2009-03-12

Alternating Links and Subdivision Rules

Brian Craig Rushton
Brigham Young University - Provo

ALTERNATING LINKS AND SUBDIVISION RULES

by

Brian Rushton

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

Department of Mathematics

Brigham Young University

April 2009

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Brian Rushton

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory.

Date	James W. Cannon, Chair
Date	Michael Dorff
Date	Jessica S. Purcell
Date	Gregory R. Conner

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Brian Rushton in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

Date	James W. Cannon Chair, Graduate Committee
Accepted for the Department	Gregory R. Conner Graduate Coordinator
Accepted for the College	Thomas Sederberg, Associate Dean College of Physical and Mathematical Sciences

ABSTRACT

ALTERNATING LINKS AND SUBDIVISION RULES

Brian Rushton

Department of Mathematics

Master of Science

The study of geometric group theory has suggested several theorems related to subdivision tilings that have a natural hyperbolic structure. However, few examples exist. We construct subdivision tilings for the complement of every nonsingular, prime alternating link and all torus links, and explore some of their properties and applications. Several examples are exhibited with color coding of tiles.

ACKNOWLEDGMENTS

This work of course would not be possible without the many years of work by my advisor Jim Cannon and his friends Bill Floyd and Walter Perry. Bill Floyd and Ken Stephenson provided an enormous amount of computer support for their subdivision and circle-packing programs, respectively. I'd also like to acknowledge the support of my roommates and friends who were forced to see many, many drafts of this thesis and explanations using Play-doh and floss. The pictures were all created using Adobe Illustrator, except for the Circlepack images found in Section 13 and elsewhere. Thanks to Kristin Gulledge for helping with some of the nastier 3-d drawings. Thanks to Dr. Steve Humphries for listening to my theories and showing me the real-life link that was the inspiration for the platonic links. Heidi Cook provided green Play-doh, without which this work would never have started. Nathan Priddis reviewed an early draft and provided a great deal of insight. And Lonette Stoddard was a great help in getting this into shape.

Contents

List of Figures

1 Introduction

This thesis attempts to bring together two beautiful and deep areas of mathematics: subdivision rules and alternating knots. Alternating knots will perhaps be more familiar to most readers, so a description of subdivision rules may be helpful. As many new terms will be defined throughout the thesis, Section 12 provides a list of definitions.

Essentially, a subdivision rule is given by a set of planar tiles together with a set of rules for subdividing each tile into several smaller tiles, each of which is in turn subdivided, etc. to produce a series of increasingly complicated tilings. They can be represented by a list of labelled polygons whose interior is divided up by a finite number of edges into smaller labelled polygons. The edges can also divided, for instance if we connect the midpoints of a triangle to subdivide it into four smaller triangles. For a more rigorous definition, see [7]. Subdivision rules in a sense are generalizations of fractals. Instead of repeating the exact same design over and over, we get slight variations in each stage, allowing a richer structure while maintaining the elegant style of fractals. An example of a subdivision rule is shown below:

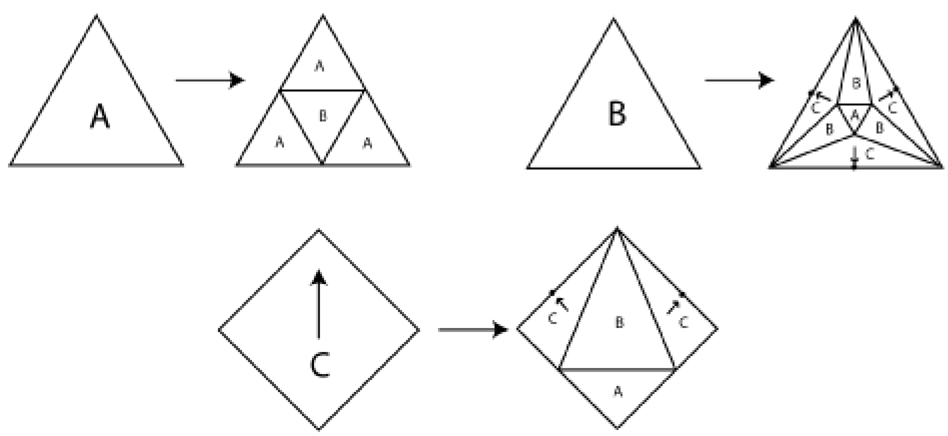

The arrows are used to indicate which direction the tiles are pointing, because tile

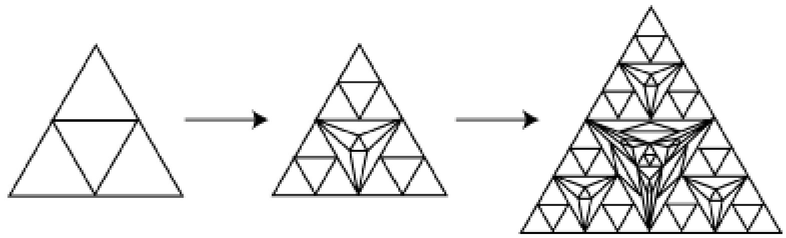

Figure 1: The first three subdivisions of a type A tile.

C is asymmetrical. Note that both edges and faces of tiles are subdivided. We can use this subdivision rule to recursively subdivide each tile. The first three subdivisions of a type A tile are shown in Figure 1.

The idea of subdivision rules follows naturally from observation of nature. Cell division is a good example. A cell naturally divides into other cells which, in turn, divide into other cells, etc. but in response to different chemicals, expressed genes, or environmental factors, the division patterns of the later cells may be different from the original. This is explored in many biology articles and books, including [17], especially Chapter 7.

Subdivision rules also arise in art, especially Islamic tile-based art. Some beautiful examples of this are found in the girih patterns in medieval mosques and mausoleums [13].

Our main mathematical motivation for subdivision rules, however, is the study of hyperbolic groups [7]. One important open question is whether every Gromov hyperbolic group with a 2-sphere as its space at infinity acts properly discontinuously, cocompactly, and isometrically on hyperbolic 3-space. Cannon and Swenson [8] showed that if a certain sequence of coverings of the Cayley graph by disks is conformal, then the group acts isometrically, properly discontinuously, and cocompactly on hyperbolic 3-space. Conformal in this context is defined in [5]. This sequence of disk coverings can best be studied through examining the properties of associated subdivision rules. Essentially, a sequence is conformal when the tiles have diameter that goes to 0 'nicely'.

Subdivision rules can also be used to measure the growth of hyperbolic groups, a topic which has received a lot of attention over the past few decades (see [9],[10]; more papers can be found in the references of [4]). They are closely related to rational maps [6] (see also Section 5.1).

Unfortunately, we don't have as many subdivision rules as we'd like, especially subdivision rules related to topological spaces. Our goal is to construct a subdivision rule for every alternating link.

The other main subject of this thesis is alternating knots. The study of knots goes back for centuries. Gauss studied knots as early as 1792, and may have begun tabulating knots [16]. Knots are particularly interesting to topologists and algebraists, as they provide many examples of 3-manifolds and infinite groups. Alternating knots are a particularly simple set of knots; they are the knots whose intersections follow an "over, under, over, under" pattern. Some of their particularly nice properties are explored in Sections 2-4. For those interested in learning more about knots, [1] is a good introductory book, written on a beginning university level, and [16] is appropriate for more advanced readers.

We bring the two subjects, knot theory and subdivision rules, together by constructing subdivision rules for the universal covers (and thus for the Cayley graphs) of the complements of alternating knots. This is a difficult process, and our initial methods fail for certain knots, as explored in Section 6. However, two large classes of knots (one a subset of the other) **do** have nice subdivision rules, the radial links and the dispersed links. These are discussed in Sections 7 and 8, respectively. By altering the method slightly, we obtain subdivision rules for all torus links, as discussed in

Section 10. By altering the method significantly, we obtain a subdivision rule for all non-singular prime alternating links. This is covered in Section 11.

There is much future work that can be done in this area. There are several closed 3-manifolds that have simple subdivision rules, and this will be covered in a future work.

2 Alternating Links

Definition. Closed graphs are images of S^1 in $\mathbb{R}^2$ such that all intersections are non-tangential. Also the number of intersections must be finite.

All results in this section are classical. Theorem 1 is found in many books, including [16], but I first found it in [1]. I developed the other proofs independently. Again, all definitions can be found in Section 12.

Theorem 1. Every closed graph can be realized as a projection of the unknot.

Proof. Given such a graph G, let $f : [0,1] \to \mathbb{R}^2$ be a parametrization of G with nonzero tangents at every point such that $f(0) = f(1)$ has a larger y-value than every other point. Define $\tilde{f} : [0,1] \to \mathbb{R}^3$ be defined by $\tilde{f}(x) = (f(x), x)$. Then $\tilde{f}[0,1] \cup \{\tilde{f}(0) \times [0,1]\}$ is a closed path. Its projection to the xy plane is G, but its projection in the yz plane has no intersections. Thus, this is the unknot. $\square$

Theorem 2. Any two closed graphs A, B intersect an even number of times.

Proof. By Theorem 1, A and B can be realized as projections of the unknot. So we can think of A and B as two unknotted loops lying on top of each other. Thus, there is a sequence of Reidemeister moves (see Section 12 for a definition) that takes the original graph to the following:

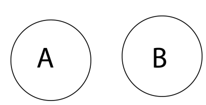

A type I move will not change the number of intersections between the loops, since it only affects self-intersections. A type II move will change the number of AB-intersections by $0, 2$ or -2 (see Figure 2).

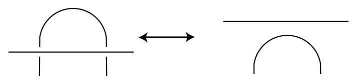

Figure 2: A type II move.

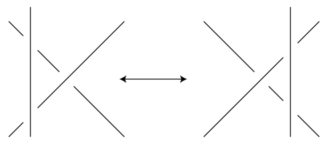

Figure 3: A type III move.

Finally, a type III move does not change the number of intersections (see Figure 3).

Thus, Reidemeister moves do not affect the parity of the number of AB-intersections. Since the final graph has 0 intersections, our original graph must have an even number of intersections. □

Theorem 3. Every closed graph can be realized as the projection of an alternating knot.

Proof. Pick a point in the graph G and an orientation. Travel along the curve, labelling each intersection you reach 1,2,3, etc. Note that every intersection will be reached twice, so every intersection receives two labels. In particular, this means that every closed graph intersects itself an even number of times, counting in this way.

Let a, b be two labels for an intersection, where $b > a$. We want to show that one is odd and one is even. The number $b - a - 1$ is exactly the number of intersections between a and b. To prove that a and b have opposite parity, we have to show that this number is even.

To see this, let A be the closed curve that begins at a and ends at b, travelling in the orientation we selected at the beginning. Let B be the closed curve that begins at b and ends at a travelling in the same orientation. We no longer consider the vertex

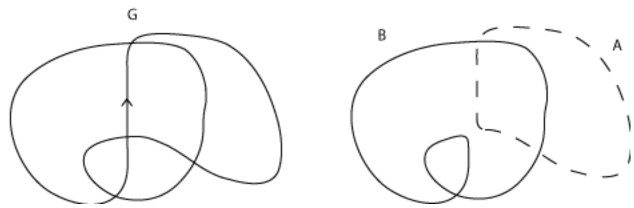

Figure 4: Decomposing G

corresponding to a and b as an intersection. Note that $A \cup B = G$ (See Figure 4). The intersections between a and b come from two sources: intersections of A with itself, and the intersections of A and B.

Since A is itself a closed graph, there is an even number of the first type of intersections. By Theorem 2, A and B intersect an even number of times. Thus, the total number of intersections between a and b is even.

Finally, to realize the graph as an alternating knot, make every even vertex in the labelling an over-crossing and every odd vertex an under-crossing. Since every intersection is labelled once with an even number and once with an odd number, this process is well defined. Since the labelling alternates between even and odd, the knot will alternate between overcrossings and undercrossings. □

The ideas in this proof can be extended to links as well:

Theorem 4. Every projection of a link with more than one component can be realized as the projection of an alternating link.

Proof. Let L be the projection of a link and let $L_1, L_2, ..., L_n$ be the simple closed curves that comprise it. If $n = 1$, we can use the theorem above, so assume $n \geq 2$. Now, there must be some region that has at least one edge from L_1 and at least one edge from another L_i. To see this, look at L_1's intersections with other simple closed curves. If any such intersections exist, then each of the four regions touching it has

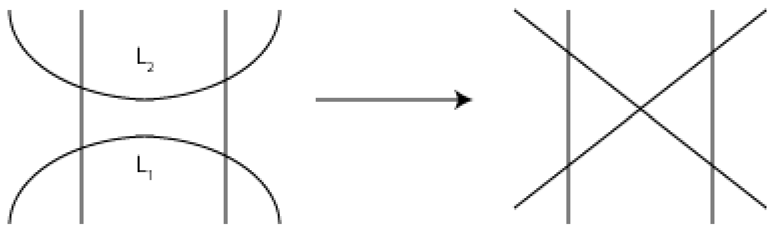

Figure 5: Combining two simple closed curves into one.

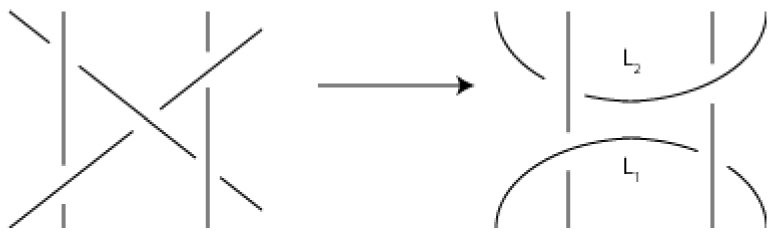

Figure 6: Slicing the curve apart again leaves two alternating loops.

at least one edge from L_1 and at least one from the L_i it intersects. If there are no intersections, then this is a split link, and so there is some region 'outside' L_1 and the other components that touches them both (using Jordan curve theorem). Now, in the proof above, we took a simple closed curve and split it into two loops. Reverse the process in this region, as shown in Figure 5. Now we have the projection of a link with one less simple closed curve. Continue this process until we have a single simple closed curve. Use Theorem 3 to make this closed graph into the projection of an alternating knot. Now reverse the splitting process, as shown in Figure 6. The two resulting loops are still alternating. If the intersection we removed was the only one in the whole loop, then we get two unknots, which we will consider alternating. If there is more than one intersection, then the situation as depicted in Figure 6 occurs; there may be more vertices in the region between L_1 and L_2 than depicted, or some of the intersections may be combined (for instance, if the original intersection can be removed by a Reidemeister I move), but the figure shows the essential situation.

Repeating this process at each of these intersections will give us an alternating link whose projection is the same as the link we started with. □

We will often talk about prime knots. A knot is said to be **prime** if there is no two-sphere in S^3 that intersects the knot exactly twice, considering all possible positions and deformations of the knot. A link is said to be **split** if there is a two-sphere containing one component of the link and not intersecting it at all, again considering all possible positions and deformations. It can be shown that, given an alternating diagram of the link, the link is split if and only if the diagram is split (i.e. not connected), and if it is non-split then it is prime if and only if the diagram is prime (i.e. no two regions share more than one edge) [15]. In this thesis, whenever I say a link is prime, I will also assume it is non-split.

3 The Complement of a Link

The goal of this section is to find a nice description of the complement in S^3 of a link. We are especially interested in the universal cover of the link complement, so we want to divide it into simply connected chunks. There are a multitude of ways to do this. The method we follow was originally described by Thurston [19] and then expanded by Menasco [14]. To begin, let's consider a link as sitting in a 3-ball, with the link generally lying in the xy-plane, making allowances for the intersections (see Figure 7).

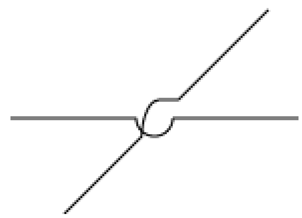

Figure 7: The link lying in the xy-plane, except for intersections.

We want to slice the complement of the link more or less along the xy-plane. Now, for unknotted curves which lie completely flat, this is all we need to do. We end up with two 'chunks', each having a copy of the curve deleted from its surface (see Figure 8).

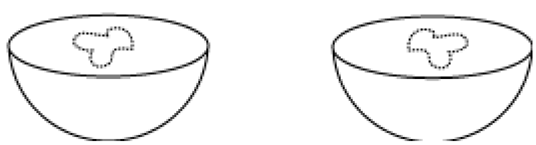

Figure 8: The two parts fit together to form the link complement

But trying to do the same trick for knotted curves runs into difficulty at the intersections. The part of the 3-ball above the intersection should go into the upper chunk, and the part below the intersection goes into the lower chunk, but what about

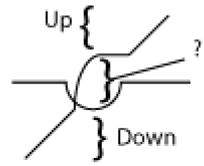

Figure 9: We want to slice the link complement in half, like in Figure 8, but what half does the middle part belong to?

the part in the middle? See Figure 9. Recall that the universal cover is by definition simply connected, which means that the link cannot tunnel through the chunks of its complement. We must find a way to split it so that the link itself is not hooked onto either chunk of the link complement.

To resolve this difficulty, we have to slice in a more complicated manner. We cut the area around the intersection into four quarters, as shown in Figure 10. Two of the quarters go up with the upper chunk, and two of the quarters go down with the lower chunk. The edges smooth out as we go away from the intersection. However, our chunks are not yet simply connected, because there is a hole in them. So we split each chunk along the line connecting its two quadrants (see Figure 11). Finally, we smooth everything flat, as shown in Figure 12, we get a very simple picture of the complement of the intersection: each half is just a 3-ball with a projection of the intersection drawn on it.

Connecting all the intersections, we get two 3-balls (one for the top and one for the bottom) which have a projection of the link drawn on the face (see Figure 13).

These processes can be hard to understand. I ended up constructing a model in real life from Play-Doh. Figure 14 is the Play-Doh version of Figure 10, and Figure 15 is the Play-doh version of Figure 11.

