How to Exhibit Toroidal Maps in Space

Dan Archdeacon
Dept. of Math. and Stat.
University of Vermont
Burlington, VT 05405 USA
dan.archdeacon@uvm.edu

C. Paul Bonnington
Department of Mathematics
University of Auckland
Auckland, New Zealand
p.bonnington@auckland.ac.nz

Joanna A. Ellis-Monaghan
Department of Mathematics
Saint Michael's College
Colchester, VT 05439 USA
jellis-monaghan@smcvt.edu

Submitted for publication.
July 18, 2004

Abstract

Steinitz's Theorem states that a graph is the 1-skeleton of a convex polyhedron if and only if it is 3-connected and planar. The polyhedron is called a geometric realization of the embedded graph. Its faces are bounded by convex polygons whose points are coplanar.

A map on the torus does not necessarily have such a geometric realization. In this paper, we relax the condition that faces are the convex hull of coplanar points. We require instead that the convex hull of the points on a face can be projected onto a plane so that the boundary of the convex hull of the projected points is the image of the boundary of the face. We also require that the interiors of the convex hulls of different faces do not intersect. Call this an exhibition of the map. A map is polyhedral if the intersection of any two closed faces is simply connected. Our main result is that every polyhedral toroidal map can be exhibited. As a corollary, every toroidal triangulation has a geometric realization.

1 Introduction

Steinitz's Theorem [23] characterizes those graphs that are 1-skeletons of convex polyhedra in 3-space ($\mathbb{R}^3$): they are the 3-connected planar graphs. We call the polyhedron a *geometric realization* of the graph (precise definitions of these and other terms are given in Section 2). Loosely speaking, the vertices of the graph are represented by points in Euclidean 3-space, the edges are straight line segments joining these points, all points in a face are coplanar, the faces are convex, and distinct faces intersect only along their common boundaries. Geometric realizations are used in rendering computer graphics, where 3-dimensional structures are represented by points in space with convex faces.

Suppose that we are given a graph embedded on a surface of higher genus. When does this embedding have a geometric realization? The problem, restricted to triangulations, was first proposed by Grünbaum ([13], Exercise 13.2.3), who conjectured that "*Every closed orientable triangulated 2-manifold without boundary can be embedded geometrically in three-dimensional Euclidean space* $\mathbb{R}^3$" (see also [6]). This conjecture was recently disproven by Bokowski and Guedes de Oliveira [4], who showed that a certain triangulation of the complete graph K_{12} on a surface of genus 6 cannot be realized geometrically. Brehm and Schild [5] showed that every triagulation of the torus does have a realization in $\mathbb{R}^4$.

In this paper we focus on graphs embedded on the torus, also called "maps" on the torus. An interesting example is the geometric realization in $\mathbb{R}^3$ of the complete graph K_7 embedded on the torus. This is commonly attributed to Császár [8], who discovered it independently, and is known as Császár's Polyhedron. Reay notes it was also known to Möbius, and models were given in Reinhardt [19]. For a popular account and details on how to build a model, see [11, 26]. Altshuler and Brehm [2] found all realizations of this polyhedron.

One of the results in this paper, Corollary 1.2, proves Grünbaum's conjecture for triangulations of the torus. This follows up on work of Altshuler [1], who found geometric embeddings of a large class of toroidal triangulations; namely, those that had a special type of Hamiltonian cycle. Similarly, Lawrenchenco (personal communication) also found geometric realizations of those toroidal triangulations where deleting two vertices gives a set of triangles that can be realized in the plane.

Some work has also been done on geometric realizations of toroidal maps that are not triangulations. For example consider the Heawood map, which is the dual of the embedding of K_7 on the torus. Figure 1 shows this map; here and in subsequent figures with dashed borders the top side is to be identified with the bottom side and the left side with the right so as to form the torus. (The role of the dashed and heavy edges in this figure will be explained in Section 7.) It has 14 vertices, each of degree 3, and each face is bounded by a hexagon. Can we realize this map in 3-space such that for each face all of the interior points are coplanar, and the interiors of the faces do not intersect? The answer is yes, as found by Szilassi ([24], as reported in [12, 27]).

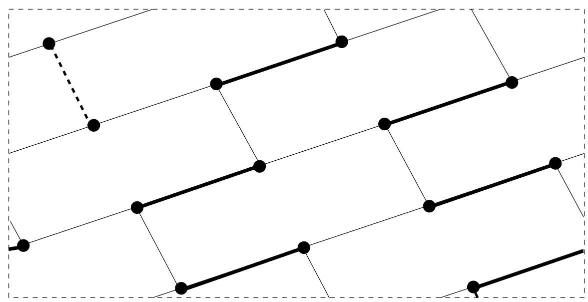

Figure 1: The Heawood Map

The faces of Szilassi's polyhedron are nonconvex. Reay [18] says "It may be shown that no such geometric realization exists with convex hexagons as faces." For completeness, we briefly outline such a proof. Suppose we had such a representation by convex faces. Let f_1 be one of the faces, and consider a representation in $\mathbb{R}^3$ with f_1 in the xy-plane. Any other face f_j, adjacent to f_1 lies entirely in the upper half-space with $z \geq 0$, or entirely in the lower half-space with $z \leq 0$. Since any two of these faces intersect, they lie, without loss of generality, in the upper half-space. Thus the entire surface lies in one half-space of the plane containing f_1. But this holds for all faces, and so the resulting surface cannot be homeomorphic to a torus.

We are interested in representing a toroidal map where we keep a form of convexity of the faces (as defined precisely in Section 2), but relax the condition that the points in a face are all coplanar. The vertices of the map are represented by points in $\mathbb{R}^3$, the edges of the map by straight-line segments, and the faces by the convex hull of their boundary vertices. We require that each face has a "viewing direction" given by a vector $\bar{v}$ such that projecting onto the plane orthogonal to $\bar{v}$ maps the boundary of the face bijectively onto the boundary of the convex hull of the projected points, consistent with their orders. In other words, the face does not necessarily have coplanar boundary vertices on the convex hull, but there is a projection where it is represented with coplanar vertices and is convex. We call this boundary cycle a “profile” of the polytope representing the face. We seek to represent the toroidal map by points in space such that each face can be individually profiled. Moreover, we require that for any two distinct faces, the interiors of their convex hulls are disjoint. The idea is that the open disk that is the face minus its boundary can be anywhere in the interior of the convex hull of its boundary points. Call such a representation an “exhibition” of the map.

We need one more concept before we give our main result. The *face-width* of a map, $fw(G)$, is $\min\{ \mid C \cap G \mid \; : \; C$ is a non-contractible cycle in the surface$\}$. A map is *polyhedral* if it is of face-width at least 3 and the underlying graph G is 3-connected. (Here, the word polyhedral is being used in its topological-graph-theoretic sense, and should not be confused with geometric polyhedra.)

