EQUILIBRIUM GRAPHS ON THE FLAT TORUS
OR
FINDING ZEN AMIDST THE BULL

BY

PATRICK LIN

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science
in the Graduate College of the
University of Illinois Urbana-Champaign, 2021

Urbana, Illinois

Doctoral Committee:

Professor Jeff Erickson, Chair
Professor Timothy Chan
Professor Chandra Chekuri
Professor Anna Lubiw

Abstract

Tutte's classical spring embedding theorem, and equilibrium graphs on the plane in general, have been a subject of study for many decades, with connections to and applications in many areas, including, but not limited to, discrete geometry, planar graph theory, graphics, surface parametrization, mechanical engineering, and graphical statics. In his 1963 paper, Tutte observed,

> *"[W]e may remark that very little is known about representations of graphs in the protective plane and higher surfaces."*

Six decades since Tutte's observation, however, our understanding of the properties and applications of equilibrium graphs on higher genus surfaces is still surprisingly limited, despite a growing body of work suggesting the utility of furthering this understanding.

Figure 1. A perfectly toroidal taurus.

In this thesis, we extend some existing structural properties and algorithmic applications of equilibrium graphs on the plane to the setting of flat tori. In particular, we consider the classical Maxwell–Cremona correspondence and a number of different planar morphing algorithms.

The Maxwell–Cremona correspondence, through a more modern computational geometry lens, can be summarized as stating that a planar graph is in positive equilibrium if and only if it is a weighted Delaunay graph of its point set. We derive some partial generalizations of this correspondence in the toroidal setting. In particular, we show that, whereas weighted Delaunay still implies positive equilibrium on flat tori, the converse is not always true; however, we give a full characterization of when the converse holds.

Next, we present generalizations of a few different techniques for morphing planar graphs. We show that techniques by Cairns and by Floater and Gotsman generalize to the toroidal setting with minor modifications; our generalization of the latter also provides a short proof of a conjecture of Connelly *et al.* for geodesic torus triangulations. On the other hand, Alamdari *et al.*'s improvement of Cairns' method uses techniques that do not seem to generalize. Instead, we obtain a similar improvement via a novel technique using toroidal spring embeddings, and then adapt said new technique to derive a new, simpler planar morphing algorithm.

To myself,
with relief,
for getting it done.

Acknowledgments

I would first like to thank my PhD advisor, Jeff Erickson. Jeff provided me with unending support, excellent guidance, honest advice, patience in the face of my shenanigans, an infinite well of historical anecdotes, and far too many dank memes. It was an honor.

Next, I would like to thank Erin W. Chambers, Chandra Chekuri, G. Carl Evans, Margaret Fleck, Sariel Har-Peled, Matus Telgarsky, and Mahesh Viswanathan for their wisdom and excellent advice about research, teaching, life, and finding humor in new places.

Special thanks to Timothy Chan, for making me aware of a preliminary version of Alamdari *et al.* [4], and the lack of a known extension to the toroidal setting: for me, it was the beginning of my dissertation work; for him, it was Monday. Thanks also to Anna Lubiw, for asking questions about Section 6.2 whose answers led to the results in Section 6.3.

I am of the firm belief that one's experience as a PhD student is shaped primarily by one's peers, so many thanks to my fellow students who passed through the theory group at UIUC during my time here. In particular, thanks to Robert Andrews, Tanvi Bajpai, Shant Boodaghians, Charlie Carlson, Hsien-Chih Chang, Spencer Gordon, Shalmoli Gupta, Mitchell Jones, Zander Kelley, Konstantinos Koiliaris, Pooja Kulkarni, Rucha Kulkarni, Vasilis Livanos, Ian Ludden, Vivek Madan, Sahand Mozaffari, Shweta Patwa, Kent Quanrud, Meghan Shanks, Alex Steiger, Chao Xu, and Xilin Yu, for intellectual stimulation, distraction, and helping me stay sane.

I would not have gotten here without the influence of the faculty at NYU(-Poly): particularly impactful were my Master's advisor Lisa Hellerstein, as well as Boris Aronov, Sylvain Cappell, John Iacono, Linda Sellie, John Sterling, and Yisong Yang. I am, of course, also indebted to the students and postdocs of the theory group at NYU-Poly during my time there: Sarah Allen, Pooya Davoodi, Devorah Kletenik, Özgur Özkan, Anirudh Reddy, and Mark Yagnatinsky.

Many, many thanks to my parents for their support and patience during my PhD, and to my sister Estella, for her artistic interpretations of my research (in particular, Figure 1).

And finally, I cannot begin to describe how grateful I am to W. Linda Xu, who was always there to provide me with intellectual, emotional, and moral support as we distracted each other from our research and other work.

Table of Contents

Chapter 1

Introduction

In his seminal 1963 paper, Tutte [213] described a method for producing visually appealing drawings of planar graphs. Specifically, his so-called spring embedding theorem states that, given a 3-connected planar graph G with a fixed convex polygon for the outer face, one can compute a *straight-line embedding* of G on the plane, wherein every edge is drawn as a straight-line segment between its endpoints and edges are not allowed to cross, in such a way that every vertex is positioned at the center of mass (or barycenter) of its neighbors and every face is convex.

Intuitively, Tutte's spring embedding theorem is based on finding *equilibrium* positions for each vertex: Imagine that each edge is a rubber band or spring that pulls inward on its endpoints with some amount of force related to some positive *stress* value on the edge. After nailing down the vertices on the outer face, we can "let go" of all the other vertices. Each edge will try to pull its endpoints together, but each vertex is pulled in many directions by its incident edges. Eventually, after maybe bouncing around for a bit, the system will reach an equilibrium, in which the forces on each unfixed vertex sum to 0 and there is no more movement. Figure 1.1 shows spring embeddings corresponding to three different sets of edge stresses.

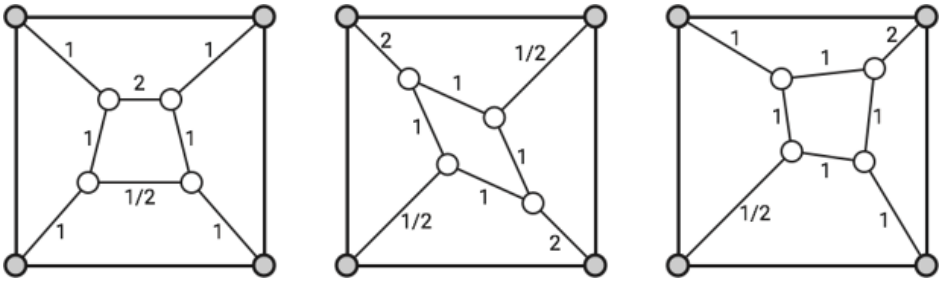

Figure 1.1. Tutte spring embeddings corresponding to three different sets of edge stresses.

Moreover, this equilibrium can be defined in terms of a linear system whose variables are the positions of the vertices. Intuitively, the system encodes the fact that at equilibrium, the sum of the forces around each vertex sums to 0; see Section 2.11.1 for a more detailed description of this system. It is well-known that this linear system can be solved in $O(n^{\omega/2}) = O(n^{1.18643})$ time, where $\omega < 2.37286$ is the matrix multiplication exponent [7, 143]. Floater later generalized this result to allow asymmetric stresses [96, 97]; intuitively, one can imagine that in this setting, edges pull on their endpoints with different amounts of force.

Tutte's spring embedding theorem turned out to be extremely important, arguably launch-

ing the research area of graph drawing, and has close relations to and applications in many areas. For example, Tutte's spring embedding theorem is closely related to many fundamental results in planar graph theory such as the Fáry–Stein–Wagner Theorem [93, 200, 216], Steinitz' theorem [198, 199], and the Koebe–Andreev–Thurston circle packing theorem [10, 11, 41, 58, 61, 95, 135, 162–164]. Tutte embeddings and their generalizations have also found extensive applications in other areas of Computer Science, including graphics [1–3], surface and mesh parametrization [76, 96, 109, 110, 114, 197, 218], and morphing [57, 99, 111, 130, 205–207]. The equilibrium systems considered by Tutte are also closely related to the classical Maxwell–Cremona correspondence [18, 68–70, 121, 156–158, 220], which has its roots in mechanical engineering; moreover, the Maxwell–Cremona correspondence can be extended [18, 221] to relate Tutte's spring embeddings to weighted Delaunay graphs, which have countless applications in computational geometry [19, 20, 27, 28, 33, 35, 42, 46, 77, 83, 84, 112, 123, 160, 194–196, 215, 222].

1.1 Outline

Given the extensive literature on the properties and applications of equilibrium graphs, it is natural to ask what aspects generalize to graphs drawn on other surfaces, and how they generalize. In this thesis, we specifically consider graphs drawn on flat tori: we investigate generalizations of the aforementioned Maxwell–Cremona(–Delaunay) correspondence and various planar morphing algorithms.

The Maxwell–Cremona correspondence relates equilibrium stresses on a straight-line drawing of a planar graph to orthogonal straight-line drawings of the planar dual, as well as dual polyhedral lifts of both drawings. In particular, when the drawing admits a *positive* equilibrium stress, then it is in fact a Tutte spring embedding, and the induced polyhedral lift is convex. In this setting, we can interpret the primal and dual *embeddings* as being dual weighted Delaunay graphs and weighted Voronoi diagrams, respectively.

Morphing between geometric objects is an important problem with applications in areas like computer graphics [107], linkage reconfiguration [74], motion planning [141], and cartography [190, 219]. In many application areas, one often uses a straight-line planar graph embedding as a *proxy* for the underlying data; computing a morph between two straight-line planar embeddings translates to a morph between the actual objects in consideration. Specifically, a *morph* between two straight-line embeddings Γ_0 and Γ_1 is a continuous family of straight-line embeddings Γ_t starting at Γ_0 and ending at Γ_1. Cairns [44] first proved the existence of morphs between straight-line embeddings in 1944; Floater, Gotsman and Surazhsky [99, 111, 205–207] later devised a morph based on Floater's asymmetric extension [96, 97]

of Tutte's spring embedding theorem.

The natural generalization of a straight-line segment on the plane to other surfaces is a *geodesic*, which is a locally-shortest path between its endpoints. Note that on the plane, there is only one geodesic between two points, which is the straight-line segment between them; in general, there can be many geodesics between two points. The natural generalization of a straight-line planar drawing to other surfaces is a *geodesic* drawing, in which every edge is drawn as one of possibly many geodesics between its endpoints.

The notion of locally-shortest implies that we have chosen a metric on our surface. The common visualization of a torus as being the surface of a donut in $\mathbb{R}^3$ does not induce a flat metric. Instead, to obtain a *flat* torus, one can identify opposite sides of a parallelogram and impose a locally Euclidean metric; see Figure 1.2. Equivalently, one can quotient out the Euclidean plane $\mathbb{R}^2$ by a lattice of translations.

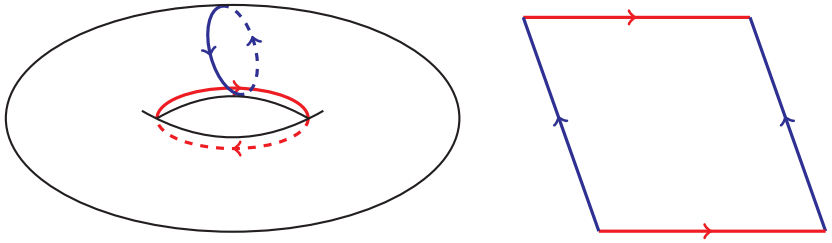

Figure 1.2. Topologically, a torus can be seen as both (a) the surface of a donut, or (b) a parallelogram with opposite sides identified.

Graphs on flat tori thus provide a natural model for biperiodic infinite planar graphs, or situations with periodic boundary conditions. Moreover, it is also well-known that the higher the genus of a surface, the larger the family of graphs that can be drawn without edge crossings; the torus thus serves as a way of working with a larger family of graphs than on the plane. Geodesic toroidal drawings have found applications across many different fields, including chemistry and physics in the study of solids since at least as far back as 1930 [101, 113, 175] and, more recently, rigidity theory [36–38, 154, 185, 186]. *Equilibrium* drawings of graphs on flat tori have previously been studied in the context of computer graphics [109, 197] as well as network visualization [53].

Tutte's spring embedding theorem was first generalized to geodesic triangulations of surfaces of non-positive curvature by Y. Colin de Verdière [60]; this result was later rediscovered by Hass and Scott [115]. The result was also extended to work for arbitrary essentially 3-connected *toroidal* graph embeddings multiple times by independent authors [73, 109, 148]. A natural extension of Floater's equilibrium system exists for toroidal graphs, but is not always solvable: the corresponding system has $2n$ equations over $2n$ variables but has rank $2n-2$. Steiner and

Fischer [197] modified the system by fixing the position of one vertex, giving a system with full rank, but solving the system does not result in a toroidal *embedding* of the graph. Very recently, Luo *et al.* [152] proved a generalization of Floater's theorem to geodesic triangulations of arbitrary surfaces with strictly negative curvature, extending the spring-embedding theorems of Colin de Verdière [60] and Hass and Scott [115] to asymmetric weights; unfortunately, their result does *not* extend to flat tori.

1.1.1 Toroidal Maxwell–Cremona Correspondences

The Maxwell–Cremona correspondence is a classical result in mechanical statics, with roots dating at least as far back as 1725 [214]. In modern terminology, the Maxwell–Cremona correspondence establishes an equivalence between three different structures on a straight-line graph drawing Γ in the plane:

- An *equilibrium stress* on Γ is an assignment of non-zero weights to the edges of Γ, such that the weighted edge vectors around every interior vertex sums to zero.
- A *reciprocal diagram* for Γ is a straight-line drawing Γ^* of the dual graph, in which every edge e^* is orthogonal to the corresponding primal edge e.
- A *polyhedral lifting* of Γ assigns z-coordinates to the vertices of Γ, so that the resulting lifted vertices in $\mathbb{R}^3$ are not all coplanar, but the lifted vertices of each face of Γ are coplanar.

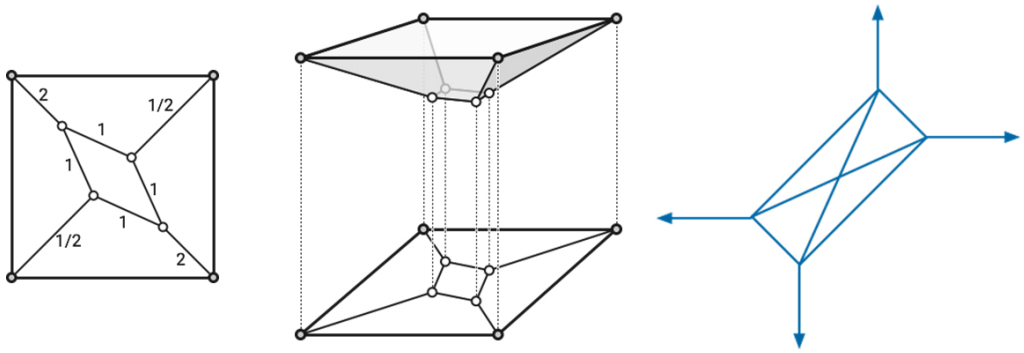

Figure 1.3. A positive equilibrium graph, the induced convex lifting, and orthogonal reciprocal embedding.

If the outer face of Γ is convex, the Maxwell–Cremona correspondence implies an equivalence between equilibrium stresses in Γ that are *positive* on every interior edge, *convex* polyhedral liftings of Γ, and reciprocal *embeddings* of Γ^*. The well-known equivalence between convex liftings and weighted Delaunay complexes implies that all three of these structures are equivalent to a fourth:

- A *Delaunay weighting* of Γ is an assignment of weights to the vertices of Γ, so that Γ is the (power-)weighted Delaunay graph of its vertices.

The correspondence has resulted in numerous applications across a wide variety of fields [18, 20, 54, 57, 65, 68, 74, 75, 81, 121, 122, 172, 173, 181, 182, 191, 202–204].

The goal of Chapter 3 is to try to extend the Maxwell–Cremona–Delaunay correspondence to flat tori. Equilibrium stresses can easily be defined for toroidal graph drawings, as can toroidal duals. Moreover, weighted Delaunay complexes and their dual Voronoi diagrams can be defined intrinsically on flat tori [33]. Borcea and Streinu investigated polyhedral lifts for biperiodic infinite planar graph drawings [38], which are equivalent to flat torus graph drawings. However, their polyhedral lifts are periodic and thus non-convex; moreover, it is not clear what a "convex" polyhedral lifting would look like in a purely toroidal setting. We describe generalizations of *parts* of the Maxwell–Cremona–Delaunay correspondence, with a full discussion of when certain parts of the generalization fail to hold. As an application, we also derive a toroidal analog of Steinitz' theorem (Theorem 3.6).

Portions of Chapter 3 are based on and adapted from joint work by Jeff Erickson [88] and an independently authored manuscript [145].

1.1.2 Toroidal Morphs

Floater and Gotsman [99] constructed a morph between two straight-line drawings of the same planar graph G, with convex faces and the same outer face, based on Floater's asymmetric extension [97] of Tutte's spring embedding theorem [213]; Gotsman and Surazhsky later generalized Floater and Gotsman's technique to arbitrary planar straight-line graphs [111, 205–207]. Floater and Gotsman's method is based on interpolating asymmetric stresses: Given asymmetric stresses describing the initial and final embeddings, one can apply Floater's theorem to convex combinations of the initial and final stresses to obtain intermediary embeddings. As we vary the stress continuously from the initial stress to the final stress, the corresponding embedding varies continuously as well, i.e., we obtain a morph.

Floater and Gotsman were not the first to construct morphs between straight-line planar graph drawings: the history of morphing between planar straight-line graphs dates back to Cairns [44, 45], who, in 1944, was the first to prove the existence of morphs between arbitrary isomorphic planar straight-line triangulations, using an inductive argument based on the idea of collapsing an edge from a low-degree vertex to one of its neighbors. Thomassen later [210] extended Cairns' proof to arbitrary planar straight-line graphs.

Floater and Gotsman's method [99, 111, 205–207] results in morphs that are natural and visually appealing but are represented implicitly; variations on Cairns' edge-collapse method [4,

44, 45, 133, 210] result in efficient explicit representations of morphs that are, unfortunately, not useful for visualization: because the morphs move vertices very close to each other, it is difficult for a user to track the motions of the vertices.

In Chapter 4, we explain how to extend Cairns' existence result to flat tori. The major roadblock towards generalizing the result is the following: Cairns' strategy relies heavily on finding vertices of degree at most 5; the existence of such vertices for planar graphs is guaranteed by Euler's formula. However, Euler's formula for torus graphs implies that the average degree of a vertex is *exactly* 6. The bulk of Chapter 4 is devoted to showing that in spite of this issue, Cairns' strategy can be adapted to flat tori with minimal modifications. Portions of Chapter 4 are based on and adapted from joint work with Erin W. Chambers, Jeff Erickson, and Salman Parsa [51].

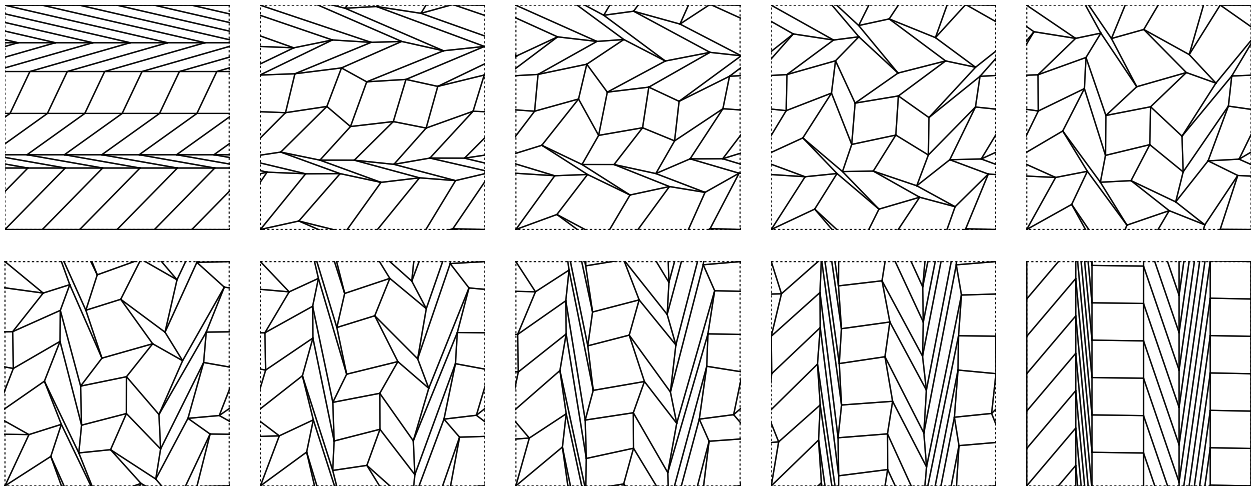

Figure 1.4. Morphing between randomly shifted toroidal grids.

In Chapter 5, we then explain how to extend Floater and Gotsman's barycentric method to flat tori. The major roadblock here is that Floater's asymmetric extension of Tutte's spring embedding theorem does not extend to flat tori: although one can derive a natural extension for the linear system, the system does not have full rank and, in general, cannot be solved. Steiner and Fischer [197] attempt to fix this issue by fixing the position of a single vertex, restoring full rank to the system, but the result of solving their system is not necessarily an embedding, meaning that the result is not a proper morph. In contrast, we describe a simple method for deriving stress vectors whose convex combinations always correspond to linear systems that *do* have solutions; moreover, these solutions correspond to actual embeddings. In other words, our method results in a proper extension of Floater and Gotsman's method to the toroidal setting. Figure 1.4 shows (snapshots of) a morph computed by the resulting algorithm. Our extension of Floater and Gotsman's method also provides a short and simple proof of a

conjecture of Connelly *et al.* [66] on the deformation space of geodesic torus triangulations. Portions of Chapter 5 are based on and adapted from joint work with Jeff Erickson [89].

1.1.3 Efficient Piecewise-Linear Morphs

Whereas Floater and Gotsman's barycentric interpolation method [99, 111, 205–207] results in morphs that are represented implicitly, variations of Cairns' edge-collapse method results in planar morphs with explicit piecewise-linear vertex trajectories. Cairns' original method [44, 45] produced piecewise-linear morphs of exponential complexity; a long series of later works [5, 13, 14, 24], culminating in papers by Alamdari *et al.* [4], and Kleist *et al.* [133], improved the method, ultimately producing efficient algorithms to construct planar morphs with explicit piecewise-linear vertex trajectories of low description complexity. (Alamdari *et al.* [4] and Roselli [184] provide more detailed history of these results.) Given any two isomorphic straight-line embeddings of the same n-vertex planar graph, these papers describe how to construct, in $O(n^{1+\omega/2})$ time, a piecewise-linear morph consisting of $O(n)$ morphing steps. Thus, each vertex moves along a piecewise-linear path with complexity $O(n)$, and the entire morph has complexity $O(n^2)$.

Unfortunately, these improvements rely on a subprocedure that does not seem to have an obvious implementation on flat tori. As such, trying to obtain similar bounds on both description and runtime complexity requires new tools. In Chapter 6, we introduce a novel technique based on Floater and Gotsman's barycentric interpolation method, that results a *unidirectional* morph, i.e., a morph in which all vertices move along parallel geodesics.

We first apply this to *symmetric* stresses to define an equilibrium counterpart to Cairns' edge collapses, allowing us to obtain a piecewise-linear morph between embeddings on flat tori. The morph consists of $O(n)$ morphing steps, and is computed in $O(n^{1+\omega/2})$ time, matching the previous state of the art on the plane. Unfortunately, because this morph is, ultimately, still based on the idea of collapsing edges, they are not useful for visualization purposes.

Next, we observe that unidirectional barycentric interpolation can also be applied to morphing *planar* graphs. We describe a "best-of-both-worlds" result that simultaneously obtains nice properties enjoyed by the two different approaches for planar morphing previously described: our algorithm is quite simple, and produces a piecewise-linear morph that, by avoiding the need to collapse edges, remains good for visualization purposes; see Figure 1.5 for an example. Remarkably, this is achieved without sacrificing description or runtime complexity: the resulting morph consists of $O(n)$ morphing steps, and is computed in $O(n^{1+\omega/2})$ time. An intermediary result also improves a result of Angelini *et al.* [15]. Unfortunately, it is not clear if we can extend the result to the toroidal setting, as it relies heavily on asymmetric stresses.

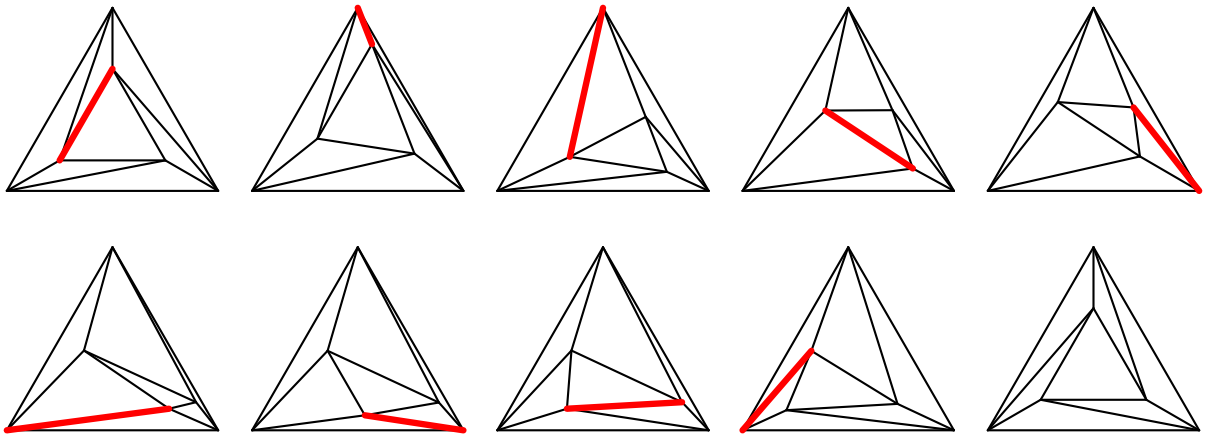

Figure 1.5. Incrementally morphing between planar graphs.

Portions of Chapter 6 are based on and adapted from joint work with Erin W. Chambers, Jeff Erickson, and Salman Parsa [51], and joint work with Jeff Erickson [89].

1.1.4 Open Questions and Miscellany

In Chapter 7, we collect some open questions and ideas related to the results presented in the thesis. Finally, in Chapter 8 we summarize some miscellaneous thoughts, including some half-baked ideas that did not pan out, as well as some useless trivia.

Chapter 2

Preliminaries

We assume some familiarity with basic terminology and definitions in graph theory, topology, and computational geometry. For useful references on graphs and topology from a mathematical viewpoint, see Mohar and Thomassen [167] and Giblin [103]; for a more algorithmic point of view, see É. Colin de Verdière [63] and Erickson [86]. Standard references for computational geometry include Devadoss and O'Rourke [77] and de Berg *et al.* [27].

2.1 Graphs

A ***graph*** is a combinatorial object consisting of a set V of ***vertices*** and a multiset E of unordered pairs of vertices called ***edges***; for those who prefer definitions that are overly concise yet still fully precise, a graph is a one-dimensional semi-simplicial set. All graphs we consider will be *undirected*, and we explicitly allow graphs with loops and parallel edges, see Figures 2.1–2.3.

Formally, we regard each edge of a graph G as a pair of opposing ***darts***, each directed from one endpoint, called the ***tail***, toward the other endpoint, called the ***head***. For each edge e, we arbitrarily label the darts e^+ and e^-; we call e^+ the ***reference dart*** of e. Given a dart d, we denote its ***reversal*** by $\mathbf{rev}(d)$; thus, for example, $\mathrm{head}(\mathrm{rev}(d)) = \mathrm{tail}(d)$ and $\mathrm{rev}(\mathrm{rev}(d)) = d$ for every dart d. Given a dart d, its edge will be denoted by $|d|$. We will sometimes write $d : u \to v$ to declare that $\mathrm{tail}(d) = u$ and $\mathrm{head}(d) = v$; at the risk of confusing the reader, we will also write $u \to v$ to denote an arbitrary *anonymous* dart with tail u and head v, and $v \to u$ for the reversal of $u \to v$. The (multi)set of darts will be referred to by D.

A graph is ***simple*** if it does not contain any self-loops or parallel edges, and ***3-connected*** if it remains connected after deleting any two vertices.

2.2 Planar Graph Drawings

A ***drawing*** Γ of a graph G on the (Euclidean) plane $\mathbb{R}^2$ is a continuous function from G as a topological space to $\mathbb{R}^2$: each vertex of G is mapped to a point in the plane, and each edge of G to a path between its endpoints.

An ***embedding*** is an *injective* drawing, mapping vertices of G to distinct points and edges

to interior-disjoint simple paths between their endpoints.[1] The faces of an embedding are the components of the complement of the image of the graph; in particular, embeddings are cellular, i.e., all faces are open disks. In any embedded graph, **left(d)** and **right(d)** denote the faces immediately to the left and right of any dart d. An embedding is a ***triangulation*** if every face is bounded by three edges, and ***convex*** if every face are convex.

We will say that a graph G is ***planar*** if there is an embedding of G on the plane.

A ***straight-line drawing*** of G on $\mathbb{R}^2$ is a drawing that maps edges to line segments, and a ***straight-line embedding*** is an embedding that maps edges to line segments.

2.3 Flat Tori

The *square* flat torus $\mathbb{T}_\square$ is the surface obtained by identifying opposite sides of the Euclidean unit square $[0,1]^2$ via the identification $(x,0) \sim (x,1)$ and $(0,y) \sim (1,y)$.

More generally, a ***flat torus*** is a surface defined by identifying opposite sides of a parallelogram in $\mathbb{R}^2$: For any nonsingular 2×2 matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, let $\mathbb{T}_M$ denote the flat torus obtained by identifying opposite edges of the ***fundamental parallelogram*** $\Diamond_M$ with vertex coordinates $\binom{0}{0}$, $\binom{a}{c}$, $\binom{b}{d}$, and $\binom{a+b}{c+d}$. In particular, $\mathbb{T}_\square = \mathbb{T}_I$. The linear map $M : \mathbb{R}^2 \to \mathbb{R}^2$ induces a homeomorphism $\underline{M} : \mathbb{T}_\square \to \mathbb{T}_M$.

Equivalently, $\mathbb{T}_M$ is the quotient space of the plane $\mathbb{R}^2$ with respect to the lattice L_M of translations generated by the columns of M; in particular, the square flat torus is the quotient space $\mathbb{R}^2/\mathbb{Z}^2$. The quotient map $\pi_M : \mathbb{R}^2 \to \mathbb{T}_M$ is called a *covering* map or ***projection***. A ***lift*** of a point $p \in \mathbb{T}_M$ is any point in the preimage $\pi_M^{-1}(p) \subset \mathbb{R}^2$.

A flat torus $\mathbb{T}_M$ naturally inherits a metric from this projection: for any two points $p, q \in \mathbb{T}_M$, their distance is equal to minimum over all lifts $\widetilde{p}$ and $\widetilde{q}$ of p and q, respectively, of the standard Euclidean distance between $\widetilde{p}$ and $\widetilde{q}$. A ***geodesic*** in $\mathbb{T}_M$ is the projection of any line segment in $\mathbb{R}^2$; we emphasize that geodesics are not necessarily shortest paths; indeed, there are infinitely many geodesics between any two points, see Figure 2.1.

2.4 Toroidal Graph Drawings

Generalizing from the planar setting, a ***drawing*** Γ of G on a flat torus $\mathbb{T}$ is a continuous function from G as a topological space to $\mathbb{T}$. Just as in the planar case, an ***embedding*** is an *injective* drawing, mapping vertices of G to distinct points and edges to interior-disjoint simple

[1] We note that this usage differs from standard terminology in many graph drawing papers, where "embedding" refers to either a homeomorphism class of (not necessarily injective) drawings or a rotation system.

paths between their endpoints. As before, the faces of an embedding are the components of the complement of the image of the graph; in particular, embeddings are cellular, i.e., all faces are open disks. **left(d)** and **right(d)** will continue denote the faces immediately to the left and right of any dart d (in a toroidal embedding, these are possibly the same face).

We will say that a graph G is ***toroidal*** if there is an embedding of G on a (flat) torus.

The ***universal cover*** $\widetilde{\Gamma}$ of a drawing Γ on any flat torus $\mathbb{T}_M$ is the unique infinite periodic graph in $\mathbb{R}^2$ such that $\pi_M(\widetilde{\Gamma}) = \Gamma$; in particular, each vertex, edge, or face of $\widetilde{\Gamma}$ projects to a vertex, edge, or face of G, respectively.

When moving from the plane to flat tori, one generalizes from speaking about line segments to geodesics; accordingly, a ***geodesic drawing*** of G on any flat torus $\mathbb{T}$ is a drawing that maps edges to geodesics, and a ***geodesic embedding*** is an embedding that maps edges to geodesics. Equivalently, a drawing (embedding) is geodesic if its universal cover $\widetilde{G}$ is a straight-line drawing (embedding) on the plane.

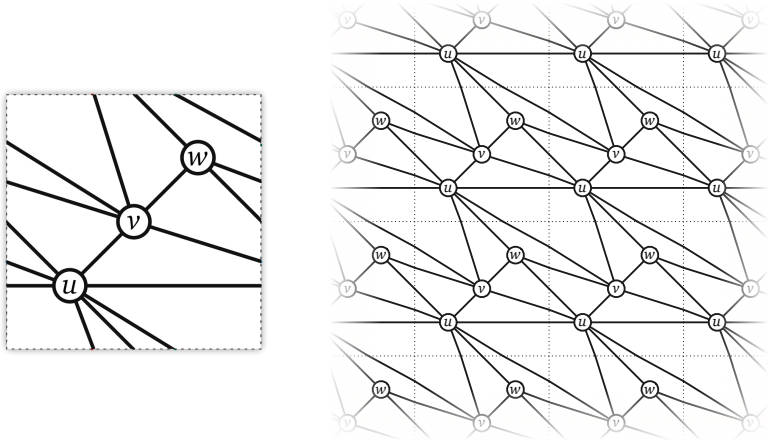

Figure 2.1. An essentially simple, essentially 3-connected geodesic graph embedding on the square flat torus (with multiple edges between pairs of vertices) and a small portion of its universal cover.

A toroidal drawing Γ is ***essentially simple*** if its universal cover $\widetilde{\Gamma}$ is simple, and ***essentially 3-connected*** if $\widetilde{\Gamma}$ is 3-connected [108, 163–166]. We emphasize that essential simplicity and essential 3-connectedness are features of *embeddings*, not of their corresponding abstract graphs. Every *geodesic* toroidal embedding is essentially simple.

An embedding Γ is a ***triangulation*** if every face of Γ is bounded by three edges; for toroidal embeddings, this is equivalent to the universal cover $\widetilde{\Gamma}$ being a planar triangulation. In particular, we do not insist that triangulations are simplicial complexes. Every *geodesic triangulation* is essentially 3-connected.

The ***link*** of a vertex $\widetilde{u}$ in the universal cover $\widetilde{\Gamma}$ is the simple polygon formed by the boundary of the union of the (closed) faces incident to $\widetilde{u}$; the vertices of the link are the neighbors of $\widetilde{u}$. We emphasize that when projecting a link down to the flat torus, the vertices and edges of the

link need not remain distinct. For a vertex u in G, we informally write "link of u" to refer to the link of an arbitrary lift $\widetilde{u}$ of u, and similarly for edges of G. Because the links of any two lifts are congruent, any property proven about one lift applies to all of the others.

Just as a flat torus inherits its geometry via projection from the plane, geometric properties of geodesics, polygons, and embeddings on a flat torus are defined by projection from the universal cover.

For example, the angle between two edges (or geodesics) e and e' at a common vertex u is equal to the angle between lifts $\widetilde{e}$ and $\widetilde{e}'$ at a common lift $\widetilde{u}$. Thus, we say a face of a toroidal embedding is convex if its lifts are convex. An embedding Γ is ***convex*** if all of its faces are convex.

Similarly, the cyclic order of edges around a vertex u of Γ is the cyclic order of the corresponding edges around an arbitrary lift $\widetilde{u}$. In particular, if u is incident to a loop, that loop appears twice in cyclic order around u, and each lift $\widetilde{u}$ of u is incident to two different lifts of that loop. Finally, angles in the link of a vertex in Γ are projections of angles in the link of an arbitrary lift $\widetilde{u}$.

2.5 Coordinate Representations

When trying to do computational geometry involving graph drawings, one needs to have a way to represent said drawings.

To represent an arbitrary straight-line embedding of a graph in the plane, it suffices to record the coordinates of each vertex; each edge in the embedding is the unique line segment between its endpoints. For geodesic graphs on flat tori, we need to specify the way that edges cross the boundary edges of the fundamental parallelogram.

We can represent a geodesic drawing Γ of a graph G on the *square* flat torus $\mathbb{T}_\square$ by associating a ***coordinate vector*** $p_v \in \mathbb{R}^2$ to every vertex v and a ***translation vector*** $\boldsymbol{\tau_d} \in \mathbb{Z}^2$ to every dart d of G. Each coordinate vector p_v records the coordinates of an arbitrary lift $\widetilde{v}$ of v to the universal cover $\widetilde{\Gamma}$, and the translation vector of each dart encodes which lifts of its endpoints are connected in $\widetilde{\Gamma}$. Specifically, for each dart d in G, the universal cover $\widetilde{\Gamma}$ contains an edge between $p_{\text{tail}(d)}$ and $p_{\text{head}(d)} + \tau_d$, and therefore also contains an edge between $p_{\text{tail}(d)} + (i, j)$ and $p_{\text{head}(d)} + \tau_d + (i, j)$ for all integers i and j. Translation vectors are antisymmetric: $\tau_d = -\tau_{\text{rev}(d)}$.

Coordinate representations are not unique; in fact, each toroidal embedding has an infinite family of equivalent representations. Two coordinate representations (p, τ) and (p', τ') with the same underlying graph are *equivalent*, meaning they represent the same geodesic

embedding (up to translation), if and only if

$$\begin{aligned}\Delta_d &:= p_{\text{head}(d)} + \tau_d - p_{\text{tail}(d)} \\ &= p'_{\text{head}(d)} + \tau'_d - p'_{\text{tail}(d)}\end{aligned} \tag{2.1}$$

for every dart d. The vector $\boldsymbol{\Delta_d}$, which we call the ***displacement vector*** of d, is the difference between the head and tail of any lift of d to $\widetilde{\Gamma}$. Reversing a dart negates its displacement: $\Delta_d = -\Delta_{\text{rev}(d)}$. We sometimes write $\boldsymbol{\Delta x_d}$ and $\boldsymbol{\Delta y_d}$ to denote the first and second coordinates of Δ_d.