Now, the black lines in Figures 11 and 12 are just guides on how to glue the two chunks together, and the dotted lines show where we removed the link. Since these

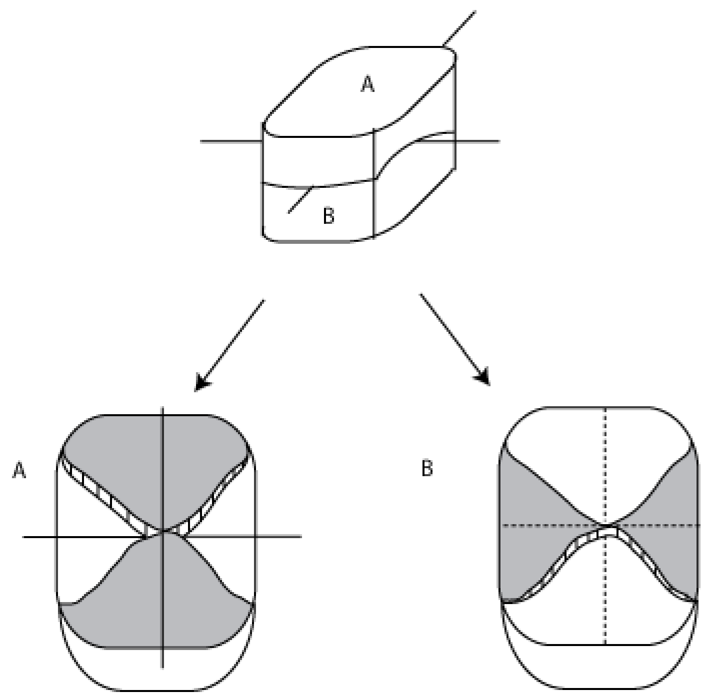

Figure 10: On the top is the complement of an intersection of a knot together with the intersection itself. We cut it into two halves, A and B. The black lines on A are the knot, and the dotted lines on B show where the knot touches it. If we flip A over to the right or the left and put it on B, we get the complement of the intersection back.

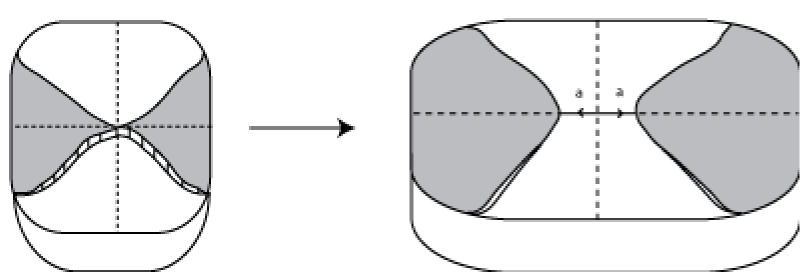

Figure 11: We split the bottom chunk along the vertical line where its two quadrants touch. The dark lines with arrows show how to glue the the surface of the solid to get the left side back. The dotted lines show where the link used to be.

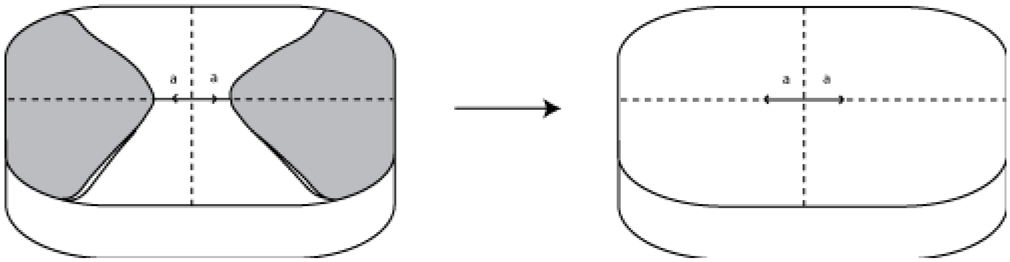

Figure 12: The ridges don't really matter, so we smooth them out. Remember, beginning in Figure 10, we've been finding the complement of a single intersection. This figure shows that that complement is made of two chunks, a top one and a bottom, both of which have a diagram of the intersection drawn on them.

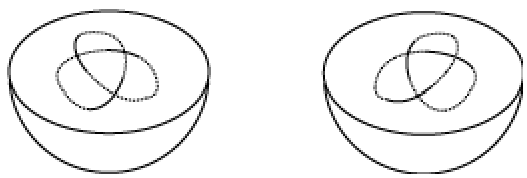

Figure 13: This is the complement of the trefoil knot.

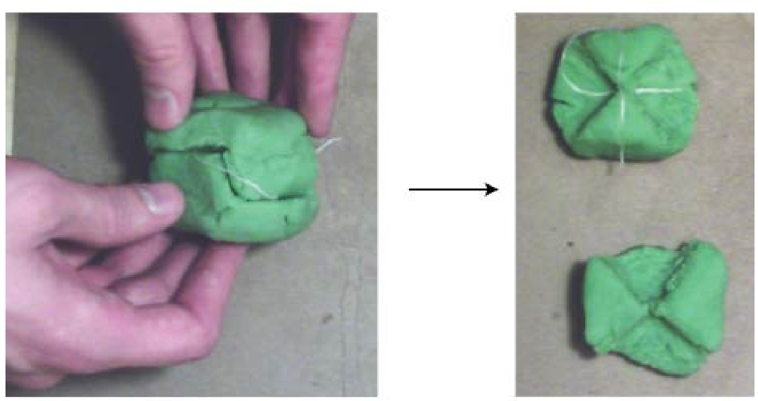

Figure 14: The string is the knot and the Play-Doh is the complement.

Figure 15: This is the Play-Doh version of Figure 11

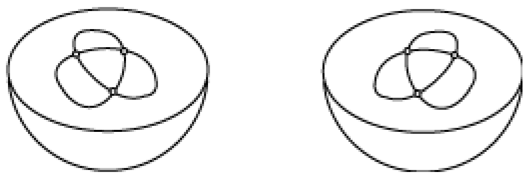

Figure 16: The 'open' parts of the complement have been shrunk to finitely many points.

dotted arcs are contractible, we can shrink them down to a point without changing the space. In the end, we get two 3-balls with only a finite number of points missing; the three balls are identified by faces and edges to get the complement of the link in the original 3-ball (see Figure 16). By identifying the outside faces, we get S^3.

3.1 Properties of the Link Complement

Let's examine some of the properties of the decomposition. First of all, it is difficult to continue to draw the chunks in three dimensions, so we squish the chunks flat to get a two dimensional image of the surface. We look at the top chunk from a 'top-down' view, so that the lines of identification of the two chunks look the same when drawn in 2-d (see Figure 17).

Notice that at each intersection, one chunk of the link complement has two arrows with the same label pointing *into* the intersection, and the other chunk has two arrows

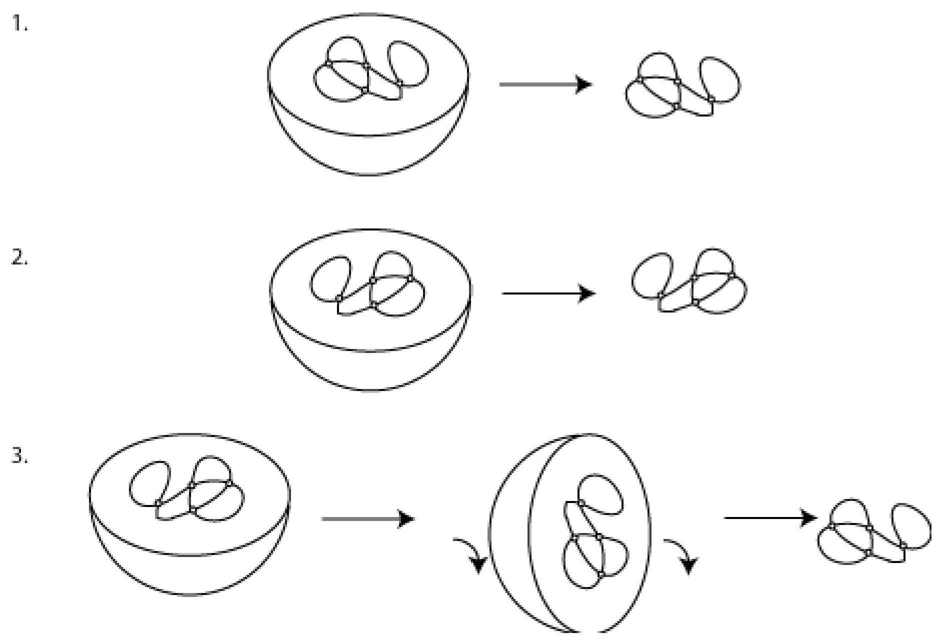

Figure 17: (1) We always project the bottom half of the link complement in this way. (2) This is the natural way of projecting the link complement, but it reverses the link. (3) By projecting from the top down, we get the two projections to look like each other. We will have occasions to use both ways of projecting the top chunk of the link complement.

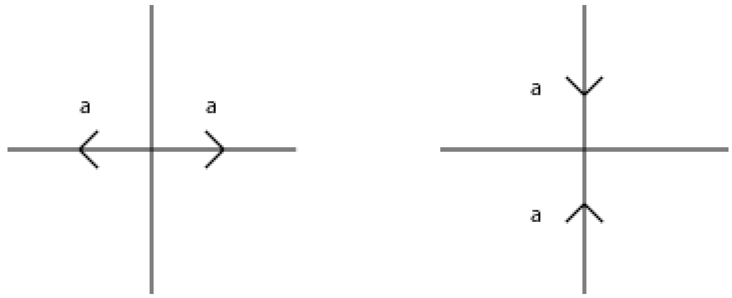

Figure 18: The left is the bottom chunk of the link complement around an intersection, and the right is the top chunk. Compare this figure with Figure 10

with the same label pointing *out* (see Figure 18).

This is because when we fold the chunks along those lines and glue them together, we want two 'tunnels' for the link to run through, one going over the other. The places where the link runs are the two intersection points, one in the top chunk of the link complement and one in the bottom. If they pointed the same way, then we would get only one tunnel. In my work, I have made the convention of having the lower chunk always point out. Also, the pair of edges pointing out in the bottom will be perpendicular to the corresponding pair of edges pointing in in the top; this can be seen by careful examination of the decomposition described earlier.

Also notice that the interior of each face is identified with one other face, and that each edge is identified with three others (one for each quadrant). So if letters are used to label the edges, each letter is used exactly four times. Throughout this text, I usually refer to the faces of the chunks of the link complement as **regions**. Regions touch in two main ways. Either they share an edge, or they **cross-touch**, by which I mean that they touch each other kitty-corner. On the standard xy-plane, the first quadrant and the third quadrant cross-touch, while the first quadrant and the second do not.

These link complements have a rich structure, much of which remains undiscovered. One property that is particularly useful is that every region is oriented. An arrow always points away from an overcrossing. If we follow the arrows on the edges around a region in an alternating link, they will always go head to tail, because the link alternates between overcrossings and undercrossings. Thus, we can think of the boundary of the region as an oriented simple closed curve. It will be either positively or negatively oriented, i.e. clockwise or counterclockwise. This phenomenon is very useful. One property of these orientations comes up so frequently that I feel it deserves its own name:

The Rotation Principle. If two regions share an edge, they have opposite orientations. If two regions cross-touch, they have the same orientation.

Proof. As seen in the construction of the universal cover of an alternating knot, the edges of a chunk of the knot complement can be given orientations in such a way that every region of the knot is 'clockwise' or 'counterclockwise'. Now, if two regions share an edge, then by examining the arrows at one of the end vertices, we can see that they must have opposite orientations. If they cross-touch, then by examining the arrows at that vertex, we can see they must have the same orientation.

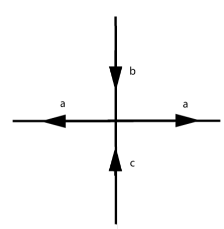

□

Corollary 1. No two regions of an alternating knot can both share an edge and cross-touch. If two regions each share an edge with a third region, they cannot share an edge with each other.

4 The Universal Cover and Replacement Rules

Now that we have a description of the link complement, we can build its universal cover. By investigating the combinatorial properties of the universal cover, we can learn more about the fundamental groups of the link complement. We construct the universal cover by gluing on copies of the fundamental domain one layer at a time; essentially, we are picking generators and looking at group elements a specific fixed 'distance' from the identity (i.e. elements of word length one, word length two, and so on). Graphically, we show this by drawing the outer surface of our universal cover as it grows. This enables us to work in two dimensions instead of three, which is easier to draw. Also, we project the polyhedra onto the plane, making one face the 'outside' of the figure. The basic idea is as follows:

The fundamental domain of the universal cover is split into two chunks, an upper chunk and a lower chunk. Draw a labelled projection of the lower chunk as described earlier in Section 3. If you are doing this by hand, it helps to draw out labelled projections of both chunks on a separate piece of paper to use as a reference (see Figure 19). Note that every region in the lower chunk has a corresponding region in the upper chunk (i.e. the edges are the same).

Pick any region. The idea is to take the upper chunk and 'rotate' it until the region you desire is on the outside, and then glue it on to the corresponding lower region. In practice, we pick a region in the lower chunk, say region 'ab'. Refer to Figure 20 as we go along. Take your labelled projection of the upper chunk and find the corresponding region. It will generally be somewhere in the middle. Pick an edge of the region, for example, edge 'a'. Look at the region sharing the edge 'a' with our initial region. In this case, it is the triangle 'acb'. So draw a triangle 'acb' inside our initial region. Now look at edge 'b' of 'acb'. It touches the region 'bc', so draw an

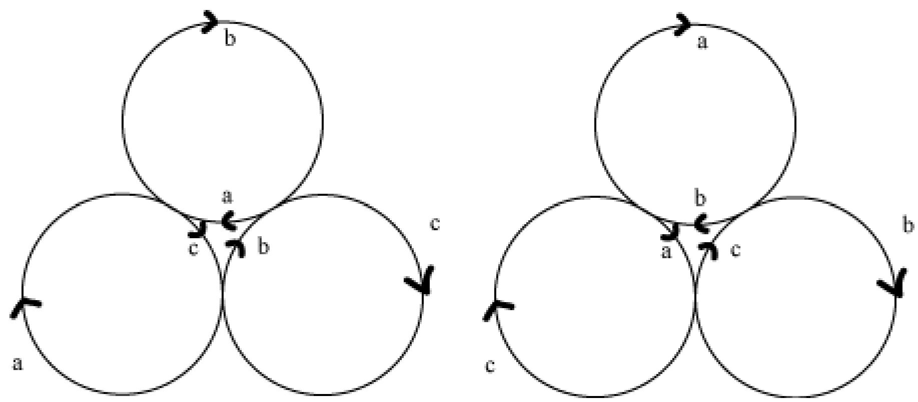

Figure 19: The left side is the upper chunk and the right side is the lower chunk. Here we're seeing them both from a 'top down' view.

edge 'c' to close off that region. Do the same with the region 'ac'. Check to see if there are any edges that have not been drawn. In this case, there are none, and we are done.

Repeat this process for every one of the initial regions. When we are done, all the open faces belong to the upper chunk of the link complement (see Figures 21 and 22). Note that even the outside 'face' has a chunk glued to it.

Now we repeat the entire process by gluing on lower chunks of the link complement, but with one fundamental difference: loaded edges. As you recall from our description of the link complement, every edge is identified exactly four times. By this stage of our process, every original edge has been identified three times. These edges are termed **loaded**. Loaded edges are shown with dotted lines. If we were to glue one copy of the lower chunk of the knot complement onto every open face, the edges would end up being identified five times, which is undesirable. So for every pair of regions sharing a loaded edge, we glue a single copy of the lower chunk of the knot complement onto both. This is done very much as we did before. Consulting Figure 22, lets add a lower chunk of the knot complement to the loaded pair 'acb' and 'ac' at the very top. Note

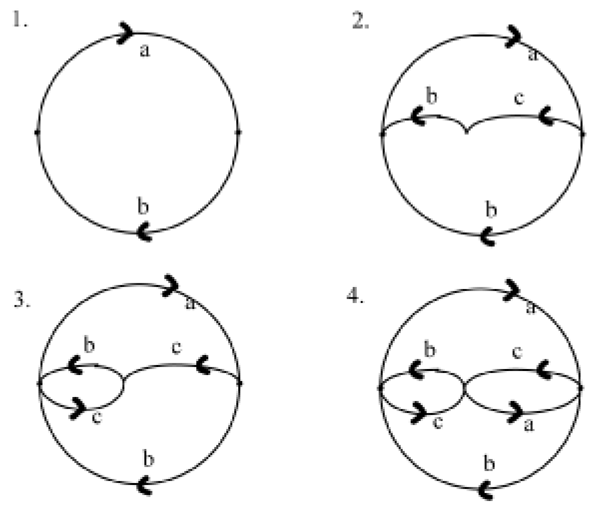

Figure 20: We begin in (1) with our initial region 'ab'. In (2), we add the triangle 'acb'. In (3),we add 'bc', and in (4) we finish by adding 'ac'.

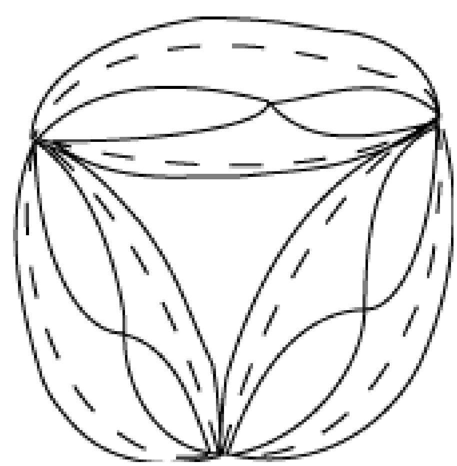

Figure 21: This is what the universal cover of the trefoil looks like after one stage of construction. The dotted lines correspond to the original lower chunk of the knot complement, and the black lines correspond to the various upper chunks of the knot complement. In the terminology of this thesis, the dotted lines are 'loaded'. Edge labels have been suppressed for clarity.

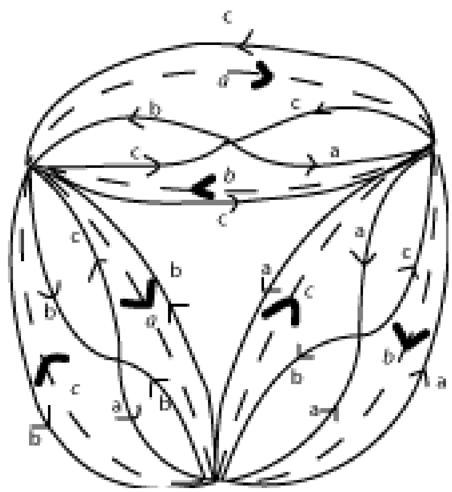

Figure 22: The same figure with edge labels applied.

that they touch over the edge ’a’. Referring to Figure 19, we see that in the lower chunk of the knot complement, the regions ’acb’ and ’ac’ do in fact touch each other over the line ’a’. Conceptually, what we are doing is rotating the lower chunk until **both** faces are on the bottom, and gluing the same chunk onto both. In practice, we draw regions line by line, just as in the non-loaded case (see Figure 23). Note that the loaded edge ’a’ is not drawn in the final picture; it is covered up by the new lower chunk. Let’s begin by looking at edge ’b’ of ’acb’. In Figure 19, we can see that ’acb’ touches the region ’ab’ over the edge ’b’, so we draw an ’a’ edge to create an ’ab’ in our picture. Looking at the edge ’c’ of ’acb’, we see that it touches ’bc’, so lets add an edge ’b’ to our diagram. Looking now at edge ’c’ of ’ac’, we see that it should touch region ’abc’. But we already have a region ’abc’ in the correct place in the picture. Looking over the picture, we see that every region is in the correct spot, and we are done.