Theorem 1.1 *Every polyhedral map on the torus can be exhibited.*

A triangular face has three boundary vertices, so all points in their convex hull are coplanar. Hence an exhibition of a triangulation is a necessarily a geometric realization (see Proposition 2.2). This answers Grünbaum's Conjecture affirmatively for toroidal triangulations:

Corollary 1.2 *Every toroidal triangulation has a geometric realization.*

This paper is organized as follows. In Section 2 we give our basic definitions and terminology. In Section 3 we give exhibitions of two special maps. These exhibitions not only illustrate our ideas, but are key in later proofs. In Section 4 we extend the exhibitions of our two special maps to exhibitions of vertex splittings of these maps. Section 5 shows how to take a profiled face and extend the exhibition to include a map inside this face. Combining these two sections, Section 6 shows that if a map G contains a surface minor H of maximum degree 4, then an exhibition of H can be modified to an exhibition of G. Section 7 shows that every toroidal map contains a vertex split of one of our two special maps. We use this and our known exhibitions to prove the main result in Section 8. We digress in Section 9 to discuss how we constructed one of our key ingredients: an exhibition of the twisted octahedron. Section 10 gives some conjectures and conclusions.

2 Definitions and terminology

The graphs in this paper are simple and connected. We consider graphs embedded on a surface; as is common, we require our maps to be *cellular*; that is, $S - G$ is a collection of disjoint open disks. We call such an embedding a *map*. A map is *circular* if every face is bounded by a simple cycle. Equivalently, the map is circular if its underlying graph is 2-connected and no non-contractible cycle in the surface (not restricted to cycles in the graph) intersects the graph in a single vertex. Fixing an orientation of the surface, we can describe any face of a circular map by a cyclic order on its incident vertices.

Recall that an embedding is polyhedral if it is 3-connected and of face-width at least 3. Polyhedral maps have the property that the intersection of any two closed faces is simply connected. Every polyhedral map is circular.

Let G be a given map. Form H from G by deleting edges and isolated vertices from the graph G while maintaining the property that H is a map. We call H a *submap* of G. If we in addition allow *smoothing* degree two vertices of G, that is, replacing a path uvw of length 2 in G with a single edge uw in H, then we call H a *topological submap* of G. If we also allow forming a map H by contracting an edge of a map G, we call H a *surface minor* of G.

We want to represent an embedding using vertices in 3-space, with the edges and faces represented by the convex hull of the points representing their incident vertices. We start with some terms about sets of points in $\mathbb{R}^3$.

Let P be a set $\{x_1, \ldots, x_k\}$ of k points in n-space. The *convex hull* of P, denoted $\langle P \rangle$, is the collection of all points of the form $c_1x_1 + c_2x_2 + \ldots + c_kx_k$ such that $c_1 + c_2 + \ldots + c_k = 1$, $0 \leq c_i \leq 1$. The convex hull is also called a (*n-dimensional*) *polytope*. For example, the convex hull of three non-colinear points in $\mathbb{R}^n$ is a triangle. Likewise, the convex hull of four non-coplanar points in $\mathbb{R}^n$ is a tetrahedron.

A point is in the *interior* of the polytope $\langle P \rangle$ if it is of the form $c_1x_1 + c_2x_2 + \ldots + c_kx_k$ such that $c_1 + c_2 + \ldots + c_k = 1$, $0 < c_i \leq 1$. The interior of a set of non-coplanar points agrees with the interior in the usual topology on $\mathbb{R}^3$, and if the points are coplanar but not colinear, it is equivalent to the interior in the usual topology on $\mathbb{R}^2$.

Fix a cyclic ordering $C = (x_1, x_2, \ldots x_k)$ of the points in P. We say that C is a *profile* of $\langle P \rangle$ if there exists a vector $\bar{v}$ such that when we project the points of P onto the plane perpendicular to $\bar{v}$, the convex hull of the projected points contain each of the points on its boundary in the clockwise order given by C. Intuitively, if we look at $\langle P \rangle$ from the direction $\bar{v}$, we see the points occur in the clockwise order given by C. We call $\bar{v}$ the *viewing direction* for the profile. Observe that if C is a profile of $\langle P \rangle$, then every point of P occurs on the boundary of the polytope $\langle P \rangle$.

Next consider a set of n points in 3-space, and a collection $P_1, P_2, \ldots, P_m$ of subsets of these points. We say that the polytopes $\langle P_i \rangle$ are (pairwise) *non-overlapping* if for each i, j, $\langle P_i \cap P_j \rangle = \langle P_i \rangle \cap \langle P_j \rangle$. In other words, the collection is non-overlapping if the interiors of the polytopes are pairwise disjoint. We will also say that a collection of profiles is *non-overlapping* if the convex hulls of their underlying point-sets are non-overlapping.

We discuss how to represent a map on a surface by points in Euclidean 3-space $\mathbb{R}^3$. A *point assignment* of a map G is an injection from $V(G)$ to $\mathbb{R}^3$. When the point assignment is understood, we freely confuse vertices in G with their associated points in $\mathbb{R}^3$. A circular map $G \subset S$ is *exhibited* by a point assignment if each facial boundary cycle C_i gives a cyclic order that is a profile of the points P_i, and if the $\langle P_i \rangle$'s are non-overlapping.

The main question we investigate is:

Problem 2.1 *Which circular embeddings can be exhibited?*

An important subclass of graph embeddings are *triangulations*, where each face is bounded by a 3-cycle. For each face, its three incident vertices lie in a common plane. The convex hull of these three points must lie in the same plane. This gives the following result.

Proposition 2.2 *A triangulation of a surface has an exhibition if and only if it has a geometric realization.*

3 Two special exhibitions

In this section we describe exhibitions of two particular toroidal maps – these exhibitions will turn out to be crucial to our main result. We first define a graph called a (3×3)-*grid*. The vertices are the elements of $\mathbb{Z}_3 \times \mathbb{Z}_3$, where $\mathbb{Z}_3$ is the cyclic group of order three; the edges join (i, j) to $(i+1, j)$ and (i, j) to $(i, j+1)$, for $i, j \in \mathbb{Z}_3$. This graph has a natural quadrilateral embedding in the torus: it is depicted in Figure 2, where a vertex (i, j) is denoted ij.

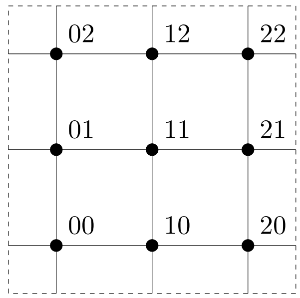

Figure 2: The (3×3)-grid

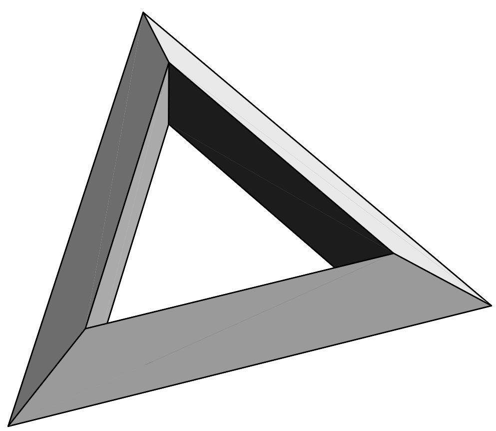

Figure 3: An exhibition of the (3×3)-grid

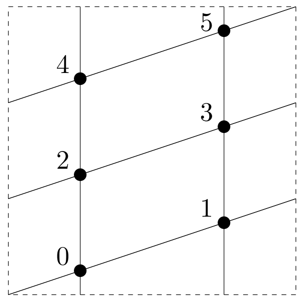

Figure 4: The octahedron embedded on the torus

Proposition 3.1 *The* (3×3)*-grid can be exhibited.*

Proof: Consider the standard model of a torus, where a circle of radius $0 < r < R$ and center $(R, 0)$ in the xy-plane is rotated around the y-axis in 3-space. We place the points equally spaced around this surface; specifically, we use the coordinates

$$([R + r\cos(\frac{2\pi i\theta}{3})]\cos(\frac{2\pi j\phi}{3}), r\sin(\frac{2\pi n\theta}{3}), [R + r\cos(\frac{2\pi n\theta}{3})]\sin(\frac{2\pi n\phi}{3}))$$

for the point (i, j), where $i, j = 0, 1, 2$. The four vertices corresponding to a face of a (3×3)-grid are coplanar. This gives our desired exhibition, as shown in Figure 3. ■

