2012-03-07

Subdivision Rules, 3-Manifolds, and Circle Packings

Brian Craig Rushton
Brigham Young University - Provo

Subdivision Rules, 3-Manifolds and Circle Packings

Brian Rushton

A dissertation submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

James W. Cannon, Chair
Jessica Purcell
Stephen Humphries
Eric Swenson
Christopher Grant

Department of Mathematics

Brigham Young University

April 2012

Abstract

Subdivision Rules, 3-Manifolds and Circle Packings

Brian Rushton

Department of Mathematics

Doctor of Philosophy

We study the relationship between subdivision rules, 3-dimensional manifolds, and circle packings. We find explicit subdivision rules for closed right-angled hyperbolic manifolds, a large family of hyperbolic manifolds with boundary, and all 3-manifolds of the $\mathbb{E}^3, \mathbb{H}^2 \times \mathbb{R}, \mathbb{S}^2 \times \mathbb{R}, \widetilde{SL_2(\mathbb{R})}$, and S^3 geometries (up to finite covers). We define subdivision rules in all dimensions and find explicit subdivision rules for the n-dimensional torus as an example in each dimension. We define a graph and space at infinity for all subdivision rules, and use that to show that all subdivision rules for non-hyperbolic manifolds have mesh not going to 0. We provide an alternate proof of the Combinatorial Riemann Mapping Theorem using circle packings (although this has been done before). We provide a new definition of conformal for subdivision rules of unbounded valence, show that the subdivision rules for the Borromean rings complement are conformal and show that barycentric subdivision is almost conformal. Finally, we show that subdivision rules can be degenerate on a dense set, while still having convergent circle packings.

Keywords: LaTeX, PDF, BYU, Math, Thesis, subdivision, rules, manifold, 3-manifold, circle, packings, infinity, space, geometries, Perelman, torus, hyperbolic, unbounded, valence

Acknowledgments

This work is for Lindsey and Leonidas. Drs. Cannon, Purcell, Humphries and Swenson provided a great deal of feedback and help.

Contents

List of Figures

Introduction

A long-standing conjecture of Cannon is that every Gromov hyperbolic group with a 2-sphere at infinity is a Kleinian group [8]. By this we mean that:

(i) the group acts on H^3 by isometries,

(ii) the action is properly discontinuous, and

(iii) the action is cocompact.

All Gromov hyperbolic groups that are already known to be 3-manifold groups are Kleinian groups by the Geometrization Conjecture, proved by Perelman [15]. However, it is not at all obvious that hyperbolic groups with a 2-sphere at infinity correspond to any manifold at all, and this is the reason the conjecture remains unsolved.

One approach, adopted by Cannon, Floyd, and Parry, among others, is to study subdivision rules [6]. All Gromov hyperbolic groups with a 2-sphere at infinity have a subdivision rule on the sphere [7]. A subdivision rule is a way of dividing the sphere into a tiling, or cell structure, with a recursive formula for dividing each tile into smaller tiles. A more rigorous definition of subdivision rule can be found in Section 3.1.

Some well-known examples of subdivision rules include barycentric subdivision and hexagonal refinement (where a triangle is chopped up into smaller triangles by connecting the midpoints of each pair of edges). Cannon has shown that, if a subdivision rule for a group is conformal (meaning that tiles don't get too distorted in the long run), then the group must be a hyperbolic 3-manifold group [8]. However, it has proven difficult to determine if a subdivision rule is conformal or not. Cannon and Swenson have proven the converse, i.e. that a hyperbolic 3-manifold groups have a conformal subdivision rule [7]. Conformality is discussed in more detail in Part II.

We will now give an outline of the dissertation. All work in this dissertation is original work, outside of Chapter 5.

This dissertation is divided into two parts. In Part I, we find a large number of explicit subdivision rules for manifolds. Previously, only a handful of explicit subdivision rules for manifolds were known, e.g. the dodecahedral orbifold depicted in [6]. In [17], [18], we found explicit subdivision rules for all non-split, prime alternating links and all torus links. Part I of this dissertation expands our collection of subdivision rules for manifolds significantly.

In Chapter 1, we construct an explicit subdivision rule for all manifolds created by gluing right-angled hyperbolic polyhedra, and for other hyperbolic manifolds created from sufficiently large polyhedra. This includes a wide variety of hyperbolic manifolds with toral boundary and hyperbolic surface boundary. This greatly expands the number of knowm conformal subdivision rules.

In Chapter 2, we construct explicit subdivision rules for all manifolds (up to finite covers) of five of the eight geometries (excluding H^3, Nil, and Sol). We use some results from Chapter 1, but many of the manifolds require new methods.

In Chapter 3, we show that the idea of a subdivision rules can be extended in interesting ways to manifolds of dimension $\neq 3$, and we construct an explicit family of examples: the n-torus.

This part culminates in Chapter 4, where we define a space at infinity and a graph associated to a subdivision rule of any dimension. The graph is like a Cayley graph, and the space, which is often non-Hausdorff, extends the idea of the space at infinity for hyperbolic groups. We show that this graph is hyperbolic whenever the space is Hausdorff. An interesting application shows that all subdivision rules for non-hyperbolic manifolds have mesh not going to 0.

Part II approaches subdivision rules from a different direction, that of circle packings. In Chapter 5, we first introduce the idea of modulus and conformality, and prove that circle packings give an accurate picture of subdivision rules and their modulus. All results and

techniques in this previous chapter were previously known, though proved independently. This proof fails for subdivision rules of unbounded valence.

In fact, many results seem to fail for subdivision rules of unbounded valence. In Chapter 6, we try an alternate way of assigning modulus to subdivision rules, and investigate a subdivision rule with exponential growth (barycentric subdivision) and one with linear growth (the Borromean rings) at each vertex. Through this, we show that the Borromean rings have conformal subdivision rules, and conjecture that all hyperbolic, alternating links have conformal subdivision rules.

Last of all, in Chapter 7, we show that subdivision rules can be degenerate in interesting ways on dense sets, even while being 'conformal' (in some sense) on a dense complement.

Part I

Subdivision Rules and Manifolds

Chapter 1. Creating Subdivision Rules From Polyhedra With Identifications

1.1 Preliminaries

As stated in the Introduction to this thesis, subdivision rules were first defined in an attempt to prove Cannon's conjecture that all Gromov hyperbolic groups with a 2-sphere at infinity are Kleinian groups. All that is necessary to prove Cannon's conjecture is to show that the subdivision rules coming from groups with a 2-sphere at infinity are conformal. This has proven difficult.

One difficulty is the lack of examples. In [6], 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 subdivision rules have been found for all non-split, prime alternating links [18], 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 definition of 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, using only 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 chapter, we construct an explicit 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. A replacement rule is different from a subdivision rule. Both give a recursive way of constructing

one tiling from another, but a subdivision rule requires the new tiling to include the old one as a subset, while a replacement rule does not. More specifically, a subdivision rule replaces one tile with several, while a replacement rule replaces several tiles with several other tiles.

We then explain how to convert these replacement rules into subdivision rules in a large number of cases. As an example of our method, the tilings in Figures 1.1 and 1.2 correspond to the same manifold that Cannon, Floyd and Parry studied in [6]; the tiling in their paper is different from the one we obtain, but strongly related.

To associate a subdivision rule to a manifold (even a non-hyperbolic one), 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. This defines a sequence of tilings of the sphere, which defines a combinatorial structure on the space at infinity. Note that $S(n)$ must be a topological sphere for all n.

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

Definition. 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. A **fan** is a chain $a_1, a_2, ..., a_n$ 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 $a_1, ..., a_n$ is the star of the vertex. We consider a single face to be a fan of size 1. See Figure 1.3 for examples of fans.

Lemma 1.1. *A fan is a topological closed disk.*

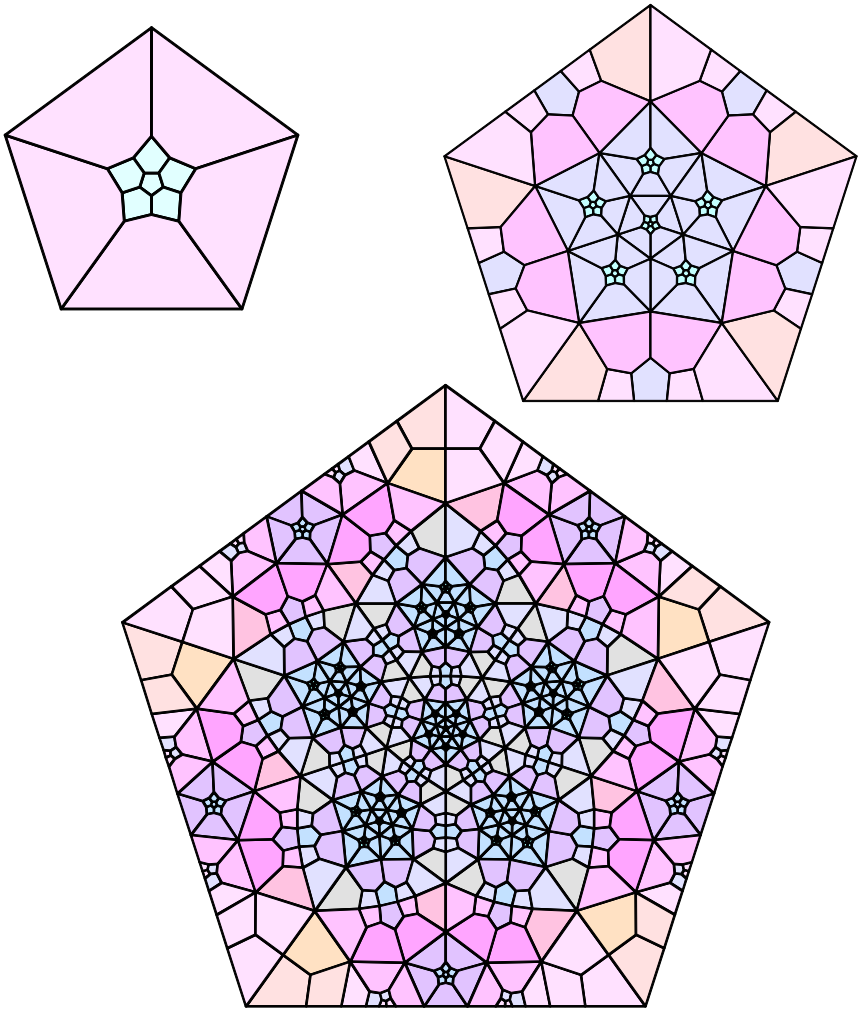

Figure 1.1: The first three subdivisions of a tile type for a hyperbolic orbifold.

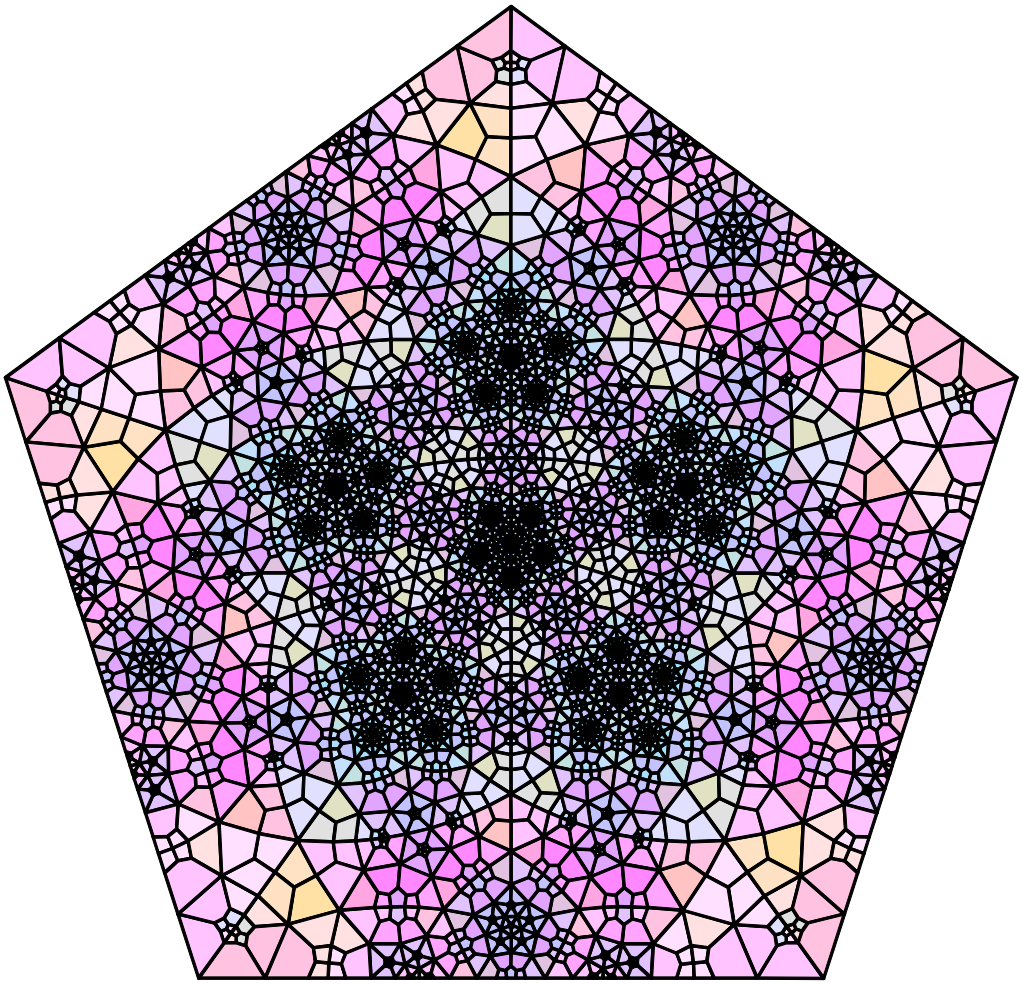

Figure 1.2: The fourth subdivision of a tile type for the same hyperbolic orbifold.

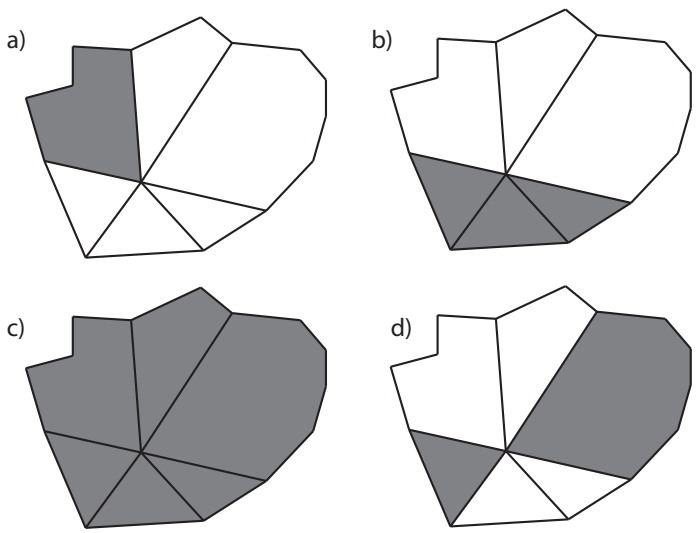

Figure 1.3: The shaded portions in a,b, and c are all examples of fans, while the shaded portion in d is not.

Proof. Each face of a polyhedron is a topological closed disk. The union of two faces over a single edge is a closed disk. Because of our definition of polyhedron, all fans of size two are the union of two faces over a single edge.

Now, assume all fans of size k are closed disks. Given a fan B of size $k+1$ with faces $a_1, ..., a_{k+1}$, we can express it as the union of the face a_{k+1} and of the fan A consisting of $a_1, ..., a_k$. By induction, the fan A is a closed disk, and we know a_{k+1} is a closed disk. Now, a_{k+1} intersects every tile of A in the common vertex of A. So, by our definition of polyhedron, the intersection of each face of A with a_{k+1} is the common vertex or an edge containing the common vertex. If a_{k+1} is an end of the fan, it will intersect A in a single edge, and so B will be a closed topological disk. If B is the star of the vertex, then a_{k+1} will intersect A in two connected edges, and B will be a closed topological disk. □

Definition. 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 two contiguous edges.

Among the platonic solids, the cube, dodecahedron, octahedron, and icosahedron are

spread out. The tetrahedron is not; if we take our fan to be the star of a vertex, all three of its edges intersect the remaining face.

A polyhedron that is spread out is analogous to a prime alternating link. 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, non-split, and prime, and these properties are very similar to the spread-out condition above.

Andreev's theorem [1] (later extended by Rivin and others [12]) implies that all convex hyperbolic polyhedra with dihedral angles equal to $\frac{\pi}{2}$ at every edge are spread out.

We will need restrictions on edge cycle lengths.

Definition. When polyhedra are glued together, each edge may be identified with many other edges. Edges identified form an equivalence class called an edge cycle. 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 $\frac{2\pi}{n}$ at an edge e, the edge cycle length of e is n. We are now ready to state the theorem.

1.2 Existence of replacement rules

Here we make a more formal definition of replacement rule. A **replacement rule** is a locally defined recursive rule that has finitely many **replacement types**. A replacement type is a labeled, connected groups of faces that is to be replaced by a different collection of faces.

Theorem 1. Let M be a manifold (possibly with boundary) that can be decomposed into polyhedra $P_1, P_2, ..., P_n$. If each polyhedron is spread out, and all edge cycles have even length ≥ 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.

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. 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. 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).

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

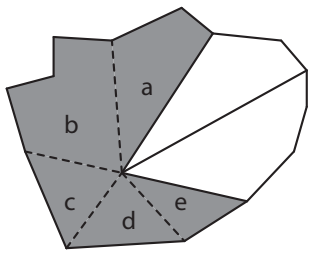

Figure 1.4: If the shaded portion is a loaded fan, then faces b,c, and d will be examples of loaded wedges, while faces a and e will be examples of half-loaded wedges. Dashed edges are loaded, and others are not.

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 1.4 for an example.

As for the second step, note that we said loaded wedges *may* group together into fans. Loaded wedges are, again, like pizza slices. If we rotate a pizza slice 180^o 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 there are two loaded edges that intersect it, one at each vertex, and these loaded edges do not border the same face. See Figure 1.5 for an example of a skew edge. Note that a skew edge can be unburdened, burdened, or loaded.

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,

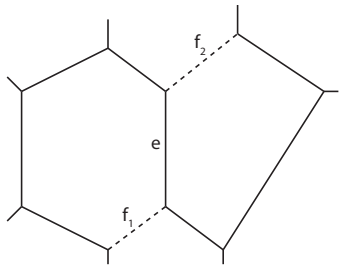

Figure 1.5: A skew edge e. Here f_1 and f_2 are loaded, intersect e in distinct vertices, but border different faces.

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 unglued faces in the interior (meaning those which do not intersect the boundary) correspond to faces on the polyhedron that don't 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 1.1), 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.

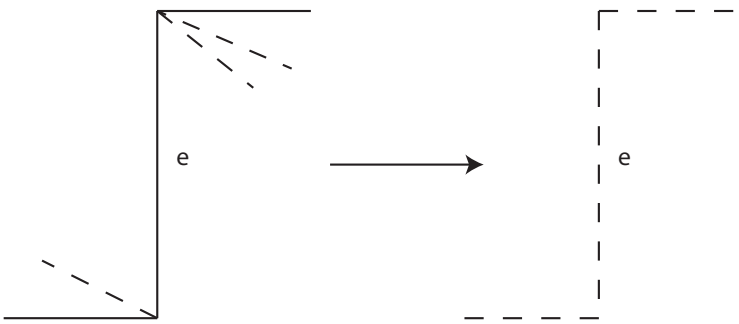

Figure 1.6: Burdened skew edges must come from other skew edges.

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 1.5), 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 1.1).

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 1.1. 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 edges. So e must have been a skew edge (burdened or unburdened) in $S(n-1)$, which is impossible by our induction hypothesis. See Figure 1.6.

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 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 1.7. Thus, the faces on either side (call them

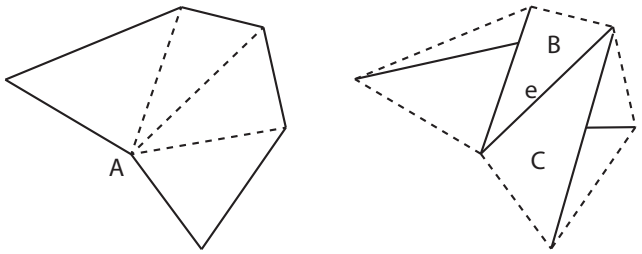

Figure 1.7: An unburdened skew edge must look something like this.

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 can't 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 1.8). 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 the universal cover, i.e. a replacement rule. □

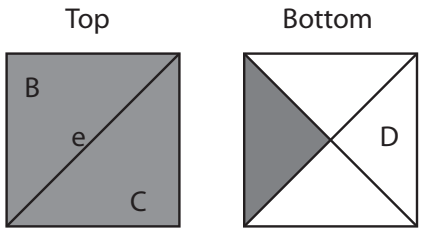

Figure 1.8: Note that the fan consisting of the shaded gray region has disconnected intersection with the face D.

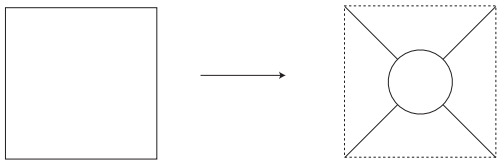

Figure 1.9: The replacement rule for an unloaded face.

1.3 Forming subdivision rules from replacement rules

Note that Theorem 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 non-hyperbolic 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 1.9, 1.10, and 1.11. They correspond to an unloaded face, a **loaded pair** of faces (i.e. a fan of two faces, sharing a common loaded edge), and a **loaded star** (i.e. 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.

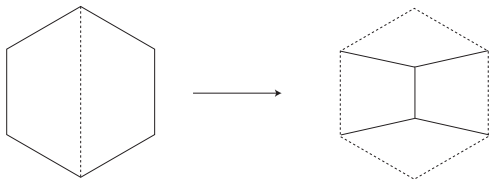

Figure 1.10: The replacement rule for a loaded pair.

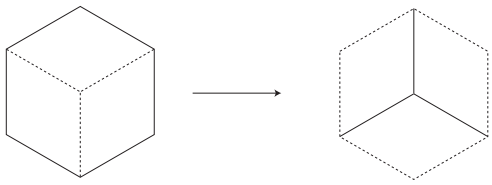

Figure 1.11: The replacement rule for a loaded fan.

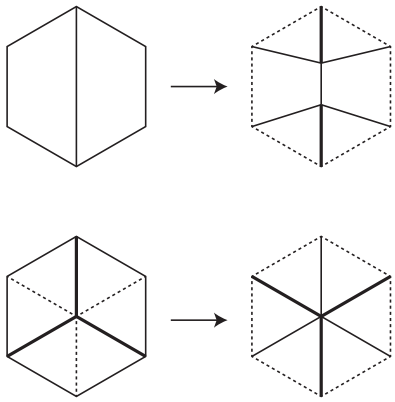

Figure 1.12: We can add lines to the two 'loaded' tile types to get a self-consistent subdivision rule.

This can be turned into a subdivision rule by a method that turns out to be very useful: adding new edges at every stage. For instance, in Figure 1.10, 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 1.12), then the new cell structure contains the old cell structure as a subset. Thus, we have a subdivision rule.

However, this 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 1.11 gives us a hexagon divided into six 'pie slices'. If we add similar lines bisecting the loaded star (as seen in Figure 1.12), we again get a hexagon divided into six triangles; thus, the subdivision on each triangle in that hexagon is just the identity.

We summarize this in Figure 1.13. Several stages of subdivision are shown in Figure 1.31 on page 35. This is a combinatorial subdivision only; the circle packed pictures are

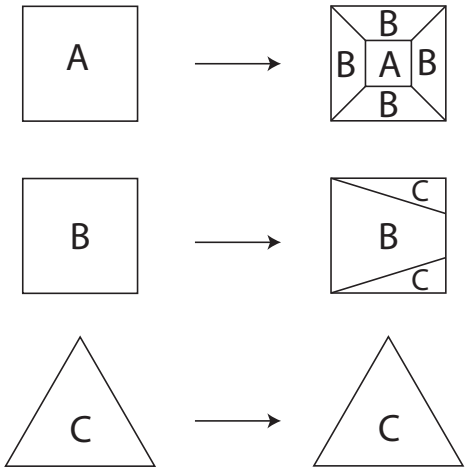

Figure 1.13: The replacement rule for the 3-dimensional torus.

not subsets of each other, because this subdivision rule is not conformal. The connection between circle packings and conformality is explained in [11].

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 tile type can cascade and force us to add more and more lines until we have infinitely many 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 $\frac{2\pi}{n}$, 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 [12]. An n-**cycle** is a set of n faces $A_1, A_2, ..., A_n$ 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 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 2. Assume a manifold M is formed by gluing together convex, hyperbolic polyhedra $P_1, .., P_n$, all of whose dihedral angles are $\frac{\pi}{2}$. Then the replacement rules in Theorem 1 can be made into subdivision rules by adding finitely many lines to replacement tile types.

Note that all such polyhedra satisfy [12]:

(i) each polyhedron P_i is spread out, and

(ii) each vertex of each polyhedron has valence three,

(iii) every edge (after gluing) has edge-cycle length four, and

(iv) there are no three-cycles or four-cycles.

The above is almost exactly the statement of Andreev's theorem in its original form [1].

Proof. We need only show that we can make each replacement of a tile a subdivision of a tile, while retaining finitely many tile types. Because every vertex has valence three, any edges that share a vertex also share an edge, which simplifies the combinatorics significantly (see Figure 1.14). Throughout the following proof, we will refer to the replacement rule for the three torus as an example. In Figure 1.15, 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.

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

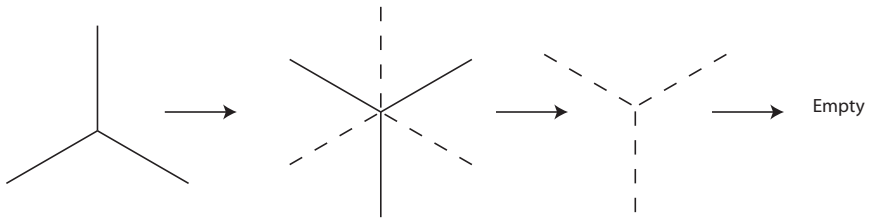

Figure 1.14: Each new vertex starts with three unburdened lines. In the next stage of replacement, three unburdened lines are added, and the old edges become loaded. The vertex is surrounded by three loaded pairs. In the next stage, there are only three edges, and the vertex is now loaded. In the next stage, the vertex and its star are entirely covered up by a new polyhedron.

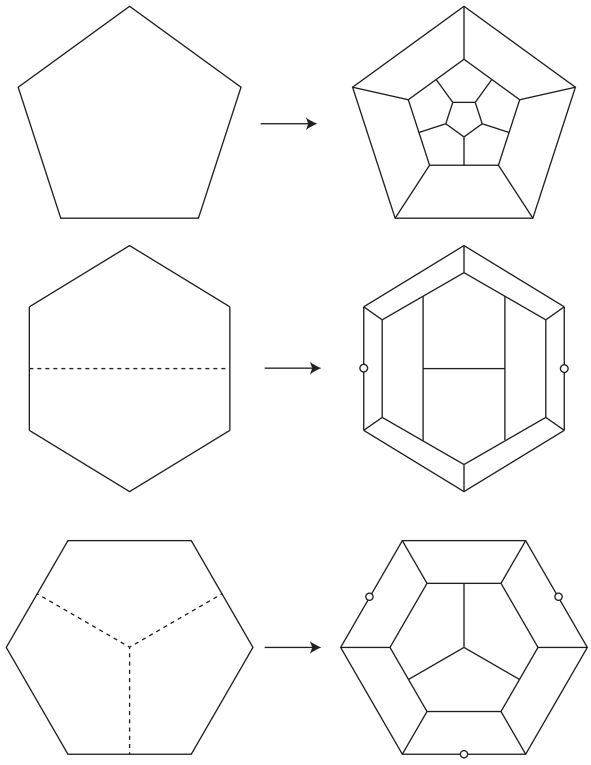

Figure 1.15: The three replacement types for the right-angled dodecahedral orbifold. They are fans of size one, two, and three.

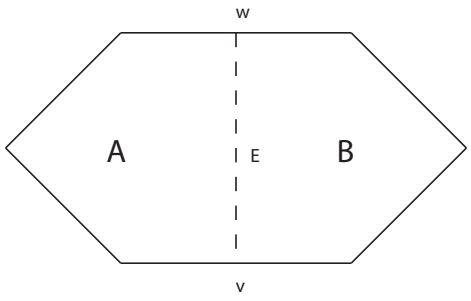

Figure 1.16: A loaded pair.

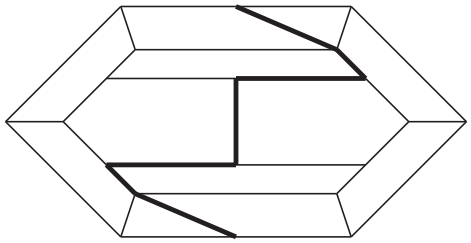

Figure 1.17: We replace the edge E from figure 1.16 with the path in bold. Note that the initial and terminal segments were not originally in P^*.

v and w (see Figure 1.16). 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 1.12 and 1.17). Each of these two vertices is contained in a unique tile of $P^* \subset S(n+1)$ (recall Figure 1.14). 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.

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 (i.e. the complement of the separating tile is disconnected, with each component of the complement containing interior edges of the separating tile). Because 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 1.1.

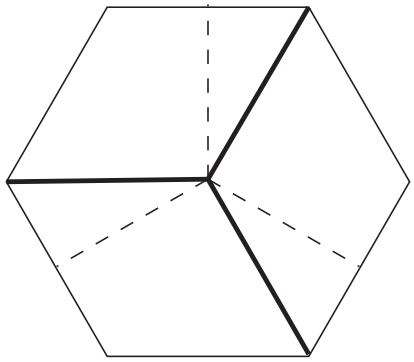

Figure 1.18: All loaded stars have three loaded lines and three added lines from the previous stage.

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 Figure 1.12 and Figure 1.17.

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 1.14). 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 Figure 1.13 and Figure 1.18). This is case III. Note that all other tiles in $P^* \subset S(n+1)$ correspond to case I or case II.

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 won't 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 1.14). 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 $\alpha_3, ..., \alpha_n$, 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 $\cup\alpha_i$.

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 1.19). 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. However, this would imply that our polyhedron contains a three cycle or a four cycle (four faces which intersect cyclically in four disjoint edges), which is impossible by Andreev's theorem [1]. See Figure 1.20. This is where the proof fails for non-hyperbolic manifolds such as the 3-torus.

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

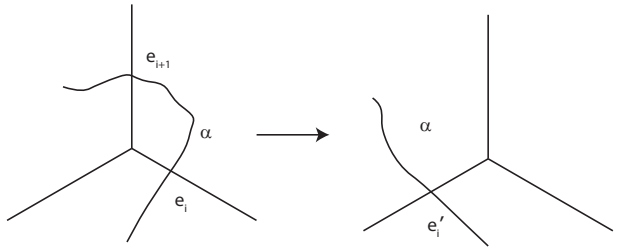

Figure 1.19: A path that goes through non-disjoint edges can be shortened.

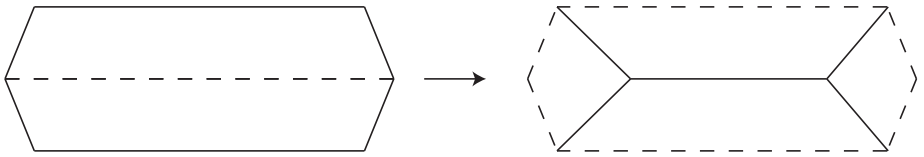

Figure 1.20: If the complement of the boundary tiles is disconnected, there must be a four cycle.

connects e to the midpoint of the segment of α contained in T. See Figure 1.21.

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 $P^* \setminus \{\alpha \cup \beta \cup T^o\}$, where T^o 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 (i.e. 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 of T_i, 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 1.22. We now identify each γ_i with the appropriate

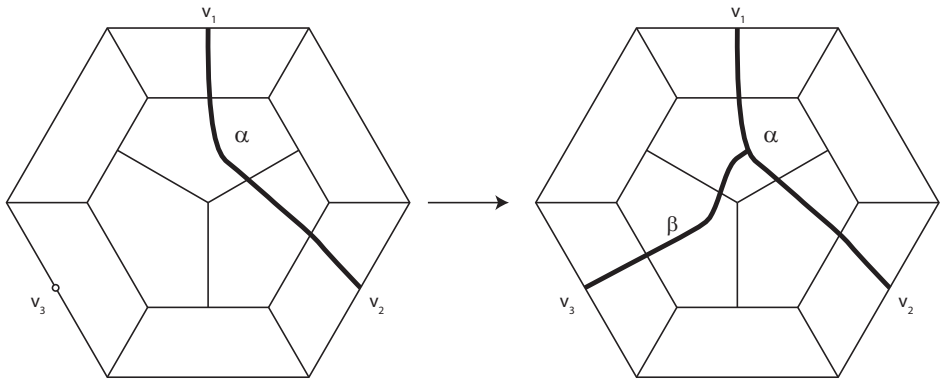

Figure 1.21: We find paths α and β to replace the loaded edges of the loaded star.

