

Constructing subdivision rules from polyhedra with identifications

BRIAN RUSHTON

Cannon, Swenson and others have proved numerous theorems about subdivision rules associated to hyperbolic groups with a 2–sphere at infinity. However, few explicit examples are known. We construct an explicit finite subdivision rule for many 3–manifolds obtained from polyhedral gluings. The manifolds that satisfy the conditions include all manifolds created from compact right angled hyperbolic polyhedra, as well as many 3–manifolds with toral or hyperbolic boundary.

20F67, 57M50

1 Introduction

A long standing conjecture of Cannon is that every Gromov hyperbolic group with a 2–sphere at infinity acts cocompactly and properly discontinuously by isometries on hyperbolic 3–space [2]. This is true for Gromov hyperbolic groups that are already known to be 3–manifold groups by the Geometrization Conjecture, proved by Perelman [8]. However, it is not at all obvious that hyperbolic groups with a 2–sphere at infinity correspond to any manifolds in the first place, and this is the reason the conjecture remains unsolved.

One approach, adopted by Cannon, Floyd, and Parry, among others, is to study subdivision rules [3]. Subdivision rules are recursive formulas for refining coverings or tilings of a surface. All Gromov hyperbolic groups with a 2–sphere at infinity have a subdivision rule on the sphere in some sense; see Cannon and Swenson [4]. Cannon has shown that, if a subdivision rule for a group gives a conformal sequence of tilings (meaning that tiles do not get too distorted in the long run), then they can be used to put a conformal structure on the sphere that is compatible with the group action [2]. Cannon and Swenson used this fact to prove that a Gromov hyperbolic group acts on hyperbolic space cocompactly, isometrically and proper discontinuously if and only if the subdivision rule for the group gives a conformal sequence of coverings on the space at infinity (see [4, Theorem 2.3.1]). However, it has proven difficult to determine if a subdivision rule is conformal or not.

The coverings created by the subdivision rules Cannon and Swenson originally found for hyperbolic groups [4] are formed of overlapping compact sets. These proved difficult to study, so Cannon, Floyd and Parry proposed a simpler class of subdivision rules: finite subdivision rules [3]. A finite subdivision rule is a way of dividing a 2–dimensional CW–complex (usually the sphere) into a labeled tiling or cell structure, together with a finite recursive formula for dividing each tile into smaller tiles. The most famous examples include barycentric subdivision and hexagonal refinement, where a triangle is chopped up into smaller triangles, which are chopped up into smaller tiles, etc.




Figure 1: The first three subdivisions of a tile type for a hyperbolic orbifold

To remove all ambiguity, we give a concrete definition of a finite subdivision rule, taken from [3].




Figure 2: The fourth subdivision of a tile type for the same hyperbolic orbifold

Definition 1.1 A finite subdivision rule R consists of the following.

- (1) A finite 2-dimensional CW-complex S_R , called the *subdivision complex*, with a fixed cell structure such that S_R is the union of its closed 2-cells. We assume that for each closed 2-cell \bar{s} of S_R there is a CW structure s on a closed 2-disk such that s has at least three vertices, the vertices and edges of s are contained in ∂s , and the characteristic map $\psi_s: s \rightarrow S_R$ which maps onto \bar{s} restricts to a homeomorphism onto each open cell.
- (2) A finite two dimensional CW-complex $R(S_R)$, which is a subdivision of S_R .
- (3) A continuous cellular map $\phi_R: R(S_R) \rightarrow S_R$ called the *subdivision map*, whose restriction to every open cell is a homeomorphism.

Each CW-complex s in the definition above (with its given characteristic map ψ_s) is called a *tile type*. Given a closed 1-cell \bar{e} in S_R , a closed 1-disk e with a characteristic map mapping it onto \bar{e} in S_R is called an *edge type*.

As the final part of our definition, we show how finite subdivision rules can act on surfaces (and 2-complexes in general). An R -complex for a subdivision rule R is a 2-dimensional CW-complex X which is the union of its closed 2-cells, together with a continuous cellular map $f: X \rightarrow S_R$ whose restriction to each open cell is a homeomorphism. We can subdivide X into a complex $R(X)$ by requiring that the induced map $f: R(X) \rightarrow R(S_R)$ restricts to a homeomorphism onto each open cell. The complex $R(X)$ is again an R -complex with map $\phi_R \circ f: R(X) \rightarrow S_R$. By repeating this process, we obtain a sequence of subdivided R -complexes $R^n(X)$ with maps $\phi_R^n \circ f: R^n(X) \rightarrow S_R$. All of the preceding definitions were adapted from [3], which contains several examples.

One difficulty in studying conformal finite subdivision rules is the lack of examples. In [3], Cannon, Floyd, and Parry describe a subdivision rule arising from a hyperbolic 3-orbifold with three tile types: a pentagon, a triangle and a quadrilateral. However, no other examples have been published. On the other hand, explicit finite subdivision rules have been found for all nonsplit, prime alternating links (see the author [10]), most of which are finite volume hyperbolic 3-manifold groups. However, none of these are Gromov hyperbolic.

Our goal in constructing these subdivision rules is to shed light on Cannon's conformality. Currently, the only ways of determining whether a subdivision rule is conformal or not require us to check infinitely many tilings. It seems likely that there is a simple way of telling if a subdivision rule is conformal or not just from the combinatorial structure of the tile types. By creating many examples, we can hope to explain why some tilings are conformal and others are not.

In this paper, we construct an explicit finite subdivision rule for a large class of 3-manifold groups, the majority of which will be hyperbolic 3-manifold groups, some closed, some finite volume, and some of infinite volume. We start by constructing a replacement rule, whose definition will be given in the following section. Essentially, a replacement rule is a subdivision rule in which some cell structure is allowed to disappear at each stage. We then explain how to convert these replacement rules into finite subdivision rules in a large number of cases. As an example of our method, the tilings in Figures 1 and 2 correspond to the same orbifold that Cannon, Floyd and Parry studied in [3]. The tiling in their paper is different from the one we obtain, but strongly related. These tilings are the complexes $R^1(s)$, $R^2(s)$, $R^3(s)$, and $R^4(s)$, where s is the pentagonal tile type in the subdivision rule R (each tile type s is an R -complex since it is equipped with an appropriate map into S_R).

