Vertex-Unfoldings of Simplicial Manifolds

Erik D. Demaine* David Eppstein† Jeff Erickson‡

George W. Hart§ Joseph O'Rourke¶

Abstract

We present an algorithm to unfold any triangulated 2-manifold (in particular, any simplicial polyhedron) into a non-overlapping, connected planar layout in linear time. The manifold is cut only along its edges. The resulting layout is connected, but it may have a disconnected interior; the triangles are connected at vertices, but not necessarily joined along edges. We extend our algorithm to establish a similar result for simplicial manifolds of arbitrary dimension.

1 Introduction

It is a long-standing open problem to determine whether every convex polyhedron can be cut along its edges and unfolded flat in one piece without overlap, that is, into a simple polygon. This type of unfolding has been termed an *edge-unfolding*; the unfolding consists of the facets of the polyhedron joined along edges. In contrast, unfolding via arbitrary cuts easily leads to nonoverlap. O'Rourke [O'R00] gives a history of the edge-unfolding problem and its applications to manufacturing. Recently it was established that not every nonconvex polyhedron can be edge-unfolded, even if the polyhedron is *simplicial*, that is, all of its faces are triangles [BDEK99, BDE$^+$01].

In this paper we loosen the meaning of "in one piece" to permit a non-overlapping connected region that (in general) does not form a simple polygon, because its interior is disconnected. We call such an unfolding a *vertex-*

*MIT Laboratory for Computer Science, 200 Technology Square, Cambridge, MA 02139, USA. edemaine@mit.edu.

†Department of Information and Computer Science, University of California, Irvine CA 92697-3425, USA. eppstein@ics.uci.edu. Supported by NSF grant CCR-9912338.

‡Department of Computer Science, University of Illinois at Urbana-Champaign; http://www.cs.uiuc.edu/~jeffe; jeffe@cs.uiuc.edu. Partially supported by a Sloan Fellowship and NSF CAREER award CCR-0093348.

§http://www.georgehart.com/; george@georgehart.com.

¶Department of Computer Science, Smith College, Northampton, MA 01063, USA. orourke@cs.smith.edu. Supported by NSF Distinguished Teaching Scholars award DUE-0123154.

unfolding: facets of the polyhedron are connected in the unfolding at common vertices, but not necessarily along edges. With this easier goal we obtain a positive result:

Theorem 1. *Every connected triangulated 2-manifold (possibly with boundary) has a vertex-unfolding, which can be computed in linear time.*

This result includes simplicial polyhedra of any genus, manifolds with any number of boundary components, and even manifolds like the Klein bottle that cannot be topologically embedded in 3-space. Our proof relies crucially on the restriction that every face is a triangle. The problem remains open for nonsimplicial polyhedra with simply connected or even convex faces; see Section 5.

We extend this result in the natural way to higher dimensions in Section 4.

2 Algorithm Overview

Let $\mathcal{M}$ be a triangulated 2-manifold, possibly with boundary. Following polyhedral terminology, we refer to the triangles of $\mathcal{M}$ as *facets*. The *(vertex-facet) incidence graph* of $\mathcal{M}$ is the bipartite graph whose nodes are the facets and vertices of $\mathcal{M}$, with an arc (v, f) whenever v is a vertex of facet f. A *facet path* is a trail $(v_0, f_1, v_1, f_2, v_2, \dots, f_k, v_k)$ in the incidence graph of $\mathcal{M}$ that includes each facet node exactly once, but may repeat vertex nodes. In any facet path, v_{i-1} and v_i are distinct vertices of facet f_i for all i. Because each facet node appears only once, no arc is repeated. A *facet cycle* is a facet path that is also a circuit, that is, where $v_0 = v_k$.

