AMPLE TWISTED FACE-PAIRING 3-MANIFOLDS

J. W. CANNON, W. J. FLOYD, AND W. R. PARRY

1. Introduction

In two previous papers [7, 8] we described and investigated a technique, called twisted face-pairing, for constructing 3-manifolds from face-pairings. The starting point for the construction is a faceted 3-ball P together with an orientation-reversing face-pairing ϵ on P. The quotient complex P/ϵ need not be a 3-manifold, and experience shows that it is unlikely that P/ϵ will be a 3-manifold if P and ϵ are chosen "randomly". The face-pairing ϵ determines an equivalence relation on the set of edges of P. The twisted face-pairing construction depends on the faceted 3-ball P, the face-pairing ϵ, and a multiplier function mul, which assigns a positive integer to each equivalence class of edges. For each multiplier function mul, twisted face-pairing produces a pair (Q, δ) consisting of a faceted 3-ball Q which is obtained from P by subdivision and a face-pairing δ on Q which is obtained from ϵ by composition with a twist. The fundamental theorem of twisted face-pairing, which is proved in [7] and in [8], is that the quotient complex $M(\epsilon, \mathrm{mul}) = Q/\delta$ is always a 3-manifold.

In this paper we specialize the construction by restricting the class of faceted 3-balls. A faceted 3-ball P is **ample** if it satisfies the following three conditions.

1. Every two distinct faces of P are either disjoint or meet in a vertex or meet in an edge.
2. Three distinct faces of P which meet each other pairwise have exactly one vertex in common.
3. No face of P is a triangle.

We refer to these conditions as **ampleness conditions** 1, 2 and 3. Ampleness condition 3 is almost superfluous; the only faceted 3-ball which satisfies ampleness conditions 1 and 2 and is not ample is a tetrahedron. We use the term ample because for our purposes ample faceted 3-balls are sufficiently roomy or spacious—ample. For example, Proposition 3.1 shows that an ample faceted 3-ball P is not so small or pinched that it contains a nontrivial simple closed edge path with at most three edges. The other results of Section 3 give more indications of the ampleness of ample faceted 3-balls.

The main theorem of this paper is Theorem 6.1, which states that if P is an ample faceted 3-ball, ϵ is an orientation-reversing face-pairing on P, and mul is a multiplier function for ϵ, then $M(\epsilon, \mathrm{mul})$ has Gromov hyperbolic fundamental group. It then follows from Theorem 5.1 of this paper and a result [2, Theorem 4.1] of Bestvina and Mess that $M(\epsilon, \mathrm{mul})$ has space at infinity a 2-sphere.

Ample faceted 3-balls exist in great profusion. We can construct an ample faceted 3-ball from any faceted 3-ball P as follows. First triangulate ∂P in any way to get a simplicial complex which subdivides ∂P. Then construct the dual cap subdivision of this triangulation of ∂P, which means that every triangle is subdivided as in Figure 1. The result is an ample

Date: March 29, 2000.

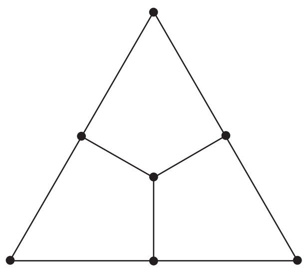

FIGURE 1. The dual cap subdivision of a triangle.

faceted 3-ball P'. A face-pairing ϵ on P naturally determines a face-pairing ϵ' on P' if P' is gotten from P in an ϵ-invariant way.

Since there is an abundance of ample faceted 3-balls, one can use twisted face-pairing to concretely construct many closed 3-manifolds whose fundamental groups are Gromov hyperbolic and have 2-spheres as spaces at infinity. It was the search for such a construction that led to the discovery of twisted face-pairings. It is hoped that these examples will be useful test examples for our attempt to prove Thurston's Hyperbolization Conjecture. For a discussion of the Hyperbolization Conjecture, see [10]. For details of our approach to resolving the conjecture, see [4], [5], [6] and [9].

In Section 2 we give what appears to be a new condition equivalent to Gromov hyperbolicity for a finitely presented group. To describe this condition, let $\mathcal{G}$ be a finitely presented group, and let Γ be a locally finite Cayley graph for $\mathcal{G}$. The condition states that there is a global bound on the length of a geodesic which lies outside a given open metric ball of Γ and whose endpoints lie on the boundary of that open metric ball. It is this condition that we verify to prove the main theorem of this paper. This condition is somewhat surprising in that the analogue for more general spaces is false. In particular, let $X = \mathbf{R}^n$, $n \geq 2$, with the standard metric. Then X is not negatively curved, but it is impossible to find two distinct points on a sphere $S = \partial B$ which are joined by a geodesic that is disjoint from the interior of B.

Let P be an ample faceted 3-ball, let ϵ be a face-pairing on P and let mul be a multiplier function for ϵ. We call $M = M(\epsilon, \mathrm{mul})$ an **ample twisted face-pairing manifold**. We maintain the notation of this paragraph throughout the next four paragraphs.

In Section 3 we develop some properties of ample faceted 3-balls. In Section 4 we set the stage for Section 5 by fixing notation, recalling facts and making definitions. In Section 5 we inductively construct a cell complex and prove that it is cellularly isomorphic to the universal covering complex $\widetilde{M}$ of M. The purpose of the construction is to obtain detailed information about the combinatorial balls of $\widetilde{M}$. Throughout Section 5 we heavily use the duality between M and its dual twisted face-pairing manifold.

The 3-torus provides an interesting example to keep in mind for Sections 3, 4 and 5. We view the 3-torus as being gotten from a cube C in $\mathbf{R}^3$ by identifying opposite faces in the usual way using translations. Then C is an ample faceted 3-ball, so the first two propositions in Section 3 apply to C, although the last two propositions in Section 3 do not apply because we are not twisting. From C we obtain a cell structure on the 3-torus, giving us a cell complex T. Just as for twisted face-pairing 3-manifolds, T has exactly one vertex and the dual of the link of that vertex is cellularly isomorphic to C (in an orientation-reversing way). Let T^* be the cell complex dual to T. The definitions in Section 4 are meaningful for the present example. The structure of the universal covering complex $\widetilde{T}$ is clear: we have the usual tesselation of $\mathbf{R}^3$ by cubes. The same holds for the dual complex $\widetilde{T}^*$, and it is clear how $\widetilde{T}$ and $\widetilde{T}^*$ fit together. Although the hypotheses of Theorem 5.1 are not meaningful, the conclusions of Theorem 5.1 all hold for the present example. So there are strong similarities between the present example and ample twisted face-pairing manifolds. With regard to Theorem 5.1, the main difference between T and M is that the intersection of two distinct 3-cells in $\widetilde{T}$ is either empty or a vertex or an edge or a face while the intersection of two distinct 3-cells in $\widetilde{M}$ is either empty or a vertex or an edge or a face or an elbow. See the middle of Section 3 for the definition of elbow. Consideration of T shows that Sections 3, 4 and 5 are not about negative curvature.

As mentioned earlier, in Section 5 we obtain detailed information about the combinatorial balls in $\widetilde{M}$. We then use this detailed information in Section 6 to prove that $G = \pi_1(M)$ is Gromov hyperbolic. A fact that plays an important role in the argument is that G is almost convex in the sense of [3] in the strongest way possible with respect to its star generating set. (See Section 4 for the definition of star generating set.) In order to state this almost convexity property more carefully, let d be the metric on G with respect to its star generating set and denote the identity element of G by 1. Let x and y be elements of G such that $d(x,1) = d(y,1) \geq 1$. Suppose that there exists an element z of G such that $d(z,x) = d(z,y) = 1$ and $d(z,1) = d(x,1) + 1$. The almost convexity property satisfied by G is that then there exists an element w of G such that $d(w,x) = d(w,y) = 1$ and $d(w,1) = d(x,1) - 1$. This property is at the heart of our argument to verify our criterion for Gromov hyperbolicity. We prove it early in the proof of Theorem 6.1.

Actually, we do not verify our criterion for Gromov hyperbolicity directly for the Cayley graph Γ_s of G with respect to its star generating set, but instead we work with what we call the bipartite graph Γ_b. We next discuss Γ_b. The Cayley graph Γ_s can be identified with the 1-skeleton of the universal covering complex $\widetilde{M}^*$ of the twisted face-pairing manifold M^* dual to M. In the same way the Cayley graph Γ_s^* of G with respect to its dual star generating set can be identified with the 1-skeleton of $\widetilde{M}$. The set of vertices of Γ_b is the union of the set of vertices of $\widetilde{M}$ and the set of vertices of $\widetilde{M}^*$. Edges of Γ_b join vertices of $\widetilde{M}$ with vertices of $\widetilde{M}^*$. A vertex u of $\widetilde{M}$ is joined by an edge of Γ_b with a vertex v of $\widetilde{M}^*$ if and only if the 3-cell of $\widetilde{M}^*$ dual to u contains v, if and only if the 3-cell of $\widetilde{M}$ dual to v contains u. Traversing one edge of Γ_s corresponds to traversing two edges of Γ_b. So combinatorial balls in Γ_b grow half as fast as combinatorial balls in Γ_s. We find that working with the half steps afforded by Γ_b greatly simplifies understanding the growth properties of Γ_s. This is the approach taken in both Section 5 and Section 6.

Section 7 is devoted to examples. We discuss three ample twisted face-pairing manifolds. In particular, we present diagrams of some of the combinatorial balls of their universal covering manifolds. We feel that these diagrams are very illuminating, especially in connection with Section 5. We have written a computer program, pairsnap.c, which provides input to J. Weeks' computer program SnapPea [13] for twisted face-pairings. Our program is freely available from http://www.math.vt.edu/people/floyd. When we apply these programs to the three ample twisted face-pairing manifolds discussed in Section 7, SnapPea says that the manifolds are hyperbolic, in agreement with the Hyperbolization Conjecture.

We maintain our conventions that faces of faceted 3-balls in figures are oriented clockwise and that vertices at corners of faces are original vertices, whereas vertices which are not at corners of faces are not original vertices. By a 3-cell we mean a closed 3-cell unless explicitly stated otherwise. As in Section 4 of [8], X_σ denotes the dual cap subdivision of a given cell complex X.

We thank Richard Canary and Peter Scott for some interesting discussions.

2. A criterion for Gromov hyperbolicity

We begin this paper by establishing a condition equivalent to Gromov hyperbolicity for a finitely presented group. To prepare for this, we first state and prove Lemma 2.1.

Lemma 2.1. *Let J be a positive integer, and let j be an integer with $j \geq J$. Let $a_0, \dots, a_j$ be integers with $a_0 = a_j$, $a_i \geq a_0$ for every $i \in \{0, \dots, j\}$ and $|a_i - a_{i-1}| \leq 1$ for every $i \in \{1, \dots, j\}$. Then there exist integers $p, q \in \{0, \dots, j\}$ such that $J \leq q - p \leq 2J$, $a_p = a_q$ and $a_i \geq a_p$ for every $i \in \{p, \dots, q\}$.*

Proof. We prove Lemma 2.1 by induction on j. If $j \leq 2J$, then we simply take $p = 0$ and $q = j$, so suppose that $j > 2J$.

Suppose that there exists an integer $i \in \{1, \dots, j-1\}$ such that $a_i = a_0$. If $i \geq J$, then $a_0, \dots, a_i$ satisfy the induction hypotheses, and so the conclusion of the lemma is true by induction. If $i < J$, then $a_i, \dots, a_j$ satisfy the induction hypotheses, and again the conclusion of the lemma is true by induction.

Hence we may assume that $a_i > a_0$ for every $i \in \{1, \dots, j-1\}$. But then $a_1, \dots, a_{j-1}$ satisfy the induction hypotheses, and again the conclusion of the lemma is true by induction.

This proves Lemma 2.1. □

The next theorem gives a criterion for proving that a finitely presented group is Gromov hyperbolic.

Theorem 2.2. *Let $\mathcal{G}$ be a finitely presented group. Let Γ be a locally finite Cayley graph for $\mathcal{G}$. Let v be a vertex of Γ, and let d be the edge path metric of Γ. Then $\mathcal{G}$ is Gromov hyperbolic if and only if there exists a positive integer J with the following property. If there exists a geodesic edge path in Γ with vertices $v_0, \dots, v_j$ in order such that $d(v_0, v) = d(v_j, v)$ and $d(v_i, v) \geq d(v_0, v)$ for every $i \in \{0, \dots, j\}$, then $j < J$.*

Proof. First suppose that $\mathcal{G}$ is Gromov hyperbolic. Then geodesic triangles in Γ are K-thin for some nonnegative integer K. We show that the property claimed to be equivalent to Gromov hyperbolicity holds for $J = 4K + 2$. Our proof is by contradiction. Let γ be a geodesic edge path in Γ with vertices $v_0, \dots, v_j$ in order such that $j \geq J$, $d(v_0, v) = d(v_j, v)$ and $d(v_i, v) \geq d(v_0, v)$ for every $i \in \{0, \dots, j\}$. Let γ_1 be a geodesic edge path in Γ joining v to v_0, and let γ_2 be a geodesic edge path in Γ joining v to v_j. Then γ_1, γ and γ_2 are the edges of a geodesic triangle in Γ. Hence the vertex v_{2K+1} lies within distance K from either γ_1 or γ_2. Suppose that v_{2K+1} lies within distance K from a vertex u of γ_1. Then

$$d(v_0, v) \leq d(v_{2K+1}, v) \leq d(v_{2K+1}, u) + d(u, v) \leq K + d(u, v).$$

Because u lies on a geodesic joining v and v_0, it follows that $d(v_0, u) \leq K$. But then

$$d(v_0, v_{2K+1}) \leq d(v_0, u) + d(u, v_{2K+1}) \leq K + K = 2K,$$

which is a contradiction. We likewise obtain a contradiction if u lies on γ_2. This proves the forward implication of Theorem 2.2.

We now prove the backward implication of Theorem 2.2 by showing that Γ satisfies a linear isoperimetric inequality. Suppose that there exists a positive integer J with the following property. There does not exist a geodesic edge path in Γ with vertices $v_0, \dots, v_j$ in order such that $j \geq J$, $d(v_0, v) = d(v_j, v)$ and $d(v_i, v) \geq d(v_0, v)$ for every $i \in \{0, \dots, j\}$. Let γ be a nontrivial closed edge path in Γ based at v. Let $v_0, v_1, v_2, \dots, v_j$ be the vertices of γ in order with $v_0 = v_j = v$. Let $a_i = d(v_i, v)$ for every $i \in \{0, \dots, j\}$. Suppose that $j \geq J$. Then Lemma 2.1 implies that there exist integers $p, q \in \{0, \dots, j\}$ such that $J \leq q - p \leq 2J$, $d(v_p, v) = d(v_q, v)$ and $d(v_i, v) \geq d(v_p, v)$ for every $i \in \{p, \dots, q\}$. By assumption the segment of γ from v_p to v_q is not a geodesic. We decompose γ into two closed edge paths γ' and γ_1 as follows. We choose a geodesic edge path α in Γ from v_p to v_q. We let γ' be the closed edge path obtained from γ by replacing the segment of γ from v_p to v_q by α. We let γ_1 be the closed edge path in Γ obtained by concatenating the inverse of α with the segment of γ from v_p to v_q. The two main features of this decomposition of γ into γ' and γ_1 are that the length of γ' is less than the length of γ and the length of γ_1 is less than $4J$. By iterating this construction with γ' instead of γ, we see that γ can be decomposed into closed edge paths $\gamma_1, \dots, \gamma_i$ such that each of $\gamma_1, \dots, \gamma_i$ has length less than $4J$ and i is at most the length of γ. This easily implies that Γ satisfies a linear isoperimetric inequality. As is well known (see [11, page 104] or [1, Theorem 2.5]), a finitely presented group which satisfies a linear isoperimetric inequality is Gromov hyperbolic. This completes the proof of Theorem 2.2. $\square$

3. Properties of ample faceted 3-balls and their twisted face-pairing subdivisions

The three ampleness conditions have a number of consequences which we explore in this section. Before doing this, we recall the definition of faceted 3-ball. A faceted 3-ball P is an oriented cell complex such that $|P|$ is a closed 3-ball, P has a single 3-cell and for each open cell the prescribed homeomorphism of an open Euclidean ball to that cell extends to a homeomorphism of the closed Euclidean ball to the closed cell.

Proposition 3.1. *Let P be an ample faceted 3-ball. Then every simple closed edge path in P consisting of three or fewer edges is trivial.*

Proof. Since each closed cell of P is homeomorphic to a closed Euclidean ball, there cannot be a closed edge path in P consisting of a single edge. Now suppose that $e_1 e_2$ is a closed edge path in P with exactly two distinct edges. Let f_{1+} and f_{1-} be the two faces that contain e_1, and let f_{2+} and f_{2-} be the two faces that contain e_2, labeled so that f_{1+} and f_{2+} are in the same component of $\partial P \setminus \{e_1 \cup e_2\}$. Since f_{1+}, f_{1-}, f_{2+}, and f_{2-} all contain the two vertices that comprise ∂e_1, by ampleness condition 2 we must have $f_{1+} = f_{2+}$ and $f_{1-} = f_{2-}$. But this is impossible by ampleness condition 1, so there cannot be a closed edge path in P with exactly two distinct edges. Now suppose that $e_1 e_2 e_3$ is a closed edge path in P with exactly three distinct edges. For $i = 1, 2, 3$ let f_{i+} and f_{i-} be the two faces that contain e_i, labeled so that f_{1+}, f_{2+}, and f_{3+} are in the same component of $\partial P \setminus \{e_1 \cup e_2 \cup e_3\}$. Suppose that one of these faces, say f_{1+}, contains all three vertices of $e_1 e_2 e_3$. Then f_{1+} and f_{2-} both contain the vertices of e_2, so ampleness condition 1 implies that $f_{1+} \cap f_{2-}$ is an edge. This edge must be e_2 because there does not exist a closed edge path in P consisting of two distinct edges. So $e_2 \subseteq f_{1+}$, and similarly $e_3 \subseteq f_{1+}$. But then f_{1+} is a triangle, which contradicts ampleness condition 3. So none of the faces which contains an edge of $e_1 e_2 e_3$ contains all three vertices of $e_1 e_2 e_3$. In particular, these faces are distinct. But then ampleness condition 2 applied to f_{1+}, f_{2-} and f_{3-} easily implies that one of these three faces contains every vertex of $e_1 e_2 e_3$, a contradiction. This proves Proposition 3.1. $\square$

Before stating Proposition 3.2, we make it clear that given cell complexes $X \subseteq Y$, we let $\mathrm{star}(X, Y)$ denote the subcomplex of Y which is the union of all the closed cells of Y which meet X.

Proposition 3.2. *Let P be an ample faceted 3-ball. Let X be either a vertex, an edge or a face of P. Then the following statements hold.*

1. *If f and g are distinct faces of $\mathrm{star}(X, \partial P)$ which both contain a vertex $u \notin X$, then $f \cap g$ is the unique edge of P which joins u and X.*
2. *Every edge of P which meets X but is not contained in X contains a vertex not in X.*
3. *The union of X and all the edges of P which meet X is contractible.*
4. *The faces of $\mathrm{star}(X, \partial P)$, other than X itself if X is a face, can be enumerated as $f_1, \dots, f_k$ such that $f_i \cap f_j$ is an edge with exactly one vertex in X if $|i - j| \in \{1, k-1\}$ and $f_i \cap f_j \subseteq X$ otherwise.*
5. *The complex $\mathrm{star}(X, \partial P)$ is a topological disk.*

Proof. We first prove statement 1. For this, let f and g be distinct faces of $\mathrm{star}(X, \partial P)$ which both contain a vertex $u \notin X$. We prove that $f \cap g$ contains a vertex of X. It is clear that if either f or g contains X, then $f \cap g$ contains a vertex of X. So suppose that neither f nor g contains X. It follows that X contains more than one vertex. Suppose that X is an edge. Ampleness condition 2 applied to f, g and a face h of P which contains X shows that f, g and h have a vertex in common. Since two faces of P which have more than a vertex in common have exactly an edge in common by ampleness condition 1 and since h is not a triangle by ampleness condition 3, it easily follows that $f \cap g$ contains a vertex of X. Suppose that X is a face. Ampleness condition 2 applied to f, g and X implies that $f \cap g$ contains a vertex of X. So in every case $f \cap g$ contains a vertex v of X. Ampleness condition 1 implies that $f \cap g$ is an edge which joins u and v. Proposition 3.1 implies that $f \cap g$ is the unique edge of P which joins u and X if X is either a vertex or an edge. We next consider the uniqueness of this edge when X is a face. Let X be a face. Suppose that e is an edge of P other than $f \cap g$ which joins u and a vertex v' of X. Proposition 3.1 implies that $v' \neq v$ and even that no edge of P joins v' and v. If $v' \in f$, then v and v' are contained in $f \cap X$, and so ampleness condition 1 implies that $f \cap X$ is an edge joining v and v', contrary to the last sentence. So $v' \notin f$ and likewise $v' \notin g$. Let h be a face of P which contains e. Then ampleness condition 2 applied to f, h and X implies that $f \cap h \cap X$ consists of a vertex w. If $w = v$, then $h \cap X$ contains v and v', which is impossible as before. If $w \neq v$, then $f \cap X$ contains v and w, and so ampleness condition 1 implies that $f \cap X$ is an edge joining v and w. Similarly, $f \cap h$ is an edge joining u and w, while $f \cap g$ is an edge joining u and v. We now have a contradiction to Proposition 3.1. This proves statement 1 of Proposition 3.2.

Now consider statement 2. Let e be an edge of P which meets X but is not contained in X. Proposition 3.1 easily implies that e contains a vertex not in X if X is either a vertex or an edge. If X is a face, then ampleness condition 2 applied to X and the two faces which contain e implies that e contains a vertex not in X. This proves statement 2 of Proposition 3.2.

Statements 1 and 2 of Proposition 3.2 easily imply statement 3.

Statements 4 and 5 of Proposition 3.2 are now easy to see.

This proves Proposition 3.2. □

Let P be a faceted 3-ball with an orientation-reversing face-pairing ϵ. Let mul be a multiplier function for ϵ. Let Q be the twisted face-pairing subdivision of P gotten by subdividing the edges of P as usual. Let δ be the twisted face-pairing for Q determined by

a) 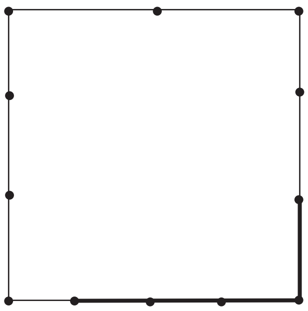 b) 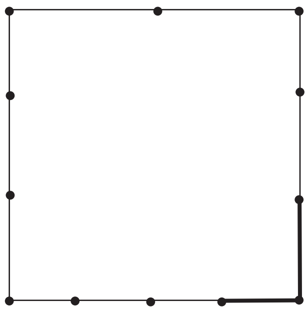

FIGURE 2. A twisted original edge and an elbow.

P, ϵ and mul. Let f be a face of Q, and let e be an original edge of f. Then we call $t = \delta_f(e)$ a **twisted original edge** of Q. Continuing, let L be a subcomplex of t consisting of some vertex of t together with all vertices and edges of t preceding (relative to f) that vertex such that L contains at least two edges. We call L an **elbow** of Q. Part a) of Figure 2 shows a twisted original edge of Q drawn with thick arcs. This twisted original edge of Q contains three elbows, and part b) of Figure 2 shows one of them drawn with thick arcs. Let L be an elbow contained in a face f of Q as above. Then there exist original edges e_1 and e_2 of Q such that $L \subseteq e_1 \cup e_2$. We call the original vertex common to e_1 and e_2 the **joint** of L. Suppose that e_1 immediately precedes e_2 (relative to f). Then we call $L \cap e_1$ the **bottom** of L, and we call $L \cap e_2$ the **top** of L. (Using clockwise orientation and viewing an elbow as the letter "L", the elbow's bottom is the horizontal segment of the letter "L" and its top is the vertical segment of the letter "L".) The bottom of an elbow always consists of one edge, and this edge might be an original edge. The top of an elbow may contain any number of edges, but it cannot be an original edge. Hence every elbow L contains at least one original vertex (its joint), and at most two; L contains two original vertices if and only if its bottom is an original edge. For emphasis we note that every elbow is contained in a twisted original edge, and every twisted original edge is an elbow unless it contains just one edge.