We now describe our second special complex; we start with the graph of the octahedron. The vertex set is $\{0, 1, 2, 3, 4, 5\}$, and the only edges not in the graph are $\{i, i+3\}$, where the vertices are read modulo 6. The octahedron embeds in the torus as shown in Figure 4. We call this embedding the *twisted octahedron*. Observe that the cyclic permutation $(0\,1\,2\,3\,4\,5)$ acts transitively on this map.

Proposition 3.2 *The twisted octahedron can be exhibited.*

Proof: The points 0,1,2,3,4,5 in the graph are mapped respectively to the following points in 3-space:

$$(128, 0, 0), (33, 27, 37), (0, 0, 128), (0, 128, 0), (31, 37, 27), (0, 0, 0).$$

It is tedious but routine to verify that this point assignment is the desired exhibition. ■

The proof given, while correct, is unsatisfying. Where did these coordinates come from? We show their derivation in Section 9 which provides an alternate proof that they give an exhibition of the twisted octahedron.

4 Exhibitions and edge contractions

We now turn our attention to exhibitions of circular maps G and exhibitions of circular maps H formed by contracting a single edge to a point. If G is a graph embedded on a surface S and e is an edge of G that is not a loop, then the edge contraction $H = G/e$ has a natural embedding on S. We call this map a *surface minor* of the original map. We believe the following is true.

Conjecture 4.1 *Let G be a circular map and let H be a circular surface minor formed by contracting a single edge of G. If H can be exhibited, then G can be exhibited.*

The idea is to take an exhibition of H and to modify it slightly at the vertex representing the contracted edge to create an exhibition of G. This seems to be difficult to do in general when the degree of the vertex representing the contracted edge is 6 or more, since modifying a single vertex of H does not necessarily give an exhibition of G. However, we can prove the conjecture in the following special case.

Theorem 4.2 *Let G be a circular map, let H be a circular map formed by contracting a single edge e incident with two degree 3 vertices of G. If H can be exhibited, then G can be exhibited.*

Proof: We label the relevent vertices, edges, and faces as shown in Figure 5. Specifically, let u_1 and u_2 be the vertices incident with the contracted edge e. Label the faces incident with an end of e by F_1, F_2, F_3, F_4 such that F_1 is incident with u_1, F_3 with u_2, and F_2, F_4 are incident with both vertices. Let e_i be the edge incident with F_i and F_{i+1}. The surface minor H is formed by contracting u_1u_2 to a single vertex u, and inherits the labeling of the incident edges and faces, with u denoting the identification of u_1 and u_2.

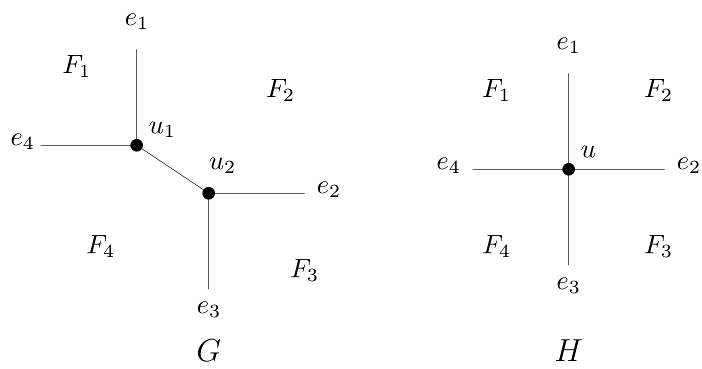

Figure 5: Contracting an edge to form a surface minor

Our goal is to modify an exhibition of H to make an exhibition of G. For convenience we will call the point in 3-space by the same name as the vertex in the graph.

Let P_i be the vertices of F_i, let C_i be the cycle profiling $\langle P_i \rangle$ in the exhibition of H, and let $\bar{v}_i$ be the viewing direction for the profile. We will construct the exhibition of G by replacing the point u by two carefully chosen points u_1 and u_2 at a small distance ϵ from u. All other points remain the same, and the vectors $\bar{v}_i$ remain the same. Replacing u by u_1 and u_2 in this way will ensure that F_1 and F_3 can be profiled with viewing direction $\bar{v}_1$ and $\bar{v}_3$ respectively. The care comes in ensuring that F_2 and F_4 can still be profiled with viewing directions $\bar{v}_2$ and $\bar{v}_4$.

In the exhibition of H, consider the projection of the profile C_2 of $\langle P_2 \rangle$ onto the plane through u perpendicular to $\bar{v}_2$. We focus on the corner of the convex polygon corresponding to u. Let ℓ_2 be the line in this plane perpendicular to the angle bisector of this corner. (See Figure 6.) Observe that if u is replaced by two points u_1, u_2 sufficiently close to u on ℓ_2, then this new set of points is still the vertices of a convex polygon. Let R_2 be the plane containing ℓ_2 and an intersecting line with direction $\bar{v}_2$. We will pick u_1 and u_2 in R_2 so that they project sufficiently close to u so that the projection is still a convex polygon.

Using the face F_4 we proceed as in the preceeding paragraph, defining ℓ_4 and R_4. The two planes R_2 and R_4 intersect in a line ℓ containing u. We will pick u_1 and u_2 on ℓ close to u.

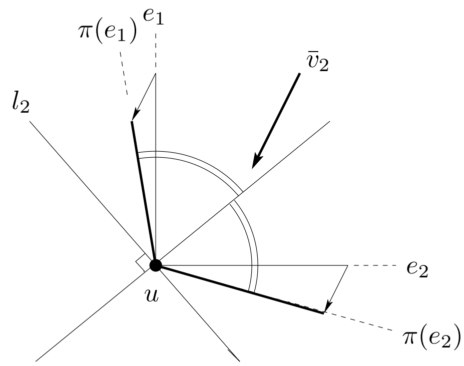

Figure 6: The line l_2 in the plane defined by $\bar{v}_2$ and u

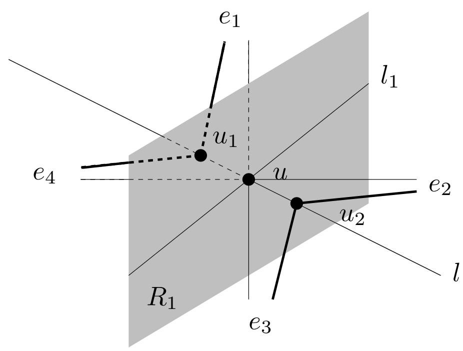

Figure 7: Choosing points to represent a split vertex

We turn our attention to the face F_1. As before, define the plane R_1. The line ℓ passes through R_1 at the point u. By the construction of R_1, both e_1 and e_4 lie on the same side of R_1. Pick the point u_1 on ℓ on this side sufficiently close to u. Pick the point u_2 on ℓ on the other side of R_1 sufficiently close to u (see Figure 7).