It may seem that this procedure is rather ad hoc; we could have started looking at the edges of ’ac’ instead of ’acb’, for instance. However, this process will always

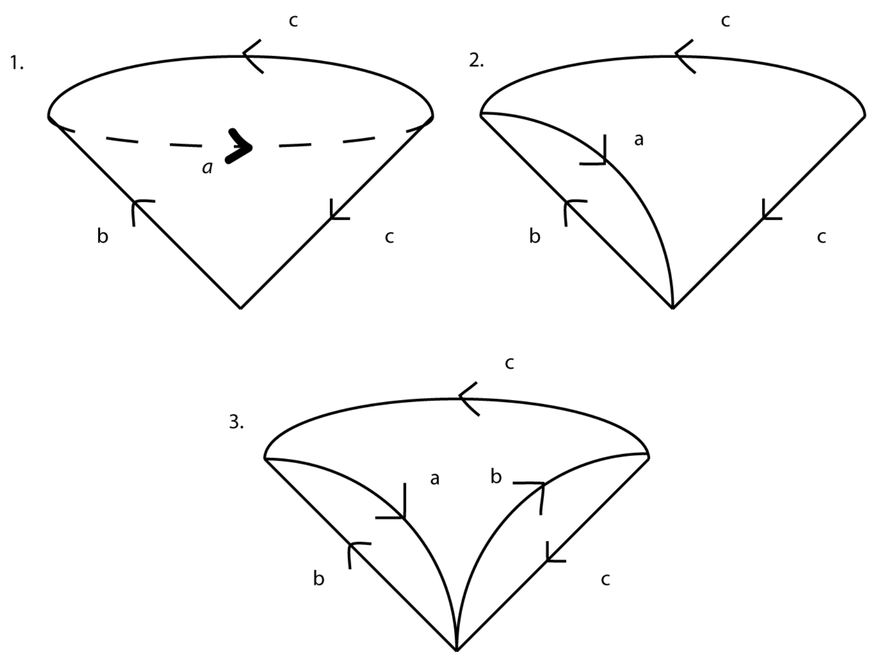

Figure 23: We start in (1) with the loaded pair 'ac' and 'acb'. The edge 'a' is covered up by the lower chunk of the knot complement, so it is not drawn in (2) or (3). In (2), we add the line 'a' to create the region 'ab', and in (3), we add the line 'b' to create the region 'bc'. This completes the drawing. Notice that we have indeed drawn the entire chunk of the knot complement; all 5 regions are present, with 'ac' and 'acb' on the bottom and 'ab', 'bc' and 'abc' on the top. Similarly, we have drawn all the edges, with one of the 'a's on the bottom and the rest visible. In the third stage of the construction, both 'c' edges and one of the 'b' edges will be loaded. All regions will be loaded, each with exactly one loaded edge.

give you a correct picture. It takes some practice to get it down. This difficulty was one motivation for the computer programs implementing Circlepak, as these help eliminate the feeling of ambiguity.

Exercise. Draw the first stage of the construction of the universal cover of the Borromean rings (see Section 12 for a definition). Pick one loaded pair and draw the second stage of the construction restricted to those two regions, just as we did above for the trefoil.

We prove a mildly interesting theorem about the universal covers:

Theorem 5. The number of non-loaded edges radiating from a given vertex increases by four at every stage of the construction of the universal cover.

Proof. Before we begin the proof, note that we only care about what's happening locally, at a single vertex. A vertex may border a loaded pair, but be far from the loaded edge. Locally, it's like the region is not loaded at all. But if the loaded edge has that vertex as an endpoint, it affects the local picture. So in this proof, we will only count pairs of regions as loaded if the loaded edge touches the vertex.

Every non-loaded edge is bordered on both sides by either a non-loaded region or a loaded pair. Call the number of loaded pairs L, and the number of non-loaded regions N. Now, in a non-loaded region, when we glue on the chunk of the link complement we add two edges, since four come out of a vertex on a single chunk and two edges were already there. This makes the original lines loaded, and all together there are three regions touching the vertex where before there were one; the two on the outside are each half of a loaded pair (see Figure 24). When we glue onto a loaded pair, all we get is two halves of new loaded pairs. So the growth of regions at a vertex is a linear function S in two variables, with $S(N) = N + 2(\frac{L}{2}) = N + L$,

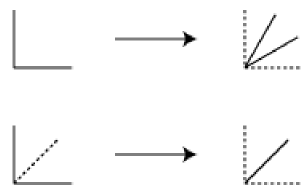

Figure 24: On the top is the behavior of a single non-loaded region at a vertex. At the bottom is the behavior of a loaded pair at a vertex. See also Figures 21 and 23.

$S(L) = 2(\frac{L}{2}) = L$. Thus, we get a matrix $\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$ We start with N=4 and L=0, since every fresh vertex has four non-loaded regions around it. Multiplying the nth power of S with this initial column vector gives the number of total non-loaded regions and loaded pairs surrounding the vertex; this is $S^n C = \begin{bmatrix} 1 & 0 \\ n & 1 \end{bmatrix} \begin{bmatrix} 4 \\ 0 \end{bmatrix} = \begin{bmatrix} 4 \\ 4\text{n} \end{bmatrix}$. Adding the two terms together gives us the total number of non-loaded edges, since each edge borders a different non-loaded region or a loaded pair on both of its sides, and each non-loaded region or loaded pair touches two different edges (we are still assuming the link is non-singular, so no edge goes back to a vertex it came from). Thus, the number of non-loaded edges coming from a vertex is $4n + 4$, where n is the number of additional layers we have added to our original fundamental domain. From this we see that the number of non-loaded regions increases by four at each stage. $\square$

5 Subdivision Rules

What we would like in the end is to find subdivision rules rather than replacement rules for the universal covers of various objects. The difference lies in the treatment of edges in the pictures we draw. In a replacement rule, such as our rule for the construction of the universal cover of the link complement falls into, edges disappear over time. Once they are identified four times, they are gone, covered by a chunk of the link complement. In a subdivision rule, edges never disappear. Regions and edges are dissected into smaller regions and edges, but old edges can always be recognized. These subdivision rules are in some sense more beautiful than replacement rules. They also arise naturally in other areas of mathematics.

Converting a replacement rule to a subdivision rule can be difficult. As we said above, lines are constantly being erased or covered up in replacement rules. To overcome this, we can nudge adjacent, fresh lines to replace the old ones. I have chosen to draw only the fresh, non-loaded edges each time. Every loaded edge is bordered by two edge subregions full of fresh edges. We take exactly one of these regions and move all of its fresh, non-loaded edges over to where they are lying exactly where the old loaded edge used to be. In general, this will add several vertices to the line, but that is acceptable and even standard in a subdivision rule. To give an example, look at the trefoil knot. The first stage in the replacement rule is shown in Figure 25 and the second stage is shown in Figure 26.

In this case, there is a particularly nice choice of edges to move, as shown in Figure 27. So we can depict the second stage of the universal cover as a subdivision of the first. What about the third? The fourth? In higher stages we have more to worry about, because we have loaded edges, and regions combine in new patterns. To obtain a real subdivision rule, we have to prove that it works in all stages of the construction of the universal cover. As an illustration, we derive a subdivision rule for all two-braid links, which includes the trefoil above. Recall, a two-braid link can be constructed by taking two strings, holding one end of each string together, twisting the strings together as many times as desired, then gluing the bottom strands to the top. They always have the form of a circle of bigons. They have a number of components equal to their number of vertices modulo 2.

Theorem 6. In each stage of constructing the universal cover of a prime, alternating link, a region can have at most one loaded edge.

Proof. This proof is inductive. Let R be a region with no loaded edges. Place a chunk of the link complement on it. The interior subregions are free from loaded edges, by definition. Each of the edge subregions touches the outside in at least one edge. By primeness of the link, an edge subregion cannot touch the outside in more than one edge. So the replacement rule creates only subregions with one loaded edge or none in this case.

Now let's pass to loaded regions. Since we place a single chunk of the link complement on two loaded regions, consider a pair of loaded regions at a time. Let R_1 and R_2 be two regions that share a loaded edge. Then place a chunk of the link complement on the two regions. In this case, both R_1 and R_2 are on the bottom. Assume a subregion L has two outside edges, i.e. two loaded edges. Then the edges

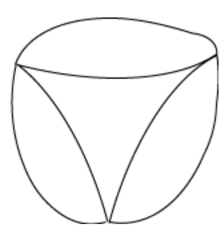

Figure 25: The first stage of the replacement rule: a single lower chunk of the link complement

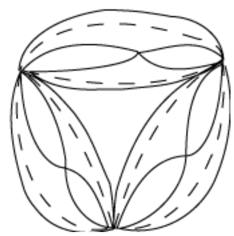

Figure 26: The second stage stage of the replacement rule: the open faces belong to five upper chunks of the link complement

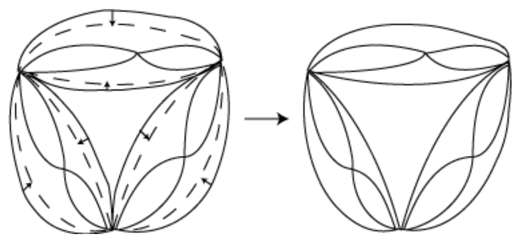

Figure 27: As shown by the arrows, the old edges are replaced by the new edges, so that the second stage of the universal cover is depicted as a subdivision of the first stage

it touches belong to R_1 or R_2. If they both come from a single region, for instance, R_1, we violate the primeness condition. If it touches one edge of each, we violate Corollary 1 of the rotation principle. Thus, even in the loaded case, only subregions containing one or zero loaded edges are created. □

Now we are ready to prove:

Theorem 7. Every two-braid has a subdivision rule.

Proof. Throughout this proof we'll use Figure 28 as our 'map'. There are only two types of regions in the first stage, the bigons and the p-gons. They have replacement rules that are as shown in Figures 29 and 30.

The two-braids have the nice property that every edge is bordered on one side by a bigon and on the other side by a p-gon. Since every loaded pair consists of

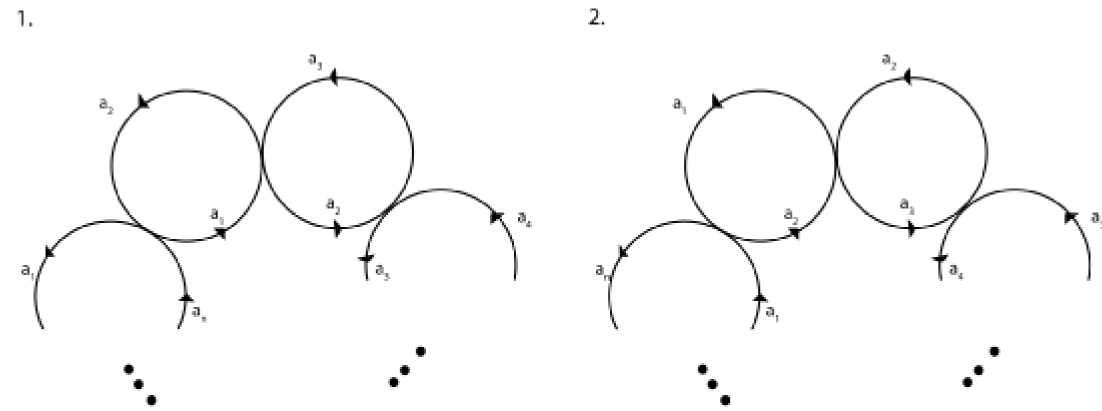

Figure 28: (1) The lower chunk of the complement of a two-braid. (2) The upper chunk.

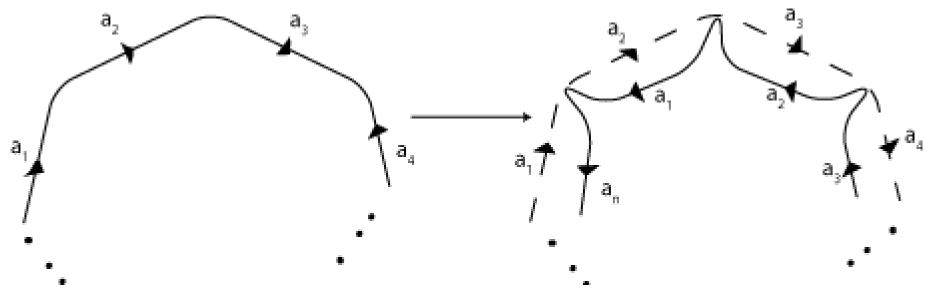

Figure 29: The replacement rule for a typical p-gon.

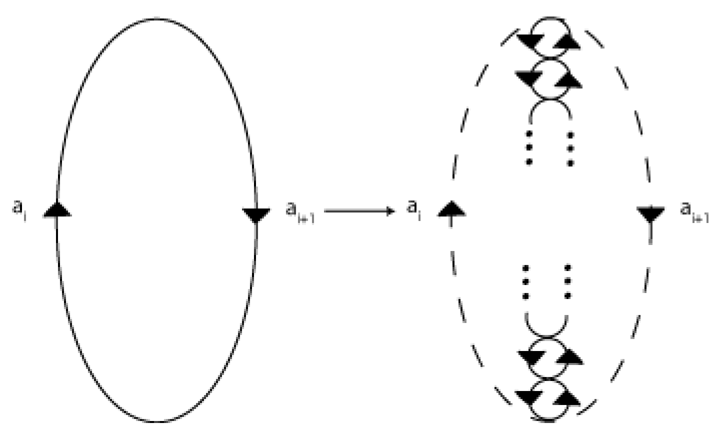

Figure 30: The replacement rule for a typical bigon. Interior edge labels have been suppressed for clarity

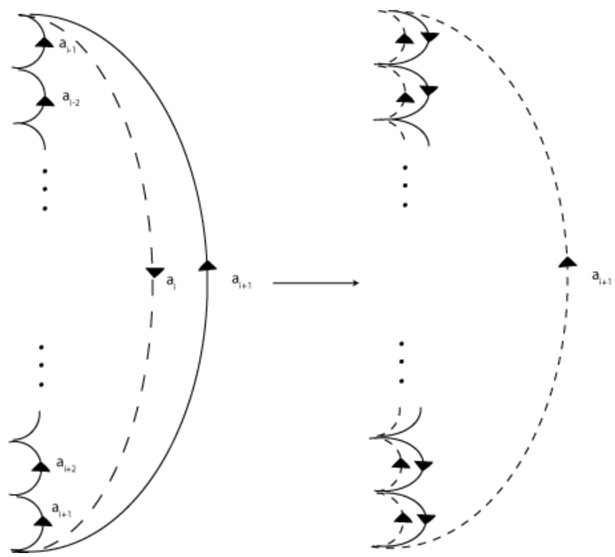

Figure 31: The replacement rule for a typical loaded pair. Interior edge labels have been suppressed for clarity

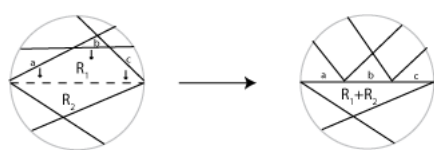

Figure 32: This picture illustrates how the lines of an edge subregion move to replace the loaded edge in a general link. We are zoomed in on one small region of the universal cover. Lines a,b, and c all move over in this example.

two regions sharing a single edge, the only loaded pairs we must consider consist of exactly one p-gon and one bigon. Their replacement rule is shown in Figure 31. Note that the pictures are essentially the same in the top chunk and the lower chunk; only the labels change. Now, to find a subdivision rule, we need to make an assignment of edges so that no matter how the regions come together in the future, exactly one of the two edge subregions touching a particular loaded edge gets assigned to move its edges over to become the new edge (see Figure 32). But this is easy; assign the unloaded lines of every edge subregion in the p-gons and in every subregion with lines that can move in the loaded pairs to replace the old edges. In the loaded pairs, all

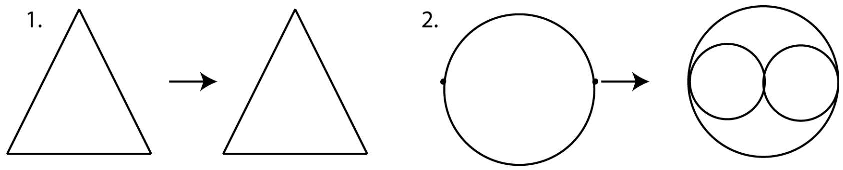

Figure 33: The subdivision rule for the trefoil complement. Note that (1) is the subdivision rule for both the unloaded triangles and the loaded pairs, even though their replacement rules were different.

but one of the old edges will be replaced by an interior edge; the only one that is not replaced is the old edge that was part of the bigon. In the bigon, no region's lines are moved; and in the p-gon, all edge subregions' lines are moved to replace the old lines.

To see that this is a well-defined subdivision rule, we must show that this rule assigns the lines of exactly one edge subregion on either side of a loaded edge to replace the loaded edge. We use the fact that every edge has a bigon on one side and a p-gon on another. Whether the bigon is part of a loaded pair or not, its edge subregion stays put, and whether the p-gon is part of a loaded pair or not, the lines of its edge subregion move to replace the old edge. So for every possible combination of regions meeting at an edge, exactly one edge subregion is absorbed, which shows that our subdivision rule is well-defined. □

The theorem above is illustrated in Figure 33 for the case $p = 3$ (i.e. the trefoil knot), and the first few subdivisions are given in Figure 34. This trefoil rule is a particularly beautiful example of a subdivision rule, and is the first of two subdivision rules for links (the trefoil and the figure-eight) that my advisor Jim Cannon discovered.