Proposition 3.3. *Let P be an ample faceted 3-ball with ϵ, mul and Q as usual. Suppose given a subcomplex $X = e \cup L \cup t$ of Q, where e is an original edge, L is either a vertex or an edge or a face or an elbow and t is a twisted original edge. Then every nontrivial simple closed edge path in X is contained in L.*

Proof. First suppose that L is an elbow. We wish to prove that every simple closed edge path in $e \cup L \cup t$ is trivial. Proposition 3.1 easily implies that every nontrivial simple closed edge path in $e \cup L \cup t$ meets each of e, L and t. Proposition 3.1 furthermore easily implies that we may assume that both L and t meet e. Statement 3 of Proposition 3.2 with $X = e$ implies that the union U of e and all the original edges of Q which meet e is contractible. Statement 4 of Proposition 3.2 with $X = e$ and the fact that every face of P has at least four sides implies that the union of U with the interior of every original edge of $\operatorname{star}(e, \partial Q)$ which meets U is also contractible. Hence every simple closed edge path in $e \cup L \cup t$ is trivial.

Now suppose that L is either a vertex or an edge or a face. Because a twisted original edge of Q contains at most one original edge, it is easy to see that to prove Proposition 3.3 we may assume that t is an original edge distinct from e. Proposition 3.1 implies that e and t have at most one vertex in common. It easily follows that we may assume that both e and

t meet L. Now statement 3 of Proposition 3.2 with $X = L$ implies that the set $e \cup L \cup t$ is contractible.

This proves Proposition 3.3. □

Before stating Proposition 3.4, we discuss some notation. Given cell complexes $X \subseteq Y$, we let $\mathrm{costar}(X, Y)$ denote the subcomplex of Y which is the union of all the closed cells of Y which are disjoint from X. Recall that X_σ denotes the dual cap subdivision of a given cell complex X.

Proposition 3.4. *Let P be an ample faceted 3-ball with ϵ, mul and Q as usual. Let L be a subcomplex of Q which is either a vertex or an edge or a face or an elbow. Then the following assertions hold.*

1. *The complex $star(L, \partial Q)$ is a topological disk.*
2. *If f and g are distinct faces of $star(L, \partial Q)$ which both contain a vertex $u \notin L$, then $f \cap g$ is the unique original edge of Q which joins u and L.*
3. *In the situation of statement 2, we have that $f \cap g \cap L$ is connected, and if t is a twisted original edge of Q which joins u and L, then $t \cap L$ is connected and the segment of t from u to L is contained in $f \cap g$.*
4. *The complex $star(L_\sigma, (\partial Q)_\sigma)$ is a topological disk.*
5. *The complex $star(L_\sigma, Q_\sigma)$ is a topological ball.*
6. *The complex $costar(L_\sigma, Q_\sigma)$ is a topological ball.*
7. *Let $\mathcal{F}$ be a set, possibly empty, of subcomplexes of $costar(L_\sigma, Q_\sigma)$ which have the form $f_\sigma \cap costar(L_\sigma, Q_\sigma)$, where f is a face of $star(L, \partial Q)$ other than possibly L. Then*
$$[star(L_\sigma, Q_\sigma) \cap costar(L_\sigma, Q_\sigma)] \cup \bigcup_{F \in \mathcal{F}} F$$
is a topological disk.

Proof. We first prove statement 1. If L is an original vertex, an original edge or a face, then this follows from statement 5 of Proposition 3.2. If L is a vertex or an edge that does not contain an original vertex, then $\mathrm{star}(L, \partial Q)$ is the union of two faces along a common original edge, and hence is a topological disk. If L is an edge which contains an original vertex v but is not an original edge, then $\mathrm{star}(L, \partial Q) = \mathrm{star}(v, \partial Q)$ and hence is a topological disk. Thus far we have proved statement 1 of Proposition 3.4 if L is either a vertex or an edge or a face. If L is an elbow, then either the bottom of L is an original edge or the joint of L is its only original vertex. In the first case $\mathrm{star}(L, \partial Q)$ is the star in ∂Q of the bottom of L, and in the second case $\mathrm{star}(L, \partial Q)$ is the star in ∂Q of the joint of L. In both cases this star is a topological disk by the above. This proves statement 1 of Proposition 3.4.

We next consider statement 2. If L does not contain an original vertex, then $f \cap g$ is the original edge that contains L and the statement is clear. Otherwise $\mathrm{star}(L, \partial Q)$ is the star in ∂Q of an original vertex, an original edge or a face, and statement 2 follows easily from Proposition 3.2. This proves statement 2 of Proposition 3.4.

To prove statement 3, first note that if L does not contain an original vertex, then $f \cap g$ is an original edge, $f \cap g \cap L = L$ and statement 3 of Proposition 3.4 follows easily. Now suppose that L contains an original vertex. Then $\mathrm{star}(L, \partial Q)$ is the star in ∂Q of an original vertex, an original edge or a face. Statement 4 of Proposition 3.2 easily implies that $f \cap g$ is an original edge with exactly one original vertex in L and hence $f \cap g \cap L$ is connected. The remaining part of statement 3 now follows from Proposition 3.3.

Statement 4 follows from a case-by-case analysis much like the proof of statement 1. Statement 5 is then immediate, since $\text{star}(L_\sigma, Q_\sigma)$ is the cone on $\text{star}(L_\sigma, (\partial Q)_\sigma)$ to the barycenter of Q. Since $\text{costar}(L_\sigma, Q_\sigma)$ is the closure of $Q_\sigma \setminus \text{star}(L_\sigma, Q_\sigma)$, statement 6 now follows from statements 4 and 5.

We finally consider statement 7. It follows from statements 4 and 5 that $\text{star}(L_\sigma, Q_\sigma) \cap \text{costar}(L_\sigma, Q_\sigma)$ is a disk. If L does not contain an original vertex, then there are exactly two faces f and g of $\text{star}(L, \partial Q)$. Each of $f_\sigma \cap \text{costar}(L_\sigma, Q_\sigma)$ and $g_\sigma \cap \text{costar}(L_\sigma, Q_\sigma)$ is a disk which intersects $\text{star}(L_\sigma, Q_\sigma) \cap \text{costar}(L_\sigma, Q_\sigma)$ in a nondegenerate arc, and their union is an annulus that has $\partial[\text{star}(L_\sigma, Q_\sigma) \cap \text{costar}(L_\sigma, Q_\sigma)]$ as one if its boundary components. Hence statement 7 holds in this case. Now suppose that L contains an original vertex and that f is a face of $\text{star}(L, \partial Q)$ with $f \neq L$. Then $f_\sigma \cap \text{costar}(L_\sigma, Q_\sigma)$ is a disk which intersects $\text{star}(L_\sigma, Q_\sigma) \cap \text{costar}(L_\sigma, Q_\sigma)$ in a nondegenerate arc. If g is another face of $\text{star}(L, \partial Q)$ with $g \neq L$, then $f_\sigma \cap g_\sigma \cap \text{costar}(L_\sigma, Q_\sigma)$ is either empty or is an arc which intersects $\text{star}(L_\sigma, Q_\sigma) \cap \text{costar}(L_\sigma, Q_\sigma)$ in exactly one of its endpoints. Statement 7 now follows easily by induction on the cardinality of $\mathcal{F}$.

This completes the proof of Proposition 3.4. □

4. Preparations

In this section we fix notation, recall some facts and make some definitions to prepare for the next section. We again draw the reader's attention to Section 7, where examples are discussed.

Let P be a faceted 3-ball, let ϵ be an orientation-reversing face-pairing on P and let mul be a multiplier function for ϵ. From this we obtain a twisted face-pairing subdivision Q of P, a face-pairing δ on Q, and a twisted face-pairing 3-manifold M. As in Section 4 of [8], we also obtain a faceted 3-ball Q^* dual to Q, meaning that there exists an orientation-reversing isomorphism $\omega : Q \to Q^*$, and we obtain a twisted face-pairing 3-manifold M^* dual to M. We let $\widetilde{M}$ denote the universal covering cell complex of M, and we let $\widetilde{M}^*$ denote the universal covering cell complex of M^*. The cell complexes M and M^* are dual to each other with $|M| = |M^*|$, and the cell complexes $\widetilde{M}$ and $\widetilde{M}^*$ are dual to each other with $|\widetilde{M}| = |\widetilde{M}^*|$. We choose a vertex $\mathcal{O}$ of $\widetilde{M}$, and we call $\mathcal{O}$ the base vertex of $\widetilde{M}$. Likewise we choose a vertex $\mathcal{O}^*$ of $\widetilde{M}^*$, and we call $\mathcal{O}^*$ the base vertex of $\widetilde{M}^*$.

Let G denote the fundamental group of M. We view G as simultaneously acting on both $\widetilde{M}$ and $\widetilde{M}^*$. We define the **star generating set** of G to be the set of all elements g in G such that $C \cap gC \neq \emptyset$, where C is the 3-cell of $\widetilde{M}$ dual to $\mathcal{O}^*$. Since the star generating set contains the geometric generating set (see Theorem 4.8 of [8] together with the material just before and after), it does indeed generate G. Let Γ_s denote the Cayley graph of G with respect to its star generating set.

We next define for every nonnegative integer k a cell complex $B(\frac{k}{2})$, which we refer to as a **combinatorial ball**. We define $B(0)$ to be simply $\mathcal{O}^*$. We proceed inductively as follows. If $B(k)$ is defined for a nonnegative integer k, then we let $B(k+\frac{1}{2})$ be the union of the 3-cells in $\widetilde{M}$ which are dual to the vertices of $B(k)$. If $B(k+\frac{1}{2})$ is defined for a nonnegative integer k, then we let $B(k+1)$ be the union of the 3-cells in $\widetilde{M}^*$ which are dual to the vertices of $B(k+\frac{1}{2})$. The 3-cells of $B(k+\frac{1}{2})$ correspond to the vertices of Γ_s which lie within distance k from the identity element for every nonnegative integer k. Equivalently, the vertices of

$B(k)$ correspond to the vertices of Γ_s which lie within distance k from the identity element for every nonnegative integer k. We find it convenient to set $B(-1) = B(-\frac{1}{2}) = \emptyset$.

We next define the **bipartite graph** Γ_b as follows. The vertices of Γ_b consist of the vertices of $\widetilde{M}$ together with the vertices of $\widetilde{M}^*$. Edges of Γ_b join vertices of $\widetilde{M}$ with vertices of $\widetilde{M}^*$. A vertex x of $\widetilde{M}$ is connected to a vertex y of $\widetilde{M}^*$ by an edge of Γ_b if and only if the 3-cell of $\widetilde{M}^*$ dual to x contains y, equivalently, the 3-cell of $\widetilde{M}$ dual to y contains x. It is clear that Γ_b is indeed bipartite. We take $\mathcal{O}^*$ to be the base vertex of Γ_b. We put an edge path metric on Γ_b so that every edge has length $\frac{1}{2}$. Then the vertices of Γ_b at distance k from $\mathcal{O}^*$ are the vertices of $B(k) \setminus B(k-1)$ for every half integer $k \geq 0$. The action of G on $\widetilde{M}$ and $\widetilde{M}^*$ determines an action of G on Γ_b.

The main theorem of this paper states that if P is ample, then G is Gromov hyperbolic. We prove this in Section 6 using our Gromov hyperbolicity criterion given in Theorem 2.2 applied to the Cayley graph Γ_s of G with respect to its star generating set. Just as in Theorem 4.10 of [8], there exists a unique G-equivariant isometry from the set of vertices of Γ_s to the bipartite graph Γ_b which maps the base vertex of Γ_s to the base vertex $\mathcal{O}^*$ of Γ_b (and whose image consists of "half" the vertices of Γ_b). This isometry allows us to transform the problem of verifying our Gromov hyperbolicity criterion for Γ_s to verifying our Gromov hyperbolicity criterion for Γ_b. We find it easier to work with Γ_b than Γ_s.

In this paragraph we define the notion of cone type for Γ_b. Let $x \in \Gamma_b$. The **cone** of the pair $(\Gamma_b, \mathcal{O}^*)$ based at x is the set $[x, \infty)$ of all points $y \in \Gamma_b$ such that there exists a geodesic from $\mathcal{O}^*$ to y which passes through x. Now let x and y be vertices of Γ_b. We say that the cones $[x, \infty)$ and $[y, \infty)$ are equivalent if there exists an element of G which maps x to y taking $[x, \infty)$ bijectively onto $[y, \infty)$. We call an equivalence class of cones a **cone type**.

5. Growth properties of the universal cover of M

This section is devoted to proving Theorem 5.1, which deals with what might be called growth properties of the universal covering cell complex of M.

Theorem 5.1. *Let P be an ample faceted 3-ball, let ϵ be an orientation-reversing face-pairing on P, and let mul be a multiplier function for ϵ. Let Q be the associated twisted face-pairing subdivision of P, and let M be the associated twisted face-pairing manifold. Then the following assertions hold.*

1. *Every lift of the quotient map from Q to M to a map from Q to the universal cover $\widetilde{M}$ of M is injective. Hence every 3-cell of $\widetilde{M}$ is canonically isomorphic to Q.*
2. *The combinatorial ball $B(k)$ is a topological ball contained in the interior of $B(k+1)$ for every positive half integer k.*
3. *The bipartite graph Γ_b has finitely many cone types.*
4. *Let x be a vertex of Γ_b, and suppose that the Γ_b-distance from x to $\mathcal{O}^*$ is k with $k > 0$. Then the dual of link$(x, B(k))$ is isomorphic to a subcomplex $L(x)$ of either Q if k is an integer or Q^* if k is not an integer, and this isomorphism is canonically defined on vertices. Furthermore, the cone type of x is determined by $L(x)$.*
5. *Maintain the setting of statement 4. Then $L(x)$ is either a vertex or an edge or a face or an elbow.*
6. *Suppose given a half integer $k \geq 1$. Then every 3-cell in $B(k)$ not in $B(k-1)$ meets $\partial B(k)$ in a connected, simply-connected complex whose interior is a topological disk. (It is a "hairy" disk.)*

7. *Suppose given a half integer $k \geq 1$, and let D be a 3-cell in $B(k)$ not in $B(k-1)$. Let $L = D \cap B(k-1)$. Then L is either a vertex or an edge or a face or an elbow of D. There exists a face f of $\partial B(k-1)$ and a 3-cell C of $B(k-1)$ with $L \subseteq f \subseteq C$ such that L is either a vertex or an edge or a face or an elbow of C. Moreover if L contains more than two vertices, then there exists just one face f of $B(k-1)$ with $L \subseteq f$.*

Remark 5.2. Statement 2 asserts that $B(k)$ is contained in the interior of $B(k+1)$. By this we mean that every point of $|B(k)|$ has a neighborhood in $|B(k+1)|$ which is homeomorphic to an open 3-ball. Recall that we defined elbows of Q, and hence Q^*, in the middle of Section 3. According to statement 1, every 3-cell of $\widetilde{M}$ is canonically isomorphic to Q. Likewise every 3-cell of $\widetilde{M}^*$ is canonically isomorphic to Q^*. This justifies the use of the word elbow in statement 7.

Proof of Theorem 5.1. Most of the proof of Theorem 5.1 consists of inductively constructing cell complexes $\mathcal{B}(0), \mathcal{B}(\frac{1}{2}), \mathcal{B}(1), \ldots$, developing their properties and eventually identifying them with $B(0), B(\frac{1}{2}), B(1), \ldots$.

We assume that P has n face-pairs so that the faces and edges of Q and Q^* are labeled with the elements of $\{1, \ldots, n\}$.

We begin by defining $\mathcal{B}(0)$ to be simply a vertex. We define $\mathcal{B}(\frac{1}{2})$ to be a copy of Q with its edge and face labels and directions. Let k be half a nonnegative integer. We assume that the cell complex $\mathcal{B}(k)$ has been defined and satisfies the following four induction hypotheses if $k > 0$.

Induction hypothesis 1. The cell complex $\mathcal{B}(k)$ is oriented and equal to the union of its 3-cells. If $k \geq 1$, then $\mathcal{B}(k)$ contains $\mathcal{B}(k-1)$ in its interior. Every edge and face of $\mathcal{B}(k)$ is labeled with an element of $\{1, \ldots, n\}$ and directed so that every 3-cell of $\mathcal{B}(k)$ is isomorphic to Q in a way which preserves orientation, labels and directions if k is not an integer and every 3-cell of $\mathcal{B}(k)$ is isomorphic to Q^* in a way which preserves orientation, labels and directions if k is an integer.

Induction hypothesis 2. The dual of the $\mathcal{B}(k)$-link of every vertex of $\mathcal{B}(k) \backslash \partial\mathcal{B}(k)$ is isomorphic to Q if k is an integer, and it is isomorphic to Q^* if k is not an integer. The dual of the $\mathcal{B}(k)$-link of every vertex of $\partial\mathcal{B}(k)$ is isomorphic to either a vertex or an edge or a face or an elbow of Q if k is an integer, and it is isomorphic to either a vertex or an edge or a face or an elbow of Q^* if k is not an integer. These isomorphisms are canonically defined on vertices. They are gotten as follows. Let v be a vertex of $\mathcal{B}(k)$. Suppose that k is not an integer. Let C be a 3-cell of $\mathcal{B}(k)$ which contains v. Then C is isomorphic to Q by induction hypothesis 1, and it is easy to see that this isomorphism is uniquely defined on vertices. The C-link of v thus determines a face of the link of the vertex of M. This map from the faces of the $\mathcal{B}(k)$-link of v to the faces of the link of the vertex of M determines an isomorphism from the $\mathcal{B}(k)$-link of v to a subcomplex of the link of the vertex of M. Once the map $\omega : Q \to Q^*$ is fixed, the dual of the link of the vertex of M is isomorphic to Q^* in a way which is canonically determined on vertices by Theorem 3.1 of [8], and so it follows that the dual of the $\mathcal{B}(k)$-link of v is isomorphic to a subcomplex of Q^* in a way which is canonical on vertices. If k is an integer, then the same holds with Q, Q^* and M replaced by Q^*, Q and M^*.

Induction hypothesis 3. There is a duality between the closed cells of $\mathcal{B}(k-\frac{1}{2})$ and the closed cells of $\mathcal{B}(k)$ which are not contained in $\partial\mathcal{B}(k)$. This duality preserves labels and directions of edges and faces.

Induction hypothesis 4. The complex $\mathcal{B}(k)$ is a topological ball.

It is clear that $\mathcal{B}(\frac{1}{2})$ satisfies all four induction hypotheses.

We construct $\mathcal{B}(k+\frac{1}{2})$ for the case in which k is a positive integer. The case in which k is not an integer is gotten from the case in which k is an integer essentially by interchanging Q and Q^*. So let k be a positive integer. We fix this value of k until the end of the verification of induction hypothesis 4 near the end of this section.

We construct $\mathcal{B}(k+\frac{1}{2})$ in this paragraph. Let V_k be the set of vertices of $\partial\mathcal{B}(k)$. For every $v \in V_k$, let Q_v be a copy of Q with its edge and face labels and directions together with orientation. Let $\mathcal{B}'(k+\frac{1}{2})$ be the disjoint union of $\mathcal{B}(k-\frac{1}{2})$ and Q_v for every $v \in V_k$. We next define what might be called a face-pairing β for $\mathcal{B}'(k+\frac{1}{2})$, although not all the faces of $\mathcal{B}'(k+\frac{1}{2})$ are paired with other faces of $\mathcal{B}'(k+\frac{1}{2})$ by maps of β. Suppose that u and v are vertices of $\partial\mathcal{B}(k)$ which are joined by an edge e of $\partial\mathcal{B}(k)$. Then the label and direction of e given by induction hypothesis 1 determine a face f_u of Q_u and a face f_v of Q_v. The isomorphisms between Q_u, Q_v and Q together with the twisted face-pairing δ on Q determine inverse cellular homeomorphisms $\beta_{f_u}: f_u \to f_v$ and $\beta_{f_v}: f_v \to f_u$. We include β_{f_u} and β_{f_v} in the face-pairing β. Suppose that v is a vertex of $\partial\mathcal{B}(k)$ and that e is an edge of $\mathcal{B}(k)$ which contains v but $e \not\subseteq \partial\mathcal{B}(k)$. Let u be the vertex of e other than v. Induction hypothesis 3 implies that e determines a face f_u of $\mathcal{B}(k-\frac{1}{2})$. Because the vertex v lies in $\partial\mathcal{B}(k)$, the face f_u lies in $\partial\mathcal{B}(k-\frac{1}{2})$ and induction hypothesis 3 implies that the 3-cell of $\mathcal{B}(k-\frac{1}{2})$ which contains f_u is dual to u and $u \in \mathcal{B}(k-1)$. As before, e determines a face f_v of Q_v and we have inverse cellular homeomorphisms $\beta_{f_u}: f_u \to f_v$ and $\beta_{f_v}: f_v \to f_u$. We include β_{f_u} and β_{f_v} in the face-pairing β. This completes the definition of β. It is easy to see that β satisfies a face-pairing compatibility condition as in Section 2 of [8]. We define $\mathcal{B}(k+\frac{1}{2})$ to be the cell complex consisting of orbits of points of $\mathcal{B}'(k+\frac{1}{2})$ under β.

We make the following definition in order to investigate $\mathcal{B}(k+\frac{1}{2})$. Let $v \in V_k$. By induction hypothesis 2, the dual of $\operatorname{link}(v,\mathcal{B}(k))$ is isomorphic to and uniquely determines either a vertex or an edge or a face or an elbow of Q. By this means we identify the dual of $\operatorname{link}(v,\mathcal{B}(k))$ with a subcomplex L_v of Q_v. We next present two lemmas which deal with the complexes L_v.

Lemma 5.3. *Let $v \in V_k$. Then the faces of Q_v which are paired with other faces of $\mathcal{B}'(k+\frac{1}{2})$ by β are precisely the faces of $star(L_v, \partial Q_v)$.*

Proof. This follows from the definitions and induction hypotheses. □

Lemma 5.4. *Let e be an edge of $\mathcal{B}(k)$ with vertices v_1 and v_2 such that $v_1 \in \partial\mathcal{B}(k)$. Let C be a 3-cell of $\mathcal{B}(k)$ which contains e. Let u_1 be the vertex of $L_{v_1} \subseteq Q_{v_1}$ determined by C. If $v_2 \in \partial\mathcal{B}(k)$, then let u_2 be the vertex of $L_{v_2} \subseteq Q_{v_2}$ determined by C. If $v_2 \notin \partial\mathcal{B}(k)$, then let Q_{v_2} be the 3-cell of $\mathcal{B}(k-\frac{1}{2})$ dual to v_2, and let u_2 be the vertex of Q_{v_2} dual to C. Then β identifies the face of Q_{v_1} corresponding to e with the face of Q_{v_2} corresponding to e, and the vertices u_1 and u_2 are identified by means of this face identification. In particular, if w_1 is a vertex of Q_{v_1}, if $v_2 \in \partial\mathcal{B}(k)$ and if w_2 is a vertex of Q_{v_2} such that w_1 and w_2 are identified by means of this face identification, then $w_1 \in L_{v_1}$ if and only if $w_2 \in L_{v_2}$.*

Proof. The vertex v_1 of C corresponds to a vertex of Q^*, which by means of the map $\omega: Q \to Q^*$ corresponds to a vertex w_1 of Q, which corresponds to the vertex u_1 of Q_{v_1}. The vertex v_2 of C corresponds to a vertex w_2 of Q and the vertex u_2 of Q_{v_2} in the same way. The vertex w_1 determines a face $f(w_1)$ of the link of the vertex of M, and the vertex w_2

determines a face $f(w_2)$ of the link of the vertex of M. Theorem 3.1 of [8] implies that the edge e determines an edge common to $f(w_1)$ and $f(w_2)$. In other words, e determines a face-pair of Q; one face of this face-pair contains w_1 and the other face of this face-pair contains w_2 so that w_1 and w_2 are identified by the corresponding face-pairing maps. It follows that u_1 and u_2 are identified. This proves Lemma 5.4. □

We are now able to verify the first induction hypothesis for $\mathcal{B}(k+\frac{1}{2})$.