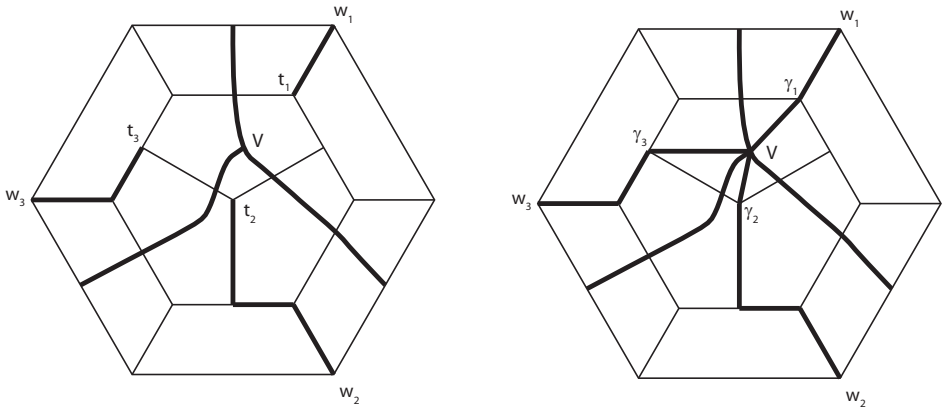

Figure 1.22: We find paths γ_1, γ_2, and γ_3 to replace the three added lines of the loaded star.

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 (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

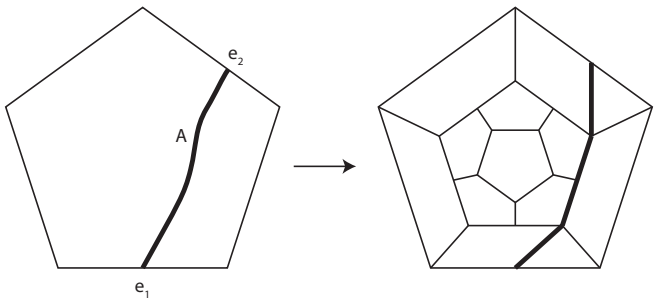

Figure 1.23: Replacing a single unloaded tile with an added line from midpoint to midpoint.

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 1.23.

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 1.24.

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 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

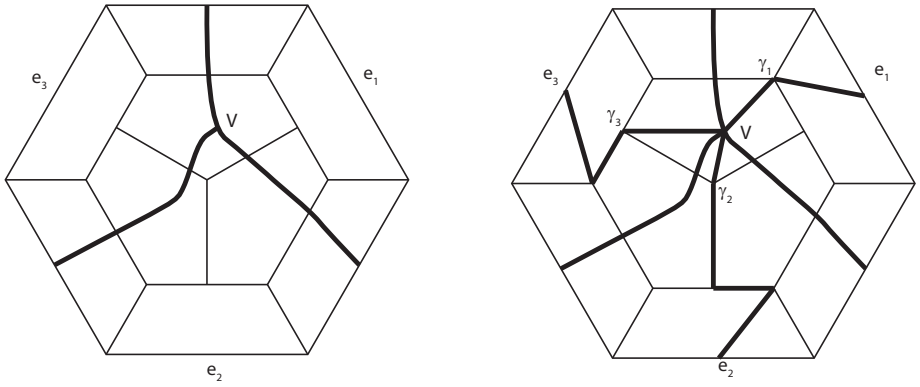

Figure 1.24: Replacing a loaded star with added edges going to midpoints instead of boundary vertices.

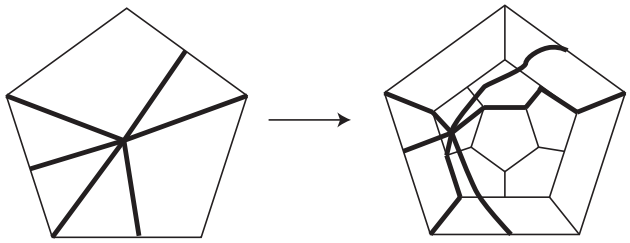

Figure 1.25: Replacing a single unloaded tile with six added lines.

α of added edges connecting disjoint edges of non-boundary 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. $S(n)$**.**

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

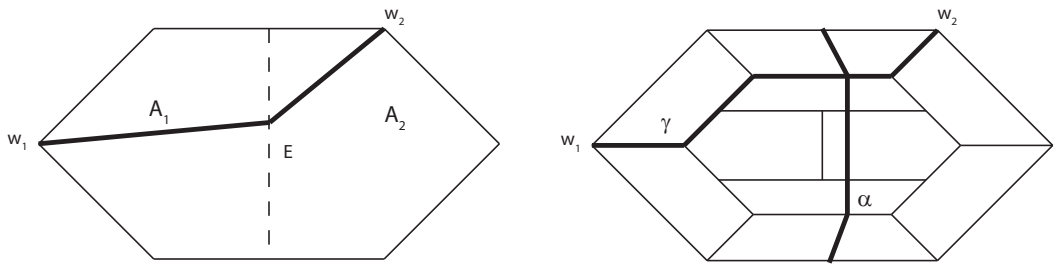

Figure 1.26: Replacing a loaded pair with added lines starting at the boundary and meeting in the midpoint of the loaded edge.

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.

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.

End of cases

So note that 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 1 for such a manifold into a subdivision rule with finitely many tile types.

□

This theorem applies only to 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 a 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), Chapter 9 of [17] ensures that only loaded pairs and unloaded tiles occur in the replacement rules for the manifold. In particular, we have the following:

Theorem 3. Let M be a 3-manifold created by gluing together polyhedra $P_1, ..., P_n$ such that each edge cycle has even length > 2. If each P_i is spread out and:

(i) all vertices have valence greater than 3, or

(ii) all edges have cycle length greater than 4, then

the replacement rules in Theorem 1 can be made into subdivision rules by adding finitely many lines to replacement tile types.

Proof. Recall that the only tile types in the replacement rule from Theorem 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 Chapter 9 of [17]. 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 (i.e. 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 Theorem 12 of [17] 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 1.1). No new tile types are created by finding such a path, and so we have a subdivision rule. □

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 [18]. 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.

The tilings in Figures 1.28 to 1.33, and Figures 1.1 and 1.2 were all created by the methods detailed in Theorems 2 and 1.3. The pictures were created using Ken Stephenson's Circlepack [21] and software by Bill Floyd [10].

Figures 1.28 and 1.29 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 1.30 is another finite-volume hyperbolic manifold: the figure-eight knot.

Figures 1.31-1.33 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 1.1 and 1.2 show the tilings from a subdivision rule created from the dodecahedral orbifold first studied by Cannon, Floyd and Parry by an adaptation of our

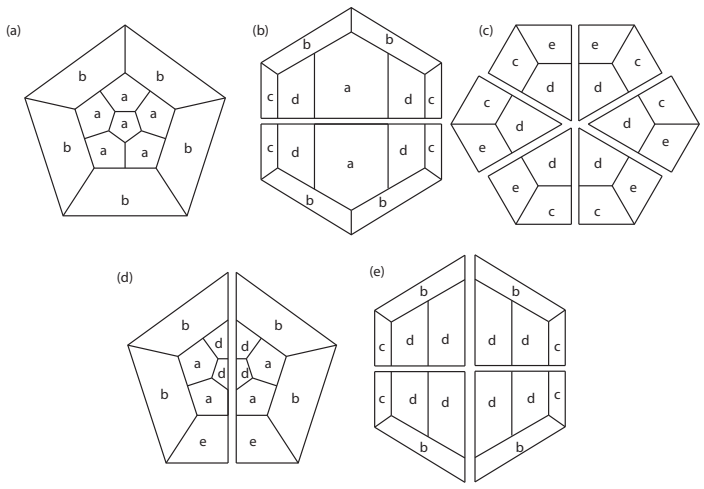

Figure 1.27: A subdivision rule obtained from the replacement rule in Figure 1.15. We use a slightly different method than the one outlined in Theorem 2 to reduce the number of tile types.

method. This is a dodecahedron with edge cycle length four. It is interesting to compare this to their original subdivision rule [6]. The tiles are shown in Figure 1.27.

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 Heegard splitting). It is unknown how many hyperbolic 3-manifolds have a decomposition into even edge-cycle length polyhedra.

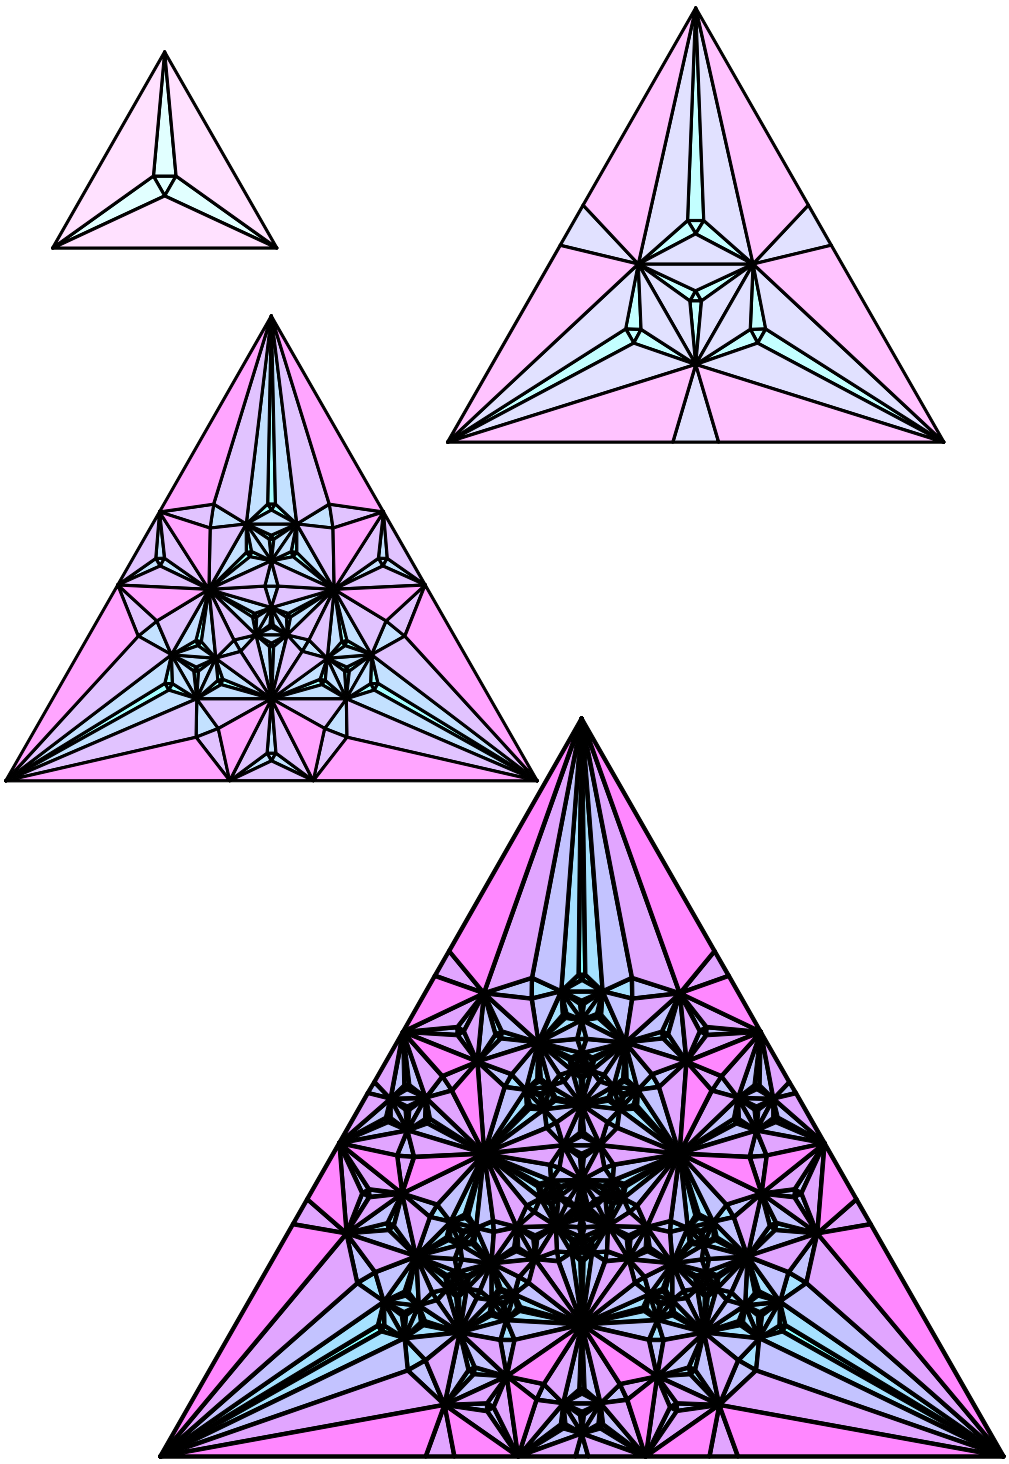

Figure 1.28: The first few subdivisions of the Borromean rings complement, which decomposes into two octahedra with edge cycle length 4.

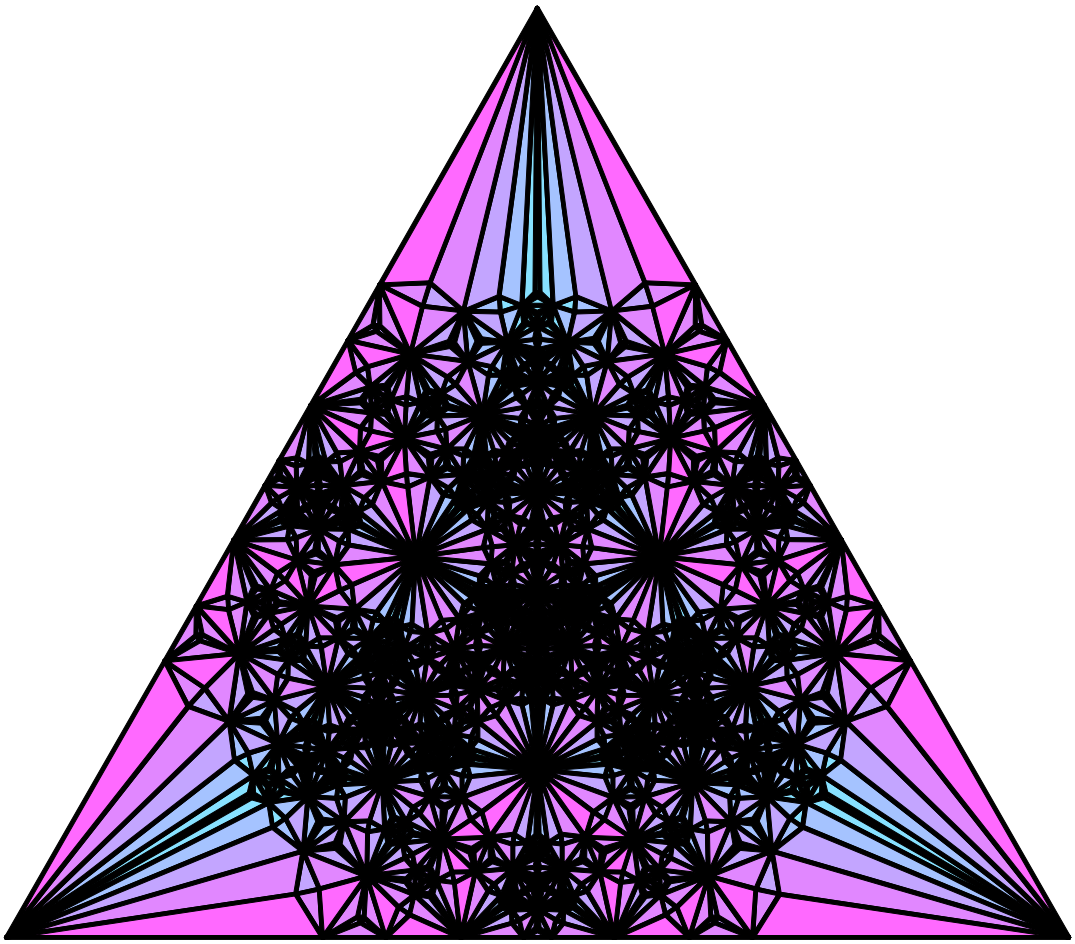

Figure 1.29: One more subdivision of the Borromean rings complement.

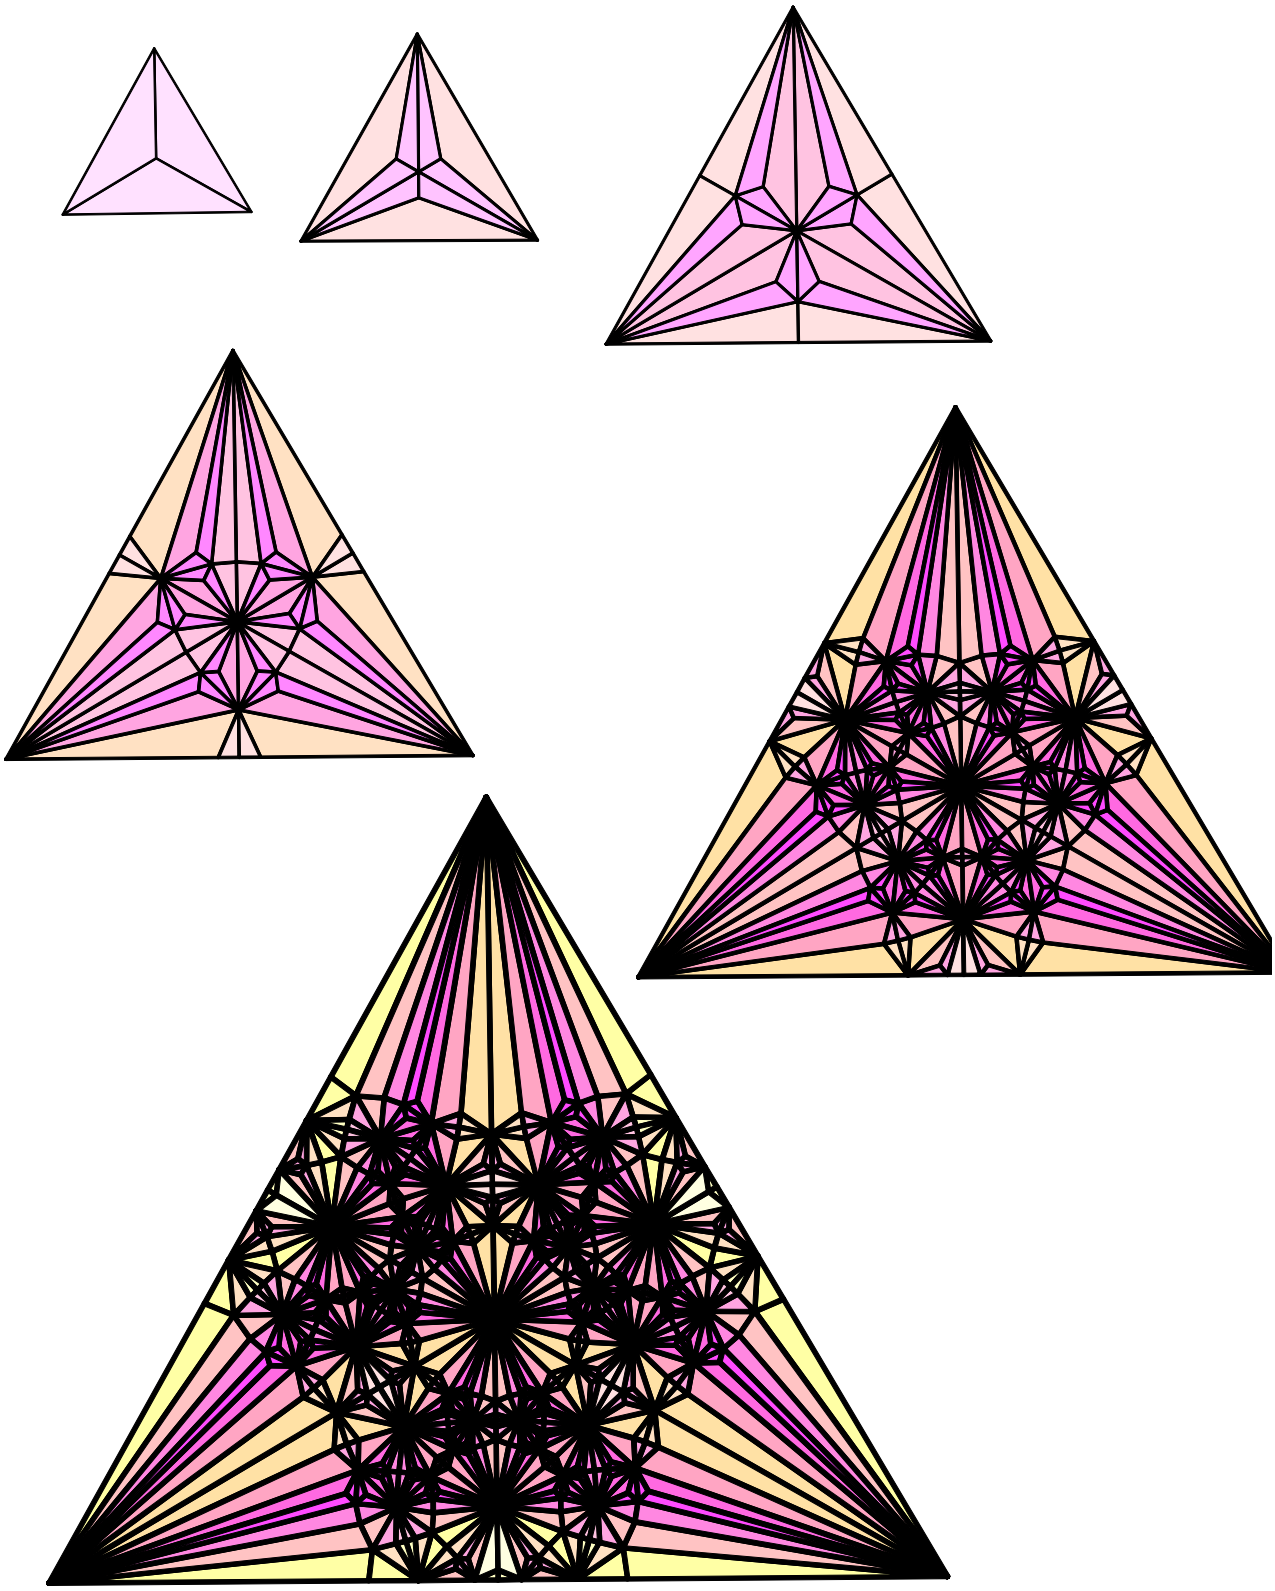

Figure 1.30: The first few subdivisions of the figure eight knot, decomposed into two regular ideal tetrahedra.

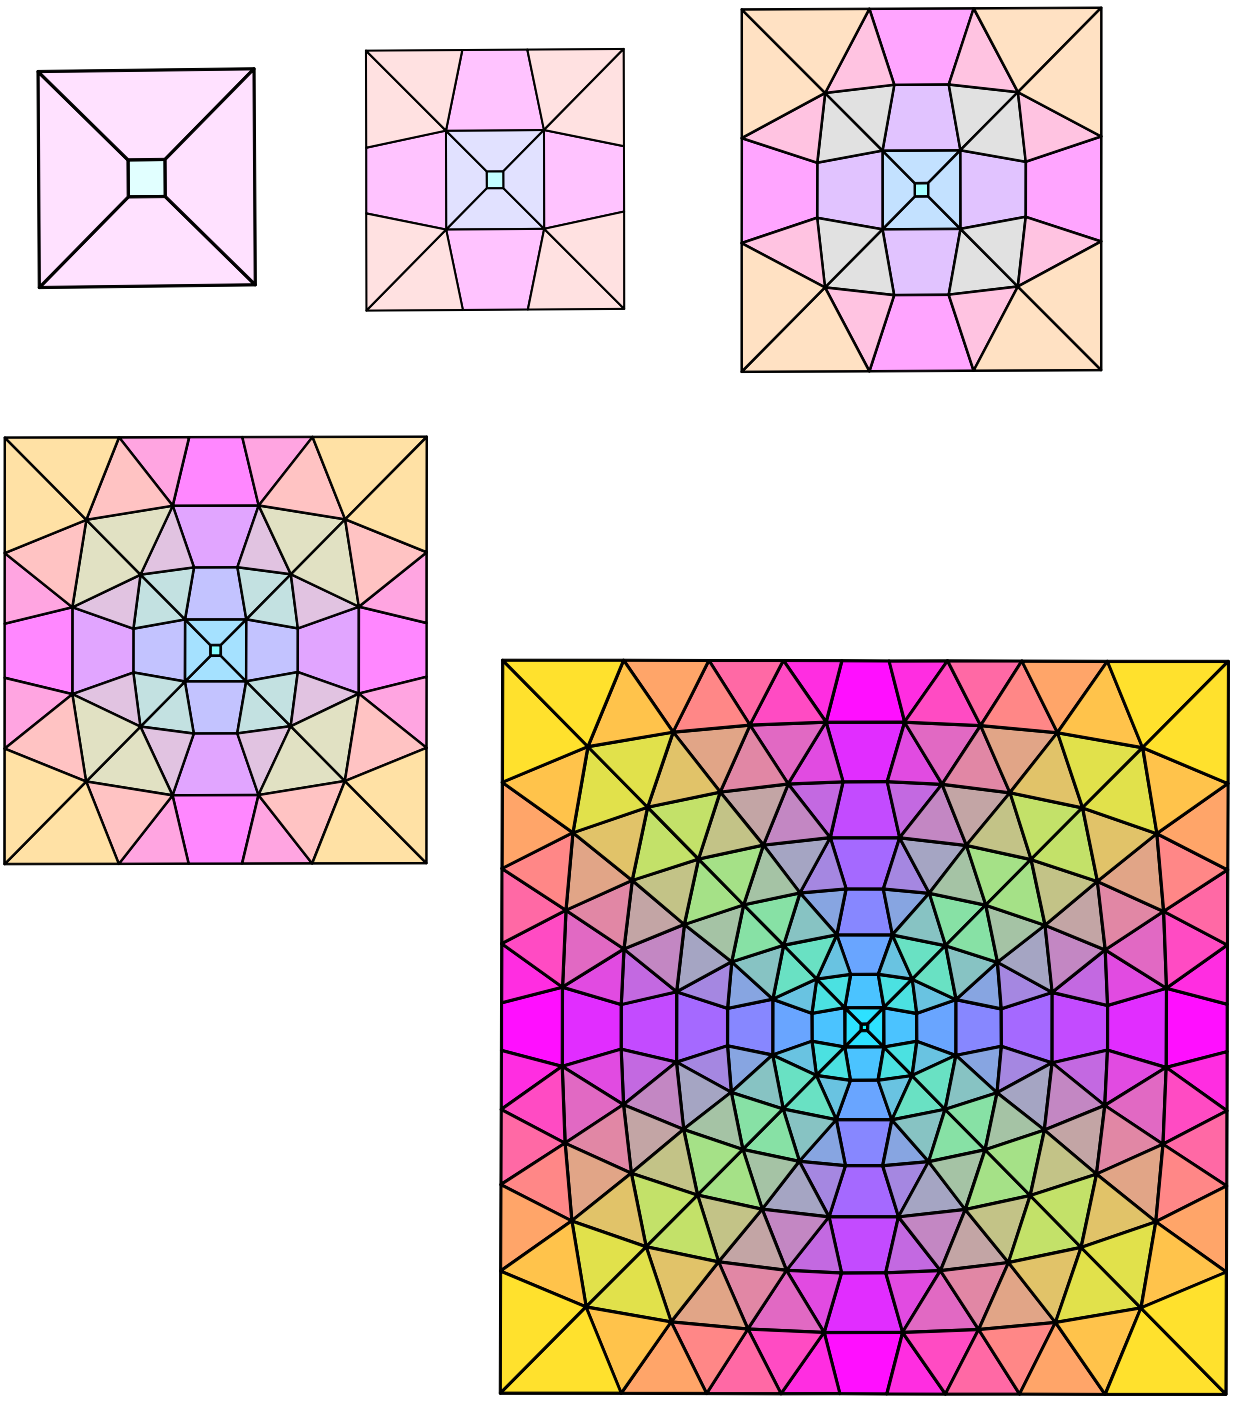

Figure 1.31: Several subdivisions of the 3-dimensional torus, with fundamental domain a Euclidean cube.

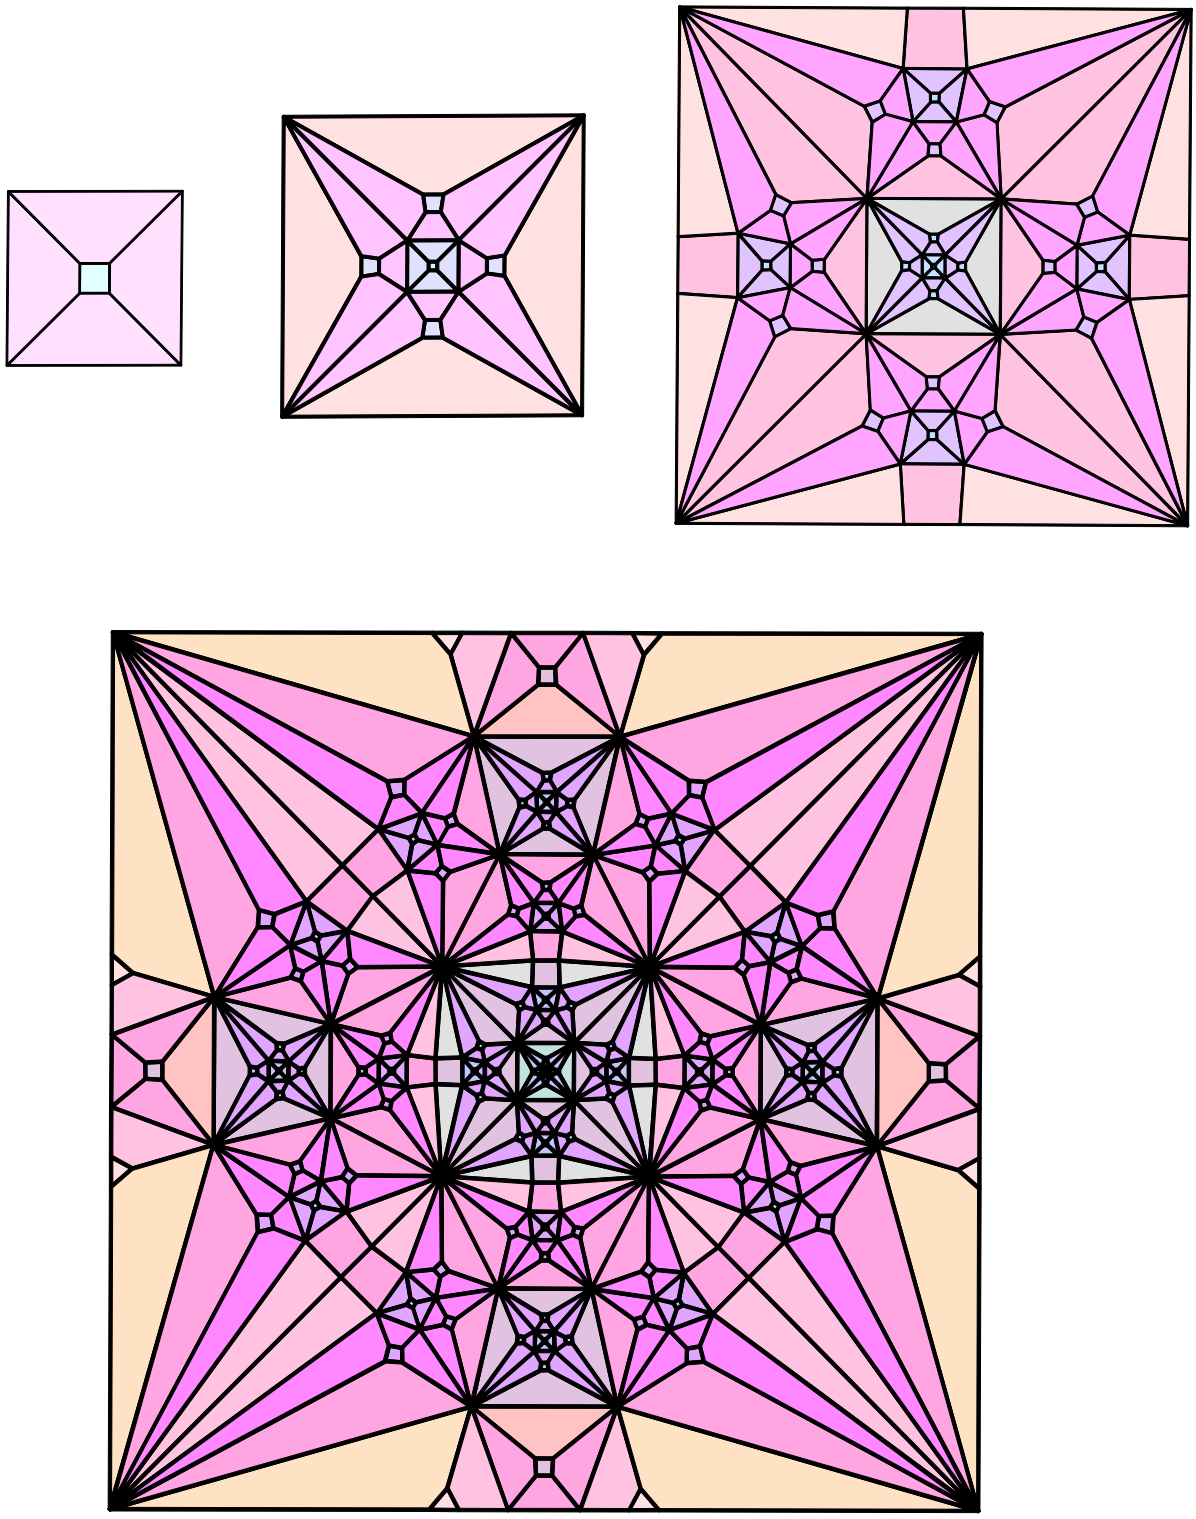

Figure 1.32: Several subdivisions for a cube with edge-cycle length 6.