Not all knots allow for a choice of subdivision rules. Take for instance the picture in Figure 35 of a loaded pair in the figure eight knot. The region A is called 'frozen', since its edges cannot be moved over to become the new line without moving one of

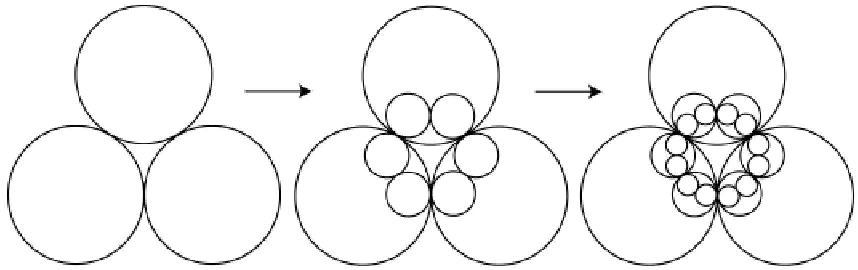

Figure 34: The first three layers of the universal cover of the trefoil complement

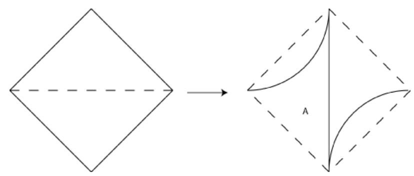

Figure 35: Region A on the right is 'frozen'. Since the original vertices remain fixed, its lines cannot move over to replace the edge.

the region's original vertices, which is undesirable. Luckily, for this particular knot, it turns out that the region adjoining A over the loaded edge is not frozen, and its edges can successfully be moved over to become the new edge. The subdivision rules for this knot will be given in Section 6.1

And finally, some links cannot be given a subdivision rule under this process. This occurs when both regions sharing an edge are 'frozen'. Although it is difficult to show, this occurs for the link in Figure 36 after five or so layers. This example is discussed in more detail in Section 6.3. However, this can be overcome by modifying the way we construct the link complement, as discussed in Section 11.

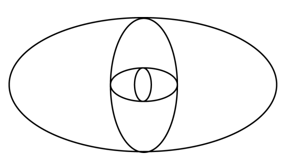

Figure 36: This link cannot be given a subdivision rule under our current method.

5.1 Application: Growth Functions

This section is not necessary for the rest of the thesis.

Subdivision rules can be used to quickly calculate the number of elements of a specific word length in the fundamental groups of the link complements, with respect to the presentation determined by our choice of fundamental domain. If we treat the number of tiles of each type present at each stage as a variable, subdivision gives a linear function on the vector space with these variables as a basis, similar to the proof of Theorem 5. The resulting matrix can be used to calculate the number of elements of a specific word length.

As an example, we take the trefoil knot. There are two tile types, triangles (T) and bigons (B). The subdivision function S satisfies $S(T) = T$ and $S(B) = 2T + 2B$, since triangles are not subdivided but bigons are split into two bigons and two triangles. This gives a matrix $S = \begin{bmatrix} 1 & 2 \\ 0 & 2 \end{bmatrix}$. It can be diagonalized as $S = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 & -2 \\ 0 & 1 \end{bmatrix}$. Now, to get the Cayley graph from the subdivision rule, we need to fix a base point; let's put it in the lower half of the link complement. Then if our initial, lower-chunk fundamental domain is the 0th stage of subdivision, only the even stages of subdivision contain chunks of the link complement that include the base-point (recall that we alternate between gluing on lower chunks and upper chunks of the link complement). We begin with an initial column vector C that gives the number and types of faces of the first fundamental domain. But the chunks we glue on to this are upper chunks of the link complement, so to get the size of the first layer of the Cayley graph, we have to look at SC, which gives the number of lower chunks of the link complement to glue on. In general, to obtain the number of elements of word length n of the Cayley graph, we look at odd powers of S. The trefoil begins with initial conditions $T = 2, B = 3$. So the number of triangles and bigons at the $2n-1$th stage of construction is $S^{2n-1}C = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 2^{2n-1} \end{bmatrix} \begin{bmatrix} 1 & -2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 1 & 2^{2n}-2 \\ 0 & 2^{2n-1} \end{bmatrix} \begin{bmatrix} 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 3(2^{2n})-4 \\ 3(2^{2n-1}) \end{bmatrix}$. Adding together, we see that the number of elements of word length n is $3(2^{2n}+2^{2n-1})-4 = 9(2^{2n-1})-4$. The formula fails for n=0, where it should be 1, since there is only one fundamental domain at that stage.

We now find the growth function, which is the power series with coefficients equal to the growth rate of the group. The first term is 1, and all other terms are given by the growth rate above. Thus, our power series is $1+\sum_{n=1}^{\infty}[9(2^{2n-1})-4]x^n = 1+9\sum_{n=1}^{\infty}(2^{2n-1})x^n - 4\sum_{n=1}^{\infty}x^n = \frac{1}{2}+\frac{9}{2}\sum_{n=0}^{\infty}(4x)^n - 4\sum_{n=0}^{\infty}x^n = \frac{1}{2}+\frac{9}{2(1-4x)}-\frac{4}{1-x} = \frac{(1-4x)(1-x)+9(1-x)-8(1-4x)}{2(1-4x)(1-x)} = \frac{4x^2+18x+2}{2(1-4x)(1-x)} = \frac{2x^2+9x+1}{(1-4x)(1-x)}$, and this is the growth function for this presentation of the trefoil complement's fundamental group.

This method can be generalized to any of the subdivision rules developed in this thesis. It needs some tweaking for the torus links in Section 10, but it works fine for the subdivisions with boundary in Section 11, which exist for every alternating link. In general, we get a square matrix of size equal to the number of tile types which we then put in Jordan canonical form, the powers of which are well-known and involve binomial coefficients. We obtain a formula for the number of elements of a certain word length, then obtain a series for it, then convert it into a rational function (adding binomial coefficients to a geometric series gives a rational function with repeated roots). Thus, we can find a growth function for every prime, nonsingular alternating link. These and other growth functions were extensively studied by Matthew Grayson, a student of William Thurston, in his 1983 Princeton thesis. His work incorporates several previous results by Cannon and Thurston that were unpublished prior to his writing.

6 Examples

6.1 The Figure Eight Knot

To illustrate some of the difficulties of finding a subdivision rule from a replacement rule, we will construct a subdivision rule for the figure eight knot.

Recall Theorem 6. This theorem is useful because it cuts down the number of tile types we need to consider. Essentially, if we can find a nice subdivision rule for all original regions and all paired regions, and if we can show that it is well-behaved for all possible placements of the tiles, then we will be done. Essentially, we are trying to replace every loaded edge by moving over the edges of one of the two edge subregions facing it. For brevity, we will abuse language and say that an edge subregion moves over or is absorbed or becomes the new line when really we mean that its edges move over.

The chunks of the knot complement are depicted in Figure 37. Note that this time the upper chunk is viewed from a bottom up view; its regions have the opposite orientation from the regions of the lower chunk. The reason I have drawn it this way is because when we are actually drawing the knots, the orientation on the regions reverses every time we glue on a new chunk. In real life, if we made a mold for the knot and cut it in half, it would look like this. I find this view best for application, and the top-down view best for theory.

All told, the figure-eight knot has 6 regions and 8 edges, giving a total of 14 tiles per chunk. However, it is rotationally symmetric, so we can cut that number down to 8 total.

The replacement rules for the different regions in the knot complement are shown in Figures 38 and 39 below. Since the figure eight knot is rotationally symmetric, the

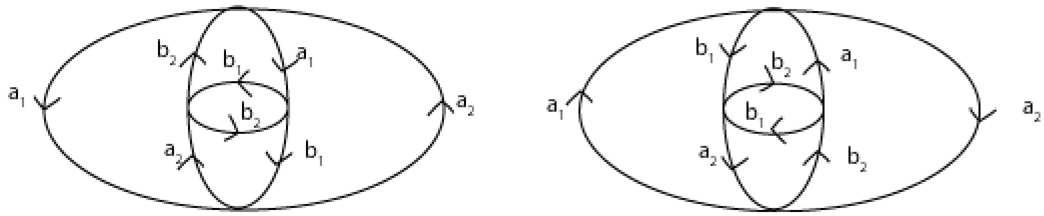

Figure 37: On the left is the bottom chunk of the knot complement, and on the right is the upper chunk, viewed from the bottom.

rules are given with labels mod 2. Note that the region a_1a_2 in both chunks appears to have an opposite orientation from the one in Figure 37. This is because it is the outside region, and while it is drawn counterclockwise in Figure 37, it is negatively oriented with respect to points in the exterior of the knot projection.

Our goal is to find a way to associate an edge subregion to each loaded edge that might possibly appear in the construction of the universal cover such that

1. Each edge is assigned exactly one edge subregion to replace it, and
2. No vertices are assigned to incompatible directions.

For instance, look at the edge a_1 in the lower half of the knot complement that goes between a_1a_2 and $a_1a_2b_2$. What are the possible regions, loaded or not, that can touch over this edge? One side will be a_1a_2 or a loaded version of it and the other side will be $a_1a_2b_2$ or a loaded version of it. The possibilities are a_1a_2 or $a_1a_2 + a_2a_1b_1$ on the left and $a_1a_2b_2$, $a_1a_2b_2+b_2a_1b_1$ or $a_1a_2b_2+a_2a_1$ on the right (why isn't $a_1a_2+a_1a_2b_2$ a possibility?). If you notice, in Figure 38, the edge subregion of the regions in the first group that touches a_1 is always $a_1b_1a_2$, and the second group always touch a_1 in the edge subregion $a_1b_1b_2$. Note that in $a_1a_2b_2+b_2a_1b_1$, the edge subregion in question (i.e. $a_1b_1b_2$) is *frozen*; it cross-touches the outside regions, and since those vertices remain fixed, it cannot move. Therefore, we choose that the other edge subregion

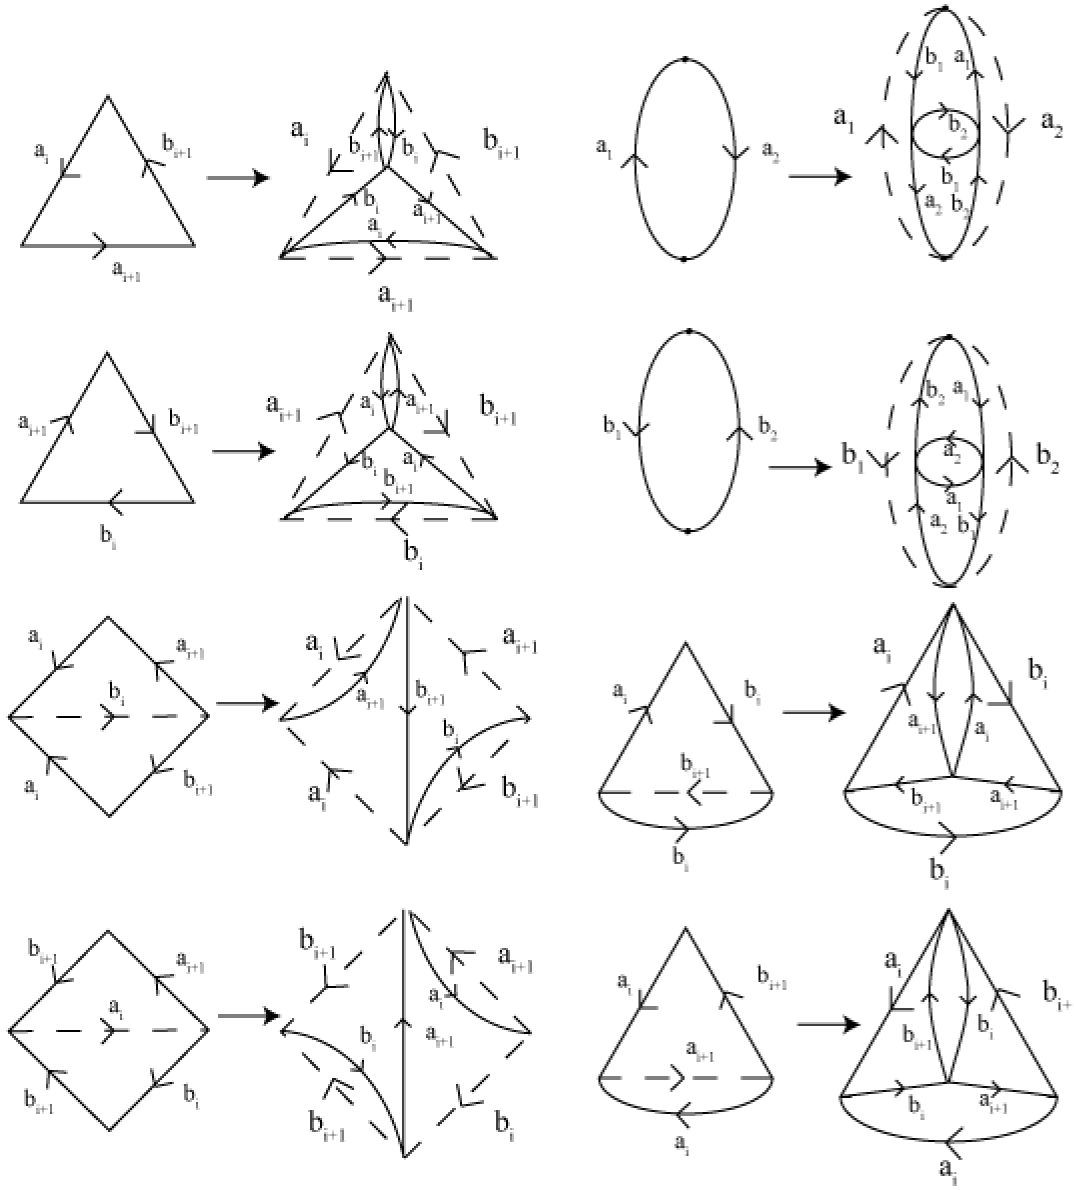

Figure 38: The replacement rules for the lower half of the figure eight knot. The subscripts are to be taken mod 2. The dotted lines represent loaded edges.

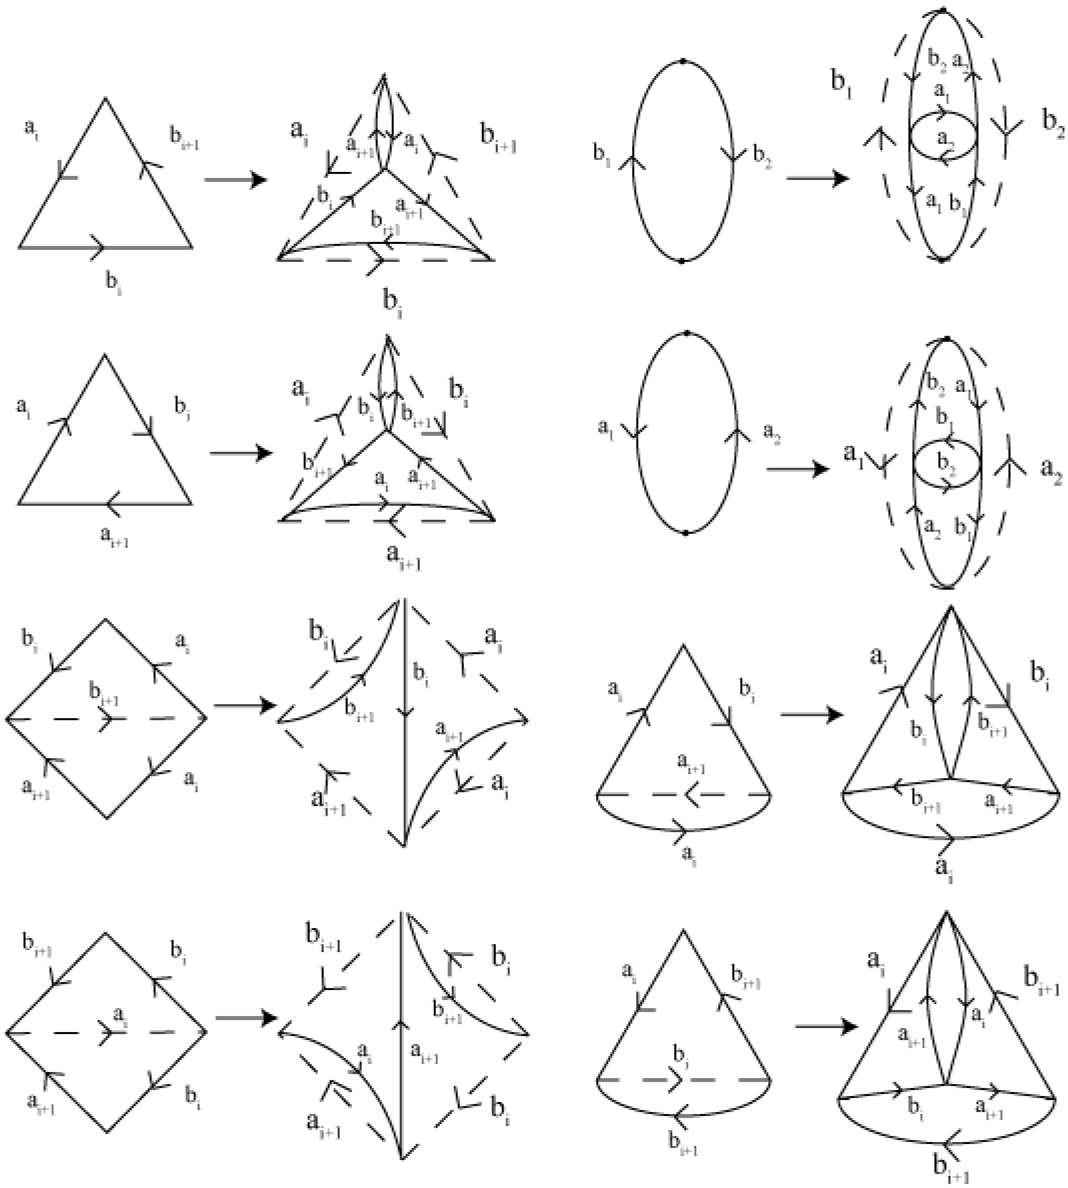

Figure 39: The replacement rules for the upper half of the figure eight knot. The subscripts are to be taken mod 2. The dotted lines represent loaded edges.

of the a_1a_2 group, $a_1b_1a_2$, always moves over. Notice that in $a_1a_2 + a_2a_1b_1$, all three edge subregions are tied together, since they cross-touch and/or share edges; now that we've moved one over, the other two are now essentially frozen as well, because that common vertex is committed. This seriously restricts our further options.

Through analysis like this, one can assign an edge subregion to move over for every edge. The other edge subregions that touch both a bigon and a triangle can be assigned the same way that we did above. Those edges that touch triangles on either sides will have a bigon as one of its edge subregions, and bigons can always be moved over without affecting other regions. Proceeding in this method, and collapsing regions that behave identically, we get the subdivision rules in Figure 40 (for both the upper and lower chunks). Thus, we have our subdivision rule. As you can see, finding the subdivision rules for even a very small knot can be a painstaking and at times frustrating process. Our theorems in later chapters will give us easy, programmable methods of finding subdivision rules for all alternating links.

