_y$. All of the vertices $\omega \notin A$ satisfying $\omega_y \leq \gamma_y$ must also satisfy $\omega_z > \gamma_z$ by the rigid region axiom. The adding rule therefore causes X_γ to point toward α after omitting \ddot{z}' . The orthogonal rays leaving vertices originally in $A(z) \setminus \gamma$ still point to the same vertices, by the scale principle. If $\omega \notin A(z) \cup \{\alpha\}$, then any orthogonal ray U_ω pointing toward $\nu \in A(z)$ before the addition still points toward the same vertex ν afterwards, because $\omega_z > \nu_z$, whence the join $\nu \vee \omega$ remains unaffected. Finally, if Z_α points toward a vertex in $A(z)$ before the addition procedure, then Z_α still points to the same vertex afterwards, by the scale principle, while Y_α remains unaffected.

PROOF OF (6)

No new phenomena occur here; see Proof of (3).

TECHNICAL LEMMA

The following result will be applied in the proofs of (7) and (8). For the proof of (8), note that it holds after any rescaling of \mathcal{V}_M as in Lemma 2.1.

Let $\omega \in \mathcal{V}_M$. If $\tau_x \geq \omega_x \geq \beta_x$, then $\omega \in \{\tau, \beta\}$ or $\omega_z \geq \gamma_z$.
 Similarly, if $\tau_y \geq \omega_y \geq \alpha_y = 0$, then $\omega \in \{\tau, \alpha\}$ or $\omega_z \geq \gamma_z$.

Proof of technical lemma. Suppose $\tau_x \geq \omega_x \geq \beta_x$, but $\omega_z < \gamma_z$. If $\omega \neq \tau$, then $\omega_y \geq \beta_y$ by the uniqueness in Lemma 3.2 of τ among vectors preceding $\tau + m\mathbf{1}$

in \mathcal{V}_M . Thus $\beta \preceq \omega$, so $\beta = \omega$. Swap the roles of x and y to prove the other statement.

PROOF OF (7)

First claim: The only orthogonal rays pointing toward β are $Y_{\dot{y}'}$, and X_ν for some vertices $\nu \in M$. These vertices ν all have $\nu_z < m + 1$.

The final sentence of the first claim is easy, because otherwise $\gamma \preceq \beta \vee \nu$. For the rest of the claim, use the inequalities $B_x = \tau_x > \beta_x \geq \gamma_x$ and $B_y = \beta_y > \tau_y \geq \gamma_y$, which follow from the hypotheses of (7). These imply $\gamma_z = m + 1 = B_z$, thanks to the region axiom. By the technical lemma, any orthogonal ray Y_ν pointing toward β in $\mathcal{S}_\mathcal{V}$ (and therefore in $\mathcal{S}_{\mathcal{V} \setminus \{\dot{x}', \dot{z}'\}}$) must have either $\nu_x \geq \tau_x$ or $\nu_z \geq m + 1$. When $\nu_x \geq \tau_x$, we get $\nu_y < \beta_y$ and hence $\dot{y}' \preceq \nu \vee \beta$, so $\nu = \dot{y}'$. The case $\nu_z \geq m + 1$ is actually impossible, for it implies $\gamma \preceq \nu \vee \beta$, so $\nu = \gamma$ is connected to β by an edge in $M \cup_\tau N$; this cannot happen in (7) if B has at least five vertices. The fact that $\beta_z = 0$ rules out Z_ν pointing toward β , completing the proof of the first claim.

Now we verify that $X_\nu \subset \mathcal{S}_{\mathcal{V} \setminus \{\dot{x}', \beta, \dot{z}'\}}$ points toward \dot{y}' whenever $X_\nu \subset \mathcal{S}_{\mathcal{V} \setminus \{\dot{x}', \dot{z}'\}}$ points toward β . In other words, we need $\omega \preceq \nu \vee \dot{y}'$ for $\omega \in \mathcal{V}$ to imply $\omega \in \{\dot{x}', \dot{z}', \dot{y}', \beta, \nu\}$. If $\omega \in \tau + \mathcal{V}_N$, then $\omega_x = \dot{y}'_x = \tau_x$ implies $\omega \in \{\dot{y}', \dot{z}'\}$. If $\omega \in \mathcal{V}_M \setminus \tau$, then either $\omega_x \leq \beta_x$, in which case $\omega \preceq \beta \vee \nu$ implies $\omega \in \{\beta, \nu\}$, or $\omega_z < m + 1$ in addition to $\tau_x \geq \omega_x \geq \beta_x$, in which case $\omega = \beta$ by the technical lemma.

It is easy to verify that \dot{y}' points toward the next vertex after β having z -coordinate zero. Note that such a next vertex must exist, since the rigid geodesic leaving β and containing Z_β strictly decreases in x . This completes the proof of (7).

SPECIAL CONSTRUCTION FOR (8)

The meet (componentwise minimum) of τ and γ is $\tau \wedge \gamma = (\gamma_x, \gamma_y, 0) = \gamma - (0, 0, m + 1)$ since Z_τ points toward γ in $\mathcal{S}_{\mathcal{V}_M}$. Observe that

$$\omega \in \mathcal{V}_M \text{ and } \tau \wedge \gamma \preceq \omega \text{ implies } \omega \in \{\tau, \gamma\}. \quad (9)$$

Indeed, if $\tau \not\preceq \omega$ but $\tau \wedge \gamma \preceq \omega$, then either $\tau_x \geq \omega_x \geq \gamma_x$ or $\tau_y \geq \omega_y \geq \gamma_y$. The hypothesis $\gamma_x > \beta_x$ of (8) plus the technical lemma imply $\omega_z \geq \gamma_z$, whence $\gamma \preceq \omega$.

It follows from (9) that the set

$$\underline{\mathcal{V}}_M = (\mathcal{V}_M \setminus \{\tau, \gamma\}) \cup \tau \wedge \gamma$$

of vectors determines a grid surface $\mathcal{S}_{\underline{\mathcal{V}}_M}$. Use the freedom afforded by Lemma 2.1 to rechoose \mathcal{V}_M and m so that $n + 3$ divides all coordinates of vectors therein, while γ_z as well as the lengths of the orthogonal rays $X_{\tau \wedge \gamma}$ and $Y_{\tau \wedge \gamma} \subset \mathcal{S}_{\underline{\mathcal{V}}_M}$ equal $m + 2$. Applying (9) again, further alter \mathcal{V}_M by moving

τ so that $\tau - \tau \wedge \gamma$ equals $(1, 0, 0)$, $(0, 1, 0)$, or $(1, 1, 0)$, depending on whether $\tau_y = \gamma_y$, $\tau_x = \gamma_x$, or neither.

Now choose \mathcal{V}_N so that all of its *nonzero* x and y -coordinates lie in the interval $[m - n, m]$, but all of its z -coordinates are no greater than n . Let

$$\underline{\mathcal{V}} = (\underline{\mathcal{V}}_M \setminus \tau \wedge \gamma) \cup (\tau \wedge \gamma + \mathcal{V}_N),$$

and denote by $\underline{\nu}$ the vector $\tau \wedge \gamma + \nu$ for $\nu \in \mathcal{V}_N$. Our (final) goal is to show that the rigid geodesics in $\mathcal{S}_{\underline{\mathcal{V}}}$ constitute an embedding of $(M \cup_{\tau} N)/e_{\underline{z}}$. After that, $e_{\underline{x}}$ and $e_{\underline{y}}$ can be contracted by omitting \underline{x} and \underline{y} , using the same arguments appearing in Scale principle, Rays leaving $\tau + \mathcal{V}_N$, Proof of (1), and Proof of (2).

PROOF OF (8)

Begin by mimicking as closely as possible the proof of Lemma 3.2. First, $\tau \wedge \gamma$ is the unique vector in $\underline{\mathcal{V}}_M$ preceding $\tau \wedge \gamma + m\mathbf{1}$, by Lemma 3.2 applied to \mathcal{V}_M ; this is why τ needs to be so close to $\tau \wedge \gamma$. The vertex axiom for $\mathcal{S}_{\underline{\mathcal{V}}}$ is immediate. Moreover, the part of $\mathcal{S}_{\underline{\mathcal{V}}}$ that precedes $\tau \wedge \gamma + m\mathbf{1}$ equals $\tau \wedge \gamma +$ (the part of $\mathcal{S}_{\mathcal{V}_N}$ preceding $m\mathbf{1}$). Thus every vertex, rigid geodesic, or bounded orthogonal ray in $N \hookrightarrow \mathcal{S}_{\mathcal{V}_N}$ gets translated by $\tau \wedge \gamma$ to the corresponding feature in $\mathcal{S}_{\underline{\mathcal{V}}}$. Similarly, the parts of $\mathcal{S}_{\underline{\mathcal{V}}}$ and $\mathcal{S}_{\underline{\mathcal{V}}_M}$ not preceded by $\tau \wedge \gamma$ agree, so any vertex, rigid geodesic, or bounded orthogonal ray in $M \hookrightarrow \mathcal{S}_{\underline{\mathcal{V}}_M}$ survives in $\mathcal{S}_{\underline{\mathcal{V}}}$ whenever it is contained in a (perhaps unbounded) region of M_{∞} not containing τ or γ , by the rigid region axiom.

The only orthogonal rays unaccounted for as yet for the rigid geodesic axiom are $X_{\underline{x}}, Y_{\underline{y}}, Y_{\alpha}, X_{\beta}, Z_{\underline{z}}$, and any orthogonal ray $U_{\nu} \subset \mathcal{S}_{\underline{\mathcal{V}}}$ such that $U_{\nu} \subset \mathcal{S}_{\mathcal{V}_M}$ points toward γ and $\nu \neq \tau$. (Neither the rigid geodesic connecting τ to γ in M nor the orthogonal rays leaving γ in $\mathcal{S}_{\mathcal{V}_M}$ play roles in this verification.) The only case requiring significant effort are the U_{ν} rays, which must point toward \underline{z} in $\mathcal{S}_{\underline{\mathcal{V}}}$.

Suppose $\omega \preceq \nu \vee \underline{z}$ for some $\omega \in \underline{\mathcal{V}}$. The technical lemma and (9) imply that $\nu_z \geq \gamma_z$ for any $\tau \neq \nu \in \mathcal{V}_M$ pointing toward γ , whence $\nu \vee \underline{z} = \nu \vee \gamma$ for any such ν . Therefore $\omega \neq \nu$ implies $\omega \in \tau \wedge \gamma + \mathcal{V}_N$. On the other hand, either $\nu_y < \beta_y$ or $\nu_x < \alpha_x$, because otherwise β or α precedes $\nu \vee \gamma = \nu \vee \underline{z}$. By the choice of scaling, $\omega_y \leq \nu_y < \beta_y$ forces $\omega_y \leq \nu_y \leq \beta_y - (n + 3) < \gamma_y + m - n$, whence $\omega = \underline{z}$ whenever $\omega \in \tau \wedge \gamma + \mathcal{V}_N$. This argument also works with the roles of x and y switched.

The above reasoning proves that the rigid geodesics in $\mathcal{S}_{\underline{\mathcal{V}}}$ embed *some* planar map L . To conclude that $L \cong (M \cup_{\tau} N)/e_{\underline{z}}$, one last item remains: show that no geodesics in M vanish. More precisely, whenever X_{γ} or Y_{γ} does not point toward τ in \mathcal{V}_M , we require it to point toward a vertex in \mathcal{V}_M that points back toward γ in $\mathcal{S}_{\mathcal{V}_M}$. Suppose $Y_{\gamma} \subset \mathcal{S}_{\mathcal{V}_M}$ does not point toward τ . Then $\gamma_x < \tau_x = B_x$, whence $\gamma_z = B_z$ by the region axiom, because $\gamma_y < \beta_y = B_y$. If $B(z) = \{\gamma\}$, then $Y_{\gamma} \subset \mathcal{S}_{\mathcal{V}_M}$ must point toward β , because $\beta_x < \gamma_x$. This is impossible whenever the region B in $M \cup_{\tau} N$ has at least five vertices, by the hypothesis of (8) stipulating that $[\underline{x}, \underline{z}]$ is an edge in N . The analogous

argument works for X_γ , but the reason why X_γ cannot point toward α is different: it is ruled out by the statement of the Proposition. \square

Recall the conventions set before the statement of Proposition 4.2.