2 Preliminaries

To associate a subdivision rule to a manifold (possibly nonhyperbolic), we create a spherical space at infinity, similar to that of word-hyperbolic groups. To do so, we approximate spheres in \mathbb{R}^3 by taking the boundary of polyhedra built up from fundamental domains. More specifically, we let $B(0)$ be a single fundamental domain, and let $S(0)$ be its boundary graph. The structure of $S(0)$ gives us a tiling of S^2 . Now, let $B(1)$ be formed from $B(0)$ by attaching polyhedra to all its exposed faces, and let $S(1)$ be its boundary, and so on. As long as $S(n)$ is always a topological sphere, this defines a sequence of tilings of the sphere, which defines a combinatorial structure on the space at infinity. For the manifolds under consideration in this paper, we will show that $S(n)$ must be a topological sphere for all n .

We now define the various terms we use in this paper.

Definition 2.1 A *replacement rule* is a finite subdivision rule coupled with a *local combination rule*. In a local combination rule, a labeled tiling $T(n)$ of a surface is made more coarse by removing some edges and vertices in a local manner. More specifically, the set of tiles in $T(n)$ is partitioned into a number of closed sets F_i with disjoint interiors which are topological closed disks, each the union of closed labeled tiles. Each F_i is replaced with a single closed labeled tile whose boundary is that of F_i . The sets F_i may be of size one (meaning that the tile is not combined with any other tile). The partition and the replacement rule are completely determined by the labels of the tiles and of their neighbors of distance less than N for a fixed constant N (here, the distance between two tiles T, T' is given by the shortest length of a chain of closed tiles $T_1 = T, T_2, \dots, T_k = T'$ such that $\bigcup T_i$ is connected). Thus, if there are two tiles in the surface with sufficiently large neighborhoods A, A' that are isomorphic (meaning they are cellularly isomorphic and the isomorphism preserves labels), there are two smaller closed neighborhoods $B \subseteq A, B' \subseteq A'$ that are isomorphic in the post combination tiling.

After each such round of combination, we subdivide the new labeled tiling according to a finite subdivision rule, and then proceed with another round of combination.

Definition 2.2 We will, for convenience, define a *polyhedron* as a topological ball with a cell structure on the boundary, where all vertices have valence three or more, all faces have at least three edges, and any two faces that intersect do so in a single edge or a vertex.

Definition 2.3 A *fan* is a chain a_1, a_2, \dots, a_n of faces in a polyhedron surrounding a single vertex v , where $a_i \cap a_{i+1}$ contains an edge coming out of v for $1 \leq i \leq n-1$.

If $a_1 \cap a_n$ does not contain an edge coming out of v , then a_1 and a_n are called the ends of the fan. If $a_1 \cap a_n$ does contain such an edge, then the fan a_1, \dots, a_n is the star of the vertex. We consider a single face to be a fan of size 1. See Figure 3 for examples of fans.




Figure 3: The shaded portions in (a), (b) and (c) are all examples of fans, while the shaded portion in (d) is not.

Definition 2.4 A polyhedron is *spread out* if, given any face A and any fan B not containing A , $A \cap B$ is a vertex, one edge, or the union of two contiguous edges.

Among the platonic solids, only the cube and the dodecahedron are spread out. In fact, no polyhedron with a triangular face T is spread out. To see this, let v be a vertex of T , and let A be the fan consisting of all faces containing v except for T . Let e be the edge of T not containing v , and let B be the other face containing e . Then $A \cap B$ contains both vertices of e without containing e itself. For our polyhedron to be spread out, $A \cap B$ would have to be two contiguous edges, with B a triangle, and it is not hard to see that $A \cup T \cup B$ would have to form a tetrahedron. A tetrahedron is not spread out, because the star of any vertex intersects the only remaining face in three edges. We thank the referee for this argument.

Manifolds constructed from spread out polyhedra are analogous to alternating links. Alternating links have a well-known checkerboard polyhedral decomposition, where

the boundary is an alternating diagram for the link. All hyperbolic alternating links have an alternating diagram that is reduced, nonsplit, and prime, and these properties are very similar to the spread out condition above.

Andreev's theorem [1] (later extended by Rivin and others [7]) implies that all convex hyperbolic polyhedra with dihedral angles equal to $\pi/2$ at every edge are spread out. In particular, such polyhedra cannot have any triangular faces. In fact, they cannot even have square faces.

We will need restrictions on edge cycle lengths.

Definition 2.5 When polyhedra are glued together, the edges are identified under the gluing maps. For a given edge e , the *edge cycle length* of e is the size of the equivalence class containing e .

In the universal cover of a manifold, the edge cycle length turns out to be the number of distinct fundamental domains surrounding lifts of the edge. Thus, if a polyhedral fundamental domain for a hyperbolic 3-manifold has dihedral angle $2\pi/n$ at an edge e , the edge cycle length of e is n . We are now ready to state the theorem.

3 Existence of replacement rules

Theorem 3.1 Let M be a manifold (possibly with boundary) that can be decomposed into polyhedra P_1, P_2, \dots, P_n . If each polyhedron is spread out, and all edge cycles have even length greater than or equal to 4, there exists a replacement rule for M . In this replacement rule, each new polyhedron is placed on exactly one fan in the previous stage of the replacement rule.

Note The fans form the partition F_i required in Definition 2.1.

Proof Let $S(n)$ represent the tiling given by gluing n layers of polyhedra to a base polyhedron. Note that in $S(0)$, all exposed faces are fans consisting of single faces. Also, all edges border fans of size 1.