Verification of induction hypothesis 1. It is clear that $\mathcal{B}(k+\frac{1}{2})$ is oriented and equal to the union of its 3-cells. The statements concerning labels and directions are also clear. It remains to show that $\mathcal{B}(k-\frac{1}{2})$ can be identified with a subcomplex in the interior of $\mathcal{B}(k+\frac{1}{2})$. For this let w be a vertex of $\partial\mathcal{B}(k-\frac{1}{2})$. Let f be a face of $\partial\mathcal{B}(k-\frac{1}{2})$ which contains w. By induction hypothesis 3 the face f is dual to an edge e of $\mathcal{B}(k)$ not contained in $\partial\mathcal{B}(k)$. Since $f\subseteq\partial\mathcal{B}(k-\frac{1}{2})$, one vertex v of e is contained in $\partial\mathcal{B}(k)$. It follows easily that $\text{link}(v,\mathcal{B}(k))$ contains a face isomorphic to f, and induction hypothesis 2 then implies that $\text{link}(v,\mathcal{B}(k))$ is in fact a face isomorphic to f. Hence the two face-pairing maps of β corresponding to e identify f with L_v. Let C be the 3-cell of $\mathcal{B}(k)$ dual to w. Lemma 5.4 implies that the vertex of L_v identified with w is the one which is determined by C. Lemma 5.4 furthermore implies that if u is any vertex in $\partial\mathcal{B}(k)$ such that w is identified with some vertex x of Q_u by a sequence of face-pairing maps of β, then $u\in C$ and x is the vertex of L_u determined by C. Since w is the unique vertex of $\partial\mathcal{B}(k-\frac{1}{2})$ dual to C, it follows that the set of vertices of $\partial\mathcal{B}(k-\frac{1}{2})$ injects into $\mathcal{B}(k+\frac{1}{2})$. Hence the set of vertices of $\mathcal{B}(k-\frac{1}{2})$ injects into $\mathcal{B}(k+\frac{1}{2})$. Because the edges and faces of Q are determined by their vertices, it easily follows that $\mathcal{B}(k-\frac{1}{2})$ injects into $\mathcal{B}(k+\frac{1}{2})$. We identify $\mathcal{B}(k-\frac{1}{2})$ with its image in $\mathcal{B}(k+\frac{1}{2})$. Given a vertex w of $\partial\mathcal{B}(k-\frac{1}{2})$, it is clear that $\mathcal{B}(k+\frac{1}{2})$ is defined to complete $\text{link}(w,\mathcal{B}(k-\frac{1}{2}))$ to a 2-sphere. Hence $\mathcal{B}(k-\frac{1}{2})$ is contained in the interior of $\mathcal{B}(k+\frac{1}{2})$. This completes the verification of induction hypothesis 1 for $\mathcal{B}(k+\frac{1}{2})$.

Let $v\in V_k$. The previous paragraph shows that if a vertex x of Q_v is identified with a vertex of $\mathcal{B}(k-\frac{1}{2})$, then $x\in L_v$. In this paragraph we show that every vertex of L_v is identified with a vertex of $\mathcal{B}(k-\frac{1}{2})$. Let x be a vertex of L_v. Let C be the 3-cell of $\mathcal{B}(k)$ which contains v corresponding to x. Then C is dual to a vertex w of $\mathcal{B}(k-\frac{1}{2})$ by induction hypothesis 3. Induction hypothesis 1 implies that some 3-cell of $\mathcal{B}(k-\frac{1}{2})$ contains w, and induction hypothesis 3 implies that such a 3-cell containing w is dual to a vertex of $\mathcal{B}(k-1)$ contained in C. Hence C meets $\mathcal{B}(k-1)$. Let γ be a minimal edge path in C from v to $\mathcal{B}(k-1)$. The definition of β and induction hypothesis 1 show that the last edge of γ contains a vertex of $\partial\mathcal{B}(k)$ and a vertex of $\mathcal{B}(k-1)$. Now Lemma 5.4 applied to the edges of γ shows that the identifications determined by β identify x with a vertex of $\partial\mathcal{B}(k-\frac{1}{2})$. Thus the vertices of L_v are exactly the vertices of Q_v which are identified with vertices of $\mathcal{B}(k-\frac{1}{2})$.

We now turn our attention to the vertices of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$, beginning with some definitions. We say that the $\mathcal{B}(k+\frac{1}{2})$-link of a vertex v of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$ is **small** if it contains either one or two faces. In this case we call v a **small link vertex**. We say that the $\mathcal{B}(k+\frac{1}{2})$-link of a vertex v of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$ is **big** if it contains at least three faces. In this case we call v a **big link vertex**. Let u be a vertex of $\mathcal{B}'(k+\frac{1}{2})$ whose image in $\mathcal{B}(k+\frac{1}{2})$ does not lie in $\mathcal{B}(k-\frac{1}{2})$. Let v be the element of V_k such that u lies in Q_v. Suppose that there exists a face f of Q_v which contains u such that f is paired with another face of $\mathcal{B}'(k+\frac{1}{2})$ by β. We say that u is **one-sided** if there exists exactly one such face f. If there exists more than one such face of Q_v, then we say that u is **two-sided**. (We will see in

Lemma 5.5 that there exist exactly two such faces if u is two-sided.) Given a subcomplex X of $\mathcal{B}'(k+\frac{1}{2})$, we let $\overline{X}$ denote the image of X in $\mathcal{B}(k+\frac{1}{2})$. In particular, we use this notation when X is a vertex. It is clear that if x is a big link vertex of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$, then there exists a two-sided vertex u of $\mathcal{B}'(k+\frac{1}{2})$ such that $\overline{u}=x$.

We next investigate the big link vertices of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$. Let x be a big link vertex of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$. Let u be a two-sided vertex of $\mathcal{B}'(k+\frac{1}{2})$ such that $\overline{u}=x$. Let v be the element of V_k such that u lies in Q_v, and let f and g be distinct faces of Q_v both of which contain u such that both f and g are paired with faces of $\mathcal{B}'(k+\frac{1}{2})$ by β. Lemma 5.3 states that both f and g are faces of $\operatorname{star}(L_v,\partial Q_v)$. Since $\overline{u}$ is not in $\mathcal{B}(k-\frac{1}{2})$, it follows that u is not in L_v. Now statement 2 of Proposition 3.4 implies that $f\cap g$ is the unique original edge of Q_v which joins u to L_v. The following lemma is one implication of this, which we record for future reference.

Lemma 5.5. *Every two-sided vertex of $\mathcal{B}'(k+\frac{1}{2})$ is contained in exactly two faces of $\mathcal{B}'(k+\frac{1}{2})$ which are paired with other faces of $\mathcal{B}'(k+\frac{1}{2})$ by the face-pairing β.*

Lemma 5.5 implies that there exist either zero or two one-sided vertices of $\mathcal{B}'(k+\frac{1}{2})$ which map to x. If there are two such one-sided vertices, then the vertices of $\mathcal{B}'(k+\frac{1}{2})$ which map to x are linearly ordered, and if there are zero such one-sided vertices, then the vertices of $\mathcal{B}'(k+\frac{1}{2})$ which map to x are cyclically ordered. Returning to the situation immediately before Lemma 5.5, we see that statement 3 of Proposition 3.4 implies that the nonempty complex $f\cap g\cap L_v$ is connected, and so there exists a vertex w of $f\cap g\cap L_v$ which is nearest u relative to $f\cap g$. We call w the **root vertex** of u, and we call the edge r of $f\cap g$ which contains w and whose interior separates u from w the **root edge** of u.

In this paragraph we show that $\overline{w}$ and $\overline{r}$ are independent of the choice of u. For this at first maintain the setting of the previous paragraph and suppose given $v'\in V_k$ for which $Q_{v'}$ contains distinct faces f' and g' such that 1) f' is paired with g by β, 2) g' is paired with some face of $\mathcal{B}'(k+\frac{1}{2})$ by β and 3) g' contains $u'=\beta_g(u)$. As in the previous paragraph, $f'\cap g'$ is the unique original edge of $Q_{v'}$ which joins u' to $L_{v'}$. The last sentence of Lemma 5.4 implies that $\beta_g(w)$ lies in $L_{v'}$. It follows that $\beta_g(f\cap g)$ is a twisted original edge of $Q_{v'}$ which joins u' to $L_{v'}$. Now statement 3 of Proposition 3.4 easily implies that $\beta_g(w)$ is the root vertex of u' and that $\beta_g(r)$ is the root edge of u'. Finally, if u and u' are any two-sided vertices of $\mathcal{B}'(k+\frac{1}{2})$ with $\overline{u}=\overline{u'}=x$, then there exist two-sided vertices $u_0=u,\ldots,u_j=u'$ of $\mathcal{B}'(k+\frac{1}{2})$ such that u_{i-1} maps to u_i by an element of β for every $i\in\{1,\ldots,j\}$. The above argument in this paragraph shows that the root vertex and root edge of u_{i-1} map to the root vertex and root edge of u_i by an element of β for every $i\in\{1,\ldots,j\}$. Thus $\overline{w}$ and $\overline{r}$ are independent of the choice of u and are uniquely determined by x.

We call $\overline{w}$ the **root vertex** of x, and we call $\overline{r}$ the **root edge** of x. It is clear that $\overline{w}$ is a vertex of $\partial\mathcal{B}(k-\frac{1}{2})$ and that the vertex of $\overline{r}$ other than $\overline{w}$ is not a vertex of $\mathcal{B}(k-\frac{1}{2})$. It is furthermore now easy to see that if a 3-cell of $\mathcal{B}(k+\frac{1}{2})$ contains a big link vertex x of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$, then that 3-cell contains the root edge and the root vertex of x.

We next investigate the edges of $\mathcal{B}(k+\frac{1}{2})$ which are root edges of big link vertices of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$. Let r be such a root edge. Let w be the root vertex contained in r. Since w is a vertex of $\partial\mathcal{B}(k-\frac{1}{2})$, induction hypothesis 3 implies that w is dual to a 3-cell C of $\mathcal{B}(k)$. It is not difficult to see that $\mathcal{B}(k+\frac{1}{2})$ is defined so that the dual of $\operatorname{link}(w,\mathcal{B}(k+\frac{1}{2}))$ is isomorphic to C. Induction hypothesis 3 implies that the faces of C which are not contained in $\partial\mathcal{B}(k)$ are dual to the edges of $\mathcal{B}(k-\frac{1}{2})$ which contain w. It follows that r is dual to a face f of C and that $f \subseteq \partial\mathcal{B}(k)$. Hence the 3-cells of $\mathcal{B}(k+\frac{1}{2})$ which contain r are exactly those which are dual to the vertices of f.

In this paragraph we prepare to describe how these 3-cells of $\mathcal{B}(k+\frac{1}{2})$ dual to the vertices of f fit together by constructing a topological space F^+ with some further structure which is closely related to part of $\mathcal{B}(k+\frac{1}{2})$. To help understand this construction, we refer the reader to the examples in Section 7. We call F^+ together with its further structure the **rooted cosubdivision** of f. Figure 3 gives a diagram of F^+. Keep in mind our conventions that in figures faces of Q and Q^* are oriented clockwise and that corners of f indicate original vertices. Figure 3 shows the face f with its edge labels and directions. The face f is part of the dual of $\mathcal{B}(k+\frac{1}{2})$, and so it is drawn with dashes. Not appearing explicitly in Figure 3 is the fact that the face label of f is 1 and that f is directed upward. We let F be a topological space which is a closed topological disk with a fixed homeomorphism to the underlying space of f. By means of this homeomorphism we speak of points of F as corresponding to points of f. In Figure 3 the space F is drawn above f. We next let r_F be a 1-cell, and we identify one of its vertices with a point in the interior of F. Let $F^+ = F \cup r_F$. We refer to r_F as the **root edge** of F^+. In Figure 3, we view r_F as passing through f. Since r_F is below F, it is drawn with dashes. We label and direct r_F in agreement with f. Let w_F be the vertex of r_F not in F. We next construct a finite tree $t^+ \subseteq F^+$ which we refer to as a **rooted face tree**. In Figure 3 the edges of t^+ are drawn as thick arcs or as a thick dashed arc. We construct t^+ as follows. The edge r_F is contained in t^+. The vertex of r_F other than w_F is called the **central vertex** of t^+. We fix an original edge e of f for the rest of this paragraph and the next paragraph. Let j be the number of edges in e. If $j > 1$, then we choose an arc in the interior of F with one endpoint the central vertex of t^+, and we subdivide this arc into $j-1$ subarcs. We view this subdivided arc as a 1-complex b and refer to it as the **branch** of t^+ associated to e. See Figure 4, where the original edge e has vertices $u_0, \dots, u_4$ and its associated branch b has vertices $v_1, \dots, v_4$. If $j = 1$, then we simply define b to be the central vertex of t^+. We choose these branches so that they meet one another only at the central vertex of t^+, hence their union is a contractible subset of F, and we let t^+ be the union of these branches together with r_F. We let $t = t^+ \cap F$, and we refer to t as simply a **face tree**. We next choose more arcs in F and label and direct the edges of the branches of t as follows. See Figure 4. Let $e_1, \dots, e_j$ be the edges of e so that e_{i-1} follows e_i for every $i \in \{2, \dots, j\}$. Let $v_1, \dots, v_j$ be the vertices of b so that v_1 is the central vertex of t and v_{i-1} is joined to v_i by an edge of b for every $i \in \{2, \dots, j\}$. We refer to v_j as the **terminal vertex** of b. We choose a point of F corresponding to a point of f in the interior of e_1, and we choose an arc in F which joins this point to the terminal vertex of the branch of t associated to the original edge of f following e. If $j > 1$, then for every $i \in \{2, \dots, j\}$ we choose a point of F corresponding to a point of f in the interior of e_i, and we choose an arc in F which joins this point to v_i. We label and direct the edge of b from v_{i-1} to v_i just as e_i is labeled and directed. These arcs from the boundary of F to t are chosen so that their union together with t is a contractible subset of F. We denote this contractible subset of F by T, and we refer to T as a **big face tree**. This completes the definition of the rooted cosubdivision F^+ of f. We refer to F together with T and the cell structure of t as the **cosubdivision** of f. We note that the cosubdivision which appears in Figure 3 also appears in the front face in Figure 34 (with more vertices and edges).

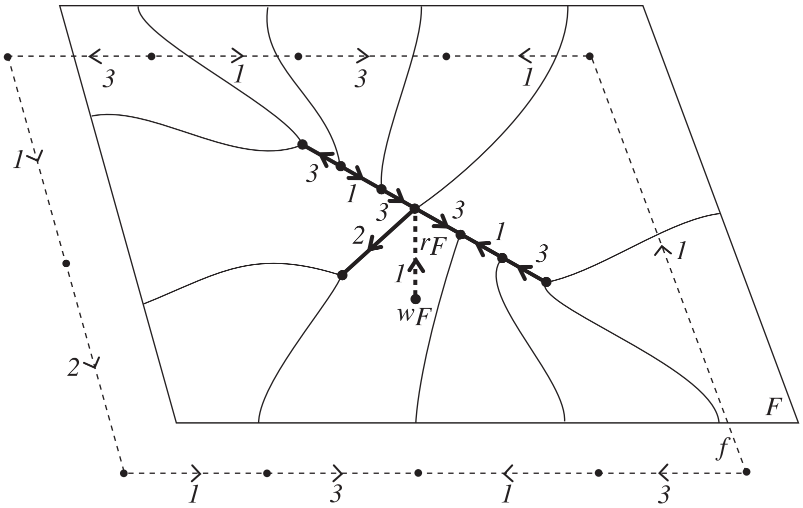

FIGURE 3. The rooted cosubdivision F^+ of f.

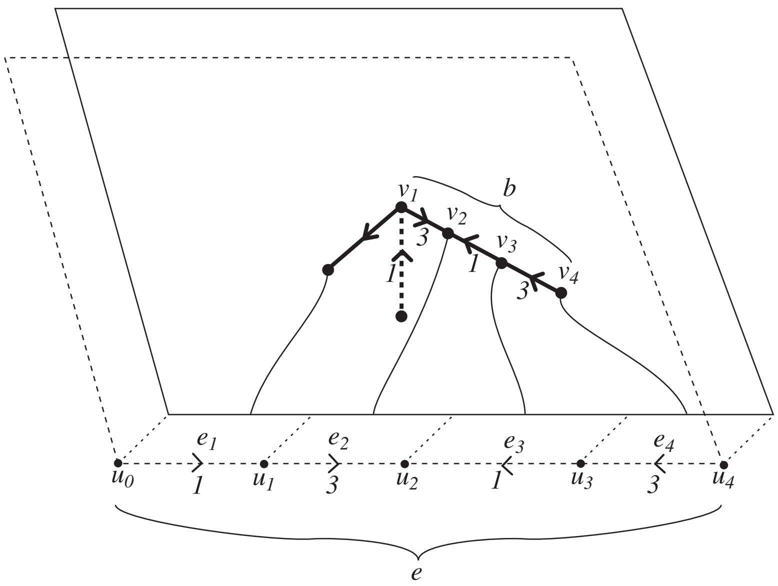

FIGURE 4. The branch b of F^+ corresponding to the original edge e.

In this paragraph we describe how the 3-cells of $\mathcal{B}(k+\frac{1}{2})$ dual to the vertices of f fit together. Verification of the assertions in this paragraph is left largely to the reader, although there is a brief discussion of the verification at the end of the paragraph. There exists a tree isomorphism from the rooted face tree t^+ to a subcomplex of $\mathcal{B}(k+\frac{1}{2})$ which takes r_F to r and respects edge labels and directions. The vertices of t map bijectively to the big link vertices of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$ with root edge r. The above tree isomorphism extends to a homeomorphism φ from F^+ to a subset of $\partial\mathcal{B}(k+\frac{1}{2})\cup r$ such that the big face tree T is mapped by φ into the 1-skeleton of $\partial\mathcal{B}(k+\frac{1}{2})$. The closure of every connected component of $F\setminus T$ is mapped by φ to a subset of $\partial\mathcal{B}(k+\frac{1}{2})$ contained in the boundary of a 3-cell of $\mathcal{B}(k+\frac{1}{2})$ containing r. Returning to the original edge e of f, we let $u_0,\dots,u_j$ denote the vertices of e ordered so that u_{i-1} follows u_i for every $i\in\{1,\dots,j\}$. Let C_i be the 3-cell of

$\mathcal{B}(k+\frac{1}{2})$ which is dual to u_i for every $i \in \{1,\dots,j\}$. Then the vertex v_i of b is mapped by φ to an original vertex of C_i for every $i \in \{1,\dots,j\}$. Furthermore, the original vertex u_j of $f \subseteq C$ corresponds to an original vertex of Q^* via the isomorphism between C and Q^*, which corresponds to an original vertex of Q via the orientation-reversing cellular homeomorphism $\omega : Q \to Q^*$, which corresponds to the original vertex $w = \varphi(w_F)$ of C_j via the isomorphism between C_j and Q. Similarly, the original vertex u_0 of $f \subseteq C$ corresponds to the original vertex $\varphi(v_j)$ of C_j. Induction hypothesis 4 implies that $\partial\mathcal{B}(k)$ is a 2-sphere. It follows that $\partial\mathcal{B}(k+\frac{1}{2})$ is a 2-sphere and as f varies over the faces of $\partial\mathcal{B}(k)$, the maps φ can be defined so that the images $\varphi(F)$ cover $\partial\mathcal{B}(k+\frac{1}{2})$ and intersect only on their boundaries. The points in the union of the images of the big face trees are precisely those points of the 1-skeleton of $\partial\mathcal{B}(k+\frac{1}{2})$ which lie in at least two 3-cells of $\mathcal{B}(k+\frac{1}{2})$. It is now clear that every vertex of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$ lies in $\partial\mathcal{B}(k+\frac{1}{2})$. We finally identify the face tree t with $\varphi(t) \subseteq \partial\mathcal{B}(k+\frac{1}{2})$, and we refer to $\varphi(t)$ as the **face tree** of f. The convention that corners indicate original vertices is not in force at vertices of the face tree in Figure 3. To verify the above assertions, we recommend to begin with C_j. Because the 3-cells dual to the vertices of C are glued together in accordance with the twisted face-pairing δ, it follows from Theorem 3.1 of [8] that the vertex u_j of C corresponds to the vertex w of C_j. Theorem 4.2 of [8] shows that just as the face f containing e is labeled with the label of e_1 and directed accordingly, the other face of C which contains e is labeled with the label of e_j and directed accordingly. The two corresponding faces of C_j are the faces of C_j dual to the edges of f which contain u_j. Given the labels and directions of the edges of e, Theorem 4.2 of [8] shows that the edges in the intersection of these two faces of C_j are labeled and directed in accordance with the labels and directions of the edges of $b \cup r_F$. Now consider the attachment of C_{j-1}, C_{j-2}, ... in succession.

With the results of the previous paragraph we can now complete the verification of the induction hypotheses for $\mathcal{B}(k+\frac{1}{2})$. We use the results of the previous paragraph freely.

Verification of induction hypothesis 2. The first sentence of induction hypothesis 2 is now clear. Let v be a vertex of $\partial\mathcal{B}(k+\frac{1}{2})$. If v is a small link vertex, then it is easy to see that the dual of $\text{link}(v,\mathcal{B}(k+\frac{1}{2}))$ is isomorphic to either a vertex or an edge of Q^* in a way which is canonical on vertices. Suppose that v is a big link vertex. Then v is a vertex of the face tree of some face f of $\partial\mathcal{B}(k)$. We return to the setting of the next-to-last paragraph. We may assume that $v \in \{\varphi(v_1),\dots,\varphi(v_j)\}$. If $v = \varphi(v_1)$, then the face f is isomorphic to a face of Q^* in a way which is canonical on vertices and the dual of $\text{link}(v,\mathcal{B}(k+\frac{1}{2}))$ is isomorphic to the face of Q^* paired with the image of f in a way which is canonical on vertices. In other words, the root edge of v with its label and direction relative to v determines a face of Q^*, and the dual of $\text{link}(v,\mathcal{B}(k+\frac{1}{2}))$ is isomorphic to this face in a way which is canonical on vertices. In general, if $j > 1$ and $v = \varphi(v_i)$ for some integer $i \in \{2,\dots,j\}$, then the dual of $\text{link}(v,\mathcal{B}(k+\frac{1}{2}))$ is isomorphic to a subcomplex of the face of Q^* which is determined by the edge of t which joins v_i with v_{i-1} in a way which is canonical on vertices. Furthermore Figures 3 and 4 show that the dual of $\text{link}(v,\mathcal{B}(k+\frac{1}{2}))$ is isomorphic to the elbow of C which is the union of $e_i,\dots,e_j$ together with the edge of f which immediately precedes e_j. The edge of f which immediately precedes e_j is the bottom of the elbow and u_j is its joint. Because v is an original vertex of C_i, it follows that u_i is an original vertex of the dual of $\text{link}(v,\mathcal{B}(k+\frac{1}{2}))$. It easily follows in this case that the dual of $\text{link}(v,\mathcal{B}(k+\frac{1}{2}))$ is isomorphic to an elbow of Q^* in a way which is canonical on vertices. The edge e_i corresponds to the bottom of the elbow, and u_i corresponds to its joint. This completes the verification of induction hypothesis 2 for $\mathcal{B}(k+\frac{1}{2})$.

Verification of induction hypothesis 3. By induction hypothesis 3 for $\mathcal{B}(k)$ there exists a bijection between the vertices of $\mathcal{B}(k)\setminus\partial\mathcal{B}(k)$ and the 3-cells of $\mathcal{B}(k-\frac{1}{2})$. This and the definition of $\mathcal{B}(k+\frac{1}{2})$ shows that there exists a bijection between the vertices of $\mathcal{B}(k)$ and the 3-cells of $\mathcal{B}(k+\frac{1}{2})$. Similarly there exists a bijection between the edges of $\mathcal{B}(k)$ and the faces of $\mathcal{B}(k+\frac{1}{2})$ which are not in $\partial\mathcal{B}(k+\frac{1}{2})$. Induction hypothesis 3 for $\mathcal{B}(k)$ and the fact that the faces of $\partial\mathcal{B}(k)$ are dual to the root edges of $\mathcal{B}(k+\frac{1}{2})$ show that there exists a bijection between the faces of $\mathcal{B}(k)$ and the edges of $\mathcal{B}(k+\frac{1}{2})$ which are not in $\partial\mathcal{B}(k+\frac{1}{2})$. By induction hypothesis 3 for $\mathcal{B}(k)$ there exists a bijection between the vertices of $\mathcal{B}(k-\frac{1}{2})$ and the 3-cells of $\mathcal{B}(k)$. Since every vertex of $\mathcal{B}(k+\frac{1}{2})\setminus\mathcal{B}(k-\frac{1}{2})$ lies in $\partial\mathcal{B}(k+\frac{1}{2})$, it follows that there exists a bijection between the vertices of $\mathcal{B}(k+\frac{1}{2})\setminus\partial\mathcal{B}(k+\frac{1}{2})$ and the 3-cells of $\mathcal{B}(k)$. It is easy to see that these bijections determine a duality between the closed cells of $\mathcal{B}(k)$ and the closed cells of $\mathcal{B}(k+\frac{1}{2})$ which are not contained in $\partial\mathcal{B}(k+\frac{1}{2})$ and that this duality preserves labels and directions of edges and faces. This completes the verification of induction hypothesis 3 for $\mathcal{B}(k+\frac{1}{2})$.