PROPOSITION 4.3 *Suppose $\tau = \dot{x}$ is trivalent. Contracting neither, either, or (if B has at least five vertices) both of e_z and $e_{\dot{y}}$ in $M \cup_\tau N$ yields a planar map possessing a rigid embedding.*

Proof. Pretend $e_{\dot{x}}$ has already been contracted, and apply the constructions in the proofs of (6)–(8) in Proposition 4.2. (In reality, only (6) and (7) are required here, given the symmetry switching the roles of y and z). \square

5 TRICONNECTIVITY AND RIGID EMBEDDING

THEOREM 5.1 (RIGID EMBEDDING) *A planar map M with given axial vertices $\dot{x}, \dot{y}, \dot{z}$ can be rigidly embedded in a grid surface if and only if the extended map $M_\infty(\dot{x}, \dot{y}, \dot{z})$ is triconnected. In particular, every triconnected planar map can be rigidly embedded.*

Proof. (\Rightarrow) Let $M \hookrightarrow \mathcal{S}_V$ be an axial rigid embedding, and delete two vertices ν, ω from M_∞ . If $\{\nu, \omega\} \subset \{\dot{x}, \dot{y}, \dot{z}\}$ then each remaining vertex has an orthogonal flow to the third axial vertex, by independence of the orthogonal flows to $\dot{x}, \dot{y}, \dot{z}$. If $\{\nu, \omega\} \not\subset \{\dot{x}, \dot{y}, \dot{z}\}$ then what remains of the exterior cycle is connected, and every vertex in the deletion still has an orthogonal flow to the exterior cycle.

(\Leftarrow) Induct on the sum of the number of regions and the number of edges in M , observing that the minimal sum of five is attained only when M is a triangle. Assume the notation of Section 3, and suppose M is not a triangle. Letting e be the edge leaving \dot{x} towards \dot{y} on the exterior cycle of M , we claim that at least one of the following occurs:

1. The endpoints of e are \dot{x} and \dot{y} .
2. The edge e does not border a triangle in $M_\infty(\dot{x}, \dot{y}, \dot{z})$, and M/e is triconnected.
3. The edge e does not contain \dot{y} , and M contains a proper ring C for which $\ddot{x} = \infty$.
4. The edge e does not contain \dot{y} , and M contains a proper ring C for which
 - (a) $\ddot{x} = \ddot{x} = \dot{x}$;
 - (b) \ddot{y} is the other endpoint of e , and $\dot{y} \neq \ddot{y}$ (that is, $\ddot{y} \neq \infty$); and
 - (c) \ddot{z} does not lie between \dot{x} and \dot{z} on the exterior cycle of M , and no single edge outside the region bounded by C has endpoints $\{\dot{x}, \ddot{z}\}$.

If the first three cases do not occur, then Proposition 1.1 produces a ring C in which \ddot{x} and \ddot{y} are the two endpoints of e , while \ddot{z} is a vertex in $F \cap F'$ as in

Proposition 1.1.2. Note that C is proper because $\ddot{y} \neq \dot{y}$. Let C be a maximal such ring.

The first half of 4(c) holds; if not, construct a ring satisfying option 3 as follows: replace the arc of C connecting \dot{x} to \dot{z} with the arc traversing the exterior cycle from \dot{x} to \ddot{z} and then $e_{\ddot{z}}$ (the latter only if $\dot{z} \neq \ddot{z}$). The second half of 4(c) also holds, for if an edge f outside C connects \dot{x} to \ddot{z} , then replace the arc of C connecting \dot{x} to \dot{z} in C by f and $e_{\dot{z}}$ (if $\ddot{z} \neq \dot{z}$). The resulting cycle is a larger ring satisfying the condition defining C , contradicting maximality. Finally, 4(a) holds by the failure of option 3: C does not contain the edge of M leaving \dot{x} toward \dot{z} on the exterior cycle of M .

Given the first option, M has a proper ring C , with $\ddot{z} = \dot{z} = \dot{z}$, containing every bounded region of M except the one containing e . Let $M \cap R \hookrightarrow \mathcal{S}_{\mathcal{V}}$ be a geodesic embedding, where R is the union of regions contained within C . Leaving \dot{x}, \dot{y} fixed while adding 1 to the z -coordinates of every other vector in \mathcal{V} yields a grid surface $\mathcal{S}_{\mathcal{V}'}$ whose rigid geodesics constitute an embedding of M ; the easy proof is omitted.

Given the second option, use Proposition 4.1. Given the third or fourth option, use Lemma 3.1 with $M = M/R$ and $N = M \cap R$, along with Proposition 4.3 for option 3, or Lemma 3.2 and Proposition 4.2 for option 4. The ‘five vertex’ conditions in Propositions 4.2 and 4.3 are always satisfied when reconstructing M from the tethered gluing $(M/R) \cup_{\tau} (M \cap R)$, because M is simple and triconnected. \square

The next corollary clarifies the close connection between grid surfaces and order dimension for posets. It shows that Theorem 5.1 generalizes the three-variable special case of [BPS98, Theorem 6.4], which is presented in the equivalent language of monomial ideals.

COROLLARY 5.2 (BRIGHTWELL–TROTTER [BT93]) *The vertices, edges, and bounded regions of any triconnected planar map form a partially ordered set of order dimension ≤ 3 .*

Proof. Theorem 5.1 and Corollary 2.5. \square

EXAMPLE 5.3 Theorem 5.1 is stronger than Corollary 5.2, even for triconnected maps. In general, every inclusion of the vertex-edge poset of M into \mathbb{N}^3 yields an inclusion of the vertex set $\mathcal{V} \hookrightarrow \mathcal{S}_{\mathcal{V}}$ such that each edge of M is a rigid geodesic in $\mathcal{S}_{\mathcal{V}}$. What fails is that there may be orthogonal rays in $\mathcal{S}_{\mathcal{V}}$ that are not contained in any edges of M . Faces such as the central face in Figure 6 are then forced to lie off of $\mathcal{S}_{\mathcal{V}}$. \square

REMARK 5.4 Rigid embeddings give a fresh perspective on a standard fact, known as Menger’s theorem: If G is a triconnected planar graph and $\nu, \omega \in G$ are distinct vertices, then there are three independent paths from ν to ω in G . To explain Menger’s theorem via Theorem 5.1, let M be a plane drawing of G , and suppose e_1, \dots, e_r are the edges of M containing ν , in cyclic order. Form




Figure 6: A vertex-edge-face poset embedding that is not a geodesic embedding

a new map M' by drawing a small circle C around ν and adding new vertices ν_1, \dots, ν_r where e_1, \dots, e_r intersect C . Then set $M^\nu = \text{del}(M'; \nu)$, with underlying graph G^ν . Clearly G^ν is triconnected.

Choose a plane drawing of G^ν in which C is the exterior cycle, and let $G^\nu \hookrightarrow \mathcal{S}$ be an axial geodesic embedding compatible with some (any) choice of axial vertices $\hat{x}, \hat{y}, \hat{z} \in C$. The orthogonal flows from ω to $\hat{x}, \hat{y}, \hat{z}$ in G^ν first intersect C at points $\nu_{i_x}, \nu_{i_y}, \nu_{i_z}$, giving rise to truncated orthogonal flows $[\omega, \nu_{i_u}]$ for $u \in \{x, y, z\}$. Connecting $[\omega, \nu_{i_u}]$ to ν via the arc in G between ν_{i_u} and ν yields independent paths in G from ω to ν .

PART II

MONOMIAL IDEALS

6 BETTI NUMBERS

Let k be a field, and consider the polynomial ring $R = k[x, y, z]$ with the \mathbb{Z}^3 -grading in which $\deg(x) = (1, 0, 0)$, $\deg(y) = (0, 1, 0)$, and $\deg(z) = (0, 0, 1)$. Use $I_{\mathcal{V}} = \langle m^\nu \mid \nu \in \mathcal{V} \rangle$ to denote the ideal generated the monomials $m^\nu = x^{\nu_x} y^{\nu_y} z^{\nu_z}$ for $\nu \in \mathcal{V}$. The integer points in $\langle \mathcal{V} \rangle$ coincide with the exponent vectors on monomials in $I_{\mathcal{V}}$, and \mathcal{V} is axial if and only if $I_{\mathcal{V}}$ is artinian, containing a power of each variable.

Any principal monomial ideal $\langle m \rangle$ is a free \mathbb{Z}^3 -graded R -module of rank 1. If ϕ is \mathbb{Z}^3 -graded homomorphism $\bigoplus \langle m_q \rangle \leftarrow \bigoplus \langle m_p \rangle$ of degree zero, then we can express ϕ as a **monomial matrix**. This is a matrix whose entries $\lambda_{pq} \in k$ are scalars, and whose p^{th} row (resp. q^{th} column) is labeled by the monomial m_p (resp. m_q) that generates the corresponding p^{th} source (resp. q^{th} target) summand. Of course, $\lambda_{pq} = 0$ whenever m_q does not divide m_p , because then there are no nonzero \mathbb{Z}^3 -graded maps $\langle m_q \rangle \leftarrow \langle m_p \rangle$. The map ϕ is called **minimal** if also $\lambda_{pq} = 0$ whenever $m_p = m_q$. See [Mil00a, Section 2] for more on monomial matrices.

We consider **free resolutions** of $I_{\mathcal{V}}$ that are exact sequences having the form

$$\mathcal{F} : \quad 0 \leftarrow I_{\mathcal{V}} \xleftarrow{\phi_0} \mathcal{F}_0 \xleftarrow{\phi_1} \mathcal{F}_1 \xleftarrow{\phi_2} \mathcal{F}_2 \leftarrow 0, \quad (10)$$

in which $\mathcal{F}_i \cong \bigoplus_p \langle m_{ip} \rangle$ for some (finite set of) monomials $m_{ip} \in R$. We call \mathcal{F} .

minimal if ϕ_1 and ϕ_2 are minimal (for any such direct sum decomposition). The **Betti number** $\beta_{i,\alpha}(I_{\mathcal{V}})$ is the number of m_{ip} equal to m^α , when \mathcal{F} is minimal. This homological data reflects the local properties of $\mathcal{S}_{\mathcal{V}}$ near the vector α via the **Koszul simplicial complex** of \mathcal{V} at $\alpha \in \mathbb{N}^3$,

$$K_\alpha(\mathcal{V}) = \{\sigma \in \{0, 1\}^3 \mid \alpha - \sigma \in \langle \mathcal{V} \rangle\},$$

which is a subcomplex of the abstract triangle $(\{0, 1\}^3, \preceq)$.

PROPOSITION 6.1 ([HOC77, ROZ70]) $\beta_{i,\alpha}(I_{\mathcal{V}}) = \dim_k \tilde{H}_{i-1}(K_\alpha(\mathcal{V}); k)$ is the dimension of the $(i-1)^{\text{st}}$ reduced simplicial homology of $K_\alpha(\mathcal{V})$ with coefficients in k .

The small number of simplicial complexes on three vertices seriously limits the possibilities for nonzero Betti numbers.

LEMMA 6.2 If $\mathcal{V} \subset \mathbb{N}^3$ and $i \in \mathbb{N}$ then $\beta_{i,\alpha}(I_{\mathcal{V}}) \neq 0$ for at most one $\alpha \in \mathbb{Z}^3$. If $\beta_{i,\alpha}$ is nonzero then $\beta_{i,\alpha} = 1$ unless $K_\alpha(\mathcal{V})$ has 3 vertices and no edges (so $\beta_{1,\alpha} = 2$).

Proof. Use the previous proposition, and list all simplicial complexes on 3 vertices. □

7 CELLULAR RESOLUTIONS

Suppose M is a cell complex (precisely, a finite CW complex) of dimension 2 whose cells P have vector labels $\alpha_P \in \mathbb{N}^3$, in such a way that $\alpha_P \preceq \alpha_{P'}$ if P lies in the closure of P' . For instance, if the vertices have natural labels, then define α_P to be the join of the labels on the vertices of P . Such a **labeled cell complex** determines monomial matrices ϕ_{vertex} , ϕ_{edge} , and ϕ_{region} for a **cellular free complex** \mathcal{F}_M , by labeling the rows and columns of matrices for the boundary map of the ordinary chain complex of M with the monomials m^{α_P} . Using P also to denote the basis vector of a rank 1 free R -module $\langle m^{\alpha_P} \rangle$, the cellular free complex \mathcal{F}_M takes the form

$$0 \leftarrow I_{\mathcal{V}} \xleftarrow{\phi_{\text{vertex}}} \bigoplus_{\text{vertices } \nu} R \cdot \nu \xleftarrow{\phi_{\text{edge}}} \bigoplus_{\text{edges } e} R \cdot e \xleftarrow{\phi_{\text{region}}} \bigoplus_{\text{regions } F} R \cdot F \leftarrow 0. \tag{11}$$

We say M **supports** \mathcal{F}_M ; see [BS98, Mil00a] for more on cellular monomial matrices.