As we glue on more and more polyhedra, each edge will be buried. This happens when we glue a number of polyhedra adjacent to it equal to its cycle length, minus one. All edges in $S(0)$ are adjacent to only one polyhedron; we say that these are *unburdened* edges. As we place polyhedra on the faces to either side, we get closer to closing up the cycle of polyhedra around the edge; we say this edge is now a *burdened* edge. More precisely, a burdened edge is any edge adjacent to more than one polyhedron. Because

the edge cycle has even length and all edges start out adjacent to a single polyhedron, when we glue a polyhedron on either side of the edge at each stage, it will eventually only have one polyhedron remaining to glue on. If we continued to glue polyhedra on either side, we would exceed the cycle length; thus, we glue a single polyhedron onto both faces neighboring this edge. Because this gluing behavior is different from that of previous stages, we say the edge is a *loaded edge*. That is, a loaded edge needs only one polyhedron to complete its cycle, and we focus on this type of edge in our proof. As we shall see, all edges start out as unburdened, become burdened, become loaded, and then disappear as we glue a polyhedron over them. Every loaded edge is burdened (because it has more than one adjacent polyhedron). We will call faces with at least one loaded edge *loaded faces*.

Our goal is to show that, at every stage, every polyhedron is glued onto a single face or a larger fan. Because there are finitely many combinatorial types of fans, this will show that there is a recursive combinatorial way of constructing $S(n+1)$ from $S(n)$. Now, when a polyhedron is glued onto a face in $S(n)$, if that face has any loaded edges, the polyhedron will be glued onto those loaded edges and onto all faces on the other side of those loaded edges; and if those faces have loaded edges, the polyhedron is glued onto those as well, and so on. To show that all polyhedra are glued onto fans and only fans, we have to show that faces with loaded edges meet up as fans.

What do fans look like? Locally, at the vertex, they look like pizzas with some number of neighboring pieces missing (so that what remains is the sector of a circle). Each piece of pizza touches exactly two neighboring pieces, except for the ends of the pizza, which touch exactly one other piece of the fan.

Therefore, if a face has exactly two loaded edges that share a vertex, we call it a *loaded wedge*. Loaded wedges may group together into fans, where the ends of the fan (if there are any) have only one loaded edge. In this case, the ends are called *half-loaded wedges* and the entire fan is a *loaded fan*. Our goal, then, is to show that all loaded faces can be grouped into loaded fans, and the first step in proving that is to prove that all loaded faces are loaded and half-loaded wedges. See Figure 4 for an example.

As for the second step, note loaded wedges *may* group together into fans. Loaded wedges are, again, like pizza slices. If we rotate a pizza slice 180° before putting it down, we get an unusual parallelogram of pizza instead of a sector of a circle. This shows that it is not sufficient to prove that all loaded faces are loaded or half-loaded wedges; we must also prove that they meet up in the correct way. An edge is *skew* if:

- (1) there are two loaded edges that intersect it, one at each vertex,
- (2) each loaded edge borders a face that the skew edge borders, but




Figure 4: If the shaded portion is a loaded fan, then faces b , c and d will be examples of loaded wedges, while faces a and e will be examples of half-loaded wedges. Dashed edges are loaded and others are not.

(3) these loaded edges do not border the same face.

See Figure 5 for an example of a skew edge.




Figure 5: A skew edge e . Here f_1 and f_2 are loaded, intersect e in distinct vertices and each border a distinct face that e also borders.

Skew edges represent the rotated pizza slices that meet up in the wrong way. The rest of this proof will be devoted to showing that all loaded faces are loaded or half-loaded wedges, and that there are no skew edges. This will show that all loaded faces group together into fans; thus, every polyhedron that is glued on to $S(n)$ is glued on to a loaded fan, and since there are only finitely many fans on each polyhedron, this implies that there are only finitely many combinatorial types of gluings.

We proceed by induction. In $S(0)$, all edges are unloaded, and so there are no skew edges or loaded faces. Assume that $S(n)$ has no skew edges and that all faces with loaded edges are loaded or half-loaded wedges; this implies that all faces are grouped into fans. Then we form $S(n+1)$ by placing polyhedra onto fans. A fan A in $S(n)$ will be replaced by several subtiles in $S(n+1)$ corresponding to unglued faces of the polyhedron. Those in the interior correspond to faces on the polyhedron that do not touch the fan A , and are formed of unburdened edges. Those touching a former edge of A correspond to faces on the polyhedron that share an edge with the fan A (where we are now considering A as a subset of the polyhedron). Because our polyhedra are spread out (recall Definition 2.4), these subtiles share either one edge or two contiguous edges with the fan. These edges are the only edges that are burdened, because they are the only edges of the polyhedron we are gluing on that will be adjacent to more than one polyhedron.

So, all faces in $S(n+1)$ have zero, one, or two contiguous burdened edges. This implies that if these faces have any loaded edges (which are a kind of burdened edge), the faces must be either loaded or half-loaded wedges. Thus, we have completed half of our induction.

Now, we must show that there are no skew edges. We will first show that there are no burdened skew edges.

If there is a burdened skew edge e in $S(n+1)$, then by definition of skew edge (recall Figure 5), the faces to either side each have two contiguous burdened edges. As we discussed earlier, faces with burdened edges in $S(n+1)$ correspond to faces in the new polyhedra that are adjacent to the fan we are gluing on to. Thus, if a face has two burdened edges, it must have been adjacent to a fan in two edges. Now, if a fan consists of a single face, then no other face can intersect it in two edges, by our definition of polyhedron (see Definition 2.2).

Thus, if a face in $S(n)$ has two burdened edges, it comes from a polyhedron glued onto a loaded fan in $S(n-1)$. Also, the two burdened edges in $S(n)$ must have touched two different faces of the loaded fan in $S(n-1)$, again by Definition 2.2. Moreover, the vertex they share must also have been the center vertex of the loaded fan in $S(n-1)$.

So, a burdened skew edge e in $S(n)$ must come from an edge in $S(n-1)$ that is adjacent to two loaded fans, and the loaded edges of those loaded fans intersect e in distinct vertices. So e must have been a skew edge (burdened or unburdened) in $S(n-1)$, which is impossible by our induction hypothesis. See Figure 6.