Verification of induction hypothesis 4. Because there exists a duality between the closed cells of $\mathcal{B}(k)$ and the closed cells of $\mathcal{B}(k+\frac{1}{2})$ which are not contained in $\partial\mathcal{B}(k+\frac{1}{2})$, it easily follows that $\mathrm{star}(\mathcal{B}(k-\frac{1}{2})_\sigma,\mathcal{B}(k+\frac{1}{2})_\sigma)$ is isomorphic as a cell complex to $\mathcal{B}(k)_\sigma$. Induction hypothesis 4 states that $\mathcal{B}(k)$ is a topological ball, and so $\mathrm{star}(\mathcal{B}(k-\frac{1}{2})_\sigma,\mathcal{B}(k+\frac{1}{2})_\sigma)$ is a topological ball. It is clear that $\mathcal{B}(k+\frac{1}{2})_\sigma$ is the union of $\mathrm{star}(\mathcal{B}(k-\frac{1}{2})_\sigma,\mathcal{B}(k+\frac{1}{2})_\sigma)$ and the complexes of the form $\overline{\mathrm{costar}((L_v)_\sigma,(Q_v)_\sigma)}$ for $v\in V_k$. Statement 6 of Proposition 3.4 implies that $\overline{\mathrm{costar}((L_v)_\sigma,(Q_v)_\sigma)}$ is a topological ball for every $v\in V_k$. If we construct a finite sequence of complexes by beginning with $\mathrm{star}(\mathcal{B}(k-\frac{1}{2})_\sigma,\mathcal{B}(k+\frac{1}{2})_\sigma)$ and adjoining the complexes $\overline{\mathrm{costar}((L_v)_\sigma,(Q_v)_\sigma)}$ one at a time in any order by identifying appropriate faces, then statement 7 of Proposition 3.4 implies that at each step we are identifying two topological balls along a topological disk. It follows that $\mathcal{B}(k+\frac{1}{2})$ is a topological ball. This completes the verification of induction hypothesis 4 for $\mathcal{B}(k+\frac{1}{2})$.

Now that the induction hypotheses have been verified for $\mathcal{B}(k+\frac{1}{2})$, it is clear that $\bigcup_{k=0}^{\infty}\mathcal{B}(k+\frac{1}{2})$ is a connected and simply-connected covering space of M. It follows that we may identify $\bigcup_{k=0}^{\infty}\mathcal{B}(k+\frac{1}{2})$ with $\widetilde{M}$, and so we may identify $\mathcal{B}(k+\frac{1}{2})$ with $B(k+\frac{1}{2})$ for every nonnegative integer k. We may likewise identify $\mathcal{B}(k)$ with $B(k)$ for every nonnegative integer k.

Having identified $\mathcal{B}(k)$ with $B(k)$ for every nonnegative half integer k, statements 2, 5 and 6 of Theorem 5.1 are now clear.

We prove statement 1 of Theorem 5.1 in this paragraph. Induction hypothesis 1 easily implies that the lift from Q to $\widetilde{M}$ with image $B(\frac{1}{2})$ is injective. Statement 1 of Theorem 5.1 now follows easily using the action of the fundamental group of M on $\widetilde{M}$.

We prove statements 2 and 3 of Theorem 5.1 in this paragraph. Let x be a vertex of Γ_b, and suppose that the Γ_b-distance from x to $\mathcal{O}^*$ is $k>0$. Then x is a vertex of $\partial B(k)$. Let C be the 3-cell of $B(k+\frac{1}{2})$ which is dual to x. The dual in C of $\mathrm{link}(x,B(k))$ determines a subcomplex $L(x)$ of either Q or Q^* as in statement 3 of Theorem 5.1. Furthermore $L(x)$ determines the isomorphism types of the faces of $\partial B(k)$ which contain x, the position of x on these faces and how these faces fit together. Our discussion of cosubdivisions of faces now shows that $L(x)$ determines which vertices of C lie in $\partial B(k+\frac{1}{2})$ and the duals of their links. It is now not difficult to see that $L(x)$ determines the cone type of x. This proves statement 3 of Theorem 5.1. Since there are only finitely many such complexes $L(x)$, it follows that Γ_b has only finitely many cone types. This proves statement 2 of Theorem 5.1.

We prove statement 7 of Theorem 5.1 in this paragraph. The first assertion is clear. The second assertion follows easily from the verification of induction hypothesis 2. It remains to prove that f is unique if L contains more than two vertices. This is clear if $L = f$. So suppose that there exists a face f of $\partial B(k-1)$ and a 3-cell C of $B(k-1)$ with $L \subseteq f \subseteq C$ such that L is an elbow of C. It follows that L contains an original vertex v of C and the two edges of f which contain v. If v is a big link vertex of $\partial B(k-1)$, then using Figure 3 it is easy to see that f is the unique face of $B(k-1)$ which contains L. If v is a small link vertex of $\partial B(k-1)$, then the uniqueness of f follows easily from the fact that distinct faces of an ample faceted 3-ball have at most one edge in common. This proves statement 7 of Theorem 5.1. ☐

6. Gromov hyperbolicity

Here is the main theorem of this paper.

Theorem 6.1. *Let P be an ample faceted 3-ball, let ϵ be an orientation-reversing face-pairing on P, and let mul be a multiplier function for ϵ. Let M be the associated twisted face-pairing manifold. Then $G = \pi_1(M)$ is Gromov hyperbolic, and its space at infinity is homeomorphic to the 2-sphere.*

Proof. We begin the proof of Theorem 6.1 by proving in this paragraph that if G is Gromov hyperbolic, then its space at infinity is homeomorphic to the 2-sphere. Suppose that G is Gromov hyperbolic. Statement 2 of Theorem 5.1 easily implies that G is infinite and that the universal cover $\widetilde{M}$ of M is irreducible. As is well known, if $\widetilde{M}$ is irreducible, then M is irreducible. Now Theorem 4.1 of [2] implies that the space at infinity of G is homeomorphic to the 2-sphere. Thus what we must do to prove Theorem 6.1 is prove that G is Gromov hyperbolic.

We wish to prove that G is Gromov hyperbolic by applying Theorem 2.2 to G and the Cayley graph Γ_s of G with respect to its star generating set. As we saw in Section 4, there exists a unique G-equivariant isometry from the set of vertices of Γ_s to the bipartite graph Γ_b which maps the base vertex of Γ_s to the base vertex $\mathcal{O}^*$ of Γ_b. It easily follows that to verify our Gromov hyperbolicity criterion for Γ_s, it suffices to verify our Gromov hyperbolicity criterion for Γ_b. In other words, it suffices to prove that there exists a positive integer J with the following property. If there exists a geodesic edge path in Γ_b with vertices $v_0, \ldots, v_j$ in order such that $d(v_0, \mathcal{O}^*) = d(v_j, \mathcal{O}^*)$ and $d(v_i, \mathcal{O}^*) \geq d(v_0, \mathcal{O}^*)$ for every $i \in \{0, \ldots, j\}$, then $j < J$, where d is an edge path metric on Γ_b for which every edge has length $\frac{1}{2}$.

In this paragraph we show that Γ_b is almost convex in the sense of [3] in the strongest way possible. For this let k be a positive half integer, and let v_0, v_1, v_2 be distinct vertices of Γ_b such that $d(v_0, \mathcal{O}^*) = d(v_2, \mathcal{O}^*) = k$, $d(v_1, \mathcal{O}^*) = k + \frac{1}{2}$ and there exist edges joining v_0 with v_1 and v_1 with v_2. See Figure 5. The vertex v_0 is dual to a 3-cell C_{v_0} of $B(k+\frac{1}{2})$, and the vertex v_2 is dual to a 3-cell C_{v_2} of $B(k+\frac{1}{2})$. If v_1 is a big link vertex, then we let v_1' be the root vertex of v_1. Since $v_1 \in C_{v_0}$ and $v_1 \in C_{v_2}$, we have that $v_1' \in C_{v_0} \cap C_{v_2}$ if v_1 is a big link vertex. Suppose that v_1 is a small link vertex. Then C_{v_0} and C_{v_2} are the only 3-cells of $B(k+\frac{1}{2})$ which contain v_1. It follows that both C_{v_0} and C_{v_2} contain a face f which

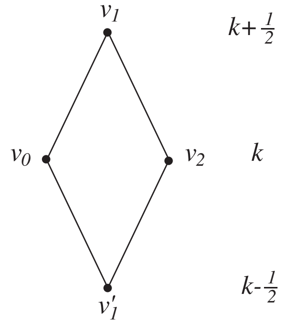

FIGURE 5. Proving that Γ_b is almost convex.

contains v_1. The face f contains two root edges of big link vertices of $\partial B(k+\frac{1}{2})$. We let v_1' be the root vertex of one of these root edges. It follows that $v_1' \in C_{v_0} \cap C_{v_2}$ if v_1 is a small link vertex. In either case we have that v_1' is a vertex of Γ_b with $d(v_1', \mathcal{O}^*) = k - \frac{1}{2}$ such that there exist edges joining v_0 with v_1' and v_1' with v_2. This is the **almost convex property** mentioned at the beginning of this paragraph. Although we will not use the following fact, we note that it easily follows that this same almost convex property holds for Γ_s.

In this paragraph we define the notion of a pull-in. For this, let k be a positive half integer. Let γ be a geodesic edge path in Γ_b with vertices $v_0, \ldots, v_j$ in order for some even positive integer j such that $d(v_i, \mathcal{O}^*) = k$ for every even integer $i \in \{0, \ldots, j\}$ and $d(v_i, \mathcal{O}^*) = k + \frac{1}{2}$ for every odd integer $i \in \{0, \ldots, j\}$. For every odd integer $i \in \{0, \ldots, j\}$ we choose as in the previous paragraph a vertex v_i' of Γ_b with $d(v_i', \mathcal{O}^*) = k - \frac{1}{2}$ such that v_i' is joined by edges of Γ_b to both v_{i-1} and v_{i+1}. Now we use these vertices to construct an edge path γ' in Γ_b with vertices $v_1', v_2, v_3', v_4, \ldots, v_{j-1}'$ in order such that $d(v_i', \mathcal{O}^*) = k - \frac{1}{2}$ for every odd integer $i \in \{1, \ldots, j-1\}$ and $d(v_i, \mathcal{O}^*) = k$ for every even integer $i \in \{1, \ldots, j-1\}$. We call γ' a **pull-in** of γ, and we say that γ **pulls in** to γ'. It is clear that γ' is a geodesic edge path.

Now let γ be a geodesic edge path in Γ_b with vertices $v_0, \ldots, v_j$ in order such that $d(v_0, \mathcal{O}^*) = d(v_j, \mathcal{O}^*) = k$ for some positive half integer k and $d(v_i, \mathcal{O}^*) \geq k$ for every $i \in \{0, \ldots, j\}$. Since we seek a bound on j, we may repeatedly apply the almost convex property of Γ_b to assume that $d(v_i, \mathcal{O}^*) = k$ for every even integer $i \in \{0, \ldots, j\}$ and $d(v_i, \mathcal{O}^*) = k + \frac{1}{2}$ for every odd integer $i \in \{0, \ldots, j\}$. Our strategy to prove that there exists a bound on j is to construct a pull-in of γ and then a pull-in of that and continue to pull in until it becomes clear that because γ is geodesic, there is a bound on the number of such pull-ins. Clearly, a bound on the number of such pull-ins gives a bound on j. To prepare to carry out this strategy, we next define and discuss tame and wild edges and after that we discuss big link vertices.

Let Q denote the twisted face-pairing subdivision of P as usual. We say that an edge of Q or Q^* is **wild** if it is an original edge. In other words, an edge of P, respectively P^*, which does not properly subdivide in Q, respectively Q^*, becomes a wild edge of Q, respectively Q^*. An edge of Q or Q^* which is not wild is **tame**. Now let C be a 3-cell of either $\widetilde{M}$ or $\widetilde{M}^*$. Statement 1 of Theorem 5.1 implies that C is canonically isomorphic to either Q or Q^* in an orientation-preserving way. We call an edge e of C **wild** or **tame** relative to C according to whether its image in Q or Q^* is wild or tame. Note that the property of e being wild or tame really is a property relative to C; e might be tame relative to C and wild relative to another 3-cell which contains it. Note that if an edge e of C is wild relative to C, then the two faces of C which contain e are paired by the face-pairing of C induced by the canonical isomorphism between C and either Q or Q^*. Lemma 6.2 follows from this because distinct faces of C meet in at most one original edge.

Lemma 6.2. *If C is a 3-cell of either $\widetilde{M}$ or $\widetilde{M}^*$ and f is a face of C, then f contains at most one edge which is wild relative to C.*

We next prove Lemma 6.3, the "dual" of Lemma 6.2.

Lemma 6.3. *If e is an edge of either $\widetilde{M}$ or $\widetilde{M}^*$, then there exists at most one 3-cell C containing e such that e is wild relative to C.*

Proof. Without loss of generality we assume that e is an edge of $\widetilde{M}$. Using statement 1 of Theorem 5.1, we see that the 3-cells of $\widetilde{M}$ which contain e bijectively correspond to the edges in an edge cycle of the twisted face-pairing on Q. Suppose that C is a 3-cell of $\widetilde{M}$ containing e such that e is wild relative to C. By means of the canonical isomorphism between C and Q, the edge e corresponds to an edge e' of Q. The edge e' is an edge of the abovementioned edge cycle. Because e is wild relative to C, Theorem 4.2 of [8] shows that the two faces of Q which contain e' have the same label as e', which has the same label as e. Since distinct faces of Q meet in at most an original edge, it follows that e determines e'. This implies that if e is wild relative to some 3-cell C of $\widetilde{M}$, then C is unique. This proves Lemma 6.3. ☐

This completes our discussion of tame and wild edges. We next discuss big link vertices.

Again let C be a 3-cell of either $\widetilde{M}$ or $\widetilde{M}^*$. Suppose that C is dual to a vertex of $\partial B(k)$ for some positive half integer k. Let $L = C \cap \partial B(k - \frac{1}{2})$. Statement 7 of Theorem 5.1 implies that L is either a vertex or an edge or a face or an elbow of C. The discussion of root vertices in Section 5 near Lemma 5.5 easily proves Lemma 6.4.

Lemma 6.4. *Let C be a 3-cell of $B(k + \frac{1}{2})$ which is dual to a vertex of $\partial B(k)$ for some positive half integer k, and let $L = C \cap \partial B(k - \frac{1}{2})$. Then the vertices of L which are root vertices of big link vertices of $C \cap \partial B(k + \frac{1}{2})$ are precisely the original vertices of C contained in L together with the vertices of L which are contained in at most one edge of L. These are the vertices of L which are contained in an edge of C not in L.*

We maintain the setting of Lemma 6.4. Let w be a vertex of L which is the root vertex of some big link vertex of $C \cap \partial B(k + \frac{1}{2})$. We next diagrammatically describe all the big link vertices of $C \cap \partial B(k + \frac{1}{2})$ which have w as root vertex. Every big link vertex of $C \cap \partial B(k + \frac{1}{2})$ which has w as root vertex is gotten as follows. Let r be an edge of C not in L which contains w. Let g and h be the faces of C which contain r as in Figure 6. The following assertions hold concerning Figure 6: 1) the vertices a and b are the vertices of an edge of g, 2) the vertex a may or may not be an original vertex of C, 3) the vertices b and c are the original vertices of an original edge of h and 4) w and b are contained in $g \cap h$, which is an original edge of C. The circular arcs drawn in some of the corners in Figure 6 indicate corners of faces of C, whereas corners in Figure 6 without circular arcs might be unions of several corners of faces of C. In particular, the vertices a,b,c need not be contained in some face of C. The vertices in Figure 6 drawn as circles are all the big link vertices of $C \cap \partial B(k + \frac{1}{2})$ which have r as root edge. There are four types of diagrams, the four types being determined by whether or not the vertices of r are original vertices of C. In part a) of Figure 6 w is an original vertex of C and the other vertex of r is not an original vertex of C; in part b) of Figure 6 w is an original

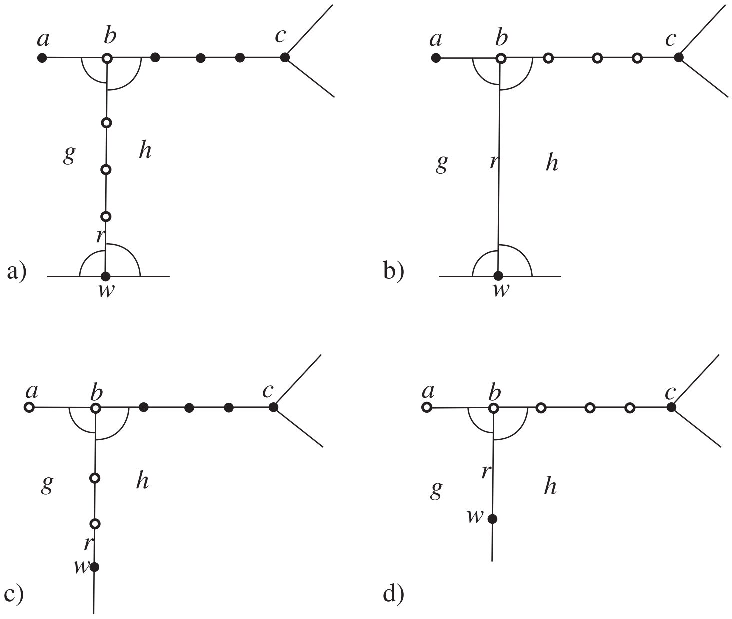

FIGURE 6. The big link vertices of $C \cap \partial B(k + \frac{1}{2})$ having root edge r.

vertex of C and the other vertex of r is also an original vertex of C (r is wild relative to C); in part c) of Figure 6 w is not an original vertex of C and the other vertex of r is also not an original vertex of C; in part d) of Figure 6 w is not an original vertex of C and the other vertex of r is an original vertex of C. We emphasize that this set of diagrams is complete in the sense that their "reflections" do not occur. For example, when r is wild relative to C the big link vertices of $C \cap \partial B(k + \frac{1}{2})$ having r as root edge all lie in the direction of c from b, not in the direction of a. The fact that this set of diagrams is complete follows easily from Figure 3 and the discussion of it in Section 5. This completes our identification of the big link vertices of $\partial B(k + \frac{1}{2})$ which have a given vertex of $\partial B(k - \frac{1}{2})$ as root vertex.

We conclude our discussion of big link vertices with Lemma 6.5.

Lemma 6.5. *Let v be a big link vertex of $\partial B(k + \frac{1}{2})$ for some positive half integer k. Let C be a 3-cell of $B(k + \frac{1}{2})$ which contains v, and let $L = C \cap \partial B(k - \frac{1}{2})$. Let r be the root edge of v, and let g and h be the faces of star$(L, \partial C)$ which contain r as in Figure 6. Let f be a face of star$(L, \partial C)$ which contains v. Then either $f = g$ or $f = h$.*

Proof. It is clear that $f = g$ or $f = h$ if v lies in the interior of $g \cap h$. Statement 1 of Proposition 3.2 implies that g and h are the only faces of star$(L, \partial C)$ which contain b. Statement 1 of Proposition 3.2 also shows that h is the only face of star$(L, \partial C)$ which contains the vertices in the interior of the original edge with endpoints b and c. Finally, suppose that a is a big link vertex and that f is a face of star$(L, \partial C)$ other than g which contains a. We are in the situation of either part c) or part d) of Figure 6. Hence the segment of ∂g which joins a and w in Figure 6 is contained in a twisted original edge t of C. Statement 1 of Proposition 3.2 implies that $f \cap g$ is an original edge e of C which joins a and L. Now Proposition 3.3 applied to e, L and t gives a contradiction. This proves Lemma 6.5. □

We return to the situation in which γ is a geodesic edge path in Γ_b for which there exists a positive half integer k such that the first and last vertices of γ have distance k from $\mathcal{O}^*$ and

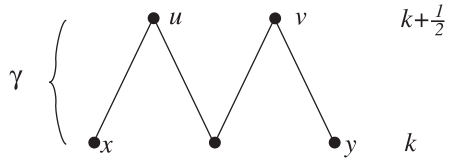

FIGURE 7. Part of γ.

every vertex of γ has distance either k or $k+\frac{1}{2}$ from $\mathcal{O}^*$. We fix γ for the rest of the proof of Theorem 6.1. To prove Theorem 6.1, we let γ_1 be a pull-in of γ, we let γ_2 be a pull-in of γ_1, we let γ_3 be a pull-in of γ_2 and so on as far as possible. Lemmas 6.6 through 6.19 show that these successive pull-ins of γ satisfy stronger and stronger properties, and finally it becomes clear that there exists a bound on the length of γ which depends only on the given ample faceted 3-ball P. This suffices to prove Theorem 6.1. In what follows we make assertions about the pull-ins $\gamma_1,\dots,\gamma_7$. Because we seek a bound on the length of γ, we may and do assume that $\gamma_1,\dots,\gamma_7$ exist.

We refer to the vertices v of γ with $d(v,\mathcal{O}^*)=k+\frac{1}{2}$ as **outer** vertices of γ, and we refer to the vertices v of γ with $d(v,\mathcal{O}^*)=k$ as **inner** vertices of γ. If u,v,w are consecutive vertices of γ such that u and w are outer vertices of γ, then we say that u and w are **adjacent outer vertices** of γ. Clearly, the notions of outer vertex, inner vertex and adjacent outer vertices are meaningful for successive pull-ins of γ.

Having finished the preliminaries, the proof of Theorem 6.1 now begins in earnest with the statement and proof of Lemma 6.6. Much of the proof of Theorem 6.1 involves analyzing the link types of the outer vertices of γ and its successive pull-ins. Lemma 6.6 handles the simplest link type.

Lemma 6.6. *If v is an outer vertex of γ, then the dual of the $B(k+\frac{1}{2})$-link of v is not a vertex.*

Proof. Let v be an outer vertex of γ. Let x and y be the inner vertices of γ adjacent to v. Then both x and y are dual to 3-cells of $B(k+\frac{1}{2})$ which contain v. Thus the dual of $\text{link}(v,B(k+\frac{1}{2}))$ contains at least two vertices. This proves Lemma 6.6. $\square$

Several following proofs by contradiction proceed by showing, given adjacent outer vertices u and v of γ, that every 3-cell of $B(k+\frac{1}{2})$ which contains u meets every 3-cell of $B(k+\frac{1}{2})$ which contains v. Lemma 6.7 shows that this is impossible.

Lemma 6.7. *Let u and v be adjacent outer vertices of γ. Then there exist disjoint 3-cells C and D of $B(k+\frac{1}{2})$ such that $u\in C$ and $v\in D$.*

Proof. Let x be the inner vertex of γ which is adjacent to u but not to v, and let y be the inner vertex of γ which is adjacent to v but not to u as in Figure 7. Let C be the 3-cell of $B(k+\frac{1}{2})$ dual to x, and let D be the 3-cell of $B(k+\frac{1}{2})$ dual to y. Then $u\in C$ and $v\in D$, but $C\cap D=\emptyset$ because $d(x,y)>1$. This proves Lemma 6.7. $\square$

We next consider Lemma 6.8, which refines statement 7 of Theorem 5.1.

Lemma 6.8. *Let u and v be adjacent outer vertices of γ_2. See Figure 8. Let C be the 3-cell of $B(k-\frac{1}{2})$ which is dual to the inner vertex of γ_2 adjacent to both u and v. Let D be the 3-cell of $B(k+\frac{1}{2})$ dual to the outer vertex of γ_1 adjacent to both u and v. Let*

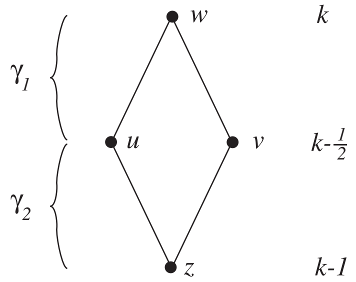

FIGURE 8. Parts of γ_1 and γ_2.

$K = D \cap B(k - \frac{1}{2})$. Then there exists a face f of $C \cap \partial B(k - \frac{1}{2})$ which contains K and hence both u and v. Moreover K is either an edge or a face or an elbow of C such that the following holds. Each of the vertices u and v is a vertex of K which either 1) is contained in at most one edge of K, 2) immediately precedes (relative to f and C) a vertex of K which is contained in at most one edge of K or 3) immediately precedes (relative to f and C) an original vertex of C in K. Furthermore, f is the only face of $B(k - \frac{1}{2})$ which contains K unless K is an edge.

Proof. Let z be the inner vertex of γ_2 adjacent to u and v, and let w be the outer vertex of γ_1 adjacent to both u and v. First suppose that w is a small link vertex. The complex K is the dual in $B(k - \frac{1}{2})$ of $\mathrm{link}(w, B(k))$. Hence K contains at most two vertices. Since u and v are contained in K, it follows from statement 5 of Theorem 5.1 that K is an edge with vertices u and v. Let C_u be the 3-cell of $B(k)$ dual to u, and let C_v be the 3-cell of $B(k)$ dual to v. Then $C_u \cap C_v$ is the face g of $B(k)$ dual to K. The vertex z is defined to be the root vertex of one of the two root edges in g. These two root edges are dual to the two faces of $\partial B(k - \frac{1}{2})$ which contain K, and C is the 3-cell which contains one of these faces f. This proves Lemma 6.8 if w is a small link vertex.