6.2 N-Chains

One particularly nice set of links that will not fall into later categories is the collection n-chains. These are alternating links with n components such that each component intersects exactly two others and each pair of components that intersect do so in exactly two points. They look like chains that have their first and last links welded together (see Figure 41). The smallest case is the two-chain, with two links that intersect each other four times. But this is just a two-braid with four twists, and we already have its subdivision rule.

We will construct the subdivision rule for these links similar to the way we constructed it for the figure eight link. We'll use the 4-link as our example, but it

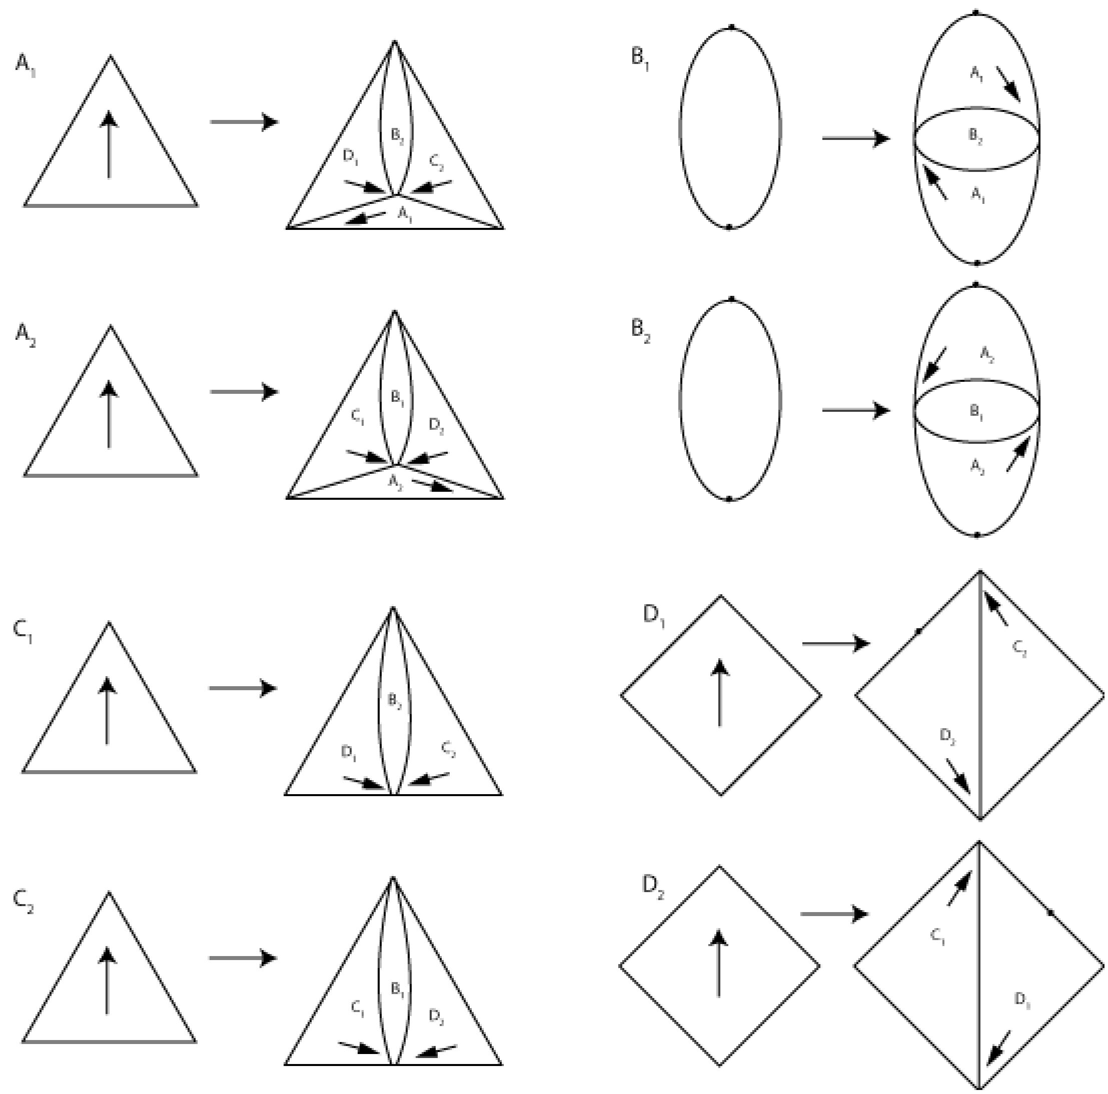

Figure 40: The tile A_1 corresponds to the tile $a_1a_2b_2$ and $a_2a_1b_1$ in both the top and bottom chunks. The tile A_2 corresponds to the tile $a_1b_1b_2$ and $a_2b_2b_1$ in both the top and bottom chunks. The tiles B_1 and B_2 correspond to a_1a_2 and b_1b_2, respectively, in both the top and the bottom. The tiles C_1 and C_2 correspond to $a_1a_2b_2 + a_1a_2$ and $a_1b_1b_2 + b_1b_2$, respectively, in both the top and the bottom. D_1 and D_2 are different for the different chunks. D_1 corresponds to $a_1a_2b_2 + a_1b_1b_2$ and $a_2a_1b_1 + a_2b_2b_1$ in the bottom chunk and $a_2b_2b_1 + a_2b_2a_1$ and $a_1b_1b_2 + a_1b_1a_2$ in the top chunk. D_2 is the same but reversed.

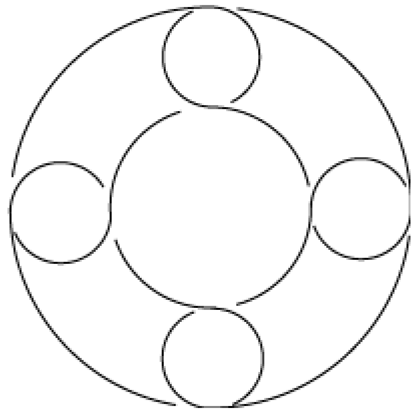

Figure 41: The 4-link

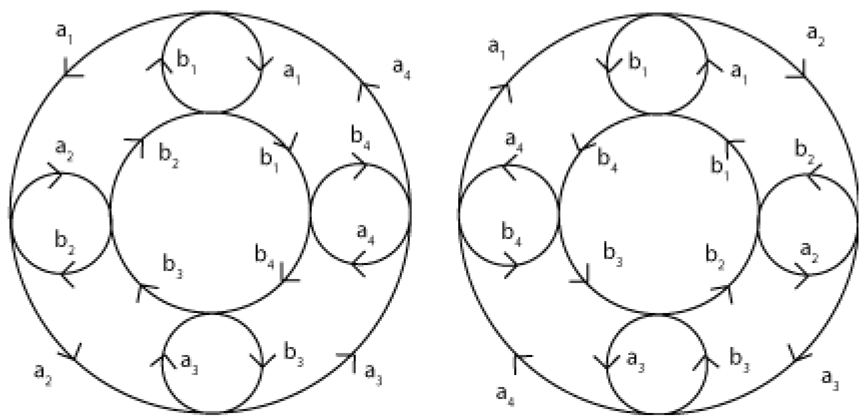

Figure 42: On the left is the lower chunk, and on the right is the upper chunk from the bottom-up view.

generalizes immediately to any n greater than one. Our map is given in Figure 42.

The replacement rules are shown in Figure 43. This time we have omitted the labels. It is a good exercise to check that these diagrams are correct and to provide labels. Notice that the replacement rules are the same for the top and bottom chunks of the link complement. This is not the case for a general knot but happens for every n-chain.

The changes to get from replacement rule to subdivision rule are not particularly difficult in this case (see Figure 44). All of the lines of edge subregions in B tiles go out

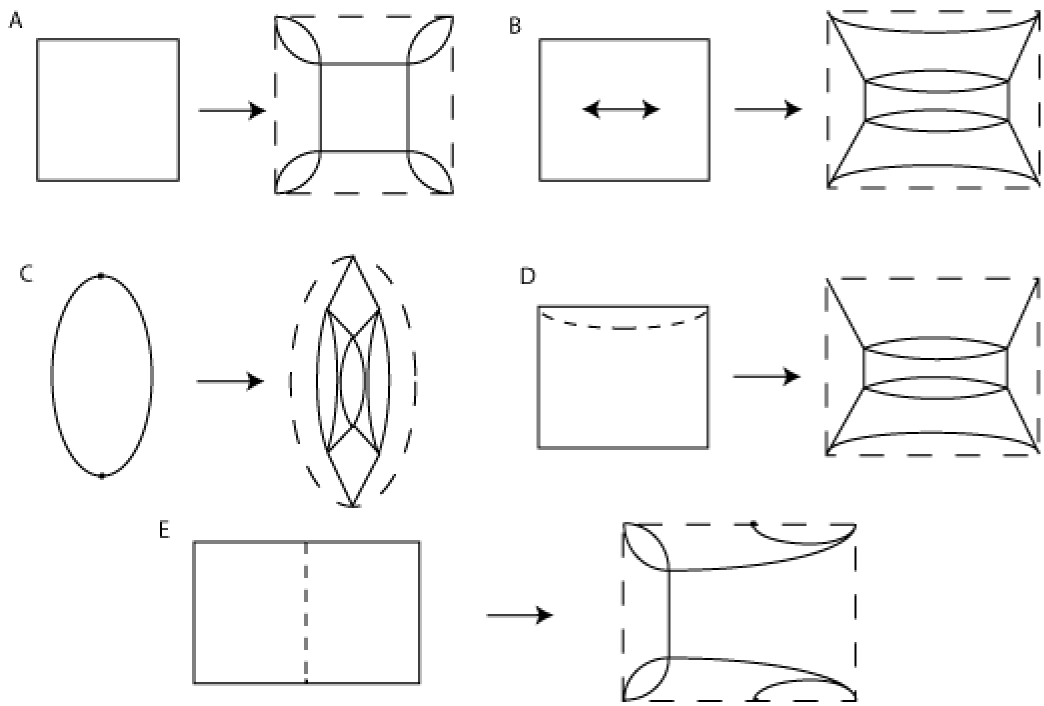

Figure 43: The replacement rules for both the top and bottom chunks of the link complement. Type A squares are those cross-touching bigons, and type B squares are those that share two edges with bigons. The arrows shown in square B point to the edges that it shares with bigons. In tile E, the square on the left is a type A and the square on the right is a type B.

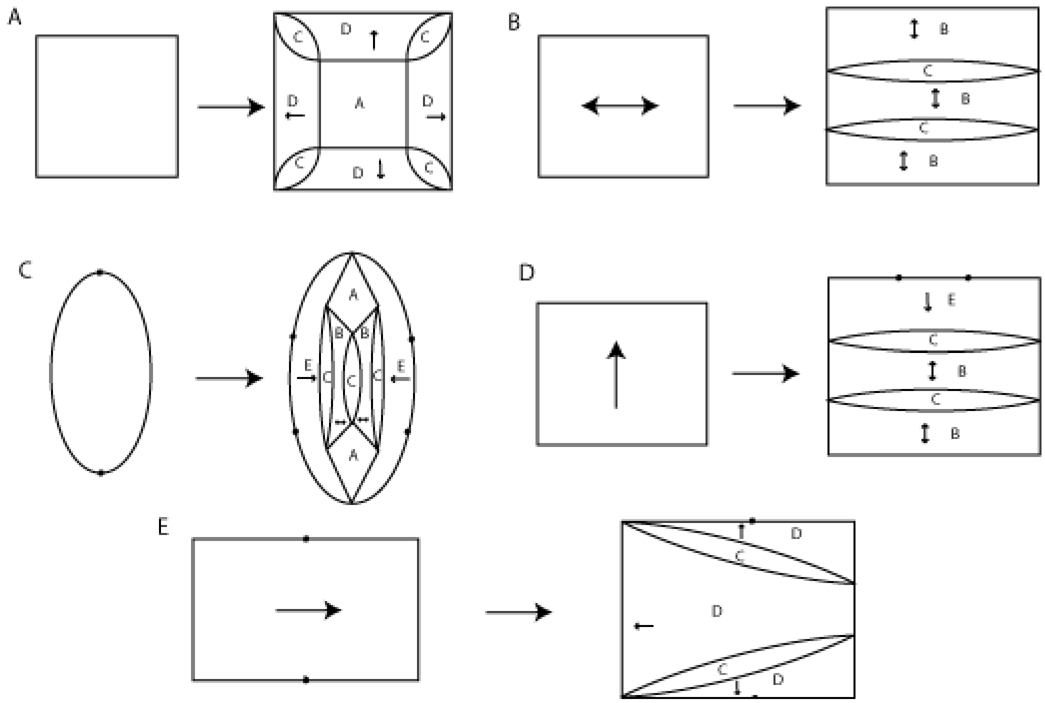

Figure 44: A subdivision rule for the 4-link

to replace the old edge, while all the lines of the edge subregions in A and C tiles stay put. The edge subregions of the loaded pairs behave in accordance with the original subregions; for example, the lines of the edge subregion belonging to the bigon in D stays put, and the lines of the other three belonging to the B-type region go out. Note that the A and C tiles have the same orientation (i.e. clockwise), and B has an orientation opposite to both A and C, so the edge subregions are absorbed or not depending entirely on their orientation. We will see in the sections on radial links, dispersed links, and subdivisions with boundary that this occurs often in creating subdivision rules from replacement rules. Note that in the figure eight knot, this principle was not and could not be followed.

Figure 45 shows the first few subdivisions given by this rule for a C tile. This is one of my favorite subdivision rules. The interested reader should try subdividing

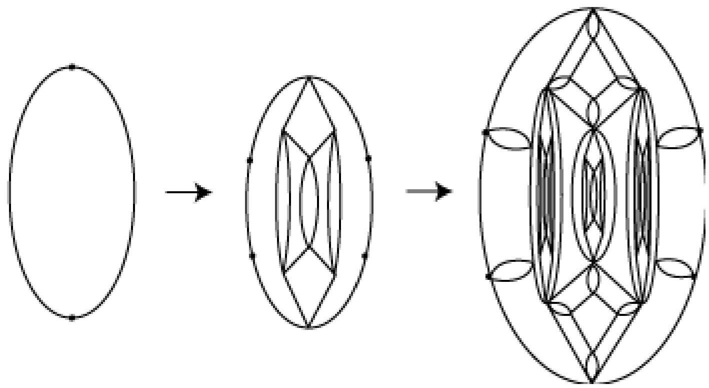

Figure 45: The first two subdivisions for a type C tile.

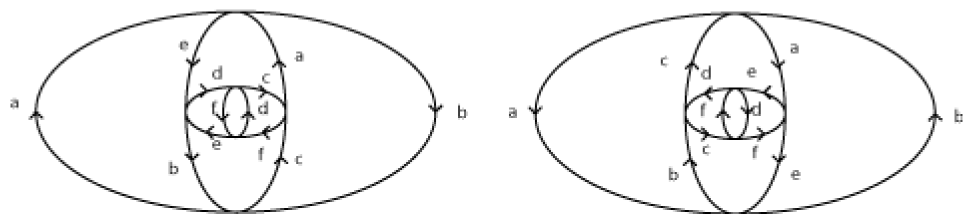

Figure 46: On the left is the bottom chunk of the link complement, and on the right is the top chunk looking from the bottom up.

other tiles and using other values of n.

6.3 A Counterexample

Finally, we give an example where no subdivision rule can be created from our replacement rule by moving subregions as we have done before. The link is pictured in Figure 46. To see that it cannot be given a subdivision rule, we must glue five layers of link complements on (see Figure 47). The reason we have to go this deep is because the loaded regions are the worst regions, and the more layers we have, the more loaded regions we have touching each other in more combinations. After five layers, we have three loaded regions touching each other, two of which have frozen edge subregions, pulling the edge subregions of the middle loaded region in two incompatible directions (see Figure 48). In Section 11, we find a subdivision rule for this link by an alternate path.

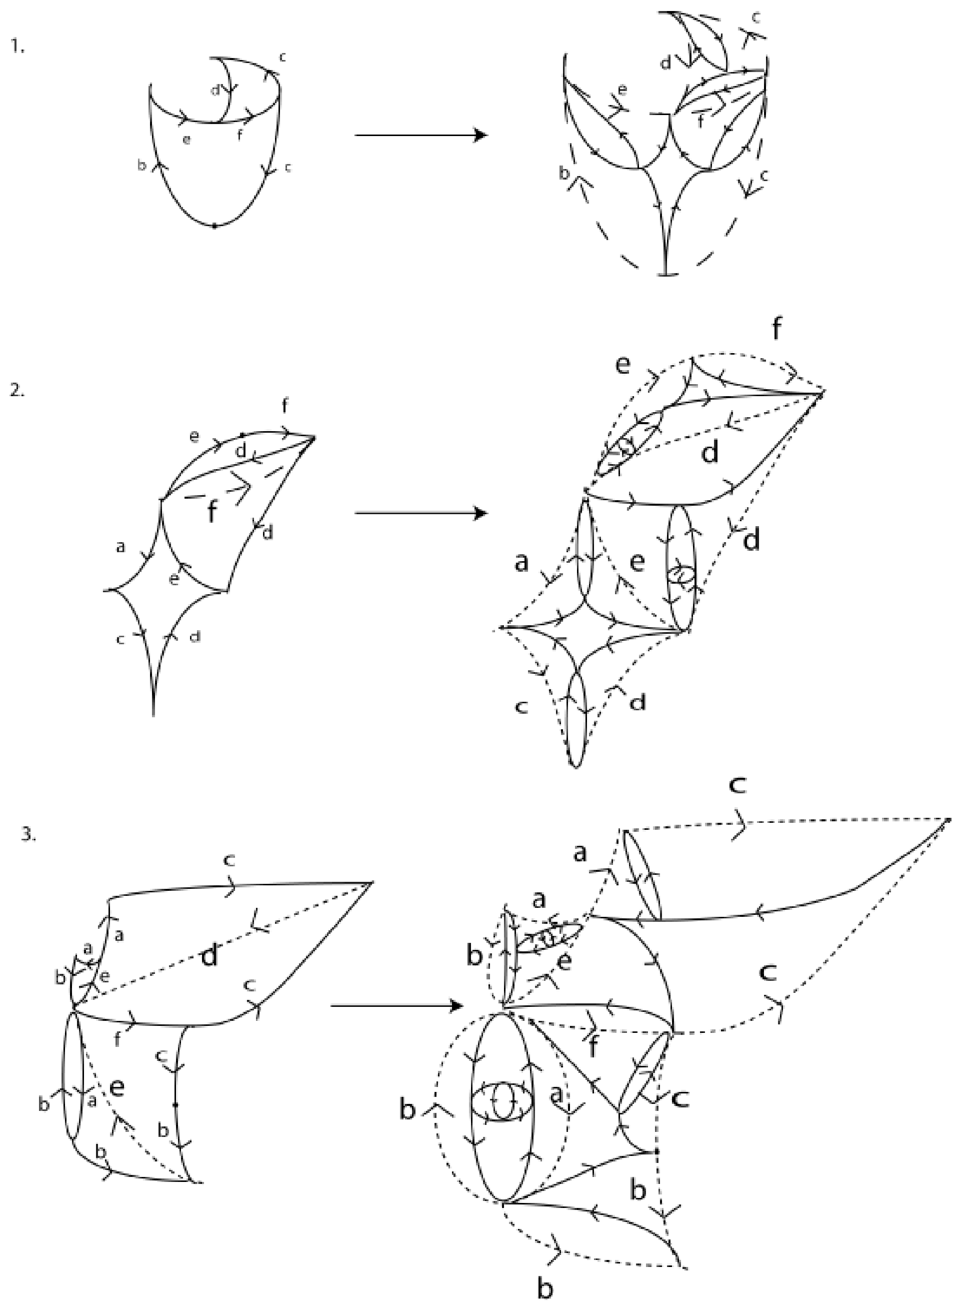

Figure 47: Each row begins with a zoom-in of the last figure of the previous row. The very first row begins with two neighboring regions in the original bottom half of the link complement. See Figure 48 for the last layer.