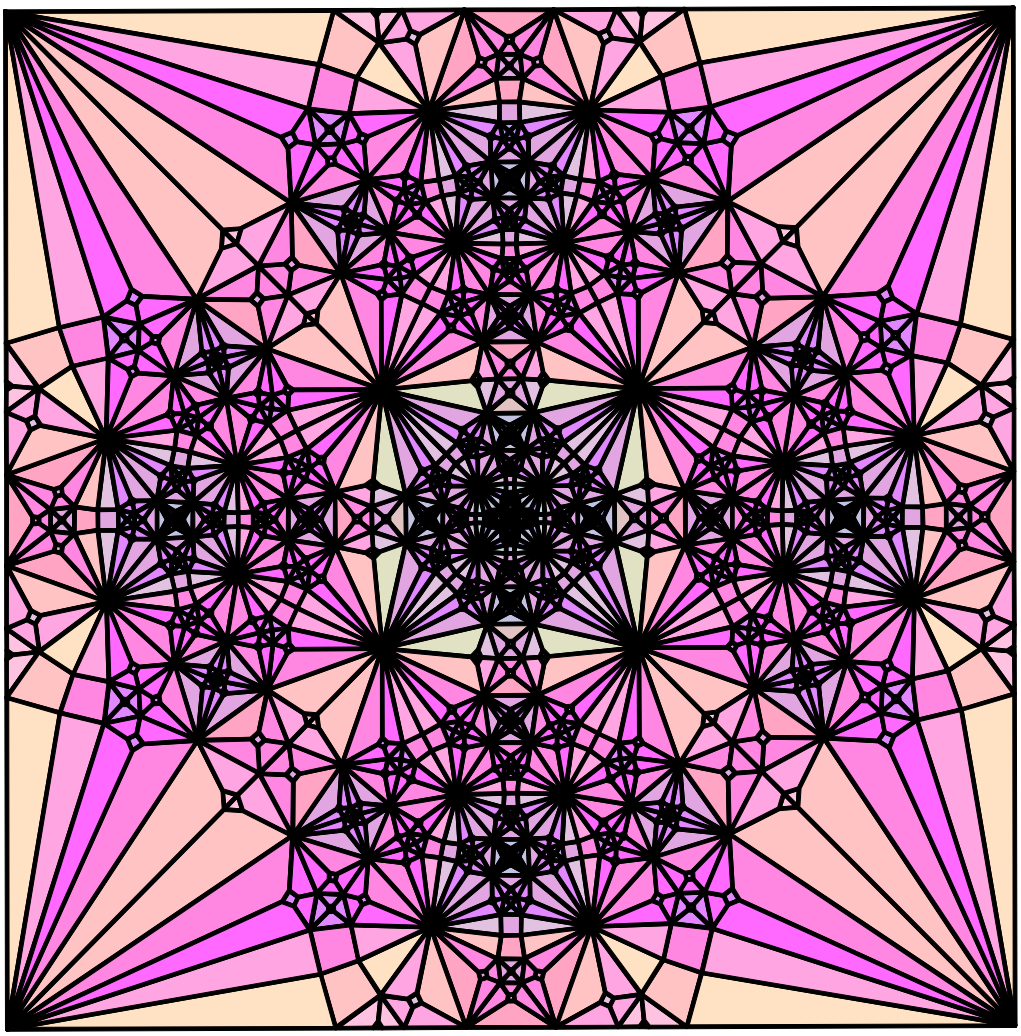

Figure 1.33: A further subdivision of the edge-cycle length 6 cube.

Chapter 2. Subdivision Rules and the Eight Geometries

2.1 Background

This chapter assumes basic knowledge of algebraic topology (fundamental group, covering spaces, etc.) and Thurston's eight geometries [23]. Peter Scott's exposition of the eight geometries [19] is especially helpful.

2.2 Preliminaries

In this chapter, we find subdivision rules for almost all closed non-hyperbolic 3-manifolds, and show that manifolds of different geometries have subdivision rules that give wildly different combinatorial spaces at infinity. At the beginning of each section, we show a circle-packed image of several stages of subdivision created with Ken Stephenson's Circlepack [21].

In the sections that follow, we will examine each of Thurston's eight geometries. We pick a representative manifold, find a subdivision rule where possible, and display circle packings of the subdivision rules using Ken Stephenson's Circlepack [21]. For four of the geometries, the representative manifold will be essentially the only manifold. In the $\mathbb{E}^3$, $\mathbb{H}^2 \times \mathbb{R}$, and $\mathbb{S}^2 \times \mathbb{R}$ geometries, every manifold is, up to a finite cover, a trivial circle bundle over a Euclidean, hyperbolic, or spherical 2-manifold, respectively [19]. We cover these three types of circle bundles in the following examples.

All $\widetilde{SL_2(}\mathbb{R}$ manifolds are, up to finite covers, unit tangent bundles over hyperbolic surfaces, which all cover the unit tangent bundle of the smallest hyperbolic surface. All $\mathbb{S}^3$ manifolds are finitely covered by $\mathbb{S}^3$ [19]. Thus, these two geometries can also each be represented by a single closed 3-manifold, which we examine. Similar statements hold for Nil and Sol, but we will see that we have no subdivision rules for these geometries.

The reason we are only interested in manifolds up to finite covers is that, if one mani-

fold finitely covers another, we can divide the manifolds into fundamental domains that are combinatorially equivalent. For instance, the Klein bottle and the torus both have fundamental domains that are squares whose vertices are identified in groups of four. Thus, in finding subdivision rules for specific manifolds in the following sections, we are really finding subdivision rules for all 3-manifolds of that geometry (except, of course, $\mathbb{H}^3$).

2.3 E^3 GEOMETRY: $S^1 \times S^1 \times S^1$

This section duplicates some of the material in the last chapter, and can be omitted if you are reading this dissertation in order. It is included here to make this chapter independent from Chapter 1, to contrast with the other geometries, and to give an alternate description of the process of creating subdivision rules from manifolds.

The first geometry we examine is Euclidean 3-space. We use $\mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$ as a representative manifold. This is the most familiar geometry, and it is perhaps the nicest geometry for visualizing universal covers. Its subdivision rule is shown in Figure 1. A fundamental domain for $S^1 \times S^1 \times S^1$ is the cube, with opposite faces identified, as shown in Figure 2.2.

The universal cover of $S^1 \times S^1 \times S^1$ is $\mathbb{R}^3$, and copies of the fundamental domain tile three-space. In this tiling, four cubes come together at an edge. We describe this by saying each edge has edge cycle length 4. In general, when edges of a fundamental domain are glued together by a map, the number of edges in the equivalence class of an edge e is the **edge cycle length** of e.

To find a subdivision rule, let's begin by constructing $S(n)$ for $S^1 \times S^1 \times S^1$. $S(1)$ is the projection of the cube. See Figure 2.3.

Adding a cube to each face, we get $S(2)$. See Figure 2.4. Dotted lines represent corners, or edges common to three copies of the fundamental domain.

This gluing is symmetric, so every face in $S(1)$ is replaced in the same way when creating $S(2)$. That is, each square face in $S(1)$ has a new cube glued on; the five remaining unglued faces of the cube form a new cell structure replacing the old square.

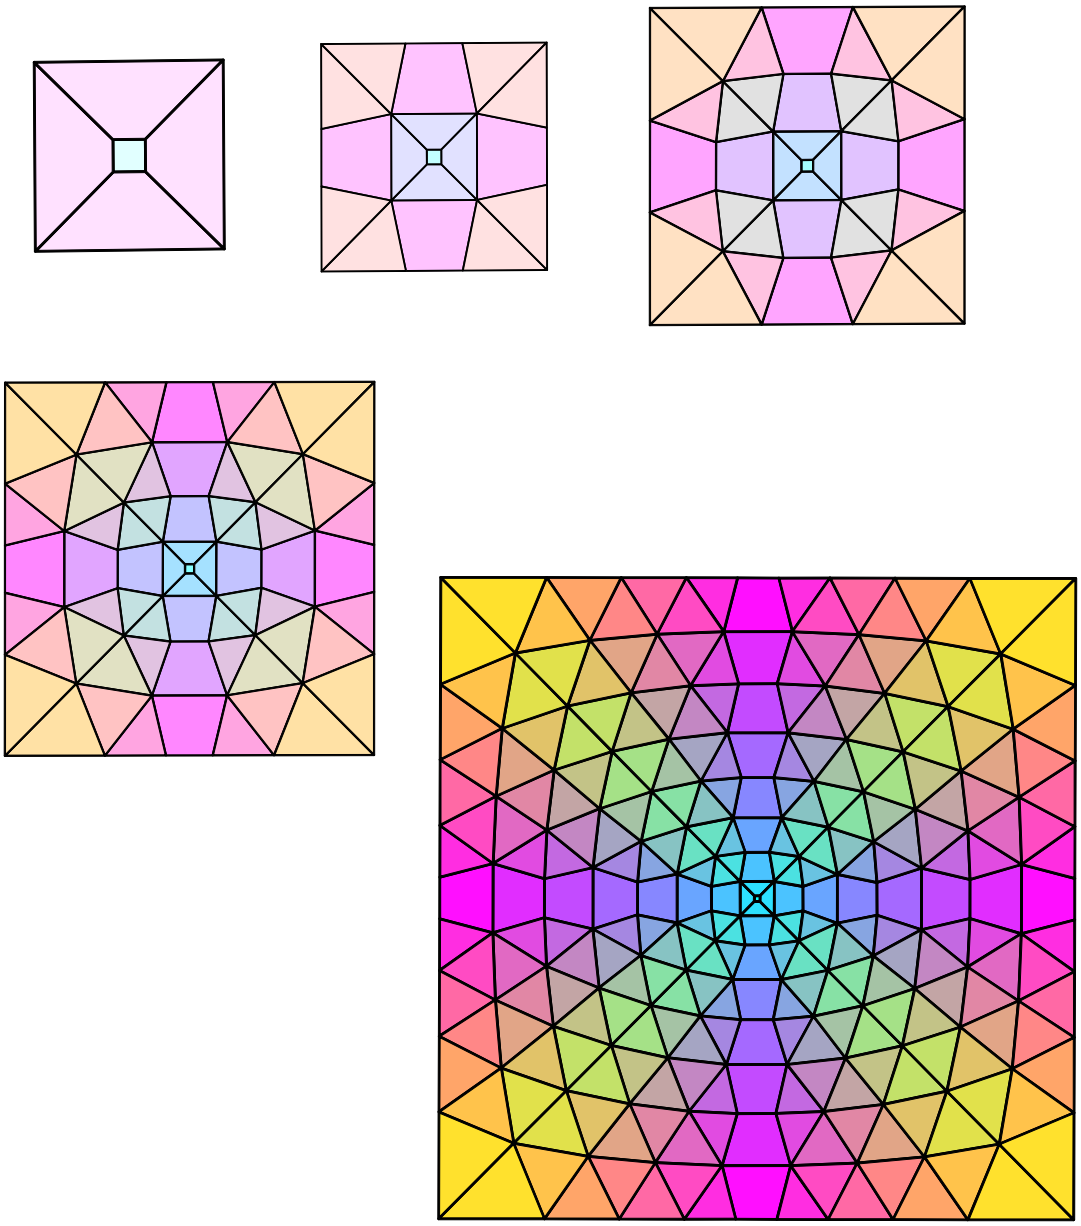

Figure 2.1: Several stages of the subdivision rule for the 3-torus.

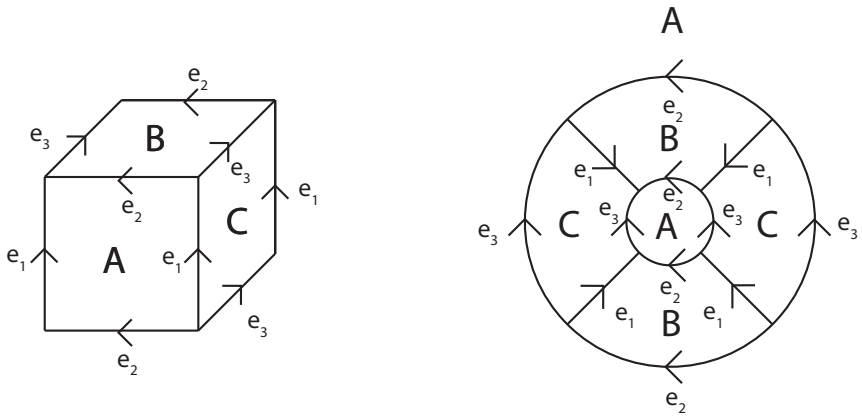

Figure 2.2: The gluing map for $S^1 \times S^1 \times S^1$.

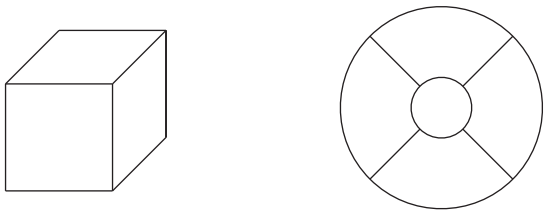

Figure 2.3: S(1)

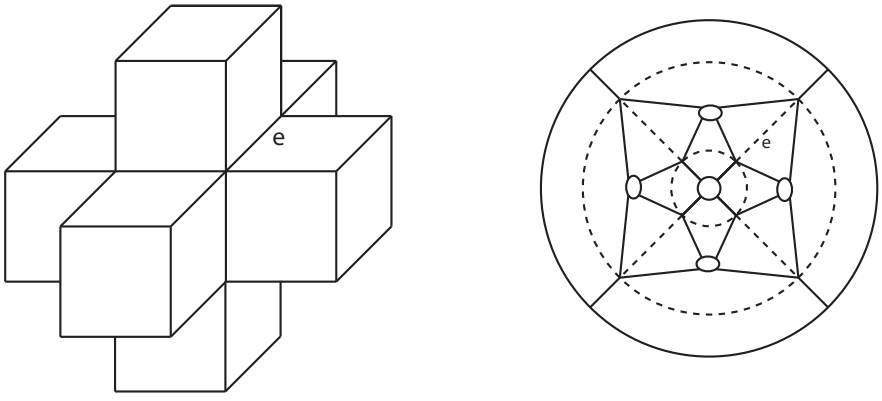

Figure 2.4: S(2)

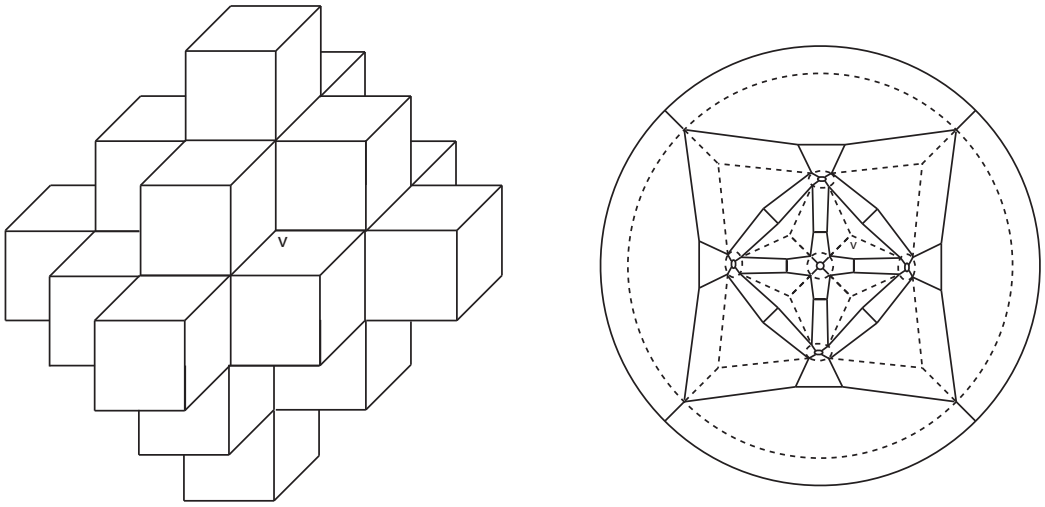

Figure 2.5: S(3)

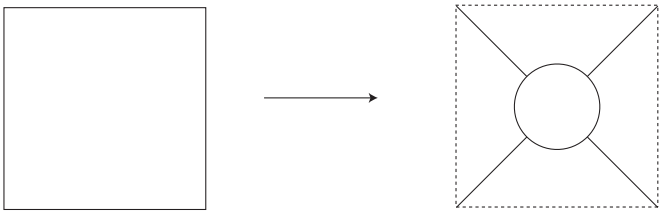

Figure 2.6: The replacement rule for a single face.

$S(3)$ is more complicated (see Figure 2.5).

Notice that, in creating $S(3)$, every face in $S(2)$ without dotted edges is replaced just as the original faces of $S(1)$ were. See Figure 2.6.

However, pairs of faces which share a dotted edge are replaced by a single polyhedron as in Figure 2.7. We will call the dotted edges, which correspond to corners, **loaded edges**.

Notice the loaded edge disappears in $S(3)$. This corresponds to a corner in $S(2)$ being covered up by a single polyhedron (see Figure 2.8). We glue a single polyhedron onto two faces because, in the universal cover, every edge should touch four cubes or copies of the fundamental domain. When an edge is loaded, it already touches three fundamental domains, and so only one more can touch that edge. Thus, one cube most glue onto both faces.

A third situation occurs in going from $S(3)$ to $S(4)$, where three loaded lines converge at

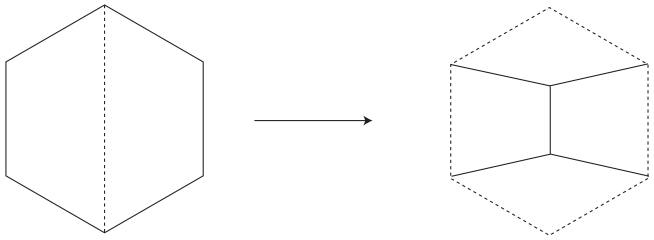

Figure 2.7: The replacement rule for a pair of faces.

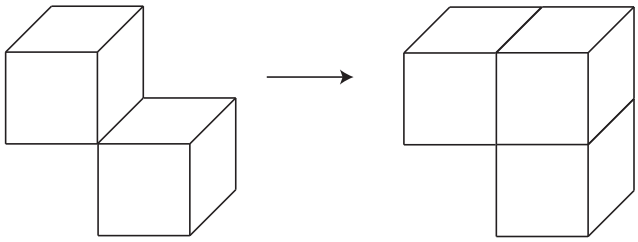

Figure 2.8: The 3D version of Figure 1.10.

a single vertex. We will call this a **loaded vertex**. In this case, a single polyhedron covers up all three, as shown in Figure 2.9. Again, this happens to satisfy the edge cycle length condition.

This corresponds to the situation in Figure 2.10.

Notice now that in $S(3)$, there are no new combinations of tiles; unloaded faces, two faces sharing a loaded edge and three faces sharing a loaded vertex are all that ever happen. Notice that this is not a subdivision rule; edges are created, disappear, reappear, etc. However, we have a replacement rule, as we know how to replace every combination of faces that appear

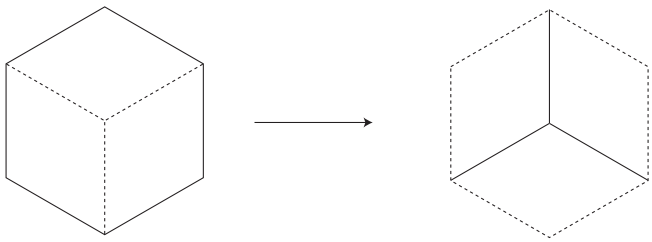

Figure 2.9: The replacement rule for three faces.

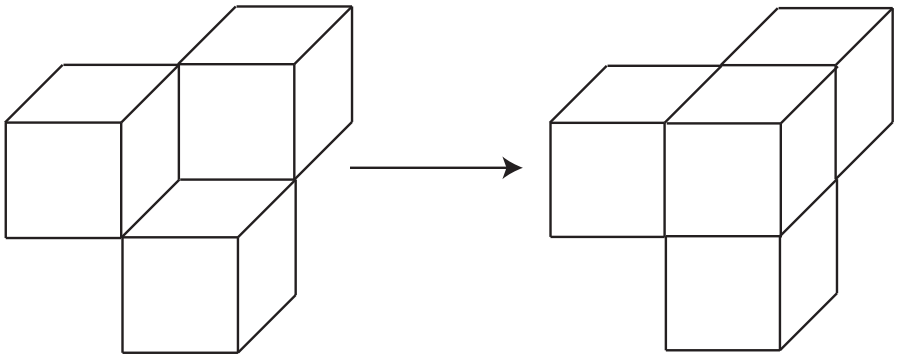

Figure 2.10: The 3D version of Figure 1.11

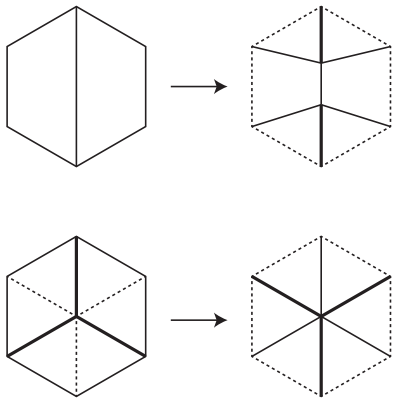

Figure 2.11: We can add lines to the two 'loaded' tile types to get a self-consistent subdivision rule.

at every stage. This can be turned into a subdivision rule by adding new edges at every stage. For instance, in Figure 2.7, 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 2.11), then the new cell structure contains the old cell structure as a subset. Thus, we have a subdivision rule.

However, this divides the top and bottom squares into two triangles each; each of these surround a loaded vertex, so we have to change the replacement rule for three squares surrounding a loaded vertex (call the three faces a **loaded star**); however, notice that adding the three lines in to the loaded star on the left of Figure 2.9 gives us a hexagon

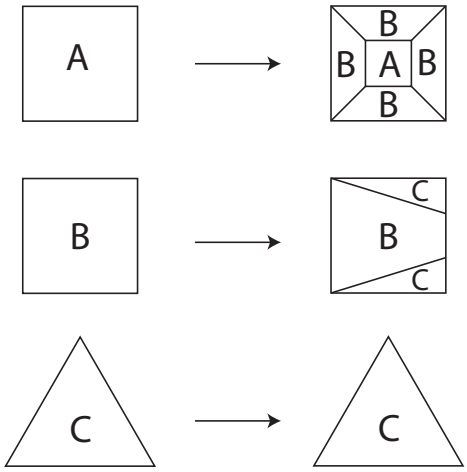

Figure 2.12: The replacement rule for the 3-dimensional torus.

divided into six 'pie slices'. If we add similar lines bisecting the loaded star (as seen in Figure 2.11), we again get a hexagon divided into six triangles; thus, the subdivision on each triangle in that hexagon is just the identity.

We summarize this in Figure 1.13. The circle packed pictures are shown on page 35. These circle packed pictures only display the subdivision rule combinatorially; the circle packed pictures are not subsets of each other, because this subdivision rule is not conformal. The connection between circle packings and conformality is explained in [11]. Notice the similarity of this subdivision rule to the subdivision rule for the Hopf link (a Euclidean knot) from [18].

2.4 $\mathbb{H}^2 \times \mathbb{R}$ GEOMETRY: $N_{-1} \times S^1$

In this section, we study the product geometry $\mathbb{H}^2 \times \mathbb{R}$, with example manifold $N_{-1} \times \mathbb{S}^1$.

Here we are using N_{-1} to represent the non-orientable surface of Euler characteristic -1, i.e. the connected sum of three projective planes. We chose this surface as our example because every other hyperbolic surface can be pieced together from it. In this sense, it is the smallest hyperbolic surface. We also chose this example because it introduces the notion of 'fragile edges', which come from manifolds with edges of odd cycle length. The circle packed

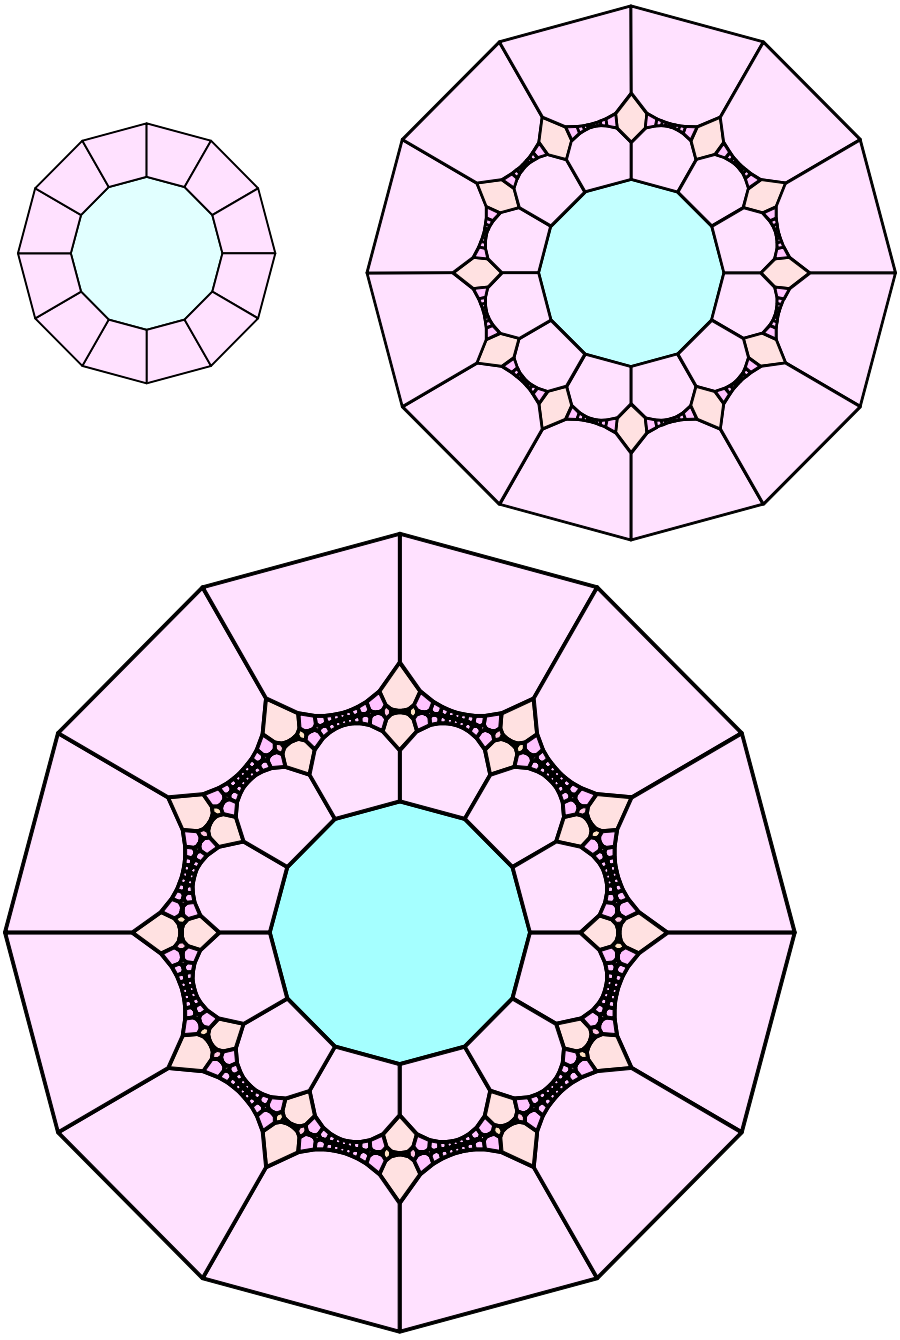

Figure 2.13: Several stages of the subdivision rule for an $H^2 \times \mathbb{R}$ manifold.

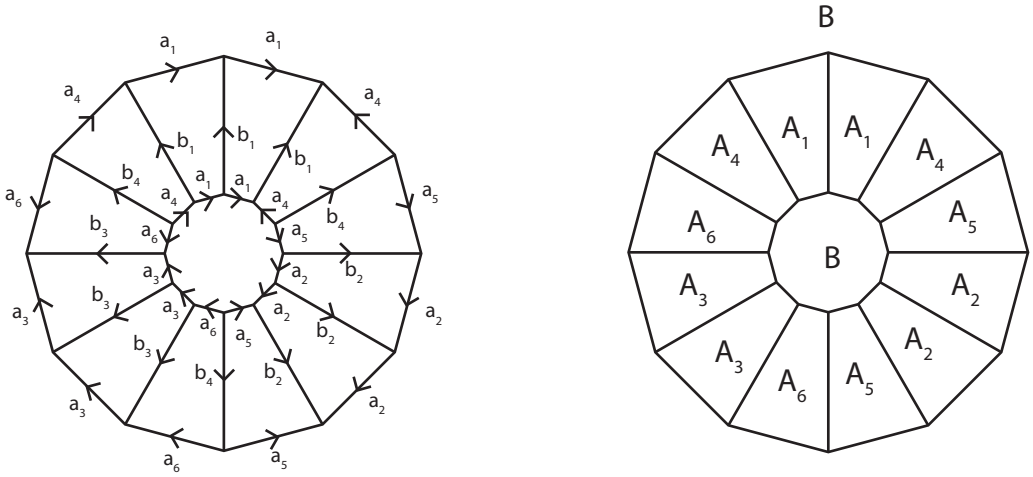

Figure 2.14: The gluings for our $N_{-1} \times S^1$ manifold.

picture of the subdivision rule we will get is shown in Figure 2.13.

The 3-manifold $N_{-1} \times S^1$ has a fundamental domain that is a dodecagonal prism (see Figure 2.14).

Notice that each a_i edge has edge cycle length 3, and each b_i edge has edge cycle length 4. This particular gluing was chosen to give these cycle lengths. Some of the face-gluing maps are orientation reversing, and others are orientation preserving. But, in creating $S(n)$, changing the orientation won't change the combinatorial structure of the fundamental domains. This is because the prism has symmetry group $D_{12} \times Z_2$ and preserves its shape under reflection.

We again let $B(1)$ be a single fundamental domain with $S(1)$ its boundary. Now that we are dealing with hyperbolic space, the balls $B(n)$ and spheres $S(n)$ are more difficult to imagine. But we can describe how each face is replaced, as we did for $S^1 \times S^1 \times S^1$, earlier. Each face should be replaced by the specific projection of our fundamental domain that sends our current face to infinity.

We call all dodecagonal faces type B, and all square faces type A. A face of type B is replaced as in Figure 2.15.

Similarly, a type A face is replaced as in Figure 2.16.

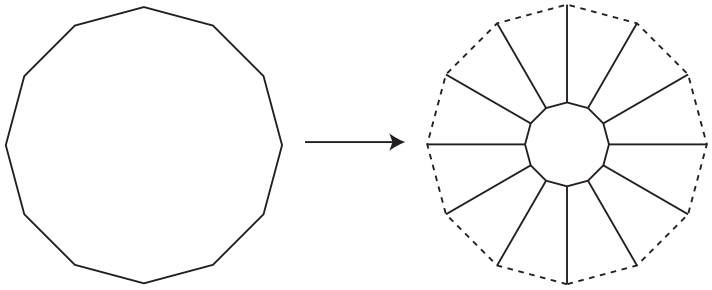

Figure 2.15: Type B face.

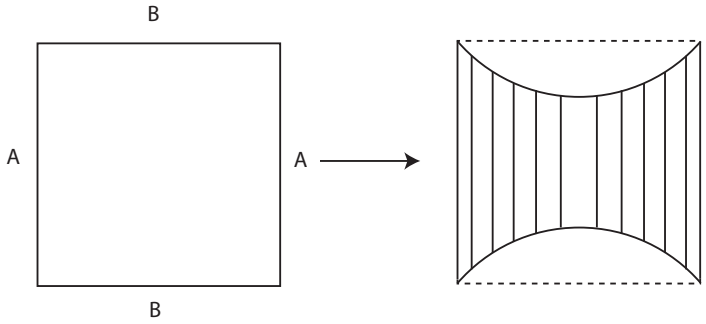

Figure 2.16: Type A face.

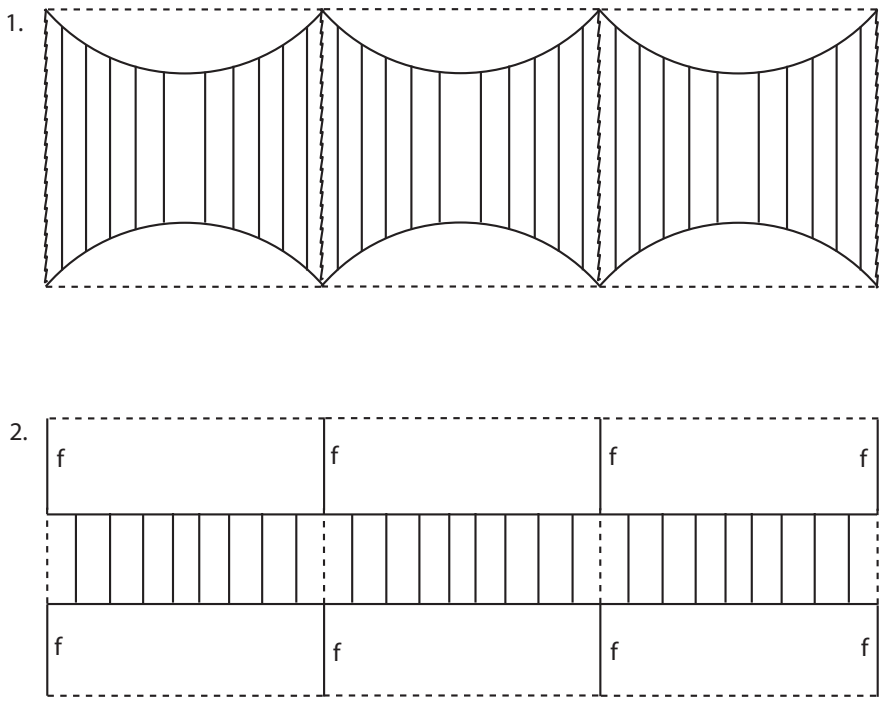

Figure 2.17: 1. The squiggly lines represent edges about to collapse. 2. After collapse. The edges marked with f are now 'fragile' (see text for definition).