Our main examples of labeled cell complexes are of course the geodesically embedded planar maps $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$, whose labels are determined by the obvious vertex labels by taking joins. We shall see in the next section that this cellular free complex is exact and minimal, so it provides a **cellular minimal free resolution** of $I_{\mathcal{V}}$. Unfortunately, there are monomial ideals whose associated grid surfaces contain no geodesically embedded map.

EXAMPLE 7.1 Let $I_{\mathcal{V}} = \langle x, y, z \rangle^2 = \langle x^2, y^2, z^2, xy, xz, yz \rangle$. The orthogonal rays X_{yz}, Y_{xz}, Z_{xy} meet at a single point not in \mathcal{V} , so there can be no planar map geodesically embedded in $\mathcal{S}_{\mathcal{V}}$. However, $\langle x, y, z \rangle^2$ still has a minimal cellular resolution: connect the midpoints of the edges of a triangle, and delete any one of the three interior edges. Label the resulting planar map M with x^2, y^2, z^2 on the corners of the outside triangle; xy between x^2 and y^2 ; yz between y^2 and z^2 ; and xz between x^2 and z^2 . Label the edges and regions of M by the joins of their vertex labels. \square

Although the map M in the above example fails to embed geodesically in $\mathcal{S}_{\mathcal{V}}$, the extended map $M_{\infty}(x^2, y^2, z^2)$ is still triconnected. This phenomenon is general. In the following proposition, we do not require planarity of M , so we use $G_{\infty}(\hat{x}, \hat{y}, \hat{z})$ to mean the abstract graph obtained from G by adding a new vertex ∞ connected to each of $\hat{x}, \hat{y}, \hat{z}$.

PROPOSITION 7.2 *If the labeled cell complex M supports a minimal free resolution of an artinian ideal $I_{\mathcal{V}}$, and the 1-skeleton of M is a graph G , then the extended graph $G_{\infty}(\hat{x}, \hat{y}, \hat{z})$ is triconnected, where $\hat{x}, \hat{y}, \hat{z}$ are the vertices whose labels lie on the axes.*

Proof. Given $\nu \in \mathcal{V}$ with $\nu \neq \hat{z}$, the orthogonal ray Z_{ν} leaving ν has its head at some vector $\alpha \in \mathbb{N}^3$ for which $K_{\alpha}(\mathcal{V})$ is disconnected; indeed, the vertex $(0, 0, 1)$ is isolated in $K_{\alpha}(\mathcal{V})$. Choose a vertex $\omega' \neq \nu$ preceding α , so that $m^{\alpha-\nu}\nu - m^{\alpha-\omega'}\omega' \in \ker(\phi_{\text{vertex}})$. Since $\alpha - \nu = Z(\nu)$ lies on the z -axis Z , there must be an edge $e \in M$ connecting ν to a vertex ω (possibly different from ω') such that $\phi_{\text{edge}}(e) = z^d\nu - m^{\nu\vee\omega-\omega}\omega$ for some $d \in \mathbb{N}$. Clearly $\omega_x \preceq \nu_x$ and $\omega_y \preceq \nu_y$.

Repeating the procedure with ω in place of ν , and with x or y in place of z , we find that M contains paths analogous to orthogonal flows from ν to each axial vertex $\hat{x}, \hat{y}, \hat{z}$. As in Section 2, these paths are independent, intersecting only at ν . \square

8 GRAPHS TO MINIMAL RESOLUTIONS

LEMMA 8.1 *If $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ is a geodesic embedding, then $\beta_{1,\alpha}(I_{\mathcal{V}}) \neq 0$ if and only if $\alpha = \nu\vee\omega$ for some elbow geodesic $[\nu, \omega] \in M$, and $\beta_{1,\nu\vee\omega}(I_{\mathcal{V}}) = 1$ in this case.*

Proof. Assume $\beta_{1,\alpha}(I_{\mathcal{V}}) \neq 0$. Then $K_{\alpha}(\mathcal{V})$ is disconnected by Proposition 6.1, and this occurs if and only if $K_{\alpha}(\mathcal{V})$ contains an isolated vertex. An isolated vertex of $K_{\alpha}(\mathcal{V})$ occurs if and only if α lies on an orthogonal ray leaving some vertex $\nu \in \mathcal{V}$. Therefore, $\alpha = \nu\vee\omega$ for some $\omega \neq \nu$ by the edge axiom. Since $[\nu, \omega]$ is an elbow geodesic, $K_{\nu\vee\omega}(\mathcal{V})$ cannot have three isolated points because three orthogonal rays cannot meet at the point $\nu\vee\omega$ in the relative interior of the edge $[\nu, \omega]$ of M . \square

Here is a result that sometimes reduces statements about arbitrary geodesic or rigid embeddings to axial ones. Its straightforward proof is omitted.

LEMMA 8.2 *Append axial vertices to \mathcal{V} by letting $\overline{\mathcal{V}} = \mathcal{V} \cup \{\dot{u} \mid \mathcal{S}_{\mathcal{V}} \cap U = \emptyset\}$ for sufficiently large $|\dot{u}|$. A planar map M is geodesically embedded in $\mathcal{S}_{\overline{\mathcal{V}}}$ if and only if $N = \text{del}(M; \overline{\mathcal{V}} \setminus \mathcal{V})$ is geodesically embedded in $\mathcal{S}_{\mathcal{V}}$. Furthermore, $M \hookrightarrow \mathcal{S}_{\overline{\mathcal{V}}}$ is rigidly embedded if and only if $N \hookrightarrow \mathcal{S}_{\mathcal{V}}$ is rigidly embedded.*

LEMMA 8.3 *Let $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ be a geodesic embedding, and $\mathcal{S}_{\mathcal{V}}^{\max}$ the set of points in $\mathcal{S}_{\mathcal{V}}$ maximal under the partial order induced by the relation \preceq on \mathbb{R}^3 . Then $\alpha \in \mathcal{S}_{\mathcal{V}}^{\max} \Leftrightarrow \beta_{2,\alpha}(I_{\mathcal{V}}) \neq 0 \Leftrightarrow \alpha = \alpha_F$ is the join of the vertices in a bounded region F of M .*

Proof. For the equivalence $\alpha \in \mathcal{S}_{\mathcal{V}}^{\max} \Leftrightarrow \beta_{2,\alpha}(I_{\mathcal{V}}) \neq 0$, use the fact that $K_{\alpha}(\mathcal{V})$ is the boundary of the triangle; the easy details are omitted. The equivalence $\alpha \in \mathcal{S}_{\mathcal{V}}^{\max} \Leftrightarrow \alpha = \alpha_F$ holds for all vertex sets \mathcal{V} if it holds when \mathcal{V} is axial. Indeed, using the notation and result of Lemma 8.2, the maximal points of $\mathcal{S}_{\mathcal{V}}$ are still maximal in $\mathcal{S}_{\overline{\mathcal{V}}}$, while the points in $\mathcal{S}_{\overline{\mathcal{V}}}^{\max} \setminus \mathcal{S}_{\mathcal{V}}^{\max}$ are exactly those having u -coordinate $|\dot{u}|$ for some u such that $\mathcal{S}_{\mathcal{V}} \cap U = \emptyset$, by the region axiom applied to $\mathcal{S}_{\overline{\mathcal{V}}}$. The bounded regions of N having such joins disappear upon deleting \dot{u} .

Assume henceforth that $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ is axial. If $\rho \in \mathcal{S}_{\mathcal{V}}$ has some coordinate $\rho_u = 0$, then $\rho \notin \mathcal{S}_{\mathcal{V}}^{\max}$ because adding ε to any other coordinate of ρ yields another point in $\mathcal{S}_{\mathcal{V}}$. Therefore each maximal point of $\mathcal{S}_{\mathcal{V}}$ lies in a bounded region of M . When F is such a bounded region, the region axiom implies $\alpha_F \in \mathcal{S}_{\mathcal{V}}^{\max}$, because some vertex $\nu \in \mathcal{V}$ strongly precedes $\alpha_F + \varepsilon\dot{u}$ for any $\varepsilon > 0$ and $u \in \{x, y, z\}$.

Every point ρ on a given elbow geodesic $[\nu, \omega]$ precedes $\nu\vee\omega$ by definition, so $\rho \preceq \nu\vee\omega \preceq \alpha_F$ whenever $[\nu, \omega] \subseteq F$. Any point σ on the line segment in \mathbb{R}^3 connecting ρ to α_F therefore satisfies $\rho \preceq \sigma \preceq \alpha_F$, whence $\sigma \in \mathcal{S}_{\mathcal{V}}$. It follows that F is the union of such line segments, so every point of F precedes α_F . \square

THEOREM 8.4 *Given a geodesic embedding $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$, the cellular free complex \mathcal{F}_M is a minimal free resolution of $I_{\mathcal{V}}$.*

Proof. Since $\beta_{i,\alpha}(I_{\mathcal{V}}) \neq 0$ if and only if $\beta_{i,\alpha}(I_{\mathcal{V}}) = 1$ by Lemmas 6.2 and 8.1, it makes sense simply to speak of the i^{th} **Betti degrees** α , for which $\beta_{i,\alpha} = 1$. The zeroth, first, and second Betti degrees are the labels on the vertices, edges, and regions of M , respectively, by Lemmas 6.2, 8.1, and 8.3. Any minimal free resolution \mathcal{F} of $I_{\mathcal{V}}$ as in (10) therefore takes the form of (11), at least as a homologically graded module; that is, $\mathcal{F} \cong \mathcal{F}_M$ abstractly as modules. We need to show that some choice of this abstract isomorphism is a homomorphism of complexes.

Identifying the homological degree zero parts of \mathcal{F} and \mathcal{F}_M , the zeroth homology of \mathcal{F}_M surjects onto $I_{\mathcal{V}}$ because the image of ϕ_{edge} is clearly contained in the kernel of ϕ_0 . Since \mathcal{F} is exact and \mathcal{F}_M is a complex of free modules, there

exists a homomorphism $\psi : \mathcal{F}_M \rightarrow \mathcal{F}$ lifting the surjection on zeroth homology and the isomorphism in homological degree zero.

Suppose $e \in \mathcal{F}_M$ maps to $\psi(e) = \sum m_j e'_j$, where each $m_j \in R$ is a monomial with nonzero scalar coefficient, and each e'_j denotes the generator of \mathcal{F}_1 corresponding to the edge $e_j \in M$. The elbow geodesic $e = [\nu, \omega]$ in M contains an orthogonal ray U_ν , so that $\pm \phi_1(\sum m_j e'_j) = m^{\nu \vee \omega - \nu} \nu - m^{\nu \vee \omega - \omega} \omega$, where the first term is $m^{\nu \vee \omega - \nu} \nu = m^{U(\nu)} \nu$. Thus $m^{U(\nu)} \nu = u^{|U_\nu|} \nu$ appears with nonzero scalar coefficient in $\phi_1(m_j e'_j)$ for some j . Since there is a *unique* first Betti degree α satisfying $\alpha - \nu \in U$, namely $\alpha_e = \nu \vee \omega$, it must be that $e'_j = e'$ and m_j is a nonzero scalar. Nakayama's lemma implies $\psi_1 : (\mathcal{F}_M)_1 \rightarrow \mathcal{F}_1$ is surjective, and hence an isomorphism by rank considerations.

No summand $R \cdot F \subset \mathcal{F}_M$ can map to zero in \mathcal{F}_2 because $\phi_{\text{region}}(F)$ is nonzero in \mathcal{F}_M , and ψ is an isomorphism in homological degree 1. On the other hand, the second Betti degrees are pairwise incomparable by Lemma 8.3. Thus $\psi(F)$ is some nonzero scalar multiple of the unique generator of \mathcal{F}_2 in degree α_F . The map ψ is therefore an isomorphism in homological degree 2, completing the proof. \square

COROLLARY 8.5 *A planar map M with axial vertices $\dot{x}, \dot{y}, \dot{z}$ supports a minimal free resolution of an artinian monomial ideal if and only if $M_\infty(\dot{x}, \dot{y}, \dot{z})$ is triconnected. In particular, every triconnected planar map supports a minimal free resolution.*

Proof. ‘Only if’ is Proposition 7.2; apply Theorem 8.4 to Theorem 5.1 for ‘if’. \square

9 UNIQUENESS VS. NONPLANARITY

Continuing with the analogy at the beginning of the Introduction, circle packings and polytopes that realize planar graphs are unique up to Möbius transformation and spherical rotation, respectively (see [Zie95] for discussion and references). Rigid embeddings $M \hookrightarrow \mathcal{S}_\mathcal{V}$ for a fixed planar map are similarly not unique: at the very least, any order-preserving bijection of \mathcal{V} as in Lemma 2.1 gives another rigid embedding. Of course, such bijections affect neither the combinatorics nor the algebra. In fact, rigid embeddings are uniquely determined by the algebraic properties of the grid surface in question, specifically the minimal free resolution of the corresponding monomial ideal.