We show that these points give a profile of each face. We begin with F_2. Since u_1 and u_2 were chosen on ℓ close to u, they project in the plane perpendicular to $\bar{v}_2$ to points that form the hull of a convex polygon. Since u_1 is on one side of R_3 and u_2 is on the other side, the cyclic order of the points along the hull of the convex polygon is the same as the cyclic order along the boundary cycle of F_2. Similarly, these points give a profile of F_4.

Finally, we note that the interior of distinct $\langle P_i \rangle$ are disjoint, which completes the proof. ■

A *split twisted octahedron* is a map embedded on the torus such that we can contract a set of pairwise non-adjacent edges to obtain the twisted octahedron. Define a *split* (3×3)-*grid* similarly. xsThe following corollary summarizes the results of this section for use in the rest of the paper.

Corollary 4.3 *If G is a split twisted octahedron or a split (3×3)-grid, then G can be exhibited.*

5 Planar patches in a profiled polytope

Consider a topological submap H of a map G. Each face f of H contains a portion G_f (including the boundary of f) of the embedded G. Given an exhibition of H with a particular face f, we want to extend this exhibition to $H \cup G_f$. We do this by taking the vertices of G_f that are not vertices of H and assigning them points in the convex hull of the points representing the boundary vertices of f. In other words, we put G_f in the polytope formed by the vertices on the boundary of f.

We begin with an examination of the structure of G_f. Consider a graph G embedded in the plane with a simple cycle C bounding the outside face. Suppose that the vertices of C are partitioned into *major* and *minor* vertices. We say that the embedded graph is *nearly* 3-*connected* if *a)* it is 2-connected, *b)* the only cut-sets $\{u, v\}$ of size 2 are non-consecutive vertices in C, and *c)* both components of $G - \{u, v\}$ contain a major vertex. Figure 8 gives an example of such a graph, where the major vertices are the corners of the bounding polygon.

Let G be a map and let H be a circular submap. Fix a face f of H, let G_f be the portion of the embedded G in f, and let the major vertices of G_f be the vertices of of degree at least 3 in H.

Lemma 5.1 *Let G be a 3-connected map, and let H be a circular submap. Then G has a circular submap H' homeomorphic to H such that for each face f of H', G_f is nearly 3-connected.*

Proof: Since G is 3-connected and the boundary of f is a cycle, G_f is 2-connected. Moreover the only cut-sets of size two are non-consecutative

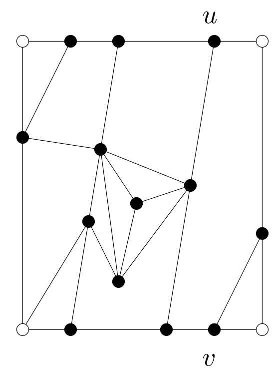

Figure 8: A nearly 3-connected graph (the nonconvex face becomes convex following the application of Proposition 5.2).

vertices $\{u, v\}$ in C. However, it may not be the case that both components of $G_f - \{u, v\}$ contain major vertices. If not, call $\{u, v\}$ a *bad 2-cut*. Pick the submap H' of G homeomorphic to H that has the minimum number of bad 2-cuts. We will show that this number is zero.

By way of contradiction, suppose that H' has a bad 2-cut. Let a be the major vertex on the boundary of G_f closest to u, and let b be the major vertex closest to v, where we allow the possibilities $a = u$ and $b = v$. Let $[u, v]$ and $[a, b]$ be the subpaths of the boundary cycle C of f, from u to v and a to b respectively, not containing an internal major vertex. The component of $G - \{u, v\}$ not containing a major vertex is bounded by a simple cycle C'. Consider the map H'' formed from H' by replacing $[a, b]$ with $([a, b] - [u, v]) \cup (C' - [u, v])$. In other words, we remove the path joining u, v in H' and replace it with the other half of C'. Note that H'' is homeomorphic to H'.

The face in H'' corresponding to f now has one fewer bad 2-cut. Let f' be the other face of H'' incident with $[u, v]$. Since G is 3-connected, there must be a path from $[u, v]$ to the other component of $G - \{a, b\}$. Hence $\{u, v\}$ is not a bad 2-cut in $G_{f'}$. It follows that H'' has fewer bad 2-cuts, a contradiction of the minimality of H'. Hence H' has no bad 2-cuts as claimed.

It follows that for each face f of H', G_f is nearly 3-connected, and the lemma is shown. ■

Knowing the structure of G_f, we examine how to realize this map.

Proposition 5.2 *Let G_f be a nearly-3-connected planar map with $k \geq 3$ major vertices and let C be the simple cycle bounding the outside face. Let D be an arbitrary convex k-gon in the plane. Then G_f can be drawn so that* 1) *C corresponds to D,* 2) *the major vertices in C are the corners of D,* 3) *minor vertices in C are evenly spaced between corners, and* 4) *every face is convex.*

Proof: If G_f is 3-connected, then the embedding exists by Tutte's Theorem [25]. This provides the base case for an induction on the number of cut-sets of order two. By Condition *b)* of nearly 3-connected, any 2-cut $\{u, v\}$ is in C. Let G_1 and G_2 denote the components of $G_f - \{u, v\}$ (there are just two since u and v are in C). We place the vertices of C along D so as to satisfy the first three conditions, and add the edge uv in the interior of D. By Condition *c)* of nearly 3-connected, u, v and the major vertices in G_1 form a convex polygon. By induction, we embed (convexly) the subgraph induced by the vertices of G_1 and $\{u, v\}$ (together with the edge uv) in this polygon. (Here, the vertices u and v play the role of major vertices.) We do the same for G_2. Now, deleting the edge uv (if necessary, when it is not in G) gives the desired embedding. ■

We next use this planar representation of G_f to find an exhibition of G_f in the interior of a polytope that is a profile of its boundary cycle C.