Let (p, τ) be any coordinate representation of Γ. Given *arbitrary* integer vector $\pi_v \in \mathbb{Z}^2$ for each vertex of G, we can define a new coordinate representation (p^π, τ^π) as follows, for every vertex v and dart d:

$$\begin{aligned}p^\pi_v &= p_v + \pi_v \\ \tau^\pi_d &= \tau_d + \pi_{\text{tail}(d)} - \pi_{\text{head}(d)}\end{aligned} \tag{2.2}$$

Easy calculation implies that the representations (p, τ) and (p^π, τ^π) are equivalent. This transformation is a multidimensional generalization of the *reweighting* or *repricing* strategy proposed by Tomizawa [211] and Edmonds and Karp [85] for minimum-cost flows, and later used by Johnson to compute shortest paths [129].

A geodesic drawing of a graph G on any other flat torus $\mathbb{T}_M$ can be specified fully by a coordinate vector and translation vector for its image on $\mathbb{T}_\square$ and the linear transformation M. For example, displacement vectors on $\mathbb{T}_M$ are linear transformations of corresponding displacement vectors on $\mathbb{T}_\square$: $\Delta_d = M\left(p_{\text{head}(d)} - p_{\text{tail}(d)} + \tau_d\right)$.

It will often be convenient to work with *canonical* coordinates. A ***(canonical) coordinate vector*** for a vertex v is a vector $p_v \in [0, 1)^2$. In this setting, we refer the translation vector for each dart d as the ***crossing*** (or ***homology***) ***vector*** $\boldsymbol{\chi_d} \in \mathbb{Z}^2$. Crossing vectors have the following intuitive interpretation: the first coordinate of χ_d is the number of times d crosses the vertical boundary rightward (if the first coordinate of χ_d is negative, then its absolute value is the number of times d crosses the vertical boundary leftward); and the second coordinate of x_d is the number of times d crosses the horizontal boundary upward (similarly, if the second coordinate is negative, then its absolute value is the number of times d crosses the horizontal boundary downward).

Crossing vectors have been used in several previous algorithms for surface graphs [48–50, 87, 90, 91] and simplicial complexes [43, 79, 80] to encode the homology classes of cycles. Translation vectors have traditionally used to model periodic (or "dynamic") graphs [55, 64, 125–127, 131, 136, 174, 179, 217] and more recently used to model periodic bar-and-joint frameworks

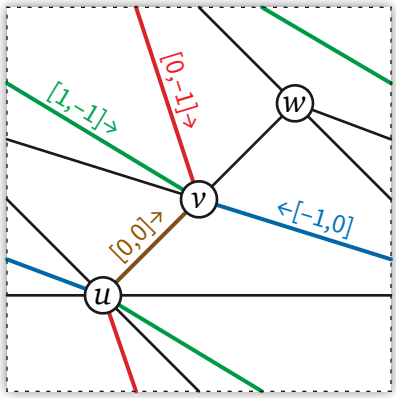

Figure 2.2. The geodesic embedding from Figure 2.1, showing the crossing vectors of all four darts from u to v.

[36, 38, 78, 132, 154, 170, 187, 188].

2.6 Homotopy, Isotopy, and Morphs

A ***homotopy*** between two drawings Γ_0 and Γ_1 of the same graph G on a surface X is a continuous function $H\colon [0,1] \times G \to X$ where $H(0,\cdot) = \Gamma_0$ and $H(1,\cdot) = \Gamma_1$. A homotopy is an ***isotopy*** if each intermediate function $H(t,\cdot)$ is injective. In other words, an isotopy is a continuous family of embeddings $(\Gamma_t)_{t\in[0,1]}$ that interpolates between Γ_0 and Γ_1. (Edges in these intermediate embeddings Γ_t are not required to be geodesics.)

The isotopy type of a planar embedding is fully determined by the combinatorial structure of the embedding, otherwise known as a *rotation system*; given a 3-connected planar graph, the rotation system is fully determined by the choice of outer face and orientation for said outer face. Thus two planar embeddings of the same 3-connected planar graph G are isotopic if and only if they are *isomorphic* as embeddings if and only if they have the same (oriented) outer face.

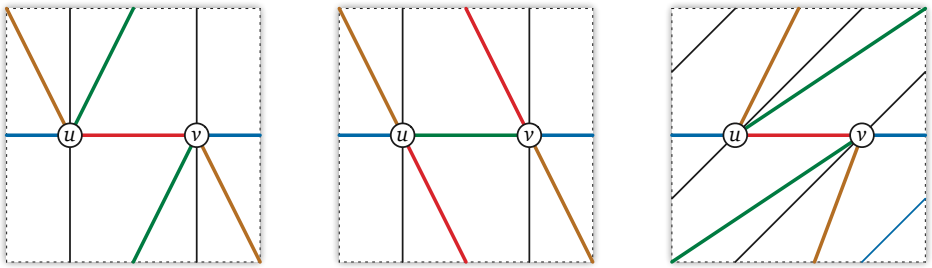

Figure 2.3. Three combinatorially equivalent but non-isotopic geodesic toroidal triangulations with parallel edges and loops.

Unlike in the case of planar embeddings, two toroidal embeddings of the same graph need

not be isotopic, even if they have the same rotation system; see Figure 2.3. A recent algorithm of É. Colin de Verdière and de Mesmay [56] can decide whether two toroidal drawings of the same graph are isotopic in linear time; we describe an arguably simpler linear-time algorithm in Section 2.8. However, neither of these algorithms actually construct an isotopy if one exists; rather, they check whether the two embeddings satisfy certain topological properties that characterize isotopy [138–140].

A ***morph*** between two isotopic straight-line (geodesic) embeddings Γ_0 and Γ_1 on the plane (torus) is a continuous family of straight-line (geodesic) embeddings $(\Gamma_t)_{t\in[0,1]}$ from Γ_0 to Γ_1, i.e., a straight-line (geodesic) isotopy between Γ_0 and Γ_1. Any planar (toroidal) morph is completely determined by the continuous motions on the vertices; straight-line (geodesic) edges update in the obvious way.

When talking about morphing between planar graph embeddings, it is standard to start by fixing the same (convex) outer face in the initial and final embeddings; for toroidal embeddings, we will instead fix the position of an arbitrary vertex.

A morph is ***linear*** if each vertex moves along a straight line (geodesic) from its initial position to its final position, and ***piecewise-linear*** if it the concatenation of linear morphs. Any planar piecewise-linear morph can be described by a finite sequence $\Gamma_0, \ldots, \Gamma_k$ of straight-line embeddings; toroidal piecewise-linear morphs also need to include a set of geodesics connecting vertices in each Γ_i to corresponding vertices in Γ_{i+1}. A linear morph is ***unidirectional*** if vertices move along parallel lines.

For the sake of convenience, we will assume the square flat torus $\mathbb{T}_\square$ as the setting for morphing toroidal graphs; however, the algorithms we describe will apply to any other flat torus $\mathbb{T}_M$ via linear transformation of the coordinates.

2.7 Homology and Circulations

The ***(integer) homology class*** $[\gamma]$ ***of a directed cycle*** γ in a toroidal drawing Γ is the sum of the crossing vectors of its forward darts.[2] A cycle is ***contractible*** if its homology class is $\binom{0}{0}$ and ***essential*** otherwise. Equivalently, a cycle on $\mathbb{T}$ is contractible if it is homotopic to a single point. In particular, the boundary cycle of each face of Γ is contractible. Two cycles on a torus are homotopic if and only if they have the same integer homology class; similarly, two embeddings of the same graph G on the same flat torus $\mathbb{T}$ are homotopic if and only if every cycle has the same homology class in both embeddings [56, 140].

A ***circulation*** ϕ in a graph G is a function from the darts of G to the reals, such that

[2] It is for this reason that crossing vectors are also called homology vectors.

$\phi(u\to v) = -\phi(v\to u)$ for every dart $u\to v$ and $\sum_{u\to v} \phi(u\to v) = 0$ for every vertex u. We represent circulations by column vectors in $\mathbb{R}^E$, indexed by the edges of G, where $\phi_e = \phi(e^+)$. Let Ξ_Γ denote the $2 \times E$ matrix whose columns are the crossing vectors of the reference darts in an embedding Γ of G. The ***homology class of a circulation*** is the matrix-vector product

$$[\phi] := \Xi_\Gamma \phi \;=\; \sum_{e\in E} \phi(e^+)x(e^+). \tag{2.3}$$

(This identity directly generalizes our earlier definition of the homology class $[\gamma]$ of a cycle γ.)

On the *square* flat torus, the integer homology class of any directed cycle is also equal to the sum of the displacement vectors of its darts:

$$[\gamma] := \sum_{u\to v\in\gamma} x(u\to v) \;=\; \sum_{u\to v\in\gamma} \Delta_{u\to v}. \tag{2.4}$$

In particular, the total displacement of any contractible cycle is zero, as expected. Define the ***displacement matrix*** $\boldsymbol{\Delta}$ of a geodesic drawing to be the $2 \times E$ matrix whose columns are the displacement vectors of the reference darts of G. Extending the identity for directed cycles to circulations by linearity gives us the following useful lemma:

Lemma 2.1. Fix a geodesic drawing of a graph G on $\mathbb{T}_\square$ with displacement matrix Δ. For any circulation ϕ in G, we have $\Delta\phi = \Xi_\Gamma\phi = [\phi]$.

More generally, for non-square tori $\mathbb{T}_M$, we have the identity $\Delta\phi = M\Xi_\Gamma\phi = M[\phi]$.

2.8 Coordinate Normalization and Isotopy Testing

Let Γ_0 and Γ_1 be two isotopic geodesic embeddings of the same graph G on $\mathbb{T}_\square$, given by coordinate representations $(p(0), \tau(0))$ and $(p(1), \tau(1))$ respectively. To simplify the presentation of our morphing algorithms, we implicitly assume that the translation vectors in both representations are identical: $\tau(0)_d = \tau(1)_d$ for every dart d. This assumption allows us to describe, reason about, and ultimately compute a morph from Γ_0 to Γ_1 entirely in terms of changes to the vertex coordinates; all translation vectors remain fixed throughout the morph.

If necessary, we can enforce this assumption in $O(n)$ time using the following *normalization* algorithm. Let $(p(0), \tau(0))$ and $(p(1), \tau(1))$ be the given coordinate representations of Γ_0 and Γ_1, respectively. First, construct an arbitrary spanning tree T of the underlying graph G, directed away from an arbitrary root vertex r. For every vertex v, let $P(v)$ denote the unique

directed path in T from r to v. For each vertex v, let

$$\pi_v = \sum_{d \in P(v)} (\tau(1)_d - \tau(0)_d). \tag{2.5}$$

We can compute the vectors π_v for all vertices in $O(n)$ time by preorder traversal of T. Finally, we replace the target representation $(p(1), \tau(1))$ with the equivalent representation $(p(1)^\pi, \tau(1)^\pi)$.

Lemma 2.2. For all darts d in T, we have $\tau(1)^\pi_d = \tau(0)_d$.

Proof. Let d be any dart in T directed from some vertex u to one of its children v in T. Straightforward calculation implies

$$\begin{aligned}\pi_u - \pi_v &= \sum_{d' \in P(u)} (\tau(1)_{d'} - \tau(0)_{d'}) \\ &\qquad - \sum_{d' \in P(v)} (\tau(1)_{d'} - \tau(0)_{d'}) \\ &= \tau(0)_d - \tau(1)_d\end{aligned}$$

and therefore $\tau(1)^\pi_d = \tau(1)_d + \pi_u - \pi_v = \tau(0)_d$. A similar calculation (or antisymmetry) implies that $\tau(1)^\pi_d = \tau(0)_d$ for every dart d directed from a vertex to its parent in T. □

Theorem 2.1. Γ_0 and Γ_1 are isotopic if and only if $\tau(1)^\pi_d = \tau(0)_d$ for every dart d.

Proof. We exploit a theorem of Ladegaillerie [138–140], which states that two embeddings are isotopic if and only if every cycle in one embedding is homotopic to the corresponding cycle in the other embedding.

For any dart d in the underlying graph G, let $\Delta_0(d)$ and $\Delta_1(d)$ denote the displacement vectors of d in Γ_0 and Γ_1, respectively. (We emphasize that displacement vectors are independent of the coordinate representation.) For any directed cycle C in G, let $\Delta_0(C)$ and $\Delta_1(C)$ denote the sum of the displacement vectors of its darts:

$$\Delta_0(C) := \sum_{d \in C} \Delta_0(d) \quad \textit{and} \quad \Delta_1(C) := \sum_{d \in C} \Delta_1(d)$$

Recall that $\Delta_0(C)$ is the integer homology class of C in Γ_0, and that two cycles on the torus are homotopic if and only if they have the same integer homology class; in particular, the image of C in Γ_0 is contractible if and only if $\Delta_0(C) = \binom{0}{0}$. Ladegaillerie's theorem implies that Γ_0 and Γ_1 are isotopic if and only if $\Delta_0(C) = \Delta_1(C)$ for every cycle C.

The spanning tree T defines a set of *fundamental cycles* that span the cycle space of G. Specifically, for each dart d that is not in T, the fundamental directed cycle $C_T(d)$ consists of d and the unique directed path in T from head(d) to tail(d). Every directed cycle in G (indeed every *circulation* in G) can be expressed as a linear combination of fundamental cycles. It follows by linearity that Γ_0 and Γ_1 are isotopic if and only if every *fundamental* cycle has the same integer homology class in both embeddings, or in other words, if and only if $\Delta_0(C_T(d)) = \Delta_1(C_T(d))$ for every dart $d \in G \setminus T$.

Straightforward calculation implies that the homology class of any cycle is also equal to the sum of the *translation* vectors of its darts with respect to *any* coordinate representation:

$$\begin{aligned}\Delta_0(C) &= \sum_{d\in C} \tau(0)_d \\ \Delta_1(C) &= \sum_{d\in C} \tau(1)_d \;=\; \sum_{d\in C} \tau(1)^{\pi}_{d}.\end{aligned}$$

In particular, for any non-tree dart $d \notin T$, we immediately have

$$\begin{aligned}\Delta_0(C_T(d)) - \Delta_1(C_T(d)) &= \sum_{d'\in C_T(d)} \left(\tau(0)_{d'} - \tau(1)^{\pi}_{d'}\right) \\ &= \tau(0)_d - \tau(1)^{\pi}_{d}\end{aligned}$$

Thus, $\tau(1)^{\pi}_{d} = \tau(0)_d$ for every dart d if and only if $\Delta_0(C) = \Delta_1(C)$ for every fundamental cycle C, which completes the proof of the theorem. □

Theorem 2.1 and our normalization algorithm immediately imply an $O(n)$-time algorithm to test whether two given coordinate representations $(p(0), \tau(0))$ and $(p(1), \tau(1))$ represent isotopic toroidal embeddings of the same graph G. Our algorithm is arguably simpler than the isotopy algorithm of É. Colin de Verdière and de Mesmay [56], which is also based on Ladegaillerie's theorem [138–140]. On the other hand, our isotopy algorithm is specific to geodesic embeddings on the flat torus, whereas É. Colin de Verdière and de Mesmay's algorithm works for arbitrary combinatorial embeddings of graphs on arbitrary 2-manifolds.

2.9 Duality and Reciprocality

Every embedding Γ on the defines a ***dual graph*** G^* whose vertices correspond to the faces of Γ, where two vertices in G^* are connected by an edge for each edge separating the corresponding pair of faces in Γ. This dual graph G^* has a natural embedding in which each vertex f^* of G^* lies in the interior of the corresponding face f of G, each edge e^* of G^* crosses

only the corresponding edge e of G, and each face p^* of G^* contains exactly one vertex p of G in its interior. Each dart d in G has a corresponding dart d^* in G^*, defined by setting $\text{head}(d^*) = \text{left}(d)^*$ and $\text{tail}(d^*) = \text{right}(d^*)$; intuitively, the dual of a dart in G is obtained by rotating the dart *counterclockwise*. See Figure 2.4.

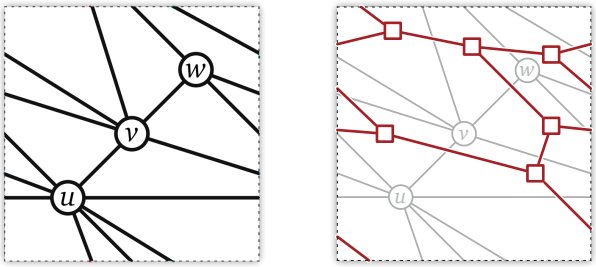

Figure 2.4. The geodesic embedding from Figure 2.1 and its dual graph.

We regard any *drawing* of G^* to be ***(orthogonal) dual*** to G if and only if it is homotopic to this natural embedding; this also defines the dual graph for any *drawing* of G that is homotopic to an embedding. Any drawing of G^* on $\mathbb{T}_\square$ is ***rotated dual*** to G if its image is homotopic to the 90° *clockwise* rotation of the image of the aforementioned natural embedding of G^* on $\mathbb{T}_\square$. More generally, a drawing of G^* is rotated dual to G on some flat torus $\mathbb{T}$ if the image of G^* on $\mathbb{T}_\square$ is rotated dual to the image of G on $\mathbb{T}_\square$.

Two dual geodesic drawings Γ and Γ^* on the same flat torus $\mathbb{T}$ are ***orthogonal reciprocal*** if every edge e in Γ is *orthogonal* to its dual edge e^* in Γ^*; Γ and Γ^* are ***parallel reciprocal*** if every edge e in G is *parallel* to its dual edge e^* in G^*.

Note that on the plane, rotated duals are homotopic to regular duals via a 90° counterclockwise rotation; in general this is not the case on flat tori. As a result, one typically speaks of planar straight-line drawings as simply being ***reciprocal*** to each other.

When working with duality, it will prove convenient to treat vertex coordinates, displacement vectors, crossing vectors, and circulations in any embedding Γ of G as *column* vectors, and their counterparts in any dual embedding Γ^* of G^* as *row* vectors. For any vector $v \in \mathbb{R}^2$ we define $\boldsymbol{v^\perp} := (Jv)^T$, where $\boldsymbol{J} := \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ is the matrix for a 90° counterclockwise rotation. Note that $J^T = J^{-1} = -J$. Similarly, for any $2 \times n$ matrix A, we define $\boldsymbol{A^\perp} := (JA)^T = -A^T J$.

A ***cocirculation*** in G a *row* vector $\theta \in \mathbb{R}^E$ whose transpose describes a circulation in G^*. The ***cohomology class*** $[\boldsymbol{\theta}]^*$ of any cocirculation with respect to a drawing Γ^* of the dual is the transpose of the homology class of the circulation θ^T in Γ^*.

Let Γ be a toroidal embedding of a graph G. Recalling that Ξ_Γ is the $2 \times E$ matrix whose columns are crossing vectors of edges in Γ, let ξ_1 and ξ_2 denote the first and second rows of Ξ_Γ.

Lemma 2.3. The row vectors ξ_1 and ξ_2 describe cocirculations in G whose cohomology classes with respect to the standard dual embedding of G^* are $[\xi_1]^* = (0,1)$ and $[\xi_2]^* = (-1,0)$.

Proof. Without loss of generality, assume that Γ is an embedding on the *square* flat torus $\mathbb{T}_\square$, with no vertices on the boundary of the fundamental square $\square$. Let γ_1 and γ_2 denote directed cycles in $\mathbb{T}_\square$ (*not* in G) induced by the boundary edges of $\square$, directed respectively rightward and upward.

Let $d_0, d_1, \ldots, d_{k-1}$ be the sequence of darts in Γ that cross γ_2 from left to right, indexed by the upward order of their intersection points. Each dart d that appears in this sequence appears exactly $\xi_1(d)$ times, once for each crossing. For each index i, we have $\text{left}(d_i) = \text{right}(d_{i+1 \bmod k})$; thus, the corresponding sequence of dual darts $d_0^*, d_1^*, \ldots, d_{k-1}^*$ describes a closed walk in G^*.

In the standard dual embedding of G^*, this closed walk can be continuously deformed to γ_2, so it has the same homology class as γ_2; see Figure 2.5. We conclude that $[\xi_1]^* = (0,1)$.

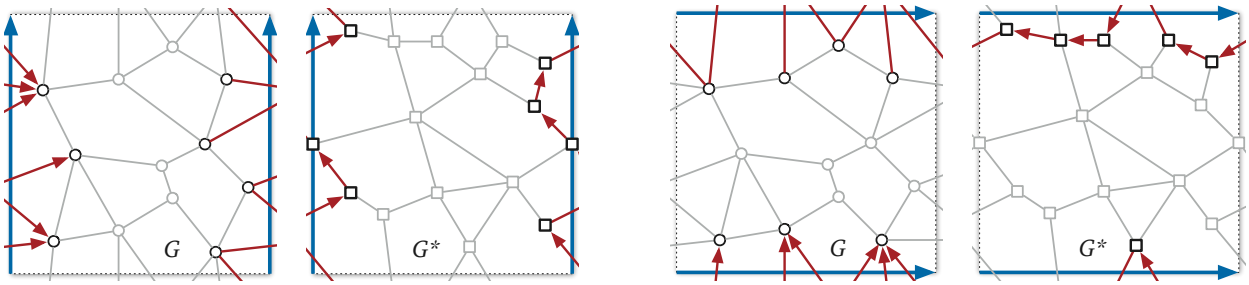

Figure 2.5. Proof of Lemma 2.3: The darts in G crossing either boundary edge of the fundamental square dualize to a closed walk in G^* parallel to that boundary edge.

Symmetrically, the darts crossing γ_1 upward define a closed walk in the standard dual embedding of G^* in the same homology class as the reversal of γ_1, and therefore $[\xi_2]^* = (-1,0)$. □

Lemma 2.4. The row vectors ξ_1 and ξ_2 describe cocirculations in G whose cohomology classes with respect to the *rotated* dual embedding of G^* are $[\xi_1]^* = (1,0)$ and $[\xi_2]^* = (0,1)$.

Proof. Just as in the proof of Lemma 2.3, we assume that Γ is an embedding on the square flat torus $\mathbb{T}_\square$, with no vertices on the boundary of the fundamental square $\square$, let γ_1 and γ_2 denote directed cycles in $\mathbb{T}_\square$ induced by the boundary edges of $\square$, directed respectively rightward and upward, and let $d_0, d_1, \ldots, d_{k-1}$ be the sequence of darts in Γ that cross γ_2 from left to right, indexed by the upward order of their intersection points; the corresponding sequence of dual darts $d_0^*, d_1^*, \ldots, d_{k-1}^*$ describes a closed walk in G^*.

In the rotated dual embedding of G^*, this closed walk can instead be continuously deformed to γ_1, so it has the same homology class as γ_1, i.e., $[\xi_1]^* = (1,0)$. See Figure 2.6.

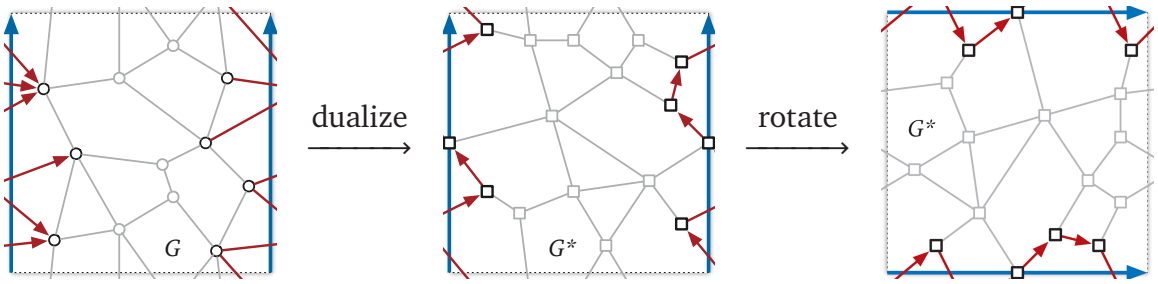

Figure 2.6. Proof of Lemma 2.4: The darts in G crossing either boundary edge of the fundamental square dualize to a closed walk in G^* parallel to *the rotation of* that boundary edge.

Symmetrically, the darts crossing γ_1 upward define a closed walk in the rotated dual embedding of G^* with the same homology class as γ_2, so $[\xi_2]^* = (0,1)$. □

2.10 Coherent Subdivisions

Let Γ be a geodesic embedding on some flat torus $\mathbb{T}_M$, and fix arbitrary real weights π_v for every vertex v of G. Let $i{\to}j$, $i{\to}k$, and $i{\to}\ell$ be three consecutive darts around a common tail p in counterclockwise order. Thus, $\text{left}(i{\to}j) = \text{right}(i{\to}k)$ and $\text{left}(i{\to}k) = \text{right}(i{\to}\ell)$. We call the edge ik ***locally Delaunay*** if the following determinant is positive:

$$\begin{vmatrix} \Delta x_{i\to j} & \Delta y_{i\to j} & \frac{1}{2}|\Delta_{i\to j}|^2 + \pi_i - \pi_j \\ \Delta x_{i\to k} & \Delta y_{i\to k} & \frac{1}{2}|\Delta_{i\to k}|^2 + \pi_i - \pi_k \\ \Delta x_{i\to \ell} & \Delta y_{i\to \ell} & \frac{1}{2}|\Delta_{i\to \ell}|^2 + \pi_i - \pi_\ell \end{vmatrix} > 0. \tag{2.6}$$

This inequality follows by elementary row operations and cofactor expansion from the standard determinant test for appropriate lifts of the vertices i, j, k, ℓ to the universal cover:

$$\begin{vmatrix} 1 & x_i & y_i & \frac{1}{2}(x_i^2 + y_i^2) - \pi_i \\ 1 & x_j & y_j & \frac{1}{2}(x_j^2 + y_j^2) - \pi_j \\ 1 & x_k & y_k & \frac{1}{2}(x_k^2 + y_k^2) - \pi_k \\ 1 & x_\ell & y_\ell & \frac{1}{2}(x_\ell^2 + y_\ell^2) - \pi_\ell \end{vmatrix} > 0. \tag{2.7}$$

(The factor $1/2$ simplifies our later calculations, and is consistent with Maxwell's construction of polyhedral liftings and reciprocal diagrams.) Similarly, we say that an edge is ***locally flat*** if the corresponding determinant is zero. Finally, Γ is the ***weighted Delaunay graph*** of its vertices if every edge of Γ is locally Delaunay and every diagonal of every non-triangular face is locally flat.

One can easily verify that this condition is equivalent to Γ being the projection of the weighted Delaunay graph of the lift $\pi_M^{-1}(V)$ of its vertices V to the universal cover. Results of Bobenko and Springborn [33] imply that any finite set of weighted points on any flat torus has a unique weighted Delaunay graph. We emphasize that weighted Delaunay graphs are *not* necessarily either simple or triangulations; however, every weighted Delaunay graphs on any flat torus is both essentially simple and essentially 3-connected. The dual ***weighted Voronoi graph*** of Γ, also known as its *power diagram* [19, 21], can be defined similarly by projection from the universal cover.

Finally, a geodesic toroidal embedding is ***coherent*** if it is the weighted Delaunay graph of its vertices, with respect to some vector of weights.

2.11 Stresses

A ***stress*** on a graph G is a real vector $\lambda \in \mathbb{R}^D$ indexed by the darts of G. A stress is ***symmetric*** if $\lambda_d = \lambda_{\text{rev}(d)}$ for each dart d; equivalently, one can view a symmetric stress as a real vector $\lambda \in \mathbb{R}^E$ indexed by the *edges* of G.

2.11.1 Equilibrium Stresses

Given a drawing Γ of G on the plane with fixed outer face, a *symmetric* stress λ is an ***equilibrium stress*** for Γ if the weighted sum of the displacement vectors around each *internal* vertex is zero:

$$\sum_{d:u\to v} \lambda_{|d|}(p_v - p_u) = \begin{pmatrix}0\\0\end{pmatrix}, \qquad \text{for each internal vertex } u. \tag{2.8}$$

If Γ instead a drawing on a flat torus, λ is an ***equilibrium stress*** for Γ if the weighted sum of the displacement vectors around *every* vertex is zero:

$$\sum_{d:u\to v} \lambda_{|d|}(p_v + \tau_d - p_u) = \begin{pmatrix}0\\0\end{pmatrix}, \qquad \text{for each vertex } u. \tag{2.9}$$

It may be helpful to imagine each stress coefficient λ_e as a linear spring constant; intuitively, each edge pulls its endpoints inward, with a force equal to the length of e times the stress coefficient λ_e.

Formally, Tutte's 1963 result about equilibrium planar graph embeddings can be summarized as follows:

Theorem 2.2 (Tutte [213])**.** Given a 3-connected planar graph drawing G with a fixed convex polygon for the outer face, assign to each edge e a stress $\lambda_e > 0$. Then the linear system (2.8) has a unique solution, corresponding to a convex straight-line planar graph embedding such that each vertex is the weighted center of mass of its neighbors.

Y. Colin de Verdière [60] extended Theorem (2.2) to all surfaces with non-negative curvature; this result was also independently discovered by Hass and Scott [115]. The cases of Y. Colin de Verdière's theorem for the flat torus has also been independently reproved by many authors [73, 109, 148]. When specialized to $\mathbb{T}_\square$, the result can be stated as follows:

Theorem 2.3 (Y. Colin de Verdière [60])**.** Given an essentially simple, essentially 3-connected toroidal drawing of a torus graph G, assign to each edge e a stress $\lambda_e > 0$. Then the linear system (2.9) is unique up to translation, and corresponds to a convex geodesic toroidal embedding of G such that each vertex is the weighted center of mass of its neighbors.

The linear system (2.8) is a symmetric matrix whose support is the planar graph G, which has balanced separators of size $O(\sqrt{n})$, and thus the linear system can be solved in $O(n^{\omega/2})$ time using the generalized nested dissection technique of Lipton *et al.* [146], where $\omega < 2.37286$ is the matrix multiplication exponent [7, 143]. The linear system (2.9) is a symmetric matrix whose support is the toroidal graph G, which also has balanced separators of size $O(\sqrt{n})$ [6], and can thus also be solved in $O(n^{\omega/2})$ time via generalized nested dissection [146].

Given a planar (toroidal) graph G and a positive symmetric stress λ, we will refer to the embedding promised by Theorem 2.2 (Theorem 2.3) as the ***Tutte embedding*** $\mathbf{\Gamma}^\lambda$.

Theorem 2.3 also provides the following sufficient condition for a displacement matrix to describe a geodesic embedding on the square torus $\mathbb{T}_\square$:

Lemma 2.5. Fix an essentially simple, essentially 3-connected graph G on $\mathbb{T}_\square$, a $2 \times E$ matrix Δ, and a symmetric stress vector λ. Suppose for every directed cycle (and therefore any circulation) ϕ in G, we have $\Delta\phi = \Lambda\phi = [\phi]$. Then Δ is the displacement matrix of a geodesic **drawing** on $\mathbb{T}_\square$ that is homotopic to G. If in addition λ is a positive equilibrium stress for that drawing, the drawing is an embedding.

Proof. A classical result of Ladegaillerie [138–140] implies that two embeddings of the same graph on the same surface are isotopic if and only if every cycle has the same homology class in both embeddings. (See É. Colin de Verdière and de Mesmay [56].) Because homology and homotopy coincide on the torus, the assumption $\Delta\phi = \Lambda\phi = [\phi]$ for every directed cycle immediately implies that Δ is the displacement matrix of a geodesic drawing that is homotopic to G.

If λ is a positive equilibrium stress for that drawing, then the uniqueness clause in Theorem 2.3 implies that the drawing is in fact an embedding. □

As observed by Delgado-Friedrichs [73], equilibrium is invariant under affine transformation. As a consequence, we have the following:

Lemma 2.6. Let G be a geodesic graph on the square flat torus $\mathbb{T}_\square$. If λ is an equilibrium stress for G, then λ is also an equilibrium stress for the image of G on any other flat torus $\mathbb{T}_M$.

2.11.2 Barycentric Stresses

Given a drawing Γ of G on the plane with fixed outer face, a (possibly *asymmetric*) *positive* stress λ is a ***barycentric stress*** for Γ if the weighted sum of the displacement vectors around each *internal* vertex is zero:

$$\sum_{d:u\to v} \lambda_d(p_v - p_u) = \begin{pmatrix}0\\0\end{pmatrix}, \qquad \text{for each internal vertex } u. \tag{2.10}$$

Floater [97] proved the following generalization of Tutte's spring embedding theorem to asymmetric stresses:

Theorem 2.4 (Floater [97])**.** Given a 3-connected planar graph drawing G with a fixed convex polygon for the outer face, assign to each dart d a stress $\lambda_d > 0$. Then the linear system (2.10) has a unique solution, corresponding to a convex straight-line planar graph embedding such that each vertex is the weighted center of mass of its neighbors.

As in the symmetric case, the support of the linear system (2.10) is the planar graph G, which has balanced separators of size $O(\sqrt{n})$ and thus the linear system can be solved in $O(n^{\omega/2})$ time using Alon and Yuster's [8] extension of the generalized nested dissection technique of Lipton *et al.* [146] to asymmetric matrices.

Given a planar graph G and a positive symmetric stress λ, we will refer to the embedding promised by Theorem 2.4 as the ***Floater embedding*** $\mathbf{\Gamma^\lambda}$.

If Γ instead a drawing on a flat torus, λ is an ***barycentric stress*** for Γ if the weighted sum of the displacement vectors around *every* vertex is zero:

$$\sum_{d:u\to v} \lambda_d(p_v + \tau_d - p_u) = \begin{pmatrix}0\\0\end{pmatrix}, \qquad \text{for each vertex } u. \tag{2.11}$$

We will frequently consider the linear system (2.11) in matrix notation as $L^\lambda P = H^\lambda$, where P is the $n \times 2$ matrix whose *rows* are the coordinate vectors p_v for each vertex v, and

$$L^\lambda_{ij} = \begin{cases} \displaystyle\sum_k \sum_{d:i\to k} \lambda_d & \text{if } i = j \\ \displaystyle\sum_{d:i\to j} -\lambda_d & \text{otherwise} \end{cases} \qquad \text{and} \quad H^\lambda_i = \sum_j \sum_{d:i\to j} \lambda_d x_d \tag{2.12}$$

The (unnormalized, asymmetric) Laplacian matrix L^λ has rank $n-1$ [197].

However, this generalization does not hold for the flat torus: Gortler *et al.* [109] showed than whenever the linear system (2.11) has a solution, then it corresponds to a convex geodesic toroidal embedding of G such that each vertex is the weighted center of mass of its neighbors; furthermore, the solution can be found in $O(n^{\omega/2})$ time via generalized nested dissection [8,146] and toroidal separators [6,104].

However, as observed by Steiner and Fischer [197], the system does not in general have a solution; Dealing with this issue is the main subject of Chapter 5. As a concrete example, let G be *any* graph on the torus, and consider the asymmetric stress λ_{bad} that assigns weight 2 to a single dart d and weight 1 to every dart other than d. (There is obviously nothing special about the values 1 and 2 here.)

Lemma 2.7. λ_{bad} is not a realizable weight vector for G.

Proof. Let $u = \text{tail}(d)$ and $v = \text{head}(d)$. For the sake of argument, suppose the linear system $L^{\lambda_{\text{bad}}} P = H^\lambda$ has a solution; let $\Gamma^{\lambda_{\text{bad}}}$ be the resulting drawing. This system remains solvable if we remove row u and arbitrarily fix p_u [197]. All dart weights in this truncated linear system are equal to 1, which implies that the drawing $\Gamma^{\lambda_{\text{bad}}}$ is identical to the Tutte drawing Γ^1 for the all-1s weight vector. Comparing the two linear systems, we conclude that $p_v - p_u + \tau_d = (0,0)$; that is, the edge of d has length zero in $\Gamma^{\lambda_{\text{bad}}} = \Gamma^1$. But this is impossible; every edge in a Tutte drawing has non-zero length [109, Lemma B.5]. □

Given a positive a torus graph G and positive stress λ, whenever the system (2.11) *does* have a solution, we will refer to the resulting embedding as the ***Floater embedding*** $\boldsymbol{\Gamma^\lambda}$. In this case, we will say that λ is ***realizable***, and that Γ^λ ***realizes*** λ. Every realizable weight vector is realized by a two-dimensional family of drawings that differ by translation; we can remove this ambiguity by arbitrarily fixing the position of one vertex [109,197].

Given a *convex* (planar or toroidal) embedding Γ, a barycentric stress can be computed in $O(n)$ time using, for example, Floater's mean-value coordinates [98,120] or, when all vertices have bounded degree, linear programming [192].

Chapter 3

Toroidal Maxwell–Cremona Correspondences

The Maxwell–Cremona correspondence is a fundamental theorem establishing an equivalence between three different structures on a straight-line graph drawing Γ in the plane:

- An *equilibrium stress* on Γ is an assignment of non-zero weights to the edges of Γ, such that the weighted edge vectors around every interior vertex sums to zero.
- A *reciprocal diagram* for Γ is a straight-line drawing Γ^* of the dual graph, in which every edge e^* is orthogonal to the corresponding primal edge e.
- A *polyhedral lifting* of Γ assigns z-coordinates to the vertices of Γ, so that the resulting lifted vertices in $\mathbb{R}^3$ are not all coplanar, but the lifted vertices of each face of Γ are coplanar.

If the outer face of Γ is convex, the Maxwell–Cremona correspondence implies an equivalence between equilibrium stresses in Γ that are *positive* on every interior edge, *convex* polyhedral liftings of Γ, and reciprocal *embeddings* of Γ^*. Moreover, as Whiteley *et al.* [221] and Aurenhammer [18] observed, the well-known equivalence between convex liftings and weighted Delaunay complexes [19, 20, 42, 84, 215] implies that all three of these structures are equivalent to a fourth:

- A *Delaunay weighting* of Γ is an assignment of weights to the vertices of Γ, so that Γ is the (power-)weighted Delaunay graph [19, 33] of its vertices.

In this chapter, we develop a few different generalizations of the Maxwell–Cremona–Delaunay correspondence to geodesic drawings of graphs on flat tori. We prove the following correspondences for any geodesic drawing Γ on any flat torus $\mathbb{T}$.

- Any reciprocal drawing Γ^* on $\mathbb{T}$ defines unique equilibrium stresses in both Γ and Γ^* (Lemma 3.1). This is true for both orthogonal and parallel reciprocality.
- Unlike in the plane, Γ may have equilibrium stresses that are not induced by reciprocal embeddings; more generally, not every equilibrium graph on $\mathbb{T}$ is reciprocal (Theorem 3.1). *Orthogonal* reciprocality depends on the *conformal* structure of $\mathbb{T}$, which is determined by the shape of its fundamental parallelogram. We derive a simple geometric condition that characterizes which equilibrium stresses are orthogonal reciprocal on $\mathbb{T}$ (Lemma 3.8). On the other hand, *parallel* reciprocality is an *affine* property: an equilibrium stress is parallel reciprocal on $\mathbb{T}$ if and only if it is parallel reciprocal on *every*

flat torus (Lemma 3.10).