We now show (by considering a possible counterexample) that there are no unburdened skew edges in $S(n+1)$, which completes the proof. Assume that there is an unburdened




Figure 6: Burdened skew edges must come from other skew edges.

skew edge e . Because e is unburdened, it must be in the interior of some polyhedron that was glued on to a fan A in $S(n + 1)$. See Figure 7. Thus, the faces on either side (call them B and C) must also be part of the same polyhedron, and both border the fan A (considered now as a subset of the polyhedron being glued on). But then B intersects A in an edge contiguous with e and in both vertices of the interior edge e . Because the intersection of A and B is connected and is one or two edges, B must be a bigon (which cannot happen) or a triangle. The same applies for C . Thus, our polyhedra consists of a fan (namely, A) and two triangles sharing an edge (namely, B and C). There can be at most four faces in the fan, as each must intersect something not in the fan (because at most two edges of each face in a fan touch other faces in the fan, and each face is a triangle or larger), and there are only four edges of $B \cup C$ intersecting the fan. If the fan is not the star of a vertex, we can add B or C to it, and this larger fan intersects the remaining triangle in all of its edges, a contradiction. So the fan must be the star of a vertex of valence 3 or 4. If it is of valence 3, then the star of a vertex of e contains 4 tiles, and the remaining tile must intersect this star in more than 2 edges, which cannot occur. If it is of valence 4, then we can pick one vertex of e and consider the fan around it consisting of B , C , and one tile of A . Then one of the tiles of A touching the other vertex of e has disconnected intersection with this fan (see Figure 8). This is a contradiction. Thus, there are no skew edges in $S(n + 1)$.

Because there are no skew edges, all faces group together in fans. Since there are finitely many types of fans, we get a finite recursive algorithm for constructing a sequence of tilings, ie a replacement rule. We claim that the underlying set consisting of unions of images of fundamental polyhedra is in fact the universal cover; since $B(n)$ is a topological ball for every n , $\bigcup B(n)$ is simply connected. Also, $B(n)$ is locally homeomorphic to our manifold exactly because our tiling required the gluing combinatorics at each face, edge, and vertex to be the same as in the ball. Thus, we have created a local replacement rule for recursively describing the universal cover. The label of a tile




Figure 7: An unburdened skew edge must look something like this.




Figure 8: Note that the fan consisting of the shaded gray region has disconnected intersection with the face D .

in this replacement rule (as mentioned in Definition 2.1) is determined by the face of the polyhedra corresponding to the tile and by the particular fan containing the tile. \square

4 Forming subdivision rules from replacement rules

Note that Theorem 3.1 does not provide a way for creating a subdivision rule. However, in many cases, a subdivision rule can be created by adding extra lines. To illustrate, consider a nonhyperbolic example, the 3–dimensional torus. It has a decomposition into a single polyhedron, a cube with edge cycle lengths all equal to 4, and so has a replacement rule as described by the theorem. The three replacement types are shown in Figures 9, 10 and 11. They correspond to an unloaded face, a *loaded pair* of faces (ie a fan of two faces, sharing a common loaded edge), and a *loaded star* (ie a maximal loaded fan, containing the entire star of a vertex). By symmetry, these are the only types.

Notice that this is only a replacement rule; edges are created, disappear, reappear, etc. This can be turned into a subdivision rule by a method that turns out to be very useful:




Figure 9: The replacement rule for an unloaded face




Figure 10: The replacement rule for a loaded pair




Figure 11: The replacement rule for a loaded fan

adding new edges at every stage. For instance, in Figure 10, the center line between two squares disappears when we glue on the new cube. However, if we add a line to the new cell structure (as shown in the top half of Figure 12), then the new cells structure is a refinement of the old cell structure, and we can interpret this refinement as a subdivision.




Figure 12: We can add lines to the two “loaded” tile types to get a self-consistent subdivision rule.

However, adding this line divides the top and bottom squares into two triangles each; each of these are part of a loaded star, so we have to change the replacement rule for a loaded star; however, notice that adding the lines in to the loaded star on the left of Figure 11 gives us a hexagon divided into six “pie slices.” If we add similar lines bisecting the replacement of the loaded star (as seen in Figure 12), we again get a hexagon divided into six triangles; thus, we can take the subdivision rule on this tile type to be the identity (meaning that the labels we gave it prior to this stage need never change, as the cell structure will never change).

We summarize this in Figure 13. Several stages of subdivision are shown in Figure 31. Note that, in the circle packed pictures, the vertices and edges move; this subdivision rule is not conformal, and so circle packings will distort the images over time. The connection between circle packings and conformality is explained by Haïssinsky in [6].

In general, we can get a subdivision rule if we can add finitely many lines to each replacement type in a way that is self-consistent. However, self-consistency can be difficult to achieve; notice that adding lines in the loaded pair forced us to add lines in the loaded star. In more complicated subdivision rules, adding lines in one local




Figure 13: The subdivision rule for the 3–dimensional torus. This can be made into CW–complex S_R mentioned in the definition of subdivision rules by identifying all vertices and all edges of the same type (there are 3 types of edges as well; can you identify them?) We can then subdivide S_R into $R(S_R)$ according to the above diagram and create the subdivision map by sending each tile of $R(S_R)$ into S_R according to its label.

replacement type can cascade and force us to add more and more lines until we have infinitely many replacement types, and thus infinitely many subdivision tile types.

There are large classes of polyhedra that do have subdivision rules, however. In the discussion that follows, we will mention the edge cycle lengths of polyhedra many times. Edge cycle length in a hyperbolic manifold can correspond to dihedral angles of polyhedral decompositions. If every dihedral angle is $2\pi/n$, then all edge cycles have length n . Polyhedral decompositions with varying dihedral angles do not follow this pattern.

It will be useful to define n –cycles, following [7]. An n –cycle is a set of n faces A_1, A_2, \dots, A_n in a polyhedron such that $A_i \cap A_{i+1}$ (subscripts taken mod n) is a single edge for each i , and all such edges in the n –cycle are pairwise disjoint.

The following two theorems give two large classes of polyhedral gluings that can be made into subdivision rules. The first always gives closed 3–manifolds; the second always gives 3–manifolds with boundary.