Our algorithm relies on this simple observation:

Lemma 2. *If $\mathcal{M}$ has a facet path, then $\mathcal{M}$ has a vertex-unfolding in which each triangle of the path occupies an otherwise empty vertical strip of the plane.*

Proof: Let p be a facet path of $\mathcal{M}$. Suppose inductively that a facet path p has been laid out in strips up to facet f_{i-1}, with all triangles left of vertex v_i, the rightmost vertex of f_{i-1}. Let (v_i, f_i, v_{i+1}) be the next few nodes in p; recall that $v_i \neq v_{i+1}$. Rotate facet f_i about vertex v_i so that v_i is leftmost and v_{i+1} rightmost, and the third vertex of f_i lies horizontally between. Such rotations exist because f_i is a triangle. Place f_i in a vertical strip with v_i and v_{i+1} on its left and right boundaries. Repeating this process for all facets in p produces a non-overlapping vertex-unfolding. □

Thus to prove Theorem 1 it suffices to prove that every connected triangulated 2-manifold has a facet path. In addition, the existence of facet paths has two applications:

1. The vertex-unfolding resulting from Lemma 2 can be viewed as a *hinged dissection* [Fre97] of the surface. Thus we demonstrate a hinged dissection for any triangulated 2-manifold.

2. A facet path also yields an "ideal rendering" of any triangulated surface on a computer graphics system with a 1-vertex cache: each triangle shares one vertex with the previous triangle in the graphics pipeline. This result is in some sense best possible: an ideal rendering for a 2-vertex cache in which every adjacent pair of triangles shares two vertices is not always achievable, because there are triangulations whose dual graphs have no Hamiltonian path [AHMS96].

It might be more pleasing to obtain a vertex-unfolding based on a *noncrossing facet path*, one that does not include a pattern $(\ldots, A, v, C, \ldots, B, v, D, \ldots)$ with the facets incident to the vertex v appearing in the cyclic order A, B, C, D. Because a facet path has either no or at most two odd nodes (its endpoints), and because any such planar graph has a noncrossing Eulerian trail, we can convert any facet path into a noncrossing facet path. Specifically, we can replace each crossing pair of vertex-to-vertex connections with one of the two alternate pair of connections, whichever alternate pair keeps the vertex-unfolding connected. See Figure 1.

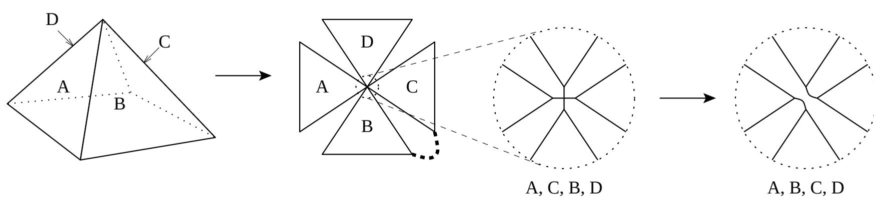

Figure 1. Vertex-unfolding the top four triangles of the regular octahedron and making the connections planar.

Figure 2 shows a vertex-unfolding of the triangulated surface of a cube, obtained from a facet path by our algorithm. Figure 3 shows several more complex examples, generated from random convex polyhedra using an earlier, less general version [DEE^{+}01] of the algorithm we present in the next section. In our examples, we permit the triangles to touch along segments at the strip boundaries, but this could easily be avoided if desired so that each strip boundary contains just the one vertex shared between the adjacent triangles.

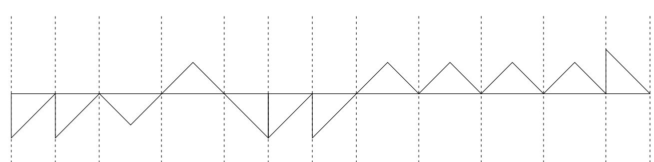

Figure 2. Laying out a facet path in vertical strips.