- We show that for any equilibrium stress on Γ, there is a flat torus $\mathbb{T}'$, unique up to rotation and scaling of its fundamental parallelogram, such that the same equilibrium stress is orthogonal reciprocal for the affine image of Γ on $\mathbb{T}'$ (Theorem 3.3). In short, every equilibrium stress for Γ is orthogonal reciprocal on *some* flat torus. This result implies a natural toroidal analogue of Steinitz's theorem (Theorem 3.6): Every essentially 3-connected torus graph Γ is homotopic to a weighted Delaunay graph on some flat torus.
- We give an alternative geometric interpretation for when an equilibrium stress induces an (orthogonal or parallel) reciprocality. Given an equilibrium stress on Γ, we can consider the induced (orthogonal or parallel) reciprocal diagram $\widetilde{\Gamma}^*$ of the universal cover $\widetilde{\Gamma}$ of Γ; $\widetilde{\Gamma}^*$ is an infinite biperiodic planar graph drawing that can be projected down to *force diagram* on a *possibly different* flat torus $\mathbb{T}'$ (Theorems 3.14 and 3.15). An equilibrium stress induces an (orthogonal or parallel) reciprocality if and only if the force diagram in fact lives on the *same* flat torus, i.e. if and only if $\mathbb{T} = \mathbb{T}'$.
- Γ has an orthogonal reciprocal embedding if and only if Γ is coherent. Specifically, each orthogonal reciprocal diagram for Γ induces an essentially unique set of Delaunay weights for the vertices of Γ (Theorem 3.5). Conversely, each set of Delaunay weights for Γ induces a *unique* orthogonal reciprocal diagram Γ^*, namely the corresponding weighted Voronoi diagram (Lemma 3.16). Thus, unlike in the plane, a reciprocal diagram Γ^* may not be a weighted Voronoi diagram of the vertices of Γ, but some unique translation of Γ^* is.

3.1 History and Related Work

Building on earlier seminal work of Varignon [214], Rankine [177, 178], and others, Maxwell [156–158] proved that any straight-line planar graph Γ with an equilibrium stress has both a reciprocal diagram and a polyhedral lifting. In particular, positive and negative stresses correspond to convex and concave edges in the polyhedral lifting, respectively. Moreover, for any equilibrium stress λ on Γ, the vector $1/\lambda$ is an equilibrium stress for the reciprocal diagram Γ^*. Finally, for any polyhedral liftings of Γ, one can obtain a polyhedral lifting of the reciprocal diagram Γ^* via projective duality. Maxwell's analysis was later extended and popularized by Cremona [69, 70] and others; the correspondence has since been rediscovered several times in other contexts [18, 121]. More recently, Whiteley [220] proved the converse of Maxwell's theorem: every reciprocal diagram and every polyhedral lift corresponds to an equilibrium stress; see also Crapo and Whiteley [68]. For modern expositions of the Maxwell–

Cremona correspondence aimed at computational geometers, see Hopcroft and Kahn [119], Richter-Gebert [182, Chapter 13], or Rote, Santos, and Streinu [189].

Among many other consequences, combining the Maxwell–Cremona correspondence [220] with Tutte's spring-embedding theorem [213] yields an elegant geometric proof of Steinitz's theorem [198, 199] that every 3-connected planar graph is the 1-skeleton of a 3-dimensional convex polytope (Richter-Gebert [182, Chapter 13] provides an exposition of this proof). The Maxwell–Cremona correspondence has been used for scene analysis of planar drawings [18, 20, 68, 121, 204], finding small grid embeddings of planar graphs and polyhedra [54, 75, 81, 122, 172, 181, 182, 191], and several linkage reconfiguration problems [65, 74, 173, 202, 203].

It is natural to ask how or whether these correspondences extend to graphs on surfaces other than the Euclidean plane. Lovász [147, Lemma 4] describes a spherical analogue of Maxwell's polyhedral lifting in terms of Colin de Verdière matrices [59, 62]; see also [137]. Izmestiev [128] provides a self-contained proof of the correspondence for planar frameworks, along with natural extensions to frameworks in the sphere and the hyperbolic plane. Finally, and most closely related to the contents of this chapter, Borcea and Streinu [38], building on their earlier study of rigidity in infinite periodic frameworks [36, 37], develop an extension of the Maxwell–Cremona correspondence to infinite periodic graphs in the plane, or equivalently, to geodesic drawings on the Euclidean flat torus. Specifically, Borcea and Streinu prove that *periodic* polyhedral liftings correspond to *periodic* stresses satisfying an additional homological constraint.[1]

Our weighted Delaunay graphs are (the duals of) *power diagrams* [19, 21] or *Laguerre-Voronoi diagrams* [123] in the intrinsic metric of the flat torus. Toroidal Delaunay triangulations are commonly used to generate finite-element meshes for simulations with periodic boundary conditions, and several efficient algorithms for constructing these triangulations are known [35, 46, 112, 160]. Building on earlier work of Rivin [183] and Indermitte *et al.* [124], Bobenko and Springborn [33] proved that on any piecewise-linear surface, intrinsic Delaunay triangulations can be constructed by an intrinsic incremental flipping algorithm, mirroring the classical planar algorithm of Lawson [142]; their analysis extends easily to intrinsic weighted Delaunay graphs. Weighted Delaunay complexes are also known as *regular* or *coherent* subdivisions [72, 222].

Finally, equilibrium and reciprocal embeddings are closely related to the celebrated Koebe-Andreev-Thurston circle-packing theorem: Every planar graph is the contact graph of a set of interior-disjoint circular disks [10, 11, 135]; see Felsner and Rote [95] for a simple proof, based in part on earlier work of Brightwell and Scheinerman [41] and Mohar [162]. The circle-

[1] Phrased in terms of toroidal frameworks, Borcea and Streinu consider only equilibrium stresses for which the corresponding reciprocal toroidal framework contains no essential cycles. The same condition was also briefly discussed by Crapo and Whiteley [68, Example 3.6].

packing theorem has been generalized to higher-genus surfaces by Colin de Verdière [58, 61] and Mohar [163, 164]. In particular, Mohar proves that any well-connected graph Γ on the torus is *homotopic* to an essentially unique circle packing *for a unique Euclidean metric* on the torus. This disk-packing representation immediately yields a weighted Delaunay graph, where the areas of the disks are the vertex weights. We revisit and extend this result in Section 3.5.

Discrete harmonic and holomorphic functions, circle packings, and intrinsic Delaunay triangulations have numerous applications in discrete differential geometry; we refer the reader to monographs by Crane [67], Lovász [149], and Stephenson [201].

3.2 An Example

Before diving into the technical details, let us first consider an example of some of the phenomena we will encounter. This example will also serve as a running example for details that are clarified or expanded on later in this chapter. Let Γ be the (unweighted) intrinsic Delaunay triangulation of the seven points $\binom{0}{0}, \binom{1/7}{3/7}, \binom{2/7}{6/7}, \binom{3/7}{2/7}, \binom{4/7}{5/7}, \binom{5/7}{1/7}, \binom{6/7}{4/7}$ on the square flat torus $\mathbb{T}_{\square}$, and let Γ^* be the corresponding intrinsic Voronoi diagram, as shown in Figure 3.1. The triangulation Γ is a highly symmetric geodesic embedding of the complete graph K_7; torus graphs isomorphic to Γ and Γ^* were studied in several early seminal works on combinatorial topology [116, 117, 161].

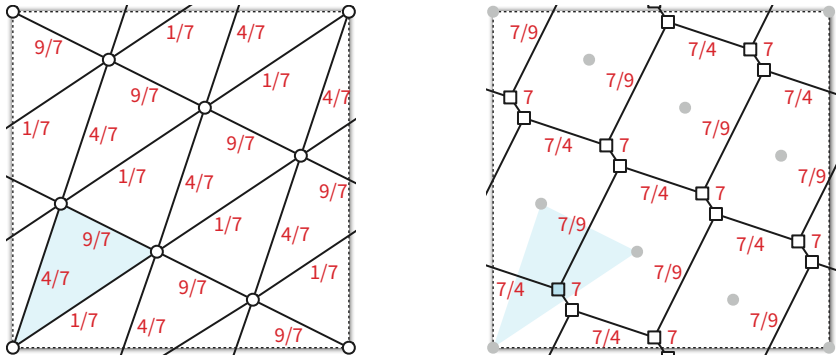

Figure 3.1. An intrinsic Delaunay triangulation, its dual Voronoi diagram, and their induced equilibrium stresses. Compare with Figures 3.2 and 3.4.

The edges of Γ fall into three equivalence classes, with slopes 3, $2/3$, $-1/2$ and lengths $\sqrt{10}/7$, $\sqrt{5}/7$, $\sqrt{14}/7$, respectively. The triangle $\binom{0}{0}, \binom{1/7}{3/7}, \binom{3/7}{2/7}$, shaded in Figure 3.1, has circumcenter $\binom{19/98}{17/98}$. Measuring slopes and distances to the nearby edge midpoints, we find that corresponding orthogonal dual edges in Γ^* have slopes $-1/3$, $-3/2$, and 2 and lengths $4\sqrt{10}/49$, $\sqrt{5}/49$, and $9\sqrt{14}/49$, respectively. These dual slopes confirm that Γ and Γ^* are reciprocal (as are any Delaunay triangulation and its dual Voronoi diagram; see Section 3.4).

Lemma 3.1 will imply that the dual edge lengths induce stress coefficients 4/7, 1/7, and 9/7 to the edges of Γ, and yield an equilibrium stress for Γ; symmetrically, the stress coefficients 7/4, 7, and 9/7 yield an equilibrium stress for Γ^*.

Of course, this is not the only equilibrium stress for Γ; indeed, symmetry implies that Γ is in equilibrium with respect to the uniform stress $\lambda \equiv 1$. However, there is no reciprocal embedding, orthogonal or parallel, Γ^* such that every edge in Γ has the same length as the corresponding dual edge in Γ^*.

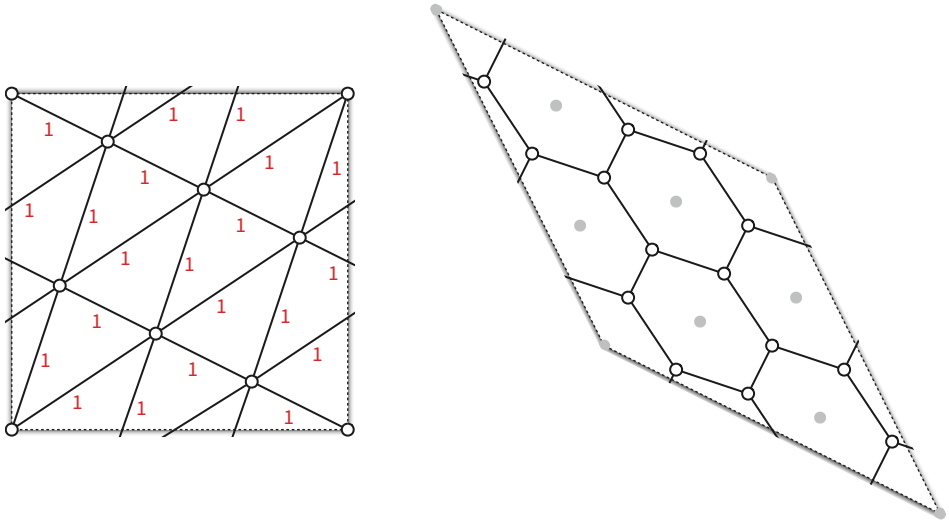

Figure 3.2. An "orthogonal reciprocal" embedding (at half scale) induced by the uniform equilibrium stress $\lambda \equiv 1$. Compare with Figures 3.1 and 3.4.

The doubly-periodic universal cover $\widetilde{\Gamma}$ is also in equilibrium with respect to the uniform stress $\lambda \equiv 1$. Thus, the classical Maxwell–Cremona correspondence implies an embedding of the dual graph $(\widetilde{\Gamma})^*$ in which every dual edge is orthogonal to and has the same length as its corresponding primal edge in $\widetilde{\Gamma}$. (Borcea and Streinu [38, Proposition 2] discuss how to solve the infinite linear system giving the heights of the corresponding polyhedral lifting of $\widetilde{\Gamma}$.) Symmetry implies that $(\widetilde{\Gamma})^*$ is doubly-periodic. Crucially, however, $\widetilde{\Gamma}$ and $(\widetilde{\Gamma})^*$ have *different period lattices*. Specifically, the period lattice of $(\widetilde{\Gamma})^*$ is generated by the vectors $\binom{2}{-1}$ and $\binom{-1}{2}$; see Figure 3.2.

Understanding which equilibrium stresses correspond to reciprocal embeddings is the topic of Section 3.3.2. In particular, in that section we describe simple necessary and sufficient conditions for an equilibrium stress to be orthogonal or parallel reciprocal, which the unit stress for Γ fails. Section 3.3.3 will then provide an interpretation of the different period lattices.

3.3 Toroidal Maxwell–Cremona...

Here we consider two different cases of the general Maxwell–Cremona correspondence, and how they generalize to flat tori. We first consider *orthogonal* reciprocal diagrams, as drawn by Maxwell. Maxwell [156–158] derived his reciprocal diagrams with dual edges drawn orthogonally to their corresponding primal edges so as to obtain the equivalence to polyhedral liftings of Γ in $\mathbb{R}^3$: Γ is the projection of this polyhedral lift, and Γ^* is the projection of its polar dual around the unit paraboloid. But following Rankine beforehand [177, 178] and Cremona afterwards [69, 70], we can also ask about what happens for for *parallel* reciprocal diagrams.

3.3.1 Reciprocal Implies Equilibrium

Lemma 3.1. Let Γ and Γ^* be (orthogonal or parallel) reciprocal geodesic drawings on some flat torus $\mathbb{T}_M$. The vector λ defined by $\lambda_e = |e^*|/|e|$ is an equilibrium stress for Γ; symmetrically, the vector λ^* defined by $\lambda^*_{e^*} = 1/\lambda_e = |e|/|e^*|$ is an equilibrium stress for Γ^*.

Proof. Let $\lambda_e = |e^*|/|e|$ and $\lambda^*_{e^*} = 1/\lambda_e = |e|/|e^*|$ for each edge e. Let Δ denote the $2 \times E$ displacement matrix of Γ, and let Δ^* denote the (transposed) displacement matrix of Γ^*. We immediately have $\Delta^*_{e^*} = \lambda_e \Delta^\perp_e$ for every edge e of Γ. The darts leaving each vertex p of Γ dualize to a cycle around the corresponding "face" u^* of Γ^*, and thus

$$\left(\sum_{v\colon uv \in E} \lambda_{uv} \Delta_{u\to v} \right)^\perp = \sum_{v\colon uv\in E} \lambda_{uv} \Delta^\perp_{u\to v} = \sum_{v\colon uv\in E} \Delta^*_{(u\to v)^*} = (0,0)\,.$$

We conclude that λ is an equilibrium stress for Γ, and thus (by swapping the roles of Γ and Γ^*) that λ^* is an equilibrium stress for Γ^*. □

A symmetric stress vector λ is an ***orthogonal reciprocal stress*** for Γ if there is an orthogonal reciprocal drawing Γ^* on the same flat torus such that $\lambda_e = |e^*|/|e|$ for each edge e; similarly, a symmetric stress vector λ is a ***parallel reciprocal stress*** for Γ if there is a parallel reciprocal drawing Γ^* on the same flat torus such that $\lambda_e = |e^*|/|e|$ for each edge e. Thus, a geodesic torus graph is orthogonal (parallel) reciprocal if and only if it has an orthogonal (parallel) reciprocal stress, and Lemma 3.1 implies that every orthogonal or parallel reciprocal stress is an equilibrium stress. The following simple construction shows that the converse of Lemma 3.1 is false.

Theorem 3.1. Not every positive equilibrium stress for Γ is an orthogonal reciprocal stress. More generally, not every equilibrium graph on $\mathbb{T}$ is reciprocal/coherent on $\mathbb{T}$.

Proof. Let Γ_1 be the geodesic triangulation in the flat square torus $\mathbb{T}_\square$ with a single vertex p and three edges, whose reference darts have displacement vectors $\binom{1}{0}$, $\binom{1}{1}$, and $\binom{2}{1}$. Every stress λ in Γ is an equilibrium stress, because the forces applied by each edge cancel out. The weighted Delaunay graph of a single point is identical for all weights, so it suffices to verify that Γ_1 is not an intrinsic Delaunay triangulation. We easily observe that the longest edge of Γ_1 is not Delaunay. See Figure 3.3.

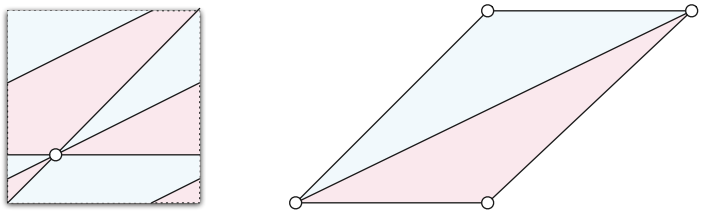

Figure 3.3. A one-vertex triangulation Γ_1 on the square flat torus, and a lift of its faces to the universal cover. Every stress in Γ_1 is an equilibrium stress, but Γ_1 is not a (weighted) intrinsic Delaunay triangulation.

More generally, for any positive integer q, let Γ_q denote the $q \times q$ covering of G_1. The vertices of G_q form a regular $q \times q$ square toroidal lattice, and the edges of Γ_q fall into three parallel families, with displacement vectors $\binom{1/q}{1/q}$, $\binom{2/q}{1/q}$, and $\binom{1/q}{0}$. Every positive stress vector where all parallel edges have equal stress coefficients is an equilibrium stress. We will show that G_q cannot be coherent; by Theorem 3.5, we conclude that G_q cannot be reciprocal.

For the sake of argument, suppose Γ_q is coherent. Let $i \to k$ be any dart with displacement vector $\binom{2/q}{1/q}$, and let j and ℓ be the vertices before and after j in clockwise order around i. The local Delaunay determinant test (2.6) implies that the weights of these four vertices satisfy the inequality $\pi_i + \pi_k + 1 < \pi_j + \pi_\ell$. Every vertex of Γ_q appears in exactly four inequalities of this form—twice on the left and twice on the right—so summing all q^2 such inequalities and canceling equal terms yields the obvious contradiction $1 < 0$. □

3.3.2 Equilibrium Implies Reciprocal, Sometimes

Now fix an essentially simple, essentially 3-connected geodesic drawing Γ on the *square* flat torus $\mathbb{T}_\square$, along with an equilibrium stress λ for Γ. In this section, we describe simple necessary and sufficient conditions for λ to be a reciprocal stress for Γ.

More generally, in the case of orthogonal reciprocality, we give necessary and sufficient conditions under which there is an essentially unique flat torus $\mathbb{T}_M$ such that a unique scalar multiple of λ is an orthogonal reciprocal stress for the image of Γ on $\mathbb{T}_M$. In the case of parallel reciprocality, however, we find that λ is a parallel reciprocal stress on $\mathbb{T}_\square$ if and only if it is parallel reciprocal stress on *every* flat torus. This is is perhaps unsurprising: orthogonality

relies on the conformal structure of the flat torus $\mathbb{T}$; parallelism is an affine property.

Let Δ be the $2 \times E$ displacement matrix of Γ, and let Λ be the $E \times E$ matrix whose diagonal entries are $\Lambda_{e,e} = \lambda_e$ and whose off-diagonal entries are all 0. The results in this section are phrased in terms of the covariance matrix $\Delta\Lambda\Delta^T = \left(\begin{smallmatrix} \alpha & \gamma \\ \gamma & \beta \end{smallmatrix}\right)$, where

$$\alpha = \sum_e \lambda_e \Delta x_e^2, \qquad \beta = \sum_e \lambda_e \Delta y_e^2, \qquad \gamma = \sum_e \lambda_e \Delta x_e \Delta y_e. \tag{3.1}$$

Recall that $A^\perp = (JA)^T$.

3.3.2.1 The Square Flat Torus

Before considering arbitrary flat tori, as a warmup we first establish necessary and sufficient conditions for λ to be a reciprocal stress for Γ on the *square* flat torus $\mathbb{T}_\square$, in terms of the parameters α, β, and γ. Although the results are identical for both orthogonal and parallel reciprocal stresses, the proofs are subtly different; for purposes of completeness we will present both versions.

Let us first consider the case of orthogonal reciprocal stresses.

Lemma 3.2. If λ is an orthogonal reciprocal stress for Γ on $\mathbb{T}_\square$, then $\Delta\Lambda\Delta^T = I$.

Proof. Suppose λ is an orthogonal reciprocal stress for Γ on $\mathbb{T}_\square$. Then there is a geodesic drawing of the dual graph Γ^* on $\mathbb{T}_\square$ where $e \perp e^*$ and $|e^*| = \lambda_e |e|$ for every edge e of Γ. Let $\Delta^* = (\Delta\Lambda)^\perp$ denote the $E \times 2$ matrix whose rows are the displacement row vectors of Γ^*.

Recall from Lemma 2.3 that the first and second rows of Ξ_Γ describe cocirculations of Γ with cohomology classes $(0,1)$ and $(-1,0)$, respectively. Applying Lemma 2.1 to Γ^* implies $\theta\Delta^* = [\theta]^*$ for any cocirculation θ in Γ. It follows immediately that $\Xi_\Gamma\Delta^* = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right) = -J$.

Because the rows of Δ^* are the displacement vectors of Γ^*, for every vertex p of Γ we have

$$\sum_{v\,:\,uv \in E} \Delta^*_{(u\to v)^*} = \sum_{d\,:\,\mathrm{tail}(d)=u} \Delta^*_{d^*} = \sum_{d\,:\,\mathrm{left}(d^*)=u^*} \Delta^*_{d^*} = (0,0). \tag{3.2}$$

It follows that the *columns* of Δ^* describe circulations in Γ. Lemma 2.1 now implies that $\Delta\Delta^* = -J$. We conclude that $\Delta\Lambda\Delta^T = \Delta\Delta^* J = I$. □

Lemma 3.3. Fix an $E \times 2$ matrix Δ^*. If $\Xi_\Gamma\Delta^* = -J$, then Δ^* is the displacement matrix of a geodesic drawing on $\mathbb{T}_\square$ that is dual to Γ. Moreover, if that drawing has a positive equilibrium stress, it is actually an embedding.

Proof. Let ξ_1 and ξ_2 denote the rows of Ξ_Γ. Rewriting the identity $\Xi_\Gamma \Delta^* = -J$ in terms of these row vectors gives us $\sum_e \Delta_e^* \lambda_{1,e} = (0,1) = [\xi_1]^*$ and $\sum_e \Delta_e^* \lambda_{2,e} = (-1,0) = [\xi_2]^*$. Extending by linearity, we have $\sum_e \Delta_e^* \theta_e = [\theta]^*$ for every cocirculation θ in Γ^*. The result now follows from Lemma 2.5. □

Lemma 3.4. If $\Delta\Lambda\Delta^T = I$, then λ is an orthogonal reciprocal stress for Γ on $\mathbb{T}_\square$. If λ is a *positive* equilibrium stress, then the orthogonal reciprocal diagram is in fact *embedded* on $\mathbb{T}_\square$

Proof. Set $\Delta^* = (\Delta\Lambda)^\perp$. Because λ is an equilibrium stress in Γ, for every vertex p of Γ we have

$$\sum_{v\,:\,uv\in E} \Delta^*_{(u\to v)^*} \;=\; \sum_{v\,:\,uv\in E} \lambda_{uv}\Delta^\perp_{u\to v} \;=\; (0,0). \tag{3.3}$$

It follows that the columns of Δ^* describe circulations in Γ, and therefore Lemma 2.1 implies $\Xi_\Gamma\Delta^* = \Delta\Delta^* = \Delta(\Delta\Lambda)^\perp = \Delta\Lambda\Delta^T J^T = -J$.

Lemma 3.3 now implies that Δ^* is the displacement matrix of a drawing Γ^* dual to Γ. Moreover, the stress vector λ^* defined by $\lambda^*_{e^*} = 1/\lambda_e$ is an equilibrium stress for Γ^*: under this stress vector, the darts leaving any dual vertex f^* are dual to the clockwise boundary cycle of face f in Γ. Thus if λ is positive, then Γ^* is in fact an embedding. By construction, each edge of Γ^* is orthogonal to the corresponding edge of Γ. □

Next, we consider the case of parallel reciprocal stresses.

Lemma 3.5. If λ is a parallel reciprocal stress for Γ on $\mathbb{T}_\square$, then $\Delta\Lambda\Delta^T = I$.

Proof. Suppose λ is a parallel reciprocal stress for Γ on $\mathbb{T}_\square$. Then there is a geodesic embedding of the dual graph Γ^* on $\mathbb{T}_\square$ where $e \parallel e^*$ and $|e^*| = \lambda_e |e|$ for every edge e of Γ. Let $\Delta^* = (\Delta\Lambda)^T$ denote the $E \times 2$ matrix whose rows are the displacement row vectors of Γ^*.

Recall from Lemma 2.4 that the first and second rows of Ξ_Γ describe cocirculations of Γ with cohomology classes $(1,0)$ and $(0,1)$, respectively. Applying Lemma 2.1 to Γ^* implies $\theta\Delta^* = [\theta]^*$ for any cocirculation θ in Γ. It follows immediately that $\Xi_\Gamma\Delta^* = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right) = I$.

Because the rows of Δ^* are the displacement vectors of Γ^*, for every vertex p of Γ we have

$$\sum_{v\,:\,uv\in E} \Delta^*_{(u\to v)^*} \;=\; \sum_{d\,:\,\mathrm{tail}(d)=u} \Delta^*_{d^*} \;=\; \sum_{d\,:\,\mathrm{left}(d^*)=u^*} \Delta^*_{d^*} \;=\; (0,0). \tag{3.4}$$

It follows that the *columns* of Δ^* describe circulations in Γ. Lemma 2.1 now implies that $\Delta\Delta^* = \Delta\Lambda\Delta^T = I$. □

Lemma 3.6. Fix an $E \times 2$ matrix Δ^*. If $\Xi_\Gamma\Delta^* = I$, then Δ^* is the displacement matrix of a geodesic drawing on $\mathbb{T}_\square$ that is dual to Γ. Moreover, if that drawing has an equilibrium stress, it is actually an embedding.

Proof. Let ξ_1 and ξ_2 denote the rows of Ξ_Γ. Rewriting the identity $\Xi_\Gamma \Delta^* = I$ in terms of these row vectors gives us $\sum_e \Delta_e^* \lambda_{1,e} = (1,0) = [\xi_1]^*$ and $\sum_e \Delta_e^* \lambda_{2,e} = (0,1) = [\xi_2]^*$. Extending by linearity, we have $\sum_e \Delta_e^* \theta_e = [\theta]^*$ for every cocirculation θ in Γ^*. The result now follows from Lemma 2.5. □

Lemma 3.7. If $\Delta\Lambda\Delta^T = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$, then λ is a parallel reciprocal stress for Γ on $\mathbb{T}_\square$.

Proof. Set $\Delta^* = (\Delta\Lambda)^T$. Because λ is an equilibrium stress in Γ, for every vertex p of Γ we have

$$\sum_{v\,:\,uv\in E} \Delta^*_{(u\to v)^*} \;=\; \sum_{v\,:\,uv\in E} \lambda_{uv}\Delta^T_{u\to v} \;=\; (0,0)\,. \tag{3.5}$$

It follows that the columns of Δ^* describe circulations in Γ, and therefore Lemma 2.1 implies $\Xi_\Gamma\Delta^* = \Delta\Delta^* = \Delta(\Delta\Lambda)^T = \Delta\Lambda\Delta^T = I$.

Lemma 3.6 now implies that Δ^* is the displacement matrix of a drawing Γ^* dual to Γ. Moreover, the stress vector λ^* defined by $\lambda^*_{e^*} = 1/\lambda_e$ is an equilibrium stress for Γ^*: under this stress vector, the darts leaving any dual vertex f^* are dual to the clockwise boundary cycle of face f in Γ. Thus if λ is positive, then Γ^* is in fact an embedding. By construction, each edge of Γ^* is parallel to the corresponding edge of Γ. □

In conclusion:

Theorem 3.2. λ is an orthogonal and parallel reciprocal stress for Γ on $\mathbb{T}_\square$ if and only if $\Delta\Lambda\Delta^T = I$. Moreover, if λ is a positive equilibrium stress, then the corresponding reciprocal diagrams are in fact *embedded* on $\mathbb{T}_\square$.

3.3.2.2 Orthogonal Reciprocality

Now we generalize our previous analysis for orthogonal reciprocality to drawings on the flat torus $\mathbb{T}_M$ defined by an arbitrary non-singular matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. These results are also stated in terms of the parameters α, β, and γ, which are still defined in terms of $\mathbb{T}_\square$, which will serve as a *reference* flat torus when talking about flat tori defined by different non-singular matrices.

Lemma 3.8. If λ is a reciprocal stress for a geodesic drawing Γ on $\mathbb{T}_M$, then $\alpha\beta - \gamma^2 = 1$; in particular, if $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, then

$$\alpha = \frac{b^2+d^2}{ad-bc}, \qquad \beta = \frac{a^2+c^2}{ad-bc}, \qquad \gamma = \frac{-(ab+cd)}{ad-bc}.$$

For example, if $M = (\rho, \sigma)$ where $\rho, \sigma \in \mathbb{R}^2$ are column vectors and $\det M = 1$, then $\Delta\Lambda\Delta^T = \left(\begin{smallmatrix} \sigma\cdot\sigma & -\rho\cdot\sigma \\ -\rho\cdot\sigma & \rho\cdot\rho \end{smallmatrix}\right)$.

Proof. Suppose λ is a reciprocal stress for Γ on $\mathbb{T}_M$. Then there is a geodesic embedding of the dual graph Γ^* on $\mathbb{T}_M$ where $e \perp e^*$ and $|e^*| = \lambda_e |e|$ for every edge e of Γ.

It will prove convenient to consider the geometry of Γ and Γ^* on the reference torus $\mathbb{T}_\square$. (The drawings of Γ and Γ^* on the reference torus $\mathbb{T}_\square$ are still dual, but not necessarily reciprocal.) Let Δ denote the $2 \times E$ *reference* displacement matrix for Γ, whose columns are the displacement vectors for Γ on the square torus $\mathbb{T}_\square$. Then the columns of $M\Delta$ are the *native* displacement vectors for Γ on the torus $\mathbb{T}_M$. Thus, the *native* displacement row vectors of Γ^* are given by the rows of the $E \times 2$ matrix $(M\Delta\Lambda)^\perp$. Finally, let $\Delta^* = (M\Delta\Lambda)^\perp (M^T)^{-1}$ denote the *reference* displacement row vectors for Γ^* on the square torus $\mathbb{T}_\square$. We can rewrite this definition as

$$\Delta^* = (M\Delta\Lambda)^\perp (M^T)^{-1} = (M\Delta\Lambda)^\perp (M^{-1})^T = (JM\Delta\Lambda)^T (M^{-1})^T = \Lambda\Delta^T M^T J^T (M^{-1})^T, \quad (3.6)$$

which implies $\Lambda\Delta^T = \Delta^* M^T J (M^{-1})^T$.

Because the rows of Δ^* are the displacement vectors for Γ^*, equation (3.2) implies that the *columns* of Δ^* describe circulations in Γ, and therefore $\Delta\Delta^* = \Xi_\Gamma \Delta^* = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right) = -J$ by Lemmas 2.1 and 2.3. We conclude that

$$\begin{aligned}
\Delta\Lambda\Delta^T &= \Delta\Delta^* M^T J (M^{-1})^T = J^T M^T J (M^{-1})^T \\
&= \frac{1}{ad-bc} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} a & c \\ b & d \end{pmatrix} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} d & -c \\ -b & a \end{pmatrix} \\
&= \frac{1}{ad-bc} \begin{pmatrix} b & d \\ -a & -c \end{pmatrix} \begin{pmatrix} b & -a \\ d & -c \end{pmatrix} \\
&= \frac{1}{ad-bc} \begin{pmatrix} b^2+d^2 & -ab-cd \\ -ab-cd & a^2+c^2 \end{pmatrix}.
\end{aligned}$$

Routine calculation now implies that $\alpha\beta - \gamma^2 = \det \Delta\Lambda\Delta^T = 1$. □

Corollary 3.1. If λ is a reciprocal stress for Γ on $\mathbb{T}_M$, then $M = \sigma R \left(\begin{smallmatrix} \beta & -\gamma \\ 0 & 1 \end{smallmatrix}\right)$ for some 2×2 rotation matrix R and some real number $\sigma > 0$.

Proof. Reciprocality is preserved by rotating and scaling the fundamental parallelogram $\Diamond_M$, so it suffices to consider the special case $M = \left(\begin{smallmatrix} a & b \\ 0 & 1 \end{smallmatrix}\right)$. In this special case, Lemma 3.8 immediately implies $\beta = a$ and $\gamma = -b$. □

Lemma 3.9. If $\alpha\beta - \gamma^2 = 1$, then λ is a reciprocal stress for Γ on $\mathbb{T}_M$ where $M = \sigma R \left(\begin{smallmatrix} \beta & -\gamma \\ 0 & 1 \end{smallmatrix}\right)$ for any 2×2 rotation matrix R and any real number $\sigma > 0$.

Proof. Suppose $\alpha\beta - \gamma^2 = 1$. Fix an arbitrary 2×2 rotation matrix R and an arbitrary real

number $\sigma > 0$, and let $M = \sigma R\left(\begin{smallmatrix}\beta & -\gamma \\ 0 & 1\end{smallmatrix}\right)$. Let Δ denote the $2 \times E$ *reference* displacement matrix for Γ on the square flat torus $\mathbb{T}_\square$, and define the $E \times 2$ matrix $\Delta^* = (M\Delta\Lambda)^\perp (M^T)^{-1}$.

Derivation (3.6) in the proof of Lemma 3.8 implies $\Delta^* = \Lambda\Delta^T (M^{-1}JM)^T$. We easily observe that $(\sigma R)^{-1} J (\sigma R) = J$, and therefore

$$\begin{aligned} M^{-1}JM &= \begin{pmatrix}\beta & -\gamma \\ 0 & 1\end{pmatrix}^{-1} \begin{pmatrix}0 & -1 \\ 1 & 0\end{pmatrix} \begin{pmatrix}\beta & -\gamma \\ 0 & 1\end{pmatrix} \\ &= \frac{1}{\beta} \begin{pmatrix}1 & \gamma \\ 0 & \beta\end{pmatrix} \begin{pmatrix}0 & -1 \\ 1 & 0\end{pmatrix} \begin{pmatrix}\beta & -\gamma \\ 0 & 1\end{pmatrix} \\ &= \frac{1}{\beta} \begin{pmatrix}\beta\gamma & -1-\gamma^2 \\ \beta^2 & -\beta\gamma\end{pmatrix} \\ &= \begin{pmatrix}\gamma & -\alpha \\ \beta & -\gamma\end{pmatrix}. \end{aligned}$$

It follows that

$$\Delta\Delta^* \;=\; \Delta\Lambda\Delta^T (M^{-1}JM)^T \;=\; \begin{pmatrix}\alpha & \gamma \\ \gamma & \beta\end{pmatrix} \begin{pmatrix}\gamma & \beta \\ -\alpha & -\gamma\end{pmatrix} \;=\; \begin{pmatrix}0 & \alpha\beta-\gamma^2 \\ \gamma^2-\alpha\beta & 0\end{pmatrix} \;=\; -J.$$

Because λ is an equilibrium stress in Γ, for every vertex p of Γ we have

$$\sum_{v\colon uv \in E} \Delta^*_{(u\to v)^*} \;=\; \sum_{v\colon uv\in E} \lambda_{uv} \Delta^\perp_{u\to v} (M^{-1}JMJ^T)^T \;=\; (0,0)(M^{-1}JMJ^T)^T \;=\; (0,0). \tag{3.7}$$

Once again, the columns of Δ^* describe circulations in Γ, so Lemma 2.1 implies $\Xi_\Gamma \Delta^* = \Delta\Delta^* = -J$. Lemma 3.3 now implies that Δ^* is the displacement matrix of a homotopic embedding of Γ^* on $\mathbb{T}_\square$. It follows that $(M\Delta\Lambda)^\perp = \Delta^* M^T$ is the displacement matrix of the image of Γ^* on $\mathbb{T}_M$. By construction, each edge of Γ^* is orthogonal to its corresponding edge of Γ. We conclude that λ is a reciprocal stress for Γ. □

The preceding lemmas directly imply the following theorem:

Theorem 3.3. Let Γ be a geodesic drawing on $\mathbb{T}_\square$ homotopic to an embedding with equilibrium stress λ. Let α, β, and γ be defined as in Equation (3.1). If $\alpha\beta - \gamma^2 = 1$, then λ is a reciprocal stress for the image of Γ on $\mathbb{T}_M$ if and only if $M = \sigma R\left(\begin{smallmatrix}\beta & -\gamma \\ 0 & 1\end{smallmatrix}\right)$ for any rotation matrix R and any real number $\sigma > 0$. On the other hand, if $\alpha\beta - \gamma^2 \neq 1$, then λ is not a reciprocal stress for the image of Γ on any flat torus.

When λ is positive, then $\alpha\beta - \gamma^2 = \frac{1}{2}\sum_{e,e'} \lambda_e \lambda_{e'} \left| \begin{smallmatrix} \Delta x_e & \Delta y_e \\ \Delta x_{e'} & \Delta y_{e'} \end{smallmatrix} \right|^2 > 0$, so in fact the requirement $\alpha\beta - \gamma^2 = 1$ is just a scaling condition: given any positive stress vector λ, the stress vector $\lambda / \sqrt{\alpha\beta - \gamma^2}$ is a positive stress vector that satisfies said requirement. If λ is non-positive, however, it is possible that $\alpha\beta - \gamma^2 < 0$, in which case no scaling of λ is an orthogonal reciprocal stress on any flat torus.

3.3.2.3 Example

Let us revisit once more the example drawing Γ from Section 3.2: the symmetric embedding of K_7 on the square flat torus $\mathbb{T}_\square$. Symmetry implies that Γ is in equilibrium with respect to the uniform stress $\lambda \equiv 1$. Straightforward calculation gives us the parameters $\alpha = \beta = 2$ and $\gamma = 1$ for this stress vector. Thus, Lemma 3.2 immediately implies that λ is not a reciprocal stress for Γ; rather, by Lemma 3.12, the force diagram of Γ with respect to λ lies on the torus $\mathbb{T}_M$, where $M = \left(\begin{smallmatrix} \beta & -\gamma \\ -\gamma & \alpha \end{smallmatrix} \right) = \left(\begin{smallmatrix} 2 & -1 \\ -1 & 2 \end{smallmatrix} \right)$. Moreover, because $\alpha\beta - \gamma^2 = 3 \neq 1$, Lemma 3.8 implies that λ is not an orthogonal reciprocal stress for the image of Γ on *any* flat torus. In short, there are no orthogonal reciprocal embeddings of Γ and Γ^* on *any* flat torus such that corresponding primal and dual edges have equal length.