In the proof that follows, a *loaded vertex* will be the center vertex of a loaded star.

Theorem 4.1 *Assume a manifold M is formed by gluing together convex, right angled hyperbolic polyhedra P_1, \dots, P_n . Then the replacement rules in Theorem 3.1 can be made into finite subdivision rules by adding finitely many lines to replacement tile types. More explicitly, we can create a finite subdivision rule R such that $S(0)$ is an R -complex, together with a map g from $R^n(S(0))$ onto $S(n)$ such that the preimage of every tile in $S(n)$ is a finite union of tiles in $R^n(S(0))$.*

Note that all convex, right angled hyperbolic polyhedra satisfy [7]:

- (1) Each polyhedron P_i is spread out.
- (2) Each vertex of each polyhedron has valence three.
- (3) Every edge (after gluing) has edge-cycle length four.
- (4) There are no three-cycles or four-cycles.

Proof We show that we can refine the series of tilings $S(n)$ into finer tilings $S'(n)$ in such a way that the natural map from $S'(n)$ to $S'(n+1)$ can be interpreted as a subdivision with only finitely many patterns of subdivision.

Because every vertex has valence three, any faces that share a vertex also share an edge, which simplifies the combinatorics significantly (see Figure 14). Throughout the following proof, we will refer to the replacement rule for the three torus as an example. In Figure 15, we give an example of a more complicated replacement rule, that of the right angled dodecahedral orbifold mentioned earlier. We will use this replacement rule as an example as well. We proceed by cases.




Figure 14: Each new vertex starts with three unburdened lines. In the next stage of replacement, three unburdened lines are added and the old edges become loaded. The vertex is surrounded by three loaded pairs. In the next stage there are only three edges and the vertex is now loaded. In the next stage the vertex and its star are entirely covered up by a new polyhedron.




Figure 15: The three replacement types for the right angled dodecahedral orbifold. They are fans of size one, two and three.

Case I Replacing a single unloaded tile in $S(n)$ The replacement rule for an unloaded tile is already a subdivision, as no edges disappear. In $S(n + 1)$, the only tile types created by this replacement are more unloaded tiles from the interior (which are also Case I) and loaded pairs from the boundary (which are Case II).

Case II Replacing a loaded pair in $S(n)$ Recall that, in creating $S(n + 1)$, we glue one polyhedron onto the entire loaded pair, which consists of two faces A and B sharing a common loaded edge E that travels between vertices v and w (see Figure 16). The cell structure of the unglued faces of the polyhedron (which we call P^*) replaces the cell structure of $A \cup B$. To make this replacement a subdivision rule, we find

or create a path from v to w in P^* that will take the place of E (as in Figures 12 and 17). Each of these two vertices is contained in a unique tile of $P^* \subset S(n + 1)$ (recall Figure 14). We can add an edge from each vertex to an interior vertex of this tile. If the two new edges meet in the same point, we can stop, and identify these two edges with the original, loaded edge E .




Figure 16: A loaded pair




Figure 17: We replace the edge E from Figure 16 with the path in bold. Note that the initial and terminal segments were not originally in P^* .

If they meet in different points, then we connect those points by a path consisting entirely of interior edges. If no such path exists, then the set of interior edges is disconnected, which implies that there must be a tile that separates the two points (ie the complement of the separating tile is disconnected, with each component of the complement containing interior edges of the separating tile). Because the valence is three, that tile must intersect the loaded pair $A \cup B$ in disjoint edges. However, this is impossible by the spread out condition, recall Definition 2.4.

Thus, we can find a path α in $P^* \subset S(n + 1)$ that we can identify with the loaded edge $E \subset A \cup B \subset S(n)$. This makes the replacement rule for a loaded pair into a subdivision rule for each half of the loaded pair, just as in Figures 12 and 17.

Note that, in finding a replacement for E in P^* , we only changed the cell structure near v and w , by adding one extra line to the tiles containing them; everywhere else, we used the existing cell structure of $P^* \subset S(n+1)$. Each tile containing v or w in $S(n+1)$ is part of a loaded star that will be replaced by a single polyhedron in $S(n+2)$ (recall Figure 14). Since we treated each loaded pair the same way, we see that every loaded star in $S(n+1)$ consists of three tiles around a vertex, each tile sharing a loaded edge with each neighbor and each tile having an added line from the center vertex to an outside, unburdened vertex (as shown in Figures 13 and 18). This is Case III. Note that all other tiles in $P^* \subset S(n+1)$ correspond to Case I or Case II.




Figure 18: All loaded stars have three loaded lines and three added lines from the previous stage.

Case III Replacing a loaded star in $S(n)$ with three added radial lines ending at boundary vertices To find a subdivision rule for these loaded stars, we must find or introduce all six of these edges in the new cell structure (which we again call $P^* \subset S(n+1)$). The method we will use is complicated, but works for all right angled hyperbolic 3-manifolds (it will not work for the three torus example we did earlier). Call the three loaded edges E_1, E_2 and E_3 , and call the added edges A_1, A_2 , and A_3 . Call the center vertex V .

We start by finding lines in P^* that represent the three loaded lines E_1, E_2 , and E_3 . Each of these lines ends at a vertex v_1, v_2 and v_3 on the boundary of the star in $S(n)$. As we glue a new polyhedron on to form $S(n+1)$, none of these vertices will have new interior edges coming off of them (they are loaded vertices; see Figure 14). So, we will add edges to P^* to make a path α in $S(n+1)$ that connects v_1 and v_2 . We begin such a path by adding edges α_1, α_2 from v_1 and v_2 to the midpoints of interior edges e_1, e_2 (we abuse terminology by referring to an arbitrary interior point of an

edge as the “midpoint” of an edge, to avoid the overuse of the word interior). We then connect the endpoints of α_1, α_2 by a series of added edges $\alpha_3, \dots, \alpha_n$, each of which goes from the midpoint of one interior edge e_i in P^* to the midpoint of another, interior edge e_j in P^* . We define the path α to be $\bigcup \alpha_i$.