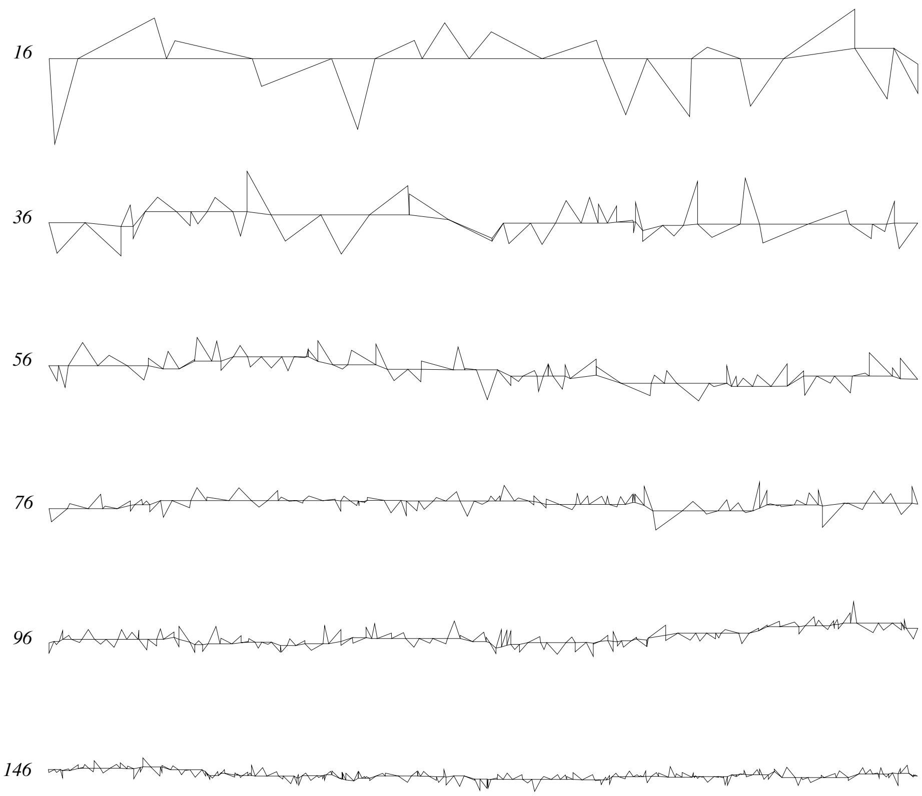

Figure 3. Vertex-unfoldings of random convex polyhedra (generated by code from O'Rourke [O'R98]). The number of triangles is indicated to the left of each unfolding.

3 Facet Paths in 2-Manifolds

In this section, we prove that every triangulated 2-manifold has a facet path, which, by Lemma 2, yields Theorem 1.

A *d-manifold* is a topological space such that every point has a neighborhood homeomorphic to the neighborhood of some point in a closed halfspace in $\mathbb{R}^d$. Interior points of a 2-manifold have neighborhoods homeomorphic to a disk, and boundary points have neighborhoods homeomorphic to a half-disk. (Thus, the vertex-unfoldings in Figures 2 and 3 are not manifolds.) A 2-manifold constructed from (topological) triangles is called a *triangulated 2-manifold*.

The *weak dual* of a triangulated 2-manifold $\mathcal{M}$ is a graph $\mathcal{M}^*$ with a node for every triangle and an arc between any pair of triangles that share a common edge. Let T^* be an arbitrary spanning tree of this dual graph. If we glue the triangles of $\mathcal{M}$ together according to the edges in T^*, we obtain a *topological unfolding* T, a simplicial complex with the topology of a triangulated polygon with no interior vertices. (If we are lucky, we can obtain a *geometric* unfolding of $\mathcal{M}$ by embedding T in the plane, but as mentioned earlier, this is not always possible.) Any facet path or facet cycle of T can be mapped to a facet path or facet cycle of $\mathcal{M}$; recall that a facet path may repeat vertex nodes.

A *scaffold* is a subgraph of the incidence graph in which every facet appears and has degree 2 and at most two vertices have odd degree; if every vertex has even degree, we call it an *even* scaffold. (See Figure 4(d) below for an example of a scaffold.) The edges of any facet path form a scaffold, and any Euler walk through a connected scaffold gives us a facet path. Thus, our goal is to find a connected scaffold for T.

First we establish a slightly weaker result:

Lemma 3. *Every triangulated polygon with no interior vertices has a (possibly disconnected) scaffold.*

Proof: Let T be a triangulated polygon with no interior vertices. We prove the lemma by induction on the number of triangles, with two base cases. If T is empty, we are done. If T is a single triangle, then a path between any two vertices is a scaffold. Henceforth, assume that T has at least two triangles.