Now suppose that w is a big link vertex. In this case z is defined to be the root vertex of w. It follows that w is a vertex of the face tree of a face f of $C \cap \partial B(k - \frac{1}{2})$. See Figure 3. It follows that $K \subseteq f$. Since w is a big link vertex, K contains at least three vertices, and so statement 7 of Theorem 5.1 implies that K is either a face or an elbow of C and that K is either a face or an elbow of D. By the definition of pull-in, u and v are root vertices of big link vertices of $D \cap \partial B(k + \frac{1}{2})$. Hence Lemma 6.4 implies that u and v are either original vertices of D or vertices of K which are contained in at most one edge of K. It easily follows that each of u and v either 1) is contained in at most one edge of K, 2) immediately precedes (relative to f and C) a vertex of K which is contained in at most one edge of K (if K is an elbow of C) or 3) immediately precedes (relative to f and C) an original vertex of C in K (if $K = f$). The uniqueness of f follows from statement 7 of Theorem 5.1.

This proves Lemma 6.8. □

If v is an outer vertex of γ_2 and u is an inner vertex of γ_2 adjacent to v, then v is contained in the 3-cell dual to u. Lemma 6.9 gives a more precise statement.

Lemma 6.9. *Let v be an outer vertex of γ_2. Let C be the 3-cell of $B(k - \frac{1}{2})$ dual to an inner vertex of γ_2 adjacent to v, and let $L = C \cap B(k - \frac{3}{2})$. Then $v \in \partial star(L, \partial C)$.*

Proof. By definition $v \in C$. Lemma 6.6 applied to γ_2 instead of γ easily implies that C is not the only 3-cell of $B(k - \frac{1}{2})$ which contains v. With this Lemma 5.3 easily implies that

$v \in \operatorname{star}(L, \partial C)$. Lemma 6.8 implies that there exists a face of $C \cap \partial B(k - \frac{1}{2})$ which contains v. It follows that $v \in \partial\operatorname{star}(L, \partial C)$. This proves Lemma 6.9. □

Lemma 6.9 deals with one outer vertex of γ_2. Lemma 6.10 sharpens the conclusion of Lemma 6.9 for two adjacent outer vertices of γ_2.

Lemma 6.10. *Let u and v be adjacent outer vertices of γ_2. Let C be the 3-cell of $B(k - \frac{1}{2})$ which is dual to the inner vertex of γ_2 adjacent to both u and v, and let $L = C \cap B(k - \frac{3}{2})$. Let f be as in Lemma 6.8. The following statements hold.*

1. *If u and v are contained in a face g of $\operatorname{star}(L, \partial C)$, then $f \cap g$ is an original edge of $\partial\operatorname{star}(L, \partial C)$ which contains both u and v and which meets a face of $\operatorname{star}(L, \partial C)$ other than g.*
2. *If there exist distinct faces g and h of $\operatorname{star}(L, \partial C)$ which have a vertex in common such that $u \in g \setminus h$ and $v \in h \setminus g$, then $f \cap g$ and $f \cap h$ are original edges of $\partial\operatorname{star}(L, \partial C)$ such that $u \in f \cap g$, $v \in f \cap h$ and $f \cap g$ meets $f \cap h$.*

Proof. We prove statement 1 of Lemma 6.10 in this paragraph. Suppose that u and v are contained in a face g of $\operatorname{star}(L, \partial C)$. Since f and g contain both u and v, ampleness condition 1 easily implies that $f \cap g$ is an original edge of C. We prove that the original edge $f \cap g$ of $\partial\operatorname{star}(L, \partial C)$ meets a face of $\operatorname{star}(L, \partial C)$ other than g by contradiction. Suppose that $f \cap g$ does not meet a face of $\operatorname{star}(L, \partial C)$ other than g. We consider Lemma 6.5 and Figure 6. It follows that $f \cap g$ contains no big link vertices of the types which appear in parts a) or b) of Figure 6. In fact $f \cap g$ contains at most one big link vertex, that vertex being of the type labeled a in parts c) and d) of Figure 6. But if at most one of u and v is a big link vertex and the vertices u and v are contained in a face g of $\operatorname{star}(L, \partial C)$, then every 3-cell of $B(k - \frac{1}{2})$ which contains u meets every 3-cell of $B(k - \frac{1}{2})$ which contains v. This contradicts Lemma 6.7, proving statement 1 of Lemma 6.10.

To prove statement 2 of Lemma 6.10, suppose that g and h are distinct faces of $\operatorname{star}(L, \partial C)$ which have a vertex in common such that $u \in g \setminus h$ and $v \in h \setminus g$. The second ampleness condition implies that f, g and h have exactly one vertex in common. Since this common vertex is neither u nor v, it follows that $f \cap g$ and $f \cap h$ both contain at least two vertices. This easily implies that $f \cap g$ and $f \cap h$ are adjacent original edges of $\partial\operatorname{star}(L, \partial C)$. This proves statement 2 of Lemma 6.10.

The proof of Lemma 6.10 is complete. □

Lemma 6.11 improves Lemma 6.10 under further hypotheses.

Lemma 6.11. *Let u and v be adjacent outer vertices of γ_2. Let C be the 3-cell of $B(k - \frac{1}{2})$ which is dual to the inner vertex of γ_2 adjacent to both u and v, and let $L = C \cap B(k - \frac{3}{2})$. Let g and h be distinct faces of $\operatorname{star}(L, \partial C)$ as in Figure 9 such that $g \cap h$ contains an original vertex x of C not in L, $g \cap h$ contains an original vertex y of L and x immediately precedes $g \cap h$ (relative to g and C). Let e_g be the original edge of $g \cap \partial\operatorname{star}(L, \partial C)$ which contains x, and let e_h be the original edge of $h \cap \partial\operatorname{star}(L, \partial C)$ which contains x. Suppose that L is not a face. Then the following statements hold.*

1. *The vertices u and v are not both original vertices of $e_g \cup e_h$.*
2. *If $g \cap h$ does not contain an edge of L and neither g nor h contains an edge of L which is wild relative to C, then u and v are not both contained in $e_g \cup e_h$.*
3. *If $g \cap h$ contains an edge of L and $h \cap L \subseteq g \cap h$, then u and v are not both contained in $e_g \cup e_h$.*

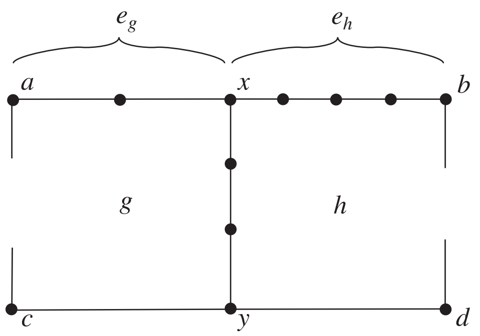

FIGURE 9. The faces g and h.

Proof. Let a be the original vertex of e_g other than x, and let b be the original vertex of e_h other than x. Let c be the original vertex of $g \setminus h$ such that c and y are the endpoints of an original edge of g, and let d be the original vertex of $h \setminus g$ such that d and y are the endpoints of an original edge of h. Let K and f be as in Lemma 6.8.

We prove statement 1 of Lemma 6.11 by contradiction: suppose that u and v are both original vertices of $e_g \cup e_h$. We have that either $\{u,v\} = \{a,b\}$ or $\{u,v\} = \{a,x\}$ or $\{u,v\} = \{b,x\}$.

Suppose that $\{u,v\} = \{b,x\}$. It follows that $f \cap h = e_h$ and that $e_h \subseteq K$. Using Lemma 6.8 we see that since K contains the original edge e_h, either e_h is wild relative to C or $K = f$. Lemma 6.8 furthermore shows that if $K = f$, then b immediately precedes (relative to f and C) an original vertex of C in K, and so e_h is wild relative to C. So we always have that e_h is wild relative to C. Now Lemma 6.2 implies that no other edge of h is wild relative to C. Because L is not a face and the original edge with endpoints d and y is not a wild edge of h, it follows that $d \notin L$. Using Lemma 6.5 and Figure 6 together with the fact that b and d are not the endpoints of a wild edge of h, we see that b is a small link vertex. It is clear that x is a big link vertex with root vertex in $g \cap h$. Hence the two 3-cells of $B(k-\frac{1}{2})$ which contain b both contain the root vertex of x. Since every 3-cell of $B(k-\frac{1}{2})$ which contains a big link vertex of $\partial B(k-\frac{1}{2})$ contains the root vertex of that big link vertex, it follows that every 3-cell of $B(k-\frac{1}{2})$ which contains u meets every 3-cell of $B(k-\frac{1}{2})$ which contains v, in contradiction to Lemma 6.7. This proves that $\{u,v\} \neq \{b,x\}$.

A similar argument proves that $\{u,v\} \neq \{a,x\}$. Finally suppose that $\{u,v\} = \{a,b\}$. Then $a \in K \subseteq f$ and $b \in K \subseteq f$. Ampleness condition 2 applied to f, g and h easily shows that $x \in f$. Hence $e_g \subseteq f$ and $e_h \subseteq f$. It easily follows that $K = f$. Now Lemma 6.8 implies that a immediately precedes (relative to f and C) an original vertex of C. Since a is an original vertex of C, it follows that a immediately precedes (relative to f and C) a wild edge of C. Likewise b immediately precedes (relative to f and C) a wild edge of C. This contradicts Lemma 6.2.

This proves statement 1 of Lemma 6.11.

We next prove statement 2 of Lemma 6.11 by contradiction: suppose that $g \cap h$ does not contain an edge of L, that neither g nor h contains an edge of L which is wild relative to C and that u and v are both contained in $e_g \cup e_h$.

Because L is not a face and g does not contain an edge of L which is wild relative to C, it easily follows that $c \notin L$. Now Lemma 6.5 and Figure 6 easily imply that every big link vertex of e_g (x is the only one) has root vertex y. If every big link vertex of e_h also has root vertex y, then every 3-cell of $B(k-\frac{1}{2})$ which contains u meets every 3-cell of $B(k-\frac{1}{2})$ which contains v, in contradiction to Lemma 6.7. Hence b is a big link vertex. Just as $c \notin L$, we also have $d \notin L$. Lemma 6.5 and Figure 6 imply that b is the only big link vertex of $e_g \cup e_h$ whose root vertex is not y and that b and d are the endpoints of an edge which is wild relative to C. Hence $b \in \{u, v\}$, and without loss of generality we assume that $b = v$. Lemma 6.2 implies that $g \cap h$ is not an edge which is wild relative to C. Hence x is the only big link vertex of e_h with root vertex y. It follows that $u \in e_g$ for otherwise u is a small link vertex, and every 3-cell of $B(k-\frac{1}{2})$ which contains u contains the root vertex of v and hence meets every 3-cell of $B(k-\frac{1}{2})$ which contains v. Since $u \in e_g$, $v = b$, $u \in K \subseteq f$ and $v \in K \subseteq f$, ampleness condition 2 applied to f, g and h easily implies that $x \in f$. This easily implies that $e_h \subseteq f$ and then that $e_h \subseteq K$. Since e_h is not an edge which is wild relative to C, it follows that $K = f$. Finally, Lemma 6.8 implies that b immediately precedes (relative to f and C) an original vertex of C. This means that e_h is an edge which is wild relative to C, a contradiction as before.

This proves statement 2 of Lemma 6.11.

We finally prove statement 3 of Lemma 6.11 by contradiction: suppose that $g \cap h$ contains an edge of L, that $h \cap L \subseteq g \cap h$ and that u and v are both contained in $e_g \cup e_h$.

Figure 6 easily shows that x is a big link vertex. Because $h \cap L \subseteq g \cap h$, Lemma 6.5 and Figure 6 easily imply that every big link vertex of e_h has the same root vertex as x. Because L contains an edge in $g \cap h$ and $L \neq g$, it is easy to see that $c \notin L$. Now Lemma 6.5 and Figure 6 easily imply that every big link vertex of e_g has the same root vertex as x. So every big link vertex of $e_g \cup e_h$ has the same root vertex as x, and so every 3-cell of $B(k-\frac{1}{2})$ which contains u meets every 3-cell of $B(k-\frac{1}{2})$ which contains v, which is a contradiction as before.

This proves statement 3 of Lemma 6.11. $\square$

Lemma 6.6 might be viewed as saying that the link types of the outer vertices of γ are not as small as possible. Lemma 6.12 states that the link types of almost all the outer vertices of γ_3 are not as big as possible.

Lemma 6.12. *If v is an outer vertex of γ_3 whose $B(k-1)$-link is dual to a face, then the distance from v to one of the endpoints of γ_3 is less than 2.*

Proof. We begin the proof of Lemma 6.12 by assuming that v is an outer vertex of γ_2, not γ_3, whose $B(k-\frac{1}{2})$-link is dual to a face. We suppose that the distance from v to the endpoints of γ_2 is at least 2. Let u and w be the inner vertices of γ_2 adjacent to v as in Figure 10. Let C_u be the 3-cell of $B(k-\frac{1}{2})$ dual to u, and let C_w be the 3-cell of $B(k-\frac{1}{2})$ dual to w. Then $v \in C_u$ and $v \in C_w$. Because v is a big link vertex, it has a root edge r. It follows that both C_u and C_w contain r. Lemma 6.3 implies that r is not wild relative to both C_u and C_w. So we assume without loss of generality that r is tame relative to C_u. Let L be the dual of $\text{link}(u, B(k-1))$ in C_u. Let u' be the outer vertex of γ_2 such that $u' \neq v$ and u' is adjacent to u, and let $C_{u'}$ be the 3-cell of $B(k)$ dual to u'. Let v' be the outer vertex of γ_1 which is adjacent to both u' and v. See Figure 10. In the following paragraphs we prove that 1) the dual L of $\text{link}(u, B(k-1))$ in C_u contains an edge which is wild relative to C_u, 2) the dual of $\text{link}(v', B(k))$ in $\partial B(k-\frac{1}{2})$ is a face and 3) the dual of $\text{link}(u', B(k-\frac{1}{2}))$ in $C_{u'}$ does not contain an edge which is wild relative to $C_{u'}$.

In this paragraph we introduce some more notation. We have that $v \in \partial\text{star}(L, \partial C_u)$ by Lemma 6.9. Using Figure 3 it is easy to see that $v \in r$. Since r is not contained in

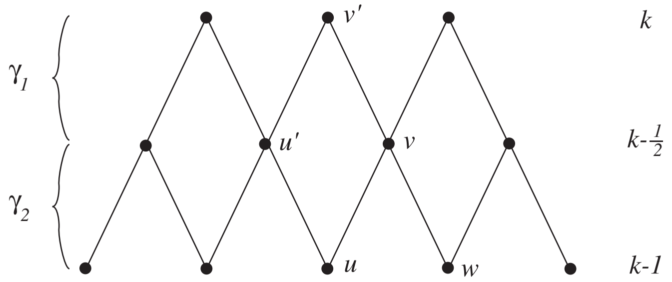

FIGURE 10. Parts of γ_1 and γ_2.

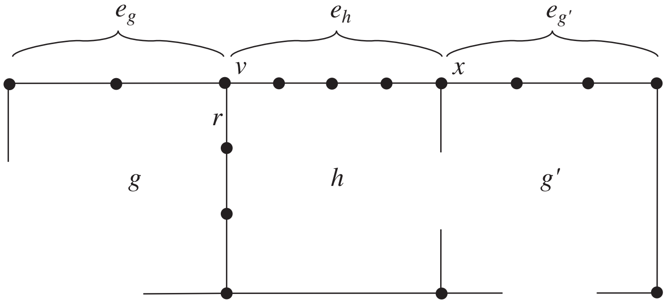

FIGURE 11. Part of the 3-cell C_u.

$\partial\text{star}(L, \partial C_u)$, it follows that v is an original vertex of C_u. Let g and h be the faces of C_u which contain r as in Figure 11 so that v immediately precedes r (relative to g and C_u). Because r is tame relative to C_u, statement 7 of Theorem 5.1 easily implies that either $L \subseteq g$ or $L \subseteq h$. Because $v \notin L$, it follows that L is not a face. It is furthermore true that h meets every face of $\text{star}(L, \partial C_u)$.

In this paragraph we prove that $u' \notin g \cup h$ by contradiction: suppose that $u' \in g \cup h$. Statement 1 of Lemma 6.10 implies that u' is in either the original edge e_g of $g \cap \partial\text{star}(L, \partial C_u)$ which contains v or the original edge e_h of $h \cap \partial\text{star}(L, C_u)$ which contains v. Statement 1 of Lemma 6.11 implies that u' is not an original vertex of $e_g \cup e_h$. With this it is easy to see using Lemma 6.5 and Figure 6 that either u' is a small link vertex or u' is a big link vertex whose root vertex equals the root vertex of v. It follows that every 3-cell of $B(k - \frac{1}{2})$ which contains u' meets every 3-cell of $B(k - \frac{1}{2})$ which contains v, in contradiction to Lemma 6.7. Thus $u' \notin g \cup h$.

In this paragraph we prove all three statements at the end of the first paragraph of the proof of Lemma 6.12. Lemma 6.8 applied to u' and v implies that there exists a face f of $C_u \cap \partial B(k - \frac{1}{2})$ which contains u' and v. According to the previous paragraph, $u' \notin g \cup h$. Lemma 6.9 implies that $u' \in \text{star}(L, C_u)$. Hence there exists a face g' of $\text{star}(L, C_u)$ which contains u'. Since $g' \neq g$ and $g' \neq h$, Lemma 6.5 implies that $v \notin g'$. Since every face of $\text{star}(L, C_u)$ meets h, we have that g' meets h, h meets f and f meets g'. Ampleness condition 2 implies that $f \cap g' \cap h$ is a vertex. Since $v \notin g'$, it follows that f and h have at least two vertices in common, and so they have exactly an original edge in common. This original edge must be e_h. It follows that if x is the original vertex of e_h other than v, then $f \cap g' \cap h = \{x\}$. Since $x \notin L$, statement 2 of Proposition 3.4 implies that $h \cap g'$ is an original edge of C_u with one endpoint in L. Since h is not a triangle, it follows that h contains two original vertices of C_u which lie in L. Since L is not a face, it follows that $L \subseteq h$ and that L contains an edge which is wild relative to C_u. This proves statement 1 of the first paragraph of the proof of Lemma 6.12. Let $e_{g'}$ be the original edge of g' other than $h \cap g'$ which contains x. Since x and u' are distinct vertices of $f \cap g'$, it follows that $f \cap g' = e_{g'}$. Let K be as in Lemma 6.8. It easily follows that e_h and at least the last edge (relative to f and C_u) of $e_{g'}$ are contained in K. It easily follows that $K = f$. This proves statement 2 of the first paragraph of the proof of Lemma 6.12. Since $L \subseteq h$ and L contains an edge which is wild relative to C_u, Lemma 6.2 implies that $h \cap g'$ is not an edge which is wild relative to C_u. It is clear that $g' \cap L$ is a vertex. Hence Lemma 6.5 and Figure 6 easily show that u' is a small link vertex. Hence the dual of $\text{link}(u', B(k - \frac{1}{2}))$ in $C_{u'}$ is an edge. Statement 1 of Lemma 6.11 implies that u' is not an original vertex of C_u. This and the fact that $u' \in f \cap g'$ imply that f and g' are the only faces of C_u which contain u'. Hence the vertex u has valence 2 in $C_{u'}$; one edge containing u in $C_{u'}$ is dual to f and one edge containing u in $C_{u'}$ is dual to g'. This means that u is not an original vertex of $C_{u'}$, and so the dual of $\text{link}(u', B(k - \frac{1}{2}))$ in $C_{u'}$ is an edge which is not wild relative to $C_{u'}$. This proves statement 3 of the first paragraph of the proof of Lemma 6.12. Thus all three statements at the end of the first paragraph of the proof of Lemma 6.12 are true.

In this paragraph we conclude the proof of Lemma 6.12. The proof is by contradiction. We assume that v is an outer vertex of γ_3 whose $B(k-1)$-link is dual to a face and that the distance from v to the endpoints of γ_3 is at least 2. Since γ_3 is the second pull-in of γ_1, we may proceed as in the first paragraph of the proof of Lemma 6.12. The three statements at the end of the first paragraph of the proof of Lemma 6.12 give the following in the present situation 1) the dual of $\text{link}(u, B(k - \frac{3}{2}))$ in C_u contains an edge which is wild relative to C_u, 2) the dual of $\text{link}(v', B(k - \frac{1}{2}))$ in $\partial B(k-1)$ is a face and 3) the dual of $\text{link}(u', B(k-1))$ in $C_{u'}$ does not contain an edge which is wild relative to $C_{u'}$. Let C_v be the 3-cell of $B(k - \frac{1}{2})$ dual to v. Then $v' \in C_v$. Since $\text{link}(v', B(k - \frac{1}{2}))$ is a face, it follows that v' is contained in its root edge. This, Lemma 6.9 applied to v' and γ_2 and the fact that $\text{link}(v, B(k-1))$ is a face easily imply that the root edge of v' is wild relative to C_v. Hence Lemma 6.3 implies that the root edge of v' is tame relative to $C_{u'}$. Since the distance from v to the endpoints of γ_3 is at least 2, the distance from v' to the endpoints of γ_2 is at least 2. Thus the three statements at the end of the first paragraph of the proof of Lemma 6.12 apply to u' and v' as well as u and v. In particular, statement 1 applied to u' and v' asserts that the dual of $\text{link}(u', B(k-1))$ in $C_{u'}$ contains an edge which is wild relative to $C_{u'}$. This contradicts statement 3 applied to u and v.

This proves Lemma 6.12. □

Now that we have Lemma 6.12, we can extend Lemma 6.11. Lemma 6.13 might be viewed as statement 4 of Lemma 6.11.

Lemma 6.13. *Let u and v be adjacent outer vertices of γ_2. Let C be the 3-cell of $B(k - \frac{1}{2})$ which is dual to the inner vertex of γ_2 adjacent to both u and v, and let $L = C \cap B(k - \frac{3}{2})$. Let g and h be distinct faces of $star(L, \partial C)$ as in Figure 9 such that $g \cap h$ contains an original vertex x of C not in L, $g \cap h$ contains an original vertex y of L and x immediately precedes $g \cap h$ (relative to g and C). Let e_g be the original edge of $g \cap \partial star(L, \partial C)$ which contains x, and let e_h be the original edge of $h \cap \partial star(L, \partial C)$ which contains x. Suppose that $g \cap h$ contains an edge of L and that u and v are both contained in $e_g \cup e_h$. Then the distance from either u or v to one of the endpoints of γ_2 is less than 2.*

Proof. Let a be the original vertex of e_g other than x, and let b be the original vertex of e_h other than x. Let c be the original vertex of $g \setminus h$ such that c and y are the endpoints of an original edge of g, and let d be the original vertex of $h \setminus g$ such that d and y are the endpoints of an original edge of h.

Because $g \cap h$ contains an edge of L and a vertex not in L, it follows that L is not a face. Statement 3 of Lemma 6.11 easily implies that h contains an edge of L not in g and in fact that $L \subseteq h$. As before we see that if u and v are either small link vertices or big link vertices with root vertices in $g \cap h$, then every 3-cell of $B(k - \frac{1}{2})$ which contains u meets every 3-cell of $B(k - \frac{1}{2})$ which contains v, contrary to Lemma 6.7. Since $c \notin L$, it is easy to see using Lemma 6.5 and Figure 6 that the root vertex of every big link vertex of e_g is contained in $g \cap h$. The previous two sentences imply that either u or v is a big link vertex of e_h with root vertex not in $g \cap h$. From this, Lemma 6.5 and Figure 6 we see that b is the only big link vertex of e_h with root vertex not in e_h, $b \in \{u, v\}$ and that h contains an edge unequal to e_h which is wild relative to C. The last assertion and Lemma 6.2 imply that e_h contains more than one edge. Let K and f be as in Lemma 6.8. It is easy to see that $e_h \subseteq f$ and that K contains an edge of e_h. From Lemma 6.8 we conclude that b either 1) is contained in at most one edge of K, 2) immediately precedes (relative to f and C) a vertex of K which is contained in at most one edge of K or 3) immediately precedes (relative to f and C) an original vertex of C in K. The third case implies that b immediately precedes (relative to f and C) an edge which is wild relative to C. But this is impossible because e_h contains more than one edge. Hence the third case is impossible. Suppose that the first case holds, namely, that b is contained in at most one edge of K. Then $K \neq f$. Since we have that b is an original vertex of C in K which is contained in at most one edge of K, that K contains an edge of e_h and that e_h contains more than one edge, it is easy to see that K is not an elbow of C. Now Lemma 6.8 implies that K is an edge. Hence u and v are the endpoints of the edge of e_h which contains b. In the second case it is easy to see that u and v are again the endpoints of the edge of e_h which contains b. Thus we have reduced the proof of Lemma 6.13 to the case in which u and v are the endpoints of the edge of e_h which contains b.