Note that we can assume that the e_i are disjoint, because if they are not disjoint, we can shorten the path α (see Figure 19). Similarly, we can assume that no tile contains more than one α_i . Finally, we can assume that α does not intersect any boundary tiles (including their interior edges) except for the two tiles containing α_1 and α_2 . We can do this because we claim that the complement of the boundary tiles in P^* is connected. To prove this claim, note that the only way the complement could be disconnected would be for two boundary tiles to intersect in an interior edge (note that because the valence of every vertex is three, any edges that intersect in a vertex intersect in an edge containing that vertex). However, this would imply that our polyhedron contains a four cycle (four faces which intersect cyclically in four disjoint edges), which is impossible by Andreev’s theorem [1]; see Figure 20. This is where the proof fails for nonhyperbolic manifolds such as the 3–torus.




Figure 19: A path that goes through nondisjoint edges can be shortened.




Figure 20: If the complement of the boundary tiles is disconnected, there must be a four cycle.

Thus, as said above, we can connect v_1 to v_2 by a path α which is formed from new edges, each joining two disjoint edges in a tile, except for the initial and terminal segments, with no tile containing two new added edges. Now, we form another path β with the same properties, but now connecting v_3 to α . Again, we start the path by

adding an edge from v_3 to an interior edge, and then continue the path by adding line segments containing disjoint edges until we reach an edge e of an interior tile T containing α . The terminal segment connects e to the midpoint of the segment of α contained in T ; see Figure 21.




Figure 21: We find paths α and β to replace the loaded edges of the loaded star.

Thus, we can add $\alpha \cup \beta$ to P^* and identify $\alpha \cup \beta \subseteq S(n+1)$ with $E_1 \cup E_2 \cup E_3 \subset S(n)$. We identify $\alpha \cap \beta$ with the vertex $V = E_1 \cap E_2 \cap E_3$. We have now replaced all loaded edges.

Now, we replace A_1, A_2 , and A_3 , which, as we recall, represent the lines added in Case II. They intersect the boundary of our star in $S(n)$ in vertices w_1, w_2 , and w_3 , respectively. Each vertex w_i lies in a distinct component C_i of $P^* \setminus \{\alpha \cup \beta \cup T^\circ\}$, where T° is the interior of the tile containing $\alpha \cap \beta$. We claim that the set of interior edges of each C_i is connected. To see this, note that any gap in the interior edges (ie having more than one connected component of edges) would be caused by a single tile of P^* being crossed by $\alpha \cup \beta$ in two or more disjoint segments. However, by construction, this never occurs, except in T . Thus, the set of interior edges in each C_i is connected.

Also, each component C_i must contain a vertex t_i of T , since β does not enter T by the same edge as α . Thus, we can find a path γ_i from each w_i to a vertex t_i of T . Extend each γ_i by an edge from t_i to V in T ; see Figure 22. We now identify each γ_i with the appropriate A_i . Thus, by adding $\alpha, \beta, \gamma_1, \gamma_2$, and γ_3 to P^* , we can embed the cell structure of the loaded star in $S(n)$ into the cell structure of the replacement in $S(n+1)$.

However, this creates three new tile types: unloaded tiles with a line added connecting midpoints of disjoint edges, loaded stars with three added radial lines that may now end at midpoints of edges instead of vertices and an unloaded tile with six added edges




Figure 22: We find paths γ_1, γ_2 and γ_3 to replace the three added lines of the loaded star.

(three intersecting the boundary in vertices, three intersecting in midpoints of edges, the two types alternating around the boundary). These are Cases IV, V and VI, respectively.

Case IV Replacing a single unloaded tile with an added line connecting midpoints of edges in $S(n)$ Call the added line $A \subseteq S(n)$. Let e_1 and e_2 be the two boundary edges connected by A . Then in $S(n + 1)$, we can add an edge going from the midpoint of each e_i to an interior vertex. We can then connect these two new edges by a path of interior edges. This creates a new case, Case VII, which consists of a loaded pair with one or two added lines going from a boundary vertex to the midpoint of the loaded edge of the loaded pair; see Figure 23.




Figure 23: Replacing a single unloaded tile with an added line from midpoint to midpoint

Case V Replacing a loaded star in $S(n)$ with three added radial lines, possibly terminating at midpoints of edges Loaded stars of this type will be replaced almost exactly as other loaded stars were earlier; in particular, there are still three loaded edges and three added edges. We still find paths α and β to replace all added edges. However, these paths cannot consist entirely of interior edges of P^* , as the midpoints

of boundary edges are not connected to interior edges. Thus, we begin α and β by adding lines from each of these boundary midpoints to an interior vertex. We then extend α and β as needed by interior edges, and the proof goes through in exactly the same way. This creates no new tile types; see Figure 24.




Figure 24: Replacing a loaded star with added edges going to midpoints instead of boundary vertices

Case VI Replacing a single unloaded tile with six added radial lines (every other one going to a boundary vertex, the rest to midpoints of edges) in $S(n)$ In this case, we are basically doing Case III for a single tile instead of a loaded star. The complement of the boundary in $S(n + 1)$ is still connected, because we would have a three-cycle if it were not. So let A_1, A_2 and A_3 be the added lines ending at vertices v_1, v_2 and v_3 and let E_1, E_2 and E_3 be the added lines ending at edges e_1, e_2 and e_3 .

As a modification of Case III, we can connect the midpoints of e_1 and e_2 with a path α of added edges connecting disjoint edges of nonboundary tiles (except for the initial and terminal segment). We can similarly construct β connecting the midpoint of e_3 and α . After this, we can construct γ_1, γ_2 and γ_3 in exactly the same way as Case III and get a subdivision. This creates more tiles of Case IV, VI, and VII.

Case VII Replacing a loaded pair with two edges coming from the boundary to the midpoint of the loaded edge of the loaded pair in $S(n)$ First, assume that the two added edges A_1, A_2 meeting in the midpoint of the loaded edge E have their other endpoints at two boundary vertices w_1 and w_2 . Call the loaded edge E .