An *ear* in T is a triangle that is adjacent to at most one other triangle. We call a triangle in T a *hat* if it is adjacent to at least one ear and at most one non-ear. If we remove all the ears from T, we obtain a new triangulated polygon T'. See Figure 4(a). If T' is nonempty, then it has at least one ear, and every ear in T' is a hat in T. On the other hand, if T' is empty, then T consists of exactly two triangles, which are both ears and hats. In either case, T contains at least one hat.

To perform the induction, we choose a hat in T, find a cycle in the facet-vertex of that hat and (at most two of) its adjacent ears, and recursively construct a scaffold for the remaining triangulation. We have two inductive cases.

Suppose T has a hat $H = qrs$ with at least two ears $E = pqr$ and $F = rst$ (a 'Mickey Mouse hat'). See Figure 4(b). We construct a cycle (r, E, q, H, s, F, r) through the facet-vertex of these three triangles, and recursively construct a scaffold for the smaller triangulation $T \setminus \{H, E, F\}$.

Otherwise, let $H = qrs$ be a hat with just one adjacent ear $E = pqr$ (a 'dunce cap'). See Figure 4(c). We construct a cycle (q, H, r, E, q) through the facet-vertex of those two triangles, and recursively construct a scaffold for the smaller triangulation $T \setminus \{H, E\}$. □

Following the proof gives us an easy linear-time algorithm, consisting of a simple depth-first traversal of the input triangulation's dual tree. Figure 4(d) shows a scaffold computed by our algorithm. The scaffold we construct may be disconnected, but we show how to make it connected using a series of local operations.

Let S be a scaffold with more than one component. If S is not actually an even scaffold, its two odd vertices must be in the same component; every other component has an Euler circuit and so must be 2-connected. Choose a pair of triangles $A = pqr$ and $B = qrs$ that lie in different components of S. Without loss of generality, suppose S contains the edges (p, A), (q, A), (r, B),

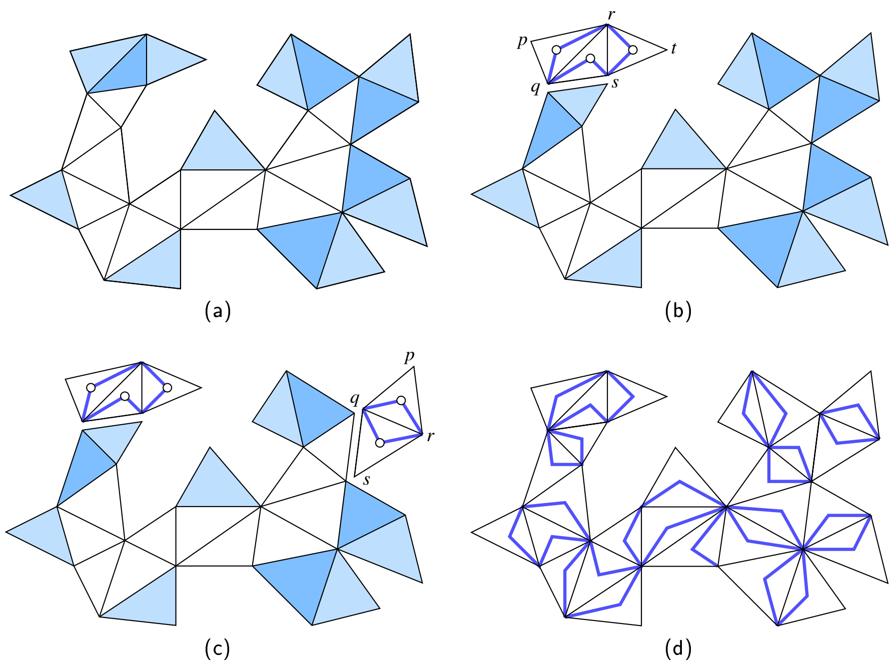

Figure 4. (a) A polygon triangulation. Lightly shaded triangles are ears; darker shaded triangles are hats. (b) Removing a Mickey House hat. (c) Removing a dunce cap. (d) A scaffold produced by our algorithm.

and (s, B). At most one of the edges (q, A) and (r, B) is a bridge (its removal would disconnect the graph). If we remove edges (q, A) and (r, B) and add edges (r, A) and (q, B), we obtain another scaffold S' with one fewer component than S. See Figure 5. Repeating this process for each adjacent pair of components gives us a connected scaffold.