Suppose that u and v are the endpoints of the edge of e_h which contains b. Without loss of generality we assume that $v = b$. As we have seen before, because u and v both lie in e_h, they are both big link vertices. Since v is the only big link vertex of e_h with root vertex not in $g \cap h$, it follows that u is a big link vertex with root vertex in $g \cap h$. Using Figure 6 it is easy to see that L contains every vertex of $g \cap h$ except x. Now let C' be the 3-cell of $B(k - \frac{1}{2})$ other than C which contains h. It follows that $L \subseteq C'$, and it is easy to see that L is an original edge of C'. So $C' \cap B(k - \frac{3}{2})$ contains an original edge of C' which consists of more than one edge. Hence statement 7 of Theorem 5.1 implies that $C' \cap B(k - \frac{3}{2})$ is a face of C'. It is easy to see that relative to C', the vertex of L adjacent to x corresponds to w, u corresponds to b and v corresponds to a in part a) of Figure 6. It follows that u and v are contained in a face of $C' \cap \partial B(k - \frac{1}{2})$. This leads us to replace γ_2 and γ_3 by slightly different edge paths γ_2' and γ_3'. Let z be the inner vertex of γ_2 adjacent to both u and v. Recall how z was chosen. We chose a face of $\partial B(k - \frac{1}{2})$ which contains u and v, then we chose the 3-cell C of $B(k - \frac{1}{2})$ containing this face and then we chose the vertex z of $\partial B(k - \frac{1}{2})$ dual to C. We define a new edge path γ_2' which has the same vertices as γ_2 except that we replace z by the vertex z' dual to C'. We also define a new edge path γ_3' which is a pull-in of γ_2' and whose vertices, except for z' and its two inner vertices adjacent to z', equal the vertices of

γ_3. Then γ_2' is a pull-in of γ_1. Since the $B(k-1)$-link of z' is a face, Lemma 6.12 implies that the distance from z' to one of the endpoints of γ_3' is less than 2. It easily follows that the distance from either u or v to one of the endpoints of γ_2 is less than 2.

This proves Lemma 6.13. □

Just as Lemma 6.6 handles vertices with smallest possible link types and Lemma 6.12 handles vertices with biggest possible link types, Lemma 6.14 handles outer vertices of γ_3 whose link duals have no wild edges.

Lemma 6.14. *Let z be an outer vertex of γ_3 as in Figure 8. Let C be the 3-cell of $B(k-\frac{1}{2})$ dual to z, and let $L = C \cap B(k-\frac{3}{2})$. Suppose that the edges of L are all tame relative to C. Then the distance from z to one of the endpoints of γ_3 is less than 2.*

Proof. Since z is an outer vertex of γ_3, it is also an inner vertex of γ_2. Let u and v be the outer vertices of γ_2 adjacent to z. Lemma 6.9 implies that $u, v \in \text{star}(L, \partial C)$. Statement 7 of Theorem 5.1 and Lemma 6.12 imply that we may assume that L is contained in a twisted original edge of C. So we assume that L is contained in a twisted original edge of C. From this and the fact that every edge of L is tame relative to C, it follows that every two faces of $\text{star}(L, \partial C)$ have a nonempty intersection. Now Lemma 6.10 easily implies that u and v are contained in the union of two adjacent original edges of $\partial\text{star}(L, \partial C)$ which are not contained in one face of $\text{star}(L, \partial C)$.

We next consider Figure 12. In Figure 12 the complex L is drawn with dashes and corners of polygons are original vertices of C. In part a) of Figure 12 is a diagram of the complex $\text{star}(L, \partial C)$ if L consists of just one edge which does not contain an original vertex of C. In part b) of Figure 12 is a diagram of the complex $\text{star}(L, \partial C)$ if L consists of just one edge which contains an original vertex of C. In part c) of Figure 12 is a diagram of the complex $\text{star}(L, \partial C)$ if L is an elbow of C. Statement 7 of Theorem 5.1 and Lemma 6.12 show that these are all the cases which we must consider.

Considering part a) of Figure 12, we see that the fact that u and v are contained in the union of two original edges of $\partial\text{star}(L, \partial C)$ which are not contained in one face of $\text{star}(L, \partial C)$ and statement 3 of Lemma 6.11 imply that L contains an original vertex of C. Similarly, considering part b) of Figure 12 and statements 2 and 3 of Lemma 6.11, we see that L contains at least two edges. Finally, considering part c) of Figure 12 and using statement 2 of Lemma 6.11 and Lemma 6.13, we see that the distance from either u or v to one of the endpoints of γ_2 is less than 2. It easily follows that the distance from z to one of the endpoints of γ_3 is less than 2.

This proves Lemma 6.14. □

Suppose that Q has no wild edges. Then Lemma 6.14 implies that the length of γ_3 is at most 3. It follows that the length of γ is at most 5. This proves Theorem 6.1 if Q has no wild edges. Thus the rest of the proof of Theorem 6.1 is devoted to handling wild edges. As the argument progresses, our bound on the length of γ gradually increases, and at the very end the bound on the length of γ is not an absolute constant, but rather it depends on P and ϵ. Both this and the amount of work remaining indicate the wildness of wild edges. We continue by improving Lemma 6.14 under further assumptions in Lemma 6.15.

Lemma 6.15. *Let z be an outer vertex of γ_4 as in Figure 14. Let C be the 3-cell of $B(k-1)$ dual to z, and let $L = C \cap B(k-2)$. Let u and v be the outer vertices of γ_3 adjacent to*

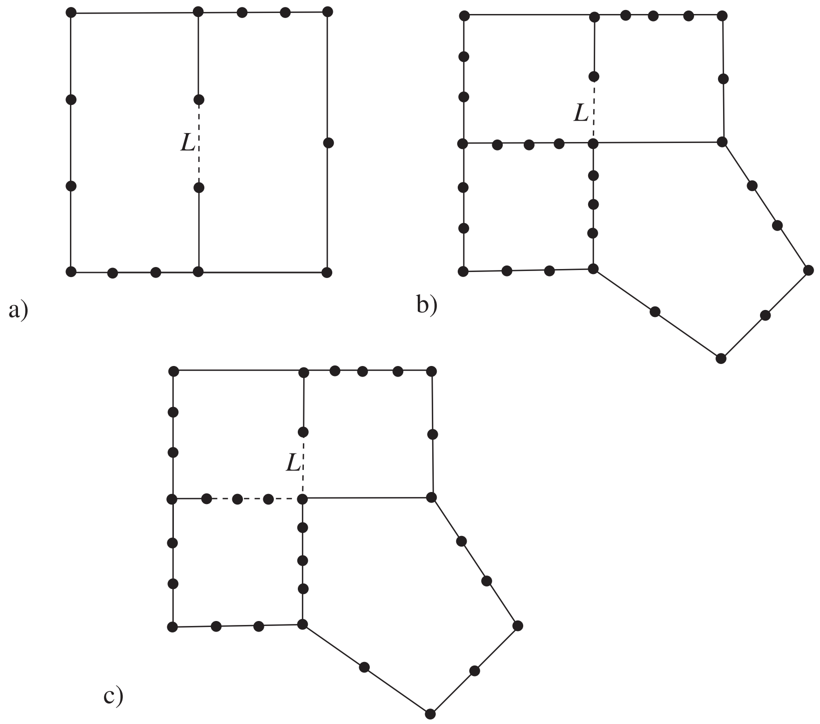

FIGURE 12. The complex star$(L, \partial C)$.

z. Suppose that u and v are both contained in the union of two adjacent original edges of ∂star$(L, \partial C)$. Then the distance from z to one of the endpoints of γ_4 is less than 2.

Proof. We begin by analyzing the case in which either u or v is a small link vertex. Suppose that u is a small link vertex. Lemma 6.9 implies that $u \in \text{star}(L, \partial C)$. Because u is a small link vertex, it follows that u is contained in exactly one face g of star$(L, \partial C)$. Applying Lemma 6.14 with the vertex z in Lemma 6.14 equal to u, it is easy to see that we may assume that the dual of the $B(k-1)$-link of u is an edge which is wild relative to the 3-cell of $B(k-\frac{1}{2})$ dual to u. It easily follows that u is an original vertex of C. It is also true that u is an original vertex of the other 3-cell of $B(k-1)$ which contains u. This and the definition of twisted face-pairing imply that the edge of g immediately following u (relative to g and C) is wild relative to C. So we may assume that if u is a small link vertex, then u is an original vertex of C which is contained in exactly one face of star$(L, \partial C)$ and the edge of this face immediately following u is wild (relative to this face and C). The same holds for v.

By Lemmas 6.12 and 6.14, we may assume that L is contained in a twisted original edge of C and that L contains an edge which is wild relative to C. We consider Figure 13, which has the same meaning as Figure 12.

Statement 1 of Lemma 6.10 easily implies that u and v are both contained in the union of two adjacent original edges of ∂star$(L, \partial C)$ which are not contained in one face of star$(L, \partial C)$. Hence we may assume that following. There exist distinct faces g and h of star$(L, \partial C)$ as in Figure 9 such that $g \cap h$ contains an original vertex x of C not in L, $g \cap h$ contains an original vertex y of L, and x immediately precedes $g \cap h$ (relative to g and C). Let e_g be the original edge of $g \cap \partial$star$(L, \partial C)$ which contains x, and let e_h be the original edge of $h \cap \partial$star$(L, \partial C)$ which contains x. We may assume that $\{u, v\} \subseteq e_g \cup e_h$. Lemma 6.13 implies that we may

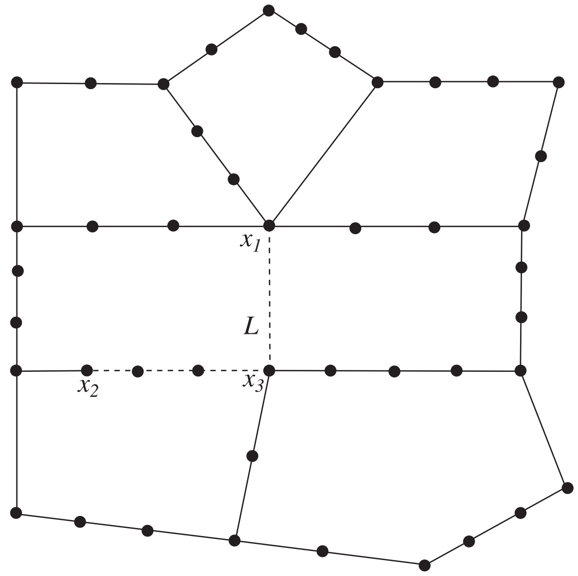

FIGURE 13. The complex $\text{star}(L, \partial C)$.

assume that $g \cap h$ does not contain an edge of L. Statement 2 of Lemma 6.11 implies that either g or h contains the edge of L which is wild relative to C.

Suppose that g contains at most one edge of L. In this case it is easy to see using Figure 13, Lemma 6.2, Lemma 6.5 and Figure 6 that every big link vertex of $e_g \cup e_h$ is an original vertex of C. By the first paragraph of this proof we may assume that if either u or v is a small link vertex, then it is an original vertex of C. Hence we may assume that both u and v are original vertices of $e_g \cup e_h$. Statement 1 of Lemma 6.11 shows that this is impossible.

We have reduced the proof of Lemma 6.15 to the situation in which g contains more than one edge of L and $g \cap h$ does not contain an edge of L. Using Figure 13, we see that $L \subseteq g$. The case in which $\{u, v\} \subseteq e_g$ reduces to the case in which $g \cap h$ contains an edge of L, so we may assume that $v \notin e_g$. As usual, it is easy to see that every vertex of $e_h \setminus e_g$ is a small link vertex. So by the first paragraph of this proof we may assume that v is the original vertex of e_h other than x and that the edge of h immediately following v (relative to h and C) is wild relative to C. This and Lemma 6.2 imply that e_h is not an edge which is wild relative to C. Just as we may assume that v is the original vertex of e_h other than x, we may assume that if $u \in e_h$, then $u = x$. In other words, we may assume that $u \in e_g$. Let K and f be as in Lemma 6.8. Lemma 6.10 easily implies that $e_h \subseteq f$, and then because $u \in e_g$ we see that $e_h \subseteq K$. As in previous arguments, it follows that $e_h \subseteq K$. Hence K contains an original edge of f which is not wild relative to C, and so Lemma 6.8 easily implies that $K = f$. Now Lemma 6.8 furthermore implies that u immediately precedes (relative to f and C) an original vertex of C in K. This means that e_h is an edge which is wild relative to C, a contradiction.

This proves Lemma 6.15. $\square$

We next consider Lemma 6.16, which we apply in the proof of Lemma 6.17.

Lemma 6.16. *Let w be an outer vertex of γ_2. Let u and v be the inner vertices of γ_2 adjacent to w as in Figure 14. Let D be the 3-cell of $B(k)$ dual to w. Suppose that u and*

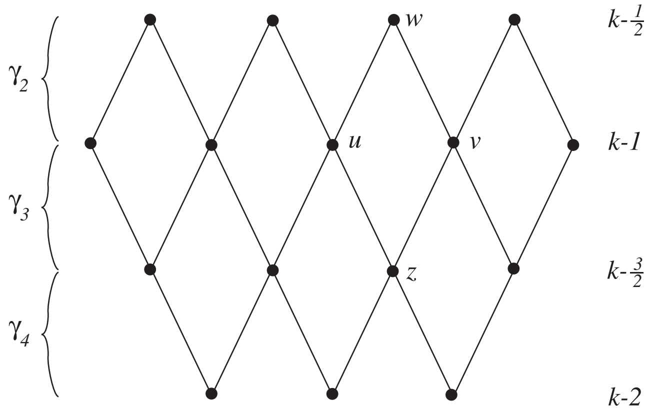

FIGURE 14. Parts of γ_2, γ_3, and γ_4.

v are the vertices of an edge of D which is wild relative to D. Then the distance from w to the endpoints of γ_2 is less than 3.

Proof. We prove Lemma 6.16 by contradiction: suppose that the distance from w to the endpoints of γ_2 is at least 3. Let e be the edge of D which contains u and v. Let z be the inner vertex of γ_3 adjacent to both u and v. Let C be the 3-cell of $B(k-1)$ dual to z, and let $L = C \cap B(k-2)$. Lemma 6.15 easily implies that e is not contained in $\partial\mathrm{star}(L, \partial C)$, and so e is like the edge e_2 in Figure 19. Lemma 6.9 implies that both u and v are contained in $\partial\mathrm{star}(L, \partial C)$. It easily follows that e is wild relative to C. But e cannot be wild relative to both C and D by Lemma 6.3. This contradiction proves Lemma 6.16. □

Continuing the related sequence Lemma 6.6, Lemma 6.12 and Lemma 6.14, we next show in Lemma 6.17 that the outer vertices of γ_3 are almost all big link vertices.

Lemma 6.17. *Let z be an outer vertex of γ_3, and suppose that z is a small link vertex. Then the distance from z to the endpoints of γ_3 is less than 3.*

Proof. To begin the proof of Lemma 6.17, we note that Lemma 6.6 implies that the dual of $\mathrm{link}(z, B(k-1))$ is an edge e. Let C be the 3-cell of $B(k-\frac{1}{2})$ dual to z. Lemma 6.14 implies that we may assume that e is wild relative to C. Now Lemma 6.16 applied to z, γ_3 and C completes the proof of Lemma 6.17. □

The main point of Lemma 6.18 is that if x and y are adjacent outer vertices of γ_5, then x and y are almost always the endpoints of the dual of the $B(k-\frac{3}{2})$-link of the outer vertex of γ_4 adjacent to both x and y.

Lemma 6.18. *Let z be an outer vertex of γ_4 such that the distance from z to the endpoints of γ_4 is at least 3. Let x and y be the inner vertices of γ_4 adjacent to z as in Figure 15. Let C be the 3-cell of $B(k-1)$ dual to z, and let $L = C \cap B(k-\frac{5}{2})$. Then L is an elbow of C such that x and y are the vertices of L which are contained in just one edge of L. Furthermore, the bottom of L is wild relative to C.*

Proof. To begin the proof of Lemma 6.18, we note that Lemma 6.17 easily implies that z is a big link vertex, and hence L contains at least two edges. Lemma 6.12 easily implies that L is not a face. It easily follows using statement 7 of Theorem 5.1 that L is an elbow of C. Lemma 6.14 easily implies that the bottom of L is wild relative to C.

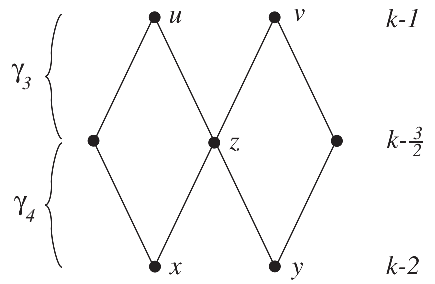

FIGURE 15. Parts of γ_3 and γ_4.

To prove Lemma 6.18, it remains to prove that x and y are the vertices of L which are contained in just one edge of L, which we do in this paragraph. Lemma 6.4 easily implies that x and y are either original vertices of L or vertices of L which are contained in at most one edge of L. Figure 13 is a diagram of $\text{star}(L, \partial C)$. The edges of L are drawn with dashes. We have that $\{x, y\} \subseteq \{x_1, x_2, x_3\}$. The case $\{x, y\} = \{x_1, x_3\}$ is ruled out by Lemma 6.16. Hence it remains to rule out the case $\{x, y\} = \{x_2, x_3\}$. Suppose that $\{x, y\} = \{x_2, x_3\}$. Let u and v be the outer vertices of γ_3 which are adjacent to z as in Figure 15. Let g and h be faces of $\text{star}(L, \partial C)$ such that $u \in g$ and $v \in h$. Since $u \in g$, the definition of pull-in easily shows that $x \in g$. Likewise the fact that $v \in h$ implies that $y \in h$. Figure 13 now easily shows that g meets h. Lemma 6.10 now implies that u and v are contained in the union of two adjacent original edges of $\partial\text{star}(L, \partial C)$. This is impossible by Lemma 6.15. This proves Lemma 6.18. $\square$

Our last lemma, Lemma 6.19, provides the setting for the conclusion of the proof of Theorem 6.1.

Lemma 6.19. *Let z be an outer vertex of γ_6 such that the distance from z to the endpoints of γ_6 is at least 3. See Figure 16. Let C be the 3-cell of $B(k-2)$ dual to z. Let u and v be the outer vertices of γ_5 adjacent to z. Then u and v are both big link vertices. Let K be the dual in $B(k-2)$ of the $B(k-\frac{3}{2})$-link of the outer vertex of γ_4 adjacent to both u and v. Then there exist adjacent original edges e_1 and e_2 of C such that $K \subseteq e_1 \cup e_2$, $e_2 \subseteq K$ and e_2 is wild relative to C. Either $u \in e_1$ and $v \in e_2$ or $u \in e_2$ and $v \in e_1$. We assume without loss of generality that $u \in e_1$ and $v \in e_2$. Then v is the original vertex of e_2 not in e_1. Furthermore if x is the original vertex of e_1 not in K, then link$(x, B(k-2))$ is dual to a face and the root vertex of x equals the root vertex of u.*

Proof. Lemma 6.17 easily implies that u and v are big link vertices. Let w be the outer vertex of γ_4 adjacent to both u and v as in Figure 16. Lemma 6.18 applied to w and γ_4 and Lemma 6.8 easily imply that K is an elbow of C and that u and v are the vertices of K which are contained in just one edge of K. It easily follows that there exist adjacent original edges e_1 and e_2 of C such that $K \subseteq e_1 \cup e_2$. Let $L = C \cap B(k-3)$. Lemma 6.9 implies that both u and v are contained in $\partial\text{star}(L, \partial C)$. Lemma 6.15 applied to z and γ_6 easily implies that u and v are not both contained in the union of two adjacent original edges of $\partial\text{star}(L, \partial C)$. It easily follows that K contains two original vertices of C. But if K contains two original vertices of C, then K contains an edge which is wild relative to C. Without loss of generality we assume that $e_2 \subseteq K$ and that e_2 is wild relative to C. Let f be the face of C which contains K. See Figure 17. It is clear that e_1 contains one original vertex

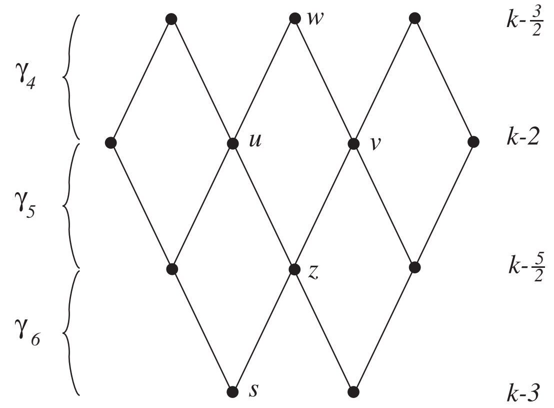

FIGURE 16. Parts of γ_4, γ_5, and γ_6.

in K and one original vertex not in K. Let x denote the original vertex of e_1 not in K, and let y denote the original vertex of e_1 in K. Since u and v are the vertices of K which are contained in just one edge of K, either u or v lies in the interior of e_1. We assume without loss of generality that u lies in the interior of e_1. It follows that v is the original vertex of e_2 not in e_1.

To prove Lemma 6.19, it remains to prove that the $B(k-2)$-link of x is dual to a face and that the root vertex of x equals the root vertex of u. Figure 6 shows that either x is a big link vertex whose root vertex equals the root vertex of u or y is a big link vertex whose root vertex equals the root vertex of u. To prove that x is a big link vertex whose root vertex equals the root vertex of u, we next prove that it is impossible for y to be a big link vertex whose root vertex equals the root vertex of u. Suppose that y is a big link vertex whose root vertex equals the root vertex of u. Considering Figure 6, we see that we are in the situation of either part c) or part d) of Figure 6, that u corresponds to the vertex in Figure 6 labeled a and that y corresponds to the vertex in Figure 6 labeled b. Let g and h be the faces of C as in Figure 6. Lemma 6.18 easily implies that L is an elbow of C and that the bottom of L is wild relative to C. It easily follows that h is the face of C which contains L. We have that $u \in g$. With this, Lemma 6.10 and Lemma 6.15 easily imply that $v \notin g \cup h$. Lemma 6.9 implies that there exists a face g' of $\mathrm{star}(L, \partial C)$ which contains v. Ampleness condition 2 applies to the faces f, g' and h, and so the faces f, g' and h have exactly one vertex t in common. Because g and h are the only two faces of $\mathrm{star}(L, \partial C)$ which contain y and g' equals neither g nor h, it follows that $t \neq y$. But then f and h contain both t and y, and so $f \cap h$ is an original edge of C. Since $f \cap h$ immediately precedes y (relative to f and C), it follows that $f \cap h = e_2$. This is impossible because $v \notin h$. This contradiction shows that y is not a big link vertex whose root vertex equals the root vertex of u. So x is a big link vertex whose root vertex equals the root vertex of u. Finally, Figure 6 shows that x is contained in its root edge, and so the dual of $\mathrm{link}(x, B(k-2))$ is a face. This proves Lemma 6.19. $\square$

In this paragraph we begin the conclusion of the proof of Theorem 6.1. We maintain the setting of Lemma 6.19. As above, we let f denote the face of C which contains e_1 and e_2, and we let y be the vertex common to e_1 and e_2. The last assertion of Lemma 6.19 implies that x is the central vertex of the face tree of some face g of some 3-cell D of $B(k-\frac{5}{2})$. Figure 18 shows D, f, the face tree of g and the big face tree of g. The edges of the face tree of g are drawn with thick arcs. As usual, the corner vertices of D in Figure 18 are

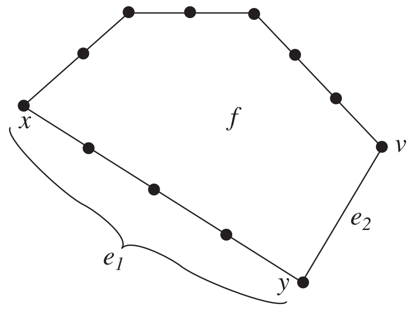

FIGURE 17. The face f.

original vertices of D. Since u is a big link vertex whose root vertex equals the root vertex of x, the vertex s dual to D is an outer vertex of γ_7 adjacent to z as in Figure 16. One of C or D is canonically isomorphic to Q in an orientation-preserving way and the other is canonically isomorphic to Q^* in an orientation-preserving way. As in two paragraphs before the verification of induction hypothesis 2, we see by means of these isomorphisms and the orientation-reversing isomorphism $\omega : Q \to Q^*$ that there exists a canonical correspondence between the vertices of C and the vertices of D. Let s' be the vertex of D corresponding to the vertex s of C. As in two paragraphs before the verification of induction hypothesis 2, we see that C is dual to s', that is, $s' = z$. Suppose that the distance from s to the endpoints of γ_7 is at least 3. Then Lemma 6.19 easily implies that the edge of g immediately following s' (relative to g and D) is in fact an edge of g which is wild relative to D. If x' is the vertex of D corresponding to the vertex x of C, then as in two paragraphs before the verification of induction hypothesis 2, we see that x' is the vertex of g immediately following s' (relative to g and D). Because of the reversal of orientation between C and D, it follows that the vertex y' in Figure 18 corresponds to y and that the vertex v' in Figure 18 corresponds to v. So the assumption that the distance from z to the endpoints of γ_6 is at least 3 implies the existence of the wild edge joining v and y and the original edge joining x and y in C. The correspondence between the vertices of C and D then gives the wild edge joining v' and y' and the original edge joining x' and y' in D. Then the assumption that the distance from s to the endpoints of γ_7 is at least 3 implies the existence of the wild edge joining s' and x' and the original edge joining x' and t' in D. It is clear that the original edge of D joining x' with y' and the original edge of D joining x' with t' correspond to edges of P which are in the same edge cycle of the model face-pairing ϵ. Denote this edge cycle by E. It follows that if z is sufficiently far from the endpoints of γ_6, then the edges of E form a closed edge path α in P. We call α an **edge cycle edge path**. Example 7.4 shows how edge cycle edge paths might arise. It is easy to see that although α might intersect itself, these self intersections are tangential, not transverse. Hence $P \setminus \alpha$ is a canonical union of two subsets, one of which we say is **inside** α and one of which we say is **outside** α. The wild edges which we obtain in the construction of α correspond to edges of P which we call **ancillary edges** of α. We call the ancillary edges of α which lie inside α **inner ancillary edges**, and we call the ancillary edges of α which lie outside α **outer ancillary edges**. The ancillary edges of α alternate between being inner and outer. Just as α cannot intersect itself transversely, if α' is any edge cycle edge path, then α and α' cannot meet transversely. Hence we may choose α so that there does not exist an edge cycle edge path inside α.