Proposition 5.3 *Let G_f be a nearly-3-connected planar map with $k \geq 3$ major vertices and let C be the simple cycle bounding the outside face. Suppose that there is a profile of the vertices of a polytope (or convex polygon) $\langle P \rangle$ in the plane orthogonal to $\bar{v}$ with C as the projected boundary. Then G_f has an exhibition such that C is the projected boundary of $\langle P \rangle$, and every face of G_f has the same viewing direction $\bar{v}$. Moreover, except for the degenerate case that all of the points of $\langle P \rangle$ are coplanar, the only vertices of G on the boundary of $\langle P \rangle$ are those in C, and the only edges on the boundary of $\langle P \rangle$ are those either in C or edges uv on a 2-cut.*

Proof: Let D denote the k-gon arising from the projection of C in direction $\bar{v}$ onto the plane perpendicular to $\bar{v}$. We begin with the drawing of G_f in the planar k-gon D given by Proposition 5.2. Let v be a vertex of this planar drawing, and let ℓ_v be the line (with direction $\bar{v}$) through v perpendicular to the plane containing D. Now ℓ_v intersects the boundary of the polytope P in exactly one point if v is in C, and in exactly two points otherwise. In the first case, we use that point v' in our exhibition of G_f in P. In the second case, we set v' as the midpoint of the two points of intersection.

If F is the cycle bounding a face in the drawing of G_f in D, then the polytope determined by the corresponding vertices of F obtained above has a profile with viewing direction $\bar{v}$. The polyhedra arising in this way from two different faces do not have a point of their interiors in common; if so, then it would project to a point in the interior of two different faces of G_f in D. Finally, the only vertices on the boundary of $\langle P \rangle$ are those corresponding to vertices of C, and the only edges on the boundary are those with both ends on C. ■

6 Surface minors of toroidal maps

In this section we examine the relation between exhibitions of a surface minor H of a map G and exhibitions of all of G. Basically, we combine the results of Sections 4 and 5 in a general easily applied form. Our goal is the following theorem.

Theorem 6.1 *Let H be a surface minor of G, where H is a circular map and G is a polyhedral map. Suppose that H has maximum degree at most 4. If H has an exhibition, then G has an exhibition.*

Proof: We can assume that G does not contain any vertices of degree 2. Since G contains H as a minor and H is of maximum degree at most 4, then G contains a submap H' together with a collection of pairwise non-adjacent edges $e_1, \ldots, e_k$ of H such that simultaneously contracting each e_i in H' gives a map homeomorphic to H. We are given that H has an exhibition, so by repeated application of Theorem 4.2 H' has an exhibition.

Consider a face f of the embedded H'. Let P_f denote the polytope that is the convex hull of the points representing the boundary of f. Let G_f be that portion of G that is embedded in a face f. By Lemma 5.1, G_f is nearly 3-connected. By Proposition 5.3, G_f can be exhibited in the interior of P_f, with the only portion of G_f on the boundary of P_f being either a portion of the cycle bounding f, or a chord connecting two non-adjacent vertices u, v on the boundary of f. In the latter case $\{u, v\}$ is a cut-set of G_f.

Consider the point representation of the vertices of G created by Proposition 5.3 applied to each of the faces of H'. This point assignment satisfies most, but not all, of the requirements for an exhibition of G. Suppose that we are given two faces f_1, f_2 of G that lie in different faces of H'. Then the interiors of the convex hulls corresponding to these faces are disjoint, since the interiors of the corresponding polyhedra are disjoint in the exhibition of H'. Similarly, if f_1, f_2 lie in the same face of H', the interiors are disjoint by Propostion 5.3.

Let f_1', f_2' be two adjacent faces of H'. They share a common boundary walk W. The vertices of $G_{f_1'}$ and $G_{f_2'}$ are evenly spaced along W, so any two faces of G incident with an edge of W intersect in only that edge.

There cannot be two disjoint faces of H' that both contain a common chord uv, since G does not contain multiple edges.

The only remaining possibility for two polytopes to intersect improperly is if H' has an edge uv, and uv also occurs as a cut-set of size two in a face not incident with that edge. For example, we refer the reader to Figure 4, where the edge 23 appears as an edge in the twisted octahedron H', but also might be a chord in the face 2134. Since G contains no parallel edges, the edge uv of H' corresponds to a path $W = v_1 v_2 \cdots v_k$ in G, where $u = v_1$ and $v = v_k$. An edge $v_i v_{i+1}$ is incident with two faces f_1, f_2 of G, and the edge uv is incident with faces f_3, f_4 of G. The convex hulls of say f_1 and f_2 both contain the edge $v_i v_{i+1}$, which cannot occur in an exhibition.

In this case we will perturb the points $v_2, \ldots, v_{k-1}$ slightly so that there is no point interior to W that also lies on the chord uv, maintaining that the interiors of the convex hulls of different points stay disjoint. We do this by perturbing a point at a time along the path W. We assume that each vertex of W is of degree 3 in G, since we can replace any vertex of degree 2 by an edge joining its two adjacent vertices. We show how to perturb v_2. Once it is perturbed, we can perturb v_3, then v_4, etc.

There are two cases to consider: either v_2 is of degree 3 or more in both f_1 and f_2, or it is of degree 3 or more in (say) f_1 and degree 2 in f_2. The first case is easier. There is an ϵ_1 such that perturbing v_2 by at most ϵ_1 gives a exhibition of G_{f_1}. Similarly define ϵ_2 for G_{f_2}. Now perturb v_2 by some $\epsilon < \min\{\epsilon_1, \epsilon_2\}$.

The second case is slightly more subtle. Since v_2 is of degree 2 in f_2, we have to perturb v_2 in a direction that does not make its incident face in G_{f_2} nonconvex. We perturb as shown in Figure 9 so that the vertex does not move into the interior of the convex hull of the points representing f_1. Such a direction always exists. ■

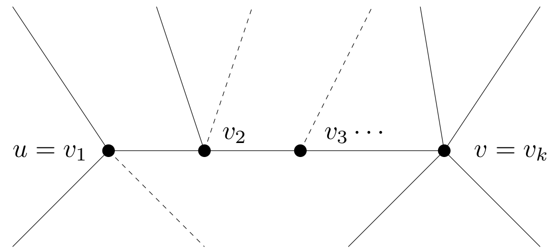

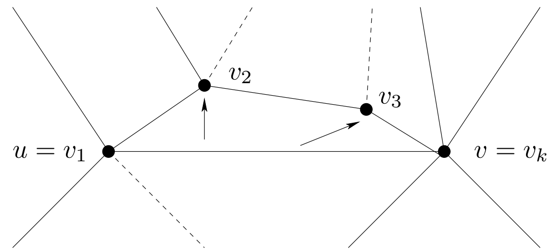

Figure 9: How to perturb intersections

7 Finding submaps of toroidal maps

In earlier sections we have given exhibitions of the split twisted octahedra and the split (3×3)-grids. We showed how such exhibitions can be extended from submaps H to the entire map G. In this section we show that any polyhedral map on the torus contains one of our special submaps (note that there are circular maps that do not contain one of our submaps, for example a subdivided (2×2)-grid). To illustrate the idea, the Heawood map in Figure 1 contains the split twisted octahedron submap. To see this, delete the dashed edge in this figure, suppressing the resulting degree two vertices, and contract the heavy edges.

We first need a preliminary lemma. It says, in essence, that if a graph on the torus contains three disjoint noncontractible cycles in two different directions, then it has a (3×3)-grid minor.