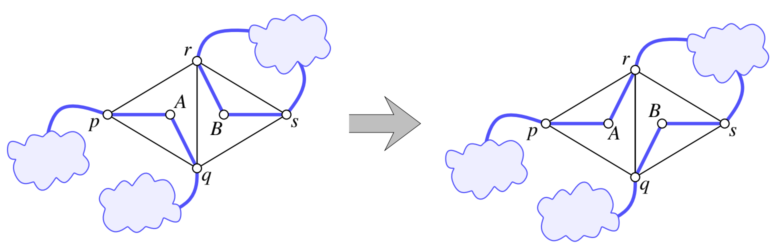

Figure 5. Joining two components of a scaffold with a flip. Clouds hide the rest of the components.

Theorem 4. *Every connected triangulated 2-manifold (possibly with boundary) has a facet path.*

3.1 Checkered Triangulations and Facet Cycles

We can strengthen this result slightly, by showing that most properly triangulated 2-manifolds actually have a facet *cycle*. This permits the strip layout of Lemma 2 to start on the left with any given triangle of the manifold.

We call a polygon triangulation *checkered* if there is a 2-coloring of the triangles so that every white triangle has three (necessarily black) neighbors. See Figure 6(a).

Lemma 5. *A polygon triangulation with no interior vertices has a facet cycle if and only if it is not checkered.*

Proof: Let T be a triangulated polygon with no interior vertices. First we prove by induction that no checkered triangulation has a facet cycle. The base case is a single (black) triangle, which clearly has no facet cycle. In any other checkered triangulation T, we can always find a Mickey Mouse hat: two black ears adjacent to a common (white) triangle. If T has a facet cycle, it must contain a subcycle of six edges inside the Mickey Mouse hat and another subcycle through the rest of the triangulation. But the rest of the triangulation is checkered, so by the induction hypothesis, it has no facet cycle. See Figure 6(b).

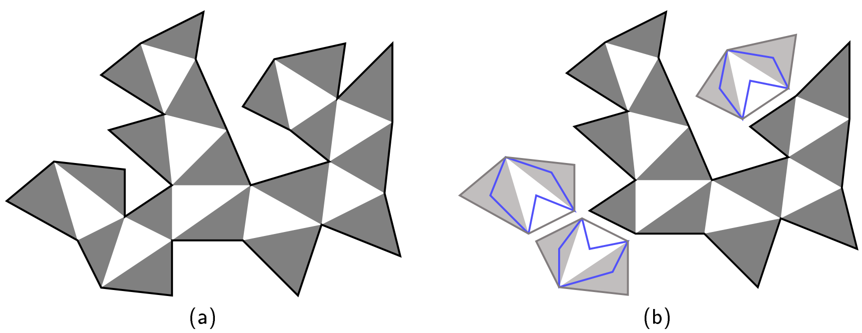

Figure 6. (a) A checkered polygon triangulation. (b) After removing three Mickey Mouse hats.

Now suppose T is a non-checkered triangulation; in particular, T has at least two triangles. To prove the lemma, it suffices to show that T has an even scaffold. Assume without loss of generality that T has *no* Mickey Mouse hats; otherwise, we can remove them as described above. (This might actually eliminate every triangle in T, but then we've computed an even scaffold!) We cannot be left with a single triangle because T is not checkered. Thus, T has at least two 'dunce caps': hats with only one adjacent ear. If we follow the algorithm in Lemma 3, removing Mickey Mouse hats whenever possible, the triangulation always contains at least one dunce cap, until either the algorithm removes every triangle or there are exactly three triangles left. Thus, the triangulation never consists of a single triangle, which implies that the algorithm constructs an even scaffold. $\square$

We say that a triangulated 2-manifold is *simplicial* if it is a simplicial complex, or equivalently, if its dual graph is simple (has no multi-edges or loops).