However, here we have a problem. The a_i edges between A faces have edge cycle length 3. Thus, only three polyhedra can touch that edge in the universal cover. In $B(2)$, three polyhedra are already touching a_i: the original polyhedron, and the two glued onto either side of the edges. There are now two unglued faces on either side of the a_i edge; if we glue new polyhedra to either one, we'll have too many polyhedra coming in at an edge. Thus, since these two faces must be glued to something, and there is no other choice that will keep our $B(2)$ simply connected, we must glue them to each other, as shown in Figure 2.17.

The essential point here was that only two more polyhedra could touch the edge, causing the faces bordering that edge to collapse. Any edge e with edge cycle length L that already touches $L-2$ polyhedra will be called **fragile**. Edges can be become fragile over time, just as they can become loaded. Note that burdened edges are different from fragile edges.

In our replacement of the A tiles above, all of the edges that were brought together now touch two polyhedra. Thus, edge cycle length 3 edges are loaded, and edge cycle length 4

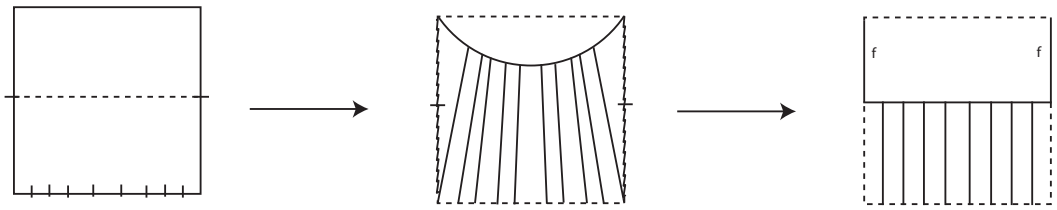

Figure 2.18: The replacement rule for an A/B pair.

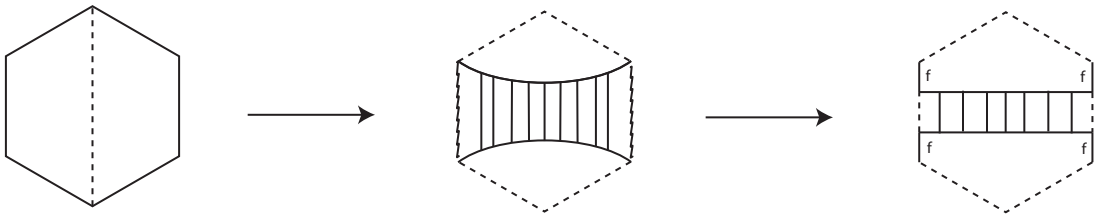

Figure 2.19: The replacement rule for an A/A pair.

edges are fragile.

These replacements have created new types of faces to deal with. Specifically, we now have: loaded pairs of B's and A's, where the B portion has two fragile edges, as well as loaded pairs of A's.

The A/B pairs are replaced as in Figure 2.18.

Notice that the leftmost and rightmost regions collapse along two fragile edges each. It is possible that this would cause confusion, as each of those faces must be identified to both faces that it borders over those two edges. However, looking carefully, we see that A/B pairs border only other A/B pairs on either side, so that the collapsing is well-defined. Note that this covers up an old vertex and adds a new one. For convenience, we've placed the new vertex directly over the old.

Pairs of A's behave as in Figure 2.19.

Together, both kinds of pairs create loaded triples consisting of two A's and a B, where the B has two fragile edges. These triples are subdivided as in Figure 2.20.

Again, the leftmost and rightmost regions collapse over two edges each. No new types of

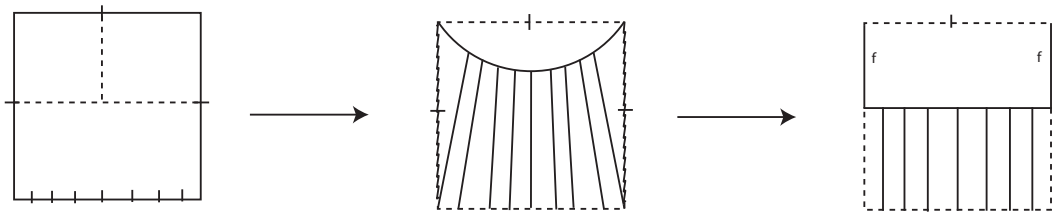

Figure 2.20: The replacement rule for an A/A/B triple.

faces have occurred, so we have found all of the replacement rules.

To convert this to a subdivision rule, we must add lines to make the new tiles a subset of the replaced tiles. These lines are added as shown in Figure 2.21. This gives us a subdivision rule. The circle packings of this subdivision rule are shown in Figure 2.13. This subdivision rule is very similar to the subdivision rule for the trefoil knot (and other 2-braid knots besides the Hopf link) as described in [18].

2.5 $\mathbb{H}^3$ geometry: Hyperbolic dodecahedral space

Subdivision rules were originally discovered for hyperbolic geometry. Cannon and Swenson's work [7] shows that all subdivision rules coming from hyperbolic 3-space are conformal.

In 1, we found replacement rules for infinitely many hyperbolic manifolds, many of which can be refined into subdivision rules. We required our polyhedra to have even edge cycle length for all edges. One example is shown in Figures 2.22 and 2.23. In general, the subdivision rules have mesh going to 0, and are the type of finite subdivision rules studied extensively by Cannon et. al., for instance in [6]. These subdivision rules are very similar in character to the subdivision rules of the hyperbolic alternating links found in [18].

2.6 $S^2 \times \mathbb{R}$ geometry: $S^2 \times S^1$

Subdivision rules for $S^2 \times \mathbb{R}$ manifolds are different than all others. For one things, the universal cover has two boundaries, which can be thought of as an 'inner' and an 'outer'

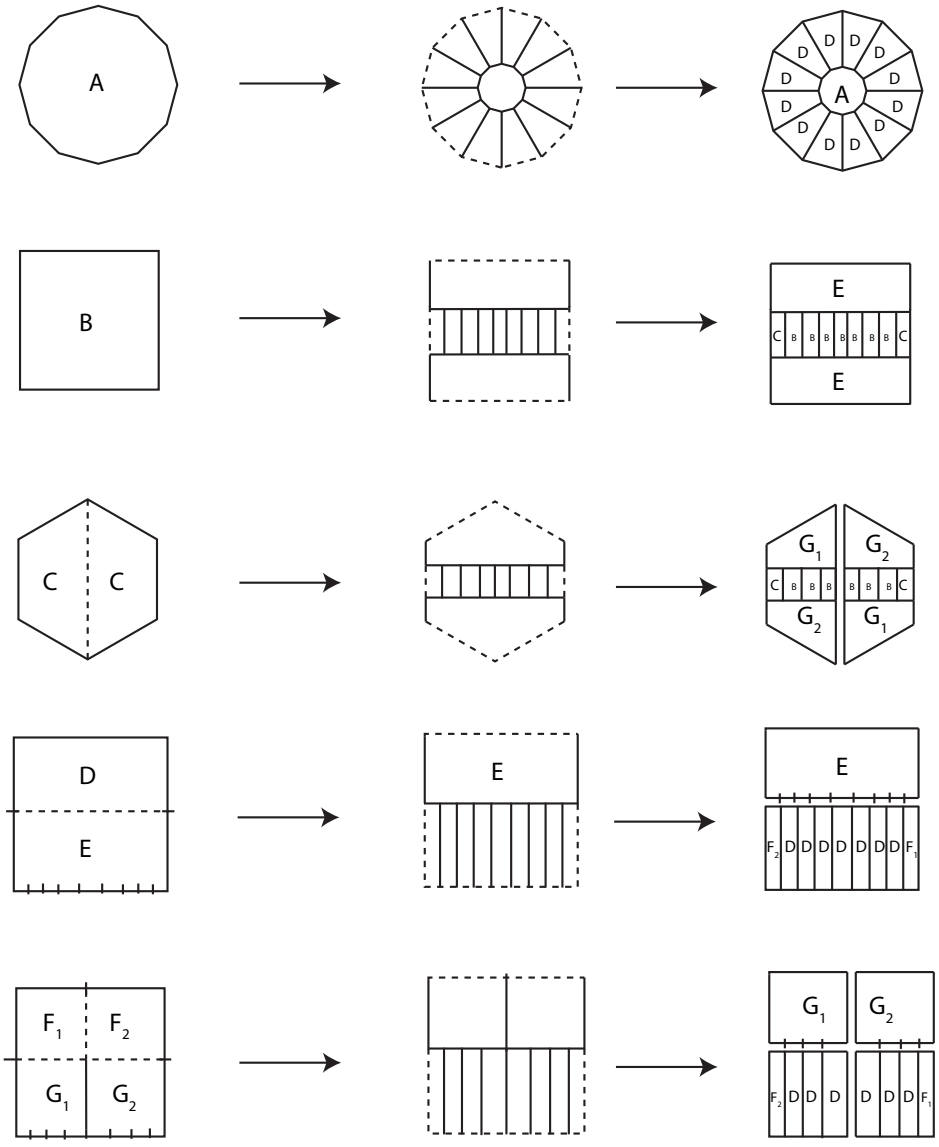

Figure 2.21: The subdivision rule for M.

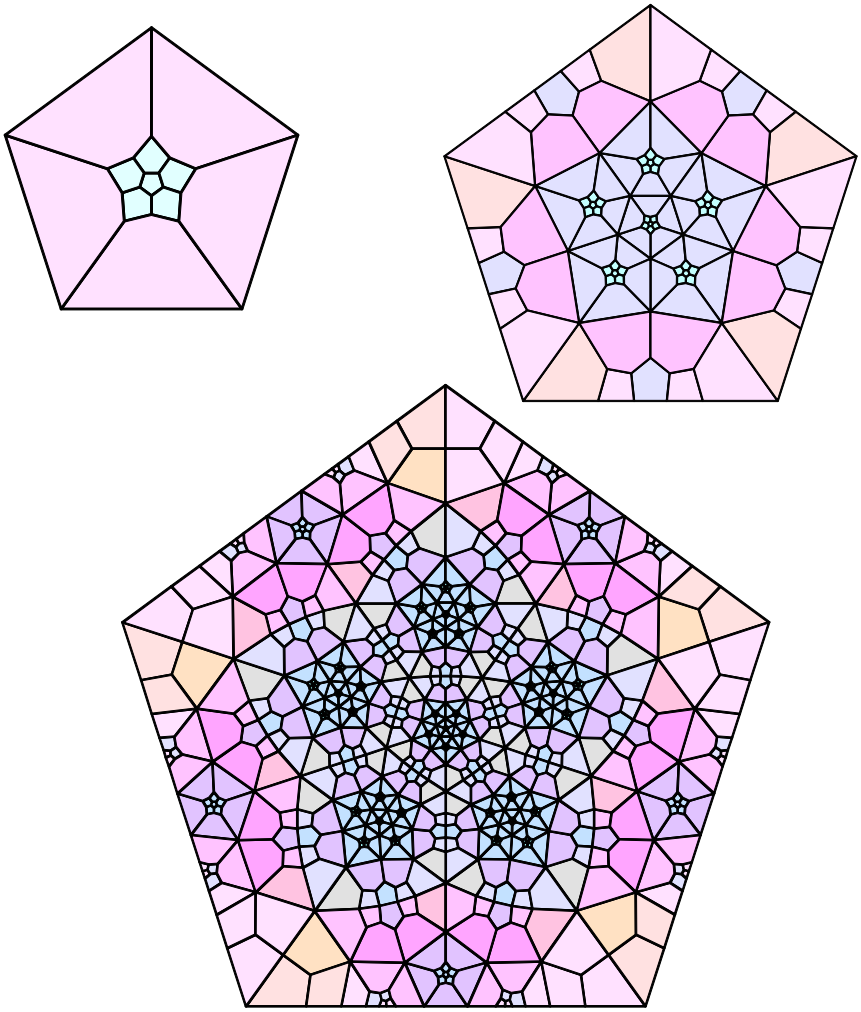

Figure 2.22: Several stages of the subdivision rule for a hyperbolic 3-manifold.

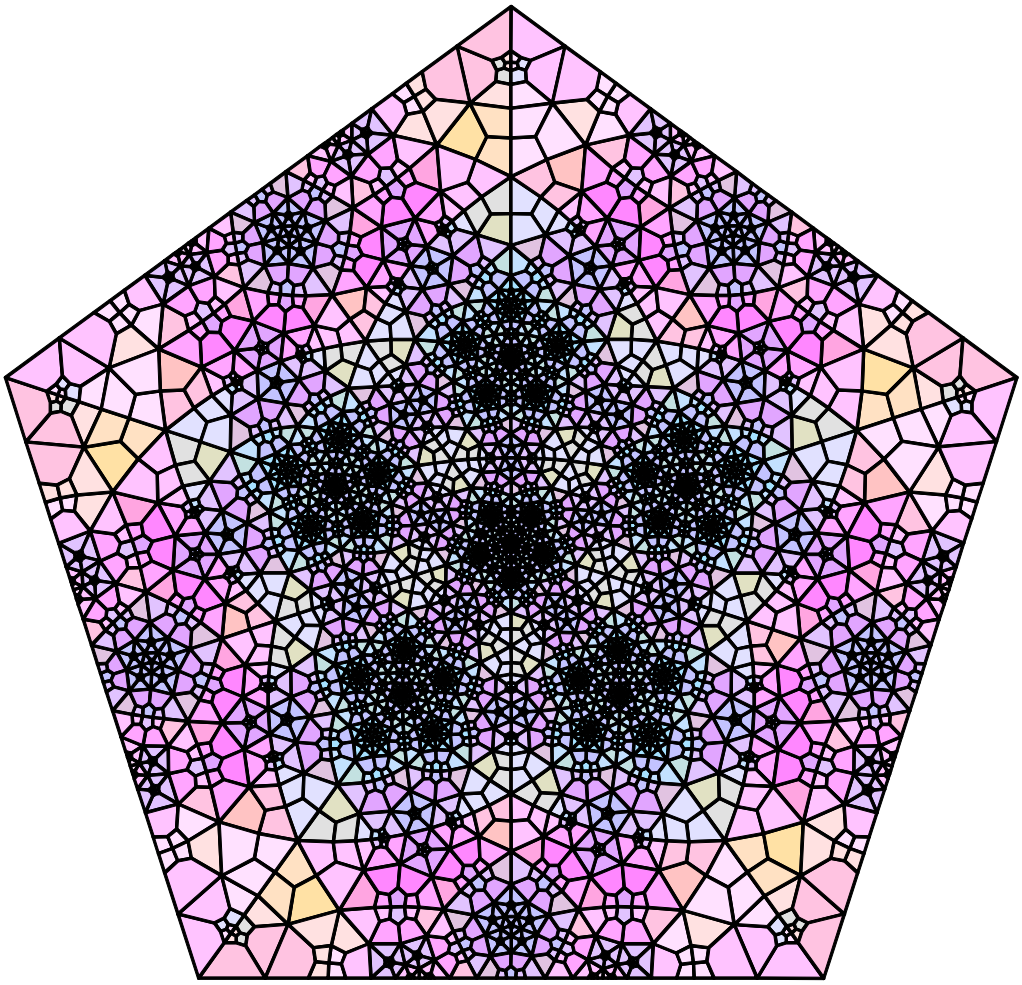

Figure 2.23: One further stage of the subdivision rule for a hyperbolic 3-manifold.

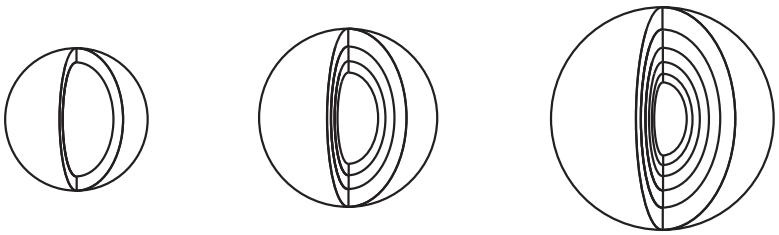

Figure 2.24: The first three stages in constructing the universal cover of $S^2 \times S^1$.

sphere at infinity.

Consider $S^2 \times S^1$. This has as a fundamental domain $S^2 \times I$, a thickened sphere. Constructing the universal cover amounts to nesting these thickened spheres. See Figure 2.24.

The figures for $S^2 \tilde{\times} S^1$ and $P^3 \oplus P^3$ are similar, as they also have a thickened sphere for their fundamental domain. $P^2 \times S^1$ is only slightly more complicated. In all of these cases, there are a fixed number of faces on each boundary sphere at each stage of subdivision, giving us a sequence of tilings $S(n)$ which is constant. The subdivision rule, then, is the identity on both spheres.

2.7 $\tilde{SL}_2(\mathbb{R})$ geometry: The unit tangent bundle of N_{-1}

Let's now consider the unit tangent bundle of N_{-1}. We choose this particular example of an $\tilde{SL}_2(\mathbb{R})$ manifold for the same reason we chose $N_{-1} \times S^1$ as an example for $H^2 \times \mathbb{R}$ earlier: it is covered by all other unit tangent bundles of hyperbolic surfaces, and these are the best known manifolds corresponding to this geometry. For convenience, we will refer to the unit tangent bundle of N_{-1} as M throughout this section.

How can we find a polyhedral fundamental domain for M? We can project M onto N_{-1}. Slice N_{-1} as we did back in Figure 2.14 and slice M along the pre-images of these edges. What remains is the unit tangent bundle of a closed disk. This is necessarily the trivial bundle, and so we can slice it along a fixed, 'horizontal' surface to get the product of a closed disk and the closed interval. But this is just a 3-ball, and so we have our polyhedron. See Figure 2.25.

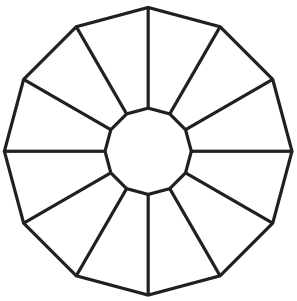

Figure 2.25: The manifold M after slicing.

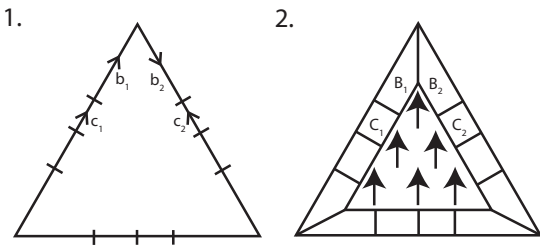

Figure 2.26: The choice of vector field for our horizontal slice.

To get the manifold back, or to understand the combinatorics, we need to understand how the edges and faces are glued. The top is glued to the bottom, but what about the square faces?

Remember that we got a polyhedron by slicing the unit tangent bundle of the disk along a horizontal surface. A horizontal surface corresponds to a fixed unit vector field $\alpha(x, y)$ on the disk. Let's make the choice of vector field explicit, as shown in Figure 2.26. Notice that we have redrawn our dodecagon on as triangle.

Then we make a correspondence between the height of a point in the polyhedron and the angle in the unit tangent bundle. See Figure 2.26. If the top is 0 and the bottom is 2π (viewing the closed interval as $[0, 2\pi]$), then the angle at a point (x, y, z) is $\alpha(x, y) + z$. This enables us to define the gluing map.

Look at the faces B_1 and B_2. They are identified together because their projections in N_{-1} are edges that are identified. Now look at the edge a_1. Relative to the face B_1, the unit vector field comes in at an angle of $\frac{\pi}{3}$, at the top of B_1 (i.e. at the edge a_1) , and at an angle

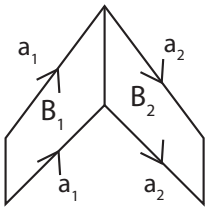

Figure 2.27: The faces B_1 and B_2 and the edges a_1 and a_2.

Figure 2.28: The angles that the vector field makes with the edges a_1 and a_2 are $\frac{\pi}{3}$ and $\frac{2\pi}{3}$.

of $\frac{2\pi}{3}$ at the top of B_2 (i.e. the edge a_2). See Figure 2.28. This means that a_1 gets sent to a line on B_2 of depth $\frac{\pi}{3}$. Similarly, a_2 gets sent to a line on B_1 with depth $\frac{5\pi}{3}$, so the faces B_1 and B_2 get mapped together as shown in Figure 2.29.

Faces C_1 and C_2 are slightly different, having opposite orientation. See Figure 2.30. The edge d_1 gets sent to a line of depth $\frac{4\pi}{3}$, while d_2 gets sent to a line at depth $\frac{2\pi}{3}$. This is shown in Figure 2.31.

Continuing all along, we get the gluing map shown in Figure 2.32. This can be simplified

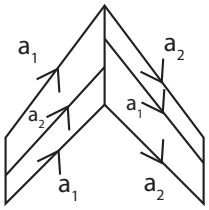

Figure 2.29: The depth in the z direction measures the change tangent direction from the horizontal field at the top. Since the vector field at a_1 comes in at an angle $\frac{\pi}{3}$ more than the vector field at a_2, it gets mapped to a line at depth $\frac{\pi}{3}$. Similarly, the edge a_2 gets mapped to a depth of $\frac{5\pi}{3}$.

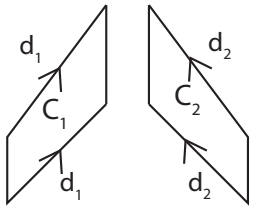

Figure 2.30: The faces C_1 and C_2 and the edges d_1 and d_2.

Figure 2.31: The angles that the vector field makes with the edges d_1 and d_2 are $\frac{4\pi}{3}$ and $\frac{2\pi}{3}$.

further by taking the resulting 'cylinder' and slicing horizontally, adding new lines. See Figure 2.33. Each 'vertical' edge has edge cycle length 3, and each horizontal edge has edge cycle length 4.

Note that the three resulting polyhedra are combinatorially identical to the polyhedron we used for $N_{-1} \times S^1$ in Section 2.4. However, in $N_{-1} \times S^1$, each polyhedron corresponded to a single group element, while in our manifold M, a group element corresponds to three polyhedra. Also, in the universal cover of M, neighboring group elements in the 'horizontal' direction are not at the same height; the middle layer of one group of three polyhedra will get glued alternately to the top and bottom layers of the neighboring groups of three.

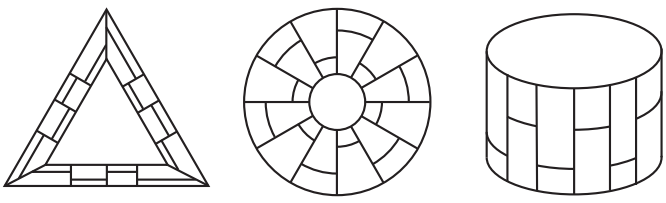

Figure 2.32: The gluing map for M. The first two are just homotopies of each other, and the last is a 3-d figure of the fundamental polyhedron.

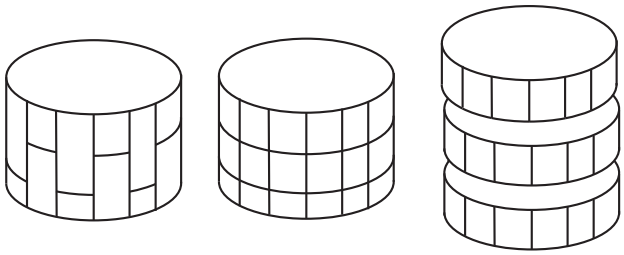

Figure 2.33: We slice the fundamental polyhedron into three smaller cylinders. This makes each horizontal edge have valence four and each vertical edge have valence three.

The reason that manifolds from these two geometries can have the same subdivision rules is that the two geometries are quasi-isometric.

2.8 The other geometries: Nil, Sol, and $\mathbb{S}^3$

There are no known subdivision rules for the last three geometries.

As all S^3 manifolds have finite fundamental group, the boundary is eventually empty, and there cannot be a subdivision rule. This is an example of how manifolds that finitely cover each other have essentially the same subdivision rule. Each manifold in this geometry is finitely covered by $\mathbb{S}^3$, so we can take this as our example manifold; this manifold has empty boundary, and the subdivision rule is the identity on the empty set.

One can obtain polyhedral gluings for both Nil and Sol manifolds by taking a thickened torus from a cube and identifying the inside and outside by a linear map (see Figures 2.34 and 2.35). However, Sol groups are not almost convex [9], and this seems to imply that they cannot have subdivision rules. On the other hand, Nil groups are almost convex [20], but finding a subdivision rule from the gluings we have seems to be quite difficult. The polyhedral gluings in [14] may help.

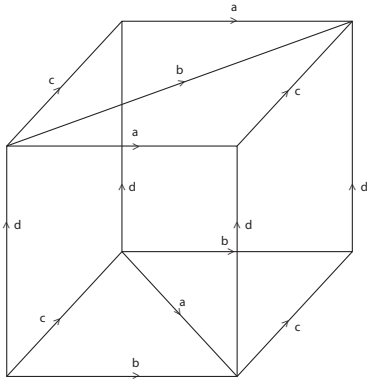

Figure 2.34: A fundamental domain for a nil manifold

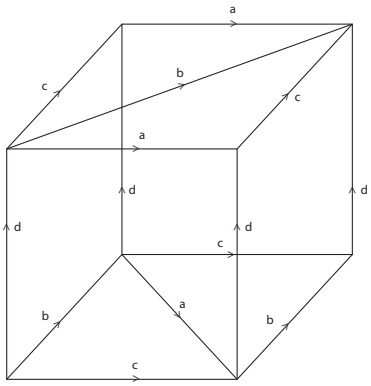

Figure 2.35: A fundamental domain for a solv manifold

Chapter 3. Creating Subdivision Rules from the n-Dimensional Torus

3.1 Preliminaries

In this chapter, we find subdivision rules for the n-dimensional torus. We define a subdivision rule in higher dimensions in a way analogous to subdivision rules in dimension 2 (following [6]). Thus, a **subdivision rule of dimension** n consists of:

(i) A finite n-dimensional CW complex which is the union of its closed n-cells, and which has the property that each closed k-cell ($k \leq n$) is the image under the characteristic map of a closed k-disk with a cell structure including all sub-cells of the image, the characteristic map restricting to a homeomorphism on each open cell. The complex is called the **subdivision complex** S_R.

(ii) A subdivision $R(S_R))$ of $S(R)$.

(iii) A **subdivision map** $\sigma_R : R(S_R) \to S_R$,which is a continuous cellular map that restricts to a homeomorphism on each open cell.

The n-cells of the subdivision complex are called the **tile types**. Essentially, a subdivision rule is just a way of refining a cell structure by a local rule that agrees on sub-cells of neighboring, larger cells.

Although we will only find subdivision rules for the n-dimensional torus, our method may possibly apply to other higher-dimensional manifolds. Our approach, as in Chapters 1 and 2 is based on the universal cover. The universal cover of any manifold can be constructed recursively by taking a copy of the fundamental domain, gluing on fundamental domains to every exposed face of the original, and repeating. More specifically, let $B(0)$ be a single copy of the fundamental domain of an n-manifold M. Let $B(k)$ be the set of all fundamental

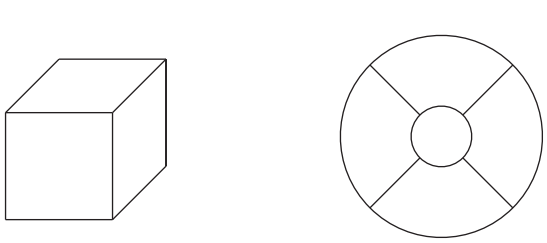

Figure 3.1: S(1)

domains that are distance $\leq k$ from $B(0)$ (in the word metric). Then for many groups and choices of generating sets, $S(k) = \partial B(k)$ will be a topological $(n-1)$-sphere for all k or for k sufficiently large. The cell structure from the fundamental domain gives a cell structure to $B(k)$ and thus to $S(k)$. This cell structure is a tiling. Thus, we get a sequence of tilings in which every tile or every group of tiles corresponds to an element of the fundamental group, and the entire group is represented at some point. We have drawn $S(1), S(2)$ and $S(3)$ for the 3-dimensional torus with the standard choice of generators in Figures 3.1 to 3.3, shown in 3-space and also as a combinatorial tiling. We are again duplicating this material (as we did in Section 2.3), as it provides the pattern for finding the subdivision rule in all other dimensions.

However, this sequence of tilings for a manifold is not necessarily created by a subdivision rule, because faces and edges are created and later covered up (it's a replacement rule). To get a recursive structure, similar to hyperbolic 3-manifolds, we need to find a way to represent $S(k)$ (or a slightly modified version of it) as a sub-tiling of $S(k+1)$ (or a modified version of it).

3.2 The n-torus

We now show how to obtain a subdivision rule for the n-dimensional torus.

In the discussion that follows, let $I = [0, 1]$, the unit interval. A q-cube is I^q, and a p-simplex is the convex hull of $p+1$ points in general position. Thus, a 1-cube is a line, a 2-cube is a square, and a 3-cube is a (standard) cube; a 1-simplex is a line, a 2-simplex is a

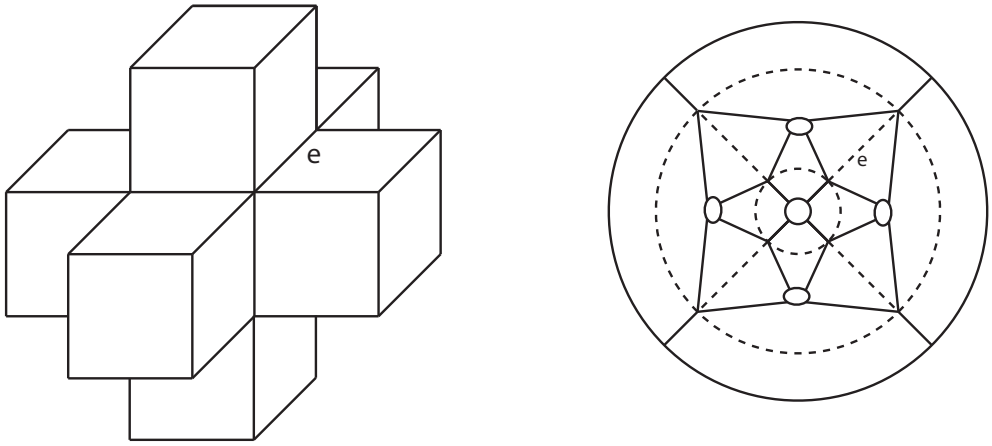

Figure 3.2: S(2)

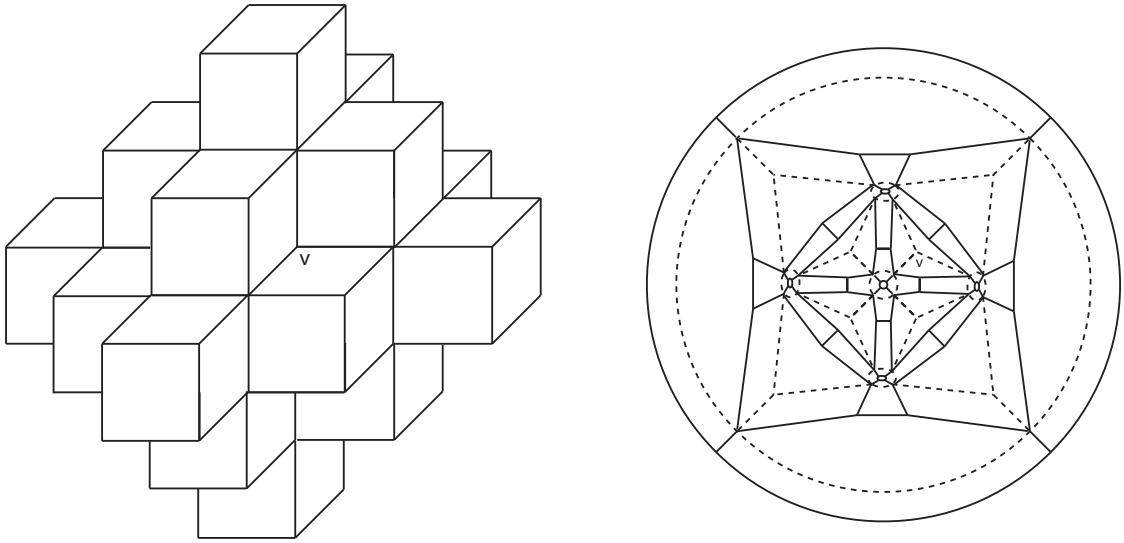

Figure 3.3: S(3)

triangle, and a 3-simplex is a tetrahedron, etc.

Theorem 4. The n-torus has a subdivision rule with n tile types. The tile types are $p-1$-simplices cross q-cubes, where $1 \leq p \leq n$ and $q = n - p$. Each such tile is subdivided into one $p-1$ simplex cross a q-cube and $2q$ p-simplices cross $q-1$ cubes.

Before we begin the proof, look at Figures 3.4 and 3.10 to see the tile types for $n = 1, 2, 3$ and 4.

Proof. The fundamental domain of the n-torus $\mathbb{T}^n = (S^1)^n$ is a hypercube of dimension n. If the generators of the fundamental group are $y_1, ..., y_n$, then every element of the fundamental group can be written uniquely as $y_1^{a_1} y_2^{a_2} ... y_n^{a_n}$.