Lemma 7.1 *Let G be a graph embedded on the torus. Suppose that G contains three pairwise disjoint cycles C_1, C_2, C_3 and three pairwise disjoint cycles C_1', C_2', C_3'. Suppose that there is a cycle C homotopic to C_1 and C' homotopic to C_1' with $|C \cap C'| = 1$. Then G has a (3×3)-grid as a topological subgraph.*

Proof: It suffices to show that $C_i \cap C'_j$ is simply connected for each i, j. Pick the submap H of G with the total minimum number of components in $C_i \cap C'_j$. If $C_i \cap C'_j$ is not simply connected for some i, j, then there is a path P_i in C_i and P'_j in C'_j such that $P_i \cup P'_j$ is contractible. Choose P_i and P'_j to enclose the minimum number of faces of the map G. Replacing P'_j with P_i in C'_j gives a submap that contradicts the minimality of H. ■

We can now present the main result of this section.

Theorem 7.2 *Every polyhedral map contains either a split twisted octahedron or a split* (3×3)*-grid.*

Proof: Let G be a polyhedral toroidal map. By Schrijver's Theorem [22] any map on the torus of face-width r contains at least $\lfloor 3r/4 \rfloor$ pairwise disjoint non-contractible cycles; these cycles are necessarily homotopic. Here $r \geq 3$, so G contains two pairwise disjoint non-contractible cycles. Call these cycles C_1 and C_2.

We next describe how to *cut an embedding along a non-contractible cycle* C. Begin with a non-contractible cycle $C = (v_1, \ldots, v_k)$ of minimal length; such a cycle can have no chords. A small neighborhood C is homeomorphic to a cylinder; direct C so it has a *left side* and a *right side*. Replace each v_i by two vertices v'_i, v''_i, replace each edge uv_i on the left side by uv'_i, each uv_i on the right by uv''_i, and add in two cycles $C' = (v'_1, \ldots, v'_k)$ and $C'' = (v''_1, \ldots, v''_k)$. (See Figure 10.)

We cut along the cycle C_1. The result is a planar graph with two distinguished faces, one corresponding to the left side of C_1 and the other to the right side. We picture it so that the first of these faces contains the origin of the plane, and the second is the outside face. Place a new vertex 0 in the interior of the first face and add edges from 0 to all vertice on the face's boundary cycle. Similarly add a new vertex ∞ in the outside face together with edges from ∞ to every vertex on the outside face's boundary. Call the resulting $\bar{G}$.

It is not hard to show that $\bar{G}$ is 3-connected. In particular, there exist three internally-disjoint paths from 0 to ∞ by Menger's Theorem [16]. Let a', b', c' denote the last vertex on these paths on the left side of C_1, and x'', y'', z'' the first vertex (respectively) on these paths on the right side of C_1. Let a, b, c, x, y, z respectively denote the vertices in C_1 corresponding to $a', b', c', x'', y'', z''$. In particular, there are pairwise disjoint ax-, by-, and cz-paths in G.

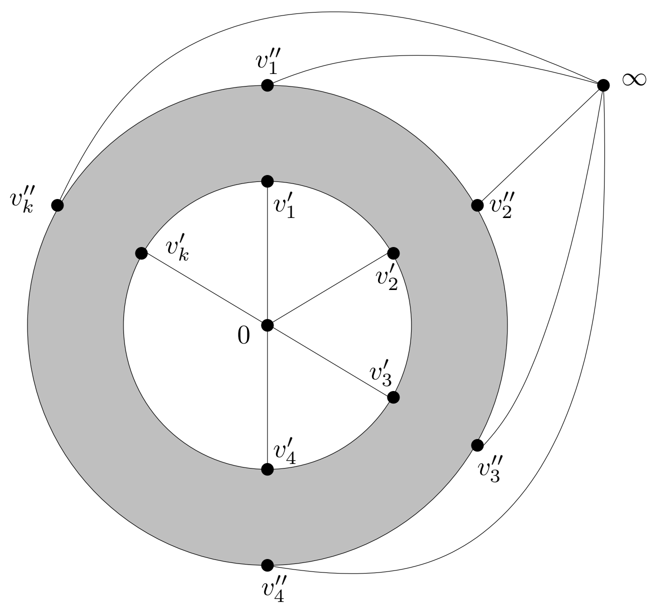

Figure 10: Cutting along a non-contractible cycle

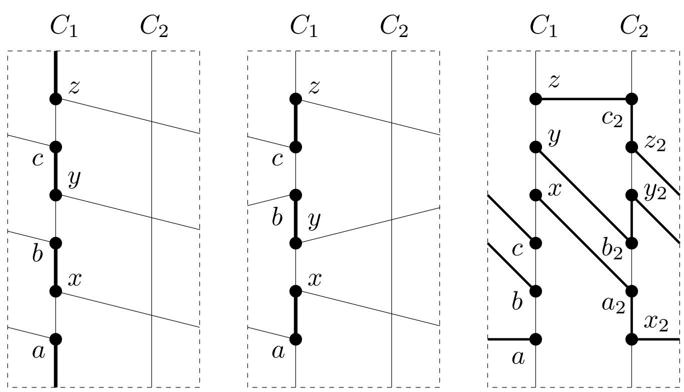

a) Case 1. b) Case 2. c) Case 3.

Figure 11: The three possibilities for the ax-, by-, and cz-paths in G

We examine the cyclic order of a, b, c, x, y, z along C_1. For convenience, consider the directed hexagon formed from C_1 by smoothing all other vertices and directing by the cyclic order induced by (a, b, c). We allow the possibility that some of $\{a, b, c\}$ are equal to some of $\{x, y, z\}$. Without loss of generality, there are three cases to consider (see Figure 11):

Case 1: *The hexagon contains arcs xb, yc, za, including the possibilities $x = b$, $y = c$, and/or $z = a$.* In this case we get a split twisted octahedron and the theorem is satisfied.

Case 2: *The hexagon contains arcs ax, yb, zc and similar subcases.* In this case, G contains three non-contractible pairwise disjoint cycles. We will examine this case later.

Case 3: *The vertices appear in cyclic order (a, b, c, x, y, z), where possibly $c = x$ and $z = a$.*

In this case, we begin again: Let a_2, b_2, c_2 denote the last vertex on the ax-, by-, and cz-paths on C_2 (see Figure 11c)), and x_2, y_2, z_2 the first vertex (respectively) on C_2. This time we cut along cycle C_2 and connect 0 *only* to the vertices a_2', b_2', c_2' corresponding to a_2, b_2, c_2 on the left of C_2 and ∞ *only* to vertices x_2'', y_2'', z_2'' corresponding to x_2, y_2, z_2 on the right of C_2 (producing $\bar{G}$).

We will show that there exist three edge-disjoint paths from 0 to ∞. Suppose instead that there is a cut-set $\{u, v\}$ of size two separating 0 and ∞ in $\bar{G}$, and let G_1 and G_2 be the components of this cut containing 0 and ∞ respectively. Recall that there are 3 pairwise vertex-disjoint paths from 0 to C_1, so C_1 is not contained entirely in G_2. Likewise, there are 3 pairwise vertex-disjoint paths from ∞ to C_1, so C_1 is not contained entirely in G_2. It follows that u, v are on C_1. No face incident with C_1 is also incident with either 0 or ∞. Hence the two faces incident with both u, v in $\bar{G}$ are not incident with 0 or ∞. In particular, they are faces of the original toroidal embedding of G. This contradicts either the fact that G is 3-connected or contradicts that the face-width is at least 3. We conclude that there are three vertex-disjoint paths from 0 to ∞. These three paths intersect C_1 so as to fall in either Case 1 or Case 2.