Now consider the scaled uniform stress $\lambda \equiv 1/\sqrt{3}$, which has parameters $\alpha = \beta = 2/\sqrt{3}$ and $\gamma = 1/\sqrt{3}$. This new stress λ is still not an orthogonal reciprocal stress for Γ; however, it does satisfy the scaling constraint $\alpha\beta - \gamma^2 = 1$. Lemma 3.8 (also Lemma 3.14) implies that λ is an orthogonal reciprocal stress for the image of Γ on the flat torus $\mathbb{T}_M$, where $M = \frac{1}{\sqrt{3}} \left(\begin{smallmatrix} 2 & -1 \\ 0 & \sqrt{3} \end{smallmatrix} \right)$. The fundamental parallelogram $\lozenge_M$ is the union of two equilateral triangles with height 1. Not surprisingly, the image of Γ on $\mathbb{T}_M$ is a Delaunay triangulation with equilateral triangle faces, and the faces of the reciprocal Voronoi diagram Γ^* (which is also the force diagram) are regular hexagons. Finally, the vector $\lambda^* \equiv \sqrt{3}$ is a reciprocal stress, and therefore an equilibrium stress, for Γ^*. See Figure 3.4.

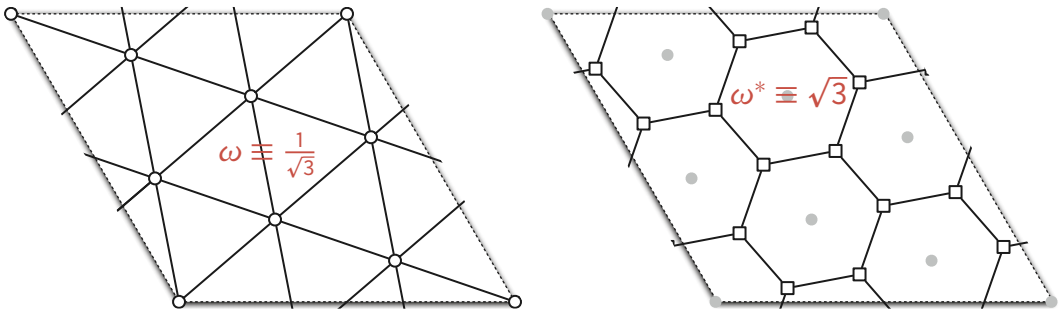

Figure 3.4. A seven-vertex Delaunay triangulation and its dual Voronoi diagram, induced by the uniform stress $1/\sqrt{3}$; compare with Figures 3.1 and 3.2.

Let us now consider a stress vector on Γ that is not strictly positive. Assigning the edges of slope 3 a stress of 2, the edges of slope $2/3$ a stress of -1, and the edges of slope $-1/2$ a stress of 3, we can verify that this indeed induces an equilibrium stress, and furthermore $\alpha\beta - \gamma^2 = 1$, so we can find a 2×2 matrix M such that the image of Γ on $\mathbb{T}_M$ has an orthogonal reciprocal diagram. However, this orthogonal reciprocal diagram has coincident vertices and overlapping edges and self-intersecting faces; see Figure 3.6.

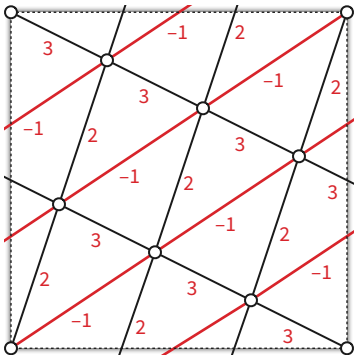

Figure 3.5. The symmetric embedding of K_7 with weights -1, 2, and 3.

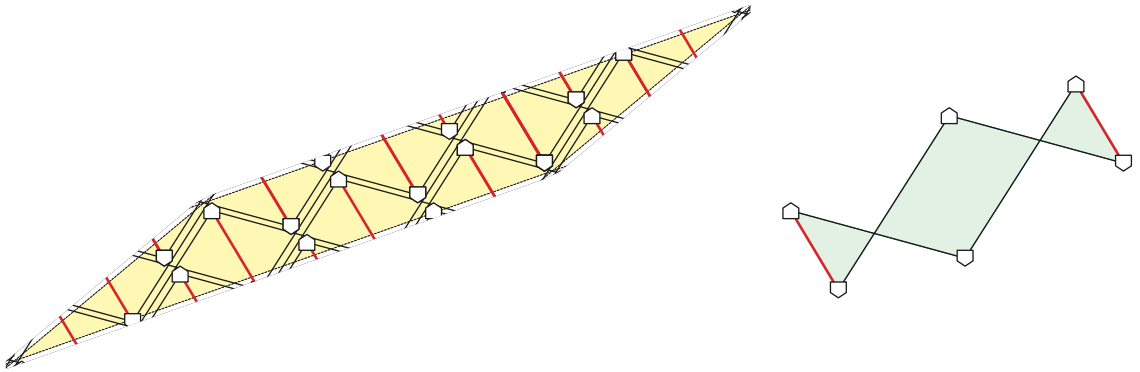

(a) (b)

Figure 3.6. (a) The reciprocal diagram of the image of K_7 on $\mathbb{T}_M$. Overlapping vertices and edges are drawn as being distinct, but vertices close to and pointing to each other actually occupy the same position. (b) One of the faces of the reciprocal diagram.

If we instead assign the edges of slope 3 and $2/3$ a stress of 1, and the edges of slope $-1/2$ a stress of -1, then this is an equilibrium stress, but $\alpha\beta - \gamma^2 = -1$, so no scaling of this stress vector can be an orthogonal reciprocal stress for Γ on any flat torus.

3.3.2.4 Parallel Reciprocality

Now we consider parallel reciprocality on the flat torus $\mathbb{T}_M$ defined by an arbitrary non-singular matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. As with the case of orthogonal reciprocality, these results are

stated in terms of the parameters α, β, and γ, which are still defined in terms of $\mathbb{T}_\square$, which will serve as a *reference* flat torus when talking about flat tori defined by different non-singular matrices.

Lemma 3.10. If λ is a parallel reciprocal stress for a geodesic graph Γ on $\mathbb{T}_M$ for some non-singular matrix M, then $\Delta\Lambda\Delta^T = I$.

Proof. Suppose λ is a parallel reciprocal stress for Γ on $\mathbb{T}_M$. Then there is a geodesic drawing of the dual graph Γ^* on $\mathbb{T}_M$ where $e \parallel e^*$ and $|e^*| = \lambda_e |e|$ for every edge e of Γ.

We will consider the geometry of Γ and Γ^* on the reference torus $\mathbb{T}_\square$. (The drawings of Γ and Γ^* on the reference torus $\mathbb{T}_\square$ are still dual, but not necessarily reciprocal.) Let Δ denote the $2 \times E$ *reference* displacement matrix for Γ, whose columns are the displacement vectors for Γ on the square torus $\mathbb{T}_\square$. Then the columns of $M\Delta$ are the *native* displacement vectors for Γ on the torus $\mathbb{T}_M$. Thus, the *native* displacement row vectors of Γ^* are given by the rows of the $E \times 2$ matrix $(M\Delta\Lambda)^T$. Finally, let $\Delta^* = (M\Delta\Lambda)^T (M^T)^{-1}$ denote the *reference* displacement row vectors for Γ^* on the square torus $\mathbb{T}_\square$. We can rewrite this definition as

$$\begin{aligned} \Delta^* &= (M\Delta\Lambda)^T (M^T)^{-1} \\ &= \Lambda\Delta^T M^T (M^T)^{-1} \\ &= \Lambda\Delta^T. \end{aligned} \tag{3.8}$$

Because the rows of Δ^* are the displacement vectors for Γ^*, equation (3.4) implies that the *columns* of Δ^* describe circulations in Γ, and therefore $\Delta\Delta^* = \Lambda\Delta^* = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right) = I$ by Lemmas 2.1 and 2.4. We conclude that $\Delta\Lambda\Delta^T = \Delta\Delta^* = I$. □

Lemma 3.11. If $\Delta\Lambda\Delta^T = I$, then λ is a parallel reciprocal stress for Γ on $\mathbb{T}_M$ where M is any non-singular 2×2 matrix. Moreover, if λ is a *positive* equilibrium stress, then the reciprocal diagram is embedded on $\mathbb{T}_M$.

Proof. Suppose $\Delta\Lambda\Delta^T = I$. Fix an arbitrary 2×2 non-singular matrix M. Let Δ denote the $2 \times E$ *reference* displacement matrix for Γ on the square flat torus $\mathbb{T}_\square$, and define the $E \times 2$ matrix $\Delta^* = (M\Delta\Lambda)^T (M^T)^{-1}$.

Derivation (3.8) in the proof of Lemma 3.10 implies $\Delta^* = \Lambda\Delta^T$. It follows that

$$\Delta\Delta^* \;=\; \Delta\Lambda\Delta^T \;=\; I.$$

Because λ is an equilibrium stress in Γ, for every vertex p of Γ we have

$$\sum_{v\,:\,uv \in E} \Delta^*_{(u\to v)^*} \;=\; \sum_{v\,:\,uv \in E} \lambda_{uv} \Delta^T_{u\to v} \;=\; (0,0). \tag{3.9}$$

Once again, the columns of Δ^* describe circulations in Γ, so Lemma 2.1 implies $\Lambda\Delta^* = \Delta\Delta^* = I$. Lemma 3.6 now implies that Δ^* is the displacement matrix of a homotopic drawing of Γ^* on $\mathbb{T}_\square$, and if λ is positive, said drawing is in fact an embedding. It follows that $(M\Delta\Lambda)^T = \Delta^* M^T$ is the displacement matrix of the image of Γ^* on $\mathbb{T}_M$. By construction, each edge of Γ^* is parallel to its corresponding edge of Γ. We conclude that λ is a parallel reciprocal stress for Γ. □

The preceding lemmas directly imply the following theorem:

Theorem 3.4. Let Γ be a geodesic graph on $\mathbb{T}_\square$ with an equilibrium stress λ. If $\Delta\Lambda\Delta^T = I$, then λ is a parallel reciprocal stress for the image of Γ on $\mathbb{T}_M$ for any non-singular matrix M; furthermore, if λ is a *positive* equilibrium stress, then the parallel reciprocal diagram is embedded on $\mathbb{T}_M$. On the other hand, if $\Delta\Lambda\Delta^T \neq I$, then λ is not a parallel reciprocal stress for the image of Γ on any flat torus.

3.3.2.5 Example

Once again consider the symmetric embedding of K_7 on the square flat torus $\mathbb{T}_\square$ shown in Figure 3.2. Recall that symmetry implies that Γ is in equilibrium with respect to the uniform stress $\lambda \equiv 1$. Straightforward calculation gives us $\Delta\Lambda\Delta^T = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 2 \end{smallmatrix}\right)$ for this stress vector. Thus, Lemma 3.5 immediately implies that λ is not a parallel reciprocal stress for Γ. (We will see later on that Lemma 3.12 implies that the parallel force diagram of Γ with respect to λ lies on the torus $\mathbb{T}_M$, where $M = \Delta\Lambda\Delta^T = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 2 \end{smallmatrix}\right)$.)

Recall that the scaled uniform stress $\lambda' \equiv 1/\sqrt{3}$ is an orthogonal reciprocal stress for Γ on the flat torus $\mathbb{T}_M$ where $M = \frac{1}{\sqrt{3}}\left(\begin{smallmatrix} 2 & -1 \\ 0 & \sqrt{3} \end{smallmatrix}\right)$. On the other hand, Theorem 3.4 implies that $\lambda \equiv 1$ and its scalings are *never* parallel reciprocal stresses.

3.3.3 Force Diagrams

The results of the previous section have a more physical interpretation that may be more intuitive. We begin once again on the unit square flat torus $\mathbb{T}_\square$.

Let Γ be any geodesic drawing on the unit square flat torus $\mathbb{T}_\square$. Recall from Section 3.2 that any equilibrium stress λ on Γ induces a equilibrium stress on its universal cover $\widetilde{\Gamma}$, which in turn induces an orthogonal (parallel) reciprocal diagram $(\widetilde{\Gamma})^*$ by the classical Maxwell–Cremona correspondence. This infinite planar drawing $(\widetilde{\Gamma})^*$ is doubly-periodic, but in general with a different period lattice from the universal cover $\widetilde{\Gamma}$.

Said differently, we can always construct another geodesic toroidal drawing F that is combinatorially dual to Γ, such that for every edge e of Γ, the corresponding edge e^* of F is orthogonal

(parallel) to e and has length $\lambda_e \cdot |e|$; however, this torus graph F does not necessarily lie on the *square* flat torus. (By construction, F is the unique torus graph whose universal cover is $(\widetilde{\Gamma})^*$, the orthogonal (parallel) reciprocal diagram of the universal cover of Γ.) We call F the ***orthogonal (parallel) force diagram*** of Γ with respect to λ. The force diagram F lies on the same flat torus $\mathbb{T}_\square$ as Γ if and only if λ is an orthogonal (parallel) reciprocal stress for Γ.

Lemma 3.12. Let Γ be a geodesic drawing in $\mathbb{T}_\square$, and let λ be an equilibrium stress for Γ. The orthogonal force diagram of Γ with respect to λ lies on the flat torus $\mathbb{T}_M$, where $M = \left(\begin{smallmatrix} \beta & -\gamma \\ -\gamma & \alpha \end{smallmatrix}\right) = J\Delta\Lambda\Delta^T J^T$.

Proof. As usual, let Δ be the displacement matrix of Γ. Let Δ^* denote the displacement matrix of the force diagram F; by definition, we have $\Delta^* = (\Delta\Lambda)^\perp = \Lambda\Delta^T J^T$. Equation (3.3) implies that the columns of Δ^* are circulations in Γ. Thus, Lemma 2.1 implies that $\Xi_\Gamma \Delta^* = \Delta\Delta^* = \Delta\Lambda\Delta^T J^T$.

Set $M = J\Delta\Delta^* = J\Delta\Lambda\Delta^T J^T = \left(\begin{smallmatrix} \beta & -\gamma \\ -\gamma & \alpha \end{smallmatrix}\right)$. We immediately have $\Xi_\Gamma \Delta^* = J^{-1}M = -JM = -JM^T$ and therefore $\Xi_\Gamma \Delta^* (M^T)^{-1} = -J$. Lemma 3.3 implies that $\Delta^*(M^T)^{-1}$ is the displacement matrix of a homotopic embedding of Γ^* on $\mathbb{T}_\square$. It follows that Δ^* is the displacement matrix of the image of Γ^* on $\mathbb{T}_M$. We conclude that F is a translation of the image of Γ^* on $\mathbb{T}_M$. □

Lemma 3.12 thus provides an intuitive geometric picture for why Theorem 3.2 is true; see Figure 3.7.

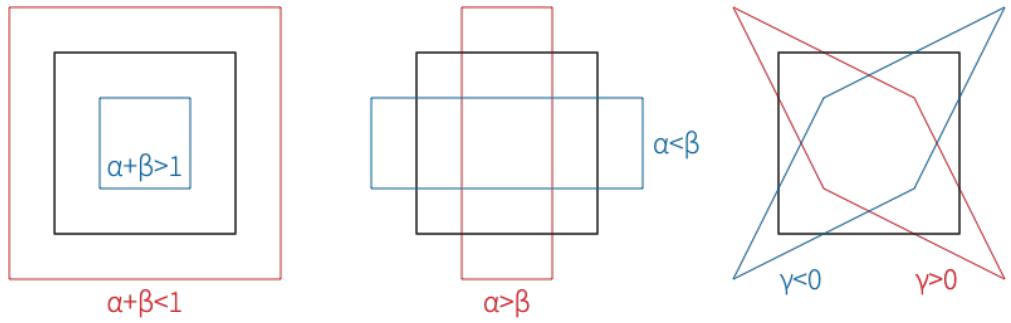

Figure 3.7. How the values α, β, γ independently affect the shape of orthogonal force diagram's native flat torus.

Lemma 3.13. Let Γ be a geodesic drawing in $\mathbb{T}_\square$, and let λ be a positive equilibrium stress for Γ. The parallel force diagram of Γ with respect to λ lies on the flat torus $\mathbb{T}_M$, where $M = \Delta\Lambda\Delta^T$.

Proof. As usual, let Δ be the displacement matrix of Γ. Let Δ^* denote the displacement matrix of the force diagram F; by definition, we have $\Delta^* = (\Delta\Lambda)^T = \Lambda\Delta^T$. Equation (3.5) implies that the columns of Δ^* are circulations in Γ. Thus, Lemma 2.1 implies that $\Lambda\Delta^* = \Delta\Delta^* = \Delta\Lambda\Delta^T$.

Set $M = \Delta\Delta^* = \Delta\Lambda\Delta^T$. We immediately have $\Lambda\Delta^* = M = M^T$ and therefore $\Lambda\Delta^*(M^T)^{-1} = I$. Lemma 3.6 implies that $\Delta^*(M^T)^{-1}$ is the displacement matrix of a homotopic embedding of G^* on $\mathbb{T}_\square$. It follows that Δ^* is the displacement matrix of the image of G^* on $\mathbb{T}_M$. We conclude that F is a translation of the image of G^* on $\mathbb{T}_M$. □

We will now extend these results to other flat tori. As before, we start with orthogonal force diagrams.

Lemma 3.14. Let Γ be a geodesic drawing on $\mathbb{T}_M$ homotopic to an embedding, and let λ be a equilibrium stress for Γ. The force diagram of Γ with respect to λ lies on the flat torus $\mathbb{T}_N$, where $N = JM\Delta\Lambda\Delta^T J^T$.

Proof. We argue exactly as in the proof of Lemma 3.12. Let Δ be the *reference* displacement matrix of (the image of) Γ on $\mathbb{T}_\square$. Then the *native* displacement matrix of the force diagram is $\Delta^* = (M\Delta\Lambda)^\perp = \Lambda\Delta^T M^T J^T$. Equation (3.7) and Lemma 2.1 imply that $\Xi_\Gamma\Delta^* = \Delta\Lambda\Delta^T M^T J^T$.

Now let $N = JM\Delta\Lambda\Delta^T J^T$. We immediately have $J^{-1}N^T = \Xi_\Gamma\Delta^*$ and thus $\Xi_\Gamma\Delta^*(N^T)^{-1} = J^{-1} = -J$. Lemma 3.3 implies that $\Delta^*(N^T)^{-1}$ is the displacement matrix of a homotopic embedding of Γ^* on $\mathbb{T}_\square$. It follows that Δ^* is the displacement matrix of the image of Γ^* on $\mathbb{T}_N$. □

Furthermore, when λ is positive, the force diagram is embedded on $\mathbb{T}_N$, and every face of the force diagram is a convex polygon; if λ is not necessarily positive, then the force diagram still lies on $\mathbb{T}_N$, but is not necessarily embedded, and faces may self-intersect.

We conclude by analyzing parallel force diagrams.

Lemma 3.15. Let Γ be a geodesic graph on $\mathbb{T}_M$, and let λ be a positive equilibrium stress for Γ. The parallel force diagram of Γ with respect to λ lies on the flat torus $\mathbb{T}_N$, where $N = M\Delta\Lambda\Delta^T$.

Proof. We argue exactly as in the proof of Lemma 3.13. Let Δ be the *reference* displacement matrix of (the image of) Γ on $\mathbb{T}_\square$. Then the *native* displacement matrix of the force diagram is $\Delta^* = (M\Delta\Lambda)^T = \Lambda\Delta^T M^T$. Equation (3.9) and Lemma 2.1 imply that $\Lambda\Delta^* = \Delta\Lambda\Delta^T M^T$.

Now let $N = M\Delta\Lambda\Delta^T$. We immediately have $N^T = \Lambda\Delta^*$ and thus $\Lambda\Delta^*(N^T)^{-1} = I$. Lemma 3.6 implies that $\Delta^*(N^T)^{-1}$ is the displacement matrix of a homotopic drawing of Γ^* on $\mathbb{T}_\square$. It follows that Δ^* is the displacement matrix of the image of Γ^* on $\mathbb{T}_N$. □

3.4 ...and Delaunay too

Unlike in the previous section, the equivalence between embeddings with orthogonal reciprocal diagrams and coherent embeddings generalizes fully from the plane to the torus. However,

there is an important difference from the planar setting. In both the plane and the torus, every translation of a reciprocal diagram is another reciprocal diagram. For a coherent plane graph Γ, *every* orthogonal reciprocal diagram is a weighted Voronoi diagram of the vertices of Γ, but *exactly one* orthogonal reciprocal diagram of a coherent geodesic embedding Γ is a weighted Voronoi diagram of the vertices of Γ. Said differently, every coherent plane graph is a weighted Delaunay graph with respect to a three-dimensional space of vertex weights, which correspond to translations of any convex polyhedral lifting, but every coherent *torus* graph is a weighted Delaunay graph with respect to only a *one*-dimensional space of vertex weights.

3.4.1 Notation

Fix a 2×2 matrix M with columns ξ, ζ and $\det M > 0$, and consider a geodesic embedding Γ on $\mathbb{T}_M$. We primarily work with the universal cover $\widetilde{\Gamma}$ of Γ; if we are given a reciprocal embedding Γ^*, we also work with its universal cover $\widetilde{\Gamma}^*$ (which is reciprocal to $\widetilde{\Gamma}$).

In this section, we will conflate vertices with their coordinates, and thus treat vertices in $\widetilde{\Gamma}$ as column vectors in $\mathbb{R}^2$.

A generic face in $\widetilde{\Gamma}$ is denoted by the letter f; the corresponding dual vertex in $\widetilde{\Gamma}^*$ is denoted f^* and interpreted as a row vector. To avoid nested subscripts when darts are indexed, we write $\Delta_i = \Delta_{d_i}$ and $\lambda_i = \lambda_{d_i}$, and therefore by Lemma 3.1, $\Delta_i^* = \lambda_i \Delta_i^\perp$. For any integers a and b, the translation $u + a\xi + b\zeta$ of any vertex u of $\widetilde{\Gamma}$ is another vertex of $\widetilde{\Gamma}$, and the translation $f + a\xi + b\zeta$ of any face f of $\widetilde{\Gamma}$ is another face of $\widetilde{\Gamma}$. It follows that $(f + a\xi + b\zeta)^* = f^* + a\xi^T + b\zeta^T$.

3.4.2 Results

The following lemma follows directly from the definitions of weighted Delaunay graphs and their dual weighted Voronoi diagrams; see, for example, Aurenhammer [19, 21].

Lemma 3.16. Let Γ be a weighted Delaunay graph on some flat torus $\mathbb{T}_M$, and let Γ^* be the corresponding weighted Voronoi diagram on $\mathbb{T}$. Every edge e of Γ is orthogonal to its dual e^*. In short, every coherent geodesic embedding is orthogonal reciprocal.

The converse of this lemma is false; unlike in the plane, an orthogonal reciprocal diagram Γ^* for a geodesic embedding Γ is *not* necessarily a weighted Voronoi diagram of the vertices of Γ. Rather, as we describe below, a unique translation of Γ^* is such a weighted Voronoi diagram.

Maxwell's theorem implies a convex polyhedral lifting $z \colon \mathbb{R}^2 \to \mathbb{R}$ of the universal cover $\widetilde{\Gamma}$ of Γ, where the gradient vector $\nabla z|_f$ within any face f is equal to the coordinate vector of the

dual vertex f^* in $\widetilde{\Gamma}^*$. To make this lifting unique, we fix a vertex o of $\widetilde{\Gamma}$ to lie at the origin $\binom{0}{0}$, and we require $z(o) = 0$.

Define the weight of each vertex $u \in \widetilde{\Gamma}$ as

$$\pi_u := \tfrac{1}{2}|u|^2 - z(u). \tag{3.10}$$

By definition, $\pi_o = 0$. The determinant conditions (2.6) and (2.7) for an edge e to be locally Delaunay are both equivalent to interpreting $\frac{1}{2}|p|^2 - \pi_u$ as a z-coordinate and requiring that the induced lifting be locally convex at e. Because z is a convex polyhedral lifting, these weights establish that $\widetilde{\Gamma}$ is the intrinsic weighted Delaunay graph of its vertex set.

Translating the universal cover $\widetilde{\Gamma}^*$ of the orthogonal reciprocal graph Γ^* adds a global linear term to the lifting function z, and therefore to the Delaunay weights π_u. The main result of this section is that there is a unique translation such that the corresponding Delaunay weights π_u are periodic.

To compute $z(q)$ for any point $q \in \mathbb{R}^2$, we choose an arbitrary face f of $\widetilde{\Gamma}$ that contains q and identify the equation $z|_f(q) = \eta q + c$ of the plane through the lift of f, where $\eta \in \mathbb{R}^2$ is a row vector and $c \in \mathbb{R}$. Borcea and Streinu [38] give a calculation for η and c, which for our setting can be written as follows:

Lemma 3.17 (Borcea and Streinu [38, Eq. 7])**.** For $q \in \mathbb{R}^2$, let f be a face containing q. The function $z|_f$ can be explicitly computed as follows:

- Pick an arbitrary **root** face f_0 incident to o.
- Pick an arbitrary path from f_0^* to f^* in $\widetilde{\Gamma}^*$, and let $d_1^*, \dots, d_\ell^*$ be the dual darts along this path. By definition, we have $f^* = f_0^* + \sum_{i=1}^{\ell} \Delta_i^*$. Set $C(f) = \sum_{i=1}^{\ell} \lambda_i \, |u_i \;\; v_i|$, where $d_i : u_i \to v_i$ and $|u_i \;\; v_i| = \det(u_i, v_i)$.
- Set $\eta = f^*$ and $c = C(f)$, implying that $z|_f(q) = f^* q + C(f)$. In particular, $C(f)$ is the intersection of this plane with the z-axis.

Orthogonal reciprocality of $\widetilde{\Gamma}^*$ implies that the actual choice of root face f_0^* and the path to f^* do not matter. We use this explicit computation to establish the existence of a translation of Γ^* such that $\pi_o = \pi_u = \pi_v = 0$. We then show that after this translation, every lift of the same vertex of Γ has the same Delaunay weight.

Lemma 3.18. There is a unique translation of $\widetilde{\Gamma}^*$ such that $\pi_o = \pi_\xi = \pi_\zeta = 0$. Specifically, this translation places the dual vertex of the root face f_0 at the point

$$f_0^* = \left(-\tfrac{1}{2}\left(|\xi|^2, |\zeta|^2\right) - \left(C(f_0 + \xi), C(f_0 + \zeta)\right)\right) M^{-1}.$$

Proof. Lemma 3.17 implies that

$$z(u) \;=\; (f_0+\xi)^*\xi + C(f_0+\xi) \;=\; f_0^*\xi + |\xi|^2 + C(f_0+\xi),$$

and by definition, $\pi_\xi = 0$ if and only if $z(\xi) = \frac{1}{2}|\xi|^2$. Thus, $\pi_\xi = 0$ if and only if $f_0^*\xi = -\frac{1}{2}|\xi|^2 - C(f_0+\xi)$. A symmetric argument implies $\pi_\zeta = 0$ if and only if $f_0^*\zeta = -\frac{1}{2}|\zeta|^2 - C(f_0+\zeta)$. □

Lemma 3.19. If $\pi_o = \pi_\xi = \pi_\zeta = 0$, then $\pi_u = \pi_{u+\xi} = \pi_{u+\zeta}$ for all $u \in V(\widetilde{\Gamma})$. In other words, all lifts of any vertex of Γ have equal weight.

Proof. Let f be any face incident to p, and let $P = d_1^*, \ldots, d_\ell^*$ be an arbitrary path from f_0^* to f^* in $\widetilde{\Gamma}^*$. We compute $C(f+\xi)$ by traversing an arbitrary path from f_0^* to $(f_0+\xi)^* = f_0^* + \xi^T$ followed by the translated path $P+\xi$ from $f_0^*+\xi^T$ to $f^*+\xi^T$. Thus by Lemma 3.17, $C(f+\xi) = C(f_0+\xi) + \sum_{i=1}^{\ell} \lambda_i\,|(u_i+\xi)\;\;(v_i+\xi)|$, and $f^* = f_0^* + \sum_{i=1}^{\ell} \Delta_i^*$. We thus have

$$\begin{aligned}
C(f+\xi) &= C(f_0+\xi) + \sum_{i=1}^{\ell} \lambda_i\,|(u_i+\xi)\;\;(v_i+\xi)| \\
&= C(f_0+\xi) + \sum_{i=1}^{\ell} \lambda_i\,|u_i\;\;v_i| - \sum_{i=1}^{\ell} \Delta_i^*\xi \\
&= C(f_0+\xi) + C(f) - \sum_{i=1}^{\ell} \Delta_i^*\xi \\
&= -\tfrac{1}{2}|\xi|^2 - f_0^*\xi + C(f) - \sum_{i=1}^{\ell} \Delta_i^*\xi \\
&= -\tfrac{1}{2}|\xi|^2 - f^*\xi + C(f).
\end{aligned}$$

It follows that

$$\begin{aligned}
\pi_{u+\xi} &= \tfrac{1}{2}|u+\xi|^2 - z(u+\xi) \\
&= \tfrac{1}{2}|u+\xi|^2 - \big(C(f+\xi) + (f^*+\xi^T)(u+\xi)\big) \\
&= \tfrac{1}{2}|u+\xi|^2 - \big(-\tfrac{1}{2}|\xi|^2 - f^*\xi + C(f) + f^*u + f^*\xi + \xi^T u + |\xi|^2\big) \\
&= \tfrac{1}{2}|u+\xi|^2 - z(u) - \tfrac{1}{2}|\xi|^2 - \xi^T u \\
&= \tfrac{1}{2}|u|^2 + \tfrac{1}{2}|\xi|^2 + \xi^T u - z(u) - \tfrac{1}{2}|\xi|^2 - \xi^T u \\
&= \tfrac{1}{2}|u|^2 - z(u) \\
&= \pi_u.
\end{aligned}$$

A similar computation implies $\pi_{u+\zeta} = \pi_u$. □

Projecting from the universal cover back to the torus, we obtain weights for the vertices of Γ, with respect to which Γ is an intrinsic weighted Delaunay complex, and a unique translation of Γ^* that is the corresponding intrinsic weighted Voronoi diagram. Moreover, these Delaunay vertex weights are unique if we fix the weight of an arbitrary vertex of Γ to be 0.

Theorem 3.5. Let Γ and Γ^* be orthogonal reciprocal embeddings on some flat torus $\mathbb{T}_M$. Then Γ is a weighted Delaunay complex, and a unique translation of Γ^* is the corresponding weighted Voronoi diagram. In short, every reciprocal embedding is coherent.

3.5 Application: A Toroidal Steinitz Theorem

In light of the results of the previous section, Theorem 2.3 and Theorem 3.3 immediately imply a natural generalization of Steinitz's theorem to graphs on the flat torus.

Theorem 3.6. Let Γ be any essentially simple, essentially 3-connected embedded graph on the square flat torus $\mathbb{T}_\square$, and let λ be *any* positive symmetric stress on the edges of Γ. Then Γ is homotopic to a geodesic embedding in $\mathbb{T}_\square$ whose image in some flat torus $\mathbb{T}_M$ is coherent.

As we mentioned above, Mohar's generalization [163] of the Koebe-Andreev-Thurston circle packing theorem already implies that every essentially simple, essentially 3-connected torus graph Γ is homotopic to *one* coherent homotopic embedding on *one* flat torus. In contrast, our results characterize *all* coherent homotopic embeddings of Γ on *all* flat tori. Every positive stress vector $\lambda \in \mathbb{R}^E$ corresponds to an essentially unique coherent homotopic embedding of Γ, which is unique up to translation, on a flat torus $\mathbb{T}_M$, which is unique up to similarity of the fundamental parallelogram $\Diamond_M$. On the other hand, Lemmas 3.1 and 3.16 imply that every coherent embedding of Γ on every flat torus corresponds to a unique positive equilibrium stress.

Chapter 4

Toroidal Morphing via Collapsing Edges

In this chapter, we consider the problem of computing morphs between two isotopic embeddings of the same graph on the same flat torus.

Recall that, for the sake of convenience, we will assume that we are working on the square flat torus $\mathbb{T}_\square$; the algorithms we describe will apply to any other flat torus $\mathbb{T}_M$ via linear transformation of the coordinates. Moreover, in light of the normalization algorithm described in Section 2.8, we assume that the inputs to our problem are specified by coordinates that have been normalized to have the same translation vectors.

Cairns [44, 45] was the first to prove the existence of morphs between arbitrary isomorphic straight-line triangulations. Cairns' approach used an inductive argument based on the idea of collapsing low-degree vertices to their neighbors. Our goal is to generalize Cairns' edge-collapse method [44] to the flat torus. To the best of our knowledge, even *existence* of such morphs was previously unknown.

Cairns starts by computing a ***pseudomorph*** between *triangulations*, in which edges remain geodesics, but vertices are allowed to become coincident (though edges may not cross and faces cannot flip orientation); the bulk of this chapter will be devoted to explaining how to extend this approach to the flat torus. We will also later see how to use this to obtain *morphs* between more general embeddings.

4.1 Direct Collapses and Cairns' Pseudomorph

Given a simple polygon P, the ***visibility kernel*** of P is the set of all points in P that can "see" all of P; more formally, the visibility kernel is $\{p \in P \mid pq \subseteq P \text{ for all } q \in P\}$. Call a vertex u ***good*** if the visibility kernel of u's link contains one of u's neighbors, say, v. Suppose we move u along the edge connecting u to v until u and v coincide, keeping all other vertices fixed. This motion, which we call a ***direct collapse***, is trivially a unidirectional linear pseudomorph.

Cairns observed that for every n-gon for $n \leq 5$, its visibility kernel contains one of its vertices. Euler's formula $V - E + F = 2$ implies that for any planar graph, the average degree of a vertex is strictly less than 6; in short, every planar graph has a vertex of degree at most 5. The happy coincidence of two observations results allowed Cairns to devise a pseudomorph between two planar triangulations by repeatedly applying direct collapses.

If u can be directly collapsed to the same neighbor v in both Γ_0 and Γ_1, then performing this operation results in two isotopic embeddings of a graph with one fewer vertex. Hence one can recursively compute a pseudomorph ψ_{n-1}, and then produce a pseudomorph ψ_n between Γ_0 and Γ_1 consisting of first directly collapsing u to v in Γ_0, performing the pseudomorph ψ_{n-1}, and then reversing a direct collapse to reobtain Γ_1. It is possible, however, that u is collapsible to one neighbor v in Γ_0, and a different neighbor w in Γ_1. Cairns observed that there exists an intermediate isotopic embedding Γ_* in which u can be collapsed to both v and w (for example, we can set Γ_* to be the Tutte embedding corresponding to the all-1s vector). Thus we can compute a pseudomorph between Γ_0 and Γ_1 by first computing Γ_* and then recursing *twice*. The result is a pseudomorph consisting of $O(2^n)$ steps.

When attempting to adapt Cairns' method to the torus, we immediately run into a major roadblock. l'Huilier's generalization [144] of Euler's formula implies that for a torus graph, the average degree of a vertex is *exactly* 6, so it is possible that the link of every vertex is a hexagon. And whereas n-gons for $n \leq 5$ were guaranteed to have a vertex in the visibility kernel, there exist simple examples of hexagons that fail to have this property; see Figure 4.1. Thus we will need to figure out how to handle 6-regular torus triangulations in order to move forward.

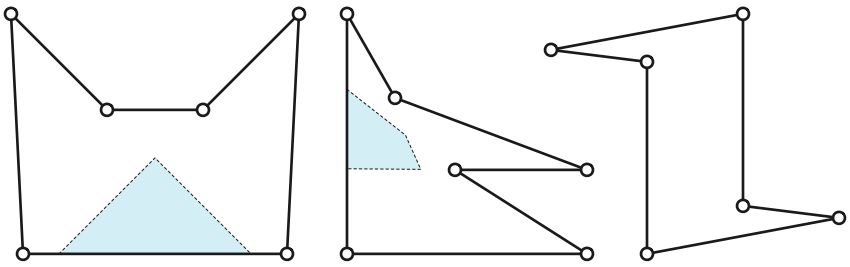

Figure 4.1. Three bad hexagons. Visibility kernels are shaded in blue; the third visibility kernel is empty.

4.2 Zippers

Even if the original input triangulations Γ_0 and Γ_1 are simple, collapsing edges eventually reduces them to triangulations with parallel edges and loops. Every loop in a geodesic toroidal triangulation is a closed geodesic. A ***zipper*** is a torus graph in which *every* vertex is incident to a loop. (This class of graphs were previously considered by Gonçalves and Lévêque [108, Fig. 44].) We will see in Lemma 4.2 that every 6-regular triangulation with a loop is a zipper.

Consider any zipper Z with n vertices, for some positive integer n. If $n = 1$, then Z consists of a single vertex, three loop edges, and two triangular faces. Otherwise, the loops in Z are disjoint closed geodesics, so they must be parallel. It follows that each vertex of Z is incident

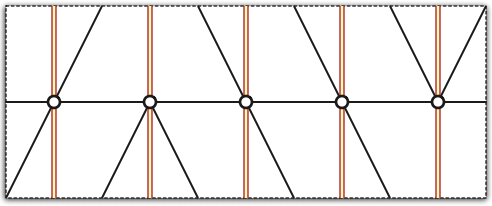

Figure 4.2. A five-vertex zipper. Doubled red edges are loops.

to exactly one loop. In either case, the n loops in Z decompose the torus into n annuli, each of which is decomposed into two triangles by two boundary-to-boundary edges. Figure 2.3 shows three two-vertex zippers, and Figure 4.2 shows a zipper with five vertices.

4.2.1 Structure

The following results will motivate our eventual choice of zippers as the base case of our generalization of Cairns' recursive pseudomorph.

Lemma 4.1. In every geodesic toroidal triangulation, every vertex incident to a loop has degree at least 6.

Proof. Let Γ be a geodesic toroidal triangulation, let v be a vertex of Γ incident to a loop, let d_0 be either of the darts of that loop.

Let ∂f denote the clockwise facial walk around the face f to the right of d_0. Because the interior of f is an open disk, ∂f is contractible. Because Γ is a triangulation, ∂f consists of exactly three darts d_0, d_1, and d_2, where $head(d_i) = tail(d_{i+1 \bmod 3})$ for each index i. Because Γ is a *geodesic* triangulation, every face is incident to three *distinct* edges. It follows that $d_2 \neq rev(d_1)$.

Symmetrically, the counterclockwise walk around the face to the left of d_0 consists of three darts d_0, d_1', d_2', where $d_2' \neq rev(d_1')$. Thus, at least six distinct darts head into v: in counterclockwise cyclic order, d_0, $rev(d_1)$, d_2, $rev(d_0)$, d_2', $rev(d_1')$.[1] □

Lemma 4.2. Every 6-regular triangulation that contains a loop is a zipper.

Proof. Again, let Γ be a 6-regular triangulation, and let v be a vertex of Γ incident to a loop ℓ. The previous proof implies that v is incident to two edges on either side of ℓ. Let e and e' be the edges incident to v on one side of ℓ; the edges ℓ, e, and e' enclose a triangular face f. Thus, e and e' share another common endpoint w. (Except in the trivial case where Γ has

[1] If v is the only vertex of Γ, then $d_2' = d_1$ and $d_1' = d_2$, and thus v is incident to only three distinct *edges*, which are all loops. Otherwise, v is incident to at least five distinct edges, only one of which is a loop.