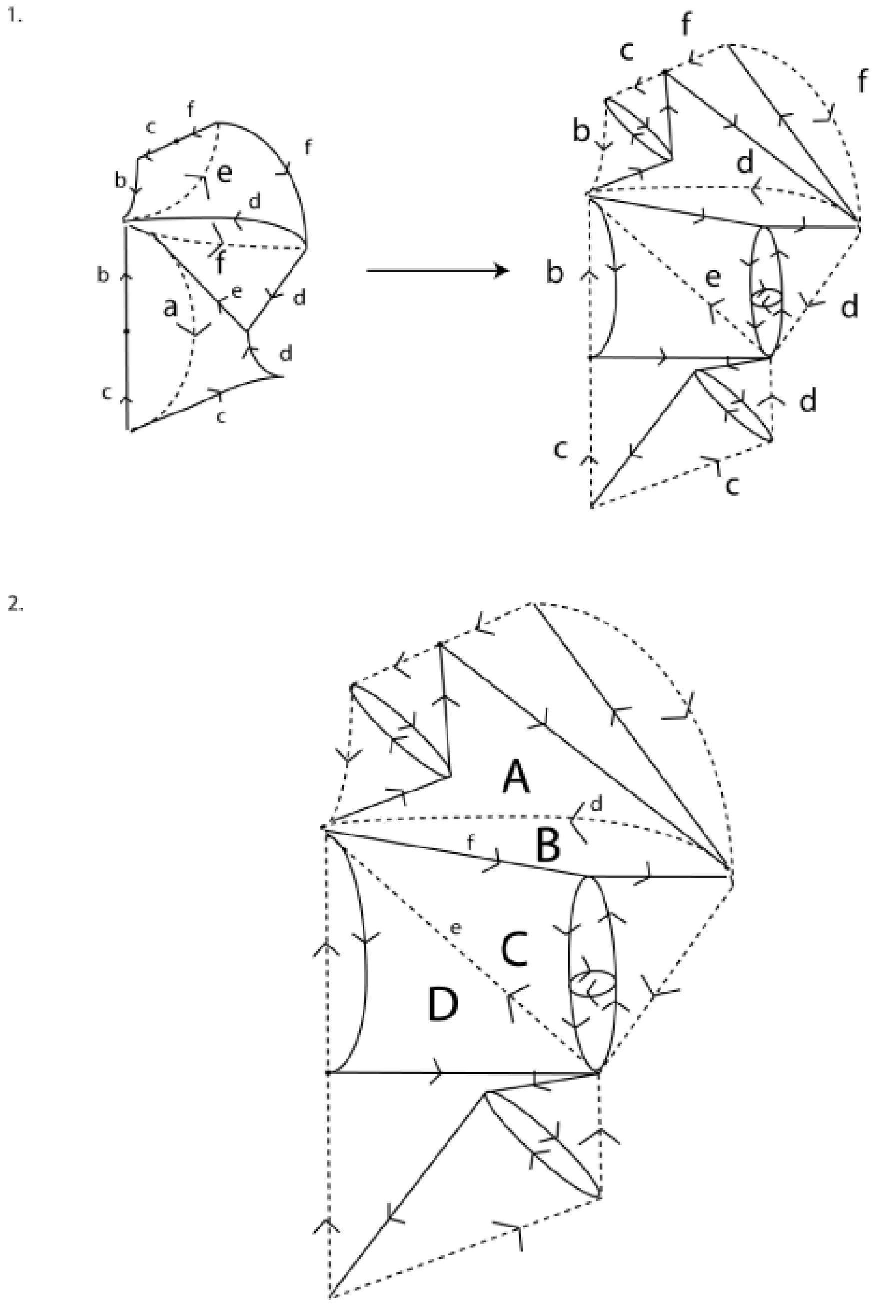

Figure 48: Note that regions A and D are frozen. This means that region B must move over to replace the old, loaded edge d, and C must move over to replace e. But B and C share the edge f, so they can't move in opposite directions.

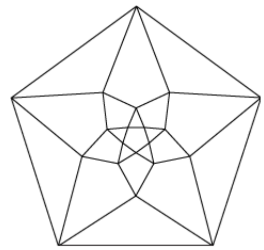

Figure 49: A projection of the (5,5)-radial link

7 Radial Links

After seeing some examples where subdivision rules are difficult to obtain or non-existent, we turn to a much nicer set of links.

Definition. A **(p,q)-radial link** is constructed from a projection made by nesting polygons. It has two parameters, p and q. They can be made by taking a standard projection of a torus knot and making it alternating. To draw a radial link, draw a p-gon. Inscribe a p-gon inside it by connecting the midpoints of the edges. Repeat this process until we have drawn q p-gons (including our first one). Turn the graph into alternating link in the usual way. like the trefoil, there will be left-hand and right-hand versions of the link. Note that the vertices of the original p-gon are not intersections of the link. The $(p, 2)$-radial links are just the two-braids themselves.

Radial knots are included in the class of knots called Turks-Head knots by sailors [3]. Some of the nicest links are radial links, including the Hopf link (the (2,2)-radial link; see Section 12 for a picture), the Borromean rings (the (3,3)-radial link), the trefoil (3,2) and all other two-braids (p,2), and the figure-eight knot (2,3). In a sense, the most 'homogeneous' links are radial links:

Theorem 8. The only alternating projections of links where every region has the same number of edges are the unknot, the Hopf link and the Borromean rings.

Proof. A link either has intersections or it does not. If it has no intersections, it consists of one or more copies of the unknot, and is homogeneous if and only if there is only one copy of the unknot.

If there are intersections, let E, F, and V be the number of edges, faces (or regions) and vertices, respectively, of a projection of a link. Assume all regions have n edges. Since each vertex has two edges "coming in", and each edge comes into exactly one vertex, $E = 2V$. Since our link projection lies on a 3-ball, $F - E + V = 2$. Since each face has n edges, and each edge touches exactly 2 faces, $F = \frac{2E}{n}$. Thus, $2 = F - E + V = \frac{2E}{n} - E + \frac{E}{2} = \frac{4E-nE}{2n} = \frac{4-n}{2n}E$. Thus, $E = \frac{4n}{4-n}$. But E is a positive integer, so $n = 2$ or 3. If $n = 2$ or 3, then the number of edges is 4 or 12, and the number of vertices is thus 2 or 6. □

Exercise. Finish the proof by showing that the only alternating link projection with 4 edges and 2 vertices where every region is a bigon is the Hopf link. Show that the only alternating link projection with 12 edges and 6 vertices where every face is a triangle is the Borromean rings.

Note that the unknot can be thought of as a radial knot with $p = 0$ and that the Hopf link and the Borromean rings are genuine radial links with $p = 2, q = 2$ and $p = 3, q = 3$, respectively. The links with $p = q > 3$ are not homogeneous since they all contain triangles.

For this next proof and all others, I will use the projection of the radial links that we used in their definition. A radial link decomposes into $q + 1$ rings, where the innermost and outermost rings are single regions with p edges. The second and qth rings consist of p triangles with one face toward the p-gons and two faces touching neighboring regions in the next ring. If $2 < i < q$, the ith ring consists of p quadrilaterals, with two consecutive edges touching two consecutive regions in the ring $i + 1$, and the same for $i-1$. Since $p > 2$, each pair of regions in a specific ring intersects once or not at all.

The main goal of this section is Theorem 9, which will later generalize to a stronger theorem: Radial links with $q > 1, p > 2$ have a subdivision rule.

Since we have previously dealt with the case $q = 2$ (i.e. the 2-braids), we will focus on the case $p, q > 2$.

Before we begin, it's important to recall the meaning of the word 'cross-touch'. When I say two regions cross-touch, I mean that they intersect 'kitty corner'. In other words, their intersection near that vertex is a point and not a line. In particular, we are interested in when two edge regions cross-touch. They always cross-touch their two neighbors (in Figure 49, for instance, all the triangles on the outside border cross-touch their neighbors). Cross-touching causes complications when it occurs more than this.

We need a definition for the following lemma. Let R_1 be a region. If R_1 and another region R_2 border a third region R_3, then they will cross-touch if the edges of R_3 that they border are next to each other. For example, the five outside triangles in Figure 49 at the beginning of the section all border the outside pentagon, and they cross-touch with their two neighbors at a vertex that lies on R_3. However, in the Figure Eight Knot (see Figure 37 in Section 6), the regions $a_1b_1b_2$, $b_1a_2b_2$ both border $a_1a_2b_2$, but they cross-touch at another vertex not part of $a_1a_2b_2$. If two regions with a common neighbor only cross-touch at vertices on the common neighbor, they form an **isolated pair** with respect to that neighbor. If a region forms an isolated pair with every region it cross-touches, it is **isolated**.

Note that an isolated region is never frozen.

Lemma 1. In a radial link with $p, q > 2$, every region is isolated.

Proof. Let R_1, R_2 border R_3. Then if R_3 is in ring i, R_1 and R_2 must individually be in row $i-1$ or $i+1$. Either they are in the same row or are two rows apart.

Case 1: R_1 and R_2 are in the same row. R_1 and R_2 cannot be in the first or $(q+1)$th row because there is only one region in each of those rows. If R_3 is also not in the first or $(q+1)$th row, then it borders exactly two regions in the nearest two rows (This is not true if it is one of the triangles that touch the innermost or outermost rows, since they border only the single region in the innermost or outermost row. But since R_1 and R_2 cannot be in those rows, we do not care about them for this proof). The two edges of R_3 that touch those regions intersect exactly once, and this intersection is also the only point of intersection of those two regions as well, which shows that R_1 and R_2 have no point of intersection except that on the boundary of R_3.

If R_3 is in the first or $(q+1)$th row, then every intersection of pairs of regions in the next row is a vertex of R_3, giving the desired result.

Case 2: R_1 and R_2 are two rows apart. Since $q > 2$, at least one of the two regions is a triangle or quadrilateral, say R_1. Each triangle or quadrilateral intersects at most one region from two rows away in either direction, so R_1 and R_2 intersect at most once. Assume they do intersect. Then R_1 borders exactly two regions in the row containing R_3, and those edges by which it borders them have as vertex the point of intersection of R_1 and R_2. Thus, the intersection of R_1 and R_2 lies on R_3. □

We are now able to prove:

Theorem 9. Radial links with $q > 1, p > 2$ have a subdivision rule.

Proof. The case $q = 2$ was dealt with previously; these are the two-braid links. Assume $q > 2$. Then we will construct an explicit subdivision rule. This is an alternating knot, so fix an orientation on the regions. Assign every clockwise edge subregion (or every counterclockwise edge subregion) at each stage to become the

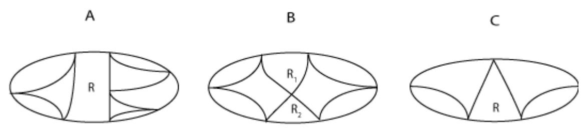

Figure 50: These show ways that the assignment of subregions could go bad. In case A, we would not be able to move the edges of R as we would like because it shares too many vertices with the outside. But radial links are clearly prime, and this doesn't happen in prime links (see the final paragraph of Section 2). In case B, the common vertex of R_1 and R_2 prevents both subregions from being assigned to cover an edge at the same time. But this never happens in radial links since regions don't cross touch. Finally, there could be frozen regions, but this never happens in radial links because every region is isolated.

new edge. Since exactly one region on either side of an edge is clockwise, this is well-defined at that edge. Now, we have to worry if it is globally well-defined; if two edge subregions that move share an edge or a vertex, then their movement must agree on their common vertex or line (see Figure 50). But since only clockwise (or only counterclockwise) regions move, no two regions that move share an edge. They can only cross-touch. Since every region is isolated, there are no frozen regions (no cross-touching with the outside regions), and if two edge subregions do cross-touch, they do so on an outside vertex, which is fixed by all movements of regions. Thus, this subdivision rule is globally well-defined. $\square$

Note that the proof of the case $q > 2$ used only the fact that all regions are isolated. Thus, every prime, nonsingular link projection where every region is isolated will admit a subdivision rule as constructed above. In fact, it will admit two subdivision rules, one for moving out clockwise regions, and one for moving out counterclockwise regions, but these may give the same subdivision rule. We give such links a name:

Definition. A link that has a nonsingular, prime projection in which every region is isolated is called **dispersed**.

The most basic example of a dispersed link is the Hopf link. In a non-loaded region, there are exactly three regions, one interior and two on the edge, both of which are isolated. In a loaded pair, there are only two subregions, and they are isolated as well. Thus, the link is dispersed, and the subdivision rule described above works.

Exercise. Find all the tiles used in the subdivision rule described above for the Hopf link. Draw out the first three or four stages.

Exercise. Find the subdivision rule given the Borromean rings by the above theorem.

Note that in both of these exercises, the two subdivision rules given by theorem 9 are the same, up to labels. Also, the first exercise shows that the Hopf link has a linear growth in the number of words of a specific length (see Section 5.1), which is understandable as its complement is homotopy equivalent to the torus.

8 Dispersed Links

In the previous section, we defined a dispersed link as a link with an alternating, prime projection in which every region is isolated. So far, the radial links are our only example of dispersed links. We'd like to find more examples, and so this section gives a method for constructing more dispersed links from old ones.

Basically, what we do is generalize radial links. In a radial link, we take a two-braid and inscribe in it a regular polygon with as many vertices as the original two-braid, then inscribe a polygon of the same size inside of that, etc. More general dispersed links can be constructed by inscribing polygons in **any** region, not just the innermost. Also, we can relax the requirements further by allowing the inscribed polygon to have fewer vertices than the region it is inscribed in. However, at most one vertex of the polygon can be placed on any one edge of the original knot. This may change the connectivity of the link (i.e. the number of components can go up or down), and that's ok.

Before stating the theorem, recall that a dispersed link is a nonsingular prime link satisfying two requirements, first that every edge subregion of a non-loaded region is isolated, and second that every edge subregion of a loaded region is isolated or paired.

Theorem 10. Let L be the projection of a dispersed link and R a region with $n > 2$ vertices. Then the graph obtained from L by inscribing a polygon with $2 < j \leq n$ vertices in R as described above is the projection of a dispersed link.

Proof. Label the interior of the inscribed polygon R', and the other regions cut out of R as R_1, R_2, etc. (see Figure 51).

Consider R'. We will show that every pair of regions touching it form an isolated pair, as in Section 7

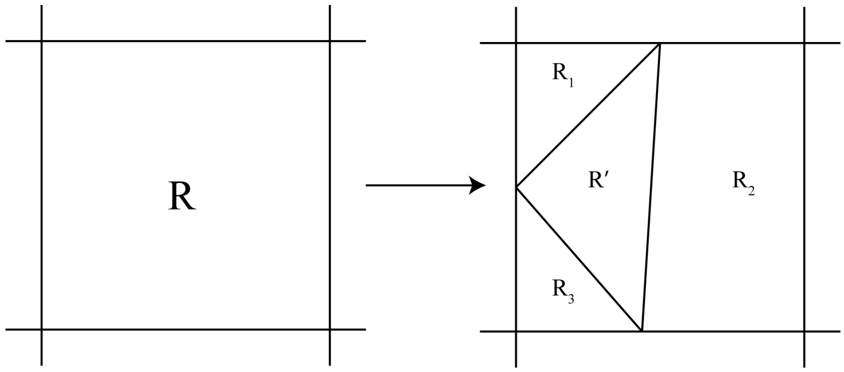

Figure 51: Inscribing a polygon

If two edge subregions of R' intersect, they do so only on the boundary of R', since the only edge subregions of R' are the R_i, and all of their intersections lie on R'.

Now consider R_i. If two regions bordering R_i (neither of which is R') cross-touch off of R_i, then before this construction, they would have bordered R and cross-touched off of R, contradicting the hypothesis that L was dispersed. R', when it intersects those regions that touch R_i, does so exactly on the border of R_i, so it causes no problems.

Finally, we consider regions A that are not R' or R_i for any i. If two regions bordering A cross-touch at a vertex not on A, they can't both be R_i's, since each pair of R_i's intersects at most once, and when they do intersect, they do so on the boundary of the unique outside region that borders both. This is true because $j > 2$. (see figure 52). If just one is an R_i, then R and the other region would have had an intersection not on A in the original knot, which would be a contradiction. And if they don't touch R' or the R_i at all, they would have intersected the same way originally.

Thus, every region is isolated, and this process gives us the projection of a dispersed link. $\square$

Exercise. Dispersed links must be prime. Complete the above proof by showing that

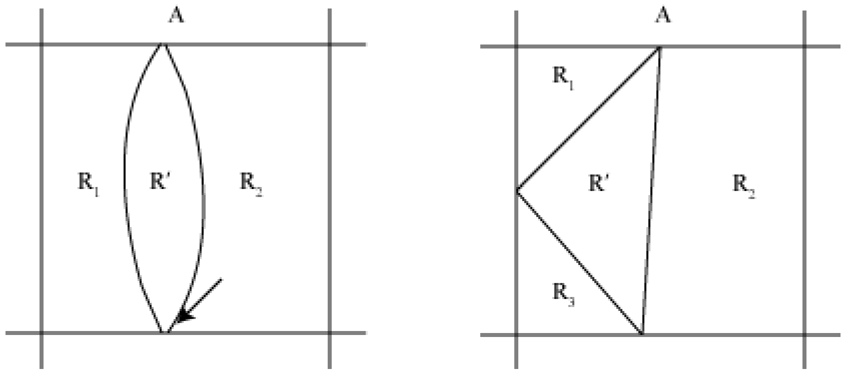

Figure 52: The figure on the left shows what would happen if we allowed $j = 2$; the region A borders both R_1 and R_2, which have an extra intersection at the arrow. The figure on the right illustrates the fact that this does not happen if $j > 2$.

the alternating projection obtained is prime.