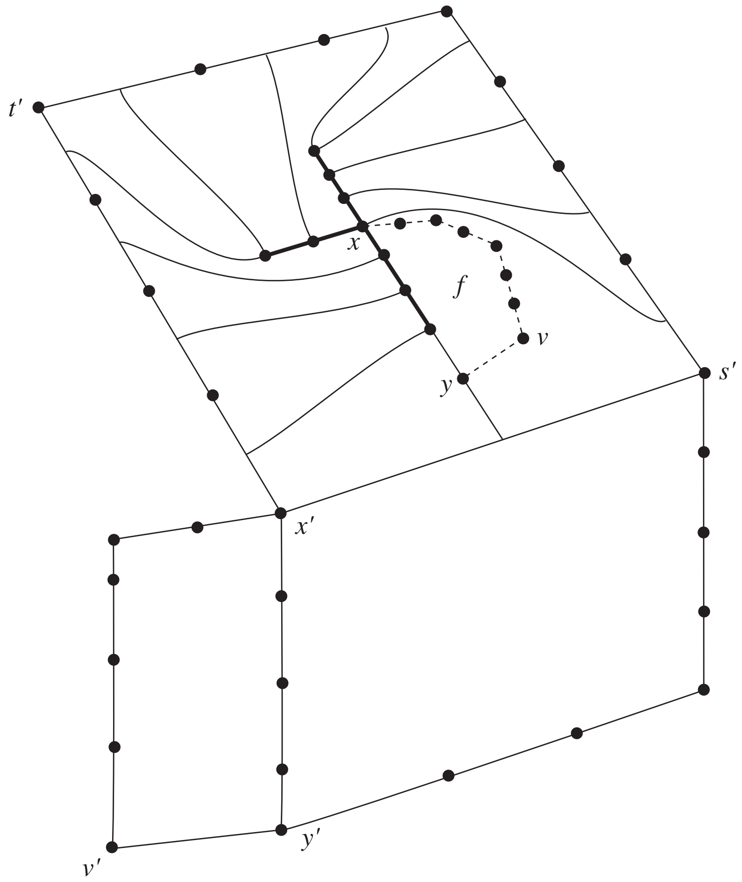

FIGURE 18. Part of the 3-cell D.

We return to the setting of Lemma 6.19. We still let L denote the dual of $\text{link}(z, B(k-\frac{5}{2}))$. Figure 19 shows $\text{star}(L, \partial C)$, where L is drawn with dashes. The original edges e_4 and e_5 are related to L just as the original edges e_1 and e_2 are related to K. Lemma 6.18 applied to z and γ_6 implies that s is a vertex of L which is contained in just one edge of L. The previous paragraph shows that s' and hence s is an original vertex of C. It follows that s is the vertex of e_5 not in e_4. Lemma 6.19 easily implies that v is a big link vertex. Lemma 6.18 applied to z and γ_6 easily implies that the root vertex of v is the vertex of $L \cap e_4$ which is contained in just one edge of L. Lemmas 6.10 and 6.15 easily imply that v is not contained in the face of C which contains L. Using this, Figure 6 and the fact that v is both a big link vertex and an original vertex, we see that v is joined to the original vertex of e_4 not in e_5 by an edge e_3 which is wild relative to C as in Figure 19. If z is sufficiently far from the endpoints of γ_6, then L gives rise to an edge cycle edge path α' just as K gives rise to the edge cycle edge path α. It follows that the original edge e_4 of C corresponds to an edge of P in the edge cycle of α'. Furthermore the edges e_3 and e_5 of Figure 19, which are wild relative to C, correspond to ancillary edges of α'. Let e_6 be the edge joining x and s. Since either e_2 or e_6 corresponds to an inner ancillary edge of α, we obtain the following conclusion. If γ is sufficiently long, then there exists an edge cycle edge path α with no edge cycle edge path inside it such that every inner ancillary edge of α either has both endpoints in α or it meets another inner ancillary edge of α. The lower bound on the length of γ in the previous sentence depends only on the maximum length of the edge cycle edge paths of P, and so it

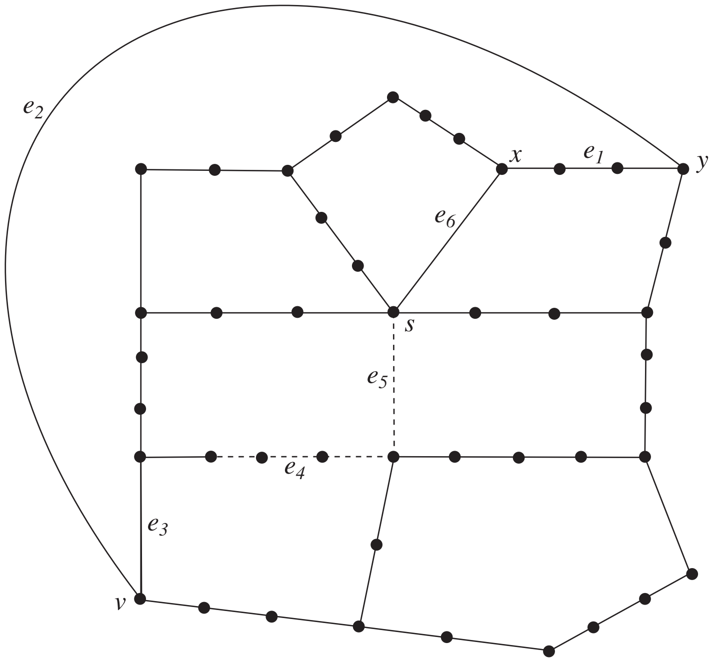

FIGURE 19. The complex star$(L, \partial C)$ and the wild edge e_2.

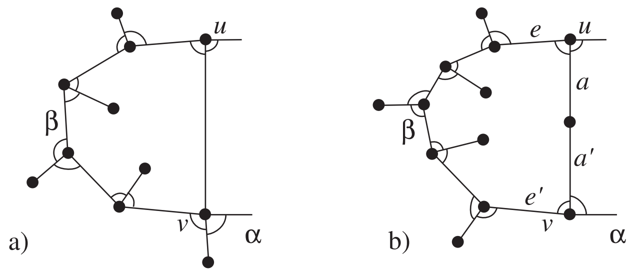

FIGURE 20. The edge subpath β of the edge cycle edge path α.

depends only on P. In the next paragraph we show that there does not exist an edge cycle edge path α with no edge cycle edge path inside it such that every inner ancillary edge of α either has both endpoints in α or it meets another inner ancillary edge of α. Hence it is impossible for γ to be longer than the above bound which depends only on P. This suffices to prove Theorem 6.1.

We conclude the proof of Theorem 6.1 in this paragraph. Let α be an edge cycle edge path of P with no edge cycle edge path inside it such that every inner ancillary edge of α either has both endpoints in α or it meets another inner ancillary edge of α. Let β be an edge subpath of α whose length is minimal with respect to the property that its endpoints are joined by either one or two inner ancillary edges of α. Let u and v be the endpoints of β. See Figure 20. We claim that β contains at least three edges. To see this, we note that if u and v are joined by one inner ancillary edge of α as in part a) of Figure 20, then this follows from Proposition 3.1, which states that there does not exist a nontrivial simple closed edge path in P consisting of three or fewer edges. If u and v are joined by two inner ancillary edges of α as in part b) of Figure 20, then the length of β is even, and so β has at least two edges. Suppose that β has exactly two edges. Denote these edges by e and e',

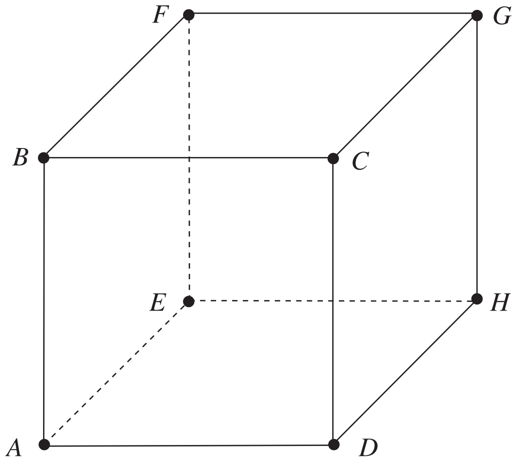

FIGURE 21. The complex P.

and denote the associated inner ancillary edges of α by a and a' as in part b) of Figure 20. Then some face f of P contains both a and e, and some face f' of P contains both a' and e'. Since a and a' are wild edges of Q, Lemma 6.2 shows that $f \neq f'$. We have that f and f' both contain the vertex common to a and a' as well as the vertex common to e and e'. Hence $f \cap f'$ is an edge joining these two vertices. But then f and f' are triangles, which is impossible. Thus β contains at least three edges. But then some inner ancillary edge of α meets β in a vertex other than u or v. Whether the endpoints of this inner ancillary edge are both in α or this inner ancillary edge meets another inner ancillary edge of α, we have a contradiction to the minimality of β.

This completes the proof Theorem 6.1. □

7. EXAMPLES

The simplest example of an ample faceted 3-ball is a cube, and for ease of discussion we consider a cube P in Euclidean 3-space with center at the origin. The most obvious choice of a face-pairing on P is the face-pairing which identifies opposite faces by translation. This is highly symmetric, and each of the three edge cycles has length four. But even if we choose each multiplier to be 1, then each face of Q is a 16-gon. In order to decrease the number of edges in Q (and hence to slow down the growth rate of our combinatorial balls in $\widetilde{M}$), we first look at the face-pairing ϵ where each face-pairing map is the antipodal map. This is still highly symmetric and seems to be the "simplest" face-pairing on P.

Example 7.1. Let P be a cube in Euclidean 3-space with center the origin, labeled as in Figure 21. Let ϵ be the face-pairing on P with each face-pairing map the antipodal map. In the notation of [7, 8], $\epsilon = \{\epsilon_1^{\pm 1}, \epsilon_2^{\pm 1}, \epsilon_3^{\pm 1}\}$, where

$$\epsilon_1 : \begin{pmatrix} A & B & C & D \\ G & H & E & F \end{pmatrix}, \ \epsilon_2 : \begin{pmatrix} E & F & B & A \\ C & D & H & G \end{pmatrix}, \text{ and } \epsilon_3 : \begin{pmatrix} E & A & D & H \\ C & G & F & B \end{pmatrix}.$$

The edge cycles for ϵ have diagrams

$$AB \xrightarrow{\epsilon_1} GH \xrightarrow{\epsilon_2^{-1}} AB, \quad BC \xrightarrow{\epsilon_1} HE \xrightarrow{\epsilon_3} BC, \quad CD \xrightarrow{\epsilon_1} EF \xrightarrow{\epsilon_2} CD,$$

$$AD \xrightarrow{\epsilon_1} GF \xrightarrow{\epsilon_3^{-1}} AD, \quad AE \xrightarrow{\epsilon_2} GC \xrightarrow{\epsilon_3^{-1}} AE, \quad \text{and} \quad BF \xrightarrow{\epsilon_2} HD \xrightarrow{\epsilon_3} BF.$$

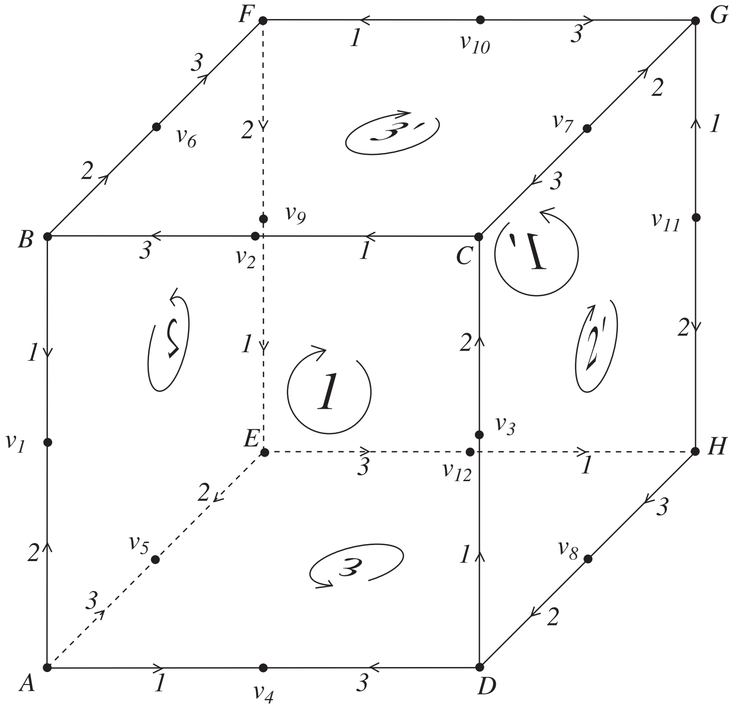

FIGURE 22. The complex Q.

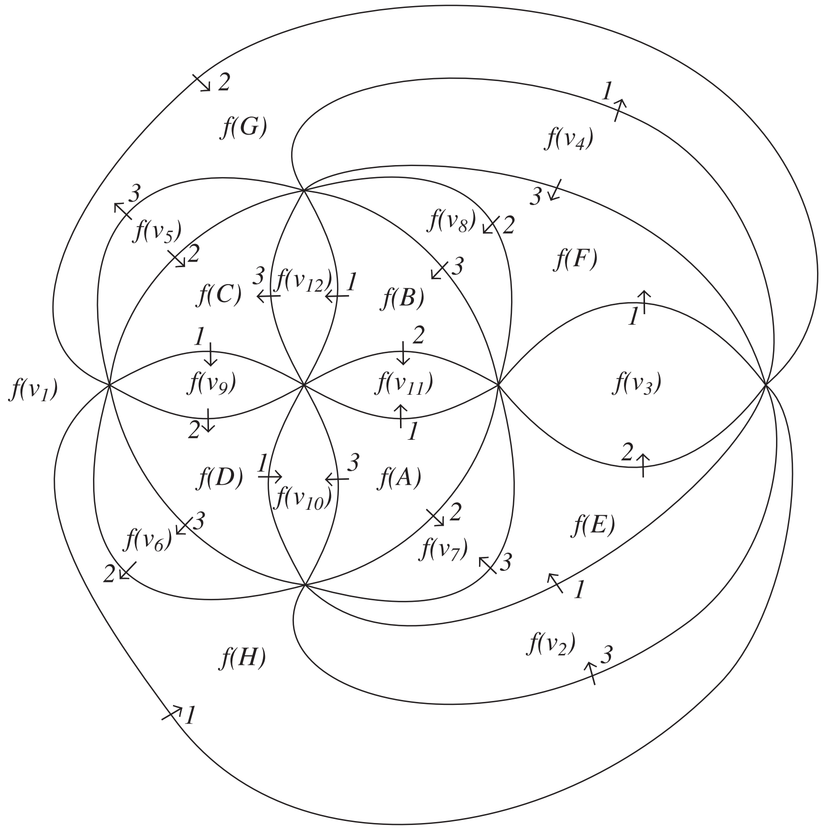

FIGURE 23. The link of the vertex of M.

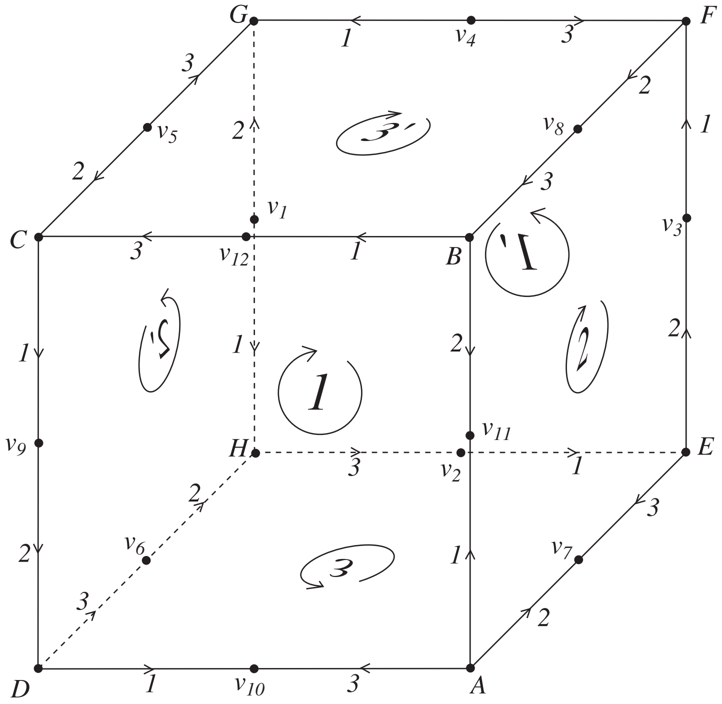

FIGURE 24. The complex Q^*.

To keep the example as simple as possible, we choose each multiplier to be 1. Then each of the twelve edges of P is subdivided into two subedges. Figure 22 shows the faceted 3-ball Q, where the new vertices have been labeled arbitrarily. The link of the vertex of M is shown in Figure 23. Conventions for both figures are as in [8].

Figure 24 shows the faceted 3-ball Q^*, and Figure 25 shows the link of the vertex of M^*. Note that Q^* is dual to the link of the vertex of M, and Q is dual to the link of the vertex of M^*. Following [8], we can read off a presentation for $\pi_1(M)$ from the labels of the edges of the faces of Q^*. Thus,

$$\begin{aligned}\pi_1(M) \cong \langle x_1, x_2, x_3 \ &: x_1x_2^{-1}x_1x_3x_1x_2x_1x_3^{-1},\ x_2x_3^{-1}x_2x_1x_2x_3x_2x_1^{-1},\\ &\ x_3x_1^{-1}x_3x_2x_3x_1x_3x_2^{-1}\rangle.\end{aligned}$$

Turning our attention to combinatorial balls, $B(0)$ is a vertex, and $B(\frac{1}{2})$ is isomorphic to a copy of Q. Figure 26 shows $B(1)$. In the figure $B(1)$ is pictured as a cube, and we see only three sides of the cube. Face trees are drawn with thick arcs, root edges are drawn as thick dashed arcs, and edges of $\partial B(1)$ that are in a single 3-cell are drawn as dotted arcs. Some thin dashed arcs are drawn. These thin dashed arcs are not part of $B(1)$. They are drawn simply to complete edges of the cube. There are 20 vertices in Q, and each of these corresponds to one of the 3-cells of $B(1)$. For example, the 3-cell $Q^*_{v_2}$ of $B(1)$ dual to the vertex v_2 of Q intersects $\partial B(1)$ in a disk which is the union of four faces. It contains six big link vertices of $\partial B(1)$, six small link vertices of $\partial B(1)$ whose links are dual to edges, and seven small link vertices of $\partial B(1)$ (on dotted edges) whose links are dual to vertices. Of course, $Q^*_{v_2}$ also contains the vertex $B(0)$, giving $Q^*_{v_2}$ a total of $6+6+7+1 = 20$ vertices. The 3-cell Q^*_C of $B(1)$ dual to the vertex C of Q intersects $\partial B(1)$ in the union of a disk (which contains three faces) and three edges which are edges of face trees (so that we have a "hairy"

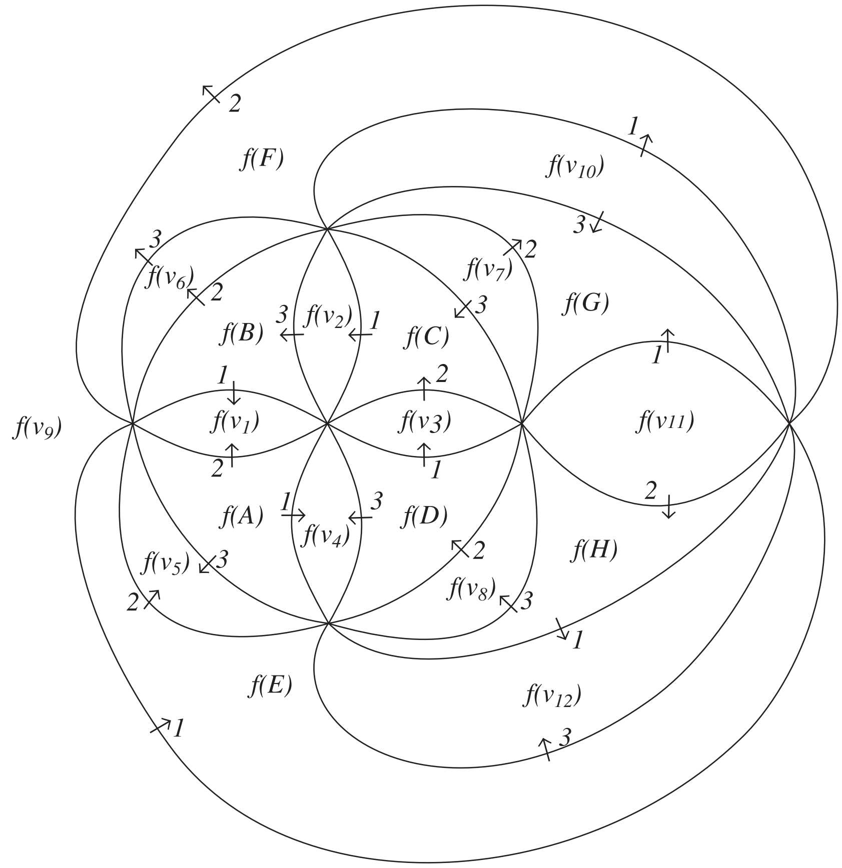

FIGURE 25. The link of the vertex of M^*.

disk as in statement 6 of Theorem 5.1). This last fact is a salient feature of combinatorial balls of universal covers of ample twisted face-pairing manifolds: not only does Q_C^* contain three edges of face trees, but Q_C^* also contains all three root edges drawn in Figure 26 as thick dashed arcs. As a result, $Q_{v_2}^*$ and $Q_{v_7}^*$ do not have a face in common as one might expect. Furthermore, Q_C^* contains six big link vertices of $\partial B(1)$, nine small link vertices of $\partial B(1)$ whose links are edges, and four small link vertices of $\partial B(1)$ whose links are vertices. Since $\pi_1(M)$ acts transitively on the 3-cells of $\widetilde{M}$, statement 7 of Theorem 5.1 with $k = \frac{3}{2}$ easily implies that the intersection of two distinct 3-cells of $\widetilde{M}$ is either empty or a vertex or an edge or a face or an elbow of both of the given 3-cells. For example, Q_C^* meets $Q_{v_2}^*$ in a face and Q_C^* meets $Q_{v_7}^*$ in a face. On the other hand, the intersection of $Q_{v_2}^*$ and $Q_{v_7}^*$ is an elbow of both $Q_{v_2}^*$ and $Q_{v_7}^*$. We next explain the meaning of the vertex labels in Figure 26 by example. Consider the central vertex v in the top of the cube in Figure 26. Next to v are the letters A, E, H and D. This means that in $Q_{v_{10}}^*$, the vertex v corresponds to the vertex A in Figure 24. In $Q_{v_7}^*$ the vertex v corresponds to the vertex E in Figure 24. In $Q_{v_2}^*$ the vertex v corresponds to the vertex H in Figure 24. In $Q_{v_6}^*$ the vertex v corresponds to the vertex D in Figure 24. As is the case for every combinatorial ball for every ample twisted face-pairing manifold, the central vertex of every face tree of $B(1)$ (in general $B(k)$) is contained in exactly one root edge and these root edges are exactly the edges of $B(1)$ (in general $B(k) \setminus B(k-1)$) which are not contained in $\partial B(1)$ (in general $\partial B(k)$).