Instead of replacing E as in Case II, we replace it by a path α as in Case III that travels from the endpoints of E and only goes through midpoints of interior edges. This creates more tiles of Case V and Case IV. Then connect w_1 and w_2 by a path γ of interior edges. This creates no new cases, and $\alpha \cup \gamma$ replaces E and A_1, A_2 .




Figure 25: Replacing a single unloaded tile with six added lines

The case when A_1 and/or A_2 have their other endpoints at midpoints of boundary edges is similar; however, we begin and end γ with segments attaching the midpoints of the boundary edge to an arbitrary interior vertex; the rest of gamma consists of interior edges. This creates no new cases.




Figure 26: Replacing a loaded pair with added lines starting at the boundary and meeting in the midpoint of the loaded edge

So, each of our seven cases of tile types generates more tiles, but only within those seven cases. However, each tile type is a slight modification of a loaded fan, of which there are finitely many. There are also only finitely many ways to modify each of these finitely many fans (by adding lines). Thus, there are finitely many tile types.

Thus, by adding finitely many lines to the replacement tile types, we can turn the replacement rule of Theorem 3.1 for such a manifold into a subdivision rule with finitely many tile types. We can make the CW-complex and subdivision map explicit in the same way we did the torus; see Figure 13. The last statement of the theorem follows because only finitely many lines were added to each fan in $S(n)$. \square

The polyhedral gluings in this theorem create only closed 3-manifolds (since the link of each vertex is a sphere). If, in some spread out polyhedral gluing with even edge cycle lengths, the valence of each polyhedron is uniformly greater than three or the edge

cycle lengths are all six or larger, we can also find a subdivision rule. However, these manifolds are never closed manifolds, as they have torus or hyperbolic surface boundary. The link of each vertex in the manifold is composed of polygons with a number of edges equal to the valence of the corresponding vertex on the polyhedron; each vertex of these polygons corresponds to an edge in the manifold, and the valence of the vertex is the edge cycle length of edge. Under these circumstances, (still assuming large valence or large edge cycle length), the author ensures that only loaded pairs and unloaded tiles occur in the replacement rules for the manifold of [9, Chapter 9]. In particular, we have the following:

Theorem 4.2 *Let M be a 3-manifold created by gluing together polyhedra P_1, \dots, P_n such that each edge cycle has even length greater than 2. If each P_i is spread out and satisfies*

- (1) *all vertices have valence greater than 3,*
- (2) *all edges have cycle length greater than 4,*

then the replacement rules in Theorem 3.1 can be made into subdivision rules by adding finitely many lines to replacement tile types. In fact, the polyhedra need not be completely spread out; we need only check the condition for fans of size two.

Note Octahedra and icosahedra satisfy this weaker spread out condition.

Proof Recall that the only tile types in the replacement rule from Theorem 3.1 are loaded fans.

The link of every vertex in our manifold is a Euclidean or hyperbolic surface, by an Euler characteristic argument. As we construct the universal cover for our manifold, we also construct, at each vertex, the universal cover for the link. Each vertex of the link corresponds to an edge of the manifold, and each edge of the link corresponds to a face of the manifold. These surfaces have replacement rules as well, with “loaded vertices” corresponding to loaded edges in the manifold. Such replacement rules were studied in [9, Chapter 9]. According to the proof of Theorem 12 of that chapter, all surfaces of sufficient size (including those coming from manifolds satisfying the conditions of this theorem) have a replacement rule in which only isolated loaded vertices occur (ie no two loaded vertices are ever adjacent). Since loaded vertices in the surface correspond to loaded edges in the manifold, this shows that every loaded edge must be bordered by unloaded edges to either side. This shows that only loaded pairs and unloaded tiles can occur.

Thus, if we can replace the loaded edge in every loaded pair, we will be done. However, the proof of [9, Theorem 12] also shows that every loaded vertex in the universal cover of the surface is replaced by one or more unloaded vertices. This implies that, in the universal cover of the manifold, the two vertices of each loaded edge in $S(n)$ will both have new interior edges coming from them in $S(n + 1)$. We can construct a path in $S(n + 1)$ between these two vertices consisting entirely of interior edges, because the only obstacle to such a path would be a tile that separates the loaded pair, which cannot occur by the spread out condition (Definition 2.4). No new tile types are created by finding such a path, and so we have a subdivision rule. \square

As you can see, it is actually easier to find a subdivision rule for 3-manifolds with boundary. Subdivision rules for alternating links (which have valence four, edge cycle length four polyhedral gluings) were found previously in [10]. Also, note that the full “spread out” condition is not needed; we need only check that no tile has disconnected intersection with single tiles or loaded pairs, or intersects in more than two edges.




Figure 27: A subdivision rule obtained from the replacement rule in Figure 15. We use a slightly different method than the one outlined in Theorem 4.1 to reduce the number of tile types.

The tilings in Figures 28–33 and Figures 1–2 were all created by the methods detailed in Theorem 4.1 and Theorem 4.2. The pictures were created using Stephenson’s Circlepack [11] and software by Floyd [5].




Figure 28: The first few subdivisions of the Borromean rings complement, which decomposes into two octahedra with edge cycle length 4.

Figure 28 and Figure 29 represent the Borromean rings, a manifold created from two octahedra, each with edge cycle lengths of 4. It is a finite volume hyperbolic manifold. Figure 30 is another finite volume hyperbolic manifold: the figure-eight knot.

Figures 31–33 are all created from cubes. The first figure is the 3–dimensional torus, with edge cycle lengths of 4. It is a Euclidean manifold. The last two figures are a cube with edge cycle lengths of six. It has torus boundary.

Finally, Figures 1 and 2 show the tilings from a subdivision rule created from the dodecahedral orbifold first studied by Cannon, Floyd and Parry by an adaptation of




Figure 29: One more subdivision of the Borromean rings complement

our method. This is a dodecahedron with edge cycle length four. It is interesting to compare this to their original subdivision rule [3]. The tiles are shown in Figure 27.