We now return to Case 2: G contains three non-contractible pairwise disjoint cycles. Call these C_1, C_2, and C_3. Again we cut along C_1. Applying the same logic as before, we either get a split twisted octahedron, or we get three non-contractible cycles in a different homotopy class. In the latter case Lemma 7.1 applies and we get a split (3×3)-grid as desired. ∎

8 Proof of the main result

We combine the results of the previous sections to prove the main result, given in Section 1.

Theorem 1.1: *Every polyhedral map on the torus can be exhibited.*

Proof: Theorem 7.2 shows that any polyhedral G has a submap H that is either a split twisted octahedron or a split (3×3)-grid. Corollary 4.3 shows that each such H has an exhibition. By Theorem 6.1 these exhibitions of H extend to exhibitions of G. Our main result follows. ■

9 Exhibiting a twisted octahedron

We earlier described an exhibition of a twisted octahedron. The description gave the specific points in 3-space, but left lacking the idea behind how we found this exhibition. Discovering this exhibition involved an interesting detour to 4-dimensional space $\mathbb{R}^4$, which we describe here. We are heavily indebted to Duke [9] and Altshuler [1], who pioneered the method combining cyclic polytopes and Schlegel diagrams. A nice exposition on cyclic polytopes is given in Grünbaum ([13], Chapter 4.7).

We begin with the *curve of moments*, which is the locus of points $(t, t^2, \ldots, t^d)$ in $\mathbb{R}^d$. Consider any d distinct nonzero numbers $t_1, \ldots, t_d$. Let M be the $(d \times d)$-matrix whose i^{th} row is given by the point coordinates $(t_i, t_i^2, \ldots, t_i^d)$, $i = 1, \ldots, d$. By Vandemonde's Theorem, the determinate of this matrix is equal to $(t_1 t_2 \ldots t_d)\Pi_{i>j}(t_i - t_j)$, and so is nonzero. This in turn implies that the row vectors are linearly independent, and so the convex hull of the d points $(t_i, t_i^2, \ldots, t_i^d)$ is a $(d-1)$-simplex.

Consider any set of v distinct nonzero points on the curve of moments in $\mathbb{R}^d$. The convex hull of these points is a simplicial d-polytope, called the *cyclic polytope* $C(v, d)$ by Grünbaum. He relates the following result ([10], see also [13] Section 4.8, and [14] Section 13.1.4) which identifies the facets of this cyclic polytope.

Lemma 9.1 Gale's Evenness Condition *A simplex S is a facet of C if and only if for every i, j, the i^{th} and j^{th} coordinates of C are separated by an even number of points t_k.* ■

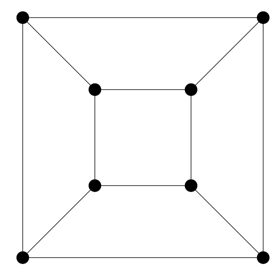

Figure 12: A Schlegel diagram of a cube

We are interested in the case $v = 6$ and $d = 4$. We choose the six points $v_1, \ldots, v_6$ on the curve of moments generated by $t = -3, -2, -1, 1, 2, 3$ respectively. These correspond to the the vertices $1, 2, 3, 4, 5, 6$ of our twisted octahedron as labeled in Figure 4. The points corresponding to the six faces

$$\{(1, 2, 4, 3), (2, 3, 5, 4), (3, 4, 6, 5), (4, 5, 1, 6), (5, 6, 2, 1), (6, 1, 5, 4)\}$$

generate six 4-simplices. Each of these satisify Gale's evenness condition, so they are facets of the cyclic polytope. It follows that they form a geometric realization of the twisted octahedron in 4-space ($\mathbb{R}^4$).

We now project from 4-space back to 3-space. We do this with a Schlegel diagram ([21], see also [13] Sec. 3.3, and [14] Sec. 13.1.6), a type of stereographic projection. Let F_0 be a facet of a d-polytope P in $\mathbb{R}^d$. Pick a point x_0 just slightly outside of P near the facet F. More precisely, chose x_0 such that the only affine hull of a facet of P that separates x_0 and P is that of the facet F_0. Project P onto the affine space generated by F_0 by rays rooted at x_0. The result is called the *Schlegel diagram of P based at F_0*. For example, Figure 12 shows a Schlegel diagram of a standard cube in 3-space from any of its faces.

The authors chose the simplex $\{v_2, v_3, v_5, v_6\}$ for the role of F_0. This satisfies Gale's Evenness Condition, so it is a facet of the cyclic polytope. It is important not to choose a face of the twisted octahedron, because the selected facet corresponds to the unbounded region of the Schlegel diagram. We calculated the coordinates of this projection using *Mathematica*. These coordinates were all rational valued. Scaling by the least common multiple of their denominators gave the coordinates of Theorem 3.2.

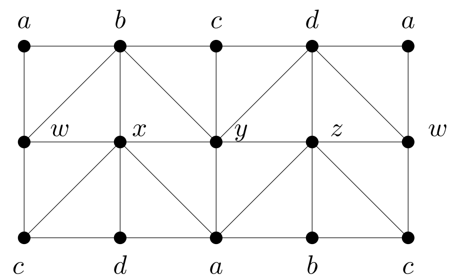

Figure 13: An interesting toroidal triangulation

10 Conjectures and conclusions

We begin with a series of comments about toroidal triangulations. If one is interested only in toroidal triangulations, instead of arbitrary polyhedral maps, the following proposition is helpful. It shows that we do not need to consider embeddings of a split (3×3)-grid.

Proposition 10.1 *Every toroidal triangulation contains a split twisted octahedron.*

Proof: Lavrenchenko [15] found the 21 toroidal triangulations that are minimal under the minor order; this was reportedly done in unpublished work by Grübaum and Duke. It is straightforward to check that each of these 21 maps has a twisted octahedron. If H is a surface minor of G and H has a twisted octahedron, then G has a split twisted octahedron. ■

We offer the following stronger conjecture.

Conjecture 10.2 *Every toroidal triangulation contains a twisted octahedron.*

It is tempting to conjecture an even stronger result: given two homotopic disjoint non-contractible cycles in a toroidal triangulation, there exists a twisted octahedron that contains both cycles. However, this stronger statement is false, as shown in Figure 13. The graph is $K_{2,2,2,2}$. The two disjoint non-contractible cycles are the horizontal 4-cycles.

We note that not every toroidal triangulation contains a split (3×3)-grid. The toroidal embedding of the complete graph K_7 is an example.

Möbius [17] observed that the unique triangulation of the torus by K_7 can be realized by gluing together seven tetrahedron along some of their faces. This motivates the following conjecture.

Conjecture 10.3 *Every triangulation of the torus can be realized geometrically by fitting together tetrahedra such that each vertex of a tetrahedron is a vertex of the triangulation.*

Altshuler and Brehm [3] showed that there is a toroidal map that is not a Schlegel diagram of a 4-polytope. Hence the techniques of this paper may not be enough to prove this conjecture, since these techniques rely on an exhibition of the twisted octahedron as a Schlegel diagram.