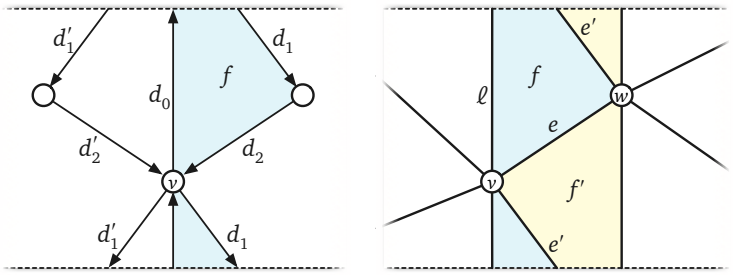

Figure 4.3. (a) Proof of Lemma 4.1. (b) Proof of Lemma 4.2

only one vertex, v and w are distinct.) Because e and e' are adjacent in cyclic order around v, there is another triangular face f' with e and e' on its boundary; the third edge of f' is a loop through w. The lemma now follows by induction. □

Corollary 4.1. Every triangulation that contains a loop but is not a zipper contains a vertex of degree at most 5 that is not incident to a loop.

4.2.2 Morphing Zippers

We next describe a straightforward approach to morphing between arbitrary isotopic zippers, which requires at most two unidirectional linear morphing steps.

Let Z and Z' be arbitrary isotopic zippers with n vertices. If $n = 1$, then Z and Z' differ only by translation, so assume otherwise. Ladegaillerie [138–140] proved that two embeddings of the same graph G on the same surface are isotopic if and only if the images of any cycle in G in both embeddings are homotopic. Two geodesic loops on the flat torus are homotopic if and only if they are parallel. Thus, Ladegaillerie's theorem implies that the loops in Z and Z' are all parallel to a common vector σ.

In the first unidirectional morphing step, we translate all vertices in Z along geodesics orthogonal to σ until each loop has the same image as the corresponding loop in Z'. Then in the second unidirectional morphing step, we translate all vertices along their respective loops to move all vertices and edges to their proper positions in Z'. See Figure 4.4 for an example. In both stages, Lemma 4.3 implies that linear interpolation between the old and new vertex coordinates yields an isotopy.

Lemma 4.3. Let p_0p_1, q_0q_1, and r_0r_1 be arbitrary parallel segments in the plane. For all real $0 \le t \le 1$, define $p_t = (1-t)p_0 + tp_1$ and $q_t = (1-t)q_0 + tq_1$ and $r_t = (1-t)r_0 + tr_1$. If the triples p_0, q_0, r_0 and p_1, q_1, r_1 are oriented counterclockwise, then for all $0 \le t \le 1$, the triple p_t, q_t, r_t is also oriented counterclockwise.

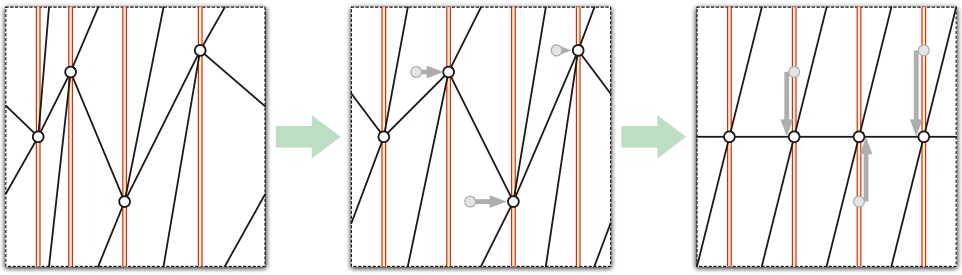

Figure 4.4. Morphing one zipper into another. Doubled red edges are loops.

Proof. Without loss of generality, assume that segments p_0p_1, q_0q_1, and r_0r_1 are horizontal. Thus, we can write $p_t = (px_t, py)$, and similarly for q_t and r_t. The triple p_t, q_t, r_t is oriented counterclockwise if and only if the following determinant is positive:

$$\Delta(t) := \begin{vmatrix} 1 & px_t & py \\ 1 & qx_t & qy \\ 1 & rx_t & ry \end{vmatrix}$$

Routine calculation implies

$$\begin{aligned} \Delta(t) &= \begin{vmatrix} 1 & (1-t)px_0 + t\,px_1 & py \\ 1 & (1-t)qx_0 + t\,qx_1 & qy \\ 1 & (1-t)rx_0 + t\,rx_1 & ry \end{vmatrix} \\ &= (1-t)\begin{vmatrix} 1 & px_0 & py \\ 1 & qx_0 & qy \\ 1 & rx_0 & ry \end{vmatrix} + t\begin{vmatrix} 1 & px_1 & py \\ 1 & qx_1 & qy \\ 1 & rx_1 & ry \end{vmatrix} \\ &= (1-t)\cdot\Delta(0) + t\cdot\Delta(1) \end{aligned}$$

Thus, the function $\Delta(t)$ has exactly one real root. It follows that if $\Delta(0) > 0$ and $\Delta(1) > 0$, then $\Delta(t) > 0$ for all $0 \le t \le 1$. □

4.3 Cats and Dogs

In this section, we consider the case of 6-regular triangulations where no vertex is incident to a loop. In essence, we will show that every geodesic toroidal triangulation without loops has a directly collapsible edge.

Recall that the ***visibility kernel*** of a simple polygon P is the set of all points in P that can "see" all of P; more formally, the visibility kernel is $\{p \in P \mid pq \subseteq P \text{ for all } q \in P\}$. The visibility

kernel is always convex. If P is the link of a vertex v in a geodesic triangulation, then v must lie in the visibility kernel of P. We call a simple polygon ***good*** if its visibility kernel contains a vertex of the polygon, and ***bad*** otherwise. All triangles, quadrilaterals, and pentagons are good [44], but some hexagons are bad; Figure 4.1 shows several examples.

Formally, a vertex u of a geodesic toroidal triangulation is ***good*** if it is not incident to a loop, and the link of any (and thus every) lift $\tilde{u}$ of u in $\widetilde{\Gamma}$ is good, and ***bad*** otherwise. A good vertex can be safely collapsed to any neighbor in the visibility kernel of its link. Finally, a geodesic toroidal triangulation Γ with no loops is ***good*** if it contains at least one good vertex, and ***bad*** otherwise. The main result of this section is that bad triangulations do not exist; that is, every geodesic triangulation without loops is good.

Lemma 4.4. Every bad triangulation is 6-regular.

Proof. Every vertex in a bad triangulation must have degree at least 6, because every vertex with degree at most 5 is good [44]. On the other hand, Euler's formula for the torus implies that the average degree is exactly 6. □

The following analysis of 6-regular triangulations without loops implicitly assumes that the vertices are in general position. Specifically, for every vertex $\tilde{v}$ in the universal cover, we assume that (1) no pair of edges incident to $\tilde{v}$ are collinear, and (2) no edge of the link of $\tilde{v}$ is collinear with another vertex of that link. Because the underlying triangulation has no loops, each nondegeneracy condition involves at least two distinct vertices of the torus triangulation, so we can enforce these conditions if necessary by perturbing the vertices.

Lemma 4.5. In every bad triangulation, the link of each vertex has exactly two reflex vertices.

Proof. If a simple polygon is convex or has exactly one reflex vertex, then it is good: the visibility kernel of a convex polygon is the polygon itself, and if there is exactly one reflex vertex, then the reflex vertex is in the visibility kernel. So each link in a bad triangulation has *at least* two reflex vertices. We argue next that the average number of reflex vertices per link is *at most* two, which implies that every link has exactly two reflex vertices.

Let Γ be any (not necessarily bad) 6-regular geodesic triangulation. A *corner* of a vertex v in Γ is the angle between two edges that are adjacent in cyclic order around v. If v is a vertex of the link of another vertex x, then the link of x contains exactly two adjacent corners of v. Moreover, if v is a *reflex* vertex of the link of x, those two corners sum to more than half a circle.

Thus, if v is reflex in the link of two neighbors x and y, the links of x and y must share a corner of v, which implies that edges vx and vy are adjacent in cyclic order around v; see

Figure 4.5. If v were reflex in the links of three neighbors x, y, and z, then all three edge pairs vx, vy and vx, vz and vy, vz would be adjacent around v, which is impossible because v has degree 6.

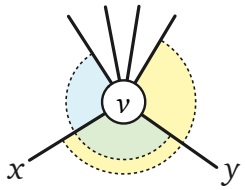

Figure 4.5. Each vertex in a 6-regular triangulation is reflex in at most two links

We conclude that each vertex in Γ is a reflex vertex of the links of at most two other vertices. It follows that the average number of reflex vertices in a link is at most two, which completes the proof. □

Lemma 4.6. A bad hexagon with two reflex vertices separated by two convex vertices has an empty visibility kernel, and thus is not the link of any vertex.

Proof. Consider a bad hexagon P whose reflex vertices are separated by two convex vertices. Label the vertices a through f in cyclic order such that c and f are the reflex vertices, as shown in Figure 4.6.

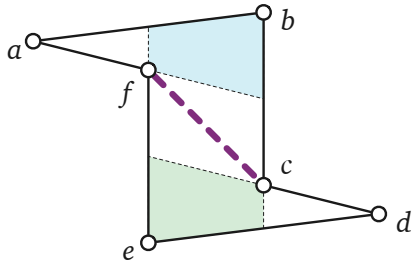

Figure 4.6. A hexagon with two reflex vertices that are separated by two convex vertices. The dashed diagonal splits the hexagon into two non-convex quadrilaterals with disjoint visibility kernels.

Because P is bad, neither c nor f lies in its visibility kernel. Vertex c can see vertices b, d, and f, so it cannot see both a and e. Without loss of generality, suppose c cannot see a. Then f is a reflex vertex of the quadrilateral $abcf$. Every quadrilateral has at most one reflex vertex, and that reflex vertex is in the visibility kernel, so f can see b. It follows that f cannot see d, and c is a reflex vertex of the quadrilateral $cdef$.

The visibility kernels of quadrilaterals $abcf$ and $cdef$ are disjoint, which implies that the visibility kernel of P is empty. □

For the remainder of the proof, we annotate the edges of any triangulation as follows. The *star* of an edge in a triangulation is the union of the faces incident to that edge. An edge is ***flippable*** if its star is convex, and non-flippable otherwise. Every non-flippable edge is incident to the unique reflex vertex of its star; we direct each non-flippable edge away from this reflex vertex.

Lemma 4.7. In any bad triangulation, every vertex is incident to exactly two incoming directed edges, exactly two outgoing directed edges, and exactly two flippable edges.

Proof. Fix a bad triangulation Γ. If w is a reflex vertex in the link of some vertex v, then the edge vw is directed toward v. So Lemma 4.5 implies that each vertex of Γ is incident to two incoming edges.

We argued in the proof of Lemma 4.5 that each vertex is reflex in the links of exactly two of its neighbors. Thus, each vertex of Γ is also incident to two outgoing edges.

Finally, because every vertex of Γ has degree 6, each vertex of Γ is incident to two flippable edges. □

Lemmas 4.6 and 4.7 imply that the links in every bad triangulation fall into two categories, which we call *cats* and *dogs*. A ***cat*** is a bad hexagon whose reflex vertices are adjacent; a ***dog*** is a bad hexagon whose reflex vertices are separated by one convex vertex. Cats are (combinatorially) symmetric; however, there are two species of dogs, which are reflections of each other.

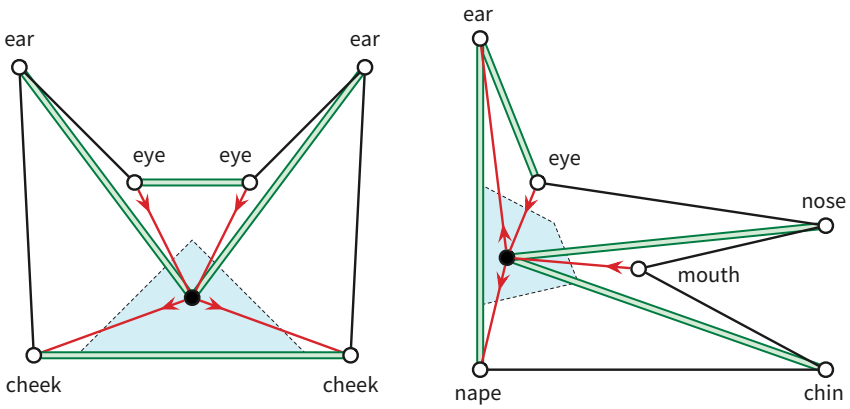

Figure 4.7. Cat and (right-facing) dog anatomy. Flippable edges are shown in double-thick green.

We label the vertices of each cat and dog mnemonically, as shown in Figure 4.7. In particular, if the link of vertex v is a dog, then $nose(v)$ is the unique convex vertex of the dog whose neighbors are reflex, and the vertex opposite $nose(v)$ in the link is $nape(v)$. Because $nape(v)$ is incident to two convex vertices, it can see every vertex of the dog except $nose(v)$. Because

neither reflex vertex lies in the visibility kernel of the dog, $nape(v)$ can see both reflex vertices. The visibility between $nape(v)$ and $nose(v)$ is blocked by one of the reflex vertices, which we call $mouth(v)$. The edge from v to $nose(v)$ is flippable, because otherwise v could not see either $ear(v)$ or $chin(v)$. A dog is ***right-facing*** if $mouth(v)$ immediately follows $nose(v)$ in clockwise order, and ***left-facing*** otherwise.

The following lemma implies that two boundary edges of each cat and dog are flippable, as shown in Figure 4.7.

Lemma 4.8. In any bad triangulation, every triangle has exactly one flippable edge.

Proof. Fix a bad triangulation Γ with n vertices. Euler's formula implies that Γ has $3n$ edges and $2n$ triangular faces. Lemma 4.7 implies that Γ has exactly n flippable edges, so the average number of flippable edges per triangle is exactly 1.

The flippable edges incident to any vertex v are separated in cyclic order around v by two non-flippable edges if the link of v is a cat, or by one non-flippable edge if the link of v is a dog. Thus, two flippable edges never appear consecutively around a common vertex. It follows that every triangle in Γ is incident to at most one flippable edge. □

Lemma 4.9. Every bad triangulation contains a cat.

Proof. Let Γ be a triangulation, and let u be any vertex in Γ whose link is a dog. We argue that the link of the nose of u must be a cat. (Mnemonically, dogs only sniff cats.)

Without loss of generality, assume that the link of u is facing right, so the triple $chin(u), nape(u), ear(u)$ is oriented clockwise, as shown in Figure 4.8.

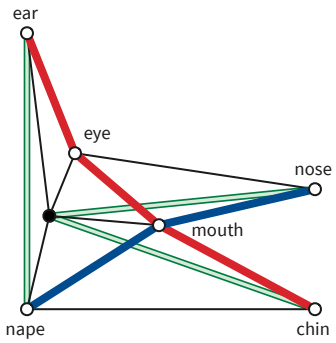

Figure 4.8. Concave chains inside any dog.

The fact that the link of u is a *bad* hexagon implies several orientation constraints on its vertices:

- The triple $ear(u), eye(u), mouth(u)$ is oriented counterclockwise; otherwise, $mouth(u)$ could see the entire dog. It follows that the triple $eye(u), mouth(u), nape(u)$ is oriented clockwise, and that $nape(u)$ and $eye(u)$ can see each other.

- The triple $eye(u), mouth(u), chin(u)$ is oriented counterclockwise; otherwise, $eye(u)$ could see the entire dog.
- Finally, the triple $nose(u), mouth(u), nape(u)$ is oriented counterclockwise; otherwise, $nape(u)$ could see the entire dog.

Now suppose for the sake of argument that the link of $v = nose(u)$ is a dog. We label the other vertices of the link of u as shown on the top row of Figure 4.9; in particular, $x = eye(u)$ and $y = mouth(u)$. Because the flippable edge uv is inside the link of v, either $u = nose(v)$ or $u = chin(v)$; each of these cases admits two subcases. The four cases are illustrated schematically in the rows of Figure 4.9.

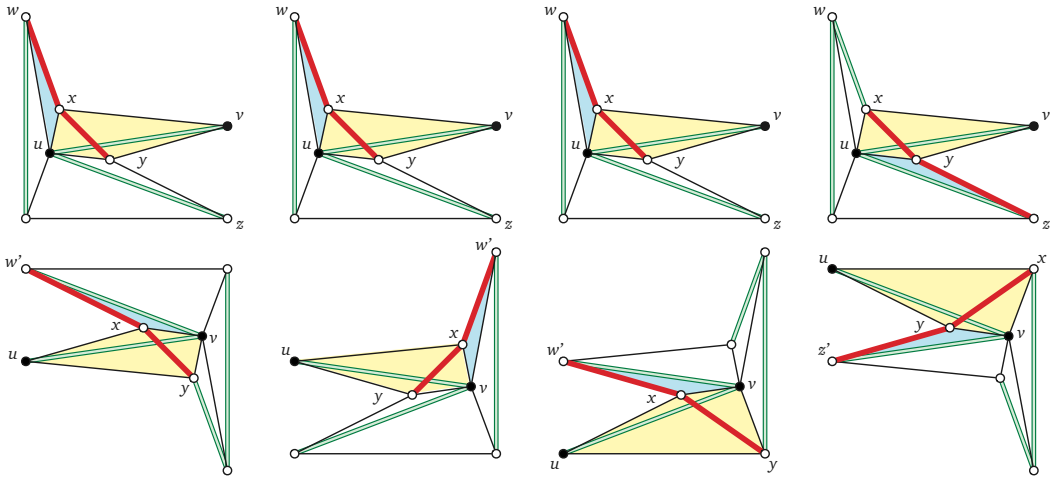

Figure 4.9. Four cases in the proof of Lemma 4.5: One dog's nose (v) cannot be the center of another dog. Each column corresponds to one case.

- Suppose $u = nose(v)$ and $x = mouth(u)$, and therefore $y = eye(u)$. Let $w' = chin(v)$. The triple w', x, y must be oriented clockwise; our earlier analysis implies that w, x, y is oriented counterclockwise. It follows that triangles $w'xv$ and wxu overlap, which is impossible.
- Suppose $u = nose(v)$ and $x = eye(u)$, and therefore $y = mouth(u)$. Let $w' = ear(v)$. The triple w', x, y must be oriented clockwise; our earlier analysis implies that w, x, y is oriented counterclockwise. It follows that triangles $w'xv$ and wxu overlap, which is impossible.
- Suppose $u = chin(v)$ and $y = nape(v)$. Let $w' = nose(v)$. The triple w', x, y is oriented clockwise; our earlier analysis implies that w, y, z is oriented counterclockwise. It follows that triangles $w'xv$ and wxu overlap, which is impossible.

- Finally, suppose $u = chin(v)$ and $y = mouth(v)$. Let $z' = nose(v)$. The triple x, y, z' is oriented clockwise; our earlier analysis implies that x, y, z is oriented counterclockwise. It follows that triangles $z'yv$ and zyu overlap, which is impossible.

In all four cases, we derive a contradiction. We conclude that the link of $nose(u)$ is actually a cat. □

Every 6-regular triangulation of the torus is isotopic to the quotient of the regular equilateral-triangle tiling of the plane by a lattice of translations [9, 40, 168]. We analyze the patterns of cats and dogs in bad triangulations by analyzing their images in this reference triangulation, and in particular, by studying the induced annotations of edges, as illustrated in Figure 4.10.

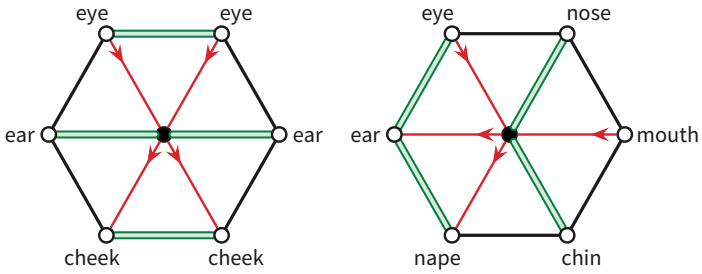

Figure 4.10. Cat and (right-facing) dog reference anatomy; compare with Figure 4.7.

A cycle in a 6-regular toroidal triangulation is ***straight*** if it corresponds to a closed geodesic in the corresponding lattice triangulation. Every straight cycle is non-contractible.

Let C be a simple piecewise geodesic closed curved on the torus (for example, a cycle in Γ), arbitrarily directed, and let u, v, w be three successive vertices of C. Let $\widetilde{v}$ be any lift of v, and let $\widetilde{u}\widetilde{v}$ and $\widetilde{v}\widetilde{w}$ be lifts of the edges of C incident to v. The ***turning angle*** of C at u is the signed counterclockwise angle, strictly between $-\pi$ and π, between the vectors $\widetilde{u}\to\widetilde{v}$ and $\widetilde{v}\to\widetilde{w}$. In other words, when walking along the cycle C in the indicated direction, the turning angle is the angle one turns left at v (or right if the angle is negative). The ***total turning angle*** of C is the sum of the turning angles of the vertices of C.

Lemma 4.10. Every simple non-contractible closed curve on the flat torus has total turning angle zero.

Lemma 4.10 follows immediately from classical results of Reinhart [180]. In short, the total turning angle (which Reinhart calls the "winding number") is an isotopy invariant of closed curves, and every simple non-contractible closed curve on the flat torus is isotopic to a closed geodesic. Lemma 4.10 implies in particular that every *straight* cycle has total turning angle zero.

Lemma 4.11. If a bad triangulation contains one dog, it contains a straight cycle of vertices whose links are dogs.

Proof. Let Γ be a bad triangulation and let v_1 be any vertex of Γ whose link is a dog. Label the vertices of v's link as shown in Figure 4.11; for example, $u_1 = nose(v_1)$, $v_2 = mouth(v_1)$, and $w_1 = chin(v_1)$.

Edges u_1v_1 and v_1w_1 are flippable, and therefore edges v_1v_2, u_1v_2, and w_1v_2 are not flippable. Thus, no opposite pair of edges incident to v_2 are both flippable. It follows that the link of v_2 is a dog, whose ear vertex is v_1.

Edges u_0v_1, v_0v_1, and v_1w_0 are not flippable, so edges u_0v_0 and v_0w_0 must be flippable. Thus, the link of v_0 is also a dog, whose mouth vertex is v_1.

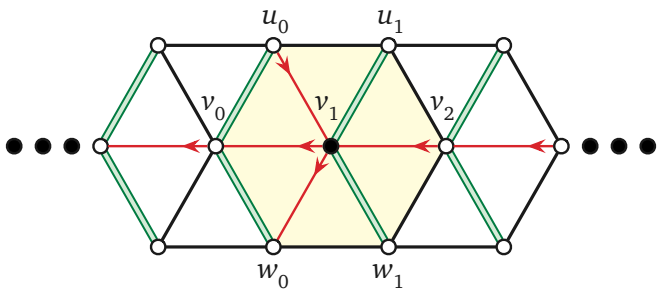

Figure 4.11. One dog induces a straight cycle of dogs.

Continuing by induction in both directions, we find a bidirectional sequence $\ldots, v_{-1}, v_0, v_1, v_2, \ldots$ of vertices whose links are dogs, where the center v_i of each dog is the mouth of the previous dog and the ear of the next dog. Because Γ is finite, this sequence must eventually repeat, forming a straight cycle. □

We can now finally prove the main result of this section.

Theorem 4.1. Bad triangulations do not exist.

Proof. Let Γ be a bad triangulation. We derive a contradiction by showing that Γ contains a non-contractible cycle whose vertices all have cat links, whose edges are all non-flippable, and finally whose turning angle is non-zero, contradicting Lemma 4.10.

First, suppose Γ contains at least one dog. Following the proof of Lemma 4.9, suppose v_1 is a vertex whose vertex is a dog and $u_1 = nose(v_1)$. Then $v_2 = mouth(v_1)$ is also a neighbor of u_1. The previous lemma now implies a straight cycle of dogs $D = \ldots, v_0, v_1, v_2, v_3, \ldots$, where $v_i = mouth(v_{i-1}) = ear(v_{i+1})$ for every index i. For each index i, the link of $u_i = nose(v_i)$ is a cat. Thus, we find a straight cycle $C = \ldots, u_0, u_1, u_2, u_3, \ldots$ of vertices, parallel to the cycle D

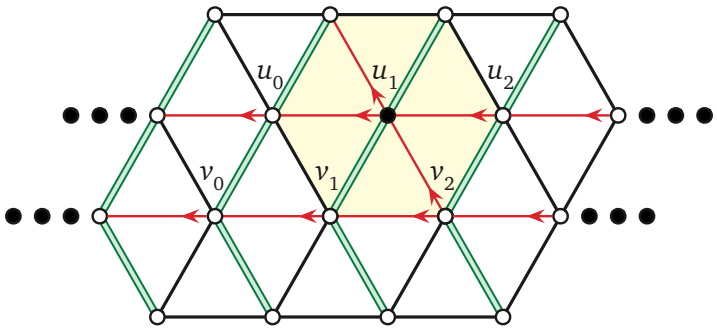

Figure 4.12. One cat induces a straight cycle of cats.

of dog vertices, all of whose links are cats, as shown in Figure 4.12. Every edge u_iu_{i+1} in C is unflippable.

On the other hand, if Γ contains no dogs, then we can construct a straight cycle $C = \ldots, u_0, u_1, u_2, \ldots$ of cat vertices starting with any unflippable edge u_0u_1. Again, every edge u_iu_{i+1} in C is unflippable.

Without loss of generality, suppose u_2 is the left eye of u_1, and thus u_0 is the right cheek of u_1. (Symmetric arguments apply if u_2 is the right eye, right cheek, or left cheek of u_1.) The triple $u_0u_1u_2$ is oriented counterclockwise, as shown in Figure 4.13; otherwise, the right cheek of u_1 would lie in the visibility kernel of the link of u_1. It follows by induction that for every index i, vertex u_{i+1} is the left eye of u_i, and thus every triple $u_{i-1}u_iu_{i+1}$ is oriented counterclockwise. We conclude that the total turning angle of C is positive, contradicting Lemma 4.10. □

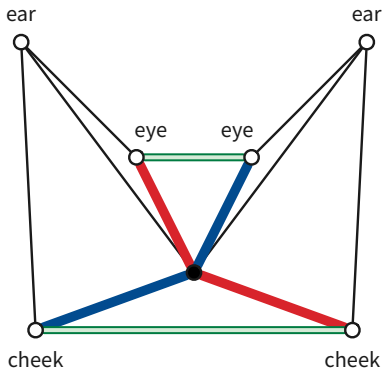

Figure 4.13. Concave chains inside any cat.

4.4 Triangulation Pseudomorphs

We saw in Section 4.2 that every triangulation with a loop is in fact a zipper, and furthermore, morphing between zippers turned out to be fairly simple. We will thus use zippers as the base

case for our extension to Cairns' recursive method for computing pseudomorphs.

Otherwise, we showed that a good vertex always exists, so we can always perform a direct collapse. Unfortunately, a worse version of the problem observed by Cairns above appears: not only is it possible that a good vertex u might not be directly collapsible to the same neighbor in both Γ_0 and Γ_1, it is possible that Γ_0 and Γ_1 do not share a good vertex! Fortunately, every 6-regular triangulation of the torus is isotopic to the quotient of the regular equilateral-triangle tiling of the plane by a lattice of translations [9, 40, 168]; equivalently, this tiling is the Tutte embedding Γ_* corresponding to the all-1s vector. In this tiling, *every* vertex is good, and directly collapsible to *every* neighbor.

Recursively constructing pseudomorphs from Γ_0 to Γ_* and from Γ_* to Γ_1 yields the desired pseudomorph. Just as with Cairns' pseudomorph, this pseudomorph consists of exponentially many steps. We will revisit this pseudomorph in Section 6.2.3, and derive a pseudomorph consisting of $O(n)$ steps.

4.5 From Triangulation Pseudomorphs to General Morphs

We first adapt a perturbation strategy of Alamdari *et al.* [4], which transforms their pseudomorphs between planar triangulations into morphs, to geodesic triangulations on the flat torus.

To explain our adaptation, we must first give a brief sketch of their algorithm. Let Γ_0 be the initial planar input triangulation. The input to their perturbation algorithm is a pseudomorph consisting of a (direct) collapse of a good vertex u to a neighbor v, a morph (not a pseudomorph) consisting of k unidirectional linear steps $\Gamma'_0 \rightsquigarrow \Gamma'_1 \rightsquigarrow \cdots \rightsquigarrow \Gamma'_k$, and finally the reverse of a collapse of u to v resulting in the final triangulation Γ_k. The output is a proper morph from Γ_0 to Γ_k, consisting of $k+2$ unidirectional linear steps.

Alamdari *et al.* proceed as follows. Let P be the link of u in the initial triangulation Γ_0. For each index i, let v_i and P_i respectively denote the images of v and P in the intermediate triangulation Γ'_i. For each index i, a position u_i is found within the visibility kernel of P_i so that for all i, the vector $u_i \to u_{i+1}$ is parallel to $v_i \to v_{i+1}$ (the direction of the unidirectional linear morph $\Gamma'_i \rightsquigarrow \Gamma'_{i+1}$). For vertices of degree 3 and 4, it is simple to place u as a certain convex combination of the vertices in P. The strategy for vertices of degree 5 is more complicated. First, a value ε is computed such that for each index i, the intersection of the disk of radius ε centered at v_i and the visibility kernel of P_i consists of a full sector of the disk; call this intersection S_i. A specific position u_i is then chosen within each region S_i.

A close examination of their paper reveals that the strategy for vertices of degree 5 generalizes to vertices u of arbitrary degree (greater than 2). In particular, the definition of u_i depends

solely on ε and the positions of the edges in P_i incident to v_i.

The radius ε is computed as follows. For each index i, we need a positive distance $\varepsilon_i > 0$ smaller than the minimum distance from v to any edge of the kernel of P_i that is not incident to v, at any time during the morphing step $\Gamma'_i \rightsquigarrow \Gamma'_{i+1}$. It suffices to compute the minimum distance from v to the lines supporting edges of P_i not incident to v. The squared distance to each of these lines at any time t can be expressed as the ratio $f(t)/g(t)$ of two quadratic polynomials f and g. Alamdari *et al.* argue that a lower bound $0 < \delta \le \min_t \sqrt{f(t)/g(t)}$ can be computed in constant time in an appropriate real RAM model [4], say, by computing roots of the derivative of $f(t)/g(t)$, whose numerator is a *cubic* polynomial. Then ε_i is the minimum of these lower bounds δ. Altogether computing ε_i takes $O(\deg(u))$ time. Finally, $\varepsilon = \min_{1\le i\le k} \varepsilon_i$.

Once the radius ε is known, computing each sector S_i in $O(\deg(u))$ time is straightforward.

The point u_0 can be chosen arbitrarily within S_0. For each index i in increasing order, Alamdari *et al.* describe how to choose a point $u_{i+1} \in S_{i+1}$ in $O(1)$ time, such that the vector $u_i \to u_{i+1}$ is parallel to the vector $v_i \to v_{i+1}$. This part of the algorithm makes no reference to the rest of the triangulation; it works entirely within the sectors S_i and S_{i+1}. Moreover, no part of this algorithm relies on u being *directly* collapsed to v, only that vertices u and v have the same image in the triangulations Γ'_i and Γ'_{i+1}.

We can thus apply this perturbation technique to our toroidal pseudomorphs with no modification.

Theorem 4.2. Given any two isotopic geodesic toroidal triangulations Γ_0 and Γ_1 with n vertices, we can compute a piecewise-linear morph from Γ_0 to Γ_1, consisting of exponentially many steps.

Finally, it remains to describe how to morph between embeddings that are not triangulations. Following existing work in the planar setting, we extend the given embeddings Γ_0 and Γ_1 to triangulations, and then invoke our earlier triangulation-morphing algorithm. The main difficulty is that it may not be possible to triangulate both Γ_0 and Γ_1 using the same diagonals, because corresponding faces, while combinatorially identical, may have different shapes.

Two different techniques have been proposed to overcome this hurdle in the planar setting. The first method subdivides each pair of corresponding faces into a compatible triangulation, introducing additional vertices if necessary [5, 111, 205, 206, 210]; however, this technique increases the complexity of the graph to $O(n^2)$ [17]. The second technique uses additional morphing steps to convexify faces to that they can be compatibly triangulated without additional vertices [4, 14]. While the subdivision technique generalizes to toroidal embeddings (at least when all faces are disks), it is unclear how to generalize existing morphing techniques.

We introduce a third technique, which avoids both subdivision and additional morphing

steps by exploiting Theorem 2.3. Our method can also be applied to 3-connected straight-line plane graphs, giving a new and arguably simpler approach for the planar case as well.

Theorem 4.3. Given any two isotopic essentially 3-connected geodesic toroidal embeddings Γ_0 and Γ_1 with n vertices, we can compute a piecewise-linear morph from Γ_0 to Γ_1, consisting of exponentially many steps.

Proof. Let Γ_* be an equilibrium embedding isotopic to Γ_0 and Γ_1 as given by Theorem 2.3. It suffices to describe how to morph from Γ_0 to Γ_*; to morph from Γ_0 to Γ_1 one can simply first morph from Γ_0 to Γ_* and then from Γ_* to Γ_1.

Arbitrarily triangulate the faces of Γ_0; this can be done in $O(n)$ time using Chazelle's algorithm [52], or in $O(n \log n)$ time in practice. Because each face of Γ_* is convex, we can triangulate Γ_* in the exact same manner. The result is two isotopic geodesic toroidal triangulations T_0 and T_*. Given a morph between T_0 and T_* as promised by Theorem 4.2, we obtain a morph between Γ_0 and Γ_* by simply ignoring the edges added when triangulating. In particular, the morph is specified by a sequence of geodesic triangulations $T_0, T_1, \ldots, T_k = T_*$, and dropping the additional edges from each triangulation T_i results in a geodesic embedding Γ_i isotopic to Γ_0. □

Finally, Theorem 4.3 immediately yields Corollary 4.2.

Corollary 4.2. Two essentially 3-connected geodesic embeddings Γ_0 and Γ_1 on the flat torus are isotopic if and only if they are isotopic through geodesic embeddings.

Chapter 5

Toroidal Morphing via Barycentric Interpolation

In this chapter, we once again consider the problem of computing morphs between two isotopic embeddings of the same graph on the same flat torus. Whereas in the previous chapter we extended Cairns' edge-collapse method on the plane [44, 45] to produce a toroidal morph, in this chapter we extend a different method, namely, the barycentric interpolation paradigm introduced by Floater and Gotsman [99], which was in turn based on Floater's extension (Theorem 2.4) to Tutte's spring embedding theorem.

Floater and Gotsman's barycentric interpolation method in morphs that are natural and visually appealing. However, the resulting morphs are represented implicitly; the vertex positions at any time during the morph can be computed in $O(n^{\omega/2})$ time by solving a linear system via nested dissection [8, 146].In contrast, Cairns' morph and its later improvements [4, 133] produce morphs with explicit piecewise-linear vertex trajectories; unfortunately, the reliance on perturbed edge collapses make these morphs not useful for visualization.

Our extension of Floater and Gotsman's method also provides a short and simple proof of a conjecture of Connelly *et al.* [66] on the deformation space of geodesic torus triangulations.

5.1 Prior Results (and why they don't immediately generalize)

Floater, Gotsman and Surazhsky [99, 111, 205–207] used Theorem 2.4 to obtain a morph between straight-line planar graph embeddings Γ_0 and Γ_1 as follows: Let $\lambda(0)$ and $\lambda(1)$ be barycentric stresses for Γ_0 and Γ_1, respectively. For $0 < t < 1$, set $\lambda(t) := (1-t)\lambda(0) + t\lambda(1)$, and let $\Gamma^{\lambda(t)}$ be the resulting Floater embedding; the function $t \mapsto \Gamma^{\lambda(t)}$ describes a morph between Γ_0 and Γ_1.

As observed by Steiner and Fischer [197], Floater and Gotsman's morphing algorithm does not directly generalize to the toroidal setting, because the linear system (2.12) is rank-deficient, and thus not all positive stresses λ are realizable; see Lemma 2.7 for a concrete example.

In particular, given arbitrary barycentric stress vectors $\lambda(0)$ and $\lambda(1)$ of two isotopic convex torus embeddings, intermediate weights $\lambda(t) := (1-t)\lambda(0) + t\lambda(1)$ are not necessarily realizable. Consider the toroidal embeddings of K_7 shown in Fig. 5.1, which differ only in the position

of vertex 2. We computed mean-value weights λ and μ for these embeddings, normalized so that the weights of all edges leaving each vertex sum to 1 [98, 120]. Writing out the linear

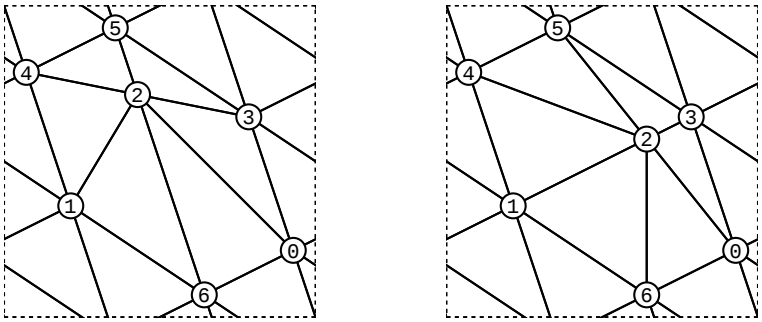

Figure 5.1. Isotopic embeddings of K_7 whose normalized mean-value weights are not morphable.

system (2.12) for the stress $(\lambda+\mu)/2$ gives

$$L^{\frac{\lambda+\mu}{2}} = \begin{pmatrix} 1 & -0.2485 & -0.0957 & -0.1305 & -0.1469 & -0.0983 & -0.28 \\ -0.2698 & 1 & -0.1974 & -0.1067 & -0.1332 & -0.1595 & -0.1333 \\ -0.0689 & -0.1045 & 1 & -0.3839 & -0.0982 & -0.2826 & -0.062 \\ -0.1744 & -0.0848 & -0.3336 & 1 & -0.2144 & -0.0662 & -0.1268 \\ -0.1486 & -0.1712 & -0.111 & -0.2513 & 1 & -0.2184 & -0.0994 \\ -0.0901 & -0.1347 & -0.2089 & -0.0806 & -0.258 & 1 & -0.2278 \\ -0.2413 & -0.1152 & -0.1191 & -0.156 & -0.1044 & -0.2639 & 1 \end{pmatrix} \tag{5.1}$$

$$H^{\frac{\lambda+\mu}{2}} = \begin{pmatrix} -0.5062 & -0.5252 \\ -0.1067 & -0.2928 \\ 0.0689 & -0.062 \\ 0.4735 & 0 \\ -0.3508 & 0.1486 \\ 0 & 0.2248 \\ 0.3457 & 0.4756 \end{pmatrix}$$

One can verify (by hand or by computer) that this system does not admit a solution.