Because our group is free abelian, the Cayley graph of the subgroup generated by any subset of the generators is contained in the Cayley graph of the fundamental group. Thus we can build the universal cover of these manifolds inductively from the universal covers of manifolds corresponding to subgroups.

We now describe how to explicitly construct the subdivision rule. It may help to follow along with the examples $n = 1, 2, 3$ and 4 starting on page 68.

To construct the universal cover, we start with a single n-cube (i.e. I^n) and begin gluing on other cubes. Faces (or cells of codimension one) correspond to generators and inverses of generators. Assume an element represented by a cube is being glued on in some stage of creating the universal cover. Assume the element can be written as $y_{k1}^{a_1} y_{k2}^{a_2} ... y_{kp}^{a_p}$, where this is a representation of minimal word length (so $1 \leq k_1 \leq k_2 \leq ... \leq k_p \leq n$ and $a_i \neq 0$). Then this element is contained in a subgroup of rank p. Let $q = n - p$. Then gluing on the cube corresponding to this element is accomplished by identifying some of its boundary to the previous stage of the universal cover. If we write the cube I^n as $I^p \times I^q$, the boundary will be $\partial I^p \times I^q \cup I^p \times \partial I^q$.

Now, because the group element has p geodesic paths into it (for instance, if a_i¿0, going to $y_{k1}^{a_1 - 1} y_{k2}^{a_2} ... y_{kp}^{a_p}$ and then going through the y_{k1}-face to our element), our cube is glued onto p faces at once. Each of the p faces represents a generator, and if one generator is represented,

its inverse is not, meaning that no pair of opposite faces is in the set. The structure of the n-cube is such that every set of p faces containing no opposing pairs determines a unique q-cell which is common to all of them (so, for instance, in a 3-cube, three non-opposing faces intersect in a vertex, two in a line, and one in a square). If we project $I^n \subseteq \mathbb{R}^n$ down onto the subspace orthogonal to this cell, we see that this set of faces projects to the star of a vertex in ∂I^p. Call this star S. Note also that every vertex in the p-cube has an opposite vertex, and the star of a vertex and its opposite have disjoint interiors and cover ∂I^p. Call the star of the opposite vertex S^*.

Thus, in gluing on I^n via ∂I^n, we glue the boundary onto $A = S \times I^q$. The faces of the ∂I^n that are not glued to anything can be written as $B = B_1 \cup B_2 = S^* \times I^q \cup I^p \times \partial I^q$. Recall that, to find a subdivision rule, we look at $S(k)$ (i.e. all exposed faces at stage k of constructing a universal cover), and $S(k+1)$ (all exposed faces at stage $k+1$),and try to find the first as a subset of the second. Therefore, our goal is to find a cell structure for A and B such that B is a refinement or subdivision of A. We use a standard simplicial decomposition of the p-cube (found, for instance, in [16], Exercise 10.18).

We now describe the decomposition. I^p is covered by the $p!$ simplices $\{[0, e_{\sigma(1)}, e_{\sigma(1)} + e_{\sigma(2)}, e_{\sigma(1)} + e_{\sigma(2)} + ... + e_{\sigma(p)}]|\sigma \in \Sigma_p\}$, each of which has disjoint interior. Here, e_i is the unit vector in the i-th direction. The symbol $[p_0, p_1, ..., p_k]$ is defined to be $\tau(Q^k)$, where τ is the affine map $\tau(x_1, ..., x_k) = p_0 + \Sigma x_i(p_i - p_0)$, and Q^k is the standard simplex $\{(x_1, ..., x_k)|0 \leq x_i$ for all i, $x_1 + ... + x_n \leq 1\}$. Each of these simplices has sub-simplices defined by deleting intermediate terms (so $[0] \subseteq [0, e_1] \subseteq [0, e_1, e_2]$, for instance).

Recall that switching two terms in the simplex (i.e. changing $[p_0, p_1, p_2]$ to $[p_2, p_1, p_0]$) gives a different map from Q^k with opposite orientation but the same image as the original map. If τ_1 and τ_2 are the map corresponding to the original simplex and the 'flipped' simplex, then $\tau_2^{-1}\tau_1$ is an orientation-reversing simplicial map.

We use this to define an involution on our simplicial cube. Define this map by switching 0 and $e_{\sigma(1)} + e_{\sigma(2)} + ... + e_{\sigma(p)}$ in every simplex. This is a simplicial map that is the identity on all

subsimplices not containing 0 or $e_{\sigma(1)}+e_{\sigma(2)}+...+e_{\sigma(p)}$. Any subsimplex that contains one of those points is sent to an opposing subsimplex that contains the other point. The existence of this map shows, in particular, that the set of all closed simplices in ∂I^p containing 0 is simplicially isomorphic to the set of all closed simplices in ∂I^p containing $e_{\sigma(1)}+e_{\sigma(2)}+...+e_{\sigma(n)}$. This means that if S and S^* are given the simplicial structure they inherit from I^p, they are isomorphic.

This means that A and B_1 have the same cell structure. If $q=0$, $B_1=B$ and A and B have the same cell structure, so our subdivision rule can be the identity.

If $q\neq 0$, it's slightly more difficult. We still give A and B_1 the simplicialized structure explained above, and give B_2 the structure of $I^p\times\partial I^q$, where I^p is given the simplicial structure of earlier. We show that B contains A as a subcomplex, with $\partial A\subseteq\partial B$. In the discussion that follows, it will be helpful to follow along with Figures 3.6-3.8 for the case $p=2$, $q=1$. Figure 3.10 gives more examples with less explanation.

So, pick a $(p-1)$-simplex Δ^{p-1} in $S\subseteq\partial I^p$. If we consider the center vertex of S as 0 and the center vertex of S^* as $e_{\sigma(1)}+e_{\sigma(2)}+...+e_{\sigma(p)}$, then there is a unique p-simplex Δ^p defined by adjoining $e_{\sigma(1)}+e_{\sigma(2)}+...+e_{\sigma(p)}$ to Δ^{p-1}. So, our goal is to show that $\Delta^{p-1}\times I^q\subseteq\Delta^{p-1}\times I^q\cup\Delta^p\times\partial I^q$. Patching together these simplices will show that $A\subseteq B_1\cup B_2=B$.

To do so, note that I^q is just the cone over the boundary of I^q (as a set, not as a complex). Thus, we look at $\Delta^{p-1}\times\partial I^q\times I$, which we will eventually collapse. Each face in ∂I^q is a $(q-1)$-cube. Given a specific face, we can embed the product $\Delta^{p-1}\times I^{q-1}\times I$ in $\mathbb{R}^{p+q-1}$ as $\{(x_1,...,x_{p-1},y_1,...,y_{q-1},z)|0\leq x_i \text{ for } 1\leq i\leq p-1, 0\leq y_j\leq 1 \text{ for } 1\leq j\leq q-1, 0\leq z\leq 1x_1+x_2+...+x_{p-1}\}$. Call this set C.

Define a family of maps $f_t:C\to C$ by

$$f_t(x_1,...,x_{p-1},y_1,...,y_{q-1},z)=(x_1,...,x_{p-1},y_1,...,y_{q-1},z(1+(x_1+...+x_{p-1}-1)\frac{t}{2})).$$

This defines an invertible homotopy (basically dragging down the corner of the top copy of

the simplex along the z-axis).

Note that

$$f_1(C) = \{(x_1, ..., x_{p-1}, y_1, ..., y_{q-1}, z) | z \leq \frac{1}{2} + \frac{x_1 + ... + x_{p-1}}{2}\}$$

and this is the same as

$$\{(x_1, ..., x_{p-1}, y_1, ..., y_{q-1}, z) | x_1 + ... + x_{p-1} + (2 - 2z) \geq 1\}.$$

The closure of its complement in C is I^{q-1} cross a p-simplex defined by $x_1 + ... + x_{p-1} + 2(1 - z) \leq 1$, where $0 \leq x_i, z \leq 1$. Thus we can write C as the union of $f_1(C) \cong C$ and $C \setminus f_1(C) \cong \Delta^p \times I^{q-1}$. The boundary of $f_1(C) \cup (C \setminus f_1(C))$ is clearly the same as C, just with a more complex cell structure, i.e. a subdivision. If we now collapse to get the cone structure mentioned earlier, the simplex we just obtained is not affected, and we still have a subdivision. Patching together all faces in ∂I^q shows that $\Delta^{p-1} \times I^q \subseteq \Delta^{p-1} \times I^q \cup \Delta^p \times \partial I^q$, as desired; since the homotopy fixed all y-coordinates, the subdivisions on each face of ∂I^q match up. Finally, we glue together all the simplices to show that $A \subseteq B_1 \cup B_2$. Note that the center vertex of S was the corner of each simplex sent to the origin in our embedding above, so when we glue our simplices together, all of those vertices are identified, and we have a well-defined subdivision rule.

There is one problem that may have arisen. First of all we are assuming that A is formed of $(p-1)$-simplices crossed with q-cubes. We also have a structure for B; but in the next stage of subdivision (or of constructing the universal cover), the new A's are formed from the old B's. Do they have the right structure? Well, note that B_1 was given the same structure as A, and the subtiles of B_1 correspond to those elements of Z^n that stay in the same subgroup of rank p. B_2 represents those elements that land in a subgroup of order $p+1$, and these are given the structure of n-cubes split into p-simplices cross $(q-1)$ cubes, and the p-simplices are grouped about the correct vertex. So, the cell structure is consistent.

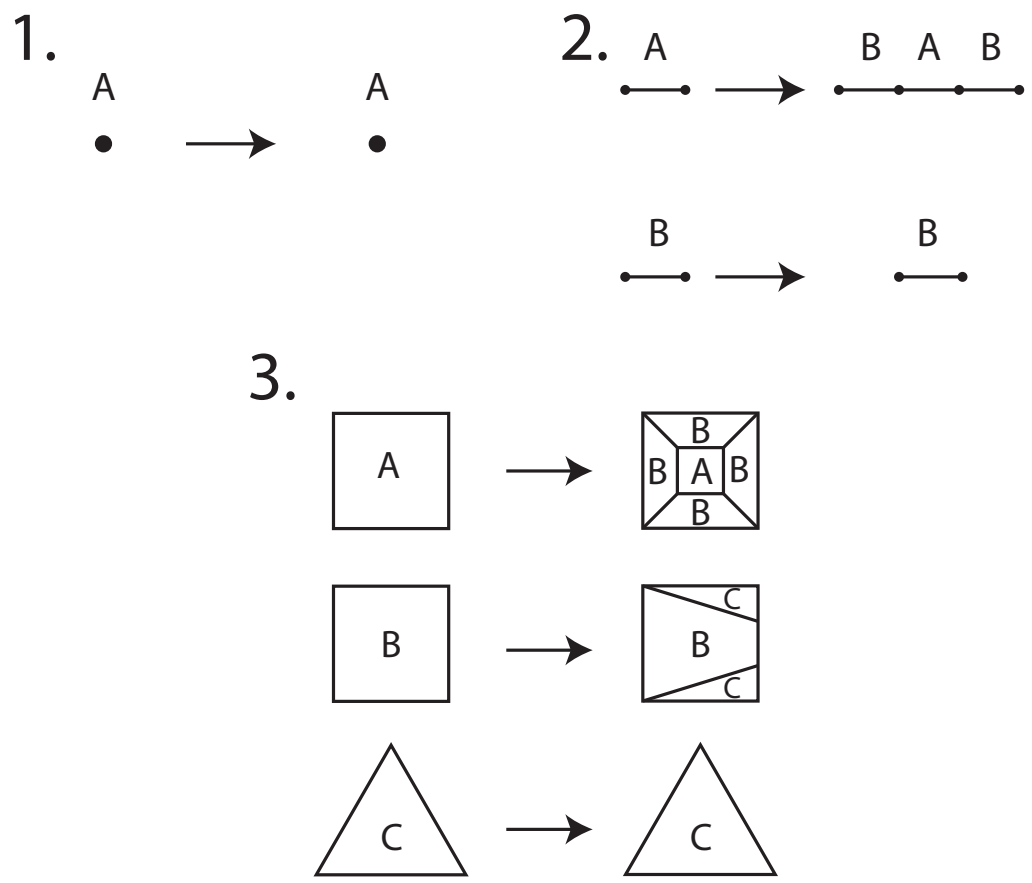

Figure 3.4: 1. The subdivision rule for the 1-torus, i.e. the circle. 2. The subdivision rule for the 2-torus, i.e. the standard torus. 3. The subdivision rule for the 3-torus.

□

We now find the subdivision rules explicitly for $n = 1, 2, 3$ and 4.

For $n = 1$, the universal cover is a line, $B(n)$ is $2n + 1$ line segments, and $S(n)$ is two points. The subdivision rule is shown in Figure 3.4.

Note that the only tile type is a point (i.e. 0-simplex cross a 0-cube), which is subdivided into one 0-simplex cross a 0-cube.

For $n = 2$, the fundamental domain is a square, $B(n + 1)$ is a topological disk, and $S(n)$ is a topological circle. The subdivision rule is shown in Figure 3.4.

Type A is a line (i.e. a 0-simplex cross a 1-cube). It's subdivided into one line(a 0-

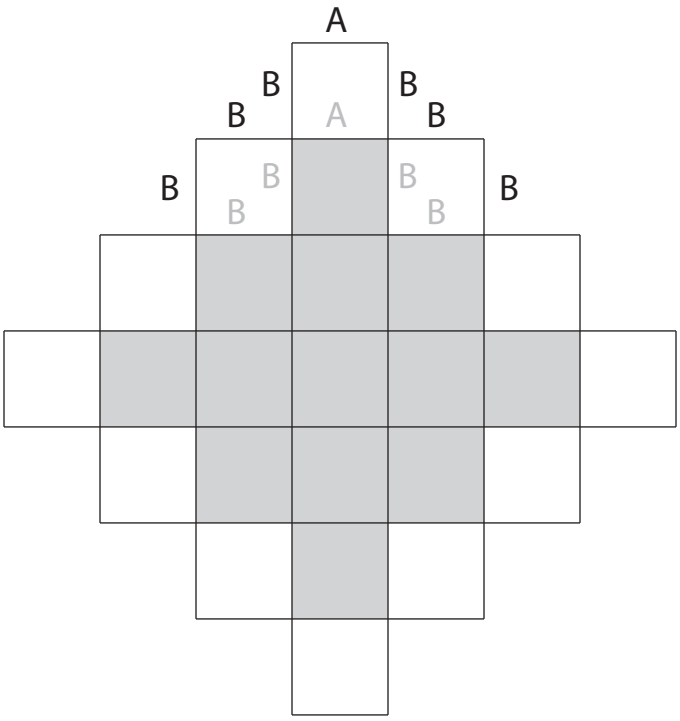

Figure 3.5: White tiles correspond to $B(3)$, and grey tiles to $B(2)$. Note that A tiles correspond to elements furthest from the origin.

simplex cross a 1-cube) and 2 more lines (1-simplices cross a 0-cube), just as the formula predicts.

Type B is a line(a 1-simplex cross a 0-cube), and represents half of a group element. Two B tiles form the star of a vertex in the boundary of the 2-cube (a square), and the subdivision rule for type B is the identity, just as the formula predicts.

As you can see in Figure 3.5, A tiles correspond to the four 'ends' of $B(n)$, or, the group elements contained in a subgroup generated by exactly one of the standard generators, while B tiles correspond to elements that must be written using both generators. Notice how two neighboring B tiles form a corner that is covered up by one square fundamental domain.

For $n = 3$, we have the 3-torus, whose universal cover is shown being constructed in Figures 2.3 to 2.5. Notice in these figures that new cubes are glued onto a single face, two neighboring faces, or three faces forming a corner. These correspond to elements whose

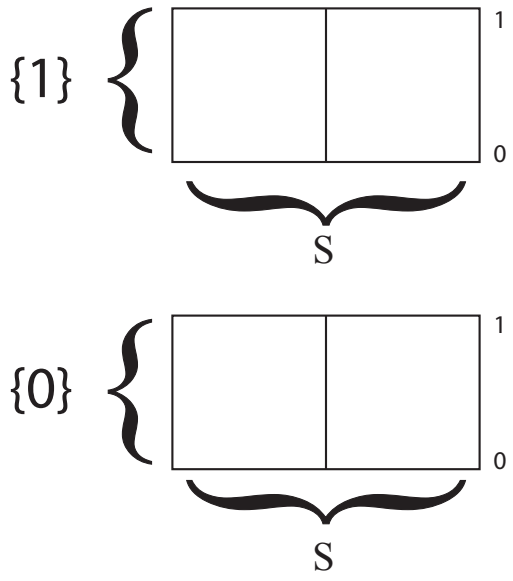

Figure 3.6: The set $S \times \partial I^1 \times I$, where S is the union of two 1-simplices.

minimal word-length representations use one, two, or three generators, respectively. Type A (corresponding to a single face) is a square (or 0-simplex cross a 2-cube), and is subdivided into one square (a 0-simplex cross a 2-cube) and 4 other squares (or 1-simplices cross 1-cubes).

Type B is a square (thought of as a 1-simplex cross a 1-cube). It is subdivided into a square (a 1-simplex cross a 1-cube), and 2 triangles (or 2-simplices cross a 0-cube). Two type B tiles correspond to the star of a vertex in the boundary of the 2-cube, which is then crossed by I.

Note that this tile shows us what happens with the homotopy portion of Theorem 4. We start with S, the star of a vertex in the boundary of $I^p = I^1$ with a simplicial structure (namely, two edges of a square sharing a vertex, each edge considered as a 1-simplex). We cross this with $I^q{=}I^1$ to get two square sharing a face. This is A. Then, we write A as a quotient of $S \times \partial I^1 \times I$, so we get Figure 3.6.

On each component (one corresponding to $S{\times}0{\times}I$, one corresponding to $S{\times}1{\times}I$), we pull down the center line by our homotopy to get Figure 3.7. Collapsing $\partial I^1 \times \{0\} = \{0,1\} \times \{0\}$

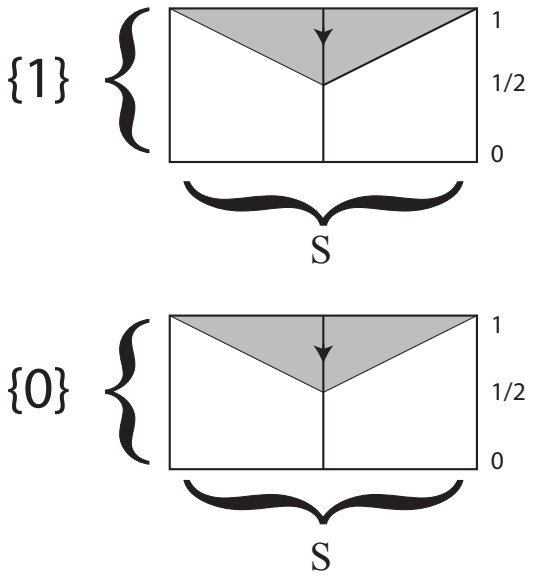

Figure 3.7: The same set as Figure 3.6, after the homotopy. The grey regions are the complement of the image of the homotopy; note that they are 2-simplices.

to a point to get a cone, we have the image in Figure 3.8.

Type C is a triangle (a 2-simplex cross a 0-cube), which is subdivided by the identity. This corresponds to a cube glued onto three adjoining faces. Each tile of type C represents one sixth of a group element, and six together form the star of a vertex in ∂I^3 when it is given a simplicial structure.

Several subdivisions of an A tile are shown in Figure 3.9. This picture was created with Ken Stephenson's Circlepack [21]. The pictures are only combinatorial subdivisions of each other; they can't be overlaid with vertices matching up. This is because the subdivision rule is not conformal. For more on the connection between conformality and circle packings, see [11].

Finally, the tile types for the four-torus are shown in Figure 3.10.

Note the homotopies in each of these tiles. In the first tile type, the homotopy is shrinking the cube in the center and dragging the cell structure with it. In the second, we only shrink down a 2-dimensional face, again dragging everything along with it. In the third tile type,

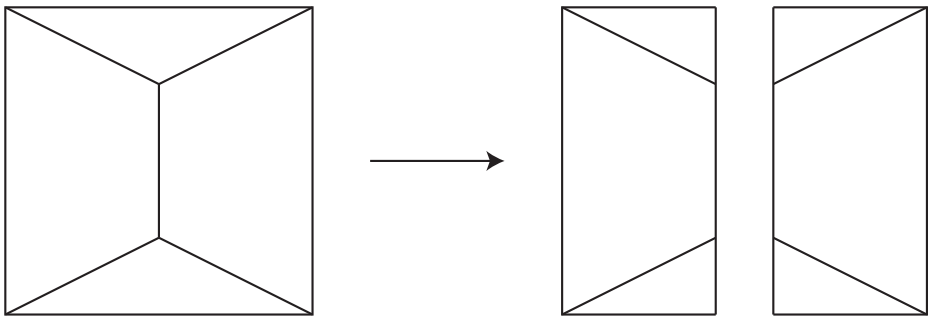

Figure 3.8: The same set as in Figure 3.7, after collapsing $\partial I^1 \times 0$. Note that the original cell structure (2 squares) is contained in the new structure.

we shrink an edge, dragging along the cell structure with it. Finally, in the fourth tile type, there is nothing to drag.

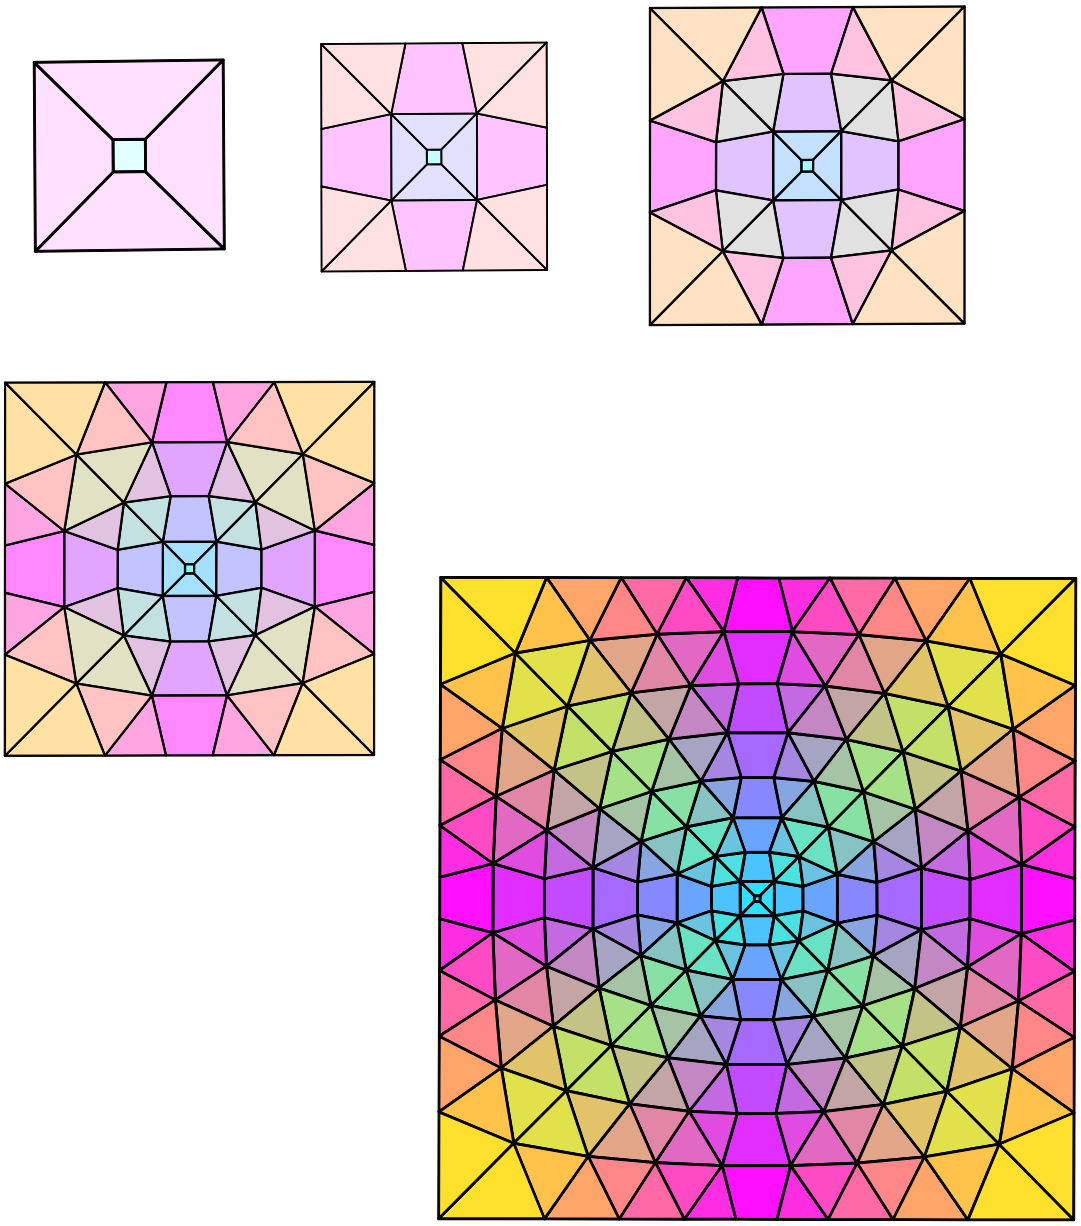

Figure 3.9: Several subdivisions of a type A tile for the 3-torus.

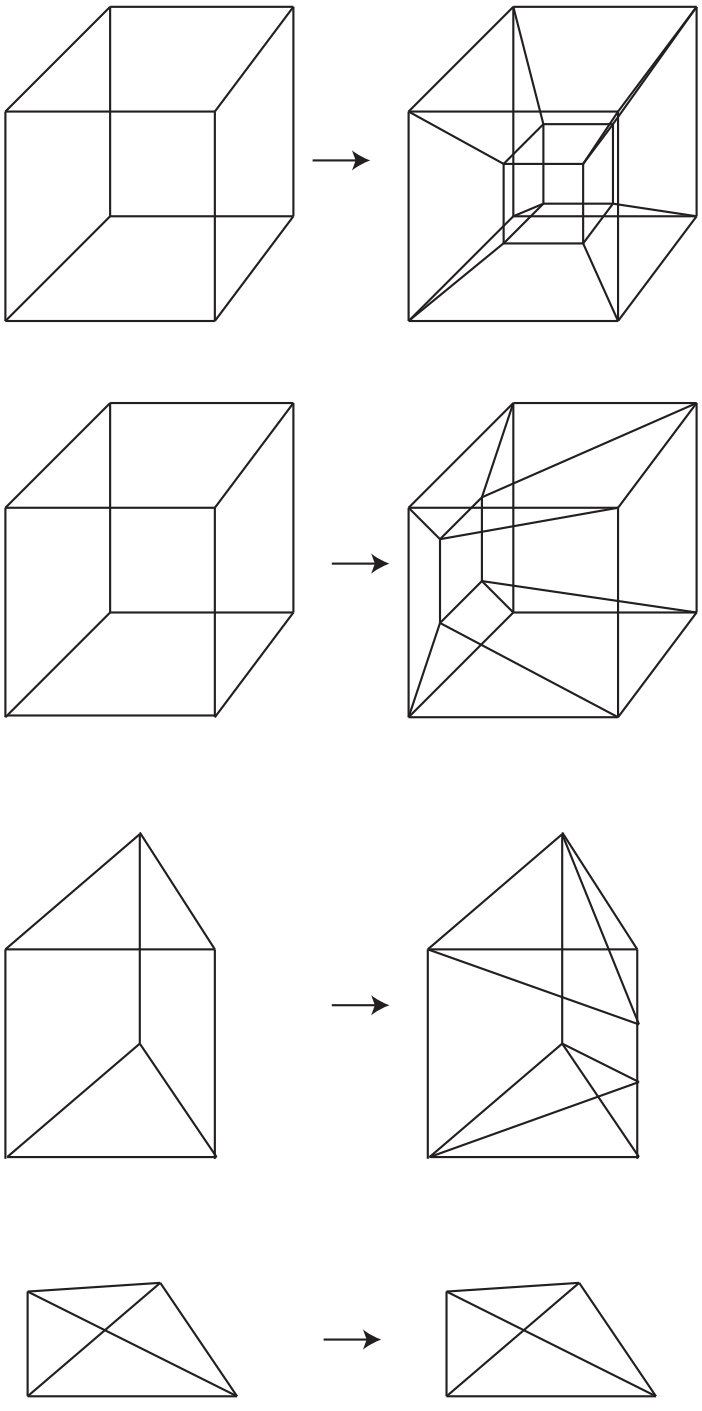

Figure 3.10: The tile types for the four-torus.

Chapter 4. The Space at Infinity of a Subdivision Rule

4.1 Definitions

A subdivision rule has **combinatorial mesh** going to 0 if i) every edge is subdivided at some point and ii) disjoint edges are separated by the subdivision, i.e. there is a stage of subdivision where no tile intersects both edges.

A subdivision rule for a hyperbolic manifold generates a basis for the sphere at infinity. In fact, any subdivision of the sphere with mesh going to 0 generates a basis for the sphere, where closed tiles form the "closed basis" (i.e. all closed sets are finite unions of arbitrary intersections of closed tiles). More generally, a subdivision rule with combinatorial mesh going to 0 can always be adjusted to have mesh going to 0 without changing the combinatorial structure.

One question is, can we generalize this to subdivision where mesh doesn't go to 0? This type of subdivision rule arises frequently when studying non-hyperbolic manifolds, as we saw in Chapter 2. In fact, I began studying non-hyperbolic manifolds to look for non-conformal subdivision rules. But they all failed to be conformal in the simplest way, by having tiles that didn't subdivide in one direction or the other. As we will show in this chapter (Corollary 4.5), only hyperbolic manifolds can have mesh going to 0.

Having so many subdivision rules with mesh not going to 0, we define a space at infinity for all subdivision rules. The topology will be generated by letting closed tiles of all stages generate a closed basis. The points will be nested sequences of cells of the same dimension: sequences of nested tiles, sequences of nested edges, and vertices. For subdivision rules of other dimensions, the points are again sequences of cells of the same dimension.

However, this definition of space at infinity does not extend the one for hyperbolic manifolds, as it is never Hausdorff or even T_1 except in trivial cases. What happens is that tiles of higher dimension can limit to tiles of lower dimension. In cases where mesh doesn't go to

0, we actually want non-Hausdorff spaces, but we want all of our spaces to be as Hausdorff as possible. We can take a quotient space by identifying points that 'converge' to each other. Two points will converge if the combinatorial mesh goes to 0 locally around them. More specifically, a sequence of nested k-cells $\{t_n\}$ converges to a sequence of nested cells $\{e_n\}$ of lower dimension if, for all sequences $\{f_n\}$ of cells of dimension $< k$ that are eventually disjoint from $\{e_n\}$, t_n contains e_n for every n but eventually does not contain f_n. Two disjoint sequences cannot be identified to the same point, as sequences of cells only converge to lower-dimension sequences that they contain.

Equivalently, we identify two points $\{f_n\}$ and $\{e_n\}$ if $\{f_n\}$ converges as a sequence of sets in the manifold to $\cap e_n$.

This definition makes more sense in lower dimensions. It means that a nested sequence of edges will converge to a vertex if each edge in the sequence is subdivided infinitely often and every edge intersects the given vertex. In dimension two, we can examine three ways of dividing a square. See Figure 4.1. In the first, the tile never subdivides, so $t_n = t_1$ for all n, and it does not converge to any sequence of edges, nor do the edges converge to any vertex; in the space at infinity, there are nine points. It is a finite model of a Hausdorff square, and has similar homotopy properties. It is contractible, and removing the center point (the only 'interior point') leaves a space homotopy equivalent to a circle, just like the Hausdorff square.

In the second example, the tile subdivides, but only in one direction. Thus, the sequence of tiles that contains the right edge converges to that edge, but that edge does not converge to either of its endpoints. In the limit, we get three Hausdorff lines glued together so that the top and bottom lines are embedded as closed subsets, but every point in the middle contains the corresponding points above and below in its closure. The resulting space, again, shares the homotopy type of the square.

In the third example, we get a Hausdorff square, as shown below.

Theorem 5. Let S be a subdivision rule on a manifold F (possibly with boundary). Let X

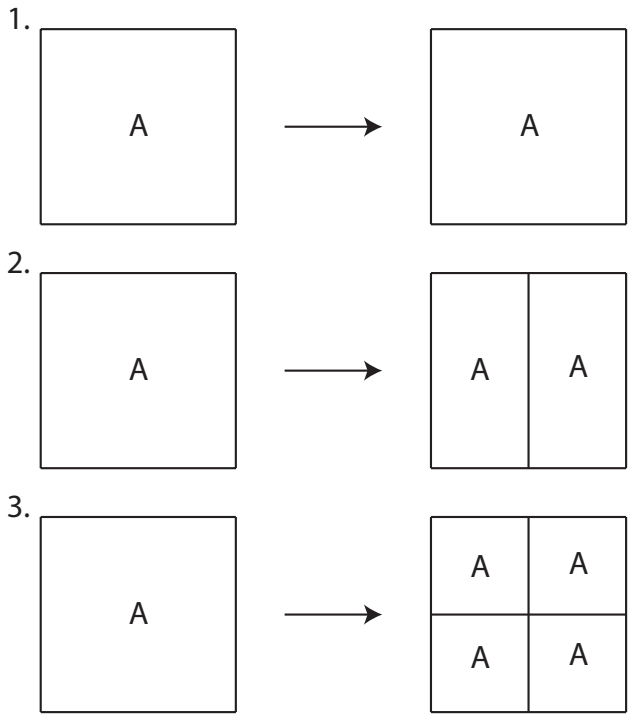

Figure 4.1: Three subdivision rules for the square.