(Hint: An alternating link with a given alternating diagram is prime iff the diagram 'looks' prime, i.e. no two regions touch at more than one edge).

9 Platonic Links

In general, the subdivision rules associated to a link have a large number of tiles. In an asymmetric link where each face and each loaded pair of faces have distinct subdivision rules, there will be $6n - 8$ tiles, where n is the number of regions in the link complement. This includes the tiles in the upper half of the link complement, which are mirror images of the bottom half. An example with 100 tiles is given in Section 13

Sometimes, symmetries will reduce the number of tiles necessary; the Borromean rings have 3 tiles instead of the expected 40, since every face is identical to every other face and has 3-fold rotational symmetry as well. Two natural questions to ask are, "What is the smallest number of tiles possible for the subdivision rule of a link?" and "Which links have this smallest number of tiles?" We will answer both questions in this section. We make the assumption that all subdivision rules are standard, in the sense that they are obtained as in previous sections, by taking the replacement rules for the link complement and assigning regions to move over and replace disappearing loaded edges. The unknot can be given a subdivision rule with a single tile type by treating its complement as a solid torus, but we don't include this in our list of candidates.

We've previously encountered links with only two tile types: the two-braid links. These are the only links with two tile types. To see this, look at the link projection. Every region in that link projection must be subdivided, so there can be at most two kinds of regions, p-gons and q-gons, with p and q not necessarily distinct. Then there will be loaded edges with $p + q - 2$ edges. To keep the number of tile types down to two, then, either $p = q$ or p or q=2. In the first case, we have a homogeneous link, which we have already classified. So assume $q = 2$. Now, we can't have p-gons touching other p-gons, since that would produce loaded regions with $2p - 1$ vertices, a contradiction. So every edge of a given p-gon has p bigons touching it, so every vertex is filled with four edges. Thus, the link must be a two-braid, if $q = 2$.

What if the link is homogeneous? The only homogeneous links besides the unknot (which we studiously ignore) are the Borromean rings and the Hopf link.

The Borromean rings cannot have only two tiles. One way of seeing this is looking at the first stage of the construction. There are 8 triangular regions; when these are replaced, we have 6 loaded edges, each with a pair of edge subregions bordering them. Thus, there is an average of 1.5 edge subregions that must move out; this is impossible to achieve with a single triangular tile. And in the next stage of construction, there are loaded pairs with four edges, giving at least three tiles.

Exercise. Prove that the Hopf link cannot have a subdivision rule with only one tile type.

Note that we have ignored subdivision rules with a single tile type. Such a link would necessarily be homogeneous, and we have shown that there are no homogeneous links with a single tile type. Thus, we have proved:

Theorem 11. The only links that can be given a standard subdivision rule with exactly two tile types are the two-braids, and this is the minimum number of tile types.

Now, what about those with three tile types? Given that we found an infinite family with two tile types, we might expect even greater variety for three tile types, but we would be pleasantly surprised. Assume a link has a subdivision rule with exactly three tile types. As above, we must have only two possible numbers of edges for regions in the link complement, since we will have p-gons, q-gons, and $p + q - 2$-gons. The cases p or $q = 2$ and $p = q$ have been dealt with previously, since they are the two-braids and the homogeneous links. Recall the proof of Theorem 8 in Section 7. If every region has more than 3 edges, then the number of faces is less than or equal to $\frac{2E}{4} = \frac{E}{2}$, so $2 \leq \frac{E}{2} - E + \frac{E}{2} = 0$, a contradiction. So p or $q = 3$. Assume $q = 3$. Then our link must be made of triangles and p-gons, where every edge has a p-gon on one side and a triangle on the other and every vertex touches four edges. The last condition holds for every link, the middle condition happens since every different combination of regions touching an edge gives a different loaded pair. This information completely determines the link, if it exists.

Let's end the suspense and resolve the question of existence. To do so, we get to construct links using two-dimensional replacement rules, which we have not yet done. Our fundamental blocks are now two-dimensional p-gons and triangles instead of three dimensional chunks of the link complement, the edges of the 2-d objects replace the faces of the 3-d objects, and the vertices take the place of the edges.

We begin constructing our link by placing a single p-gon:

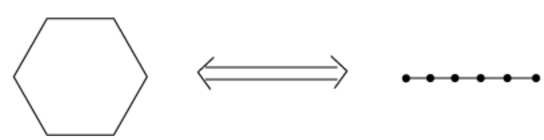

Just as we squashed the 3-d chunks into a 2-d figure with one or more invisible regions, we squash the 2-d p-gon into a 1-d line with one or more invisible segments. We go down a dimension because we are only interested in the boundary of each stage, as before. The replacement rules for this link are shown in Figure 53. Note that there are loaded vertices, just as there were loaded edges before. Since four regions meet at a vertex, the same as when four chunks touched an edge, loaded vertices behave the same as loaded edges.

In Figure 53, we had $p = 6$, but for any larger p, the subdivision rules are essen-

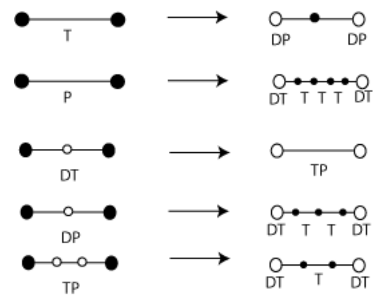

Figure 53: T stands for triangle, P for p-gon, DP for double p-gon and TP for triple p-gon. The labels represent what will be glued on **next**. The white dots are loaded vertices, the black dots are non-loaded. When the edges of a subdivision are labelled DP, that means that they are half of a loaded pair or are the endpoints of a loaded triple, and when they're labelled TP, they're the center of a loaded triple.

tially the same, but P now has $p-2$ vertices in the interior, DP has $p-3$ and TP has $p-4$. These tile types shown are the only ones that occur for $p=6$. This is because every tile type except DT adds enough black dots to the interior to space out the loaded vertices from each other. The only complication would be if we had two DT's right next to each other, but from the picture we see that that does not occur. Thus, if we try to construct a link with only triangles and p-gons, we can never close it off, and it grows bigger and bigger until it tiles the plane.

However, such a link, with only three tile types, does exist for $n = 3, 4$, or 5. For $n = 3$, it is the Borromean rings (the stages of its construction can be seen in the answer for the first exercise of Section 7); for $n = 4$, it is the (4,4)-radial link; and for $n = 5$, it is a new dispersed link. It is pictured below: Note that the rows alternate between pentagons and triangles. This is not a radial link (cf. Figure 49), but it is related to them:

Exercise. Show that this link can be obtained by inscribing and/or circumscribing four pentagons to a two-braid knot with ten vertices.

Thus, we have proved:

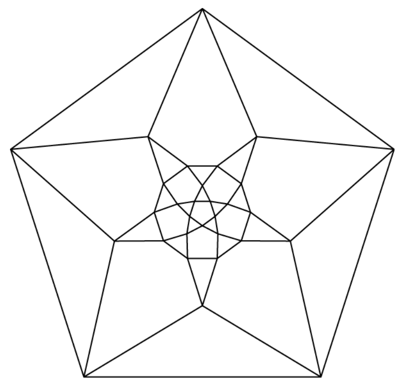

Theorem 12. There are only three links whose minimum number of standard subdivision tiles is three.

The existence of these special links has been known for some time (for example, see [2] or [11]); I found these articles after discovering the links independently, and I have yet to find a name for these links as a group, so I feel justified in calling them the Platonic links. This is because they are strongly related to the five platonic solids. Take a projection of any platonic solid, and on each face, inscribe a polygon by connecting the midpoints of each edge. If you erase the old edges, what remains is the projection of one of the three platonic links! The dodecahedron and icosahedron both give the big platonic link, which has twelve pentagonal faces and twenty triangular faces. The octahedron and cube both give the (4,4)-radial link, with six square faces and eight triangular faces. Finally, the self-dual tetrahedron yields the Borromean rings, which have two sets of four triangular faces (also, they form a standard octahedron).

Thus, there is a deep, unexpected connection between alternating links, subdivision rules, and platonic solids.

10 Torus Links

Consider a standard depiction of a (p, q) torus link on the torus represented as a rectangle with opposite edges identified (all examples in this section will have $p = 4$ and $q = 3$):

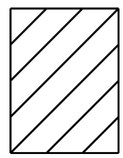

The slope of the diagonal lines is one. The link intersects the left edge p times and the top edge q times. Note that this depiction is somewhat asymmetric, dividing the rectangle into triangles and trapezoids. We will make it more symmetric by adding $p - 1$ horizontal lines and $q - 1$ vertical lines:

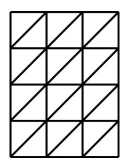

Notice that the torus is divided into $2pq$ regions, each identical.

Now, if we wanted to embed this torus in space, we could identify the top and bottom edges first and then the two sides, or the two sides first and then the top and bottom. Create one torus of each kind, and fill them in to become two solid tori. In one solid torus, the vertical lines become meridians, and in the other solid torus, the horizontal lines become meridians. The diagonal lines become a (p, q) torus link (which is the same as a (q, p) torus link). Identifying the two solid tori by the identity map on their boundary tori yields S^3.

We want a nice decomposition of the complement of the torus link into simply connected, easily manageable parts. Cut each of the solid tori into cylinders by slicing

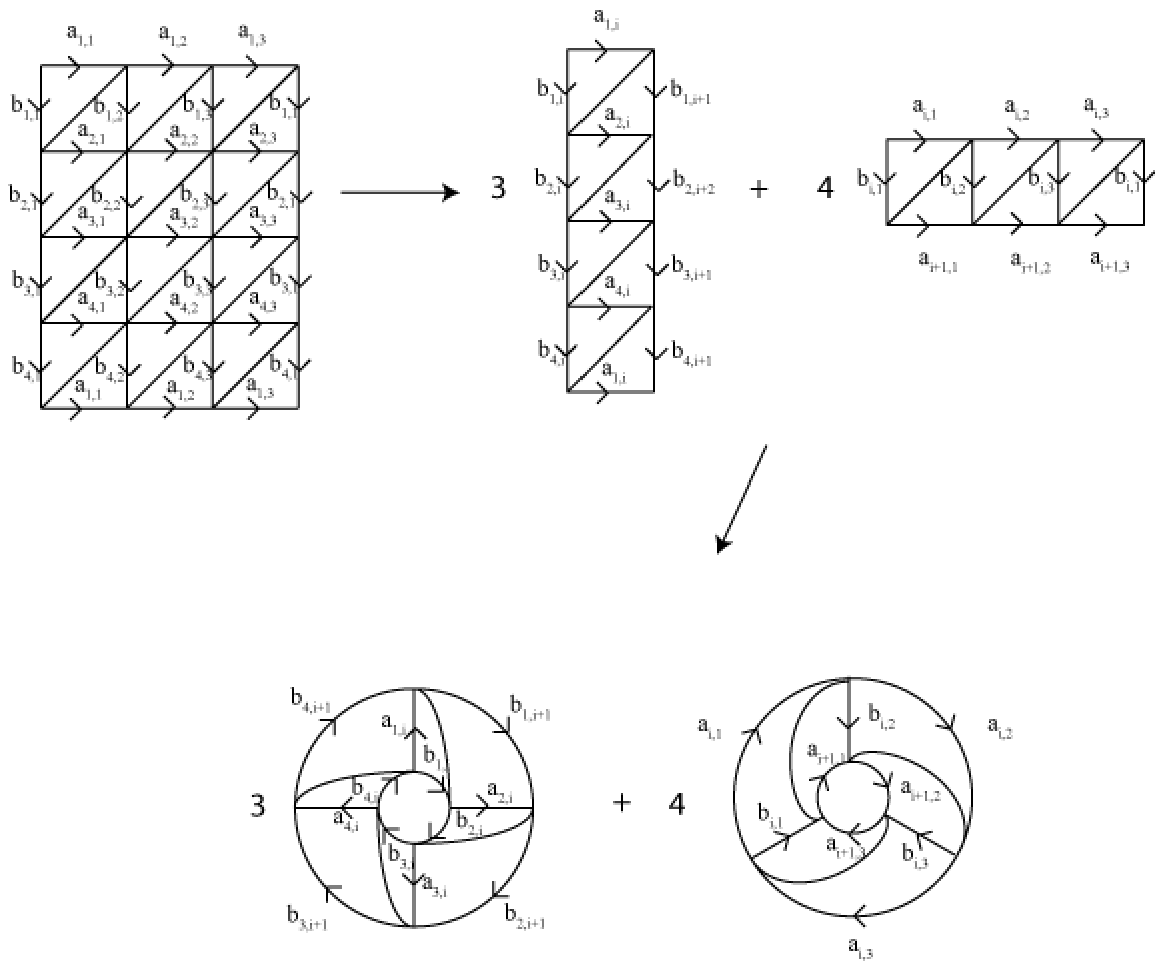

Figure 54: Decomposing the complement of a torus link. Note that the circles at the bottom are cylinders; the inside and outside faces are the two ends of the cylinder.

along every meridian that was a line in our original diagram. Doing so gives us p cylinders with q squares around the circumference and q cylinders with p squares around the circumference. They can also be obtained by cutting apart our original rectangle diagram into vertical and horizontal strips and identifying their narrow edges, then filling in (see Figure 54).

If we now contract the arcs corresponding to the link to points, we get chains of theta shapes (see Figure 55).

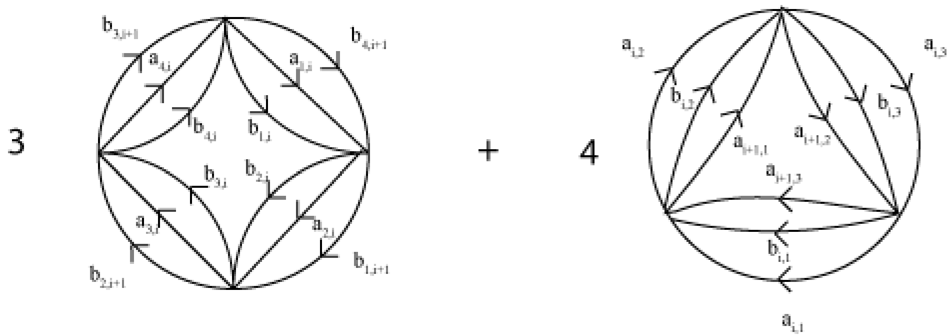

Figure 55: The cylinders after shrinking the arcs of the link to points.

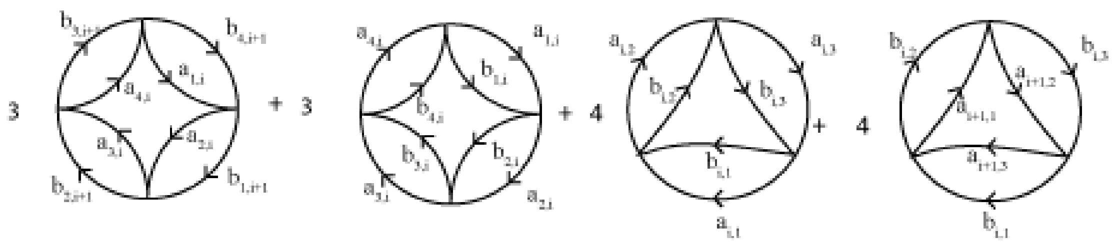

Figure 56: The cylinders after they have been split in half.

Now, some edges are identified three times, and some are identified four. If we cut each cylinder along a disc bounded by the center of the three circles, we get chains of bigons (see Figure 56), and now every edge is identified four times.

Note the similarities between this decomposition and the decomposition we obtained earlier for the two-braids, i.e. the torus knots with $p = 2$. This decomposition leads to a very clean and simple subdivision rule. Let's begin with the replacement rules. Every p-gon or q-gon is glued to another p-gon or q-gon of the same size on a different bigon chain:

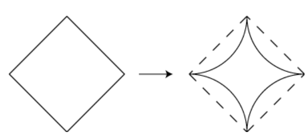

Each bigon that is part of a chain of p bigons is attached to a bigon on a chain of q bigons, and vice versa:

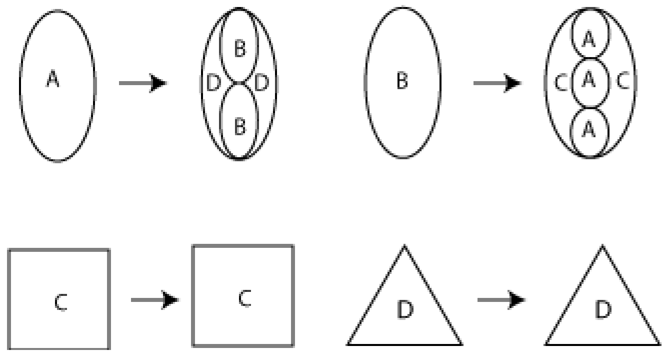

Figure 57: The subdivision rules for the $(3, 4)$ or $(4, 3)$ torus knot

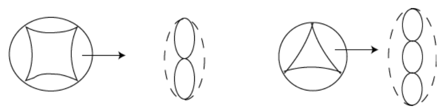

Each edge is identified four times, so we get loaded edges, just as before. Since every fresh, non-loaded edge has a bigon on one side and a p-gon on another side, by consulting the replacement rules above we see that the only loaded regions in the first stage of gluing are those that consist of one p-gon (or q-gon) and one bigon:

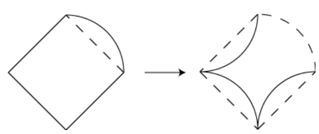

Notice that the only new loaded regions created are, again, p-gons (or q-gons) with a single bigon attached. Thus, we have found all the replacement rules. Noting the similarities of these rules with those of the two-braids found earlier, we can construct a subdivision rule on the same pattern: every edge subregion that is a bigon moves out to be absorbed into the new edge, and every other edge subregion stays put. The final subdivision rule is shown in Figure 57.

Exercise. Draw the first few stages of the subdivision rule for a two-braid link according to both methods described in this thesis. What similarities and differences do you notice?