Closed hyperbolic manifolds with even edge cycle length are easy to create: Start with a set of glued together polyhedra with large valence at each vertex and even edge cycle length greater than 2. These manifolds will usually have hyperbolic surfaces at the boundary. If we expand each vertex into a face, we get a subdivision rule with valence three, and certain faces that never subdivide (these represent the boundary at infinity). If we double the manifold over its boundary, all edges of the blown up vertices will have edge-cycle length four, and the other edges retain their original edge cycle length. This gives a closed hyperbolic manifold with even edge-cycle length, as desired.

As a final note, odd edge-cycle length polyhedra are more difficult to work with, especially length 3, but many of the same principles apply. One open problem is to find a concise set of conditions on odd edge-cycle length polyhedral gluings that ensures a subdivision rule exists. This is interesting, because all 3-manifolds have a decomposition into valence 3, edge-cycle length 3 polyhedra (by virtue of the Heegaard splitting). It is unknown how many hyperbolic 3-manifolds have a decomposition into even edge-cycle length polyhedra.

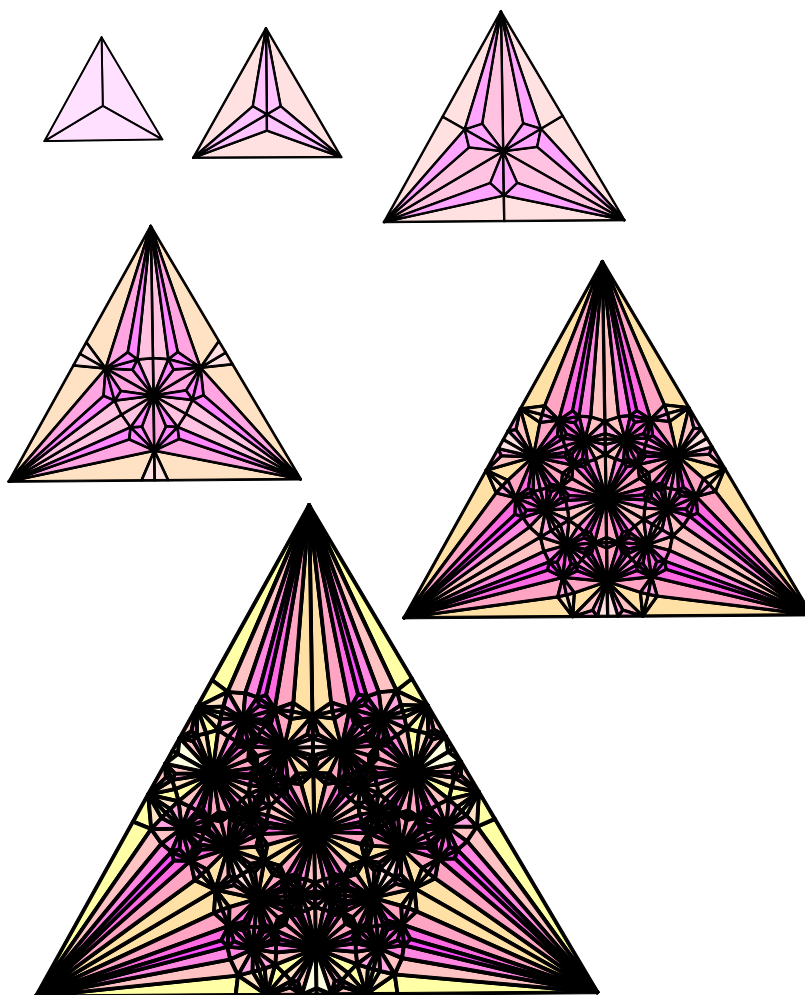


Figure 30: The first few subdivisions of the figure eight knot, decomposed into two regular ideal tetrahedra

References

- [1] **E M Andreev**, *Convex polyhedra in Lobačevskiĭspaces (English transl.)*, Math USSR Sbornik 10 (1970) 413–440 MR0259734
- [2] **J W Cannon**, *The combinatorial structure of cocompact discrete hyperbolic groups*, Geom. Dedicata 16 (1984) 123–148 MR758901




Figure 31: Several subdivisions of the 3–dimensional torus, with fundamental domain a Euclidean cube

- [3] **J W Cannon, W J Floyd, W R Parry**, *Finite subdivision rules*, *Conform. Geom. Dyn.* 5 (2001) 153–196 (electronic) MR1875951
- [4] **J W Cannon, E L Swenson**, *Recognizing constant curvature discrete groups in dimension 3*, *Trans. Amer. Math. Soc.* 350 (1998) 809–849 MR1458317
- [5] **W J Floyd**, *tilepack.c, tilepackhistory.c, subdivide.c and subdividehistory.c*, Software Available at <http://www.math.vt.edu/people/floyd/>
- [6] **P Haïssinsky**, *Empilements de cercles et modules combinatoires*, *Ann. Inst. Fourier (Grenoble)* 59 (2009) 2175–2222 MR2640918
- [7] **C D Hodgson, I Rivin, W Smith**, *A characterization of convex hyperbolic polyhedra and of convex polyhedra inscribed in the sphere*, *Bull. Amer. Math. Soc.* 27 (1992) 246–251 MR1149872




Figure 32: Several subdivisions for a cube with edge-cycle length 6

- [8] **G Perelman**, *Finite extinction time for the solutions to the Ricci flow on certain three-manifolds* arXiv:math.DG/0307245
- [9] **B Rushton**, *Alternating Links and Subdivision Rules*, Master's thesis, Brigham Young University (2009)
- [10] **B Rushton**, *Creating subdivision rules from alternating links*, *Conform. Geom. Dyn.* 14 (2010) 1–13
- [11] **K Stephenson**, *Circlepack*, Software Available at <http://www.math.utk.edu/~kens>




Figure 33: A further subdivision of the edge-cycle length 6 cube

Department of Mathematics, Temple University
Room 638 Wachman Hall, 1805 N. Broad Street, Philadelphia, PA 19122, USA

`lindianr@gmail.com`

<https://sites.google.com/site/rushtonresearch/home>

Received: 24 January 2012 Revised: 5 June 2012