COROLLARY 9.1 *When $M \hookrightarrow \mathcal{S}_\mathcal{V}$ is rigidly embedded, M is the unique cell complex supporting a minimal cellular free resolution of $I_\mathcal{V}$.*

Proof. Let N be a labeled cell complex supporting a minimal cellular free resolution of $I_\mathcal{V}$. The abstract graph underlying N (the 1-skeleton) coincides with that of M by Theorem 8.4 and rigidity. The label α_F on any region F of N is a second Betti degree of $I_\mathcal{V}$, and hence coincides with the label on a region F' of M by Theorem 8.4. The boundary of F in N is a cycle of edges




Figure 7: Nonplanar minimal free resolution

whose degrees precede α_F . The only such cycle consists of all the edges whose degrees precede α_F , by the rigid region axiom for F' in M (by Lemma 8.2 the rigid region axiom holds for nonaxial grid surfaces). \square

Nonrigid monomial ideals can have many distinct isomorphism classes of minimal cellular resolutions; [MS99, Figure 4] depicts an example of this phenomenon. In fact, the bad behavior gets much worse.

EXAMPLE 9.2 Minimal cellular resolutions of ideals in $k[x, y, z]$ need not be supported planar cell complexes. In fact, explicit examples crop up with even the smallest violations of rigidity. For instance, let

$$\mathcal{V} = \{(4, 3, 0), (3, 4, 0), (3, 0, 2), (2, 1, 1), (1, 2, 1), (0, 3, 2)\}$$

and $\bar{\mathcal{V}} = \mathcal{V} \cup \{(5, 0, 0), (0, 5, 0), (0, 0, 3)\}$. The cell complex M depicted in Figure 7 consists of five triangles in addition to the three quadrilaterals with vertices

$$\{500, 430, 121, 211\}, \{430, 121, 211, 341\}, \{050, 121, 211, 340\}.$$

Label the edges and regions by the joins of their vertex labels. That M supports a minimal free resolution of $I_{\bar{\mathcal{V}}}$ can be checked by verifying for each $\alpha \preceq (5, 5, 3)$ that $M_{\preceq \alpha}$ is acyclic [BS98, Corollary 1.3]. That M cannot be planar follows by contracting the edges labeled 530 and 350 while deleting the edges labeled 312 and 132 to get the complete graph K_5 as a minor of the 1-skeleton. \square

10 DEFORMATION AND GENERICITY

Given a finite subset $\mathcal{W} \subset \mathbb{R}^3$, write $\vee \mathcal{W}$ for the join of the vectors in \mathcal{W} , and set $m^{\mathcal{W}} = m^{\vee \mathcal{W}}$. Following [BPS98], define the **Scarf complex** of \mathcal{V} ,

$$\Delta_{\mathcal{V}} = \{\mathcal{W} \subseteq \mathcal{V} \mid \text{if } \vee \mathcal{W}' = \vee \mathcal{W} \text{ for some } \mathcal{W}' \subseteq \mathcal{V} \text{ then } \mathcal{W}' = \mathcal{W}\},$$

to consist of the subsets whose joins are uniquely attained. It is an easy (but not obvious) fact that $\Delta_{\mathcal{V}}$ is a simplicial complex. Each face $\mathcal{W} \in \Delta_{\mathcal{V}}$ comes

with a natural label $\vee\mathcal{W}$, and the resulting cellular free complex $\mathcal{F}_{\Delta_{\mathcal{V}}}$ is called the **free Scarf complex** of $I_{\mathcal{V}}$. The Scarf complex is planar by virtue of its containment in the **hull complex** [BS98], so the union of its edges and vertices is a planar map. This planar map has already appeared, in Section 2: two vertices $\nu, \omega \in \mathcal{V}$ are connected by a rigid geodesic if and only if $\{\nu, \omega\} \in \Delta_{\mathcal{V}}$. Under special circumstances, the Scarf complex is rigidly embedded in $\mathcal{S}_{\mathcal{V}}$. To be precise, call \mathcal{V} **strongly generic** if no two distinct elements of \mathcal{V} share a nonzero coordinate. In other words, $\nu_u = \omega_u \neq 0$ for some $u \in \{x, y, z\}$ implies $\nu = \omega$.

COROLLARY 10.1 (BAYER–PEEVA–STURMFELS [BPS98, §3]) *The free Scarf complex $\mathcal{F}_{\Delta_{\mathcal{V}}}$ minimally resolves $I_{\mathcal{V}}$ when \mathcal{V} is strongly generic.*

Proof. Strong genericity easily implies that every orthogonal ray is contained in a rigid geodesic, so the Scarf graph rigidly embeds in $\mathcal{S}_{\mathcal{V}}$. It is straightforward to verify that the labels on triangles (2-dimensional faces) in $\Delta_{\mathcal{V}}$ are maximal in $\mathcal{S}_{\mathcal{V}}$. Furthermore, every maximal point of $\mathcal{S}_{\mathcal{V}}$ has exactly three vectors in \mathcal{V} preceding it by Lemma 8.3, the region axiom, and strong genericity. Therefore, all maximal points are labels on regions in $\Delta_{\mathcal{V}}$, and the result holds by Theorem 8.4. \square

Since the definition of the Scarf complex depends only on the coordinatewise order of the exponents of the generators, it also makes sense for (formal) monomials with real exponents in \mathbb{R}^n . This makes way for the following definition. Let \mathbb{Q} denote the rational numbers. A **deformation** ϵ of \mathcal{V} is a choice of vectors $\epsilon^{\nu} = (\epsilon_x^{\nu}, \epsilon_y^{\nu}, \epsilon_z^{\nu}) \in \mathbb{Q}^3$ for each $\nu \in \mathcal{V}$ satisfying

$$\nu_u < \omega_u \Rightarrow \nu_u + \epsilon_u^{\nu} < \omega_u + \epsilon_u^{\omega}, \quad \text{and} \quad \nu_u = 0 \Rightarrow \epsilon_u^{\nu} = 0$$

for $u \in \{x, y, z\}$. In practice, everything we do is invariant under scaling of \mathcal{V} , so we will always assume $\mathcal{V}^{\epsilon} = \{\nu + \epsilon^{\nu} \mid \nu \in \mathcal{V}\}$ consists of integer vectors. Set $\nu^{\epsilon} = \nu + \epsilon^{\nu}$.

The sole purpose of the ϵ vectors is to break ties one way or the other between equal nonzero coordinates of vectors in \mathcal{V} . In this manner, deformations of \mathcal{V} are closer to being generic than \mathcal{V} is. The verb **specialize** is used here to indicate that a deformation (generization) is being reversed; thus \mathcal{V} is a **specialization** of \mathcal{V}^{ϵ} if the latter is a deformation of the former.

One particular deformation will play a key role in the coming sections. To define it, let $\mathcal{V}(u, a) = \{\nu \in \mathcal{V} \mid \nu_u = a\}$ for each $0 < a \in \mathbb{N}$ and $u \in \{x, y, z\}$. Up to order-preserving bijection (as in Lemma 2.1), there is a unique deformation ϵ satisfying the following condition as well as its analogues via cyclic permutation of x, y, z :

$$\text{If the elements of } \mathcal{V}(z, a) \text{ satisfy } \nu_x > \cdots > \omega_x, \text{ then } a = \nu_z^{\epsilon} < \cdots < \omega_z^{\epsilon}. \quad (12)$$

Note that $\nu_x > \cdots > \omega_x$ is equivalent to $\nu_y < \cdots < \omega_y$ for elements of $\mathcal{V}(z, a)$, by pairwise incomparability. Thus, looking down the x -axis, ϵ raises the vectors in $\mathcal{V}(z, a)$ higher as they move to the right.

11 IDEALS TO GRAPHS: ALGORITHM

Arbitrary monomial ideals in more than three variables need not have minimal cellular free resolutions [RW01], but limiting to three variables forces better behavior.

THEOREM 11.1 *Any monomial ideal $I_{\mathcal{V}} \subset k[x, y, z]$ has a cellular minimal free resolution supported on a labeled planar map M .*

The proof (at the end of Section 12) will reduce to the artinian case, for which Algorithm 11.2 produces M . The justification of Algorithm 11.2 appears in Section 12. Heuristically, the idea is to apply the deformation ϵ of (12), and show that specializing $\mathcal{V}^\epsilon \rightsquigarrow \mathcal{V}$ step by step makes the spurious edges in the Scarf triangulation disappear one at a time.

More precisely, the algorithm specializes \mathcal{V}^ϵ back to \mathcal{V} by making strict inequalities $\nu_u^\epsilon < \omega_u^\epsilon$ into equalities $\nu_u = \omega_u$, judiciously and one at a time. Before each specialization step, the (already partially specialized) ideal has a cellular minimal resolution by induction; after each specialization step, the same planar map still supports a cellular free resolution, although it may not be minimal. However, in the nonminimal case, minimality is achieved by removing exactly one edge.

ALGORITHM 11.2

```

INPUT  an artinian ideal  $I_{\mathcal{V}} \subset k[x, y, z]$ 
OUTPUT a planar map  $M$  supporting a cellular minimal free resolution of  $I_{\mathcal{V}}$ 
INITIALIZE   $\epsilon :=$  the deformation of  $I_{\mathcal{V}}$  in (12) to a strongly generic ideal  $I_{\mathcal{V}^\epsilon}$ 
             $M :=$  Scarf complex of  $\mathcal{V}^\epsilon$ 
WHILE   $\mathcal{V}^\epsilon \neq \mathcal{V}$  DO
    CHOOSE  $\nu \in \mathcal{V}$  and  $u \in \{x, y, z\}$  such that  $\nu_u^\epsilon \neq \nu_u$  and  $\nu_u^\epsilon$  is minimal;
           for ease of notation, assume  $u = x$  by applying a cyclic
           permutation of  $\{x, y, z\}$  translating  $u$  to  $x$ , if necessary
            $\gamma_x := \nu_x$ 
            $\gamma_y := \nu_y^\epsilon + |Y_{\nu^\epsilon}|$ 
            $\gamma_z := \min_{\omega \neq \nu} \{ \omega_z^\epsilon \mid \omega_x^\epsilon \leq \nu_x \text{ and } \omega_y^\epsilon < \gamma_y \}$ 
            $\rho :=$  the element of  $\{ \omega^\epsilon \in \mathcal{V}^\epsilon \mid \omega_y^\epsilon = \gamma_y \text{ and } \omega_z^\epsilon < \gamma_z \}$  with
           maximal  $\omega_z^\epsilon$ 
    REDEFINE  $\epsilon$  by replacing the  $x$ -coordinate  $\epsilon_x^\nu$  with 0
             $M$  by changing the label on  $\nu^\epsilon$  accordingly
            IF  $\rho_x = \nu_x$ 
                THEN redefine  $M$  by removing edge labeled  $\gamma = (\gamma_x, \gamma_y, \gamma_z)$ 
                ELSE leave  $M$  unchanged
            END IF-THEN-ELSE
    END WHILE-DO
OUTPUT   $M$ 

```

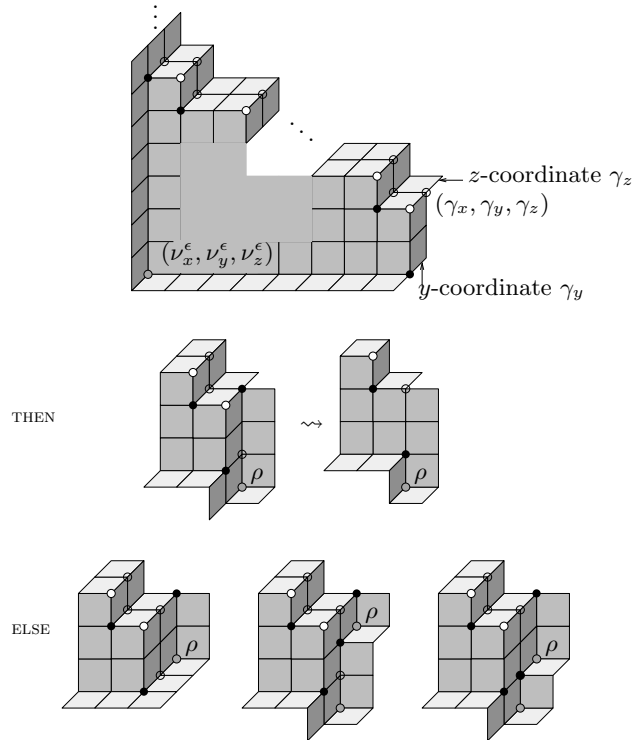


Figure 8: The geometry of Algorithm 11.2

REMARK 11.3 Here are some elementary observations to aid in parsing the algorithm. See also Figure 8, which illustrates the geometry.