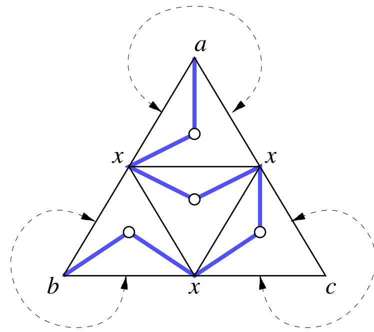

Figure 7. An improper triangulation of the sphere with a facet path but no facet cycle. Edges with identically labeled endpoints are identified.

Every manifold constructed from geometric triangles is simplicial. (A non-simplicial manifold is shown below in Figure 7.)

Lemma 6. *Every connected simplicial* 2*-manifold (possibly with boundary) has a non-checkered topological unfolding, except a checkered polygon triangulation.*

Proof: Let $\mathcal{M}$ be a simplicial 2-manifold that either is multiply connected or has interior vertices. Assume without loss of generality that $\mathcal{M}$ has a checkered topological unfolding T, because otherwise we have nothing to prove. This immediately implies that $\mathcal{M}$ is not 2-colorable.

Color the triangles of T black and white, so that adjacent triangles have opposite colors and every boundary edge of T lies on a black triangle. Because the dual 1-skeleton of $\mathcal{M}$ is a simple graph, we can cut T into two simple polygons along some edge, and then reglue those pieces along some other pair of edges, to obtain another combinatorial unfolding T' of $\mathcal{M}$. Because $\mathcal{M}$ is not 2-colorable, we must reverse the colors of one of those pieces to obtain a proper 2-coloring of T'. Because each piece has at least three edges, each piece has at least one edge that is on the boundary of both T and T'. It follows that T' has boundary edges adjacent to both black and white triangles, so T' is not checkered. $\square$

Combining the previous two lemmas, we conclude the following:

Theorem 7. *Every connected simplicial* 2*-manifold (possibly with boundary) has a facet cycle, except a checkered polygon triangulation.*

This theorem requires that we start with a simplicial complex. There are triangulated but non-simplicial 2-manifolds that have no facet cycle, like the triangulation of the sphere shown in Figure 7. However, even improperly triangulated 2-manifolds have facet paths.

4 Higher Dimensions

In this final section, we generalize our results to higher dimensions. Lemma 2 generalizes in the obvious way, yielding a vertex-unfolding into parallel "slabs" from any facet path. We will show that any simplicial polyhedron has a facet cycle, and thus a vertex-unfolding. Like the results in the previous section, the proofs are almost purely topological, and thus actually apply to arbitrary triangulated manifolds, possibly with boundary, independent of any embedding.

First let us establish some general terminology. A *simplicial d-manifold* is a simplicial complex[1] homeomorphic to a d-dimensional manifold, possibly with boundary. For example, any simplicial convex d-polytope is a simplicial manifold homeomorphic to the sphere $\mathbb{S}^{d-1}$. Following standard polytope nomenclature, we call the full d-dimensional simplices *facets* and the codimension-1 simplices (i.e., the simplices of dimension $d-1$) *ridges*.

The *dual 1-skeleton* $\mathcal{M}^*$ of a simplicial manifold $\mathcal{M}$ is a simple graph, with a node for each facet and an arc between any two facets that share a ridge. We call a simplicial manifold $\mathcal{M}$ *unfolded* if its dual 1-skeleton $\mathcal{M}^*$ is a tree, or equivalently, if it is conencted and its codimension-2 simplices all lie on the boundary. Every unfolded simplicial manifold is homeomorphic to a ball. An unfolded 3-manifold has the topology of a triangulated polyhedron in $\mathbb{R}^3$ with no diagonals or interior points.