Steiner and Fischer [197] modify the system by fixing a single vertex, restoring full rank. However, solving this modified linear system does not necessarily yield an embedding, because the fixed vertex may not lie in the convex hull of its neighbors. The top row of Fig. 5.2 shows two isotopic embeddings of a 12×12 toroidal grid, one with a single row of vertices shifted $1/2$ to the left, the other with a single column of vertices shifted $1/2$ downward. Let $\lambda^{\leftarrow}$ and $\lambda^{\downarrow}$ respectively denote the normalized mean-value weights for these embeddings [98, 120]. The bottom left image in Fig. 5.2 shows the Steiner-Fischer *drawing* for the weight $\lambda = (2\lambda^{\leftarrow}+\lambda^{\downarrow})/3$,

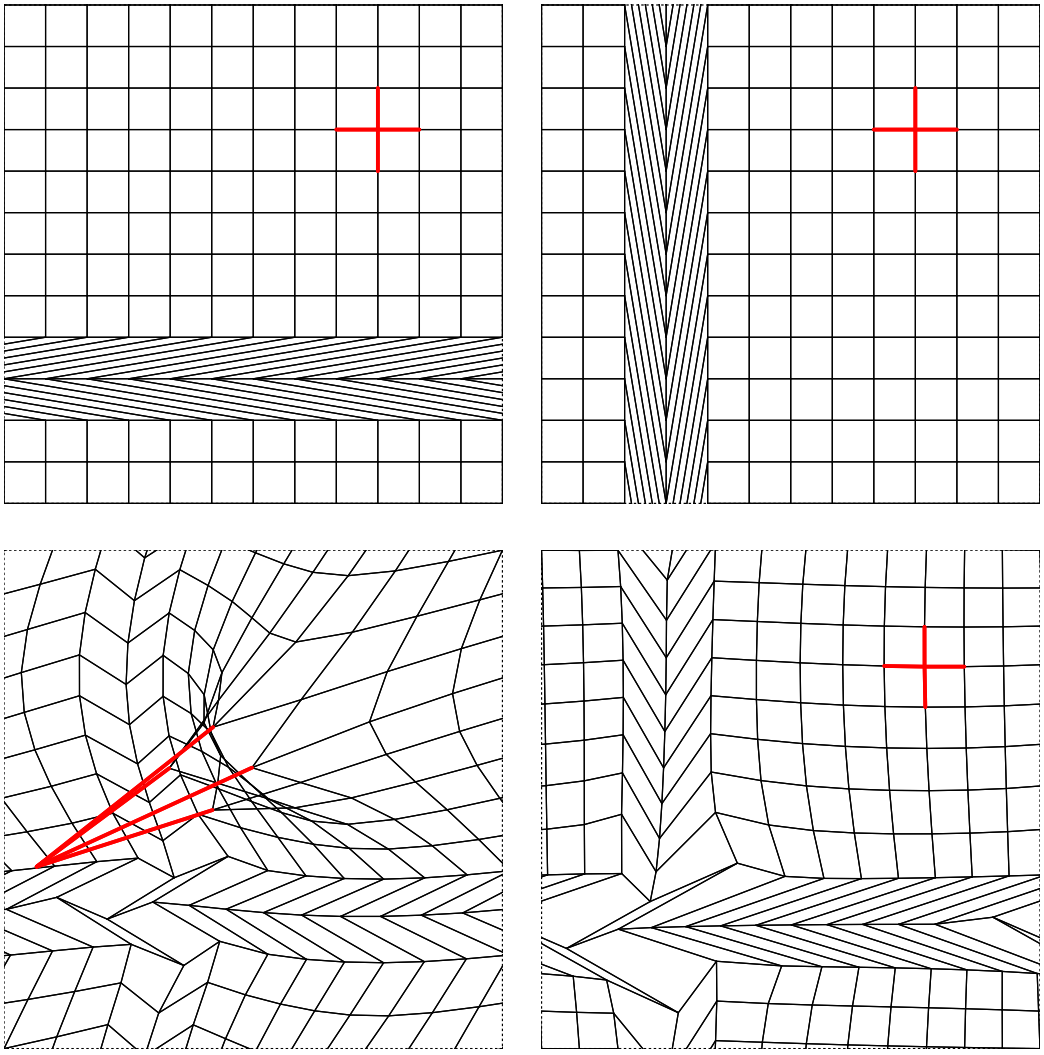

Figure 5.2. A bad example of fixed-vertex weight interpolation; see the text for explanation.

with the red edges indicating the fixed vertex. This drawing is clearly not an embedding; it also follows that the weight vector λ is not realizable.

The bottom right of Fig. 5.2 shows the corresponding Floater embedding for the realizable weight vector $\mu = (2\mu^{\leftarrow} + \mu^{\downarrow})/3$, where $\mu^{\leftarrow}$ and $\mu^{\downarrow}$ are *morphable* weights derived by rescaling $\lambda^{\leftarrow}$ and $\lambda^{\downarrow}$, as described in Lemma 5.3.

Steiner and Fischer also claim [197, Section 2.2.1] that the resulting drawing has no "foldovers" except at the fixed vertex and its neighbors. Here a "foldover" is a vertex in the drawing whose incident faces overlap (where formally, faces are determined by the reference drawing used to define the translation vectors). Close examination of Fig. 5.2 shows that this claim is incorrect: foldovers occur at more than just the fixed vertex and its neighbors.

Very recently, Luo *et al.* [152] proved a generalization of Floater's theorem to geodesic

triangulations of arbitrary closed Riemannian 2-manifolds with strictly negative curvature, extending the spring-embedding theorems of Colin de Verdière [60] and Hass and Scott [115] to asymmetric weights. Their result immediately implies that if two geodesic triangulations of such a surface are homotopic, then linearly interpolating the dart weights yields a morph. Their result applies only to surfaces with negative Euler characteristic; alas, the torus has Euler characteristic 0. The result derived in this chapter thus fills an important gap in the existing literature.

5.2 Morphable Stresses

To get around these issues, we identify a subspace of *morphable* stresses, such that every convex torus embeddings has a morphable barycentric stress, every morphable stress is realizable, and convex combinations of morphable stresses are morphable. Specifically, a *positive* stress λ is ***morphable*** if each *column* of the matrices L^λ and H^λ as defined in Equation (2.12) sums to 0.

The following lemma follows immediately from the definition of morphable stresses.

Lemma 5.1. Convex combinations of morphable stresses are morphable.

Lemma 5.2. Every morphable stress is realizable.

Proof. If λ is a morphable stress, then the nth row of the linear system $L^\lambda P = H^\lambda$ is implied by the other $n-1$ rows, so we can remove it. The resulting abbreviated linear system still has rank $n-1$, so it has a solution. □

Lemma 5.3. Given a barycentric stress λ for a convex torus embedding Γ, a *morphable* barycentric stress for Γ can be computed in $O(n^{\omega/2})$ time.

Proof. The matrix L^λ has rank $n-1$, so there is a one-dimensional space of (row) vectors $\alpha = (\alpha_1, \ldots, \alpha_n)$ such that $\alpha L^\lambda = (0, \ldots, 0)$. We can compute a non-zero vector α in $O(n^{\omega/2})$ time using general nested dissection [8, 146] and toroidal separators [6, 104].

Moreover, a directed version of the matrix tree theorem [39, 71, 212] implies that we can choose all α_i to be positive. Specifically, let $G^\pm$ be the weighted directed graph whose weighted arcs correspond to the weighted darts of G. An *inward directed spanning tree* is an acyclic spanning subgraph of $G^\pm$ where every vertex except one (called the *root*) has out-degree 1. The weight of an inward directed spanning tree is the product of the weights of its arcs. For each i, let α_i be the sum of the weights of all inward directed spanning trees rooted at vertex i;

we have $\alpha_i > 0$ because all dart weights are positive. The directed matrix tree theorem implies that $\alpha L = 0$, as required; for an elementary proof, see De Leenheer [71, Theorem 3].

Define a new positive stress μ by setting $\mu_d := \alpha_{\text{tail}(d)}\lambda_d$ for each dart d. For each i, we immediately have $L_i^\mu P = \alpha_i L_i^\lambda P = \alpha_i H_i^\lambda = H_i^\mu$, where P is the position matrix for Γ, so μ is in fact a barycentric stress for Γ. Finally, we easily observe that $(1,\dots,1)L^\mu = \alpha L^\lambda = (0,\dots,0)$ and $(1,\dots,1)H^\mu = \alpha H^\lambda = \alpha L^\lambda P = (0,\dots,0)P = (0,0)$, which imply that μ is morphable. □

Putting the preceding lemmas together gives us the following theorem, which, as a corollary, offers an alternative proof of Corollary 4.2.

Theorem 5.1. Given two isotopic essentially 3-connected geodesic torus embeddings Γ_0 and Γ_1, we can efficiently compute a morph from Γ_0 to Γ_1, whose intermediate embeddings are given by solving the linear system (2.10). Specifically, after $O(n^{\omega/2})$ preprocessing time, we can compute any intermediate embedding during the morph in $O(n^{\omega/2})$ time.

Proof. If Γ_0 and Γ_1 are convex embeddings, we proceed as follows. First, if necessary, we normalize the given coordinate representations so that their translation vectors agree, in $O(n)$ time (see Section 2.8). Then we find barycentric stresses $\lambda(0)$ and $\lambda(1)$ for Γ_0 and Γ_1, respectively, in $O(n)$ time, for example using Floater's mean-value coordinates [98, 120]. Following Lemma 5.3, we derive morphable weights $\mu(0)$ and $\mu(1)$ from $\lambda(0)$ and $\lambda(1)$, respectively, in $O(n^{\omega/2})$ time. Finally, given any real number $0 < t < 1$, we set $\mu(t) := (1-t)\mu(0) + t\mu(1)$ and solve the linear system $L^{\mu(t)}P^{(t)} = H^{\mu(t)}$ for the position matrix $P^{(t)}$ of an intermediate embedding $\Gamma^{\mu(t)}$; Lemmas 5.1 and 5.2 imply that this system is solvable. The function $t \mapsto \Gamma^{\mu(t)}$ is a morph between Γ_0 and Γ_1.

If Γ_0 or Γ_1 is not convex, we can employ the following reduction, which is similar to the one employed in the proof of Theorem 4.3: Let Γ_* be the Tutte embedding of G obtained by setting every dart weight to 1. Compute an arbitrary triangulation T_0 of Γ_0. Because every face of Γ_* is convex, we can triangulate Γ_* using the same diagonals, yielding a triangulation T_* isotopic to T_0. Assign weight 0 to the darts of each edge in $T_* \setminus \Gamma_*$ to obtain a barycentric weight vector μ_* for T_*, which is symmetric and therefore morphable. Derive morphable weights μ_0 for T_0 using, say, mean-value coordinates [98, 120] and Lemma 5.3. Then we can morph from T_0 to T_* by weight interpolation, using the weight vector $\mu(t) := (1-2t)\mu_0 + 2t\mu_*$ for any $0 \le t \le 1/2$. Ignoring the additional triangulation edges gives us a morph from Γ_0 to Γ_*. Finally, a symmetric procedure gives us a morph from Γ_* to Γ_1. □

5.3 Deformation Spaces of Geodesic Torus Triangulations

Our formulation of morphable stresses provides a straightforward solution to a conjecture of Connelly *et al.* [66] on the deformation space of geodesic triangulations. Bloch, Connelly, and Henderson [31] proved that for any planar straight-line triangulation Γ of a convex polygon P, the space of all planar straight-line triangulations of P that are homeomorphic to Γ is contractible. (Cairns' morphing theorem [44, 45] asserts only that this space is connected.) Simpler proofs of this theorem were recently given by Cerf [47] and Luo [151]; in particular, Luo observed that the Bloch–Connelly–Henderson theorem follows immediately from Floater's barycentric embedding theorem [96, 97].

Connelly *et al.* [66] conjectured that every isotopy class of geodesic triangulations on any surface S with constant curvature is homotopy-equivalent to the group $\mathrm{Isom}_0(S)$ of isometries of S that are homotopic to the identity. In particular, $\mathrm{Isom}_0(S^2)$ is the rotation group $SO(3)$ and $\mathrm{Isom}_0(\mathbb{T})$ is the translation group $S^1 \times S^1$; for every other orientable surface S without boundary, $\mathrm{Isom}_0(S)$ is trivial [171]. Very recently, Luo *et al.* proved this conjecture for all surfaces of genus at least 2 [152] and for flat tori [153]; both proofs use nontrivial extensions of Floater's theorem [96, 97].

Here we offer a simpler proof for flat tori; in fact, we prove a more general result about convex embeddings instead of just triangulations.

Theorem 5.2. For any convex embedding Γ on a flat torus $\mathbb{T}$, the space of all convex embeddings isotopic to Γ is homotopy equivalent to $\mathbb{T}$.

Proof. Fix a convex embedding Γ of a graph G with n vertices and m edges; without loss of generality, assume some vertex v is positioned at $(0,0)$. Let $X = X(\Gamma)$ denote the space of all convex s of G isotopic to Γ, and let $X_0 = X_0(\Gamma)$ be the subspace of embeddings in X where vertex v is positioned at $(0,0)$. Every embedding in X is a translation of a unique embedding in X_0, so $X = X_0 \times S^1 \times S^1$. Thus, to prove the theorem, it suffices to prove that X_0 is contractible.

Call a positive asymmetric stress $\lambda \in \mathbb{R}^{2m}_+$ for Γ *normalized* if $\sum_d \lambda_d = 1$, where the sum is over all darts of Γ. Let $R = R(\Gamma)$ denote the set of of all realizable stresses for T, and let $\overline{M} = \overline{M}(\Gamma)$ denote the set of all normalized morphable stresses for T.

Lemma 5.1 implies that $\overline{M}$ is convex and therefore contractible. (Specifically, $\overline{M}$ is the interior of a $(2m-n-2)$-dimensional convex polytope in $\mathbb{R}^{2m}$.)

Call two realizable stresses $\lambda, \lambda' \in R$ *equivalent* if there is a scaling vector $\alpha \in \mathbb{R}^n_+$ such that $\lambda'_d = \alpha_{\mathrm{tail}(d)} \lambda_d$ for every dart d. Because the Laplacian matrix L^λ has rank $n-1$, Lemma 5.3 implies that every realizable stress $\lambda \in R$ is equivalent to a *unique* normalized morphable stress $\mu \in \overline{M}$. It follows that R is homeomorphic to $\overline{M} \times \mathbb{R}^n_+$ and therefore contractible.

Now we follow the proof of Theorem 1.4 in Luo *et al.* [152]. Because every morphable stress is realizable, solving linear system (2.12) gives us a continuous map $\Phi : R \to X_0$. Floater's mean-value weights [98,120] give us a continuous map $\Psi : X_0 \to R$ such that $\Phi \circ \Psi$ is the identity map on X_0. Because R is contractible, the function $\Psi \circ \Phi$ is homotopic to the identity map on R. We conclude that X_0 is homotopy equivalent to R and therefore contractible. □

Chapter 6

Efficient Piecewise-Linear Morphs

In this chapter we introduce a new tool for computing piecewise-linear morphs: Let Γ be a convex embedding (planar or toroidal) and λ be a barycentric stress vector for Γ. We show that changing the stresses on the darts along a single edge results in a unidirectional morph.

We first use this to define a new kind of edge collapse, which we call a *spring collapse*. This new tool will allow us to improve the pseudomorph described in Section 4.4. Whereas the pseudomorph described in Section 4.4 produced a a piecewise-linear morph between toroidal embeddings, consisting of $O(2^n)$ steps, here the result is a piecewise-linear morph consisting of just $O(n)$ steps; moreover this morph can be computed in $O(n^{1+\omega/2})$ time, matching the previous state of the art for planar piecewise-linear morphs [133]. Previous methods [4, 133] used to improve Cairns' edge-collapse based pseudomorph for planar graphs [44, 45] do not obviously generalize to the flat torus, necessitating this new technology.

Next, this unidirectional barycentric interpolation allows us to design a new method for computing piecewise-linear morphs *on the plane*. Previously, Kleist *et al.* [133] improved a previous algorithm of Alamdari *et al.* [4] by reimplementing a key operation via (a simulation of) barycentric interpolation; our new method replaces edge collapses entirely with unidirectional barycentric interpolation. Avoiding edge collapses allows us to produce piecewise-linear morphs that are still good for visualization. Our resulting morph also consists of $O(n)$ steps, computed in $O(n^{1+\omega/2})$ time, matching the previous state of the art [133]. Unfortunately, this method does not appear to immediately generalize to flat tori.

6.1 Unidirectional Barycentric Interpolation

Fix a planar graph G and a convex polygon for the outer face. Let p_v^λ denote the position of vertex v in the Floater embedding Γ^λ with respect to weight vector λ. Intuitively, the following lemma states that changing the (possibly asymmetric) stresses on the darts of a single edge e moves each vertex in the Floater embedding along lines parallel to e; see Figure 6.1.

Lemma 6.1. Let λ and μ be arbitrary stress vectors such that $\lambda_d \neq \mu_d$ or $\lambda_{\mathrm{rev}(d)} \neq \mu_{\mathrm{rev}(d)}$ for some dart d, but $\lambda_{d'} = \mu_{d'}$ for all darts $d' \notin \{d, \mathrm{rev}(d)\}$. For each vertex w, the vector $p_w^\mu - p_w^\lambda$ is parallel to the embedding of d in Γ^λ.

Proof. Suppose d has tail u and head v, and (by rotating the embedding if necessary) that d is

drawn in parallel to the x-axis. For each vertex i, let y_i^λ and y_i^μ be the projections of p_i^λ and p_i^μ, respectively, along the y-axis, i.e., the direction orthogonal to d. Note that by assumption $y_u^\lambda = y_v^\lambda$. We need to prove that $y_w^\lambda = y_w^\mu$ for every vertex w.

Projecting the linear system (2.10) for λ along the y-axis gives us

$$\sum_{i \to j} \lambda_{i \to j}(y_j^\lambda - y_i^\lambda) = 0 \qquad \text{for each vertex } i. \tag{6.1}$$

Swapping entries of λ with corresponding entries of μ in the system (6.1) only changes at most two constraints, namely, the ones corresponding to the two endpoints u and v of d. Moreover, in each affected constraint, the single changed coefficient is multiplied by $y_u^\lambda - y_v^\lambda = y_v^\lambda - y_u^\lambda = 0$, so the y_i^λ's also solve the corresponding system for μ. Since the system (6.1) and its counterpart for μ each have a unique solution, we conclude that $y_w^\lambda = y_w^\mu$ for every vertex w. □

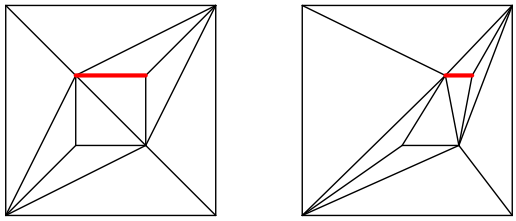

Figure 6.1. Changing the dart weights along one edge moves all vertices along lines parallel to that edge.

One can generalize Lemma 6.1 to the toroidal setting, for *realizable* stresses.

Lemma 6.2. Let λ and μ be arbitrary realizable stresses such that $\lambda_d \neq \mu_d$ or $\lambda_{\mathrm{rev}(d)} \neq \mu_{\mathrm{rev}(d)}$ for some dart d, and $\lambda_{d'} = \mu_{d'}$ for all darts $d' \notin \{d, \mathrm{rev}(d)\}$. For every vertex w, the vector $p_w^\mu - p_w^\lambda$ is parallel to the embedding of d in Γ^λ.

Proof. Suppose d has tail u and head v. By the definition of translation vectors, dart d appears in Γ^λ as the projection of a dart in the universal cover from p_u^λ to $p_v^\lambda + \tau_d$. Fix a non-zero vector $\sigma \in \mathbb{R}^2$ orthogonal to the vector $p_v^\lambda - p_u^\lambda + \tau_d$ and thus orthogonal to darts $\{d, \mathrm{rev}(d)\}$ in Γ^λ. For each vertex i, let $z_i^\lambda = p_i^\lambda \cdot \sigma$ and $z_i^\mu = p_i^\mu \cdot \sigma$, and for each dart d', let $\xi_{d'} = \tau_{d'} \cdot \sigma$. Our choice of σ implies that $\xi_d = z_u^\lambda - z_v^\lambda$. We need to prove that $z_i^\lambda = z_i^\mu$ for every vertex i.

Consider the Laplacian linear systems (2.12) corresponding to λ and μ. Let $X^\lambda = H^\lambda \cdot \sigma$ and $X^\mu = H^\mu \cdot \sigma$. The real vector $Z^\lambda = (z_i^\lambda)_i$ is a solution to the linear system $L^\lambda Z = X^\lambda$; in fact, Z^λ is the *unique* solution such that $z_n^\lambda = 0$. Similarly, $Z^\mu = (z_i^\mu)_i$ is the unique solution to an analogous equation $L^\mu Z = X^\mu$ with $z_n^\mu = 0$. We will prove that $L^\mu Z^\lambda = X^\mu$, so that in fact $Z^\lambda = Z^\mu$.

Let $\delta = \mu_d - \lambda_d$ and $\varepsilon = \mu_{\text{rev}(d)} - \lambda_{\text{rev}(d)}$. The matrices L^λ and L^μ differ in only four locations:

$$L^\mu_{ij} - L^\lambda_{ij} = \begin{cases} \delta & \text{if } (i,j) = (u,u) \\ -\delta & \text{if } (i,j) = (u,v) \\ -\varepsilon & \text{if } (i,j) = (v,u) \\ \varepsilon & \text{if } (i,j) = (v,v) \\ 0 & \text{otherwise} \end{cases}$$

More concisely, we have $L^\mu = L^\lambda + (\delta \mathbf{e}_u - \varepsilon \mathbf{e}_v)(\mathbf{e}_u - \mathbf{e}_v)^T$, where $\mathbf{e}_i$ denotes the ith standard coordinate vector. Similarly, $H^\mu = H^\lambda + \tau_d(\delta \mathbf{e}_u - \varepsilon \mathbf{e}_v)$ and therefore $X^\mu = X^\lambda + \xi_d(\delta \mathbf{e}_u - \varepsilon \mathbf{e}_v)$. It follows that

$$\begin{aligned} L^\mu Z^\lambda &= L^\lambda Z^\lambda + (\delta \mathbf{e}_u - \varepsilon \mathbf{e}_v)(\mathbf{e}_u - \mathbf{e}_v)^T Z^\lambda \\ &= X^\lambda + (\delta \mathbf{e}_u - \varepsilon \mathbf{e}_v)(z^\lambda_u - z^\lambda_v) \\ &= X^\lambda + (\delta \mathbf{e}_u - \varepsilon \mathbf{e}_v)\,\xi_d \\ &= X^\mu, \end{aligned}$$

completing the proof. □

By Lemma 4.3, linear (geodesic) interpolation of vertex positions between Γ^λ and Γ^μ where λ and μ satisfy the assumptions of Lemma 6.1 or Lemma 6.2, results in a unidirectional linear morph.

6.2 Efficient Piecewise-Linear Toroidal Morphs

In this section, we will use Lemma 6.2 to improve the piecewise-linear morph obtained in Section 4.4. Specifically, we will obtain a morph consisting of $O(n)$ unidirectional morphing steps, in $O(n^{1+\omega/2})$ time, matching the previous state of the art in the planar setting [133] (we provide a simpler planar morphing algorithm with the same bounds in Section 6.3).

6.2.1 Prior Results (and why they don't immediately generalize)

Recall that Cairns [44, 45] proved the existence of morphs between arbitrary isomorphic planar straight-line triangulations, using an inductive argument based on the idea of collapsing an edge from a low-degree vertex to one of its neighbors. Thomassen [210] extended Cairns'

proof to arbitrary planar straight-line graphs. Cairns and Thomassen's proofs are constructive, but are built from pseudomorphs consisting of an exponential number of steps.

A long series of later works [5, 13, 14, 24], culminating in papers by Alamdari *et al.* [4], and Kleist *et al.* [133], resulted in efficient algorithms to construct planar morphs with explicit piecewise-linear vertex trajectories, all ultimately based on Cairns' inductive edge-collapsing strategy. (Alamdari *et al.* [4] and Roselli [184] provide more detailed history of these results.) Given any two isomorphic straight-line embeddings of the same n-vertex planar graph, these papers describe how to construct a morph consisting of $O(n)$ *unidirectional* morphing steps. Thus, each vertex moves along a piecewise-linear path with complexity $O(n)$, and the entire morph has complexity $O(n^2)$. Moreover, such a morph can be computed in $O(n^{1+\omega/2})$. These results require several delicate arguments; in particular, to morph the initial and final embeddings so that all faces are convex, and to perturb the pseudomorphs defined by direct collapses and their reversals into true morphs. The resulting morphs contract all vertices into an exponentially small neighborhood and then expand them again, so they are not useful for visualization.

At its core, Alamdari *et al.*'s algorithm [4] relies on a procedure we will call ***unidirectional convexification***: Let ℓ be an arbitrary line in the plane. Call a polygon ℓ-monotone if its boundary consists of two paths that project injectively onto the line ℓ. If each face in an embedding Γ is ℓ-monotone, then the procedure moves each vertex along lines orthogonal to ℓ in such a way that the resulting embedding is convex.

By strategically removing edges temporarily from a triangulation, Alamdari *et al.* use unidirectional convexification to partially convexify the links of vertices of degree at most 5, so that said a vertex can be direct collapsed along the same edge in both the initial and final embeddings. This allows them to avoid the double recursion of Cairns' original pseudomorph, resulting in a pseudomorph consisting of $O(n)$ steps. The running time of their algorithm is dominated by $O(n)$ calls to Hong and Nagamochi's implementation [118] of unidirectional convexification, which takes $O(n^2)$ time each, for a total of $O(n^3)$ time.

Kleist *et al.* [133] observed that unidirectional convexification can be implemented by simulating barycentric interpolation, via solving a single instance of the linear system (2.10); Chrobak *et al.* [54] previously gave an implementation involving solving a single instance of the linear system (2.8). (Klemz [134] provides a more detailed history of these implementations.) Either implementation reduces the running time of unidirectional convexification to $O(n^{\omega/2})$, and, accordingly, reduces the running time of the algorithm of Alamdari *et al.* to $O(n^{1+\omega/2})$. Unfortunately, it is unclear how to reproduce unidirectional convexification on flat tori; see Section 7.4.

6.2.2 Spring Collapses

Fix a toroidal triangulation Γ and a *symmetric* stress λ. Let e_{uv} be an arbitrary edge of Γ with endpoints u and v, and let Γ' be the result of direct collapsing u to v along e_{uv}. Let a, b, c, d be the edges in the star of uv in Γ, as shown in Figure 6.2. Edges a and b collapse to a single edge ab in Γ', and edges c and d collapse to a single edge cd in Γ'. Now define a stress λ' for the edges of Γ' as follows:

$$\lambda'_e := \begin{cases} \lambda_a + \lambda_b & \text{if } e = ab \\ \lambda_c + \lambda_d & \text{if } e = cd \\ \lambda_e & \text{otherwise} \end{cases} \tag{6.2}$$

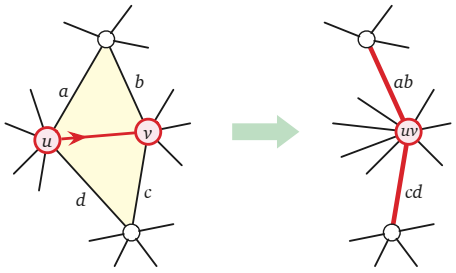

Figure 6.2. Collapsing u to v.

The new equilibrium embedding $(\Gamma')^{\lambda'}$ has the same image as the limit of Γ^λ as we increase $\lambda_{e_{uv}}$ to infinity and keep all other edge weights fixed. Since Theorem 2.3 implies that symmetric stresses are always realizable, Lemma 6.2 implies that as we continuously increase $\lambda_{e_{uv}}$, the vertices of Γ^λ move continuously along geodesics parallel to e_{uv}.

This continuous deformation from Γ^λ to $(\Gamma')^{\lambda'}$ is a unidirectional pseudomorph, but not a *linear* pseudomorph, because the vertices do not necessarily move at fixed speeds. We can simply simulate the operation by moving each vertex at constant speed along its corresponding geodesic instead; Lemma 4.3 immediately implies that this is a unidirectional linear pseudomorph. The resulting linear pseudomorph will be referred to as a ***spring collapse***.

6.2.3 Efficient Pseudomorph Between Triangulations

Let us first recall the pseudomorph derived in Section 4.4. Given isotopic geodesic triangulations Γ_0 and Γ_1, there exists some good vertex in each triangulation. However, it is possible that no vertex is good in both Γ_0 and Γ_1. More subtly, even if some vertex u is good in both triangulations, that vertex may be collapsible along a unique edge e_0 in Γ_0 but along a different unique edge e_1 in Γ_1. We thus introduced an intermediate triangulation Γ_* in which a good

vertex u of Γ_0 can be directly collapsed along both e_0 and e_1. Recursively constructing pseudomorphs from Γ_0 to Γ_* and from Γ_* to Γ_1 yields a pseudomorph from Γ_0 to Γ_1 with exponentially many steps.

We will improve this pseudomorph by using spring collapses, in a way that will allow us to avoid the exponential blowup of Cairns' algorithm. Suppose we are given two isotopic geodesic triangulations Γ_0 and Γ_1. We explain how to compute a pseudomorph Ψ_0 from Γ_0 to any isotopic equilibrium triangulation Γ_*, whose existence is promised by Theorem 2.3. The same algorithm gives a pseudomorph Ψ_1 from Γ_1 to Γ_*, and concatenating Ψ_0 with the reversal of Ψ_1 yields the desired pseudomorph from Γ_0 to Γ_1.

If Γ_0 is a zipper, we morph directly between Γ_0 and an equilibrium zipper Γ_* using at most two unidirectional linear morphs. This is the base case of our recursive algorithm.

If Γ_0 is not a zipper, then it contains a good vertex u. By definition, u can be directly collapsed along some edge e to another vertex v, without introducing edge crossings. This direct collapse gives us a unidirectional linear pseudomorph from Γ_0 to Γ_0', a geodesic toroidal triangulation whose underlying graph G' has $n-1$ vertices. On the other hand, performing a spring collapse in Γ_* by increasing the weight of the same edge e to ∞ leads to a *drawing* where u and v coincide, that is, an equilibrium triangulation Γ_*' of G' that is isotopic to Γ_0'. Finally, because Γ_0' and Γ_*' are isotopic embeddings of the same graph G', we can compute a pseudomorph from Γ_0' to Γ_*' recursively.

Our full pseudomorph from Γ_0 to Γ_* thus consists of the direct collapse from Γ_0 to Γ_0', followed by the recursive pseudomorph from Γ_0' to Γ_*', followed by the reverse of the spring collapse from Γ_* to Γ_*'. See Figure 6.3.

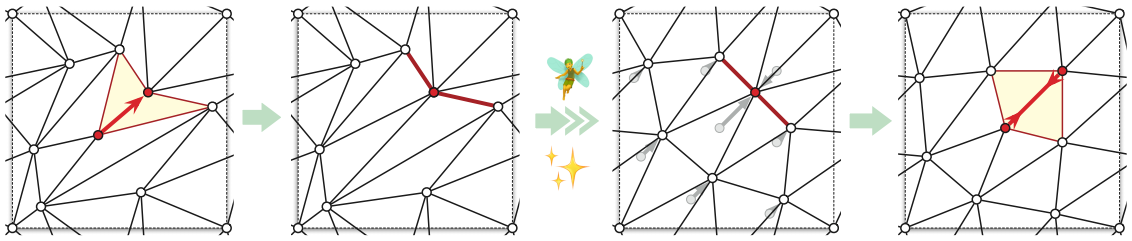

Figure 6.3. Our pseudomorph consists of a direct collapse, a recursive pseudomorph, and a reversed spring collapse.

Altogether our pseudomorph consists of $O(n)$ unidirectional linear pseudomorphs: at most n direct collapses to reach a zipper, $O(1)$ unidirectional linear pseudomorphs to reach an equilibrium zipper, and finally at most n reversed spring collapses. The time to compute the overall pseudomorph is dominated by the time needed to compute the $O(n)$ equilibrium embeddings; the overall running time is $O(n^{1+\omega/2})$.

6.2.4 From Triangulation Pseudomorphs to General Morphs, Again

We will once again use the perturbation technique of Alamdari *et al.* [4]. Recall that our analysis of the technique in Section 4.5 revealed that no part of their algorithm relied on vertex u being *directly* collapsed to v, only that vertices u and v have the same image in the triangulations Γ'_i and Γ'_{i+1}. We can thus apply this technique to our spring collapses as well.

Our pseudomorph consists of a direct collapse, a recursively computed pseudomorph, and a reversed spring collapse. First we (recursively) perturb the recursive pseudomorph into a proper morph $\Gamma'_0 \rightsquigarrow \cdots \rightsquigarrow \Gamma'_k$ consisting of k unidirectional linear morphs. We then compute the sectors S_i and the radius ε exactly as described in Section 4.5. To perturb the initial direct collapse from u to v, we move u to an arbitrary point in the intersection of S_0 and the edge e along which u is collapsed. We compute the intermediate positions u_i for u exactly as described before, working entirely within the local coordinates of the sectors S_i. Finally, to perturb the reversed spring collapse, we first move u from u_k to a new point u'_k in the visibility kernel so that the image of the collapsing edge e becomes parallel to the direction of the original spring collapse, after which we simply interpolate to the final triangulation. Because the vertices move along parallel geodesics, Lemma 4.3 implies that this final interpolation is a unidirectional linear morph. Altogether, we obtain a morph consisting of $k+3$ unidirectional linear morphs. We emphasize that the additional step moving u_k to u'_k is the only significant difference from the algorithm presented by Alamdari *et al.* [4].

Unrolling the recursion, we can perturb our pseudomorph between two n-vertex toroidal triangulations into a proper morph consisting of $O(n)$ unidirectional linear morphing steps in $O(n^2)$ time. The overall time to compute this morph is still dominated by the time needed to compute $O(n)$ equilibrium triangulations for the spring collapses.

Theorem 6.1. Given any two isotopic geodesic toroidal triangulations Γ_0 and Γ_1 with n vertices, we can compute a piecewise-linear morph from Γ_0 to Γ_1, consisting of $O(n)$ unidirectional steps, in $O(n^{1+\omega/2})$ time.

Finally, we will generalize from triangulations to general embeddings, using the same technique as in the proof of Theorem 4.3.

Theorem 6.2. Given any two isotopic essentially 3-connected geodesic toroidal embeddings Γ_0 and Γ_1 with n vertices, we can compute a piecewise-linear morph from Γ_0 to Γ_1, consisting of $O(n)$ unidirectional steps, in $O(n^{1+\omega/2})$ time.

Proof. Let Γ_* be an equilibrium embedding isotopic to Γ_0 and Γ_1 as given by Theorem 2.3. It suffices to describe how to morph from Γ_0 to Γ_*; to morph from Γ_0 to Γ_1 one can simply first morph from Γ_0 to Γ_* and then from Γ_* to Γ_1.

Arbitrarily triangulate the faces of Γ_0; this can be done in $O(n)$ time using Chazelle's algorithm [52], or in $O(n \log n)$ time in practice. Because each face of Γ_* is convex, we can triangulate Γ_* in the exact same manner. The result is two isotopic geodesic toroidal triangulations T_0 and T_*. Given a morph between T_0 and T_* as promised by Theorem 6.1, we obtain a morph between Γ_0 and Γ_* by simply ignoring the edges added when triangulating. In particular, the morph is specified by a sequence of geodesic triangulations $T_0, T_1, \ldots, T_k = T_*$, and dropping the additional edges from each triangulation T_i results in a geodesic embedding Γ_i isotopic to Γ_0.

The number of unidirectional morphing steps remains $O(n)$, and the running time is dominated by the computation of the morph between T_0 and T_*, which is $O(n^{1+\omega/2})$ by Theorem 6.1. □

6.3 Morphing Planar Graphs Edge by Edge

In this section, we describe a simple morph between planar straight-line graphs that combines the benefits of both the Floater and Gotsman style [99, 111, 205–207] and the Cairns style [4, 44, 45, 133, 210] of morphs. In particular, we construct a morph with explicit piecewise-linear vertex trajectories, while still being good for visualization by avoiding edge collapses. Our morph consists of at most $4n - 12 = O(n)$ morphing steps, computed in $O(n^{1+\omega/2})$ time, matching the previous state of the art [133]. Our algorithm is also *significantly* simpler than previous methods for computing piecewise-linear morphs.

The discussion of Section 6.1 implies that we can morph between isomorphic convex embeddings one edge at a time, where for each edge we get a unidirectional linear morphing step. The resulting algorithm is shown in Figure 6.4. The initial and final weight vectors can be found in $O(n)$ time using, for example, Floater's mean-value method [98, 120]. Each intermediate embedding can be computed in $O(n^{\omega/2})$ time using generalized nested dissection and planar separators [8, 146].

Lemma 4.3 implies that the algorithm presented in Figure 6.4 actually produces a *convexity-preserving* piecewise-linear morph, i.e., all faces remain convex throughout the morph. The existence of convexity-preserving morphs was first proven by Thomassen [210]. Angelini *et al.* [15] described an algorithm that produces a convexity-preserving piecewise-linear morph consisting of $O(n)$ steps; however, they do not analyze the runtime of their algorithm. Our algorithm is significantly simpler than that of Angelini *et al.* [15]; moreover, since Euler's formula implies that a planar graph at most $3n - 6$ edges, of which at least three must bound the outer face, our algorithm produces a morph consisting of at most $3n - 9$ steps, in $O(n^{1+\omega/2})$

MorphConvex($\Gamma_{\text{start}}, \Gamma_{\text{end}}$):
  $\lambda \leftarrow$ barycentric stress for Γ_{start}
  $\mu \leftarrow$ barycentric stress for Γ_{end}
  $k \leftarrow 0$
  for each internal edge e
    $k \leftarrow k+1$
    $d \leftarrow$ a dart of e
    $\lambda_d \leftarrow \mu_d$
    $\lambda_{\text{rev}(d)} \leftarrow \mu_{\text{rev}(d)}$
    $\Gamma_k \leftarrow \Gamma^{\lambda}$
  return $\Gamma_{\text{start}}, \Gamma_1, \Gamma_2, \ldots, \Gamma_k$ $\langle\langle = \Gamma_{\text{end}} \rangle\rangle$

Figure 6.4. Algorithm for morphing between convex planar embedding.

time.

Theorem 6.3. Given any two isomorphic convex planar drawings Γ_0 and Γ_1 with n vertices and the same convex outer face, we can compute a morph from Γ_0 to Γ_1 consisting of at most $3n-9$ unidirectional linear morphing steps, in $O(n^{1+\omega/2})$ time.

The algorithm can easily be extended to embeddings with *non-convex* faces by first morphing to embeddings with convex faces: Add edges to decompose every face into convex polygons, compute barycentric weights for the resulting embedding, and then reduce the weights of each added edge (one-by-one) to zero; after dropping the added edges, this results in a piecewise-linear morph to a convex embedding, thus reducing to the previously solved case. (This is the same idea as used in the proofs of Theorems 4.3, 5.1, and 6.2.) As with the convex case, each intermediate embedding can be computed in $O(n^{\omega/2})$ time. The full procedure is shown in Figure 6.5.