- (i) It may be necessary to scale \mathcal{V} in order to choose $\epsilon^\nu \in \mathbb{N}^3$ for all ν , because of the condition $a = \nu_u^\epsilon$.
- (ii) Note that $\gamma_x = \nu_x$, not ν_x^ϵ .
- (iii) The orthogonal ray Y_{ν^ϵ} used to define γ_y is bounded because $\nu^\epsilon \neq \nu$, so that ν cannot be the axial vertex \dot{y} .
- (iv) The set used to define γ_z is nonempty because the axial vector \dot{z} is in the set; indeed, $\nu^\epsilon \neq \nu$ implies $\nu \neq \dot{z}$.
- (v) $\gamma_z > \nu_z^\epsilon$ because \mathcal{V}^ϵ consists of pairwise incomparable vectors.
- (vi) The set defining ρ is nonempty because $\nu^\epsilon + Y_{\nu^\epsilon} = \nu^\epsilon \vee \omega^\epsilon$ for some ω^ϵ in this set. Uniqueness of ρ follows from pairwise incomparability of elements in \mathcal{V}^ϵ .
- (vii) Relabeling M at the REDEFINE step yields a cellular free resolution of the resulting ideal with the new ϵ by [GPW00, Theorem 3.3], though it need not be minimal.

EXAMPLE 11.4 If $I_{\mathcal{V}} = \langle x^4, x^2y^2, x^2z^2, y^4, y^2z^2, z^4 \rangle$ is obtained from Example 7.1 by a scale factor of 2, then the generic deformation $I_{\mathcal{V}^\epsilon} = \langle x^4, x^2y^3, x^3z^2, y^4, y^2z^3, x^4 \rangle$ satisfies the condition of Algorithm 11.2. Furthermore, the Scarf complex of $I_{\mathcal{V}^\epsilon}$ is the triangle with its edge-midpoints connected, as in Example 7.1. If Algorithm 11.2 is run on this $I_{\mathcal{V}}$, then one of the three nonminimal edges is removed on the first iteration of the WHILE loop. Precisely which of the nonminimal edges is removed depends on which $u \in \{x, y, z\}$ is chosen first; any u will work, not just $u = x$. In the remaining two iterations of the WHILE loop, no further edges are removed. It is instructive to work out this example by hand; see Figure 8 for general pictures. \square

12 IDEALS TO GRAPHS: PROOF

The gruntwork in proving that the algorithm accomplishes its goal is contained in the following two technical lemmas, whose hypotheses are designed to be satisfied by the deformation taking place in one pass of the WHILE loop (after a cyclic permutation of (x, y, z) and an order-preserving bijection as in Lemma 2.1, perhaps).

LEMMA 12.1 *Suppose $I_{\mathcal{V}} \subset k[x, y, z]$ is artinian, and that $\nu \in \mathcal{V}$ has $\nu_x \neq 0$ and satisfies $\omega_y \geq \nu_y$ whenever $\omega_x = \nu_x$. Suppose further that $\epsilon = \{\epsilon^\omega\}_{\omega \in \mathcal{V}}$ is a deformation of \mathcal{V} with $\epsilon^\omega = \mathbf{0}$ for $\omega \neq \nu$ and $\epsilon^\nu = (1, 0, 0)$. Let $\gamma_y = \nu_y^\epsilon + |Y_{\nu^\epsilon}|$. If $\alpha \in \mathbb{N}^3$ then $K_\alpha(\mathcal{V}^\epsilon) = K_\alpha(\mathcal{V})$ unless*

$$\nu_z \leq \alpha_z \quad \text{and} \quad \nu_y \leq \alpha_y \leq \gamma_y \quad \text{and} \quad \alpha_x \in \{\nu_x, 1 + \nu_x\}. \tag{13}$$

If $\alpha_y \neq \gamma_y$ and α satisfies (13) with $\alpha_x = \nu_x$, then $K_{\alpha+\epsilon^\nu}(\mathcal{V}^\epsilon) = K_\alpha(\mathcal{V})$ while both $K_\alpha(\mathcal{V}^\epsilon)$ and $K_{\alpha+\epsilon^\nu}(\mathcal{V})$ have no reduced homology.

The last sentence takes care of the case where α satisfies (13) and $\alpha_x = 1 + \nu_x$, because $\alpha + \epsilon^\nu$ has x -coordinate $1 + \nu_x$; the case $\alpha_y = \gamma_y$ will be covered in Lemma 12.2.

The idea comes from Figure 8, where the grey dots represent elements of \mathcal{V}^ϵ , the white dots represent maximal points of $\mathcal{S}_{\mathcal{V}^\epsilon}$ (= second syzygies of $I_{\mathcal{V}^\epsilon}$ = irreducible components of $I_{\mathcal{V}^\epsilon}$), and the black dots represent first syzygies of $I_{\mathcal{V}^\epsilon}$. Looking from far down the x -axis, the vector ν^ϵ has a vertical plateau behind it: the big medium-grey wall, parallel to the yz -plane. Pushing ν^ϵ back to ν moves that vertical wall back a single unit. The only places where the topology of $K_\alpha(\mathcal{V})$ can possibly change are at lattice points α that sit either on the original wall in $\mathcal{S}_{\mathcal{V}^\epsilon}$ or its pushed-back image in $\mathcal{S}_{\mathcal{V}}$; these are the vectors α described in (13).

For vectors $\alpha + \epsilon^\nu$ that sit on the original wall but to the left of its right-hand edge (i.e. those with $\alpha_y \neq \gamma_y$), the Koszul simplicial complex $K_{\alpha+\epsilon^\nu}(\mathcal{V}^\epsilon)$ gets carried along for the ride to $K_\alpha(\mathcal{V})$ as the wall gets pushed back; the empty circles denote where the filled (black and white) dots get moved to. On the other hand, if α sits on the pushed-back image of the wall in $\mathcal{S}_{\mathcal{V}}$ (as the empty

circles strictly to the left of γ_y do), then $\mathcal{S}_{\mathcal{V}^\epsilon}$ is translation-invariant in the x -direction near α , making $K_\alpha(\mathcal{V}^\epsilon)$ a cone; the same goes for $K_{\alpha+\epsilon^\nu}(\mathcal{V})$. Cones have no reduced homology. Having this geometry in mind, here's the official proof of the lemma.

Proof. First assume α fails to satisfy (13). Use σ to denote an element of $\{0, 1\}^3$. To start off with, $\omega \preceq \alpha - \sigma$ if and only if $\omega^\epsilon \preceq \alpha - \sigma$ whenever $\omega \in \mathcal{V} \setminus \nu$, so the only possible differences in the simplicial complexes $K_\alpha(\mathcal{V})$ and $K_\alpha(\mathcal{V}^\epsilon)$ come from the placement of ν and ν^ϵ relative to the vectors $\alpha - \sigma$. If $\alpha_x \notin \{\nu_x, 1 + \nu_x\}$ then clearly $\nu \preceq \alpha - \sigma$ if and only if $\nu^\epsilon \preceq \alpha - \sigma$. And if $\alpha_y < \nu_y$ or $\alpha_z < \nu_z$, then neither ν nor ν^ϵ precedes $\alpha - \sigma$. The only remaining case of α not satisfying (13) has $\alpha_x \geq \nu_x$ and $\alpha_y > \gamma_y$ and $\alpha_z \geq \nu_z$. Suppose $\nu^\epsilon + Y_{\nu^\epsilon} = \nu^\epsilon \vee \omega$. Then $\nu^\epsilon \preceq \alpha \Rightarrow \omega \preceq \alpha$, and also $\nu \preceq \alpha \Rightarrow \omega \preceq \alpha$, so $K_\alpha(\mathcal{V})$ and $K_\alpha(\mathcal{V}^\epsilon)$ do not depend on ν or ν^ϵ .

Now assume α satisfies (13) and $\alpha_y < \gamma_y$. Then $\omega \preceq \alpha - \sigma$ if and only if $\omega \preceq \alpha - (\sigma \cup \epsilon^\nu)$ whenever $\omega \in \mathcal{V} \setminus \nu$ by the assumption ' $\omega_y \geq \nu_y$ whenever $\omega_x = \nu_x$ '. If $\alpha_x = \nu_x$ then $\nu^\epsilon \not\preceq \alpha$ and thus $K_\alpha(\mathcal{V}^\epsilon)$ is a cone with vertex ϵ^ν . If $\alpha_x = 1 + \nu_x$ then $\nu \preceq \alpha - \sigma$ if and only if $\nu \preceq \alpha - (\sigma \cup \epsilon^\nu)$, so $K_\alpha(\mathcal{V})$ is another cone with vertex ϵ^ν . Finally, suppose that $\alpha_x = \nu_x$ in addition to (13) and $\alpha_y < \gamma_y$. Then $\omega \preceq \alpha + \epsilon^\nu - \sigma$ if and only if $\omega \preceq \alpha - \sigma$ for $\omega \in \mathcal{V} \setminus \nu$ by the assumption ' $\omega_y \geq \nu_y$ whenever $\omega_x = \nu_x$ '. And clearly $\nu^\epsilon \preceq \alpha + \epsilon^\nu - \sigma$ if and only if $\nu \preceq \alpha - \sigma$, since $\nu^\epsilon = \nu + \epsilon^\nu$. Therefore $K_{\alpha+\epsilon^\nu}(\mathcal{V}^\epsilon) = K_\alpha(\mathcal{V})$ in this case. \square

LEMMA 12.2 *Assume the hypotheses and notation from Lemma 12.1, let $\gamma_x = \nu_x$, and let*

$$\gamma_z = \min_{\omega \neq \nu} \{\omega_z \mid \omega_x \leq \nu_x \text{ and } \omega_y < \gamma_y\},$$

which exists because $I_{\mathcal{V}}$ is artinian. Denote by \mathcal{F}^ϵ a minimal free resolution of $I_{\mathcal{V}^\epsilon}$ and by \mathcal{F} the specialized free resolution of $I_{\mathcal{V}}$ via [GPW00, Theorem 3.3]. (This amounts to the REDEFINE step applied to any cellular resolution supported on a complex with vertex set \mathcal{V}^ϵ .) Then at most two syzygies of \mathcal{F}^ϵ become nonminimal in \mathcal{F} : a second syzygy s_2^ϵ in degree $\gamma + \epsilon^\nu = (1 + \gamma_x, \gamma_y, \gamma_z)$ and a first syzygy s_1^ϵ in degree γ . Choose the unique $\rho \in \mathcal{V}$ with $\rho_y = \gamma_y$ such that $\rho_z < \gamma_z$ is maximal. Then the specializations (s_1, s_2) of $(s_1^\epsilon, s_2^\epsilon)$ are nonminimal if and only if $\rho_x = \gamma_x$.