An *unfolding* of a connected simplicial manifold $\mathcal{M}$ is a pair (U, f), where U is an unfolded simplicial manifold, and $f : U \to \mathcal{M}$ is a simplicial map that is onto and whose restriction to the interior of U is one-to one. The function f is called the *folding map*, and the set of points in $\mathcal{M}$ with more than one preimage in U is called the *cut set* of the unfolding. We will refer to U as an unfolding of $\mathcal{M}$ when the folding map is clear from context. We easily observe that any spanning tree of $\mathcal{M}^*$ corresponds to a unique unfolding of $\mathcal{M}$. It is an open problem, for every $d \geq 3$, whether every d-polytope has an unfolding that embeds geometrically in $\mathbb{R}^{d-1}$ without overlap.

The *(vertex-facet) incidence graph* of a simplicial manifold has a node for every vertex and every facet, and an arc (v, f) whenever v is a vertex of facet f. A *facet cycle* is a circuit in the incidence graph that passes through every facet exactly once. More generally, an *even scaffold* is a subgraph of the incidence graph in which every facet has degree 2 and every vertex has even degree. Any Euler tour of a connected even scaffold is a facet cycle.

Lemma 8. *For all $d \geq 3$, every unfolded simplicial d-manifold has an even scaffold, except a single d-simplex.*

Proof: Let U be an unfolded simplicial d-manifold with $d \geq 3$ and with more than one facet. An *ear* of U is a facet that is adjacent to only one other facet.

[1] In fact, our results apply to *pseudo-manifold Δ-complexes*—sets of d-simplices whose facets (our ridges) are glued together in pairs [Hat01]. A pseudo-manifold Δ-complex is not necessarily a manifold, even with boundary, but it becomes a manifold if all faces of dimension $d-2$ or less are deleted.

A *hat* is a facet that is adjacent to at least one ear and at most one non-ear. Just as in the two-dimensional case, every unfolded simplicial d-manifold with more than one facet has at least two ears and at least one hat (and with only $d-1$ exceptions, at least two hats).

We prove the lemma by induction. If U is non-empty, we identify a small collection of simplices in U, find a cycle in the incidence graph of those simplices, and recursively construct an even scaffold for the remaining complex, which is still an unfolded manifold. It is fairly easy to construct an even scaffold for any subcomplex consisting of a hat and its ears, and to decompose any unfolded complex into a sequence of such complexes by removing a hat (and its ears) and recursing [Gei58]. To keep things simple, however, we will consider only the four following cases.

1. If U is empty, there is nothing to do.

2. Suppose U consists of exactly three simplices: a hat H and two ears E and F. E and F share at least one vertex p; E and H share at least one vertex $q \neq p$; and H and F share at least one vertex $r \neq p, q$. (In fact, we have $d-2$ choices for each of these three vertices.) Then (p, E, q, H, r, F, p) is a cycle in the incidence graph of U.

3. Suppose some hat H is adjacent to just one ear E. H and E share an edge pq. (In fact, they share an entire ridge.) We recursively construct an even scaffold for the subcomplex $U \setminus \{H, E\}$, and add the cycle (p, H, q, E, p).

4. Finally, suppose some hat H is adjacent to more than one ear and U has more than three facets. Let E and F be any two ears adjacent to H. E and F share an edge pq. (In fact, they share an entire face of dimension $d-2 > 0$; this is the only step of the proof that requires $d \geq 3$.) We recursively construct an even scaffold for the subcomplex $U \setminus \{E, F\}$ and add the cycle (p, E, q, F, p).

The only unfolded complex that does not fall into one of these four cases is a single simplex. □

Once we have an even scaffold, we can make it connected using local flip operations as in the two-dimensional case; in fact the proof is slightly simpler because every component of an even scaffold is 2-connected. Let A and B be adjacent simplices that lie in different components of the even scaffold, and suppose the scaffold contains edges (p, A), (q, A), (r, B), and (s, B). The ridge $A \cap B$ contains all but one vertex of A and all but one vertex of B, so without loss of generality, q and r are both in $A \cap B$. If we replace edges (q, A) and (r, B) with edges (q, B) and (r, A), we obtain a new even scaffold. Any node in the old component of A is still connected to p, then A, then q, and then to any node in the old component of B. Thus, the new even scaffold has one fewer component. Repeating this process for each adjacent pair of components, we obtain a connected even scaffold.