Convexify(Γ):
  $\Gamma' \leftarrow$ convex decomposition of Γ
  $\lambda \leftarrow$ barycentric stress for Γ'
  $k \leftarrow 0$
  for each *new* edge e in Γ'
    $k \leftarrow k+1$
    $d \leftarrow$ a dart of e
    $\lambda_d \leftarrow 0$
    $\lambda_{\text{rev}(d)} \leftarrow 0$
    $\Gamma_k \leftarrow \Gamma^{\lambda}$ without new edges
  return $\Gamma, \Gamma_1, \Gamma_2, \ldots, \Gamma_k$

Morph($\Gamma_{\text{start}}, \Gamma_{\text{end}}$):
  $S_{\text{before}} \leftarrow$ Convexify(Γ_{start})
  $\Gamma_{\text{before}} \leftarrow$ last embedding in S_{before}
  $S_{\text{after}} \leftarrow$ Convexify(Γ_{end})
  $\Gamma_{\text{after}} \leftarrow$ last embedding in S_{after}
  $S_{\text{convex}} \leftarrow$ MorphConvex($\Gamma_{\text{before}}, \Gamma_{\text{after}}$)
  return $S_{\text{before}}, S_{\text{convex}}, \text{reverse}(S_{\text{after}})$

Figure 6.5. Algorithm for morphing between general planar straight-line embedding.

In total, we perform one morphing step for each internal edge of G, plus at most $2(k-3)$

morphing steps for each bounded face with degree k. More specifically, to convexify the initial and final faces, it suffices to add at most one additional edge for each reflex angle; on the other hand, each internal vertex can only support at most one reflex angle around it, so we need at most $n-3$ additional edges. Our morph thus consists of at most $4n-12$ linear morphing steps. In summary:

Theorem 6.4. Given any two isomorphic 3-connected planar straight-line embeddings Γ_0 and Γ_1 with n vertices and the same convex polygon for its outer face, we can compute a morph from Γ_0 to Γ_1 consisting of at most $4n-12$ unidirectional linear morphing steps, in $O(n^{1+\omega/2})$ time.

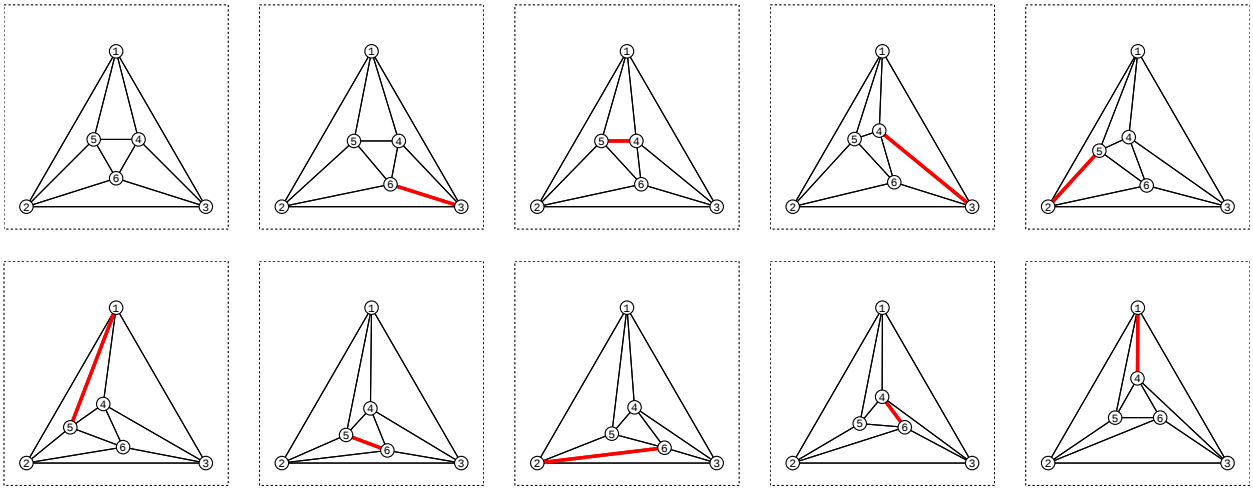

Figure 6.6. Morphing between twists of the Schönhardt triangulation edge-by-edge. In each step, the edge whose stresses are being modified are shown in red.

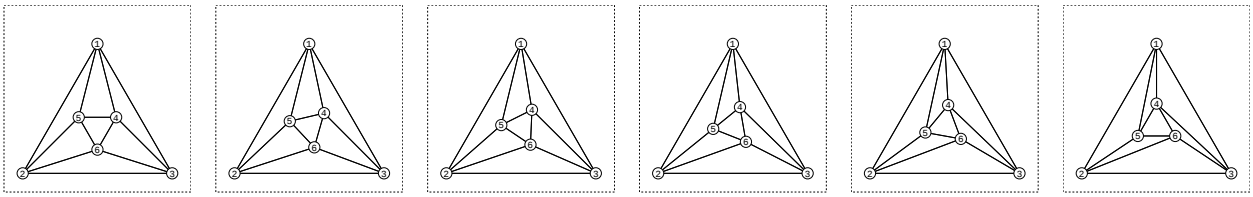

Figure 6.7. Morphing between twists of the Schönhardt triangulation via barycentric interpolation.

Figures 6.6 and 6.7 display morphs between planar graph embeddings obtained by the unidirectional barycentric interpolation method described in this section, and Floater and Gotsman's barycentric interpolation method as discussed in Chapter 5, respectively.

6.4 Models of Computation

We conclude with a comment about the models of computation implicitly assumed by the algorithm of Alamdari *et al.* [4] and the algorithms described in this chapter.

As explained in Section 4.5, the perturbation strategy of Alamdari *et al.* for converting pseudomorphs into morphs require a slightly nonstandard real RAM model of computation that supports exact square and exact cube roots. Since the toroidal morphing algorithms of Chapter 4 and Section 6.2 use this perturbation as a subroutine, they also require such a model.

In contrast, Floater's mean-value weights can be expressed in terms of areas and Euclidean lengths [120], which require only square roots to evaluate exactly. If initial and final barycentric weights are given, Floater and Gotsman's morphing algorithm [99], our toroidal extension in Chapter 5, and the algorithm described in Section 6.3 use only basic arithmetic operations: addition, subtraction, multiplication, and division.

Even without exact roots, any integer-RAM or floating-point implementation of our morphing algorithm must contend with precision issues. A careful implementation of Alon and Yuster's nested dissection algorithm [8] solves Floater's linear system (2.10) exactly in $O(n^{1+\omega/2} \operatorname{polylog} n)$ bit operations, assuming all dart weights $\lambda_{u\to v}$ are $O(\log n)$-bit integers [8, 22]. Thus, at least then all weights are given as part of the input, an exact implementation of our morphing algorithm runs in $O(n^{2+\omega/2} \operatorname{polylog} n)$ time on a standard integer RAM. However, coordinates of Tutte/Floater drawings can require $\Omega(n)$ bits of precision to avoid collapsing or crossing edges; see Eades and Garvan [81] or Figure 6.8. Thus, the near-linear cost of exact arithmetic is unavoidable in the worst case.

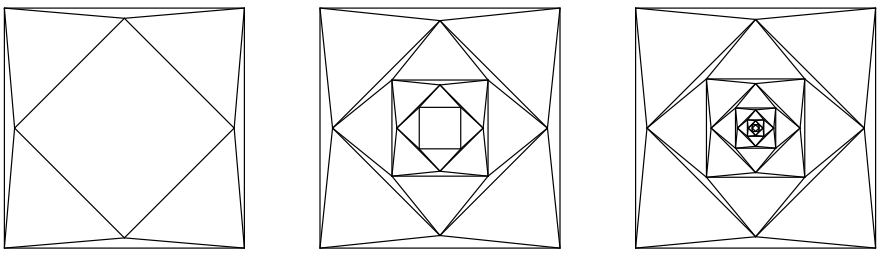

Figure 6.8. A family of graph embeddings consisting of nested squares, with 2, 5, and 10 layers, respectively.

Shen, Jiang, Zorin, and Panozzo [193] observe that floating-point implementations of Tutte's algorithm suffer from robustness issues in practice. Shen *et al.* describe an iterative procedure to repair floating-point-Tutte drawings; however, it is unclear whether a similar procedure can be used to avoid precision issues in our algorithm while maintaining continuity of the resulting morph. It is also unclear whether precision issues in our algorithm can be avoided, or at least minimized, by carefully choosing the order in which edge weights are changed and/or by

morphing through a carefully chosen intermediate drawing.

Chapter 7

Open Questions

Research is never done: every result discovered only uncovers more questions. As such, we collect in this chapter a small list of open questions that arise from the results previously discussed in this thesis.

7.1 Variations on a Maxwell–Cremona–Delaunay Theme

One can make one of many minor tweaks to the question of when there exists a correspondence between equilibrium stress and some form of reciprocal diagram.

In the plane, the angle between an edge and its corresponding edge in a reciprocal diagram is irrelevant; one can always rotate. We found that on flat tori, the picture for orthogonal reciprocal diagrams and parallel reciprocal diagrams are quite different; in particular, we suspect that a more general theory can be derived that fully characterizes when equilibrium stresses are reciprocal for all angles between 0° (parallel) and 90° (orthogonal).

We can also look at the other parts of the Maxwell–Cremona(–Delaunay) correspondence.

For the most part, Chapter 3 ignored polyhedral liftings of flat torus graphs. Borcea and Streinu [38] investigated equilibrium stresses that induce liftings into $\mathbb{T}\times\mathbb{R}$, and gave necessary and sufficient conditions for this to occur. A natural question to ask is what the reciprocal (or rather, force) diagrams look like for equilibrium stresses that induce such liftings. It is easily observable that a periodic polyhedral lift cannot be convex, and so some edge stresses must be negative. Another question is whether an equilibrium stress that induces a lifting into $\mathbb{T}_{\square}\times\mathbb{R}$ also induces a lifting into $\mathbb{T}_M\times\mathbb{R}$ for any other 2×2 non-singular matrix M.

We also pose as a line of inquiry whether there exists a more general correspondence between parallel (or more general) reciprocality and some variant of coherence, although, as previously noted, parallel reciprocality is an affine property (whereas orthogonal reciprocality is a conformal property), and so any form of coherence corresponding to parallel reciprocality will also need to be an affine property.

The results in Chapter 3 summarize some important first steps towards a generalization of the Maxwell–Cremona–Delaunay correspondence to higher genus, namely, the torus. As previously noted, Y. Colin de Verdière [60] and Hass and Scott [115] actually extended Tutte's spring embedding theorem [213] to all higher genus surfaces equipped with a metric of non-

positive curvature. It is well-known that every surface of genus greater than 1 can be endowed with a hyperbolic metric, of negative curvature (see, e.g., Farb and Margalit [92]), and it was recently shown that weighted Delaunay triangulations can be defined naturally in such settings [34, 35, 82]. Finally, Izmestiev recently described a natural extension of the planar Maxwell–Cremona correspondence to the hyperbolic plane [128], giving a nice starting point for trying to look for generalization to hyperbolic surfaces.

In terms of higher dimension, Maxwell had already extended his result to graphs drawn in $\mathbb{R}^3$ [158]; notions of weighted Delaunay *tetrahedralizations* and their corresponding weighted Voronoi diagrams have also been studied (see, e.g., Edelsbrunner [83]). In rigidity theory, frameworks in d-dimensional tori have been studied [185]. It is reasonable to wonder, then, what Maxwell–Cremona theory looks like in said d-dimensional tori.

7.2 On Morphing Less Connected Graphs

Because they rely heavily on Theorems 2.3 and 2.4, the algorithms described in Chapters 5 and 6 require that the input embeddings be (essentially) 3-connected. If a given embedding Γ is *not* (essentially) 3-connected, an isotopic equilibrium drawing Γ_* still exists, but it may not be an embedding; nontrivial subgraphs can collapse to geodesics or even points in Γ_*. Morphing less connected toroidal embeddings remains an open problem.

On the other hand, one of the accomplishments of Alamdari *et al.*'s planar morphing algorithm [4] was the ability to morph graphs that are *disconnected*, as long as the connected components have the same nesting structure. Their method is based on a first morphing so that a compatible triangulation can be found; this method relies heavily on unidirectional convexification. This method can be used to augment the planar morphing algorithm described in Section 6.3, but, as stated before, the operation does not seem to have an obvious generalization to the toroidal setting, meaning that we would need other methods to solve the problem of morphing disconnected graphs on flat tori.

One of the most obvious approaches is to arbitrarily triangulate and then morph towards an equilibrium embedding, possibly wherein the added edges have a stress of 0; this idea was used in the proofs of Theorems 4.2, 5.1, 6.2, and 6.4 to reduce to the case of triangulations.

7.3 On Morphing Graphs Edge by Edge

In Section 6.3, we obtained a simple, efficient "best-of-both-worlds" planar morph that simultaneously obtains nice properties enjoyed by both the Floater and Gotsman style [99, 111,

205–207] and the Cairns style [4, 44, 45, 133, 210] of morphs. However, in the toroidal case, while we were able to obtain both a Floater and Gotsman style morph (Chapter 5) and Cairns style morph (Section 6.2), it remains an open question if we can combine the two techniques into another "best-of-both-worlds" result for flat tori. We make some partial progress towards this result in Section 8.1.3.

A detail that we did not elaborate on in Section 6.3 is the fact that, in no part of the description of the algorithm did we specify an ordering on the edges. Indeed, correctness of Theorem 6.4 does not depend on the ordering; however, different orderings on the edges produce different morphs, and the subjective quality of the resulting morph may depend on the ordering.

Figure 7.1(a) shows snapshots from morphing via a random permutation on the edges. In Figure 7.1(b), on the other hand, the ordering on the edges was hand-picked: edges are ordered from innermost to outermost, in counterclockwise order within each "layer". For reference, intermediary snapshots from Floater and Gotsman's barycentric interpolation method is shown in Figure 7.2.

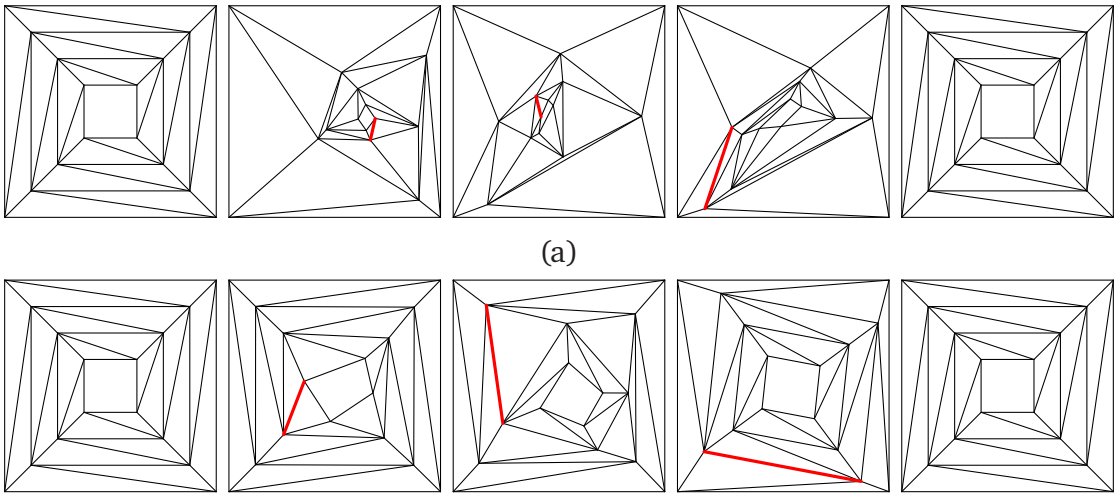

(a)

(b)

Figure 7.1. Morphing edge-by-edge, (a) in random order, and (b) from inside out.

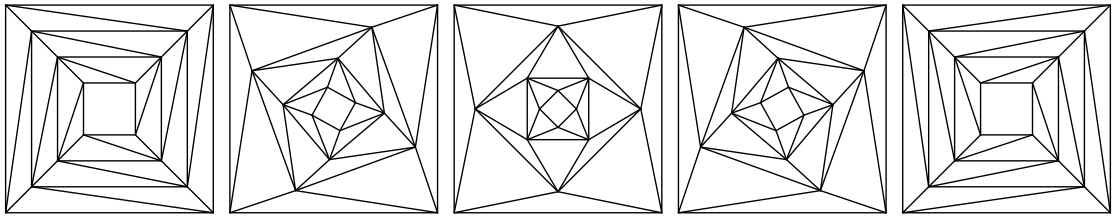

Figure 7.2. Morphing via barycentric interpolation.

An interesting practical question would be to determine what, if any, an ideal ordering on

the edges would be for visualization purposes, and how to compute such an ordering.

7.4 Toroidal Unidirectional Convexification

Recall that the algorithm of Alamdari *et al.* [4] made repeated use of a procedure we call *unidirectional convexification* to partially convexify the links of vertices, in order to avoid the double recursion of Cairns' original method. Intuitively, the procedure produces a unidirectional convexifying morph wherein vertices move along lines parallel to an *arbitrarily chosen* valid direction (see Section 6.2.1 for a description of the valid directions).

Although the use of spring collapses obviated the need to find a toroidal implementation of unidirectional convexification in Section 6.2, possible extensions of the procedure to a toroidal setting would be of independent interest. For example, motivated by studies that show that orthogonal movements are good for visualization purposes [16, 155, 176], Kleist *et al.* [133] made use of unidirectional convexification to produce a convexifying morph consisting of unidirectional morphing steps along only the horizontal and vertical directions; it would be nice to be able to produce such a morph in a toroidal setting. Section 8.1.2 describes a less ambitious application for unidirectional convexification on flat tori. Currently, we have no tools for producing a unidirectional morph on flat tori that wherein vertices do not move in parallel to the direction of an edge in a given embedding.

Hong and Nagamochi's implementation [118] of unidirectional convexification explicitly assumes that the initial planar embedding has a convex outer face, and compute new positions for each vertex by decomposing the embedding into three convex pieces and recursing on each piece; there is no obvious corresponding substructure for toroidal embeddings. The implementation of Kleist *et al.* [133] involves setting up a very specific instance of the linear system (2.10); as previously discussed, the corresponding system (2.11) for flat tori need not be solvable.

Finally, Chrobak *et al.*'s original implementation [54] of the procedure involves setting up a specfic instance of the linear system (2.8). Since the corresponding toroidal system (2.9) is solvable, this implementation seems to provide a more approachable avenue for generalization. The implementation works as follows: The input to the procedure is a planar embedding with a fixed convex polygon for its outer face and an arbitrary line ℓ; by rotating the embedding if necessary, we may assume that ℓ is the y-axis. Chrobak *et al.* compute an positive symmetric stress such that the linear system (2.8), when projected along the x-axis, is solved by the x-coordinates of the vertices. Thus solving the unprojected system gives new y-coordinates for the interior vertices, and by Lemma 4.3, linear interpolation between the old and new coordinates gives a unidirectional linear morph. (Kleist *et al.*'s implementation [133] instead

computes a positive *asymmetric* stress such that projection of the linear system (2.10) is satisfied by the x-coordinates.)

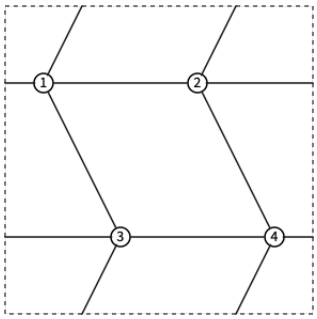

Figure 7.3. A toroidal embedding that does not admit an equilibrium stress describing its x-coordinates.

The issue in the toroidal setting is that, unlike on the plane, it is not always possible to compute an equilibrium stress describing (say) the x-coordinates. Suppose for the sake of argument that such a stress exists for the graph shown in Figure 7.3. The two edges connecting vertex 1 to vertex 3 pull towards the right. Thus the edge connecting vertex 1 to vertex 2 towards the left (corresponding to the dart $1\to2$ with homology vector $(-1,0)$) must have a larger stress than the edge connecting vertex 1 to vertex 2 towards the right (corresponding to the dart $1\to2$ with homology vector $(0,0)$). On the other hand, one can give a symmetric argument for vertex 2: since the edges connecting vertex 2 to vertex 4 pull towards the right, the edge connecting vertex 2 to vertex 1 (corresponding to the dart $2\to1$ with homology vector $(0,0)$) must have a larger stress than the edge connecting vertex 2 to vertex 1 towards the right (corresponding to the dart $2\to1$ with homology vector $(1,0)$). But both of these constraints cannot be simultaneously satisfied, so the stress cannot exist.

7.5 Generalizing Other Planar Morphs

In this thesis, we generalized two kinds of planar morphs; specifically, Floater and Gotsman's barycentric interpolation, and Cairns style edge-collapse based morphs. However, many other kinds of planar morphing techniques exist that we did not discuss, and it is worth asking if any of them generalize. We list a few of these other techniques below.

Barerra-Cruz *et al.* [25] describe an algorithm to morph between two isomorphic *weighted Schnyder drawings* of the same triangulation, each determined by a Schnyder wood together with an assignment of positive weights to the faces. The resulting morph consists of $O(n^2)$ steps, where after each step, all vertices lie on a $6n \times 6n$ integer grid. The algorithm relies crucially on the fact that the set of Schnyder woods of a planar triangulation is a distributive

lattice [94]; although toroidal versions of Schnyder woods exist [108], they do not share this property. Furthermore, despite some initial progress by Barerra-Cruz [23], it is still an open question whether this algorithm can be extended to arbitrary planar triangulations, or even to arbitrary planar straight-line graphs.

Biedl *et al.* [29] initiated a line of work [30, 105, 106] investigating the case of morphing *orthogonal* planar graph embeddings, i.e., embeddings in which edges must be drawn as horizontal or vertical line segments; naturally, morphs involved must preserve orthogonality. A number of recent results have appeared about how to find *orthoradial* embeddings, i.e., orthogonal embeddings on a flat cylinder [26, 169]; as such, one possible line of inquiry is to compute morphs between orthoradial embeddings. However, before asking for morphs between orthogonal toroidal embeddings, one must first investigate how to *obtain* such embeddings: standard methods for finding orthogonal planar embeddings [102, 208, 209] do not seem to generalize.

All of the previously mentioned morphing algorithms restricted our attention to finding straight-line deformations between straight-line embeddings. Lubiw and Petrick [150] relaxed the straight-line condition, constructing morphs that allowed edges to "bend", as long as they remained *piecewise* linear. On the other hand, Angelini *et al.* [12] recently relaxed the *embedding* condition: they considered a small family of straight-line *drawings*, and produced straight-line morphs that preserve edge crossings, i.e., two edges cross at some point during the morph if and only if they remain crossed at *all* points during the morph. It would be interesting to generalize either relaxation (or any others) to flat tori or beyond.

Chapter 8

Flotsam and Jetsam

In this chapter we collect some miscellaneous thoughts that do not fit elsewhere in the thesis. In particular, we include some half-baked ideas that did not pan out, as well as some useless quips and trivia.

8.1 Half-Baked or Incomplete Ideas

8.1.1 Toroidal Coherent Morphs

One of our motivations for investigating toroidal Maxwell–Cremona correspondences was based on an observation of É. Colin de Verdière *et al.* [57], who considered restricting Floater and Gotsman's method for morphing planar graphs to *symmetric* stresses. In other words, they required that the given embeddings Γ_0 and Γ_1 admit equilibrium stresses, and since convex combinations of symmetric stresses are symmetric, the resulting morph is a morph through *Tutte* embeddings. É. Colin de Verdière *et al.* noted that in light of the planar Maxwell–Cremona–Delaunay correspondence, this results in a ***coherent morph***, i.e., a morph between coherent (or weighted Delaunay) graphs such that, at all times during the morph, the embeddings remain coherent.

Naturally, one can ask if, in light of the results of Chapter 3, the same operation results in a coherent morph on a flat torus. Specifically, suppose we are given Γ_0 and Γ_1 on some flat torus $\mathbb{T}$ with equilibrium stresses $\lambda(0)$ and $\lambda(1)$, respectively. For $t \in (0,1)$, set $\lambda(t) = (1-t)\lambda(0) + t\lambda(1)$. Each $\lambda(t)$ is a positive symmetric stress. By Theorem 2.3, such stresses are *always* realizable, giving us a morph between Γ_0 and Γ_1 through Tutte embeddings.

However, even when Γ_0 and Γ_1 are coherent and $\lambda(0)$ and $\lambda(1)$ are *orthogonal reciprocal* stresses for Γ_0 and Γ_1, respectively, it is not clear that each $\lambda(t)$ would be an orthogonal reciprocal stress for the corresponding intermediate embedding $\Gamma^{\lambda(t)}$. In particular, it is unknown how the values $\alpha(t), \beta(t), \gamma(t)$ as defined in Equation (3.1) corresponding to each $\lambda(t)$ depend on the parameter t.

8.1.2 Convexification of Star-Shaped Hexagons

Recall that one of the major roadblocks in trying to generalize the algorithm of Alamdari *et al.* [4] to the toroidal setting is the inability to implement unidirectional convexification on flat tori; see Section 6.2.1 for more details. In this section, we suppose the existence of a toroidal implementation of unidirectional convexification, and describe a potential way to use said hypothetical implementation in order to generalize Alamdari *et al.*'s result. Had this idea panned out, it would have obviated the need for the "Cats and Dogs" analysis in Section 4.3, and given direct extension of Alamdari *et al.*'s algorithm instead.

Recall that when trying to recursively compute a pseudomorph between triangulations Γ_0 and Γ_1, Cairns' method requires two recursive calls whenever a vertex u can be direct collapsed to some neighbor v in Γ_0, but can only be direct collapsed to a different neighbor w in Γ_1. Alamdari *et al.* make use of unidirectional convexification to bypass this issue as follows: Since u cannot be collapsed to vertex w in Γ_0, it must be because there is some vertex x blocking u from collapsing to w. First, collapse u to v, then identify a quadrilateral containing v, w, and x. Temporarily delete the internal chord of this quadrilateral, and apply unidirectional convexification in the direction of, say, the temporarily deleted chord: this convexifies the quadrilateral. Now uncollapse u from v; since the problematic quadrilateral has been convexified, we can now collapse u to w. In summary, the potential extra recursive call is replaced by at most three unidirectional linear (pseudo)morphing steps.

Whereas Alamdari *et al.* used unidirectional convexification to convexify *quadrilaterals*, we now explain how to use it to convexify *hexagons*. Effectively, we would choose an *arbitrary* vertex u of degree 6 and convexify its link in both Γ_0 and Γ_1, thus allowing us (hypothetically) to collapse u to an arbitrary neighbor v. Specifically, we analyze the possible links of a degree 6 vertex u, and identify a rotation such that the link is y-monotone. We can then collapse u to any valid neighbor, temporarily delete all chords of the hexagon, apply unidirectional convexification to convexify the entire hexagon, and then uncollapse u.

A shape is y-monotone if and only if all local extrema in the y direction are convex. If there are no reflex angles, then there is nothing to do. Otherwise, note that a hexagon cannot have more than three reflex angles, and so we proceed by case analysis.

8.1.2.1 One reflex angle

First suppose that the link of u has one reflex angle. Refer to Figure 8.1. Rotate the hexagon so that ua is almost horizontal and a lies to the right of u. Since b must lie above a and f must lie below a, we conclude that all extrema are convex, and thus the hexagon is y-monotone.

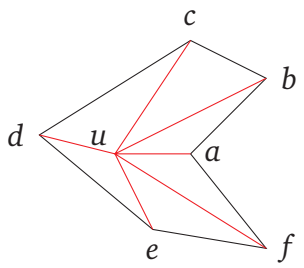

Figure 8.1. A hexagon with one reflex angle

8.1.2.2 Two reflex angles

When the link of u has two reflex angles, there are three possibilities: the reflex angles are adjacent, when they are one apart, and when they are two apart. Refer to Figure 8.2.

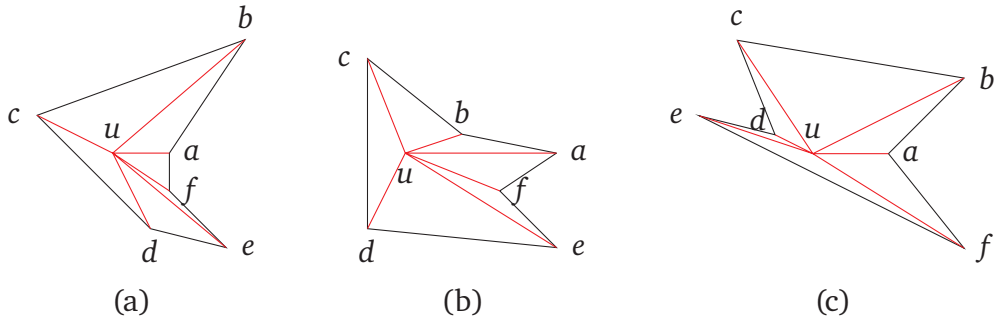

Figure 8.2. The three possibilities for two reflex angles: (a) the reflex angles are adjacent, (b) the reflex angles are one apart, and (c) the reflex angles are two apart.

In cases (a) and (b), we rotate the hexagon so that ua is almost horizontal, and a lies to the right of u. It immediately follows that b lies above a and f lies below a since the edges ub and uf are internal. In case (a), the reflex angles are a and f. Vertex a cannot be an extremum, and since ue is also internal and f lies below u, e must lie below f, so that f is also not an extremum. In case (b), the reflex angles are b and f. Since b lies above u and f lies below u and uc and ue are internal, c lies above b and e lies below f, meaning neither b nor f are internal.

Finally, in case (c), the reflex angles are a and d. We rotate the hexagon so that the edge bc is horizontal, with u below bc. Then a and d lie below b and c. If e and f both lie below a and d then we are done. Since ue and uf are internal, e and f cannot both lie above the line connecting a and d. Without loss of generality assume that e lies above the line connecting a and d. Rotate the hexagon so that ef is horizontal and u lies above it. since e lies above the line connecting a and d, u must lie below that line, meaning that c cannot lie below d and b cannot lie below a, meaning neither reflex angle is an extremum.

8.1.2.3 Three reflex angles

When the link of u has three reflex angles, there are three possibilities: all reflex angles are adjacent, when just two are adjacent, and when none of them are adjacent. Refer to Figure 8.3.

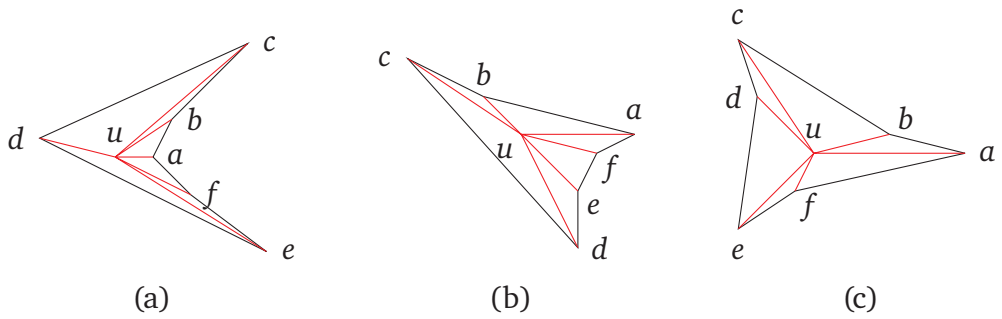

Figure 8.3. The three possibilities for three reflex angles: (a) all reflex angles are adjacent, (b) two reflex angles are adjacent, and (c) no reflex angles are adjacent.

In all three cases, we rotate the hexagon so that ua is almost horizontal, and a lies to the right of u. It immediately follows that b lies above a and f lies below a, as before.

In case (a), the reflex angles are a, b, and f. We know a is not an extremum. Next, c must lie above b and e must lie below f in order for the edges ub and ue to be internal, so b and f are not extrema either.

Next, in case (b) the reflex angles are b, e, and f. If c were the reflex angle instead of b, then we flip the picture so that a lies to the left of u, and then the analysis proceeds without change. Since the edges ue and ud are internal, e lies below f and d lies below e, so neither e nor f are extrema. Next, c must lie above d since uc is also internal, so b is also not an extremum.

Finally, in case (c), the reflex angles are b, d, and f. Since uc and ue are internal, c lies above b and e lies below f. But since d is reflex and ud is internal, the height of d must lie between that of c and that of e, meaning that d cannot be a local extremum. We conclude that no reflex angle is an extremum.

8.1.3 Morphing Toroidal Graphs Edge by Edge

In Section 6.3, we obtained a simple, efficient "best-of-both-worlds" planar morph that simultaneously obtains nice properties enjoyed by both the Floater and Gotsman style [99, 111, 205–207] and the Cairns style [4, 44, 45, 133, 210] of morphs. However, in the toroidal case, while we were able to obtain both a Floater and Gotsman style morph (Chapter 5) and Cairns style morph (Section 6.2), it remains an open question if we can combine the two techniques into another "best-of-both-worlds" result for flat tori.

The problem is that, given a realizable stress vector λ, it is unclear when changing the stresses along the darts of a single edge results in another realizable stress, even if λ were morphable too. We have some partial results. For example, if λ is symmetric, and we change the stresses along the darts of a single edge by the same amount, then the result is also symmetric; in particular, this was the basis for the spring collapses introduced in Section 6.2.2. We can push this idea slightly further:

Lemma 8.1. Let λ be a realizable weight vector, and let α be a positive row vector such that $\alpha L^\lambda = (0,\ldots,0)$ and $\alpha H^\lambda = (0,0)$. Let μ be another positive weight vector such that $\lambda_d \neq \mu_d$ or $\lambda_{\text{rev}(d)} \neq \mu_{\text{rev}(d)}$ for some dart d, and $\lambda_{d'} = \mu_{d'}$ for all darts $d' \notin \{d, \text{rev}(d)\}$. Set $\delta := \mu_d - \lambda_d$ and $\varepsilon := \mu_{\text{rev}(d)} - \lambda_{\text{rev}(d)}$. If $\delta\alpha_{\text{tail}(d)} = \varepsilon\alpha_{\text{head}(d)}$, then μ is realizable.

Proof. Suppose d has tail u and head v. The analysis of Lemma 6.2 gives us

$$L^\mu = L^\lambda + (\delta\mathbf{e}_u - \varepsilon\mathbf{e}_v)(\mathbf{e}_u - \mathbf{e}_v)^T$$
$$H^\mu = H^\lambda + \tau_d(\delta\mathbf{e}_u - \varepsilon\mathbf{e}_v)$$

Because $\alpha L^\lambda = (0,\ldots,0)$ and $\alpha H^\lambda = (0,0)$, we immediately have $\alpha L^\mu = (\delta\alpha_u - \varepsilon\alpha_v)(\mathbf{e}_u - \mathbf{e}_v)^T$ and $\alpha H^\mu = x_d(\delta\alpha_u - \varepsilon\alpha_v)$.

If $\delta\alpha_u = \varepsilon\alpha_v$, then $\alpha L^\mu = (0,\ldots,0)$ and $\alpha X^\mu = (0,0)$. It follows that a suitable scaling of μ is morphable, and therefore realizable, which implies that μ itself is also realizable. □

In particular, when λ is symmetric, then we can choose α to be the all-1s vector, which implies that any weight μ satisfying the conditions of Lemma 8.1 is also symmetric.

8.2 Trivia

Two of the people who worked on planar morphing ended up at UIUC after their initial work on the problem: Stewart S. Cairns published his initial results in 1944 [44,45], and subsequently moved to UIUC in 1948 to serve as the Head of the Mathematics department. Timothy Chan was an author on a preliminary version [5] of Alamdari *et al.* [4] that appeared in 2013, and subsequently moved to UIUC in 2016, joining the Computer Science department.

I was able to obtain organic citations within the main text of this thesis for all but one of the people in my (rather short) academic lineage. In order from most recent to least: Jeff Erickson [48, 49, 51, 86–91], Raimund Seidel [17, 84, 192], John Gilbert [28, 104], and Bob Tarjan [104, 146]. Unfortunately, I did not use any of the potentially related results [32, 100] of my academic lineage's progenitor, Bob Floyd.

Both schools that I attended for my higher education changed their names in my last year of attendance: I enrolled at the Polytechnic Institute of NYU (NYU-Poly) in 2011, but in 2014 the school was renamed the NYU Polytechnic School of Engineering; I was part of the *only* graduating class under that name, for in 2015 the school was once again renamed to the NYU Tandon School of Engineering. I then enrolled at the University of Illinois *at* Urbana-Champaign in 2015; in 2020 the school formally dropped the "at" from the name, becoming the University of Illinois Urbana-Champaign. If one were to extrapolate from a single previous data point, one may reasonably expect UIUC to rename itself again in 2022.

References

[1] Noam Aigerman, Shahar Z. Kovalsky, and Yaron Lipman. Spherical orbifold tutte embeddings. *ACM Trans. Graph.* 36(4):90:1–90:13, 2017.

[2] Noam Aigerman and Yaron Lipman. Orbifold Tutte embeddings. *ACM Trans. Graphics* 34(6):190:1–190:12, 2015.

[3] Noam Aigerman and Yaron Lipman. Hyperbolic orbifold Tutte embeddings. *ACM Trans. Graphics* 35(6):217:1–217:14, 2016.

[4] Soroush Alamdari, Patrizio Angelini, Fidel Barrera-Cruz, Timothy M. Chan, Giordano Da Lozzo, Giuseppe Di Battista, Fabrizio Frati, Penny Haxell, Anna Lubiw, Maurizio Patrignani, Vincenzo Roselli, Sahil Singla, and Bryan T. Wilkinson. How to morph planar graph drawings. *SIAM J. Comput.* 46(2):824–852, 2017.

[5] Soroush Alamdari, Patrizio Angelini, Timothy M. Chan, Giuseppe Di Battista, Fabrizio Frati, Anna Lubiw, Maurizio Patrignani, Vincenzo Roselli, Sahil Singla, and Bryan T. Wilkinson. Morphing planar graph drawings with a polynomial number of steps. *Proc. 24th Ann. ACM-SIAM Symp. Discrete Algorithms*, 1656–1667, 2013.

[6] Lyudmil Aleksandrov and Hristo Djidjev. Linear algorithms for partitioning embedded graphs of bounded genus. *SIAM J. Discrete Math.* 9(1):129–150, 1996.

[7] Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. *Proc. 32nd Ann. ACM-SIAM Symp. Discrete Algorithms*, 2021. arXiv:2010.05846.

[8] Noga Alon and Raphael Yuster. Matrix sparsification and nested dissection over arbitrary fields. *J. ACM* 60(4):25:1–25:8, 2013.

[9] Amos Altshuler. Construction and enumeration of regular maps on the torus. *Discrete Math.* 4:201–217, 1973.

[10] E. M. Andreev. Convex polyhedra in Lobačevskiĭ space. *Mat. Sbornik* 10(3):413–440, 1970.

[11] E. M. Andreev. On convex polyhedra of finite volume in Lobačevskiĭ space. *Mat. Sbornik* 12(2):270–259, 1970.

[12] Patrizio Angelini, Michael A. Bekos, Fabrizio Montecchiani, and Maximilian Pfister. On morphing 1-planar drawings. Preprint, May 2021. arXiv:2105.13040.

[13] Patrizio Angelini, Fabrizio Frati, Maurizio Patrignani, and Vincenzo Roselli. Morphing planar graph drawings efficiently. *Proc. 21st Int. Symp. Graph Drawing*, 49–60, 2013. Lecture Notes Comput. Sci. 8242, Springer.