Proof. The only possible nonminimal summands of \mathcal{F} occur in degrees $(\nu_x, \gamma_y, \alpha_z)$ or $(\nu_x^\epsilon, \gamma_y, \alpha_z)$ for some value of $\alpha_z \geq \nu_z$, because the other Betti numbers of $I_{\mathcal{V}}$ and $I_{\mathcal{V}^\epsilon}$ are in bijection by Lemma 12.1. Furthermore, nonminimal summands cannot come from zeroth syzygies of $I_{\mathcal{V}^\epsilon}$, since these are in bijection with those of $I_{\mathcal{V}}$ (no elements of \mathcal{V}^ϵ disappear when the ϵ is removed). Therefore, nonminimal syzygies in \mathcal{F} can only be first or second syzygies. It is a general fact about nonminimal free resolutions that nonminimal summands come in pairs (s_1, s_2) consisting of a first and second syzygy. In the

present case, such pairs arise from minimal first and second syzygies $(s_1^\epsilon, s_2^\epsilon)$ of $I_{\mathcal{V}^\epsilon}$. Since $\deg(s_1) = \deg(s_2)$ but $\deg(s_1^\epsilon) \neq \deg(s_2^\epsilon)$, and the only change is occurring in the x -direction, it must be that $\deg(s_2^\epsilon) = \epsilon^\nu + \deg(s_1^\epsilon)$. Lemma 12.1 therefore implies that any minimal syzygy s_2^ϵ becoming nonminimal in \mathcal{F} . must have $\deg(s_2^\epsilon)$ along the vertical ray $(\nu_x^\epsilon, \gamma_y, \alpha_z)$ for varying $\alpha_z \geq \nu_z$. Furthermore, there can be at most one value γ_z for α_z , since there can be only one second syzygy along any line parallel to an axis. This proves all but the last sentence.

The two specialized syzygies are nonminimal if and only if either one of them is. The specialization of s_2^ϵ is a second syzygy in degree γ that is minimal if and only if $K_\gamma(\mathcal{V})$ is the boundary of a triangle by Proposition 6.1. In any case, $\gamma - (1, 1, 1) \notin \langle \mathcal{V} \rangle$ by minimality of γ_z . Suppose $\rho_x \neq \gamma_x$. Then $(0, 1, 1) \in K_\gamma(\mathcal{V})$ because $\gamma_z > \nu_z$ and $\gamma_y > \nu_y$. Any vector $\omega \in \mathcal{V}$ whose z -coordinate was used to define γ_z has $\omega_x < \nu_x$ by the assumption ' $\omega_y \geq \nu_y$ whenever $\omega_x = \nu_x$ '; thus $(1, 1, 0) \in K_\gamma(\mathcal{V})$. And $(1, 0, 1) \in K_\gamma(\mathcal{V})$ because of ρ . Therefore, $K_\gamma(\mathcal{V})$ is the boundary of the triangle when $\rho_x \neq \gamma_x$ (in this case, there is no first syzygy of $I_{\mathcal{V}^\epsilon}$ in degree γ waiting to cancel s_2^ϵ as it specializes to s_2). Finally, if $\rho_x = \gamma_x$, then $(1, 0, 1) \notin K_\gamma(\mathcal{V})$, whence $K_\gamma(\mathcal{V})$ cannot be the boundary of a triangle. \square

EXAMPLE 12.3 Some possible combinatorial types for $K_\alpha(\mathcal{V})$, where $\alpha = \deg(s_1)$ is the degree of the specialized first syzygy of Lemma 12.2, are depicted in Figure 8. The headings 'ELSE' and 'THEN' correspond to the cases in Algorithm 11.2. Observe that in the single THEN case, the white dot s_2^ϵ at $(1 + \gamma_x, \gamma_y, \gamma_z)$ gets smashed into the vertical plane during specialization and cancels the black dot s_1^ϵ at $(\gamma_x, \gamma_y, \gamma_z)$. On the other hand, the topology remains constant in the first two ELSE cases. In the final ELSE case, two of the black dots merge to become a "double" black dot, since the resulting Koszul simplicial complex (3 disjoint vertices), has 2-dimensional \tilde{H}_0 after the wall is pushed back. \square

PROPOSITION 12.4 *At every iteration of the line END WHILE-DO in Algorithm 11.2, the labeled map M provides a minimal cellular free resolution of $I_{\mathcal{V}^\epsilon}$.*

Proof. This has two parts, of course: THEN and ELSE. Both follow from Lemma 12.2, given Remark 11.3(vii). Indeed, removing the unique nonminimal edge automatically destroys the unique nonminimal region by merging it with an adjacent region.

This argument implicitly uses Proposition 7.2, which guarantees that the deleted edge equals the entire intersection of the two regions containing it, so that matrices for the ordinary boundary complex of the deletion are obtained from those for M by removing the appropriate rows and columns. It should also be reiterated that we can choose the deformation in Lemma 12.1 to be the one occurring in each pass of the WHILE loop; indeed it is here that the precise condition (12) on the deformation ϵ in Algorithm 11.2 is used in an essential way. \square

Proof of Theorem 11.1. It remains only to reduce to the artinian case. Let $\overline{\mathcal{V}} = \mathcal{V} \cup \{\dot{u} \mid \mathcal{S}_{\mathcal{V}} \cap U = \emptyset\}$ for sufficiently large $|\dot{u}|$, as in Lemma 8.2. Given any labeled cell complex \overline{M} supporting a minimal cellular resolution of $I_{\overline{\mathcal{V}}}$, taking the subcomplex $M \subseteq \overline{M}$ whose labels precede the join $\vee \mathcal{V}$ produces a cell complex supporting a minimal resolution of M . This is the content of [BS98, Corollary 1.3]. \square

PART III

PLANAR MAPS REVISITED

13 ORTHOGONAL COLORING

Let M be a planar map with vertex set \mathcal{V} and M_{∞} an extended map. Since M_{∞} is embedded in a surface S homeomorphic to the plane, it makes sense to order the angles at any of its vertices cyclically, and to say that a list of angles at a vertex is **consecutive**. The same comment applies as well to the angles in any bounded region of M_{∞} , to the four angles having any fixed bounded edge as a leg (two at each vertex), and to the unbounded edges (read as the hands on an analog clock). With these cyclically ordered sets in mind, let A be any finite set of objects arranged cyclically in the surface S . Given three **colors** x, y, z , the set A is **trichromatic** if

- there is an element in A colored u , for each $u = x, y, z$;
- the elements in A colored u are consecutive, for each $u = x, y, z$; and
- the block of elements colored z is immediately counterclockwise from the block of elements colored y .

Deleting ‘counter’ from the last item defines **clockwise trichromatic** instead.

An **orthogonal coloring** \mathcal{O} of M_{∞} is a labeling of the angles in M_{∞} at every vertex in M by three **colors** x, y, z such that

- (i) all vertices of M are trichromatic in M_{∞} ;
- (ii) all bounded edges of M are clockwise trichromatic in M_{∞} ;
- (iii) all bounded regions of M are trichromatic in M_{∞} ;
- (iv) the two angles adjacent to each unbounded edge have different colors; and
- (v) attaching to each unbounded edge the color missing from its two angles makes the set of unbounded edges in M_{∞} trichromatic.

Suppose, in addition, that M has axial vertices $\dot{x}, \dot{y}, \dot{z}$, so $M_{\infty} = M_{\infty}(\dot{x}, \dot{y}, \dot{z})$. The angles interior to the three unbounded regions of M_{∞} are called **exterior angles** of M_{∞} ; besides the obvious pair of angles at each of $\dot{x}, \dot{y}, \dot{z}$, they include one angle at each nonaxial vertex lying on the exterior cycle. The **interior angles** at vertices on the exterior cycle are the angles lying inside M —that is, the non-exterior angles. Call an orthogonal coloring of M **axial** if the following boundary conditions, which imply axioms (iv) and (v), are satisfied for $u = x, y, z$:

- (iv)' all interior angles at the axial vertex \dot{u} are colored u ;
- (v)' all exterior angles of M_∞ in the unbounded region *not* touching \dot{u} are colored u .

For example, the orthogonal coloring in Figure 1 is axial. Roughly speaking, even axioms (iv)' and (v)' say that something is trichromatic: (iv)' says that the exterior cycle, thought of as a triangle with vertices $\dot{x}, \dot{y}, \dot{z}$, is trichromatic; while (v)' says that the set of exterior angles is trichromatic. This observation makes it particularly easy to remember the axioms defining an orthogonal coloring.

Felsner independently defined what he called **Schnyder colorings** in [Fel01, Section 1]. These are the *same thing* as axial orthogonal colorings. Felsner also pointed out that axiom (ii) follows from the others [Fel01, Lemma 1]. When the planar map is a triangulation of a simplex with no new vertices on the boundary, axial orthogonal coloring reduces to the angle-coloring of Schnyder [Tro92, Theorem 6.2.1], justifying Felsner's terminology. The next result and its proof justify the adjective "orthogonal".

PROPOSITION 13.1 *Any (axial) geodesic embedding $M \hookrightarrow \mathcal{S}_V$ induces an (axial) orthogonal coloring of $M_\infty = M \cup$ (unbounded orthogonal rays).*

Proof. Any angle in a geodesic embedding, whether between bounded or unbounded edges or both, locally lies in a plane $u = \text{constant}$ for some $u \in \{x, y, z\}$. Color such an angle by u . The orthogonality axioms follow readily from the definition of axial geodesic embedding and the region axiom (which holds for nonaxial grid surfaces by Lemma 8.2) in Section 2. \square

14 ORTHOGONAL FLOWS

In this section the planar map M has axial vertices $\dot{x}, \dot{y}, \dot{z}$, and $M_\infty(\dot{x}, \dot{y}, \dot{z})$ is orthogonally colored. We derive properties of orthogonal flows in grid surfaces (Section 2) from the axioms for axial orthogonal coloring, for comparison with [BT93]. In an earlier version of this paper, the current section was intended to serve as a possible proof method for Conjecture 16.3, below. Felsner in fact carried out this program [Fel02], having independently found the results in this section already [Fel01].

To begin, interpret an orthogonal coloring as a family of three **orthogonal vector fields** on M_∞ : for each vertex of M_∞ , assign precisely three arrows pointing away from it—one of each color—along the edges separating the blocks of differently colored angles. Thus, for example, the z -colored arrows point upward along an edge if and only if the angles around the edge have colors $\begin{smallmatrix} z \\ y|x, y|x, z|x \end{smallmatrix}$, or $y|x$ (the z -axis). The first of these edges has no arrow pointing downward from its top vertex, while the second and third have downward arrows colored x and y , respectively.

The u -colored vector fields for $u = x, y, z$ can be "integrated" to get **orthogonal flow lines**: the u -colored flow line from $\nu \in M$ is a directed path in M_∞ , beginning with ν , that is a union of edges underlying u -colored arrows. Thus

the next vertex after ν is at the other end of the edge whose u -colored arrow points away from ν .

Orthogonal flow lines can only meet in certain orientations. To make a precise statement, let L be a directed path passing through a vertex ν , and K a directed path containing an edge \vec{e} pointing toward ν . Then K **approaches L from the left** at ν if \vec{e} is distinct from L 's two arrows at ν , and these three arrows are oriented as in Eq. (14), below (ignoring the labels for the moment). Similarly, given the mirror orientations, K approaches L from the right.

LEMMA 14.1 *A flow line colored x never approaches a flow line colored z from the left. A flow line colored y never approaches a flow line colored z from the right. These statements remain true for cyclic permutations of x, y, z .*

Proof. Suppose \vec{e} approaches the z -colored flow line L from the left at ν . The angle coloration around ν looks like the following diagram,

$$\begin{array}{ccc}
 & & L \\
 & & \uparrow \\
 & y & | \ x \\
 \overset{c}{\leftarrow} \vec{e} & \rightarrow & \nu \\
 \underset{b}{\leftarrow} & & \uparrow \\
 & a & | \ x \\
 & & y
 \end{array} \tag{14}$$

in which $a \neq x$ by edge trichromatics. (There may be other edges containing ν but not shown.) Since none of the angles going clockwise between a and y at ν can be x -colored (by vertex trichromatics), it is impossible to have $\{b, c\} \subseteq \{y, z\}$, by edge trichromatics. The other case is similar, and the symmetry is obvious. \square

The observation in Lemma 14.1 imposes useful conditions on flow lines. The next proposition says that orthogonal flow lines satisfy conditions slightly stronger than the five ‘‘path properties’’ defining a **normal family of paths**, as introduced by Brightwell and Trotter [BT93] (see [Tro92, Chapter 6] for an exposition). A set of paths whose pairwise intersections consist of a single vertex is called **independent**.