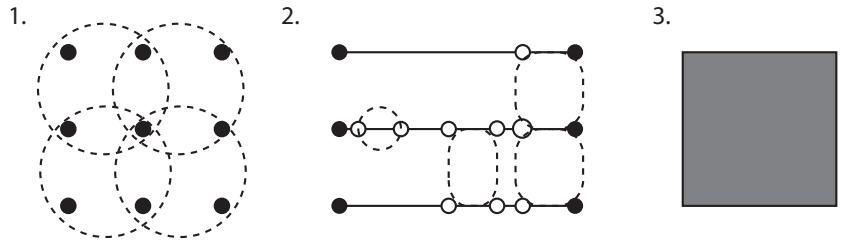

Figure 4.2: The space at infinity for each of the subdivision rules in Figure 4.1.

be the space at infinity. Then X is homeomorphic to F if S has (topological) mesh going to 0, and if X is homeomorphic to F, then X has combinatorial mesh going to 0.

Proof. We begin with the second statement. If X does not have combinatorial mesh going to 0, then there is a sequence of k-cells t_n and two disjoint sequences of m_1-cells e_n and m_2-cells f_n (with $m_i \leq k$) such that t_n contains e_n and f_n for every n. But then t_n is not identified with e_n or f_n, but contains both in its closure. In either case, X is not Hausdorff, and cannot be homeomorphic to F.

To show the first statement in the theorem, define a map $f : X \to F$ by sending each sequence of nested tiles $\{t_n\}$ to $\cap t_n \subseteq F$. Because mesh goes to 0, there is exactly one point in this intersection, making the map one-to-one, if it is well-defined. But, as stated above, if two sequences of cells are identified with each other, their intersection must have been non-empty at every stage of subdivision. Since mesh goes to 0, the two sequences get sent to the same point, and the map is well defined. It is continuous, because closed tiles in F are pulled back to closed basis elements in X.

The function f has an inverse defined by sending a point p in F to a sequence of tiles t_n (in any dimension) where every tile in the sequence contains p. This map is well-defined, because any two sequences containing the same point will be identified with each other (since combinatorial mesh goes to 0). It pulls closed basis elements back to closed tiles, and so is continuous. Thus, f is a homeomorphism. □

4.2 Low Dimensional Examples

In dimension 0, the only manifolds are collections of points, and these cannot subdivide. Thus, the space at infinity is a discrete set of points.

In dimension 1, we can start with a single closed tile. Let our subdivision rule be 4.4. Then the sequences of nested edges (before identification, and not including vertices) form a Cantor set, while the vertices form a countable set with one vertex between each pair of endpoints of the Cantor set. Identifying convergent sequences identifies the pairs of endpoints

Figure 4.3: The initial tile.

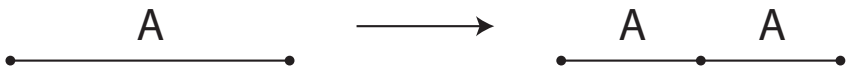

Figure 4.4: A binary subdivision rule.

with the vertex and with each other, and we end up with a set homeomorphic to a closed interval.

Now, consider the same tile with the rule shown in Figure 4.5. After several iterates, we have the picture shown in Figure 4.6. In this case, a sequence of edges with any B in it stops subdividing and does not converge to anything, and we have a countable topological space composed of "chains", with the topology in Figure 4.7. Note that points coming from vertices are closed, and points coming from B edges contain the two neighboring vertices in their closure. Note also that the single point in the middle is the only point whose neighborhoods all contain infinitely other many points. It corresponds to the unique sequence of nested A tiles.

These two examples are the archetypes for dimension 1 subdivision rules. In particular, consider a hyperbolic 2-manifold. As in Section 2.7, we can consider it as a covering of a 'pair of pants', which is made from dodecagons with vertices identified in groups of three.

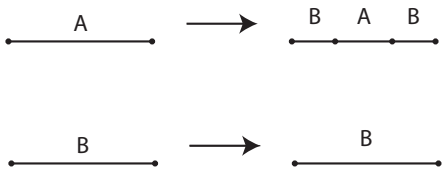

Figure 4.5: A subdivision rule coming from the torus.

Figure 4.6: A tile of type A after repeated subdivision.

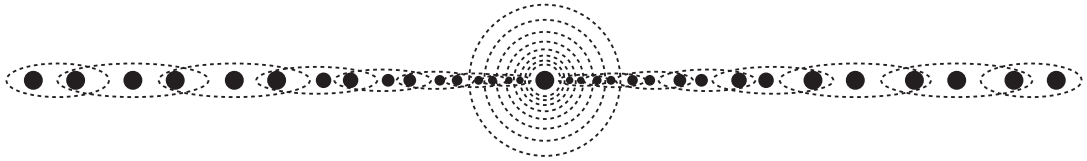

Figure 4.7: The space at infinity associated to the subdivision rule in Figure 4.5.

The associated subdivision rule is show in Figure 4.8. The space at infinity is a Hausdorff circle.

Now consider the torus. One subdivision rule for it is the one described earlier in Figure 4.5. The associated space is made up of four intervals as described above glued in cyclic fashion, and the resulting space is a countable non-Hausdorff space with almost all the properties of the circle. The group $\mathbb{Z} \times \mathbb{Z}$ acts faithfully on this space; however, there are other subdivision rules for the torus where the group does not act faithfully, or does not act at all.

An interesting variant on these surfaces is the punctured torus, which has a fundamental domain of the form shown in Figure 4.9, where the vertices are missing. The subdivision rule has a single tile, as shown in Figure 4.10. Because the vertices are deleted, we don't include them in the space at infinity. Thus, the space we get is a Cantor set, the same as the Gromov space at infinity.

4.3 Examples from 3-manifolds

The situation is more complicated for 3-manifolds. Theorem 5 shows that subdivision rules for hyperbolic manifolds (which have combinatorial mesh going to 0) have a 2-sphere at infinity.

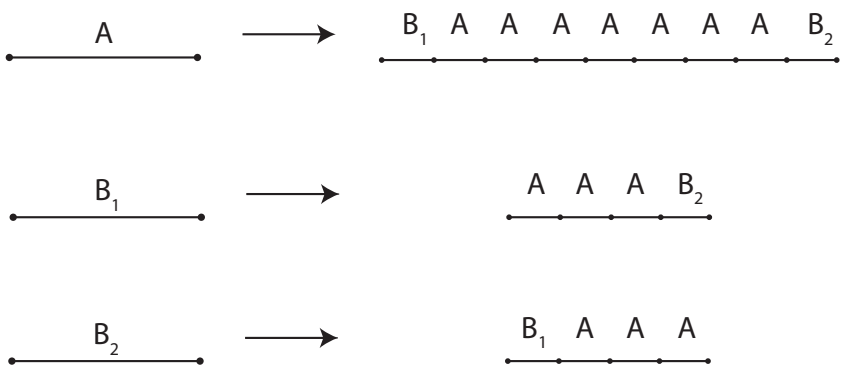

Figure 4.8: A subdivision rule coming from a hyperbolic surface.

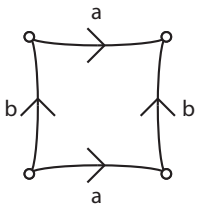

Figure 4.9: A fundamental domain for the punctured torus.

Figure 4.10: A subdivision rule coming from the punctured torus.

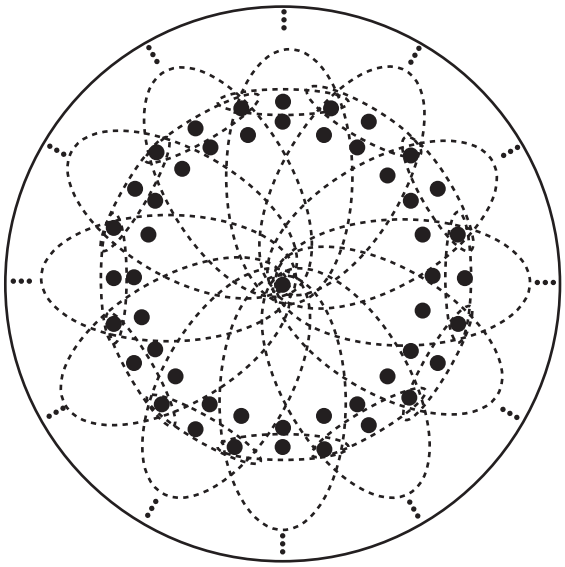

Figure 4.11: The upper hemisphere for the space at infinity of an $\mathbb{H}^2 \times R$ manifold.

For $S^2 \times S^1$, each hollow sphere can be considered as a single 'tile', and the space at infinity is two points, which is as it should be, since the group is $\mathbb{Z}$, which is a hyperbolic group with two points at infinity.

For an $\mathbb{H}^2 \times S^1$ manifold or $S\tilde{L}_2\mathbb{R}$ manifold, the subdivision rule we obtained earlier gives a space that is a non-Hausdorff suspension of a circle. Specifically, there is an embedded Hausdorff circle at the equator, plus 2 countable sets, each homotopy equivalent to a 2-disk. The points of each countable disk limit onto the equator, so that all neighborhoods of a point on the equator contain infinitely many points of the disk. See Figure 4.11.

For the Euclidean manifold $S^1 \times S^1 \times S^1$, we get a belt around the equator that is the same as the space at infinity for $S^1 \times S^1$. The space for $S^1 \times S^1 \times S^1$ is also a non-Hausdorff suspension of the equator, which is what we generally get when crossing a group with $\mathbb{Z}$. See Figure 4.12.

We can also define a space at infinity for non-closed 3-manifolds. In [18], I found an explicit subdivision rule for every alternating knot complement. As for the punctured torus, we don't consider the vertices as part of the space at infinity. For hyperbolic knots (like the

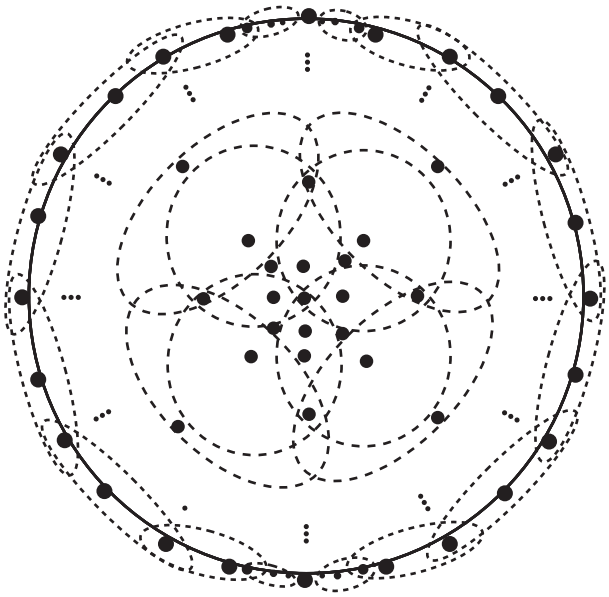

Figure 4.12: The upper hemisphere of the space at infinity for the three-torus.

Borromean rings),we get a Sierpinski carpet. Because each vertex is missing, each tangent vector at the vertex is a different, unidentified point, and we get a deleted open disk at each vertex. For the trefoil knot, we get a non-Hausdorff suspension of the Cantor set. For the Hopf link, we get the same space as the torus; if we look closely, we see that the complement of the Hopf link is an interval cross the torus. Its fundamental group is $\mathbb{Z} \times \mathbb{Z}$, a two-dimensional group, and the space at infinity somehow detects that.

4.4 Higher Dimensional Spaces

The subdivision rules for the n-cubes that we found in Chapter 3 give us an example of a subdivision rule in each dimension with mesh not going to 0 (except trivially when $n = 1$). The space can be defined inductively by first finding the space for each subgroup generated by $(n-1)$ of the standard generators, and then gluing them together along the subgroups generated by $(n-2)$ generators. Everything that is not an element of one of these subgroups has the tile type of an $(n-1)$ simplex which never subdivides. Each such simplex will be, at infinity, a collection of points, one for each subsimplex, with each point containing in its

closure all subsimplices of the corresponding simplex. These leftover group elements form a countable set that span the holes left by the $(n-1)$ generator subgroups, and together form a space homotopy equivalent to an $(n-1)$-sphere. Every neighborhood of a point in the $(n-1)$-generator subgroups contains infinitely many elements corresponding to these $(n-1)$-simplices. This corresponds to the fact that every tile in the subdivision rule given by Theorem 4 'sheds' simplices of dimension one higher than themselves, until we get $(n-1)$-simplices, which do not subdivide.

For $n = 1$, we get two points. Note that this is a Gromov hyperbolic group (namely, $\mathbb{Z}$), and that its space at infinity in that setting is also two points.

For $n = 2$, our manifold is the torus, and our space is exactly the space described earlier, formed of four copies of Figure 4.7. Note that the four 'central' points or A tiles correspond to the subgroups generated by a single generator.

For $n = 3$, our manifold is the 3-torus, although obtained in a different way than the one in Chapter 2. The figure, then, will be very similar to the picture in Figure 4.12, except the equator will have four central points, instead of two, and there will be two more embedded spaces homeomorphic to the torus space, one perpendicular to each coordinate axis. The resulting picture looks like a polka-dot gyroscope.

The picture is similar in higher dimensions.

Note that we can glue this space at infinity onto the Cayley graph or onto E^n to give a compactification or 'quasicompactification', where quasicompact means compact without the Hausdorff condition. In this compactification, every isometry of the Cayley graph gives a homeomorphism of the boundary, because the isometries are composed of reflections in a coordinate hyperplane (which easily gives a homeomorphism of the space at infinity) or translation by a generator, which also acts on the space at infinity by fixing the the points corresponding to that generator, fixing the equator perpendicular to the axis between those two points, and shifting everything else one spot closer to the generator and one spot away from its inverse.

For hyperbolic groups, the Combinatorial Riemann Mapping Theorem shows that a transformation that takes the combinatorial structure to itself is a conformal map. In our situation, the mapping theorem does not apply, but all of our isometries of Z^n extend to maps on the space at infinity that preserve the combinatorial structure. Thus, we can say that all isometries of Z^n lift to a conformal map on the boundary at infinity.

Note, however, that their are 'conformal' maps on the boundary that do not correspond to isometries of Z^n. In particular, we can take any quadrant of the space at infinity of Z^2 and shift it independently of the other three, and this does not correspond to an isometry of Z^2.

4.5 Hausdorff subsets are hyperbolic

Given any subdivision rule, we define a graph G in the following way:

Definition. The **n-th layer** G_n of the graph G consists of one vertex for each tile of $S(n)$, with one edge connecting vertices for each edge in $S(n)$ separating the corresponding faces. The edges of G_n are called **horizontal edges**.

Definition. The graph G also has **vertical edges** connecting each vertex in $G(n)$ representing a tile in $S(n)$ to each vertex in G_{n+1} representing a tile in $S(n+1)$.

Definition. If a is in G_n and $m < n$, let b in G_m correspond to the unique tile of $S(m)$ containing the tile of $S(n)$ corresponding to b. Then we define functions **transition functions** $f_{m,n} : G_n \to G_m$ by $f_{m,n}(a) = b$.

Lemma 4.1. *If a subdivision rule is bounded valence with combinatorial mesh going to 0, then there are integers $k, M > 0$ such that any arc α of length $l(\alpha) > M$ in G_{n+k} has image arc $f_{m,m+k}(\alpha)$ in G_m with length $< l(\alpha) - 2$.*

Proof. Because combinatorial mesh is going to 0, there is a $k > 0$ such that any path in $S(n)$ joining disjoint edges crosses at least two new edges in $S(n+k)$.

Now, any arc of sufficient size in $G(n)$ corresponds to a path in $S(n)$ that joins disjoint edges. To see this, note that if every edge that the path crosses intersects the first edge, then they must share a vertex with the first edge, and there are no more than $2V$ such edges, where V is the maximum valence of the subdivision rule. Thus, if $f_{n,n+k}(\alpha)$ has length $> 2V$, then its corresponding path in $S(n)$ joins disjoint edges, and so the path in $S(n+k)$ corresponding to α must cross at least two more edges.

Thus, if $f_{n,n+k}(\alpha)$ has length $> 2V$, then $f_{n,n+k}(\alpha)$ has length $< l(\alpha) - 2$. Also, if $f_{n,n+k}(\alpha)$ has length $\leq 2V$ and α has length $> 2V + 2$ then $f_{n,n+k}(\alpha)$ is also strictly shorter then $l(\alpha) - 2$. Combining these two statements, we see that if α has length greater than $M = 2V + 2$, then $f_{n,n+k}(\alpha)$ has length strictly less than $l(\alpha) - 2$. □

Now, we define a standard path.

Definition. A **standard path** in G from a point a to a point b is a geodesic that consists of a vertical, upward path beginning at a, a purely horizontal path of size $< M$, followed by a downward vertical path ending at b.

Lemma 4.2. *For a finite, bounded-valence subdivision rule with combinatorial meshing going to 0, every geodesic α is within ϵ of a standard path, where ϵ is a global constant.*

Proof. Let x in G_m be the initial point of α, and let y in G_n be the terminal point. Then let $A = \{$all vertical, upward edges$\}$, $B = \{$all horizontal edges$\}$, $C = \{$all vertical, downward edges$\}$. Note that $m + |A| = n + |C|$; call this number k. Now let β be the path consisting of $|A|$ upward edges, a horizontal path of length $\leq |B|$ from $f_{k,m}(a)$ to $f_{k,n}(b)$, and a downward path of $|C|$ downward edges.

The horizontal path in the middle can be taken to be the image of all elements of B under the appropriate transition functions. If this horizontal path has length $> M$, then β can be shortened by adding one more vertical edge to the vertical path, following the image of the horizontal path under $f_{k-1,k}$, going down a vertical edge, and then following the original downward path. Because α is a geodesic, β cannot be shortened, so the horizontal path has length $\leq M$ and β is a standard path.

Now we show that α and β stay within some fixed constant ϵ of each other, depending only on the subdivision rule. We begin by showing that all upward edges occur in α before all downward edges. Assume the geodesic goes down one edge, follows a horizontal path in G_p for some p, then goes up a vertical edge. The image of the horizontal segment is no shorter in G_{p-1}, and so removing the vertical segments while keeping the horizontal path (mapped into G_{p-1}) gives a strictly shorter path.

Thus, we have shown that α consists of a path where all vertical segments occur before all downward segments, and all horizontal segments total no more than M in length.

Now, compare α and β. They both start at a in G_m, go upward to some points in G_k (with α possibly taking horizontal detours of total length $\leq M$), take a horizontal path of length $\leq M$ to some other point in G_k, the go downward to b in G_n (again with possible detours for α). Since distances in G_p are no longer than distances in G_q for $p < q$, the path β stays within M of α on the upward segment. Then the endpoints of the last upward edges of α and β lie in G_k and have distance $< M$. The horizontal paths in G_k have length $< M$, and so the paths lie within $2M$ of each other at all times. By symmetry, the downward segments of α and β remain within M of each other. Thus, if $\epsilon = 2M$, all geodesics are within ϵ of a geodesic with the same starting points that follows a standard path. □

Using this last lemma, we can now show that G is Gromov hyperbolic by considering triangles of geodesics that follow standard paths. Standard paths are useful, because they allow us to focus on the upward and downward segments. Vertical geodesics connecting equivalent points in different layers are unique; no horizontal moves are necessary, and each vertex in a layer has a unique upward path coming from it. This uniqueness of vertical geodesics gives our graph a tree-like structure.

Theorem 6. If a subdivision rule has bounded valence and combinatorial mesh going to 0, the associated graph G is Gromov hyperbolic.

Proof. We will show that geodesic triangles are delta-thin. By Lemma 4.2, we can assume

the geodesics in a triangle follow standard paths by changing distances a bounded amount. Now, let x, y, and z be points lying in G_m, G_n, and G_k, respectively, with geodesics $\alpha xy, \alpha xz$, and αyz connecting them. Because these are standard paths, they are vertical except for a horizontal portion lying entirely in some layer of the graph. Let $h(x, y), h(x, z)$ and $h(y, z)$ be the height of the horizontal portions of the corresponding geodesics (so, for instance, the horizontal part of αxy lies in $G_{h(x,y)}$.

Without loss of generality, assume $H(x, y) \leq H(x, z) \leq H(y, z)$. We first show that α_{xy} remains close to the other two geodesics. It is the same path as α_{xz} until they reach $G_{H(x,y)}$; it then follows a path of length $\leq M$ (where M is, again, the constant from Lemma 4.1), then follows a downward segment to z that is the same as that followed by α_{yz}. Thus, α_{xy} is within M of the other paths at all times.

Now, the horizontal path of α_{xy} actually connected the images of x and y under the appropriate transition functions. So, in $G_{H(x,y)}$, the vertical segments of α_{xz} and α_{yz} are no more than M apart, and they are no further apart in higher levels. Thus, they are within M of each other until $H(x, z)$.

Now, in $H(x, z)$, the images of x and z are $\leq M$ apart, and so are those of y and x (because $H(x, y) \leq H(x, z)$). Thus, y and z are no more than $2M$ apart. Thus, the path of α_{xz} (which goes up, over and down from the image of x in this stage to the image of z in this stage) must be less than $2M$ in length. Thus, it is never more than $3M$ away from α_{xz}. Finally, both end with the same downward segment from the image of z in $G_H(y, z)$ down to z itself. Thus, all geodesics in the triangle are no more than $3M$ apart at any point. Our assumption that our geodesics were standard paths shifted each geodesic by no more than ϵ (the constant from Lemma 4.2. Thus, every edge in a geodesic triangle is within $\delta = 3M + 2\epsilon$ of the union of the other edges, and our graph G is Gromov hyperbolic. □

Corollary 1. No bounded valence subdivision rule associated to a non-hyperbolic 3-manifold has combinatorial mesh going to 0.

Part II

Circle packings and subdivision rules

Chapter 5. A Circle Packing Proof of the Combinatorial Riemann Mapping Theorem

5.1 The Theorem

The proof in this section is essentially the same as that of [11], although it was developed independently.

We begin with the statement of our version of the Combinatorial Riemann Mapping theorem, first proved in [2]. We first recall the definitions involved.

A **tiling** is a decomposition of a topological sphere into polygons that have disjoint interiors and that meet in an edge or a vertex. They can be triangulated by putting a vertex in each face and adding an edge from the new vertex to each of the old vertices.

A tiling T of a ring R (i.e., a closed annulus) gives two invariants, $M_{sup}(R,T)$ and $m_{inf}(R,T)$, called **approximate moduli**. These are similar to the classical modulus of a ring. They are defined by the use of **weight functions**. A weight function ρ assigns a non-negative number called a **weight** to each vertex of T. Every edge-path in R can be given a length, defined to be the sum of the weights of all vertices in the path. We define the **height** $H(\rho)$ of R under ρ to be the infimum of the length of all possible edge-paths connecting the inner boundary of R to the outer boundary. The **circumference** $C(\rho)$ of R under ρ is the infimum of the length of all possible edge-paths circling the ring (i.e. not nullhomotopic in R). The **area** $A(\rho)$ of R under ρ is defined to be the sum of the squares of all weights in R. Then we define $M_{sup}(R,T) = \sup_{\rho} \frac{H(\rho)^2}{A(\rho)}$ and $m_{inf}(R,T) = \inf_{\rho} \frac{A(\rho)}{C(\rho)^2}$. Note that they are invariant under scaling of the metric.

A sequence $T_1, T_2, ...$ of tilings is **conformal (K)** if:

(i) mesh approaches 0 (note that this is independent of topological metric);

(ii) for each ring R, the approximate moduli $M_{sup}(R,T_i)$ and $m_{inf}(R,T_i)$, for all i suffi-

ciently large, lie in a single interval of the form $[r, Kr]$; and

(iii) given a point x in the surface, a neighborhood N of x, and an integer I, there is a ring R in $N \setminus \{x\}$ separating x from the complement of N, such that for all large i the approximate moduli of R are all greater than I.

Triangulations are examples of tilings. The most familiar example of a sequence of tilings is iterated barycentric subdivision of a triangulation. Barycentric subdivision is not conformal [6] (see also Chapter 6.2), but many common sequences of tilings are. A simple example of a conformal sequence is connecting the midpoints of all triangles in a triangulation with new edges; this has been called hexagonal refinement [22]. In general, all sequences of tilings with tiles that are uniformly 'almost round' in some metric are conformal [2].

The original Combinatorial Riemann Mapping Theorem was proved in a more general setting, where tiles are allowed to overlap. Also, there were no restrictions on the tilings besides conformality. However, our use of circle packings limits us to sequences of tilings of **bounded size**, which means that the number of edges of each face and the valence of each edge are uniformly bounded above for the entire sequence. This is equivalent to the triangulations of the tilings having uniformly bounded valence. The motivating problem for the Combinatorial Riemann Mapping Theorem is showing that natural disk covers of the sphere at infinity for certain Gromov-hyperbolic groups give the sphere an analytic structure that is comparable to the combinatorial one. In this setting, the disk covers can be replaced by tilings of bounded size, and we lose nothing by restricting to this case.

Theorem 7 (The Combinatorial Riemann Mapping Theorem for bounded-valence triangulations)**.** Let $T_1, T_2, ...$ be a conformal sequence (K) of triangulations of the topological 2-sphere S, with the valence of all vertices bounded above by V. There exists a homeomorphism identifying S with the Riemann sphere $\mathbb{P}$ and a constant $K' = K'(K, V)$ such that classical and asymptotic moduli are K'-comparable. That is, if R is any ring in S, then there is an integer I and an interval of the form $[r, K'r]$ such that the classical analytic modulus of R and the approximate moduli $M_{sup}(R, T_i)$ all lie in $[r, K'r]$ for each $i \geq I$.

Corollary 2 (The Combinatorial Riemann Mapping Theorem for bounded-size tilings)**.** Let $T_1, T_2, ...$ be a bounded-size conformal sequence (K) of tilings of the topological 2-sphere S. There exists a homeomorphism identifying S with the Riemann sphere $\mathbb{P}$ and a constant K' (depending only on K and the size of the tiling) such that classical and asymptotic moduli are K'-comparable. That is, if R is any ring formed of closed tiles in S, then there is an integer I and an interval of the form $[r, K'r]$ such that the classical analytic modulus of T and the approximate moduli $M_{sup}(R, T_i)$ all lie in $[r, K'r]$ for each $i \geq I$.

Proof. This follows from Theorem 4.3.1 of [7], which states that two tilings with bounded overlap have comparable moduli. Two tilings T, T' have bounded overlap if there is a constant L such that each shingle in T intersects at most L shingles in T' and vice versa. A tiling of bounded size has bounded overlap with its triangulation, so moduli in the tiling are comparable to those in its triangulation, which are comparable to the analytic moduli. □

Circle packings are analogous to quasiconformal maps. Instead of preserving the classical modulus of rings up to a fixed constant, circle packings push combinatorial rings with no analytic structure into the Riemann sphere, giving them a classical modulus uniformly comparable to the asymptotic moduli. Many of the classic quasiconformal theorems carry over directly to this setting, as we will show.

What do we mean by a circle packing?. Given a triangulation T of the sphere, there is a unique circle packing (up to Möbius transformation) of circles on the sphere whose adjacency graph is combinatorially equivalent to T. By fixing the image of three points, we get a unique circle packing. We then define a map from S to $\mathbb{P}$ by sending 0-,1-, and 2- cells to 0-, 1- and 2- cells by arbitrary orientation-preserving homeomorphisms and call the result a **circle packing of T**.

5.2 An Important Lemma

A **combinatorial annulus** is a closed, triangulated annulus. We say that a combinatorial annulus is **induced** by a triangulation T of S if every triangle in the annulus is a triangle of T.

We use the following result (found in [22]) to show that circle packings respect moduli.

Lemma 5.1. *Suppose R is a combinatorial annulus, and Q is a univalent spherical circle packing for K, and that $\Omega = carr(Q) \subset \mathbb{P}$. Then $Mod(\Omega)$ (classical modulus), $M_{sup}(R)$, and $m_{inf}(R)$ are comparable by a constant c which depends only on maximum valence of R.*

Now, we begin to draw analogies between quasiconformal maps and these maps that respect moduli.

5.3 Analogues of Classical Theorems

Our first theorem also applies to all conformal tilings, whether bounded valence or not.

Theorem 8. Given a domain G with a conformal sequence $T_1, T_2, ...$ of triangulations with circle packings $w_1, w_2, ...$, if every w_n omits 2 values p_1, p_2 whose spherical distance is greater than some fixed $t > 0$, then $\{w_n\}$ is equicontinuous in G. Compare with Theorem II.4.1 of [13]

Proof. Given z_0 in G and ϵ, with $0 < \epsilon < t$, choose a ring R that separates z_0 from p_1, p_2, with $m_{inf}(R, T_n) \geq \frac{32}{\epsilon^2}$ for all $n \geq N$ for some N. Choose another annulus R' so that $\overline{R}$ is contained in the interior of R' and R' misses z_0, p_1 and p_2. For all $n \geq N'$ for some $N' \geq N$, there is a combinatorial annulus R_n induced by T_n with $R \subset R_n \subset R'$. Fix a tiling T_n and its packing w_n, with $n \geq N'$. Give the tiling $w_n(R_n)$ the weighting ρ associated to the circle packing, so that each vertex is assigned the radius of its associated circle. Then by definition of the discrete modulus, $\frac{C^2(\rho)}{A(\rho)} \leq \frac{1}{m_{inf}(R_n, T_n)} \leq \frac{1}{m_{inf}(R, T_n)}$, so $C(\rho) \leq \sqrt{\frac{A(\rho)}{m_{inf}(R)}}$. But $A(\rho) = \sum_{C \in w_n(R_n)} (radius(C))^2 \leq \sum_{C \in w_n(R_n)} \frac{2}{\pi} area(C) \leq 8$, because the area of the sphere

is 4π. Thus, $C(\rho) \leq \sqrt{\frac{8}{m_{inf}(R,T_n)}} < \frac{\epsilon}{2}$. Pick a path in $w_n(R_n)$ with length $2C(\rho)$. Then this path lies in a disk of radius $\leq \epsilon$, and one of the components of $\mathbb{P} \setminus w_n(R_n)$ lies in this disk. Because $\epsilon < t$, it must be the component containing $w_n(z_0)$. Thus, the component of $S \setminus R_n$ containing z_0 has image radius less than ϵ under every packing w_n for $n \geq N'$. Thus, the component of $S \setminus R'$ containing z_0 will also have image radius less than ϵ under every packing w_n for $n \geq N'$. This gives equicontinuity of the maps $w_1, w_2, ...$, and we are done. □

If we normalize our packings so that three points z_1, z_2, and z_3 have fixed, distinct images in the sphere, then we have equicontinuity of the sequence by applying the above to S minus each pair of points in turn. Because the images lie in $\mathbb{P}$, a compact set, this is a normal family and thus converges uniformly to a continuous function w on S, whose image contains 3 distinct points. We now use the following theorem to show the map is univalent, and thus a homeomorphism.

Theorem 9. The limit function w of a sequence w_n of circle packings (of a conformal sequence (K) $T_1, T_2, ...$ of bounded-valence triangulations) which converges uniformly in S is either univalent or a constant map. Compare with Theorem II.5.3 in [13].

Proof. Assume there are a, b in G such that $w(a) = w(b)$. Eventually, we will show that w is a constant map. But first, we will show that every neighborhood of a contains a point $z \neq a$ such that $w(z) = w(a)$.

Choose a closed disk neighborhood D of a omitting b and lying in G. Then $w_n(\partial D)$ separates $w_n(a)$ and $w_n(b)$ for every n. The infimum of $d(w_n(z), w_n(a))$ on ∂D is $0 = d(w(a), w(b))$, so our first assertion is proved.

Now we will show that every point z_0 in G possesses a neighborhood $U(z_0) \subset (G)$ in which w is either one-to-one or constant. In particular, w will be constant on $U(a)$ (i.e. $w(z) = w(a)$), since w is not one-to-one on any neighborhood of a by our first step.

Choose a disk neighborhood $U(z_0) \subset G$ of z_0, having image diameter less than $\frac{\pi}{4}$ for all w_n. This is to keep the outer boundary of an annulus we will form from being to small.

If $U(z_0)$ does not have our desired property, then there are z_1, z_2, and z_3 in $U(z_0)$ with $w(z_1) \neq w(z_2) = w(z_3)$ (because w is neither constant nor one-to-one).

Choose a closed annulus $R \subset U(z_0)$ that separates z_1, z_2 from z_3, and another closed annulus R' contained in the interior of R. Then for all $n \geq N$ for some N, there is a combinatorial annulus R_n induced by T_n with $R' \subset R_n \subset R$. The image $w_n(R_n)$ lies in the disk $d(w_n(z), w_n(z_0)) < \frac{\pi}{4}$. It separates $w_n(z_1), w_n(z_2)$ from $w_n(z_3)$ and the circle $d(w, w_n(z_0)) = \frac{\pi}{4}$. Also, by Axiom 1, all moduli $M_{sup}(R', T_n)$ lie in a single interval $[r, Kr], r > 0$.

Now we have $\lim_{n\to\infty} d(w_n(z_2, z_3)) = 0$, $\lim_{n\to\infty} d(w_n(z_1, z_2)) > 0$, and both boundary components of $w_n(R)$ have diameters greater than $d(w_n(z_1), w_n(z_2))$ for all n. By Lemma I.6.2 of [13] we have $\lim_{n\to\infty} Mod(w_n(R)) = 0$ (classical modulus in $\mathbb{P}$). But, $Mod(w_n(R)) \geq Mod(w_n(R_n)) \geq \frac{1}{c} M_{sup}(w_n(R_n), T_n) \geq \frac{M_{sup}(R', T_n)}{c} \geq \frac{r}{c}$, where c is the constant from Lemma 5.1. Thus, the moduli $Mod(w_n(R))$ are bounded below, and we have a contradiction.