[14] Patrizio Angelini, Giordano Da Lozzo, Giuseppe Di Battista, Fabrizio Frati, Maurizio Patrignani, and Vincenzo Roselli. Morphing planar graph drawings optimally. *Proc. 41st Int. Colloq. Automata Lang. Prog.*, 126–137, 2014. Lecture Notes Comput. Sci. 8572, Springer.

[15] Patrizio Angelini, Giordano Da Lozzo, Fabrizio Frati, Anna Lubiw, Maurizio Patrignani, and Vincenzo Roselli. Optimal morphs of convex drawings. *Proc. 31st Int. Symp. Comput. Geom.*, 126–140, 2015. Leibniz Int. Proc. Informatics 34. arXiv:1503.09021.

[16] Stuart Appelle. Perception and discrimination as a function of stimulus orientation: the "oblique effect" in man and animals. *Psychological bulletin* 78(4):266. American Psychological Association, 1972.

[17] Boris Aronov, Raimund Seidel, and Diane Souvaine. On compatible triangulations of simple polygons. *Comput. Geom. Theory Appl.* 3(1):27–35, 1993.

[18] Franz Aurenhammer. A criterion for the affine equivalence of cell complexes in R^d and convex polyhedra in R^{d+1}. *Discrete Comput. Geom.* 2(1):49–64, 1987.

[19] Franz Aurenhammer. Power diagrams: Properties, algorithms and applications. *SIAM J. Comput.* 16(1):78–96, 1987.

[20] Franz Aurenhammer. Recognising polytopical cell complexes and constructing projection polyhedra. *J. Symb. Comput.* 3(3):249–255, 1987.

[21] Franz Aurenhammer and Hiroshi Imai. Geometric relations among Voronoi diagrams. *Geom. Dedicata* 27(1):65–75, 1988.

[22] Erwin H. Bareiss. Sylvester's identity and multistep integer-preserving Gaussian elimination. *Math. Comput.* 22(103):565–578, 1968.

[23] Fidel Barrera-Cruz. *Morphing Planar Triangulations*. Ph.D. thesis, Univ. Waterloo, 2014. ⟨http://hdl.handle.net/10012/8518⟩.

[24] Fidel Barrera-Cruz, Penny Haxell, and Anna Lubiw. Morphing planar graph drawings with unidirectional moves. Preprint, November 2014. arXiv:1411.6185.

[25] Fidel Barrera-Cruz, Penny Haxell, and Anna Lubiw. Morphing Schnyder drawings of planar triangulations. *Discrete Comput. Geom.* 61(1):161–184, 2019. arXiv:1411.6186.

[26] Lukas Barth, Benjamin Niedermann, Ignaz Rutter, and Matthias Wolf. Towards a Topology-Shape-Metrics Framework for Ortho-Radial Drawings. *33rd International Symposium on Computational Geometry (SoCG 2017)*, 14:1–14:16, 2017. Leibniz International Proceedings in Informatics (LIPIcs) 77, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ⟨http://drops.dagstuhl.de/opus/volltexte/2017/7223⟩.

[27] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*, 3rd edition. Springer-Verlag, 2008.

[28] M. Bern, D. Eppstein, and J. Gilbert. Provably good mesh generation. *Proc. 31st Annu. IEEE Sympos. Found. Comput. Sci.*, 231–241, 1990.

[29] Therese Biedl, Anna Lubiw, and Michael J. Spriggs. Morphing planar graphs while preserving edge directions. *Proc. 13th Int. Symp. Graph Drawing*, 13–24, 2005. Lecture Notes Comput. Sci. 3843, Springer.

[30] Therese C. Biedl, Anna Lubiw, Mark Petrick, and Michael J. Spriggs. Morphing orthogonal planar graph drawings. *ACM Trans. Algorithms* 9(4):29:1–29:24, 2013.

[31] Ethan D. Bloch, Robert Connelly, and David W. Henderson. The space of simplexwise linear homeomorphisms of a convex 2-disk. *Topology* 23(2):161–175, 1984.

[32] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan. Time bounds for selection. *J. Comput. Syst. Sci.* 7:448–461, 1973.

[33] Alexander I. Bobenko and Boris A. Springborn. A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete Comput. Geom.* 38(4):740–756, 2007. arXiv:math/0503219.

[34] Mikhail Bogdanov, Olivier Devillers, and Monique Teillaud. Hyperbolic delaunay complexes and voronoi diagrams made practical. *Journal of Computational Geometry* 5(1):56–85, 2014.

[35] Mikhail Bogdanov, Monique Teillaud, and Gert Vegter. Delaunay triangulations on orientable surfaces of low genus. *Proc. 32nd Int. Symp. Comput. Geom.*, 20:1–20:17, 2016. Leibniz Int. Proc. Informatics 51.

[36] Ciprian Borcea and Ileana Streinu. Periodic frameworks and flexibility. *Proc. Royal Soc. A* 466(2121):2633–2649, 2010.

[37] Ciprian Borcea and Ileana Streinu. Minimally rigid periodic graphs. *Bull. London Math. Soc.* 43(6):1093–1103, 2011.

[38] Ciprian Borcea and Ileana Streinu. Liftings and stresses for planar periodic frameworks. *Discrete Comput. Geom.* 53(4):747–782, 2015. arXiv:1501.03549.

[39] Carl W. Borchardt. Ueber eine der interpolation entsprechende Darstellung der Eliminations-Resultante. *J. Reine Angew. Math.* 57:111–121, 1860.

[40] Ulrich Brehm and Wolfgang Kühnel. Equivelar maps on the torus. *Europ. J. Combin.* 29(8):1843–1861, 2008.

[41] Graham R. Brightwell and Edward R. Scheinerman. Representations of planar graphs. *SIAM J. Discrete Math.* 6(2):214–229, 1993.

[42] Kevin Q. Brown. Voronoi diagrams from convex hulls. *Inform. Process. Lett.* 9(5):223–228, 1979.

[43] Oleksiy Busaryev, Sergio Cabello, Chao Chen, Tamal K. Dey, and Yusu Wang. Annotating simplices with a homology basis and its applications. *Proc. 13th Scand. Workshop Algorithm Theory*, 189–200, 2012. Lecture Notes Comput. Sci. 7357, Springer. arXiv:1107.3793.

[44] Stewart S. Cairns. Deformations of plane rectilinear complexes. *Amer. Math. Monthly* 51(5):247–252, 1944.

[45] Stewart S. Cairns. Isotopic deformations of geodesic complexes on the 2-sphere and on the plane. *Ann. Math.* 45(2):207–217, 1944.

[46] Manuel Caroli and Monique Teillaud. Delaunay triangulations of closed Euclidean d-orbifolds. *Discrete Comput. Geom.* 55(4):827–853, 2016.

[47] Jean Cerf. About the Bloch–Connelly–Henderson theorem on the simplexwise linear homeomorphisms of a convex 2-disk. Preprint, October 2019. arXiv:1910.00240.

[48] Erin W. Chambers, Éric Colin de Verdière, Jeff Erickson, Francis Lazarus, and Kim Whittlesey. Splitting (complicated) surfaces is hard. *Comput. Geom. Theory Appl.* 41(1–2):94–110, 2008.

[49] Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Minimum cuts and shortest homologous cycles. *Proc. 25th Ann. Symp. Comput. Geom.*, 377–385, 2009.

[50] Erin W. Chambers, Kyle Fox, and Amir Nayyeri. Counting and sampling minimum cuts in genus g graphs. *Discrete Comput. Geom.* 52(3):450–475, 2014.

[51] Erin Wolf Chambers, Jeff Erickson, Patrick Lin, and Salman Parsa. How to morph graphs on the torus. *Proc. 32nd Ann. ACM-SIAM Symp. Discrete Algorithms*, 2759–2778, 2021. arXiv:2007.07927.

[52] Bernard Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.* 6(5):485–524, 1991.

[53] Kun-Ting Chen, Tim Dwyer, Kim Marriott, and Benjamin Bach. Doughnets: Visualising networks using torus wrapping. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, p. 1–11, 2020. CHI '20, Association for Computing Machinery.

[54] Marek Chrobak, Michael T. Goodrich, and Roberto Tamassia. Convex drawings of graphs in two and three dimensions (preliminary version). *Proc. 12th Ann. Symp. Comput. Geom.*, 319–328, 1996.

[55] Edith Cohen and Nimrod Megiddo. Strongly polynomial-time and NC algorithms for detecting cycles in periodic graphs. *J. Assoc. Comput. Mach.* 40(4):791–830, 1993.

[56] Éric Colin de Verdière and Arnaud de Mesmay. Testing graph isotopy on surfaces. *Discrete Comput. Geom.* 51(1):171–206, 2014.

[57] Éric Colin de Verdière, Michel Pocchiola, and Gert Vegter. Tutte's barycenter method applied to isotopies. *Comput. Geom. Theory Appl.* 26(1):81–97, 2003.

[58] Yves Colin de Verdière. Empilements de cercles: Convergence d'une méthode de point fixe. *Forum Math.* 1(1):395–402, 1989.

[59] Yves Colin de Verdière. Sur un nouvel invariant des graphes et un critère de planarité. *J. Comb. Theory Ser. B* 50(1):11–21, 1990. In French, English translation in [62].

[60] Yves Colin de Verdière. Comment rendre géodésique une triangulation d'une surface? *L'Enseignment Mathématique* 37:201–212, 1991.

[61] Yves Colin de Verdière. Un principe variationnel pour les empilements de cercles. *Invent. Math.* 104(1):655–669, 1991.

[62] Yves Colin de Verdière. On a new graph invariant and a criterion for planarity. *Graph Structure Theory*, 137–147, 1993. Contemporary Mathematics 147, Amer. Math. Soc. English translation of [59] by Neil Calkin.

[63] Éric Colin de Verdière. Computational topology of graphs on surfaces. *Handbook of Discrete and Computational Geometry*, 3rd edition, chapter 23, 605–636, 2018. CRC Press LLC.

[64] Lothar Collatz. Spektren periodischer Graphen. *Resultate Math.* 1(1–2):42–53, 1978.

[65] Robert Connelly, Erik D. Demaine, and Günter Rote. Infinitesimally locked self-touching linkages with applications to locked trees. *Physical Knots: Knotting, Linking, and Folding of Geometric Objects in* $\mathbb{R}^3$, 287–311, 2002. Amer. Math. Soc.

[66] Robert Connelly, David W. Henderson, Chung Wu Ho, and Michael Starbird. On the problems related to linear homeomorphisms, embeddings, and isotopies. *Continua, Decompositions, Manifolds: Proc. Texas Topology Symposium, 1980*, 229–239, 1983. Univ. Texas Press.

[67] Keenan Crane. *Discrete Differential Geometry: An Applied Introduction*. 2019. ⟨http://www.cs.cmu.edu/~kmcrane/Projects/DDG/paper.pdf⟩.

[68] Henry Crapo and Walter Whiteley. Plane self stresses and projected polyhedra I: The basic pattern. *Topologie structurale / Structural Topology* 20:55–77, 1993. ⟨http://hdl.handle.net/2099/1091⟩.

[69] Luigi Cremona. *Le figure reciproche nella statica grafica*. Tipografia di Giuseppe Bernardoni, 1872. ⟨http://www.luigi-cremona.it/download/Scritti_matematici/1872_statica_grafica.pdf⟩. English translation in [70].

[70] Luigi Cremona. *Graphical Statics*. Oxford Univ. Press, 1890. ⟨https://archive.org/details/graphicalstatico2cremgoog⟩. English translation of [69] by Thomas Hudson Beare.

[71] Patrick De Leenheer. An elementary proof of a matrix tree theorem for directed graphs. *SIAM Review* 62(3):716–726, 2020. arXiv:1904.12221.

[72] Jesús De Loera, Jörg Rambau, and Francisco Santos. *Triangulations: Structures for Algorithms and Applications*. Algorithms and Computation in Mathematics 25. Springer, 2010.

[73] Olaf Delgado-Friedrichs. Equilibrium placement of periodic graphs and convexity of plane tilings. *Discrete Comput. Geom.* 33(1):67–81, 2004.

[74] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge Univ. Press, 2007.

[75] Erik D. Demaine and André Schulz. Embedding stacked polytopes on a polynomial-size grid. *Discrete Comput. Geom.* 57(4):782–809, 2017. arXiv:1403.7980.

[76] Mathieu Desbrun, Mark Meyer, and Pierre Alliez. Intrinsic parameterizations of surface meshes. *Comput. Graph. Forum* 21(3):209–218, 2002.

[77] Satyan L. Devadoss and Joseph O'Rourke. *Discrete and Computational Geometry*. Princeton University Press, 2011. ⟨http://press.princeton.edu/titles/9489.html⟩.

[78] Sean Dewar. Flexible placements of periodic graphs in the plane. Preprint, November 2019. arXiv:1911.05634.

[79] Tamal K. Dey, Tianqi Li, and Yusu Wang. Efficient algorithms for computing a minimal homology basis. *Proc. 13th Latin American Symp. Theoretical Informatics*, 376–398, 2018. Lecture Notes Comput. Sci. 10807, Springer. arXiv:1801.06759.

[80] Tamal K. Dey, Jian Sun, and Yusu Wang. Approximating loops in a shortest homology basis from point data. *Proc. 26th Ann. Symp. Comput. Geom.*, 166–175, 2010. arXiv:0909.5654.

[81] Peter Eades and Patrick Garvan. Drawing stressed planar graphs in three dimensions. *Proc. 2nd Symp. Graph Drawing*, 212–223, 1995. Lecture Notes Comput. Sci. 1027.

[82] Matthijs Ebbens. *Delaunay triangulations on hyperbolic surfaces*. Ph.D. thesis, Uniiv. Groningen, 2017. ⟨http://fse.studenttheses.ub.rug.nl/id/eprint/15727⟩.

[83] Herbert Edelsbrunner. *Geometry and Topology of Mesh Generation*. Cambridge Univ. Press, 2006.

[84] Herbert Edelsbrunner and Raimund Seidel. Voronoi diagrams and arrangements. *Discrete Comput. Geom.* 1(1):25–44, 1986.

[85] Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency of network flow problems. *J. Assoc. Comput. Mach.* 19(2):248–264, 1972.

[86] Jeff Erickson. Lecture Notes: Computational Topology, 2021. ⟨http://jeffe.cs.illinois.edu/teaching/topology20/schedule.html⟩.

[87] Jeff Erickson, Kyle Fox, and Luvsandondov Lkhamsuren. Holiest minimum-cost paths and flows in surface graphs. *Proc. 50th Ann. ACM Symp. Theory Comput.*, 1319–1332, 2018. arXiv:1804.01045.

[88] Jeff Erickson and Patrick Lin. A toroidal Maxwell-Cremona-Delaunay correspondence. *Proc. 36th Int. Symp. Comput. Geom.*, 40:1–40:17, 2020. Leibniz Int. Proc. Informatics 164, Schloss Dagstuhl–Leibniz-Zentrum für Informatik. arXiv:2003.10057.

[89] Jeff Erickson and Patrick Lin. Planar and toroidal morphs made easier. Preprint, June 2021. arXiv:2106.14086.

[90] Jeff Erickson and Amir Nayyeri. Minimum cuts and shortest non-separating cycles via homology covers. *Proc. 22nd Ann. ACM-SIAM Symp. Discrete Algorithms*, 1166–1176, 2011.

[91] Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. *Proc. 16th Ann. ACM-SIAM Symp. Discrete Algorithms*, 1038–1046, 2005.

[92] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups*. Princeton Math. Series 50. Princeton Univ. Press, 2011.

[93] István Fáry. On straight lines representation of planar graphs. *Acta Sci. Math. Szeged* 11:229–233, 1948.

[94] Stefan Felsner. Lattice structures from planar graphs. *Elec. J. Combin.* 11(1):#R15, 2004.

[95] Stefan Felsner and Günter Rote. On Primal-Dual Circle Representations. *Proc. 2nd Symp. Simplicity in Algorithms*, 8:1–8:18, 2018. OpenAccess Ser. Inform. 69, Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[96] Michael S. Floater. Parametrization and smooth approximation of surface triangulations. *Comp. Aided Geom. Design* 14(3):231–250, 1997.

[97] Michael S. Floater. Parametric tilings and scattered data approximation. *Int. J. Shape Modeling* 4(3–4):165–182, 1998.

[98] Michael S. Floater. Mean value coordinates. *Comput. Aided Geom. Design* 20(1):19–27, 2003.

[99] Michael S. Floater and Craig Gotsman. How to morph tilings injectively. *J. Comput. Appl. Math.* 101(1–2):117–129, 1999.

[100] Robert W. Floyd. Algorithm 97: Shortest path. *Commun. ACM* 5(6):345, 1962. ⟨https://doi.org/10.1145/367766.368168⟩.

[101] Patrick W. Fowler and Simon D. Guest. Symmetry and states of self-stress in triangulated toroidal frames. *International Journal of Solids and Structures* 39:4385–4393, 2002.

[102] Ashim Garg and Roberto Tamassia. A new minimum cost flow algorithm with applications to graph drawing. *Graph Drawing (Proc. GD '96)*, 201–216, 1997. Lecture Notes Comput. Sci. 1190, Springer-Verlag.

[103] Peter Giblin. *Graphs, Surfaces and Homology*, 3rd edition. Cambridge Univ. Press, 2010.

[104] John R. Gilbert, Joan P. Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *J. Algorithms* 5(3):391–407, 1984.

[105] Arthur van Goethem, Bettina Speckmann, and Kevin Verbeek. Optimal morphs of planar orthogonal drawings II. *Proc. 27th Int. Symp. Graph Drawing Network Vis.*, 33–45, 2019. Lecture Notes Comput. Sci. 1904, Springer. arXiv:1908.08365.

[106] Arthur van Goethem and Kevin Verbeek. Optimal morphs of planar orthogonal drawings. *Proc. 34th Int. Symp. Comput. Geom.*, 42:1–42:14, 2018. Leibniz Int. Proc. Informatics 99, Schloss Dagstuhl–Leibniz-Zentrum für Informatik. arXiv:1801.02455.

[107] Jonas Gomes, Lucia Darsa, Bruno Costa, and Luiz Velho. *Warping & morphing of graphical objects*. Morgan Kaufmann, 1999.

[108] Daniel Gonçalves and Benjamin Lévêque. Toroidal maps: Schnyder woods, orthogonal surfaces and straight-line representations. *Discrete Comput. Geom.* 51(1):67–131, 2014. arXiv:1202.0911.

[109] Steven J. Gortler, Craig Gotsman, and Dylan Thurston. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Comput. Aided Geom. Design* 23(2):83–112, 2006.

[110] Craig Gotsman, Xianfeng Gu, and Alla Sheffer. Fundamentals of spherical parameterization for 3d meshes. *ACM Trans. Graph.* 22(3):358–363, 2003.

[111] Craig Gotsman and Vitaly Surazhsky. Guaranteed intersection-free polygon morphing. *Computers & Graphics* 25(1):67–75, 2001.

[112] Clara I. Grima and Alberto Márquez. *Computational Geometry on Surfaces*. Springer, 2001.

[113] Simon D. Guest and John W. Hutchinson. On the determinancy of repetitive structures. *Journal of the mechanics and physics of solids* 51:383–391, 2003.

[114] Xianfeng Gu and Shing-Tung Yau. Global conformal surface parameterization. *Proc. Eurographics/ACM SIGGRAPH Symp. Geom. Process.*, 127–137, 2003.

[115] Joel Hass and Peter Scott. Simplicial energy and simplicial harmonic maps. *Asian J. Math.* 19(4):593–636, 2015. arXiv:1206.2574.

[116] Percy John Heawood. Map colour theorems. *Quart. J. Pure Appl. Math.* 24:332–338, 1890. ⟨https://babel.hathitrust.org/cgi/pt?id=inu.30000050138159&seq=344⟩.

[117] Lothar Heffter. Ueber das Problem der Nachbargebiete. *Math. Ann.* 38(4):477–508, 1891.

[118] Seok-Hee Hong and Hiroshi Nagamochi. Convex drawings of hierarchical planar graphs and clustered planar graphs. *J. Discrete Algorithms* 8(3):282–295, 2010.

[119] John E. Hopcroft and Peter J. Kahn. A paradigm for robust geometric algorithms. *Algorithmica* 7(1–6):339–380, 1992.

[120] Kai Hormann and Michael S. Floater. Mean value coordinates for arbitrary planar polygons. *ACM Trans. Graphics* 25(4):1424–1441, 2006.

[121] David Huffman. A duality concept for the analysis of polyhedral scenes. *Machine Intelligence*, vol. 8, 475–492, 1977. Ellis Horwood Ltd. and John Wiley & Sons.

[122] Alexander Igambardiev and André Schulz. A duality transform for constructing small grid embeddings of 3d polytopes. *Comput. Geom. Theory Appl.* 56:19–36, 2016. arXiv:1402.1660.

[123] Hiroshi Imai, Masao Iri, and Kazuo Murota. Voronoi diagram in the Laguerre geometry and its applications. *SIAM J. Comput.* 14(1):93–105, 1985.

[124] Clause Indermitte, Thomas M. Liebling, Marc Troyanov, and Heinz Clemençon. Voronoi diagrams on piecewise flat surfaces and an application to biological growth. *Theoret. Comput. Sci.* 263(1–2):263–274, 2001.

[125] Kazuo Iwano. *Two-Dimensional Dynamic Graphs and their VLSI Applications*. Ph.D. thesis, Princeton Univ., October 1987. ⟨https://search.proquest.com/docview/303485362⟩.

[126] Kazuo Iwano and Kenneth Steiglitz. Testing for cycles in infintie graphs with periodic structure. *Proc. 19th Ann. ACM Symp. Theory Comput.*, 46–55, 1987.

[127] Kazuo Iwano and Kenneth Steiglitz. A semiring on convex polygons and zero-sum cycle problems. *SIAM J. Comput.* 19(5):883–901, 1990.

[128] Ivan Izmestiev. Statics and kinematics of frameworks in Euclidean and non-Euclidean geometry. *Eighteen Essays in Non-Euclidean Geometry*, 2019. IRMA Lectures in Mathematics and Theoretical Physics 29, Europ. Math. Soc. arXiv:1707.02172.

[129] Donald B. Johnson. Efficient algorithms for shortest paths in sparse networks. *J. Assoc. Comput. Mach.* 24(1):1–13, 1977.

[130] Takashi Kanai, Hiromasa Suzuki, and Fumihiko Kimura. Metamorphosis of arbitrary triangular meshes. *IEEE Computer Graphics and Applications* 20(2):62–75, 2000.

[131] Richard M. Karp, Raymond E. Miller, and Shmuel Winograd. The organization of computations for uniform recurrence equations. *J. Assoc. Comput. Mach.* 14(3):563–590, 1967.

[132] Viktória E. Kaszanitzky, Bernd Schulze, and Shin-ichi Tanigawa. Global rigidity of periodic graphs under fixed-lattice representations. *J. Comb. Theory Ser. B* 146:176–218, 2021.

[133] Linda Kleist, Boris Klemz, Anna Lubiw, Lena Schlipf, Frank Staals, and Darren Strash. Convexity-increasing morphs of planar graphs. *Comput. Geom. Theory Appl.* 84:69–88, 2019.

[134] Boris Klemz. *Facets of Planar Graph Drawing*. Ph.D. thesis, Freie Universität Berlin, 2020. ⟨https://refubium.fu-berlin.de/handle/fub188/28715⟩.

[135] Paul Koebe. Kontaktprobleme der Konformen Abbildung. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl.* 88:141–164, 1936.

[136] S. Rao Kosaraju and Gregory F. Sullivan. Detecting cycles in dynamic graphs in polynomial time (preliminary version). *Proc. 20th ACM Symp. Theory. Comput.*, 398–406, 1988.

[137] Andrew Kotlov, László Lovász, and Santosh Vempala. The Colin de Verdière number and sphere representations of a graph. *Combinatorica* 17(4):483–521, 1997.

[138] Yves Ladegaillerie. Classification topologique des plongements des 1-complexes compacts dans les surfaces. *C. R. Acad. Sci. Paris A* 278:1401–1403, 1974. ⟨https://gallica.bnf.fr/ark:/12148/bpt6k6236784g/f179⟩.

[139] Yves Ladegaillerie. Classification topologique des plongements des 1-complexes compacts dans les surfaces. *C. R. Acad. Sci. Paris A* 279:129–132, 1974. ⟨https://gallica.bnf.fr/ark:/12148/bpt6k6238171d/f143⟩.

[140] Yves Ladegaillerie. Classes d'isotopie de plongements de 1-complexes dans les surfaces. *Topology* 23(3):303–311, 1984.

[141] Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.

[142] Charles L. Lawson. Transforming triangulations. *Discrete Math.* 3(4):365–372, 1972.

[143] François Le Gall. Powers of tensors and fast matrix multiplication. *Proc. 25th int. Symb. Alg. Comput.*, 296–303, 2014.

[144] Simon Antoine Jean l'Huillier. Mémoire sur la polyédrométrie contenant une démonstration directe du théorème d'Euler sur les polyédres, et un examen des diverses exceptions auxquelles ce théorème est assujetti. *Annales de Mathématiques Pures et Appliquées [Annales de Gergonne]* 3:169–189, 1813. Summarized by Joseph Diaz Gergonne.

[145] Patrick Lin. A note on toroidal maxwell-cremona correspondences. *CoRR* abs/2009.12205, 2020. arXiv:2009.12205.

[146] Richard J. Lipton, Donald J. Rose, and Robert Endre Tarjan. Generalized nested dissection. *SIAM J. Numer. Anal.* 16:346–358, 1979.

[147] László Lovász. Representations of polyhedra and the Colin de Verdière number. *J. Comb. Theory Ser. B* 82(2):223–236, 2001.

[148] László Lovász. Discrete analytic functions: An exposition. *Eigenvalues of Laplacians and other geometric operators*, 241–273, 2004. Surveys in Differential Geometry 9, Int. Press.

[149] László Lovász. *Graphs and Geometry*. Colloquium Publications 69. Amer. Math. Soc., 2019.

[150] Anna Lubiw and Mark Petrick. Morphing planar graph drawings with bent edges. *J. Graph Algorithms Appl.* 15(2):205–227, 2011.

[151] Yanwen Luo. Spaces of geodesic triangulations of surfaces. Preprint, August 2020. arXiv:1910.03070v3.

[152] Yanwen Luo, Tianqi Wu, and Xiaoping Zhu. The deformation space of geodesic triangulations and generalized Tutte's embedding theorem. Preprint, May 2021. arXiv:2105.00612.

[153] Yanwen Luo, Tianqi Wu, and Xiaoping Zhu. The deformation space of geodesic triangulations of flat tori. Preprint, July 2021. arXiv:2107.05159.

[154] Justin Malestein and Louis Theran. Generic combinatorial rigidity of periodic frameworks. *Advances in Mathematics* 233(1):291–331, 2013. arXiv:1008.1837.

[155] Kim Marriott, Helen Purchase, Michael Wybrow, and Cagatay Goncu. Memorability of visual features in network diagrams. *IEEE Transactions on Visualization and Computer Graphics* 18(12):2477–2485. IEEE, 2012.

[156] James Clerk Maxwell. On reciprocal figures and diagrams of forces. *Phil. Mag. (Ser. 4)* 27(182):250–261, 1864.

[157] James Clerk Maxwell. On the application of the theory of reciprocal polar figures to the construction of diagrams of forces. *Engineer* 24:402, 1867. Reprinted in [159, pp. 313–316].

[158] James Clerk Maxwell. On reciprocal figures, frames, and diagrams of forces. *Trans. Royal Soc. Edinburgh* 26(1):1–40, 1870.

[159] James Clerk Maxwell. *The Scientific Letters and Papers of James Clerk Maxwell. Volume 2: 1862–1873*. Cambridge Univ. Press, 2009.

[160] Maria Mazón and Tomás Recio. Voronoi diagrams on orbifolds. *Comput. Geom. Theory Appl.* 8(5):219–230, 1997.

[161] August F. Möbius. Zur Theorie der Polyëder und der Elementarverwandtschaft [Nachlass]. *Gesammelte Werke*, vol. 2, 515–559, 1886. Hirzel, Leipzig. ⟨http://gallica.bnf.fr/ark:/12148/bpt6k994243/f524⟩.

[162] Bojan Mohar. A polynomial time circle packing algorithm. *Discrete Math.* 117(1–3):257–263, 1993.

[163] Bojan Mohar. Circle packings of maps—The Euclidean case. *Rend. Sem. Mat. Fis. Milano* 67(1):191–206, 1997.

[164] Bojan Mohar. Circle packings of maps in polynomial time. *Europ. J. Combin.* 18(7):785–805, 1997.

[165] Bojan Mohar and Pierre Rosenstiehl. Tessellation and visibility representations of maps on the torus. *Discrete Comput. Geom.* 19(2):249–263, 1998.

[166] Bojan Mohar and Alexander Schrijver. Blocking nonorientability of a surface. *J. Comb. Theory Ser. B* 87(1):2–16, 2003.

[167] Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. Johns Hopkins series in the mathematical sciences. Johns Hopkins University Press, 2001.

[168] Seiya Negami. Uniqueness and faithfulness of embedding of toroidal graphs. *Discrete Math.* 44(2):161–180, 1983.

[169] Benjamin Niedermann, Ignaz Rutter, and Matthias Wolf. Efficient Algorithms for Ortho-Radial Graph Drawing. *35th International Symposium on Computational Geometry (SoCG 2019)*, 53:1–53:14, 2019. Leibniz International Proceedings in Informatics (LIPIcs) 129, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ⟨http://drops.dagstuhl.de/opus/volltexte/2019/10457⟩.

[170] Anthony Nixon and Elissa Ross. Periodic rigidity on a variable torus using inductive constructions. *Elec. J. Combin.* 22(1):P1.1, 2015. arXiv:1204.1349.

[171] Douglas A. Norris. Isometries homotopic to the identity. *Proc. Amer. Math. Soc.* 105(3):692–696, 1989.

[172] Shmuel Onn and Bernd Sturmfels. A quantitative Steinitz' theorem. *Beitr. Algebra Geom.* 35(1):125–129, 1994. ⟨https://www.emis.de/journals/BAG/vol.35/no.1/⟩.

[173] David Orden, Günter Rote, Fransisco Santos, Brigitte Servatius, Herman Servatius, and Walter Whiteley. Non-crossing frameworks with non-crossing reciprocals. *Discrete Comput. Geom.* 32(4):567–600, 2004. arXiv:math/0309156.

[174] James B. Orlin. Some problems on dynamic/periodic graphs. *Progress in Combinatorial Optimization*, 273–294, 1984. Academic Press. ⟨https://dspace.mit.edu/handle/1721.1/2033⟩.

[175] Linus Pauling. Xxii. the structure of sodalite and helvite. *Zeitschrift für Kristallographie - Crystalline Materials* 74(1):213 – 225. De Gruyter, 01 Dec. 1930. ⟨https://www.degruyter.com/view/journals/zkri/74/1/article-p213.xml⟩.

[176] Helen C Purchase, Christopher Pilcher, and Beryl Plimmer. Graph drawing aesthetics—created by users, not algorithms. *IEEE Transactions on Visualization and Computer Graphics* 18(1):81–92. IEEE, 2010.

[177] W. J. Macquorn Rankine. Principle of the equilibrium of polyhedral frames. *London, Edinburgh, and Dublin Phil. Mag J. Sci.* 27(180):92, 1864.

[178] William John Macquorn Rankine. *A Manual of Applied Mechanics*. Richard Griffin and Co., 1858. ⟨https://archive.org/details/manualappmechaoorankrich⟩.

[179] Sailesh K. Rao and Thomas Kailath. Regular iterative algorithms and their implementation on processor arrays. *Proc. IEEE* 76(3):259–269, 1988.

[180] Bruce L. Reinhart. Line elements on the torus. *Amer. J. Math.* 81(3):617–631, 1959.

[181] Ares Ribó Mor, Günter Rote, and André Schulz. Small grid embeddings of 3-polytopes. *Discrete Comput. Geom.* 45(1):65–87, 2011.

[182] Jürgen Richter-Gebert. *Realization spaces of polytopes*. Lecture Notes Math. 1643. Springer, 1996.

[183] Igor Rivin. Euclidean structures on simplicial surfaces and hyperbolic volume. *Ann. Math.* 139:553–580, 1994.

[184] Vincenzo Roselli. *Morphing and visiting drawings of graphs*. Ph.D. thesis, Univ. Roma Tre, June 2014. ⟨http://hdl.handle.net/2307/4428⟩.

[185] Elissa Ross. *Geometric and combinatorial rigidity of periodic frameworks as graphs on the torus*. Ph.D. thesis, York University, 2011.

[186] Elissa Ross. The rigidity of periodic frameworks as graphs on a fixed torus. *Contrib. Discrete Math.* 9(1):11–45, 2014. ⟨https://cdm.ucalgary.ca/cdm/index.php/cdm/article/view/329/179⟩. arXiv:1202.6652.

[187] Elissa Ross. The rigidity of periodic frameworks as graphs on a fixed torus. *Contrib. Discrete Math.* 9(1):11–45, 2014. arXiv:1202.6652.

[188] Elissa Ross. Inductive constructions for frameworks on a two-dimensional fixed torus. *Discrete Comput. Geom.* 54(1):78–109, 2015. arXiv:1203.6561.

[189] Günter Rote, Fransisco Santos, and Ileana Streinu. Pseudo-triangulations—a survey. *Essays on Discrete and Computational Geometry: Twenty Years Later*, 343–410, 2008. Contemporary Mathematics 453, Amer. Math. Soc.

[190] Alan Saalfeld. Joint triangulations and triangulation maps. *Proc. 3rd Ann. Symp. Comput. Geom.*, 195–204, 1987.

[191] André Schulz. Drawing 3-polytopes with good vertex resolution. *J. Graph Algorithms Appl.* 15(1):33–52, 2011.

[192] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete Comput. Geom.* 6:423–434, 1991.

[193] Hanxiao Shen, Zhongshi Jiang, Denis Zorin, and Daniele Panozzo. Progressive embedding. *ACM Trans. Graphics* 38(4):32:1–32:13, 2019. Proc. SIGGRAPH 2019.

[194] Jonathan Richard Shewchuk. Triangle: Engineering a 2d quality mesh generator and Delaunay triangulator. *First Workshop on Applied Computational Geometry*, 1996.

[195] Jonathan Richard Shewchuk. A condition guaranteeing the existence of higher-dimensional constrained Delaunay triangulations. *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, 76–85, 1998.

[196] Jonathan Richard Shewchuk. Sweep algorithms for constructing higher-dimensional constrained Delaunay triangulations. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 350–359, 2000.

[197] Dvir Steiner and Anath Fischer. Planar parameterization for closed 2-manifold genus-1 meshes. *Proc. 9th ACM Symp. Solid Modeling Appl.*, 83–91, 2004.

[198] Ernst Steinitz. Polyeder und Raumeinteilungen. *Enzyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen* III.AB(12):1–139, 1916.

[199] Ernst Steinitz and Hans Rademacher. *Vorlesungen über die Theorie der Polyeder: unter Einschluß der Elemente der Topologie*. Grundlehren der mathematischen Wissenschaften 41. Springer-Verlag, 1934. Reprinted 1976.

[200] Sherman K. Stein. Convex maps. *Proc. Amer. Math. Soc.* 2(3):464–466, 1951.

[201] Kenneth Stephenson. *Introduction to Circle Packing: The Theory of Discrete Analytic Functions*. Cambridge Univ. Press, 2005.

[202] Ileana Streinu. Erratum to "Pseudo-triangulations, rigidity and motion planning". *Discrete Comput. Geom.* 35(2):358, 2006.

[203] Ileana Streinu. Pseudo-triangulations, rigidity and motion planning. *Discrete Comput. Geom.* 34(4):587–635, 2006. Publisher's erratum in [202].

[204] Kokichi Sugihara. Realizability of polyhedrons from line drawings. *Computational Morphology: A Computational Geometric Approach to the Analysis Of Form*, 177–206, 1988. Machine Intelligence and Pattern Recognition 6.

[205] Vitaly Surazhsky and Craig Gotsman. Controllable morphing of compatible planar triangulations. *ACM Trans. Graphics* 20(4):203–231, 2001.

[206] Vitaly Surazhsky and Craig Gotsman. Morphing stick figures using optimal compatible triangulations. *Proc. 9th Pacific Conf. Comput. Graphics Appl.*, 40–49, 2001.

[207] Vitaly Surazhsky and Craig Gotsman. Intrinsic morphing of compatible triangulations. *Int. J. Shape Modeling* 9(2):191–201, 2003.

[208] Roberto Tamassia. On embedding a graph in the grid with the minimum numberof bends. *SIAM J. Comput.* 16(3):421–444, 1987.

[209] Roberto Tamassia and Ioannis G. Tollis. Planar grid embedding in linear time. *IEEE Trans. Circ. Syst.* CAS-36(9):1230–1234, 1989.

[210] Carsten Thomassen. Deformations of plane graphs. *J. Comb. Theory Ser. B* 34(3):244–257, 1983.

[211] Nobuaki Tomizawa. On some techniques useful for solution of transportation network problems. *Networks* 1:173–194, 1971.

[212] William T. Tutte. The dissection of equilateral triangles into equilateral triangles. *Math. Proc. Cambridge Phil. Soc.* 44(4):463–482, 1948.

[213] William T. Tutte. How to draw a graph. *Proc. London Math. Soc.* 13(3):743–768, 1963.

[214] Pierre Varignon. *Nouvelle mechanique ou statique, dont le projet fut donné en M.DC.LXXVII*. Claude Jombert, Paris, 1725. ⟨https://gallica.bnf.fr/ark:/12148/bpt6k5652714w.texteImage⟩.

[215] Georges F. Voronoï. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire. Recherches sur les parallélloèdres primitifs. *J. Reine Angew. Math.* 134:198–287, 1908. ⟨https://eudml.org/doc/149291⟩.

[216] Klaus Wagner. Bemerkungen zum Vierfarbenproblem. *Jahresbericht Deutschen Math.-Verein.* 46:26–32, 1936.

[217] William M. Waite. Path detection in multidimensional iterative arrays. *J. Assoc. Comput. Mach.* 14(2):300–310, 1967.

[218] Ofir Weber and Denis Zorin. Locally injective parametrization with arbitrary fixed boundaries. *ACM Trans. Graph.* 33(4):75:1–75:12, 2014.

[219] Marvin S. White Jr and Patricia Griffin. Piecewise linear rubber-sheet map transformation. *The American Cartographer* 12(2):123–131. Taylor & Francis, 1985.

[220] Walter Whiteley. Motion and stresses of projected polyhedra. *Topologie structurale / Structural Topology* 7:13–38, 1982. ⟨http://hdl.handle.net/2099/989⟩.

[221] Walter Whiteley, Peter F. Ash, Ethan Poiker, and Henry Crapo. Convex polyhedra, Dirichlet tesselations, and spider webs. *Shaping Space: Exploring Polyhedra in Nature, Art, and the Geometrical Imagination*, chapter 18, 231–251, 2013. Springer.

[222] Günter M. Ziegler. *Lectures on Polytopes*. Graduate Texts in Mathematics 152. Springer, 1995.