PROPOSITION 14.2 (PATH PROPERTIES) *Endow $M_\infty(\dot{x}, \dot{y}, \dot{z})$ with an axial orthogonal coloring. Suppose $\nu \in M$ is a vertex and $u \in \{x, y, z\}$ is a color.*

1. *There is a unique u -colored flow line beginning at ν ; it connects ν to ∞ via \dot{u} .*
2. *If $[\nu, \dot{u}]$ denotes the part of the u -colored flow line starting with ν and ending with \dot{u} , then $[\nu, \dot{x}]$, $[\nu, \dot{y}]$, and $[\nu, \dot{z}]$ are independent paths from ν .*
3. *The six paths $[\dot{x}, \dot{y}]$, $[\dot{y}, \dot{x}]$, $[\dot{y}, \dot{z}]$, $[\dot{z}, \dot{y}]$, $[\dot{x}, \dot{z}]$, $[\dot{z}, \dot{x}]$ are on the exterior cycle of M .*
4. *If $\omega \in M$ is a vertex such that $\omega \in [\nu, \dot{u}]$, then $[\omega, \dot{u}] \subseteq [\nu, \dot{u}]$.*

5. If ω is in the union $[\dot{x}, \nu, \dot{y}]$ of $[\nu, \dot{x}]$, $[\nu, \dot{y}]$, $[\dot{x}, \dot{y}]$, and any regions they enclose, then $[\dot{x}, \omega, \dot{y}] \subseteq [\dot{x}, \nu, \dot{y}]$; the same holds with \dot{z} in place of \dot{x} or \dot{y} .

Proof. Existence and uniqueness in part 1 are obvious. Since the only arrows in M_∞ pointing out of M are on the axes by orthogonality axioms (iv)' and (v)', part 1 is equivalent to statement that flow lines contain no cycles. Any flow cycle C containing a vertex in the interior of the region it bounds contains an entire flow cycle of some other color in its interior: one of the other colors cannot escape by Lemma 14.1. Thus we may assume C has no vertices interior to it. But then the argument in the proof of Lemma 14.1 forces the coloration of the angles in the interior of C to omit one color entirely, violating region trichromatics.

Now suppose two flow lines from ν -colored x and y , say—intersect, and consider the cycle C formed by their arcs connecting ν to their first intersection point. Assume for the moment that ω lies interior to C . Depending on the orientations of the two flow lines around C , either z -flow lines cannot escape C , or the u -flow line from ω exits through the u -colored arc of C , for $u = x, y$. The first case contradicts acyclicity, while the second case produces a smaller cycle C . Again, we may therefore assume C contains no vertices interior to it. In the first orientation, no interior angle of C is colored z , while in the second orientation, all interior angles of C are colored z .

Part 3 follows easily by applying orthogonality axiom (ii) to the edges on the exterior cycle of M , each of which has two of its four colors specified by (v)'.

Part 4 is obvious from the definition of flow line.

To prove part 5, first observe that a flow line colored x or y cannot escape $[\dot{x}, \nu, \dot{y}]$ if it originates at a vertex $\omega \in [\dot{x}, \nu, \dot{y}]$ that is on neither $[\nu, \dot{x}]$ nor $[\nu, \dot{y}]$. Indeed, if $[\omega, \dot{x}]$ intersects $[\nu, \dot{x}]$, then these two flow lines agree thereafter; and $[\omega, \dot{x}]$ cannot even approach $[\nu, \dot{y}]$, thanks to Lemma 14.1. On the other hand, if $\nu \neq \omega \in [\nu, \dot{y}]$, say, then $[\omega, \dot{y}] \subset [\nu, \dot{y}]$ by part 4. Moreover, vertex trichromatics force the first edge of $[\omega, \dot{x}]$ to exit ω clockwise from the y -colored arrow pointing away ω , but counterclockwise from the other edge leaving ω and in $[\nu, \dot{y}]$. Therefore, $[\omega, \dot{x}]$ remains inside $[\dot{x}, \nu, \dot{y}]$ either by part 4 or the first sentence of this paragraph. \square

EXAMPLE 14.3 Orthogonal flow lines are better behaved than arbitrary normal families of paths, since their strong local properties imply Path Property 2, and especially the crucial Path Property 5, which are global. Figure 9 depicts two triples of vector fields determined by trichromatic angle colorings whose corresponding flow lines satisfy the conclusion of Proposition 14.2, and therefore constitute normal families of paths. The vector fields in the left diagram are not orthogonal because of the edge connecting the rightmost interior vertices, the edge connecting the remaining interior vertex to \dot{x} , the interior region, and the bottom region bordering $[\dot{x}, \dot{z}]$. Recoloring the angles at the leftmost interior vertices yields the orthogonal vector fields at right. \square




Figure 9: Nonorthogonal and orthogonal vector fields

COROLLARY 14.4 *Let M be a planar with axial vertices $\hat{x}, \hat{y}, \hat{z}$. The extended map $M_\infty(\hat{x}, \hat{y}, \hat{z})$ can be orthogonally colored if and only if it is triconnected.*

Proof. Proposition 14.2.2 proves the ‘only if’ direction, while Theorem 5.1 and Proposition 13.1 prove the ‘if’ direction. \square

REMARK 14.5 The results of this section can be massaged to work for nonaxial orthogonally colored extended maps M_∞ , as well. In fact, everything reduces to the axial case: draw a large triangle containing all of M —not M_∞ —in its interior, and call its three vertices $\hat{x}, \hat{y}, \hat{z}$, in counterclockwise order. Then connect each u -colored unbounded edge to \hat{u} . It is straightforward to verify the axioms for axial orthogonal colorings, given the ordinary axioms for M_∞ . Observe the analogy with Lemma 8.2.

15 DUALITY FOR GEODESIC EMBEDDINGS

Let G be a graph embedded in the sphere $S \cup \{\infty\}$, where $S \cong \mathbb{R}^2$, and assume that $G \cap S$ is a planar map or extended map. For this paragraph only, we allow graphs and planar maps to have multiple edges, although we assume G has no loops, and that the edges of G in each one of its regions form a simple cycle. Define the **spherical dual** \hat{G} of G as usual: place a vertex \hat{A} in each region A of G , and draw an edge connecting \hat{A} to \hat{B} through each edge contained in $A \cap B$. Then \hat{G} also satisfies the no-loop and simple-cycle conditions, which are dual to each other. Assume that ∞ is a vertex of either G or \hat{G} . When $\infty \notin G$, so G is a planar map $M \subset S$, then \hat{G} is an extended map that we denote by \hat{M}_∞ (with associated planar map $\text{del}(\hat{G}; \infty) = \hat{M}$) and call the **planar dual** of M . When $\infty \in G$, so $G \cap S = N_\infty$ is an extended map satisfying $N = \text{del}(G; \infty)$, then \hat{G} is a planar map that we denote by \hat{N} and call the **planar dual** of N_∞ . For an arbitrary grid surface $\mathcal{S}_\mathcal{V}$, let \hat{a} be any vector preceded by $\mathbf{1}$ + the join of the vectors in \mathcal{V} , where $\mathbf{1} = (1, 1, 1)$. Define $\bar{\mathcal{V}}$ by throwing in axial vertices missing from \mathcal{V} :

$$\bar{\mathcal{V}} = \mathcal{V} \cup \{\hat{u} \mid \mathcal{S}_\mathcal{V} \cap U = \emptyset\}, \quad \text{where } \hat{u} \in U \text{ has length } |\hat{u}| = \hat{a}_u.$$

This notation agrees with that in Lemma 8.2, but specifies the length of \hat{u} . Define the **Alexander dual** grid surface $\mathcal{S}_{\hat{\mathcal{V}}}$ by

$$\hat{\mathcal{V}} = \{\hat{\alpha} - \rho \mid \rho \in \mathcal{S}_{\overline{\mathcal{V}}}^{\max}\},$$

where the notation comes from Lemma 8.3. Neither $\mathcal{S}_{\hat{\mathcal{V}}}$ nor $\mathcal{S}_{\overline{\mathcal{V}}}$ depends combinatorially on $\hat{\alpha}$, in the sense of Lemma 2.1.

By definition, $\overline{\mathcal{V}} = \mathcal{V}$ and $\hat{\alpha} \succeq \hat{x} + \hat{y} + \hat{z} + \mathbf{1}$ when $\mathcal{S}_{\mathcal{V}}$ is axial. Define $\mathcal{S}_{\mathcal{V}}$ to be **radial** if $\nu_u \neq 0$ for all $\nu \in \mathcal{V}$ and $u \in \{x, y, z\}$. In particular, $\overline{\mathcal{V}} = \mathcal{V} \cup \{\hat{x}, \hat{y}, \hat{z}\}$ is a disjoint union when $\mathcal{S}_{\mathcal{V}}$ is radial. Although the following duality theorem is stated only for axial and radial grid surfaces, similar (but slightly harder to state) considerations apply to arbitrary geodesic and rigid embeddings; the definition of Alexander dual grid surface extends verbatim. Recall that $M_\infty = M \cup$ (unbounded orthogonal rays in $\mathcal{S}_{\mathcal{V}}$) for geodesic embeddings $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$.

THEOREM 15.1 *Let $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ be an axial geodesic embedding, and $N \hookrightarrow \mathcal{S}_{\mathcal{W}}$ a radial geodesic embedding.*

1. $M \cong \hat{N}$ if and only if $\hat{M} \cong N$.
2. There is a natural radial geodesic embedding $\hat{M} \hookrightarrow \mathcal{S}_{\hat{\mathcal{V}}}$ that is rigid if and only if $M \hookrightarrow \mathcal{S}_{\mathcal{V}}$ is rigid.
3. There is a natural axial geodesic embedding $\hat{N} \hookrightarrow \mathcal{S}_{\hat{\mathcal{W}}}$ that is rigid if and only if $N \hookrightarrow \mathcal{S}_{\mathcal{W}}$ is rigid.

Proof. Part 1 is simply duality for spherical maps, as in the first paragraph of this section. We prove part 2, since part 3 is similar and can even be simplified by using parts 1 and 2. The vertex axiom for $\mathcal{S}_{\hat{\mathcal{V}}}$ follows immediately from Lemma 8.3. The main point for the rest of the proof is that

$$\{\sigma \in \mathcal{S}_{\hat{\mathcal{V}}} \mid \sigma \preceq \hat{\alpha}\} = \{\hat{\alpha} - \sigma \mid \sigma \in \mathcal{S}_{\mathcal{V}} \text{ and } \sigma \preceq \hat{\alpha}\}.$$

In other words, $\{\sigma \in \mathcal{S}_{\mathcal{V}} \mid \sigma \preceq \hat{\alpha}\}$ lies in the topological boundary of $\mathbb{R}^3 \setminus \langle \mathcal{V} \rangle$. Every elbow geodesic $[\nu, \omega]$ in M borders precisely two regions of M_∞ because M_∞ is triconnected (apply Proposition 7.2 along with Theorem 8.4, or Proposition 13.1 along with Corollary 14.4). If these two regions A and B are bounded, then the maximal points \hat{A} and \hat{B} in them (Lemma 8.3) connect via straight line segments in $\mathcal{S}_{\mathcal{V}}$ to $\nu \vee \omega$ by the region axiom (see the proof of Lemma 8.3). Furthermore, one of these segments must transform into an orthogonal ray in $\mathcal{S}_{\hat{\mathcal{V}}}$ via $\sigma \mapsto \hat{\alpha} - \sigma$, because $\nu \vee \omega = \hat{A} \wedge \hat{B}$ (by the region axiom: for each $u \in \{x, y, z\}$ one of the two vectors ν and ω shares its u -coordinate with one of the vectors \hat{A} and \hat{B}). Thus every bounded edge of \hat{M}_∞ is an elbow geodesic in $\mathcal{S}_{\hat{\mathcal{V}}}$.

If $[\nu, \omega]$ borders a bounded region A and an unbounded region, then $\nu_u = \omega_u = 0$ for some $u \in \{x, y, z\}$. The vector $\nu \vee \omega$ shares both of its nonzero coordinates with \hat{A} in this case, so \hat{A} is the unique maximal point of $\mathcal{S}_{\mathcal{V}}$ preceded by $\nu \vee \omega$. This forces the negative ray $\hat{A} - U$ passing through \hat{A} and $\nu \vee \omega$ to transform




Figure 10: Duality for geodesic embeddings

into an unbounded orthogonal ray $\hat{\alpha} - (\hat{A} - U) = (\hat{\alpha} - \hat{A}) + U$ in $\mathcal{S}_{\hat{\gamma}}$. Thus the unbounded edges of \hat{M}_{∞} are unbounded orthogonal rays in $\mathcal{S}_{\hat{\gamma}}$.

Now we show that every orthogonal ray in $\mathcal{S}_{\hat{\gamma}}$ is contained in an edge of \hat{M} . Equivalently, for each maximal point $\hat{A} \in \mathcal{S}_{\hat{\gamma}}$ and $U \in \{X, Y, Z\}$, there is an elbow geodesic $[\nu, \omega] \in M$ such that $\hat{A} - \nu \vee \omega \in U$. By symmetry, set $U = Z$. Choose the edge $[\nu, \omega] \subset A$ so that $\nu_x = \hat{A}_x$ and $\omega_y = \hat{A}_y$; such an edge exists by orthogonality property (iii) and Proposition 13.1.