Putting the pieces together yields the following:

Theorem 9. *For any $d \geq 3$, every connected simplicial d-manifold (possibly with boundary) has a facet cycle, except a single d-simplex.*

As we only need a path for the slab construction, we immediately obtain:

Corollary 10. *Every connected simplicial manifold (possibly with boundary) has a vertex-unfolding, which can be computed in linear time.*

5 Discussion

The obvious question left open by our work is whether the restriction to simplicial facets is necessary. Does every three-dimensional polyhedron with simply-connected facets have a non-overlapping vertex-unfolding? What if we require the facets to be convex?

Our strip construction fails for polyhedra with non-triangular convex facets, because such polyhedra may not have facet paths. For example, the truncated cube has no facet path: no pair of its eight triangles can be adjacent in a path, but its six octagons are not enough to separate the triangles.

If facets are permitted to have holes, then there are polyhedra that cannot be vertex-unfolded at all, for example, the box-on-top-of-a-box construction of Biedl et al. (Figure 7 of [BDD$^+$98]).

A related reverse problem is, given a collection of polygons glued together at vertices, can they be glued along their edges to form a polyhedron? In other words, do they form a vertex-unfolding of some polyhedron? What about convex polyhedra? What is the complexity of these decision problems?

Acknowledgments We thank Anna Lubiw for a clarifying discussion, and Allison Baird, Dessislava Michaylova, and Amanda Toop for assisting with the implementation.

References

[AHMS96] E. M. Arkin, M. Held, J. S. B. Mitchell, and S. S. Skiena. Hamiltonian triangulations for fast rendering. *Visual Comput.* 12(6):429–444, 1996.

[BDD$^+$98] T. Biedl, E. Demaine, M. Demaine, A. Lubiw, J. O'Rourke, M. Overmars, S. Robbins, and S. Whitesides. Unfolding some classes of orthogonal polyhedra. In *Proc. 10th Canad. Conf. Comput. Geom.*, pages 70–71, 1998. Fuller version in *Elec. Proc.*: http://cgm.cs.mcgill.ca/cccg98/proceedings/welcome.hml.

[BDE$^+$01] M. Bern, E. D. Demaine, D. Eppstein, E. Kuo, A. Mantler, and J. Snoeyink. Ununfoldable polyhedra with convex faces. *Comput. Geom. Theory Appl.*, to appear, 2001.

[BDEK99] M. Bern, E. D. Demaine, D. Eppstein, and E. Kuo. Ununfoldable polyhedra. In *Proc. 11th Canad. Conf. Comput. Geom.*, pages 13–16, 1999. Full version: http://arXiv.org/abs/cs/9908003/.

[DEE+01] E. D. Demaine, D. Eppstein, J. Erickson, G. W. Hart, and J. O'Rourke. Vertex-unfoldings of simplicial polyhedra. Technical Report 071, Dept. Comput. Sci., Smith College, Northampton, MA, USA, July 2001. http://arXiv.org/abs/cs/0107023/.

[Fre97] G. Frederickson. *Dissections: Plane and Fancy.* Cambridge University Press, 1997.

[Gei58] Dr. Seuss [Theodore Seuss Geisel]. *The Cat in the Hat Comes Back.* Beginner Books. Random House, 1958.

[Hat01] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2001. http://www.math.cornell.edu/~hatcher/.

[O'R98] J. O'Rourke. *Computational Geometry in C.* Cambridge University Press, 2nd edition, 1998. http://cs.smith.edu/~orourke/books/compgeom.html.

[O'R00] J. O'Rourke. Folding and unfolding in computational geometry. In *Discrete and Computational Geometry*, volume 1763 of *Lecture Notes Comput. Sci.*, pages 258–266. Springer-Verlag, 2000. Papers from the Japan Conf. Discrete Comput. Geom., Tokyo, Dec. 1998.