Thus, S is the disjoint union of two open sets: U, where w is locally constant, and V where w is locally univalent. Because a is in U, U is all of S, and w is a constant map.

□

Because we normalized our packings, the image of w contains at least three points, so we have a univalent, continuous map w of S into $\mathbb{P}$ which is thus a homeomorphism. All that remains is to show that the approximate and classical moduli are comparable. We know that it is the uniform limit of maps w_n satisfying this property, so all that remains is our last lemma.

A sequence of rings R_n is said to **converge to R from the inside** if $\overline{R_n} \subset R$ for every n and if for every $\epsilon > 0$ there is an n_ϵ such that for $n \geq n_\epsilon$, every point of the inner (resp. outer) boundary of R_n has spherical distance $< \epsilon$ from the corresponding boundary of R.

Lemma 5.2. *Let w_n be a sequence of circle packings of a conformal sequence (K) of bounded-valence triangulations $T_1, T_2, ...$ of a topological sphere S. Let R be ring in S with asymptotic moduli lying in $[r, Kr]$. If the packings converge uniformly to a map w on S, then the asymp-*

totic moduli $M_{sup}(R, T_n)$, $M_{inf}(R, T_n))$ and $Mod(w(R))$ all lie in a K' interval $[r', K'r']$, where K' depends only on K and the constant c from Lemma 5.1. Compare this to Theorem I.5.2 of [13].

Proof. First we show that $Mod(w(R)) \leq cKr$ for any annulus $R \subset S$. Let R_n be a nested sequence of combinatorial annuli whose union is the interior of R, where each R_n is induced by T_n. For each R_n, there is a k such that $w_m(R_n)$ is contained in $w(R)$ for $m \geq k$. We can get $w_n(R_n) \subset w(R)$ by replacing w_n with a subsequence. The w_n converge uniformly in R to w, which is uniformly continuous, so the $w_n(R_n)$ converge to $w(R)$ from the inside. Then by lemma I.6.4 of [13], $\lim_{n\to\infty} Mod(w_n(R_n)) = Mod(w(R))$. But for sufficiently large n, $Mod(w_n(R_n)) \leq Mod(w_n(R)) \leq cKr$, so in the limit, $Mod(w(R)) \leq cKr$

By surrounding R with nested combinatorial annuli, we can similarly obtain $Mod(w(R)) \geq \frac{r}{c}$. Thus, with $r' = \frac{r}{c}, K' = Kc^2$, all of our moduli of both types are comparable.

□

Thus, our sequence of normalized circle packings converges to our desired homeomorphism, and Theorem 7 is proved.

5.4 Further Work

One obvious generalization would be to extend the theorem to a sequence of tilings where the valence is unbounded. Our key lemma is missing, but there may be ways to work around this, such as blowing up the vertices (replacing each vertex with a tile).

Also, normalized circle packings may converge for a sequence that is not conformal. The easiest way to construct such sequences of tilings is by artificially forcing the mesh to not be 0; but other techniques allow us to create convergent circle packings that don't satisfy axiom 1. For instance, we can take the binary square subdivision rule and add rings around every vertex that follow the 2-3 subdivision rule. These rings are degenerate, and collapse in the limit, and the binary subdivision rule dominates.

The converse of our main theorem is clearly true; if the circle packings converge to a homeomorphism, the classical moduli are comparable with the approximate moduli, and the classical modulus can be made arbitrarily large in arbitrarily small neighborhoods of points. But a stronger converse may be true, where we drop the condition that approximate moduli and classical moduli are comparable, by using the properties of circle packings.

Chapter 6. Modulus of Subdivision Rules of Unbounded Valence

6.1 Introduction

As described in the introduction to this dissertation, a subdivision rule is a recursive way of combinatorially dividing a tiling on a surface into a smaller, more refined tiling. Barycentric subdivision is an example, as is hexagonal refinement, both of which we will explore later. Subdivision rules have been studied extensively by Cannon, Floyd, and Parry in an attempt to prove that all hyperbolic groups with a 2-sphere at infinity are hyperbolic 3-manifold groups; as a corollary, this would give a simple proof of the hyperbolization conjecture, now proved by Perelman.

While investigating these hyperbolic groups with a 2-sphere at infinity, they showed (in [7]) that all such groups have an associated finite subdivision rule on the sphere. Cannon showed in [2] that, if a subdivision rule associated to a group is conformal (meaning it satisfies two simple axioms), then the group acts on the sphere by Möbius transformations and on hyperbolic 3-space by isometries, cocompactly and properly discontinuously.

Cannon, Floyd and Parry have studied finite subdivision rules extensively under the hypothesis of bounded valence at every vertex. They have showed that barycentric subdivision did not satisfy either of the two axioms mentioned above. However, they noted (in [3], p.25) that barycentric subdivision should have many of the same properties as conformal subdivision rules, and called for an alternate definition of the axioms that would better handle subdivision rules of unbounded valence. We will describe two ways of altering the axioms in Section 6.2.

To state the axioms, we need to define combinatorial modulus. Combinatorial modulus is a direct analog of the modulus of a topological annulus or quadrilateral in complex analysis.

In that setting, every topological annulus or quadrilateral is conformally equivalent to a standard annulus or quadrilateral, and the modulus is defined by how thick or thin the resulting annulus or quadrilateral is.

A conformal map can be thought of as a change in metric; finding the classical modulus, then, is finding an optimal metric in some sense. Combinatorial modulus is defined to mirror this.

A tiling T of a ring R (i.e., a closed annulus) gives two invariants, $M_{sup}(R,T)$ and $m_{inf}(R,T)$, called **approximate moduli**. These are similar to the classical modulus of a ring. They are defined by the use of **weight functions**. A weight function ρ assigns a non-negative number called a **weight** to each tile of T. Every path in R can be given a length, defined to be the sum of the weights of all tiles in the path. We define the **height** $H(\rho)$ of R under ρ to be the infimum of the length of all possible paths connecting the inner boundary of R to the outer boundary. The **circumference** $C(\rho)$ of R under ρ is the infimum of the length of all possible paths circling the ring (i.e. not nullhomotopic in R). The **area** $A(\rho)$ of R under ρ is defined to be the sum of the squares of all weights in R. Then we define $M_{sup}(R,T) = \sup_{\rho} \frac{H(\rho)^2}{A(\rho)}$ and $m_{inf}(R,T) = \inf_{\rho} \frac{A(\rho)}{C(\rho)^2}$. Note that they are invariant under scaling of the metric.

A sequence $T_1, T_2, ...$ of tilings is **conformal (K)** if mesh approaches 0 and:

(i) for each ring R, the approximate moduli $M_{sup}(R,T_i)$ and $m_{inf}(R,T_i)$, for all i sufficiently large, lie in a single interval of the form $[r, Kr]$; and

(ii) given a point x in the surface, a neighborhood N of x, and an integer I, there is a ring R in $N \setminus \{x\}$ separating x from the complement of N, such that for all large i the approximate moduli of R are all greater than I.

Note that mesh approaching 0 is independent of topological metric.

These are the two axioms mentioned earlier. The first axiom is similar to equicontinuity of maps, and the second shows that points do not blow up. In Section 6.2, we give an example of a subdivision rule satisfying axiom 1 but not axiom 2.

One way to achieve an intuitive understanding of these axioms is by circle packings. Haïssinsky showed that, if a sequence of tilings is conformal, then some subsequence of its circle packings will converge to a limit map (if we normalize the circle packings by fixing the image of three points). Thus, throughout this paper, we will use circle packings as a way of studying modulus.

6.2 Barycentric Subdivision

Cannon and Swenson have shown [7] that the subdivision rules arising from closed hyperbolic manifolds are conformal. Later, Cannon, Floyd and Parry [6] showed that barycentric subdivision does not satisfy Axiom 1. We show that by altering the definition of weightings, we can make barycentric subdivision satisfy Axiom 1 on a basis of annuli, but not Axiom 2.

We will use the Layer Theorem [6]. This theorem says that the modulus of an annulus is greater than the sum of the moduli of any disjoint collection of essential sub-annuli. In other words, if you stack several rings together, you get a ring at least as big as all of them put together.

The modulus of a quadrilateral or annulus corresponds to an optimal weighting of the shinglings of the quadrilateral. The reason that unbounded valence subdivision rules have been difficult to study before is that the optimal weight vectors corresponding to a quadrilateral or annulus with tilings of large valence are asymmetrical. Specifically, Cannon, Floyd and Parry showed [4] that optimizing a quadrilateral to measure its height in the traditional, 'fat' way (counting all shingles that intersect a path) measures its width in a 'skinny' way (counting any one set of shingles that covers the path). See Figure 6.1. If many edges come into a single vertex, the fat height paths have to go all the way around the vertex, while the skinny width paths can sneak through the vertex. Thus, even rotationally symmetric tilings with high valence will have degenerate moduli.

My idea is to show that barycentric subdivision satisfies Axiom 1 on a basis of annuli if we change how we place weights. We either place weights at the vertices, or blow up

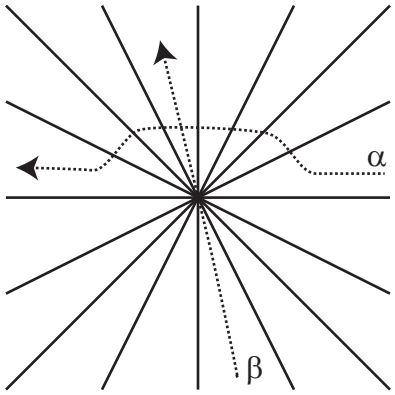

Figure 6.1: Fat paths go around vertices, skinny paths go through them.

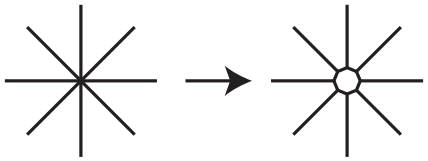

Figure 6.2: Blowing up a vertex.

the vertices by replacing each vertex with a closed disk (see Figure 6.2). Blowing up the vertices gives a new tiling with all vertices have valence 3. Placing weights at the vertices (and requiring paths to be edge paths) is equivalent to taking the dual tiling, where, again, all vertices have valence 3 (because all tiles of barycentric subdivision are triangles). In these settings, 'skinny' and 'fat' sets of shingles are the same, because any two tiles sharing a vertex share an edge. Thus, the height and width are measured in a symmetrical way. Because of this symmetry, Theorem 2.4.5.1 of [4] shows that $m_{inf} = M_{sup}$ for all rings at all stages of subdivision. Thus, we can speak of 'the' combinatorial modulus.

Our future calculations will not depend on the method we choose, as long as fat and skinny are the same. One motivation in putting weights on the vertices is circle packings; a circle packing assigns a weight to each vertex (i.e the radius of the circle).

Now, we perform some calculations. Consider two adjacent triangular tiles an any stage of barycentric subdivision. Together, they form a square. Notice in Figures 6.3, 6.4 that reflection about the common edge preserves the tiling (i.e. cell structure) and swaps cuts

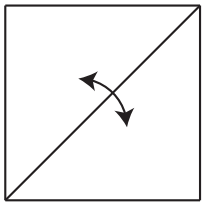

Figure 6.3: A square in a triangulation.

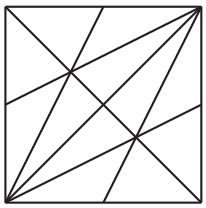

Figure 6.4: The same square after one subdivision.

and flows; in the language of Cannon, Floyd, and Parry, there is a weak isomorphism taking cuts to flows and vice versa [6]. So the optimal height and length of the square must be equal, and either combinatorial modulus of the square must be exactly one.

Now, consider the star of any vertex. It consists of some number n of triangles arranged about the vertex. When we subdivide, we can remove the new star about the vertex (consisting of $2n$ tiles) to get an annulus. This annulus consists of $2n$ square of the type considered above, arranged with alternating orientations.

Let α be a cut in the annulus (i.e. an arc connecting the ends of the annulus). Then pick any square in the annulus, and notice that reflecting α every time it touches the boundary of the square gives us an arc that remains entirely in the square, while having the exact same

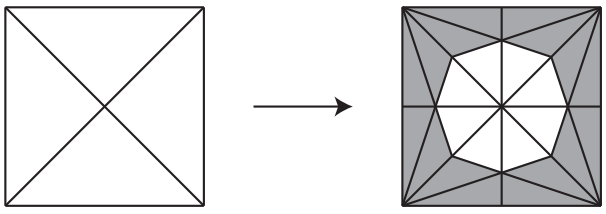

Figure 6.5: The star of every vertex contains an annulus made of squares.

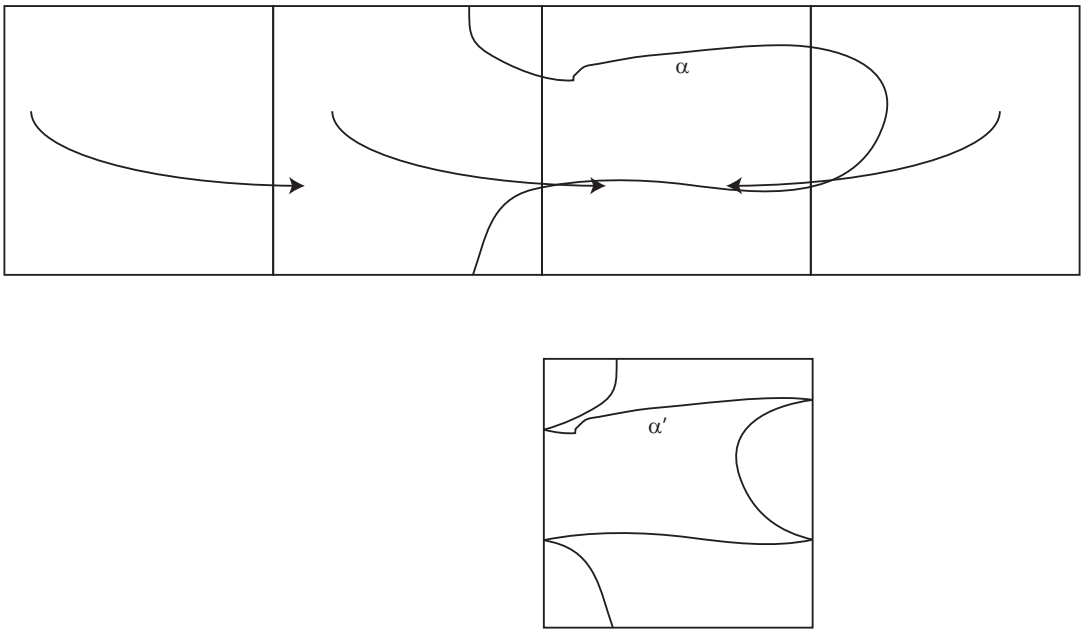

Figure 6.6: If we consider the annulus as four squares with edges identified, we can reflect any cut of the annulus so that it lies in one given square.

length (see Figure 6.6). Thus, the height of the annulus is equal to the height of the square, and the are is equal to n times the area of the square (after we multiply by an appropriate constant). Thus, both moduli are $\frac{1}{n}$. This class of annuli, therefore, satisfy Axiom 1.

However, we cannot layer these annuli to get larger annuli which we know to be of unbounded modulus. To see this, notice that the next smaller annulus has $2n$ squares and has modulus $\frac{1}{2n}$, and, in general, the $(k+1)$th annulus has modulus $\frac{1}{2^k n}$, and layering all of these together gives us a larger annulus which has $\sum_{k=0}^{\infty} \frac{1}{2^k n} = \frac{2}{n}$ as a lower bound for its modulus. The exponential growth about the vertex makes the modulus shrink too quickly. This contrasts strongly with the finite valence case, in which we have nested annuli with identical moduli that sum to infinity. In the next section, we will consider subdivision rules that have linear growth at each vertex, and which thus have annuli of unbounded size.

While we don't have an upper bound for the size of the annuli, Ken Stephenson's Circlepack [21] suggests that such a limit exists. If a basis of annuli satisfy Axiom 1 and Axiom

2, the circle packings will converge to a limit function on the sphere [11]. However, the packings for barycentric subdivision seem to converge to a relation in which the set of vertices is mapped to a dense set of disks. See the Figure 6.2. Note that all points that are *not* vertices are contained in the double star of a new vertex at every stage of subdivision. The star of each of these vertices contains an annulus surrounding the point of modulus $\frac{1}{12}$ or $\frac{1}{8}$ for all sufficiently refined stages of subdivision. Thus, we can layer to get unbounded moduli in this case. This means that, under a sequence of circle packings, the complement of the vertices is mapped homeomorphically (by Arzela-Ascoli and the fact that locally conformal circle packings are locally equicontinuous [11]) to a subset of the sphere, i.e. the complement of the disks, which is similar to a Julia set.

This is related to the work of Cannon, Floyd and Parry on subdivision rules and rational maps [3], in which they showed that many subdivision rules can be realized as a rational map on the sphere, where the subdivision rule is obtained by pulling back a cell structure on the sphere. Conformal subdivision rules have a Julia set that is the entire sphere, but the map associated to barycentric subdivision has a Sierpsinski carpet as its Julia set.

6.3 Borromean Rings

In contrast to barycentric subdivision, a subdivision rule associated with the Borromean rings is conformal, although the valence at each vertex remains unbounded.

The tile types are shown in Figure 6.7. The dot indicates the orientation of the B tiles.

Let's examine the star of a vertex. Each new vertex in the subdivision rule, after all tiles around it have been subdivided once, has a star of the form shown in Figure 6.8.

After another subdivision, we have Figure 6.9. Note the marked annulus.

Notice that this annulus can be reflected twice in between neighboring B tiles to get a quadrilateral of the form shown in Figure 6.10. Several B tiles have had their bottom halves cut off.

But we can fold this up into itself as shown in Figure 6.11

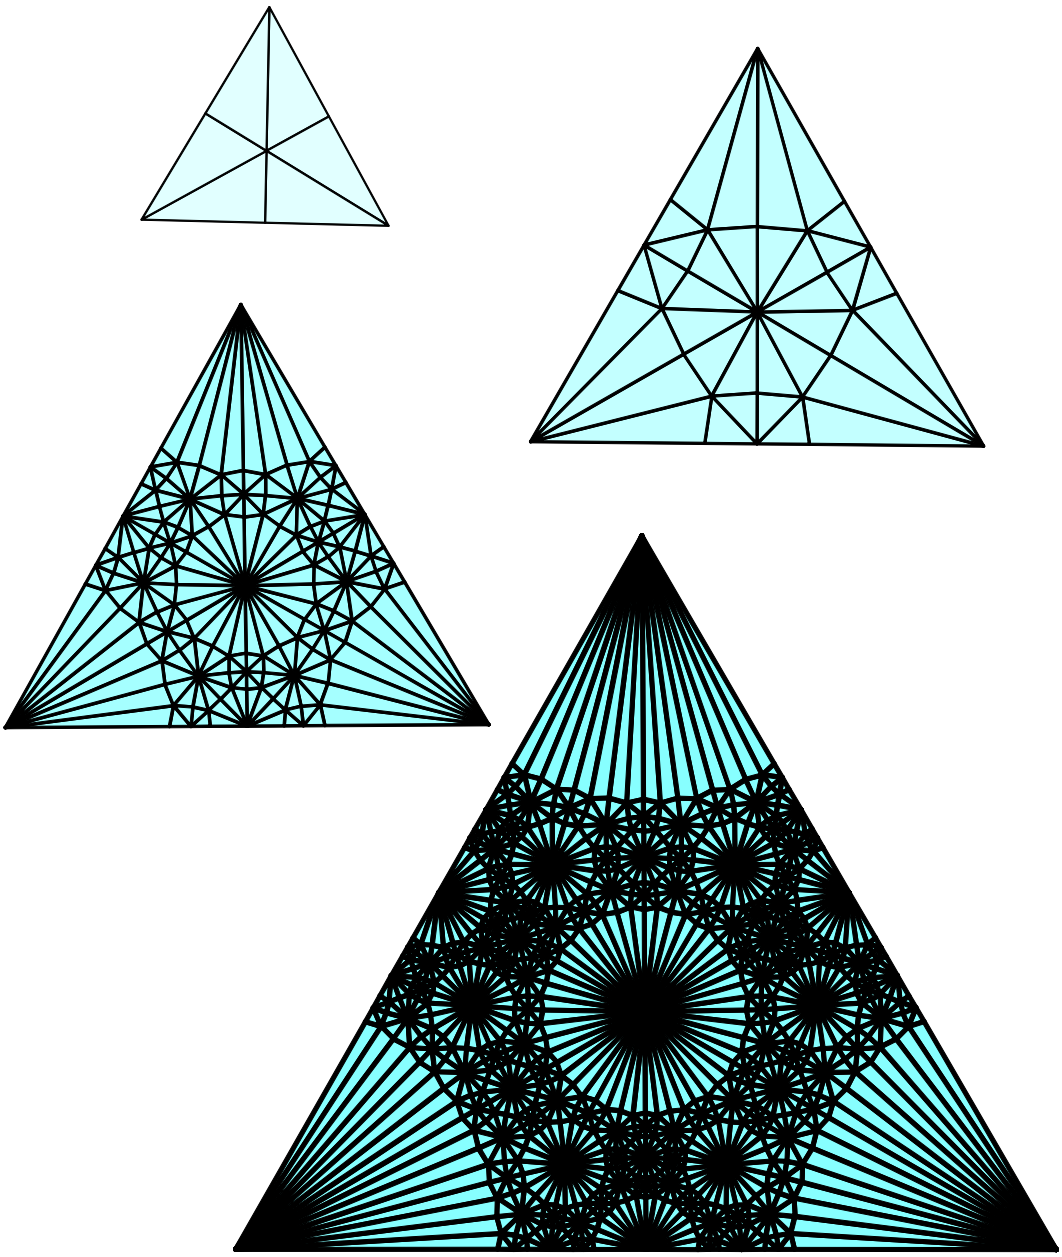

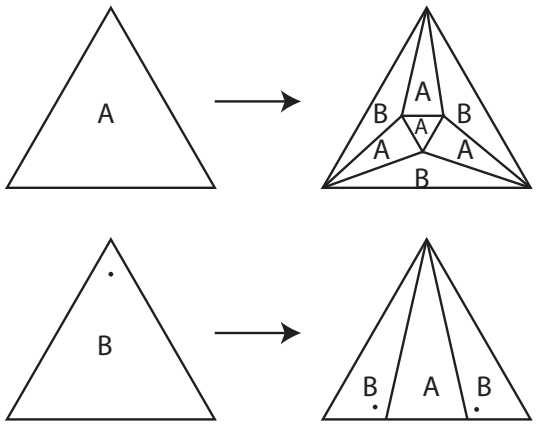

Figure 6.7: One subdivision rule for the Borromean rings. All polyhedra are octahedra.

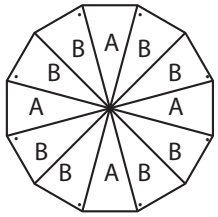

Figure 6.8: The star of a vertex in the Borromean rings' subdivision rule.

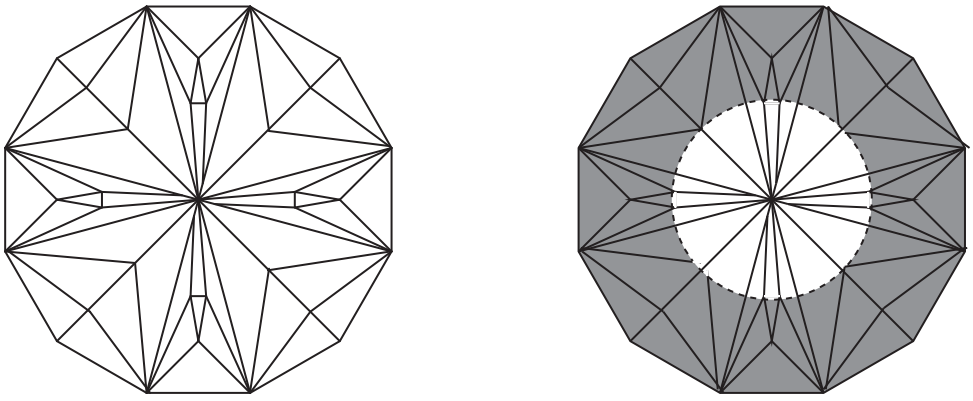

Figure 6.9: We find an annulus in the subdivided star whose modulus is easy to calculate.

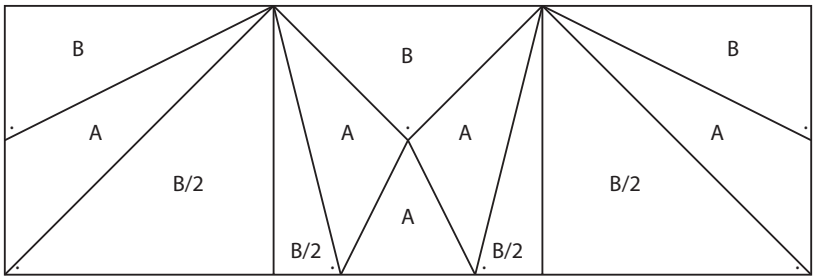

Figure 6.10: One of the four quadrilaterals that make up the annulus in Figure 6.9. Note that the two pairs of $B\backslash 2$ tiles don't form two whole B tiles; their other halves are below the picture.

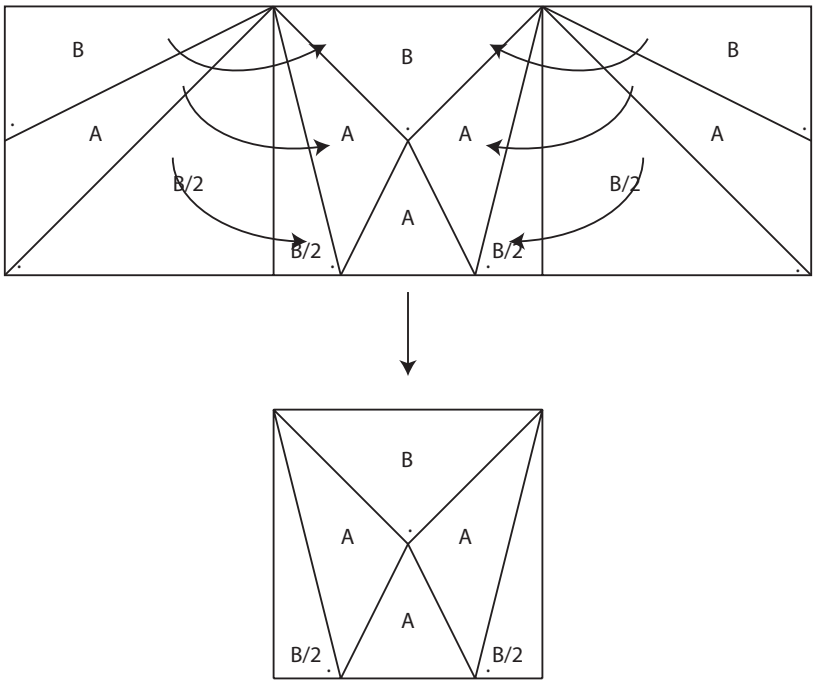

Figure 6.11: There is a weak cellular folding (for definition, see [6]) that takes the two outer tiles into inner tiles, with tops and bottoms going to tops and bottoms.

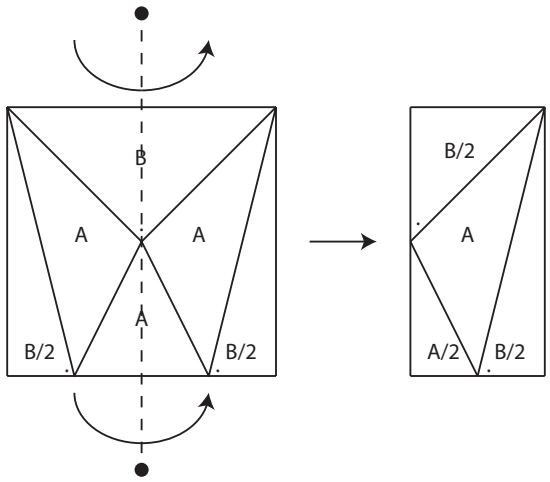

Figure 6.12: The tile is folded along a line of symmetry.

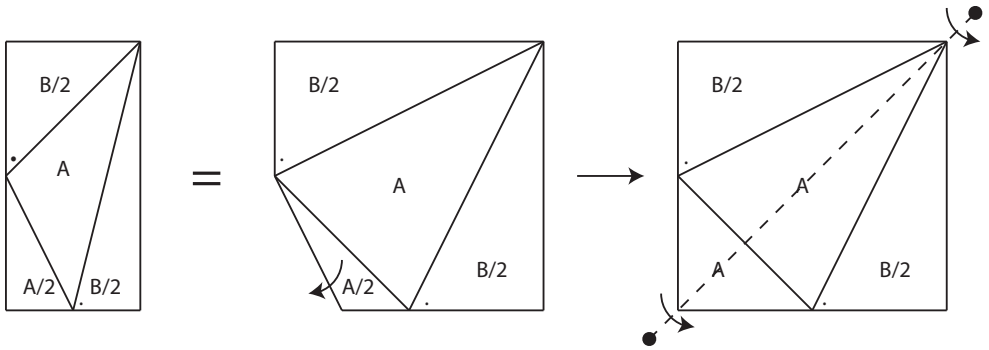

Figure 6.13: We first redraw our shape to show how it is to be unfolded, and then unfold it. The shape on the right has modulus 1, by symmetry.

This tile, in turn, can be folded in on itself, as shown in Figure 6.12.

We almost have something whose modulus we can calculate. In one final step, we unfold our shape to get something with modulus 1, as shown in Figure 6.13.

This final shape (call it Q) has modulus 1 in all stages of subdivision, as it has a symmetry interchanging the top/bottom with the left/right. We can now estimate the modulus of the original annulus by "pulling back" the weighting given to this tile. We do this by undoing each step.

Putting back the fold in Figure 6.13, we get something with the same height as Q, but

with less area than that of Q. Thus, the modulus is at least 1. Undoing the steps in Figure 6.12 at most doubles the area, giving a quadrilateral with modulus at least $\frac{1}{2}$. Undoing Figure 6.11 gives us modulus at least $\frac{1}{6}$, since area at most triples. Unfolding this into the original annulus gives modulus at least $\frac{1}{24}$.

We can find similar annuli in the star of each vertex at any stage of subdivision, whose modulus can be calculated in the same way; the annulus at each stage can be split into four quadrilaterals which have more squares in them, as in Figure 6.14. By reflection, we can fold in the extra tiles, and counting arguments show that the modulus is at least $\frac{1}{2n}$, where n is the number of squares in the star. The annuli thus obtained are disjoint, and so we can layer them to get annuli whose moduli are bounded below by the partial sums of a harmonic series. Assuming that the actual modulus is not much greater than this, we have logarithmic growth of modulus at each vertex. This is supported by circle packings, like Figures 6.3 and 6.3, where the stars of each vertex seem to shrink slowly.

6.4 Hexagonal refinement

Hexagonal refinement is a classic example of a conformal subdivision rule. The subdivision rule takes triangles and subdivides them into four smaller triangles by connecting midpoints of edges (see Figure 6.4). While it is easy to calculate the exact modulus of annuli in this sequence of tilings, we give lower bounds for the modulus as in our other examples.

At every stage of subdivision, every vertex is surrounded by an annulus of 18 triangles, as shown in Figure 6.15. By reflection, we can consider a single quadrilateral as shown in Figure 6.16. Note that, by the reflection shown in that figure, we can assume that any cuts remain in the right two triangles. These last triangles, by the reflection shown in Figure 6.17, have modulus 1. Working backwards, we see that the quadrilateral in Figure 6.16 has modulus $\geq \frac{1}{2}$. Thus, the entire annulus in Figure 6.15 has modulus $\geq \frac{1}{12}$.