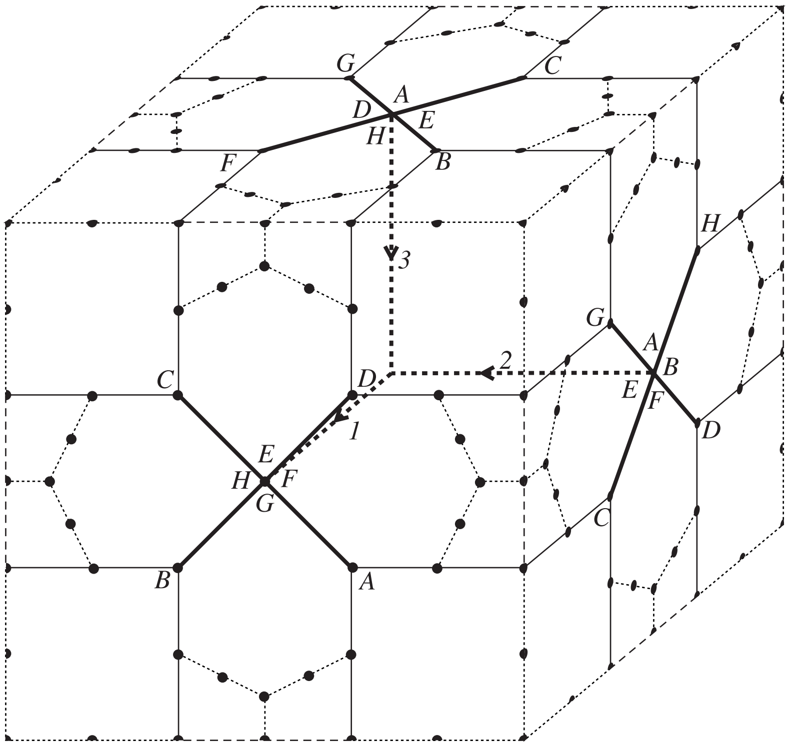

FIGURE 26. $B(1)$.

There are 218 vertices in $\partial B(1)$, and they parametrize the 3-cells of $B(\frac{3}{2})$ which meet $\partial B(\frac{3}{2})$. Hence there are 218 elements in the star generating set for $\pi_1(M)$. It follows easily from statement 3 of Theorem 5.1 that the growth function of $\pi_1(M)$ with respect to the star generating set is rational. With effort, one can calculate the recursion and show that the growth function is

$$f(z) = \frac{z^3 + 119z^2 + 119z + 1}{(1-z)(z^2 - 98z + 1)}.$$

The growth exponent is $49 + 20\sqrt{6} \approx 98$. There are 21,602 elements of $\pi_1(M)$ of length 2, and 2,117,018 elements of $\pi_1(M)$ of length 3. So even in this example, which we believe is the simplest ample example, the group is growing rapidly.

The growth exponent for $\pi_1(M)$ is the square of the growth exponent $5 + 2\sqrt{6}$ for the number of vertices of $\partial B(\frac{k}{2})$. Figure 27 shows $\partial B(\frac{3}{2})$. The figure was drawn using Kenneth Stephenson's circle packing program, CirclePack [12]. To improve clarity in this and several following figures, vertices have not been drawn. In particular, vertices of valence 2 have not been drawn, and so $\partial B(\frac{3}{2})$ contains many more vertices than is apparent. Similarly, Figure 28 shows $\partial B(2)$ and Figure 29 shows $\partial B(\frac{5}{2})$. According to SnapPea, M is a hyperbolic manifold with volume approximately 5.333. Even though M has relatively small volume, the length of its shortest closed geodesic is relatively large. The shortest closed geodesic in M has length approximately 1.266.

Example 7.2. In this example we again choose P to be a cube in Euclidean 3-space, labeled as in Figure 21. Let ϵ be the face-pairing on P with each face-pairing map a translation. In

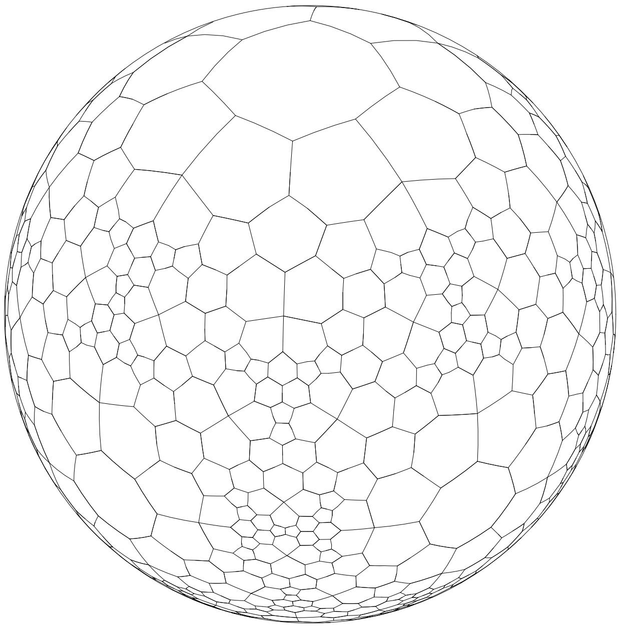

FIGURE 27. $\partial B(\frac{3}{2})$.

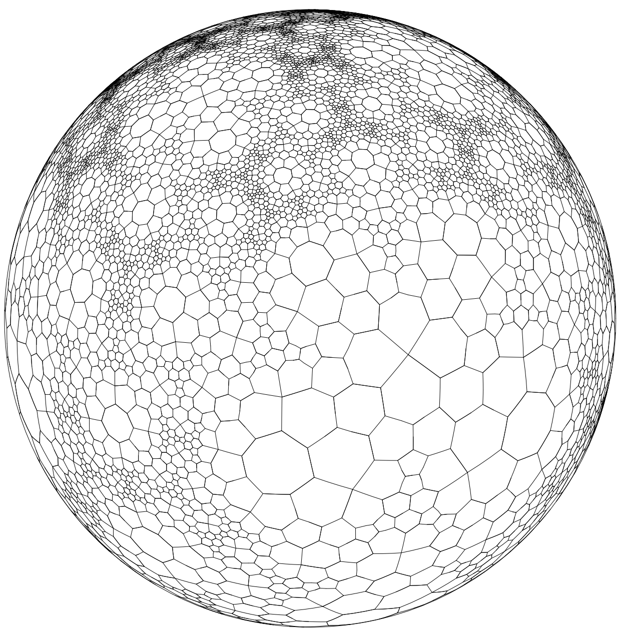

FIGURE 28. $\partial B(2)$.

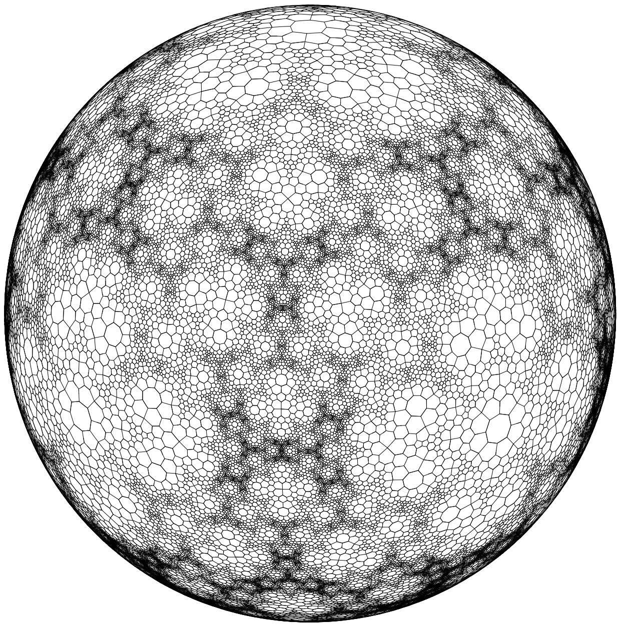

FIGURE 29. $\partial B(\frac{5}{2})$.

the notation of [7, 8], $\epsilon = \{\epsilon_1^{\pm 1}, \epsilon_2^{\pm 1}, \epsilon_3^{\pm 1}\}$, where

$$\epsilon_1 : \begin{pmatrix} A & B & C & D \\ E & F & G & H \end{pmatrix}, \ \epsilon_2 : \begin{pmatrix} E & F & B & A \\ H & G & C & D \end{pmatrix}, \text{ and } \epsilon_3 : \begin{pmatrix} E & A & D & H \\ F & B & C & G \end{pmatrix}.$$

The edge cycles for ϵ have diagrams

$$AB \xrightarrow{\epsilon_1} EF \xrightarrow{\epsilon_2} HG \xrightarrow{\epsilon_1^{-1}} DC \xrightarrow{\epsilon_2^{-1}} AB,$$

$$AD \xrightarrow{\epsilon_3} BC \xrightarrow{\epsilon_1} FG \xrightarrow{\epsilon_3^{-1}} EH \xrightarrow{\epsilon_1^{-1}} AD,$$

and

$$AE \xrightarrow{\epsilon_3} BF \xrightarrow{\epsilon_2} CG \xrightarrow{\epsilon_3^{-1}} DH \xrightarrow{\epsilon_2^{-1}} AE.$$

We again choose each multiplier to be 1. Then each of the twelve edges of P is subdivided into four subedges. Figure 30 shows $\partial B(1)$, where again we have drawn the face trees with thick arcs and drawn edges in a single 3-cell with dotted arcs. There are 1,178 vertices in $\partial B(1)$, and they parametrize the 3-cells in $B(\frac{3}{2})$ that meet $\partial B(\frac{3}{2})$. Hence they correspond to the generators of the star generating set for $\pi_1(M)$. Once again, it easily follows from statement 3 of Theorem 5.1 that $\pi_1(M)$ has a rational growth function with respect to the star generating set. With considerable effort, one can show that the growth function is

$$f(z) = \frac{z^3 + 503z^2 + 503z + 1}{(1-z)(z^2 - 674z + 1)}.$$

The growth exponent is $337 + 52\sqrt{42} \approx 674$. According to SnapPea, M is a hyperbolic manifold with volume approximately 12.84.

Figure 31, which was drawn using CirclePack, shows the boundary of $B(\frac{3}{2})$.

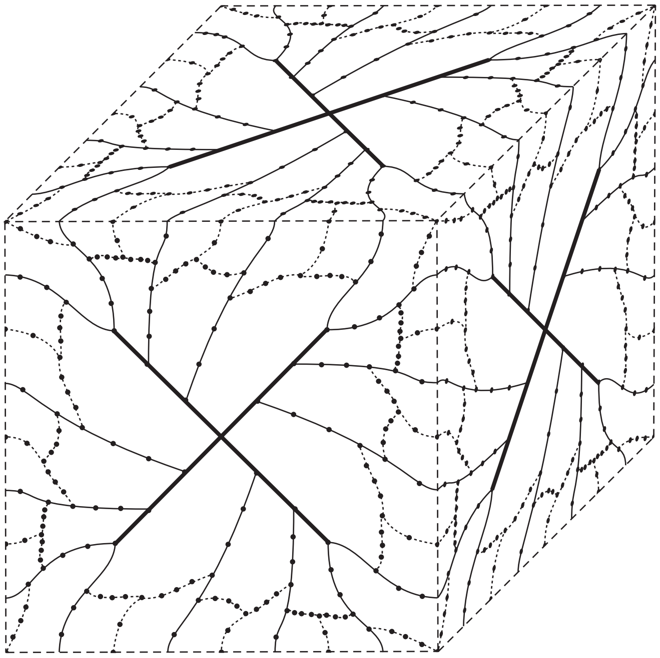

FIGURE 30. $\partial B(1)$.

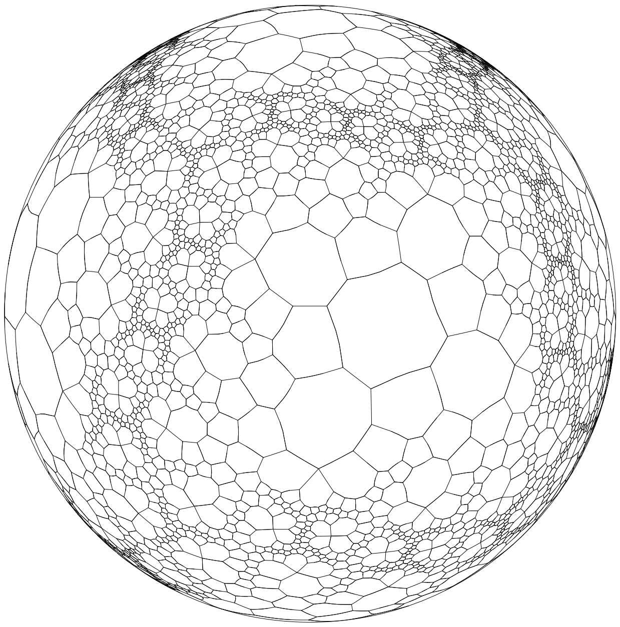

FIGURE 31. $\partial B(\frac{3}{2})$.

Example 7.3. We have seen that the behavior of wild edges is rather pathological. In particular, when we defined cosubdivisions of faces in Section 5, we saw that the branch of a face tree corresponding to a wild edge degenerates to a vertex. We now present an example with wild edges. Again let P be a cube in Euclidean 3-space, labeled as in Figure 21. This

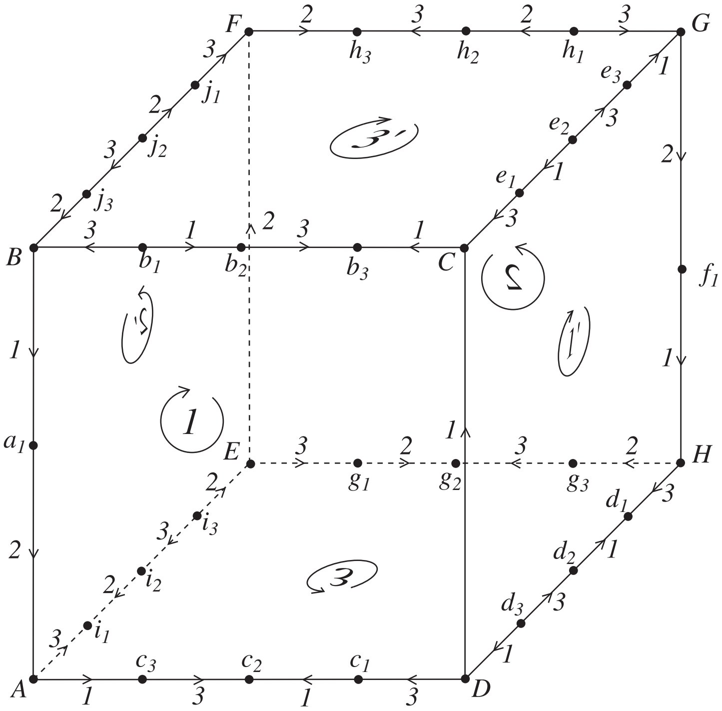

FIGURE 32. The complex Q.

time let the face-pairing be $\epsilon = \{\epsilon_1^{\pm 1}, \epsilon_2^{\pm 1}, \epsilon_3^{\pm 1}\}$, where

$$\epsilon_1 : \begin{pmatrix} A & B & C & D \\ H & G & C & D \end{pmatrix}, \quad \epsilon_2 : \begin{pmatrix} F & E & H & G \\ F & E & A & B \end{pmatrix}, \text{ and } \epsilon_3 : \begin{pmatrix} E & A & D & H \\ F & B & C & G \end{pmatrix}.$$

The edge cycles for ϵ have diagrams

$$CD \xrightarrow{\epsilon_1} CD, \quad EF \xrightarrow{\epsilon_2} EF, \quad AB \xrightarrow{\epsilon_1} HG \xrightarrow{\epsilon_2} AB,$$

$$BC \xrightarrow{\epsilon_1} GC \xrightarrow{\epsilon_3^{-1}} HD \xrightarrow{\epsilon_1^{-1}} AD \xrightarrow{\epsilon_3} BC,$$

and

$$GF \xrightarrow{\epsilon_2} BF \xrightarrow{\epsilon_3^{-1}} AE \xrightarrow{\epsilon_2^{-1}} HE \xrightarrow{\epsilon_3} GF.$$

Again, we choose each multiplier to be 1. The complexes Q and Q^* are shown in Figure 32 and Figure 33. The edges CD and EF of Q (and of Q^*) are wild edges. Two of the six faces of Q are 16-gons and the other four faces of Q are 11-gons.

The boundary of $B(1)$ is shown in Figure 34. The rooted cosubdivision of the front face of Q was shown earlier in Figure 3. According to SnapPea, M is a hyperbolic manifold with volume approximately 8.793.

Example 7.4. At the end of the proof of Theorem 6.1, after Lemma 6.19, there is a discussion of edge cycle edge paths. This last example shows one way in which they occur for ample faceted 3-balls. The model faceted 3-ball P is shown in Figure 35. The face-pairing is defined so that each thick edge is in an edge cycle of length one (face-pairing maps "reflect" across thick edges) and the top is identified with the bottom by translation. For convenience choose each multiplier to be 1. Then each thick edge becomes a wild edge in Q, and every

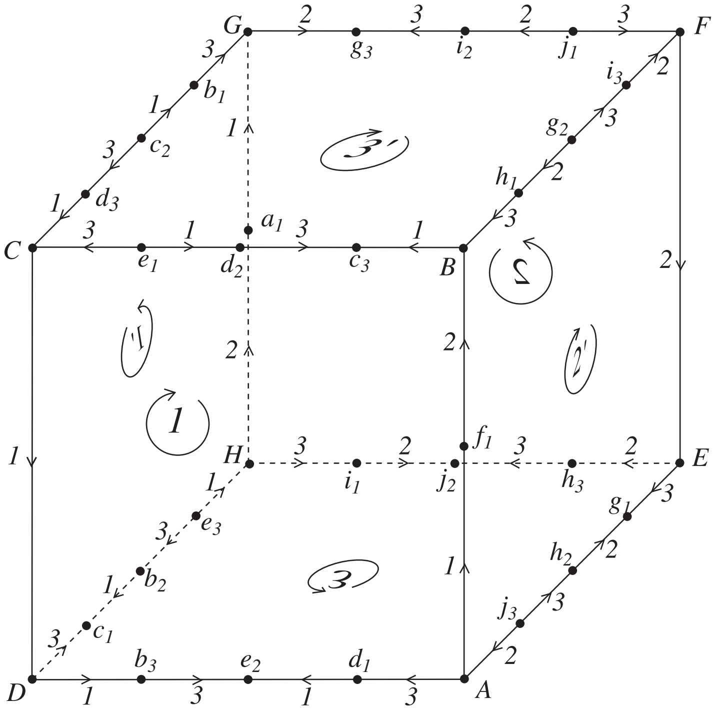

FIGURE 33. The complex Q^*.

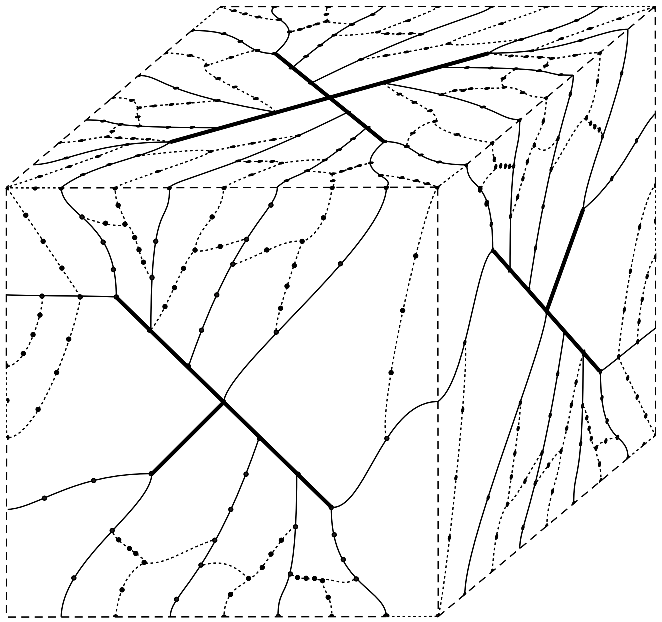

FIGURE 34. $\partial B(1)$.

other edge is properly subdivided in Q. The three horizontal loops other than the top and the bottom are edge cycle edge paths.

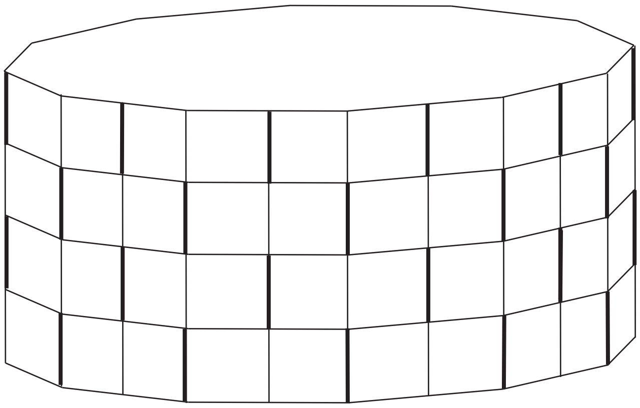

FIGURE 35. The complex P.

REFERENCES

[1] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, Notes on word hyperbolic groups, in *Group Theory from a Geometrical Viewpoint*, edited by E. Ghys, A. Haefliger and A. Verjovsky, (World Scientific, Singapore 1991) 3–63.

[2] M. Bestvina and G. Mess, The boundary of negatively curved groups, J. Amer. Math. Soc. **4,** (1991) 469–481.

[3] J. W. Cannon, Almost convex groups, Geom. Dedicata **22,** (1987) 197–210.

[4] J. W. Cannon, The combinatorial Riemann mapping theorem, Acta Math. **173,** (1994) 155–234.

[5] J. W. Cannon, W. J. Floyd, and W. R. Parry, Squaring rectangles: the finite Riemann mapping theorem, in *The Mathematical Heritage of Wilhelm Magnus–Groups, Geometry and Special Functions*, Contemporary Mathematics **169**, (Amer. Math. Soc., Providence 1994) 133–212.

[6] J. W. Cannon, W. J. Floyd, and W. R. Parry, Sufficiently rich families of planar rings, Ann. Acad. Sci. Fenn. **24,** (1999) 265-304.

[7] J. W. Cannon, W. J. Floyd, and W. R. Parry, Introduction to twisted face-pairings, preprint.

[8] J. W. Cannon, W. J. Floyd, and W. R. Parry, Twisted face-pairing 3-manifolds, preprint.

[9] J. W. Cannon and E. L. Swenson, Recognizing constant curvature discrete groups in dimension 3, Trans. Amer. Math. Soc. **350,** (1998) 809–849.

[10] D. Gabai, Quasi-minimal semi-Euclidean laminations in 3-manifolds, in *Surveys in Differential Geometry, Vol. III (Cambridge, MA, 1996)*, (Int. Press, Boston 1998) 195–242.

[11] M. Gromov, Hyperbolic groups, in *Essays in Group Theory*, edited by S. M. Gersten, M.S.R.I. Publications, (Springer-Verlag, New York 1987) 75–263.

[12] K. Stephenson, CirclePack, computer software, available from http://www.math.utk.edu/~kens.

[13] J. Weeks, SnapPea: A computer program for creating and studying hyperbolic 3-manifolds, available by anonymous ftp from geom.umn.edu/pub/software/snappea/.

DEPARTMENT OF MATHEMATICS, BRIGHAM YOUNG UNIVERSITY, PROVO, UT 84602, U.S.A.
E-mail address: `cannon@math.byu.edu`

DEPARTMENT OF MATHEMATICS, VIRGINIA TECH, BLACKSBURG, VA 24061, U.S.A.
E-mail address: `floyd@math.vt.edu`
URL: `http://www.math.vt.edu/people/floyd`

DEPARTMENT OF MATHEMATICS, EASTERN MICHIGAN UNIVERSITY, YPSILANTI, MI 48197, U.S.A.
E-mail address: `walter.parry@emich.edu`