For the statement about rigidity, use the rigid region axiom: $\nu \vee \omega \preceq \hat{A} \Leftrightarrow \nu \preceq \hat{A}$ and $\omega \preceq \hat{A}$, and this occurs precisely when $\nu \in A$ and $\omega \in A$. Thus, when $[\nu, \omega] \subset \mathcal{S}_{\hat{\gamma}}$ is a rigid geodesic, $\nu \vee \omega$ precedes the maximal points in exactly two regions of M : the regions A and B containing both ν and ω . That $\nu \vee \omega = \hat{A} \wedge \hat{B}$ was shown above. \square

EXAMPLE 15.2 Figure 10 illustrates Theorem 15.1 for a particular (nonrigid) geodesic embedding. Turning the picture upside-down yields two pictures of the Alexander dual grid surface, with the radial embedding appearing the right way out and the axial embedding backwards. \square

Considering the Theorem 8.4, the duality result here essentially falls under the umbrella of duality for resolutions of monomial ideals [Mil00a, Section 4.2]. Although the above proof can be simplified greatly by applying duality for resolutions, particularly in concert with [Mil00a, Proposition 3.20], it seemed appropriate to keep Theorem 15.1 as self-contained as possible.

16 OPEN PROBLEMS

16.1 PLANAR MAPS SUPPORTING RESOLUTIONS OF NONARTINIAN IDEALS

PROBLEM 16.1 *Characterize those planar maps that support minimal resolutions of trivariate monomial ideals.*

It would be nice to clean up the statement of Theorem B by proving $3 \Rightarrow 2$. Unfortunately, the most direct proof attempt fails.

To be precise, suppose the planar map N supports a minimal cellular free resolution of some ideal $I_{\mathcal{V}}$. Following the procedure in the proof of Proposition 7.2, the underlying graph H of N has paths analogous to orthogonal flows. However, since \mathcal{V} may not be axial, these flows do not reach axial vertices, but instead reach unbounded orthogonal rays. As an abstract graph, we can define \overline{H} by adding three new vertices $\hat{x}, \hat{y}, \hat{z}$ to H , and then connecting each vertex $\nu \in N$ to \hat{u} if the orthogonal ray U_{ν} is unbounded. Judging from Theorems A and B, one might expect that \overline{H} is planar and triconnected, supporting a minimal free resolution of some artinian approximation to $I_{\mathcal{V}}$. But although \overline{H} is obviously triconnected, it need not be planar!

Take, for instance, the set \mathcal{V} from Example 9.2, without the axial vectors. Deleting $\hat{x}, \hat{y}, \hat{z}$ from the nonplanar map M there yields a planar map N supporting a minimal free resolution of $I_{\mathcal{V}}$, but reconnecting to $\hat{x}, \hat{y}, \hat{z}$ as above returns M again.

16.2 NONRIGID GEODESIC EMBEDDINGS

CONJECTURE 16.2 *Let the axial grid surface $S_{\mathcal{V}}$ have vertices $\nu_1, \nu_2, \nu_3, \nu_4 \in \mathcal{V}$, each with no coordinate zero, such that $[\nu_i, \nu_j]$ is an elbow geodesic whenever $i \neq j$. Then the ideal $I_{\mathcal{V}}$ possesses a nonplanar cellular minimal free resolution.*

Thus relatively minor violation of rigidity for $S_{\mathcal{V}}$ not only implies nonuniqueness of minimal cellular resolutions (cf. Corollary 9.1), but should even imply nonplanarity. Intuitively, it should be possible to construct K_5 out of elbow geodesics (two of which cross) using two of the ν_i and $\hat{x}, \hat{y}, \hat{z}$, as in Example 9.2. Given such a configuration, one is tempted to “fill in” the resulting 1-skeleton to form a cell complex minimally resolving $I_{\mathcal{V}}$. This is, in fact, how Example 9.2 was constructed.

16.3 ORTHOGONAL COLORING TO RIGID EMBEDDING

CONJECTURE 16.3 (FELSNER’S THEOREM [FEL02]) *Every orthogonal coloring on a planar map is induced by a rigid embedding, as in Proposition 13.1.*

This converse to Proposition 13.1 reduces easily to the axial case, via Lemma 8.2 and Remark 14.5. Since the axial case was proved by Felsner [Fel02] in response to seeing the conjecture here, this is in fact no longer an open problem. But see [Fel02] for Felsner’s open questions regarding the set of orthogonal colorings. Besides its applications to the questions discussed in the next subsection, the motivation behind formulating the above statement was that it reduces Theorem 5.1 to verifying that every planar map there has an orthogonal coloring. This actually follows by the same induction used for rigid embeddings, but

the details become much simpler; moreover, Felsner already proved existence of orthogonal colorings in [Fel01]. Thus we get a substantially more palatable proof of Theorem 5.1 via [Fel01, Fel02].

16.4 THE SCARF STRATIFICATION

Part of the motivation for the results presented here stems from the desire to understand not just how to assign a minimal free resolution to any particular monomial ideal, but to understand the collection of all minimal free or injective resolutions of monomial ideals. Some recent results, notably those in [GPW00], aim to classify monomial ideals according to whether their minimal resolutions are in some sense isomorphic, by ascertaining what data determines the minimal resolutions. Other results, such as those in [MSY00], raise the question of which deformations (as in Section 10) preserve minimal resolutions. This latter idea originated from [BPS98, PS98], which was in turn based upon work of H. Scarf on related classifications for integer programs.

It seems that the most universal approach for minimal resolutions of monomial ideals should combine the two types of classification above. Heuristically, the question becomes, ‘What are all possible ways of deforming continuously between monomial ideals having “nearby” isomorphism classes of resolutions?’ Ideally one would like a ‘fine moduli space’ for minimal resolutions, in the sense that arcs in that space correspond to families of monomial ideals whose minimal resolutions deform continuously. Of course, there are only finitely many deformation classes of minimal resolutions, so the space should interact well with the poset of monomial ideals under deformation.

One tempting candidate for ‘fine moduli space’ can be defined as follows. Giving r monomials in n variables is the same as giving an $r \times n$ matrix of exponents. The **generic** monomial ideals partition an open dense subset of the nonnegative orthant in this matrix space, by [BPS98, MSY00]. Taking the cells formed by intersections of the closures of the generic loci yields a decomposition that we propose to call the **Scarf stratification**. It should be a rational polyhedral fan, if life is fair; but at least there should be a subdividing fan with finitely many maximal cones such that the maximal Scarf cells are unions of maximal cones. Any classifying space such as the Scarf stratification will feel rather more like an algebraic stack than a fine moduli space, because even if it classifies deformations of minimal resolutions, the actual *set* of isomorphism classes of minimal resolutions would be a quotient by some finite group (containing the symmetric group on the variables, at least) of the set of strata: two monomial ideals differing by a permutation of the variables might be far from each other in the stratification.

Whatever the correct space of minimal resolutions ends up being, the methods introduced here can be applied to elucidate its combinatorial structure, for $n = 3$. Note that Felsner’s theorem (Conjecture 16.3) classifies the maximal strata for the case of three variables—that is, the generic trivariate monomial

ideals—by [MSY00]:

COROLLARY 16.4 *Deformation classes of generic artinian trivariate monomial ideals correspond bijectively to orthogonally colored axial planar triangulations.*

The triangulations referred to here are orthogonally colored triangulations of the simplex *with new vertices on the boundary*, so they are not Schnyder normal colorings. Corollary 16.4 and Corollary 9.1 prompt the following:

QUESTION 16.5 *In terms of monomial ideals, what do deformation classes of axial rigid embeddings correspond to in general, for non-triangulations?*

For instance, do they correspond to certain Scarf strata? If the Scarf stratification is a rational polyhedral fan, what linear equations define these cones?

We note that arbitrary geodesic embeddings have considerably more freedom than do rigid embeddings, from the point of view of deformations (hence the adjective ‘rigid’). The possible application of the material in this paper to the classification of deformations of minimal resolutions was one of the motivations for stating as many results as possible in the context of arbitrary geodesic embeddings.

Our final remark concerns the bias in this paper toward artinian monomial ideals. Combinatorial considerations such as triconnectivity notwithstanding, the bias also makes sense algebraically. Briefly, the homological characterization of genericity for arbitrary monomial ideals [MSY00, Theorem 1.5 and Remark 1.7] is a statement about graded injective resolutions under deformation; and any result concerning \mathbb{Z}^3 -graded injective resolutions of arbitrary monomial ideals has an equivalent statement in terms of \mathbb{Z}^3 -graded free resolutions of artinian monomial ideals, by the duality results in [Mil00a]. We emphasize that the theory surrounding injective resolutions played a crucial role in properly formulating the graph-theoretic results in Part I, as well as the algebraic results in Part II (the exposition in Sections 10–12 is based on [Mil00b, Theorem 5.60]). Moreover, concentrating on injective resolutions (equivalently, artinian monomial ideals) should ease the nonplanarity difficulties raised in Section 16.1, by applying [Mil00a, Theorem 4.5.5 and Example 4.8.5], which says how to recover free resolutions from injective resolutions.

REFERENCES

- [BPS98] Dave Bayer, Irena Peeva, and Bernd Sturmfels, *Monomial resolutions*, Math. Res. Lett. 5 (1998), no. 1-2, 31–46.
- [BS98] Dave Bayer and Bernd Sturmfels, *Cellular resolutions of monomial modules*, J. Reine Angew. Math. 502 (1998), 123–140.
- [BT93] Graham Brightwell and William T. Trotter, *The order dimension of convex polytopes*, SIAM J. Discrete Math. 6 (1993), no. 2, 230–245.

- [Fel01] Stefan Felsner, *Convex drawings of planar graphs and the order dimension of 3-polytopes*, Order 18 (2001), no. 1, 19–37.
- [Fel02] Stefan Felsner, *Geodesic embeddings and planar graphs*, Preprint, 2002.
- [GPW00] Vesselin Gasharov, Irena Peeva, and Volkmar Welker, *The LCM-lattice in monomial resolutions*, Math. Res. Lett. 6 (1999), no. 5–6, 521–532.
- [Hoc77] Melvin Hochster, *Cohen-Macaulay rings, combinatorics, and simplicial complexes*, Ring theory, II (Proc. Second Conf., Univ. Oklahoma, Norman, Okla., 1975) (B. R. McDonald and R. Morris, eds.), Lect. Notes in Pure and Appl. Math., no. 26, Dekker, New York, 1977, pp. 171–223. Lecture Notes in Pure and Appl. Math., Vol. 26.
- [Mil00a] Ezra Miller, *The Alexander duality functors and local duality with monomial support*, J. Algebra 231 (2000), 180–234.
- [Mil00b] Ezra Miller, *Resolutions and duality for monomial ideals*, Ph.D. thesis, University of California at Berkeley, 2000.
- [MS99] Ezra Miller and Bernd Sturmfels, *Monomial ideals and planar graphs*, Applied Algebra, Algebraic Algorithms and Error-Correcting Codes (M. Fossorier, H. Imai, S. Lin, and A. Poli, eds.), Springer Lecture Notes in Computer Science, no. 1719, Springer Verlag, 1999, Proceedings of AAECC-13 (Honolulu, Nov. 1999), pp. 19–28.
- [MSY00] Ezra Miller, Bernd Sturmfels, and Kohji Yanagawa, *Generic and cogeneric monomial ideals*, J. Symbolic Comp. 29 (2000), 691–708.
- [PS98] Irena Peeva and Bernd Sturmfels, *Generic lattice ideals*, J. Amer. Math. Soc. 11 (1998), no. 2, 363–373.
- [Roz70] I. Z. Rozenknop, *Polynomial ideals that are generated by monomials*, Moskov. Oblast. Ped. Inst. Učen. Zap. 282 (1970), 151–159, (in Russian).
- [RW01] Victor Reiner and Volkmar Welker, *On the linear syzygies of a Stanley–Reisner ideal*, Math. Scand. 89 (2001), no. 1, 117–132.
- [Tro92] William T. Trotter, *Combinatorics and partially ordered sets: Dimension theory*, Johns Hopkins University Press, Baltimore, MD, 1992.
- [Zie95] Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, New York, 1995.

Ezra Miller
Department of Mathematics
M.I.T.
Cambridge, Massachusetts
ezra@math.mit.edu