In contrast to barycentric subdivision and the Borromean rings, the annuli surrounding a given vertex at each stage have a fixed, constant size $C \geq \frac{1}{12}$. Thus, layering gives us annuli

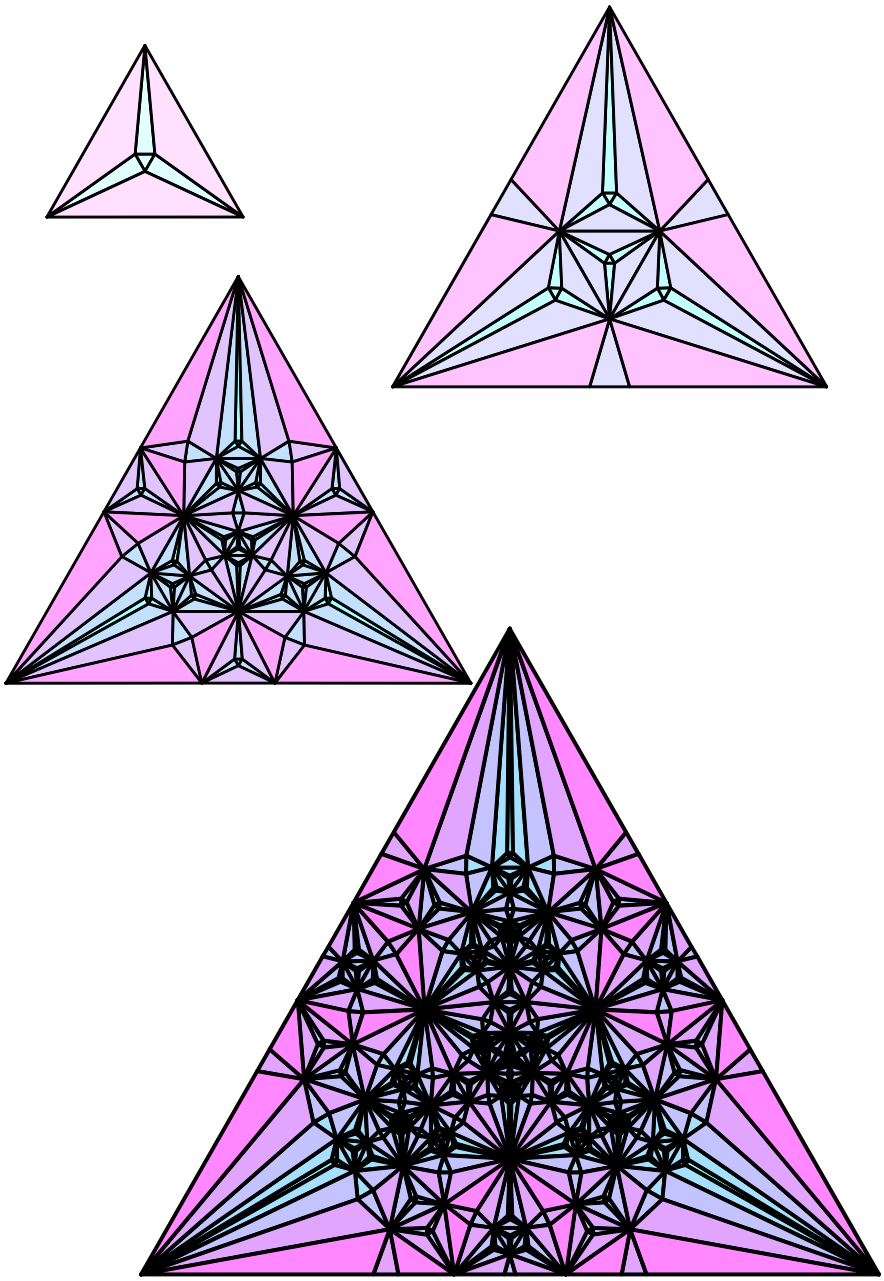

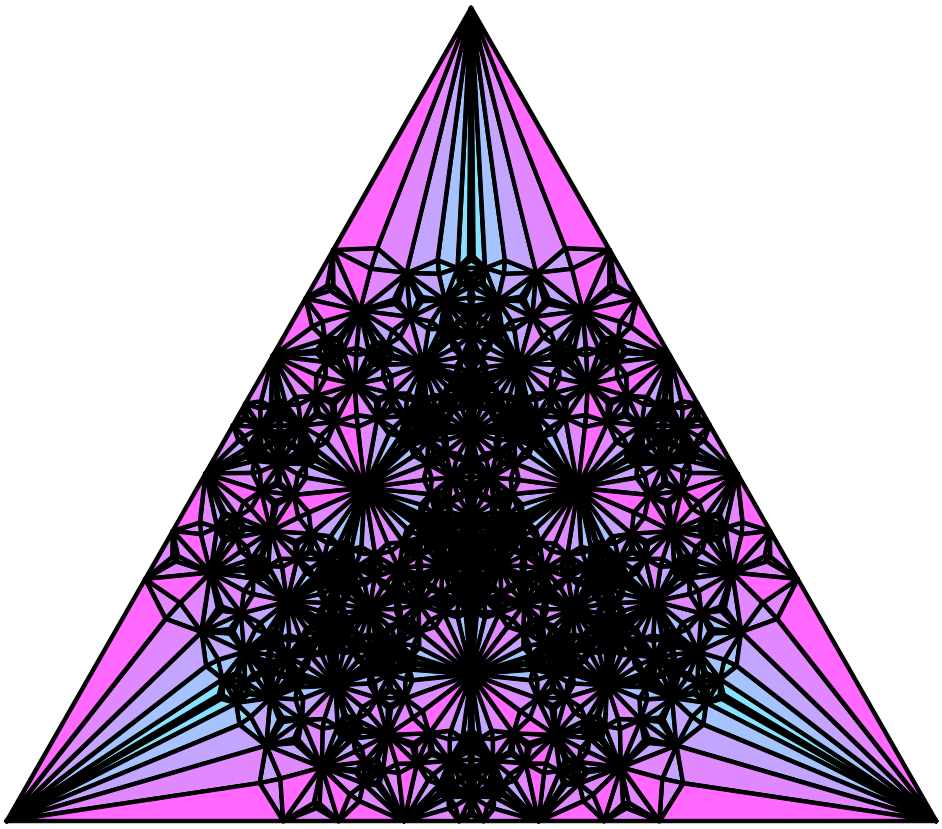

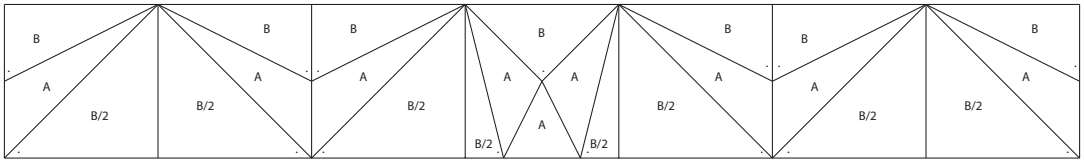

Figure 6.14: A figure similar to Figure 6.10, but for the star of a vertex at a later stage of subdivision.

of modulus $\geq Cn$ at stage n, and so axioms 1 and 2 are satisfied at each vertex. This easy layering for finite valence subdivision rules is a small part of why Cannon, Floyd, and Parry were able to replace axioms 1 and 2 with a simpler axiom 0 for finite valence subdivision rules [5]:

Axiom 0. Given $x \in Y$ and a neighborhood N of x, there is a ring R in N surrounding x such that the moduli $m(R, T_i)$ are bounded away from 0.

Note that the estimate for every vertex in hexagonal refinement is the same at every stage; this is reflected in its circle packing on page 113.

6.5 Generalizations

Notice that the key ideas in the estimates for our three subdivision rules were:

(i) Finding pieces of modulus 1, and

(ii) Showing that we need only consider cuts that remain in these pieces.

One powerful technique related to these ideas is the 1,2,3-tile criterion of Cannon, Floyd, and Parry [6]. Essentially, a subdivision rule satisfies the criterion if every 'test quadrilateral' formed from one or two tiles in a row (or three if the middle one is a triangle) has modulus uniformly bounded away from 0. This criterion allows us to argue as we did in the three cases above to provide estimates.

There are three types of test quadrilaterals:

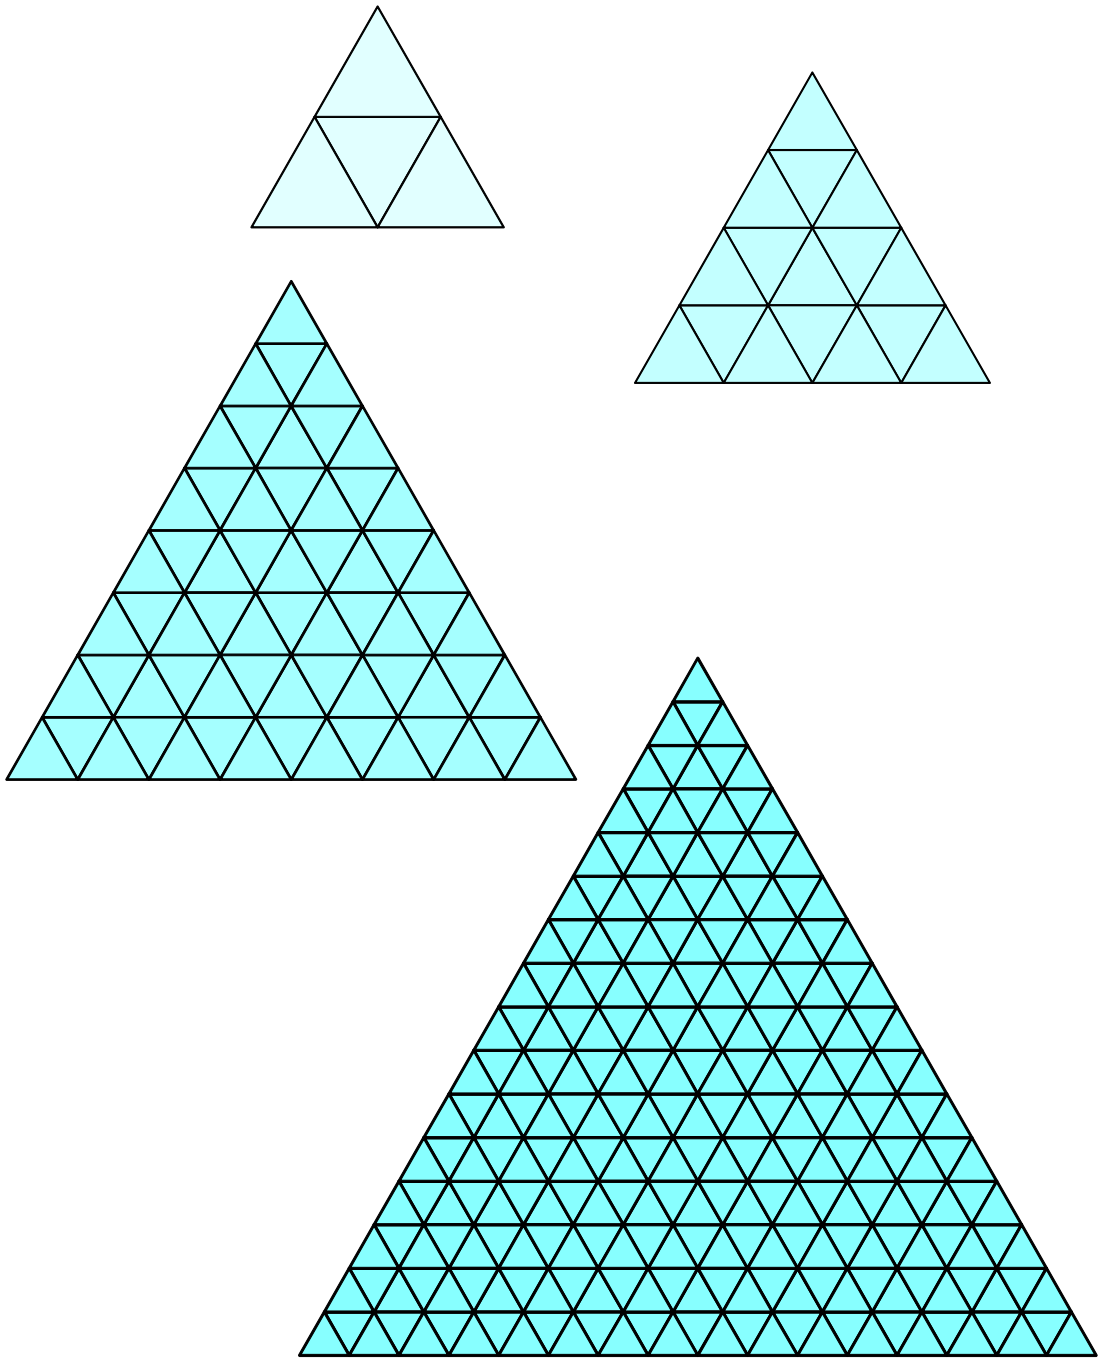

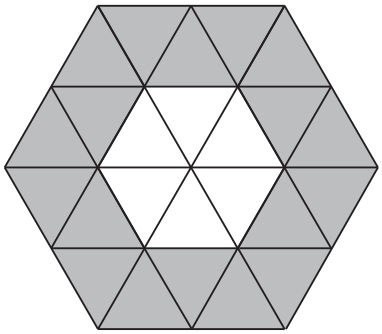

Figure 6.15: Eventually, at every stage of subdivision, every vertex is surrounded by an annulus of this form.

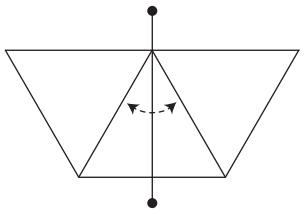

Figure 6.16: The annulus in Figure 6.15 is made of 6 quadrilaterals of this form. By reflection, we can assume that all cuts stay in two neighboring triangles.

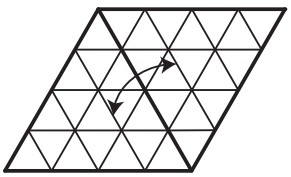

Figure 6.17: This square has modulus one, even after several stages of subdivision.

Type I. The quadrilateral Q consists of one tile, and each of the ends of Q consists of one edge.

Type II. The quadrilateral Q consists of two tiles whose intersection consists of one edge f, and each of the ends of Q consists of one edge which meets f. We call f the **interior edge** of Q.

Type III. The quadrilateral Q consists of three tiles t_1, t_2 and t_3, where t_2 is a triangle with edges f_1, f_2, and f_3. The intersection of t_1 and t_2 is f_1, and the intersection of t_2 and t_3 is f_3. The intersection of t_1 and t_3 is a vertex v. The top of Q consists of an edge containing v, and the bottom of Q is f_2. We call f_1 and f_3 the **interior edges** of Q.

Cannon, Floyd and Parry showed that a uniform lower bound on the modulus of all test quadrilaterals (call it M) gives a lower bound for the modulus of a ring R in our surface, depending on the size of the ring. Specifically, in the proof of Theorem 5.1 in [6], they take the star of a simple closed curve α in some subdivision R^i of the ring, where α misses every vertex, and show that any flow curve from the bottom of this star to the top must join the ends of a test quadrilateral (i.e. its image in some test quadrilateral is a cut). They then give the star of α a weight vector that is the sum of all optimal weightings of test quadrilaterals, normalized to have area 1. The star of α, then, has modulus $\geq \frac{M}{Ak}$, where k is the number of tiles in the star of α in R^i and A is the maximum area a single tile can have under the weighted sum (which is seen to be bounded, as the number of test quadrilaterals a single tile can lie in is uniformly bounded).

When the star of a vertex v in a subdivided surface S is a disk, we can take its boundary and push it slightly outward to be such an α. The number of tiles in the star of α will then be $\leq B_i(v)val_i(v)$, where $B_i(v)$ is the maximum valence of vertices surrounding v. In a subdivision rule with combinatorial mesh going to 0, each vertex v has a uniform bound on $B_i(v)$, as each edge surrounding the vertex is subdivided periodically. Thus, we see that any annulus containing the star of α must have modulus $\geq \frac{C}{val_i(v)}$ for some constant C.

By nesting this sequence of annuli, we can construct an annulus of larger modulus. When valence grows exponentially at each vertex, as in barycentric subdivision, we get a geometric sequence of estimates of the size of annuli, and we cannot construct annuli whose modulus estimates are unbounded. I conjecture that every subdivision rule with exponential growth at each vertex which satisfies the 1,2,3-tile criterion will satisfy axiom 1 but not axiom 2, just like barycentric subdivision. In this case, the limit of circle packings will be a continuous relation instead of a continuous function onto the sphere, where points are blown up into disks. In 1, there are subdivision rules for hyperbolic 3-manifolds with hyperbolic surface boundary; I conjecture that all such subdivision rules satisfy the 1,2,3-tile criterion and behave like barycentric subdivision at each vertex corresponding to a hyperbolic surface in the boundary.

Nesting annuli in a subdivision rule of linear growth gives us a harmonic series for our estimates, which shows that all such subdivision rules that satisfy the 1,2,3-tile criterion are conformal. I conjecture that all the subdivision rules arising from alternating links in [18] satisfy this criterion, and are conformal.

Finally, subdivision rules of bounded valence give a linear modulus estimate, and those that satisfy the 1,2,3-tile criterion are conformal, which was why Cannon, Floyd and Parry originally proved the criterion in [6].

Chapter 7. Non-Conformal Subdivision Rules and Circle Packings

We now construct an interesting class of subdivision rules. Consider the binary square subdivision rule, shown in Figure 7.1. This subdivision rule is conformal. We can 'inflate' this subdivision rules by adding extra tiles about each vertex, in several different ways, as shown in Figures 7.2 and 7.3.

These three variants behave differently. Each takes the vertices of the binary subdivision rule and replace them with a series of nested rings, each ring of modulus going to 0, ∞, and $c > 0$ for some c, respectively. We would expect widely varying behavior when we circle pack them. However, when we construct the packings, we get almost identical pictures. See the Figure 7.4.

What is going on? In the binary subdivision rule, each vertex was surrounded by annuli of arbitrarily large modulus. Inflating these vertices did not decrease the size of these annuli. Each B tile (or C tile) is surrounded by infinitely many disjoint annuli of modulus $\frac{1}{4}$, each made up of four type A tiles. No matter what subdivision occurs within these annuli, the circle packings will push it all to one point, by the equicontinuity lemma in Chapter 5.

This is an example of a sequence of tilings whose normalized circle packings converge,

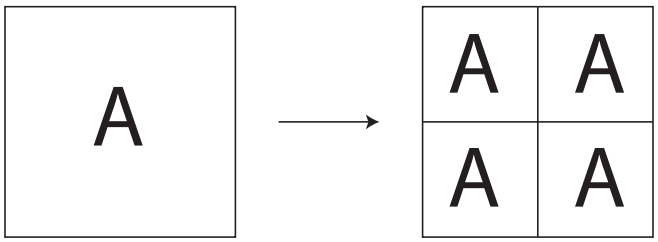

Figure 7.1: The binary square subdivision rule.

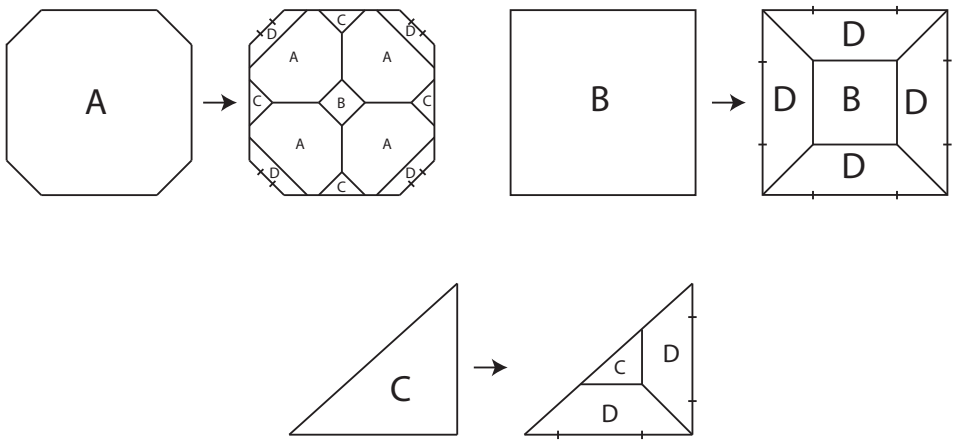

Figure 7.2: Tile types for the inflated binary subdivision rules. All three subdivision rules share these types in common. Note that if all B,C,D, and E tiles are collapsed, we recover the original binary subdivision rule.

but not to a homeomorphism. All three varieties are conformal on a dense set (i.e. all points in the set have a family of annuli surrounding them satisfying axioms 1 and 2), and are not conformal on the complement. In all three examples, each point in the 'bad' set has large annuli around it with modulus going to infinity. In one example, each point in the interior of the bad set has a neighborhood where all annuli in it satisfy axioms 1 and 2; in one example, each point in the interior of the bad set has annuli with modulus going to 0; and in the other example, each interior point has annuli with modulus going to 0 except for the exact center, whose moduli all go to infinity. See Figure 7.5.

Clearly, other conformal subdivision rules can be blown up in this same way.

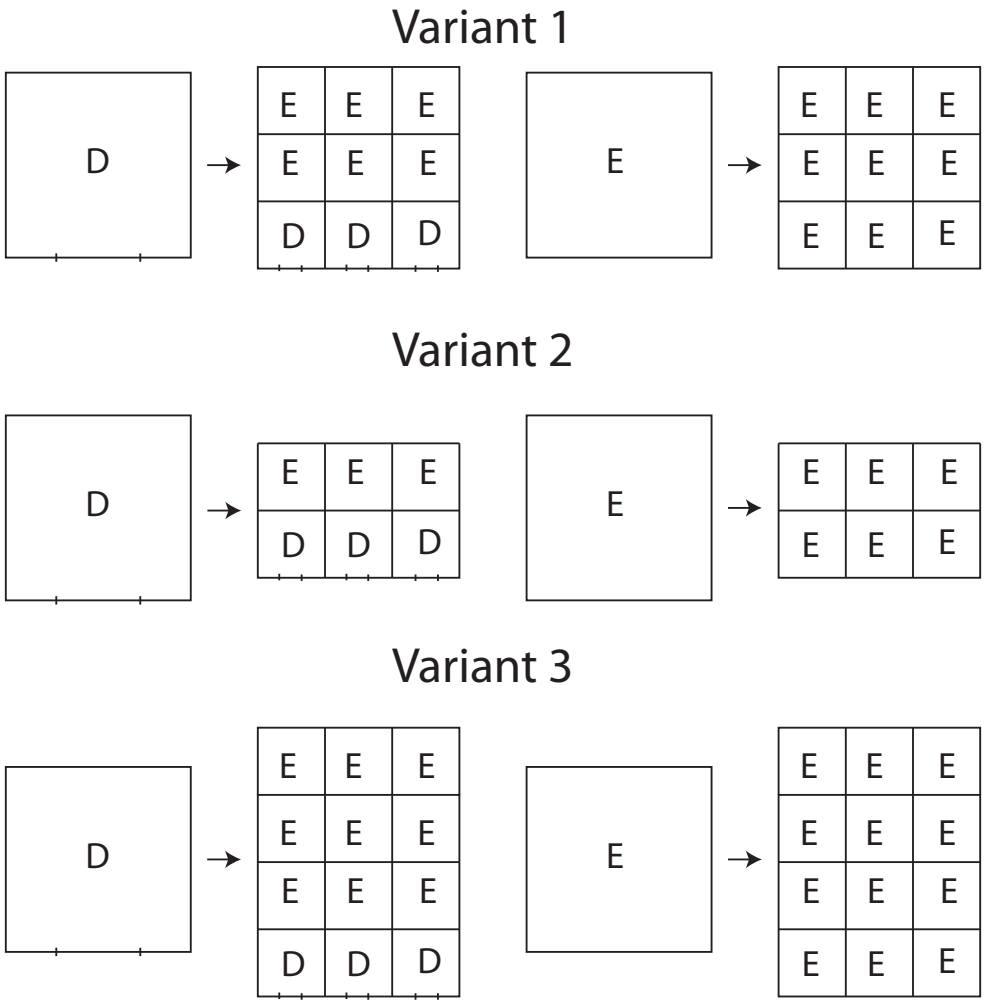

Figure 7.3: The last two tile types for each of the inflated binary subdivision rules.

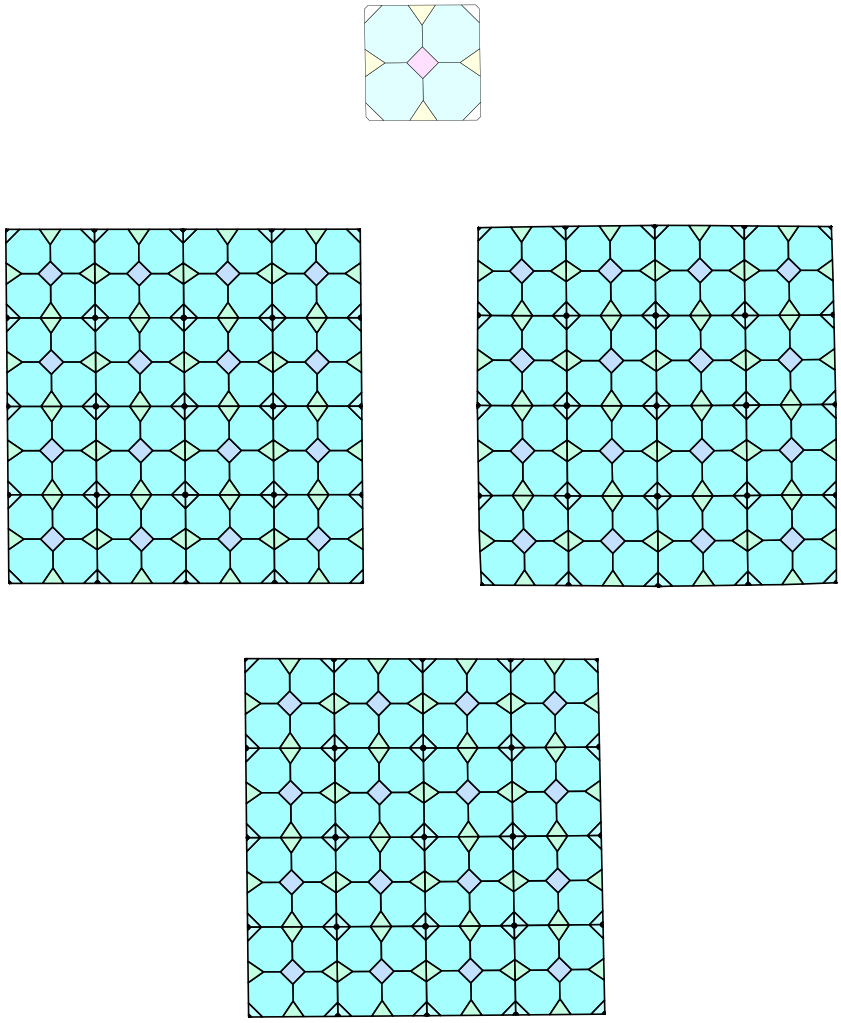

Figure 7.4: The first and third stages of the octagonal tile type for the three variants of our non-conformal subdivision rule in Chapter 7. Notice that the third stages are indistinguishable at this level of detail.

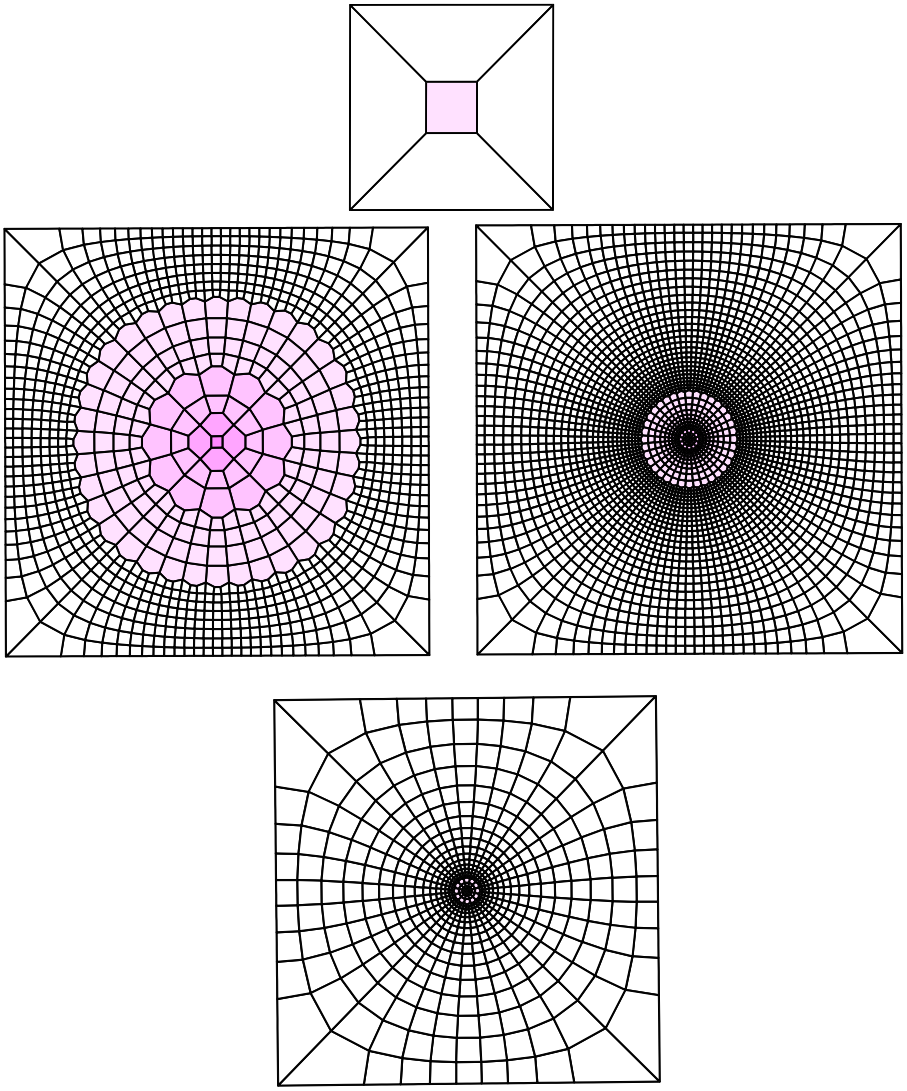

Figure 7.5: The same subdivision rules as Figure 7.4, but now the first and fourth stages of the square tile type (i.e. the blown-up vertices). The pink center of the top square is carried to the pink centers in the other squares. In the 3x2 subdivision rule on the left, the pink center will eventually grow to fill the whole interior of the square. In the 3x3 subdivision rule on the right, its position is stable. In the 3x4 subdivision rule on the bottom (which we have only shown the third stage of, as the fourth is much too large), the pink center contracts down to a point.

Part III

Future Work and References

Chapter 8. Future Work

We would like to find a concise set of conditions on odd edge-cycle length polyhedral gluings that ensures a subdivision rule exists.

We would still like to find subdivision rules for Nil and Sol geometries.

We would like to show that composites of alternating links have subdivision rules like those in Chapter 7, with a non-Hausdorff 2-sphere at infinity that has n Hausdorff spheres as quotients, one for each summand of the link (as experiments have suggested).

We would like to show that all subdivision rules we have found for hyperbolic manifolds with boundary are conformal in the sense that the barycentric subdivision or the subdivision rule for the Borromean rings are conformal.

We would like to prove that conformal unbounded valence subdivision rules are associated to rational maps, and to extend other results of Cannon, Floyd and Parry.

We would like to find symmetries in subdivision rules of all hyperbolic groups with 2-spheres at infinity to prove that they are conformal.

Bibliography

[1] E. M. Andreev, *Convex polyhedra in lobačevskiĭ spaces (english transl.)*, Math USSR Sbornik **10** (1970), 413–440.

[2] J. W. Cannon, *The combinatorial riemann mapping theorem*, Acta Mathematica **173** (1994), no. 2, 155–234.

[3] J. W. Cannon, W. J. Floyd, R. Kenyon, and W. R. Parry, *Constructing rational maps from subdivision rules*, Conformal Geometry and Dynamics **7** (2003), 76–102, (electronic).

[4] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Squaring rectangles: the finite riemann mapping theorem*, The Mathematical Legacy of Wilhelm Magnus: Groups, Geometry and Special Functions, Contemp. Math., vol. 169, Amer. Math. Soc., May 1992, pp. 33–50.

[5] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Sufficiently rich families of planar rings*, Annales Academiæ Scientiarum Fennicæ Mathematica **24** (1999), 265–304.

[6] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Finite subdivision rules*, Conformal Geometry and Dynamics **5** (2001), 153–196.

[7] J. W. Cannon and E. L. Swenson, *Recognizing constant curvature discrete groups in dimension 3*, Transactions of the American Mathematical Society **350** (1998), no. 2, 809–849.

[8] J.W. Cannon, *The combinatorial structure of cocompact discrete hyperbolic groups*, Geom. Dedicata **16** (1984), 123–148.

[9] J.W. Cannon, W.J. Floyd, M.A. Grayson, and W.P. Thurston, *Solvgroups are not almost convex*, Geom. Dedicata **31** (1989), no. 3, 291–300.

[10] W. J. Floyd, *tilepack.c, tilepackhistory.c, subdivide.c and subdividehistory.c*, Software, available from http://www.math.vt.edu/people/floyd/.

[11] P. Haissinsky, *Empilements de cercles et modules combinatoires*, Annales de l'Institut Fourier **59** (2009), no. 6, 2175–2222.

[12] C. D. Hodgson, I. Rivin, and W. Smith, *A characterization of convex hyperbolic polyhedra and of convex polyhedra inscribed in the sphere*, Bull. Amer. Math. Soc. **27** (1992), no. 2, 246–251.

[13] O. Lehto and K.I. Virtanen, *Quasiconformal mappings in the plane*, second ed., Springer-Verlag, 1973.

[14] D. Papp, *Visualization of nil-geometry*, Proceedings of the Dresden Geometry Symposium, November 2003.

[15] G. Perelman, *Finite extinction time for the solutions to the ricci flow on certain three-manifolds*, 2003, arXiv:math.DG/0307245.

[16] W. Rudin, *Principles of mathematical analysis*, International series in pure and applied mathematics, McGraw-Hill, 1976.

[17] B. Rushton, *Alternating links and subdivision rules*, Master's thesis, Brigham Young University, 2009.

[18] _____, *Creating subdivision rules from alternating links*, Conformal Geometry and Dynamics **14** (2010), 1–13.

[19] P. Scott, *The geometries of 3-manifolds*, Bulletin of the London Mathematical Society **15** (1983), 401–487.

[20] M. Shapiro and M. Stein, *Almost convex groups and the eight geometries*, Geom. Dedicata **55** (1993), no. 2, 125–140.

[21] K. Stephenson, *Circlepack*, Software, available from http://www.math.utk.edu/∼kens.

[22] _____, *Introduction to circle packing:the theory of discrete analytic functions*, Cambridge University Press, 2005.

[23] W. P. Thurston, *Three-dimensional geometry and topology*, vol. 1, Princeton University Press, 1997.