What about surfaces other than the torus? Recall that Bokowski and Guedes de Oliveira [4] showed that there is a triangulation of the surface of genus 6 that has no geometric realization, and hence has no exhibition. It is an easy modification of their example to find a trianguation of genus g for each $g \geq 6$ that has no geometric realization. The existence of such examples is open for g between 2 and 5. We ask about the simplest open genus.

Question 10.4 *Does every triangulation of the double torus have a geometric realization?*

We turn our attention to non-triangulations. Note that for these maps on the torus we needed both the embeddings of the split twisted octahedron and of the split (3×3)-grid. In particular, the twisted octahedron does not contain a split (3×3)-grid, and (3×3)-grid does not contain a split twisted octahedron.

The situation is different for maps of large face-width.

Theorem 10.5 *There is a number $f(g)$ such that every 3-connected map on the surface of genus g having face-width at least $f(g)$ has a geometric realization, in particular, it has an exhibition.*

Proof: Nakamoto notes (personal communication) that for each fixed surface, there is a cubic map H that has a geometric realization. By a result of Robertson and Seymour [20] every map G of sufficiently large face-width contains H as a surface minor. Since H is cubic, the map G must also contain H as a topological minor. It follows that, by our methods, G has a geometric realization, which is also an exhibition. ■

The preceeding theorem is especially interesting for triangulatons. Having shown that such a function $f(g)$ exists for general polyhedral maps, it is natural to ask the following.

Question 10.6 *What is its asymptotic behavior of $f(g)$? Could it be constant?*

We believe the answer to the second question is no.

Finally, we consider the nonorientable surfaces. Triangulations of these surfaces do not have geometric realizations, since they do not embed in $\mathbb{R}^3$. However, suppose that we remove the interior of one face from a triangulation of the projective plane. The result is a triangulation of the Möbius band. Not every triangulation of the Möbius band has a geometric realization in $\mathbb{R}^3$, more strongly, they may not have geometric immersions (allowing self-intersections) [7].

Question 10.7 *Which triangulations of the Möbius band have a geometric realization?*

For other unsolved problems along these lines see Reay [18].

ACKNOWLEDGEMENTS: The second author is thankful for support from the Marsden Fund administered by the Royal Society of New Zealand. We also thank Serge Lawrencenko for introducing us to this problem and for helpful discussions.

References

[1] A. Altshuler, Polyhedral realization in $\mathbb{R}^3$ of triangulations of the torus and 2-manifolds in cyclic 4-polytopes, *Discrete Math.* **1** (1971) 211–238.

[2] A. Altshuler and U. Brehm, All realizations of Möbius' torus with seven vertices, *Structural Topology*, **20** (1991) 59–78.

[3] A. Altshuler and U. Brehm, A non-Schlegelian map on the torus, *Mathematika* **31** no. 1 (1984) 83–88.

[4] J. Bokowski and A. Guedes de Oliveira, On the generation of oriented matroids, *Discrete Comput. Geom.* **24** (2000) 197–208.

[5] U. Brehm and G. Schild, Realizability of of the torus and the projective plane in $\mathbb{R}^4$, *Israel J. Math.* **91** (1995) 249–251.

[6] U. Brehm and E. Schulte, Polyhedral maps, in "Handbook of Discrete and Computational Geometry", pp. 345–358, (eds. J.E. Goodman and J. O'Rourke), CRC Press, Boca Raton, FL, 1977.

[7] U. Brehm and J.M. Wills, Polyhedral manifolds, in "Handbook of Convex Geometry", pp. 535-554, (eds. P.M. Gruber and J.M. Wills), North Holland, Amsterdam, 1993.

[8] Á. Császár, A polyhedron without diagonals, *Acta. Sci. Math. (Szeged)*, **13** (1949) 140–142.

[9] R.A. Duke, Geometric embeddings of complexes, *Amer. Math. Monthly*, **77** (1970) 597-603.

[10] D. Gale, Neighborly and cyclic polytopes, *Proc. Symp. Pure Math.*, **7** (Convexity) (1963) 225–232.

[11] M. Gardner, On the remarkable Császár polyhedron and its applications, *Scientific American* (March 1975) 102–107.

[12] M. Gardner, Modern minimal art, *Scientific American* (November 1978) 20–24.

[13] B. Grünbaum, "Convex Polytopes", *Pure and Applied Mathematics* Vol. 16, Interscience-Wiley, New York, 1967.

[14] M. Henk, J. Richter-Gebert, and G. Ziegler, Basic properties of convex polytopes, in "Handbook of Discrete and Computational Geometry", pp. 243–270, (eds. J.E. Goodman and J. O'Rourke), CRC Press, Boca Raton, FL, 1977.

[15] S. Lavrenchenko, Irreducible triangulations of a torus (Russian), Ukrain. Geom. Sb. **30** (1987) 52–62. Translation in *J. Soviet Math.* **51-5** (1990) 2537–2543.

[16] K. Menger, Zur allgemeinen Kurventheorie, *Fund. Math.*, **10**, (1927), 95–115.

[17] A.F. Möbius, Über die bestimmung des inhalts eines polyeders, *Berichte der K. Sachs. Ges. Wiss., Math. Phys. Klass* **17** (1865) 31–68.

[18] J. Reay, Can neighborly polyhedra be realized geometrically?, in "Shaping Space: a Polyhedral Approach" pp. 251–253, (eds. M. Senechal and G. Fleck), Birkhäuser, Boston, 1988.

[19] C. Reinhardt, Zu Möbius polyedertheorie, *Berichte der K. Sachs. Ges. Wiss. Math. Phys. Klasse* (March 1885).

[20] N. Robertson and P. D. Seymour, Graph minors. VII. Disjoint paths on a surface, *J. Combin. Theory Ser. B*, **45**, no. 2, (1988) 212–254.

[21] V. Schlegel, Theorie der homogen zusammengesetzten Raumgebilde, *Nova Acta Leop. Carol.*, **44** (1883) 343–459.

[22] A. Schrijver, Graphs on the torus and the geometry of numbers, *J. Combin. Theory Ser. B* **58**(1993) 147–158

[23] E. Steinitz, Polyeder un Raumeinteilunger, *Enzykl. Math. Wiss.* Vol. 3, Teil 3A612 (1922) 1–139.

[24] L. Szilassi, A polyhedron in which any two faces are contiguous [in Hungarian, with a Russian summary], *A juhasz Gyula Tanarkepzo Foiskola Tudomanyos Kozlemenyei*, Szeged **2** (1977).

[25] W. T. Tutte, How to draw a graph, *Proc. London Math. Soc.* (3), **13** (1963) 743–767.

[26] E. W. Weisstein, Császár Polyhedron, *From MathWorld - A Wolfram Web Resource*,
`http://mathworld.wolfram.com/CsaszarPolyhedron.html`.

[27] E. W. Weisstein, Szilassi Polyhedron, *From MathWorld - A Wolfram Web Resource*,
`http://mathworld.wolfram.com/SzilassiPolyhedron.html`.