11 Subdivisions With Boundary

To this point, we have considered links as 1-dimensional objects lying in S^3. Now, we bulk them up by considering a regular tubular neighborhood of the link. Deleting such a neighborhood leaves a 3-manifold with torus boundary components. We can use the same polyhedral decomposition as before on this new 3-manifold, with a slight difference at the vertices. A careful examination of the description of the decomposition shows that at each vertex of the link becomes a square with four lines of identification radiating from it.

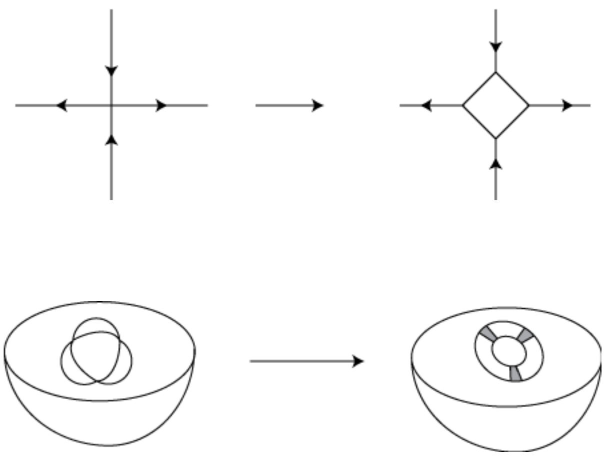

Note that collapsing the boundary squares gives the same picture as the original decomposition. A subdivision rule that uses this fatter link will be called a **subdivision rule with boundary**.

These boundary squares are never subdivided or replaced, and are not dual to the Cayley graph of the fundamental group, as the other faces are. However, they add great flexibility to the possible subdivision rules for the knot. In particular, they prevent the occurrence of frozen regions.

Adding the boundary squares does not change the combinatorics of the replacement rules of the links. Regions are attached to the same other regions as before and loaded pairs occur in exactly the same places. It's only drawn differently, with these boundary squares slowly growing out of each vertex. It is not really necessary to label the edges of the boundary squares, because regions are still determined by the edge labels from the complement of the 1-dimensional link.

Theorem 13. Every prime, nonsingular alternating link admits a subdivision rule with boundary.

Proof. By primeness and nonsingularity, no region has more than one loaded edge, as with complements of 1-dimensional links. Assign an orientation to each region like before. The boundary components are not oriented and aren't counted as faces. A clockwise region shares an edge only with boundary squares and counterclockwise regions. Because only three lines come in at each vertex now, a clockwise region only touches the vertices of regions it shares an edge with. The boundary squares act as buffers to prevent cross-touching.

This means that we can use the subdivision rule created for dispersed links: always pick the clockwise region (or counterclockwise) to replace an edge. This works since every edge touches exactly one clockwise region and no two clockwise regions share any edges or vertices. Thus, the assignment is well-defined and gives a subdivision rule for every alternating link complement. □

Let's look at some examples. The trefoil is the easiest example. Just as before, there are three tile types for the trefoil (four if you count the boundary squares). The replacement tiles are shown below, along with the subdivision rule we get by moving all bigons out to replace loaded edges. Edge labels are omitted, and grey tiles represent boundary squares. Note that this is the same subdivision rule we used in the case of a 1-dimensional link, if we collapse grey squares.

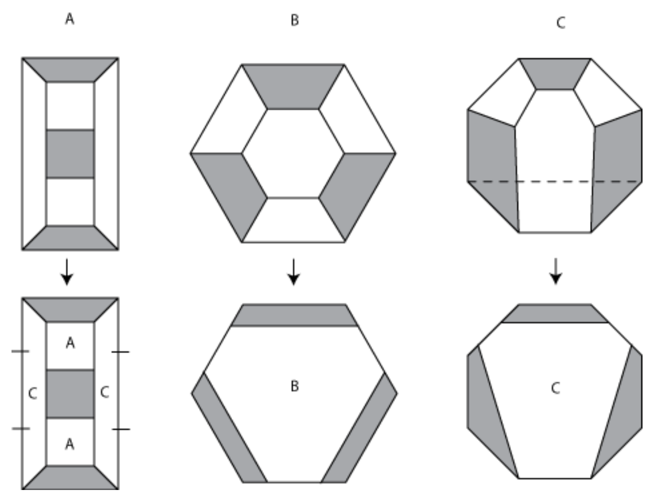

A color-coded picture of the first few subdivisions of tile A can be found in Section 13. These subdivision rules should be interpreted using a few key facts about subdivisions with boundary. For one thing, the boundary squares on the edge of the tiles should always be placed to line up with the boundary squares of the previously placed regions. What's happening is that we're building the universal cover of the torus at each boundary square.

Also, the subdivision rule for loaded pairs is intimately connected to the subdivision rules for their constituents. In general, the tile types inside the subdivision can be read off the clockwise constituent, and the placement of interior edges can be read off the counterclockwise tile (if clockwise regions move out in the subdivision). This is true in the case of a 1-dimensional link as well. If you look at the subdivision rule for the Borromean rings, tile C is the exact same as tile B except that we allow one of the edges of tile B to 'bow out'. Tiles B and C in the trefoil rule above give another example of how counterclockwise tiles are related to loaded pairs.

Also, note that if every boundary square tile is contracted to a point, we get the exact same subdivision rule as in the one-dimensional case.

A more illuminating example is given by the (2,4)-radial link, which was our counterexample in Section 6.3. As every link has a subdivision rule with boundary, we can now get an explicit subdivision rule for this link. The various replacement tile types and their corresponding subdivision tiles are shown in Figure 58.

Again, remember to place the boundary squares of an edge of a new tile so that they touch boundary squares of the old tiles. Note again the relationship between loaded regions and their constituents. Tiles D,E, and F have the exact same interior edges as tile C, the only counterclockwise tile, and the types of their interior tiles can be found in roughly the same places as in tiles A,B, and B, respectively.

Section 13 contains the first few subdivisions of tile A, with different coloring patterns to emphasize different aspects of the subdivision. Especially illustrating is the image that highlights boundary squares. There, you can see the universal cover of the torus growing at each vertex of the link in a beautiful and edifying pattern.

As a final note, the subdivision tiles show why this particular subdivision rule cannot be used in the one-dimensional case. If we try to contract the boundary squares to points as we did for the trefoil, we get singularities from the boundary squares that stretch from one outer edge to another in tiles C,D,E, and F. However, we can still calculate the growth function of the link complement, and we can quickly calculate the combinatorics of each layer by using the subdivision rule with boundary until we reach the stage we want, then contracting squares to points. Thus, by adding additional structure we have been able to simplify the theory, just as with torus links.

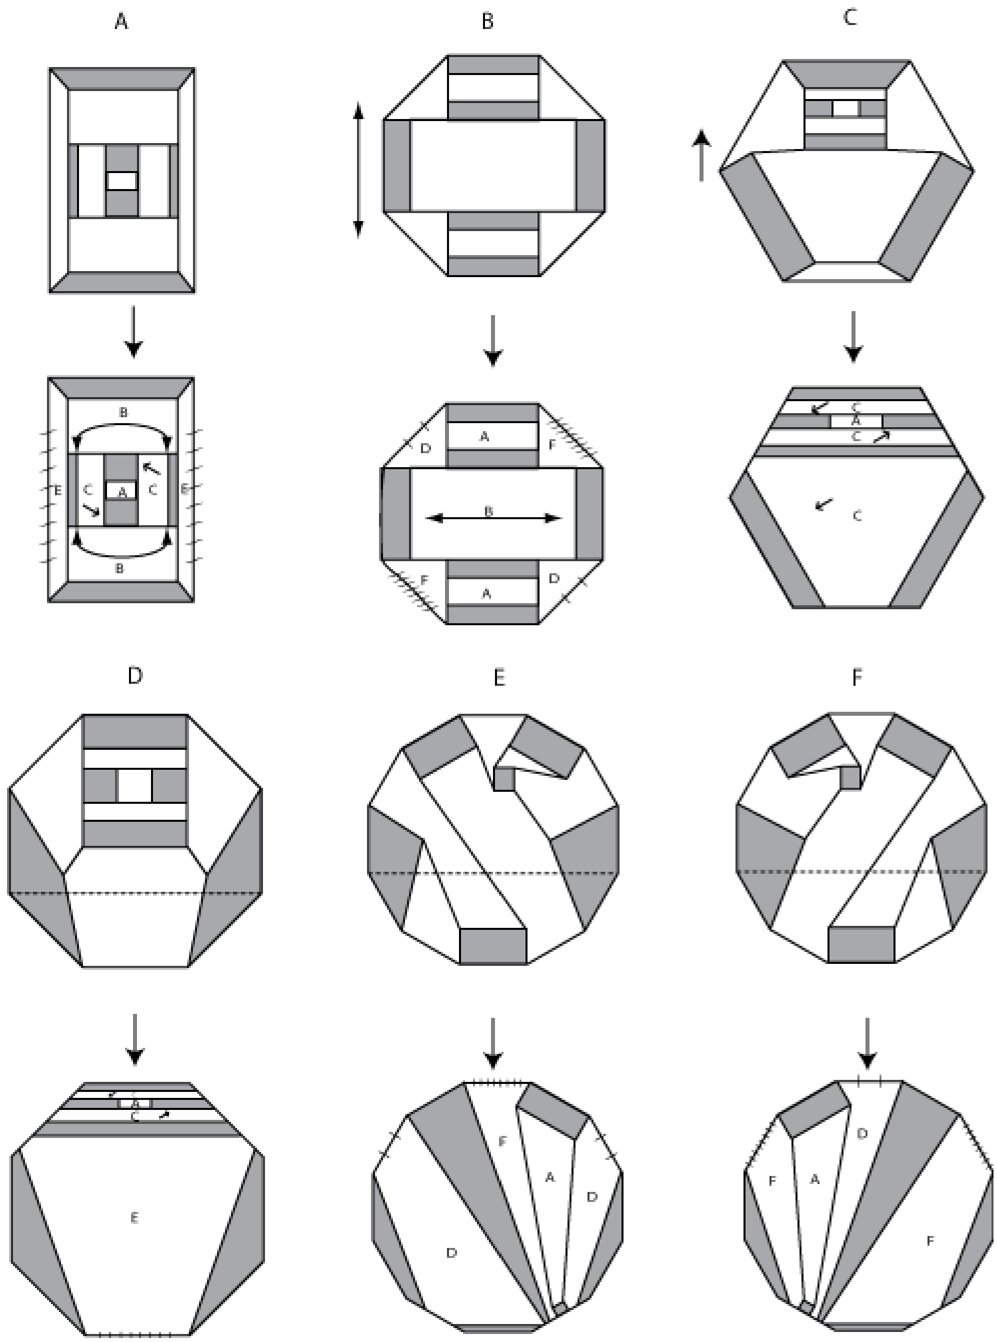

Figure 58: The tile types for the (2,4)-radial link. Grey tiles are boundary squares.

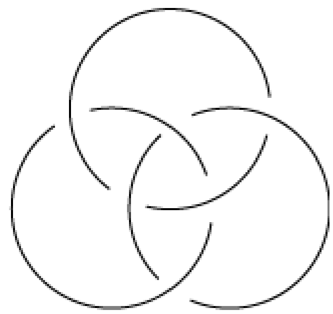

Figure 59: The Borromean rings.

12 Definitions

Absorbs: See **Moves over**.

Alternating Link: A link where crossings alternate between overcrossings and undercrossings on each component of the link.

Borromean Rings: See Figure 59. It is the **platonic link** corresponding to the self-dual tetrahedron. It is the (3,3)-**radial link**, and is **dispersed**. It is **homogeneous**.

Chunk of the Knot Complement: One of the two halves of the knot complement used in the construction of the universal cover.

Clockwise, Counterclockwise, Orientation: In drawing the chunks of the knot complement, the arrows that show how edges are identified form natural patterns. The most obvious pattern is orientation. The arrows around the edges of any region go in a clockwise or counterclockwise pattern. Reversing the orientation on all edges of both chunks of the knot complement does not change the construction at all, so there is no natural way of identifying which regions are clockwise and which are counterclockwise. We just pick an orientation and stick with it. I prefer having the edges corresponding to overcrossings pointing away in the bottom half of the knot complement. This determines an orientation.

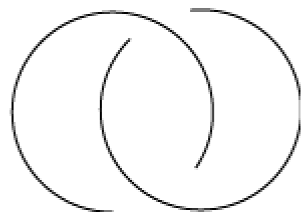

Figure 60: The Hopf Link.

Composite Link, Prime Link: Some links can be twisted into a position where a two sphere can separate the three sphere and only intersect the knot in two points. Such a link is called composite if the two pieces that the link is cut into are both knotted (i.e. neither component of the link complement minus the 2-sphere is a solid torus). A link that is not composite is prime. In this thesis, a prime link is also assumed to be non-**split.**

Cross-touch: Two regions cross-touch when they touch 'diagonally'. For instance, in a trefoil, the inner triangle and outer triangle cross-touch, but no triangle cross-touches a bigon. See **Rotation Principle**.

Dispersed Knot: A nonsingular prime knot in which every region is isolated.

textbfEdge Subregion: See **Subregion**.

Frozen: An edge subregion is frozen if it cross-touches the outside region(s). This occurs only in loaded regions. See Figure 35 in Section 5.

Homogeneous Link: A link in which every region has the same number of edges.

Hopf Link: The non-singular link with fewest vertices. See Figure 60. It is **homogeneous**, a **two-braid link**, the **(2,2)-radial link**, and is **dispersed**. Its subdivision rule has a linear growth rate.

Interior Subregion: See **Subregion**.

Isolated Pair, Isolated: If two regions with a common neighbor only cross-touch at vertices on the common neighbor, they form an isolated pair with respect to that neighbor. If a region forms an isolated pair with every region it cross-touches with respect to every common neighbor, it is isolated.

Knot: A tame embedding of the 1-sphere in the 3-sphere. Tame in this sense means more or less that it is not infinitely knotted.

Link: A tame embedding of one or more 1-spheres in a 3-sphere. Tame in this sense means more of less that they are not infinitely knotted. Knots are links with one component, or strand.

Overcrossing, Undercrossing: When projecting a link onto the plane, we get singular points or vertices where the original link crossed itself. If we pick a point on the original link and travel in a given direction, then those vertices that we pass from above are called overcrossings and those that we pass from below are called undercrossings. Every vertex is both an overcrossing and an undercrossing from some point of view.

Loaded Region, Loaded Subregion, Loaded Edge, Loaded Pair: In the construction of the universal cover, each edge can only be identified with three other edges. We begin with a fresh edge that has never been identified. We then glue chunks of the knot complement onto the regions on either side of the edge. The edge has now been identified with two others, so it admits only one more identification. Thus, one chunk of the knot complement must cover both regions that touch the edge. These two regions are called loaded regions (or loaded subregions, depending on the circumstances). The two regions together are called a loaded pair, and the edge is called a loaded edge.

Moves Over: When we say that an edge subregion moves over or is absorbed or becomes a new line or replaces a new line, we mean that its edges are homotoped to occupy the same position as the disappearing loaded edge.

Orientation: See **Clockwise**

Platonic Link: One of the three links based off of the pairs of dual platonic solids. See Section 9. They are all one-skeletons of Archimedean solids.

Projection: When this is used without further qualifiers, it means the projection of a link onto a plane. I generally assume that all projections are 'nice', i.e. the fiber over any point in the projection contains at most two points and arcs intersect nontangentially.

Radial Link, (p,q)-Radial Link: A radial link is formed from a projection given by a set of nested polygons. In general, there are q layers of p-sided polygons. See Section 7. The **Borromean rings** and the **two-braids** are examples. All radial links are **dispersed**. They can be obtained by making projections of torus links alternating.

Region: One of the connected components of the plane minus the graph of a link. It can also be thought of as a face of one of the 3-cells that form the universal cover of an alternating link.

Reidemeister Move: One of three ways of changing a projection of a link without changing the link itself. It is a theorem [1] that every projection of a given link can be changed into every other projection of the same link by a finite sequence of Reidemeister moves.

Rotation Principle: If two regions share an edge, they have opposite **orientations**. If two regions cross-touch, they have the same orientation. See Section 3.

Split: A link L is split if there is a 2-sphere in $S^3 - L$ that separates one component of the link from another.

Subdivision Rule with Boundary: A subdivision rule that includes boundary components which never subdivide. Obtained from the complement of a regular tubular neighborhood of a link.

Subregion, Edge Subregion, Interior Subregion: When you draw a knot in the plane with a specific desired region on the outside, all the other regions are referred to as subregions. We can also think of it in terms of the universal cover. When constructing the universal cover, we take a region and glue one chunk of the knot complement to it, and leaving the others, for now, unglued. Those free regions are called subregions of the original region.

Subregions come in two 'flavors', so to speak. Those that share an edge with the original region are called edge subregions. Those that don't are called interior subregions.

Two-braid Link: A link that is made up of a chain of bigons. They can be constructed by taking two strings, holding the ends together, twisting as many times as desired, then gluing the bottom strands to the top. They are the (2,p)-**radial links**, and are the only links with a standard subdivision rule with two tile types.

Unknot: An unknotted loop; the only knot whose complement in the 3-sphere is a torus. In general, it is difficult to tell from a projection if a knot is the unknot.

13 Computer Generated Images

This section consists of computer images of subdivision rules. Many programs were involved, including Bill Floyd's 'subdivide', 'subdividehistory', 'tilepack', and 'tilepack-history' programs [12] and Ken Stephenson's Circlepack [18].

Pages 77 to 81 contain the subdivision rules for the three platonic solids. They are color coded; a red tile is subdivided according to the tile shown in the column 'Red', for example. The other figures in these pages are the first few subdivisions of the blue tiles. Color is inherited, so a blue tile that was subdivided out of a red tile will be purple.

Page 82 gives an example of a non-symmetric dispersed link (with two components) where no collapsing of tiles occur at all. It has 16 vertices and 100 distinct tile types, the full amount needed (one for each face and one for each edge, which determines a loaded pair, and then all doubled by the mirror symmetry of the top chunk of the link complement). Compare this to the highly symmetric big platonic link, which has 30 vertices and only 3 distinct tile types.

Page 83 shows the first, second, third, and seventh subdivisions of the bigon on a torus using subdivisions with boundary. Blue is tile type C, red is tile type A, and yellow is the torus boundary. Type B never occurs in subdivisions of the bigon. See Section 11 for the tile types.

Finally, pages 84 and 85 show the subdivision rules with boundary for the (2,4)-radial link using different colorings. On page 84, tile type A is colored blue, tile type B is colored red, and tile type C is colored yellow. All loaded tiles (D,E, and F) and torus boundaries are white. On page 85, the torus boundary is colored blue and all other tiles are colored white. See Section 11 for details.

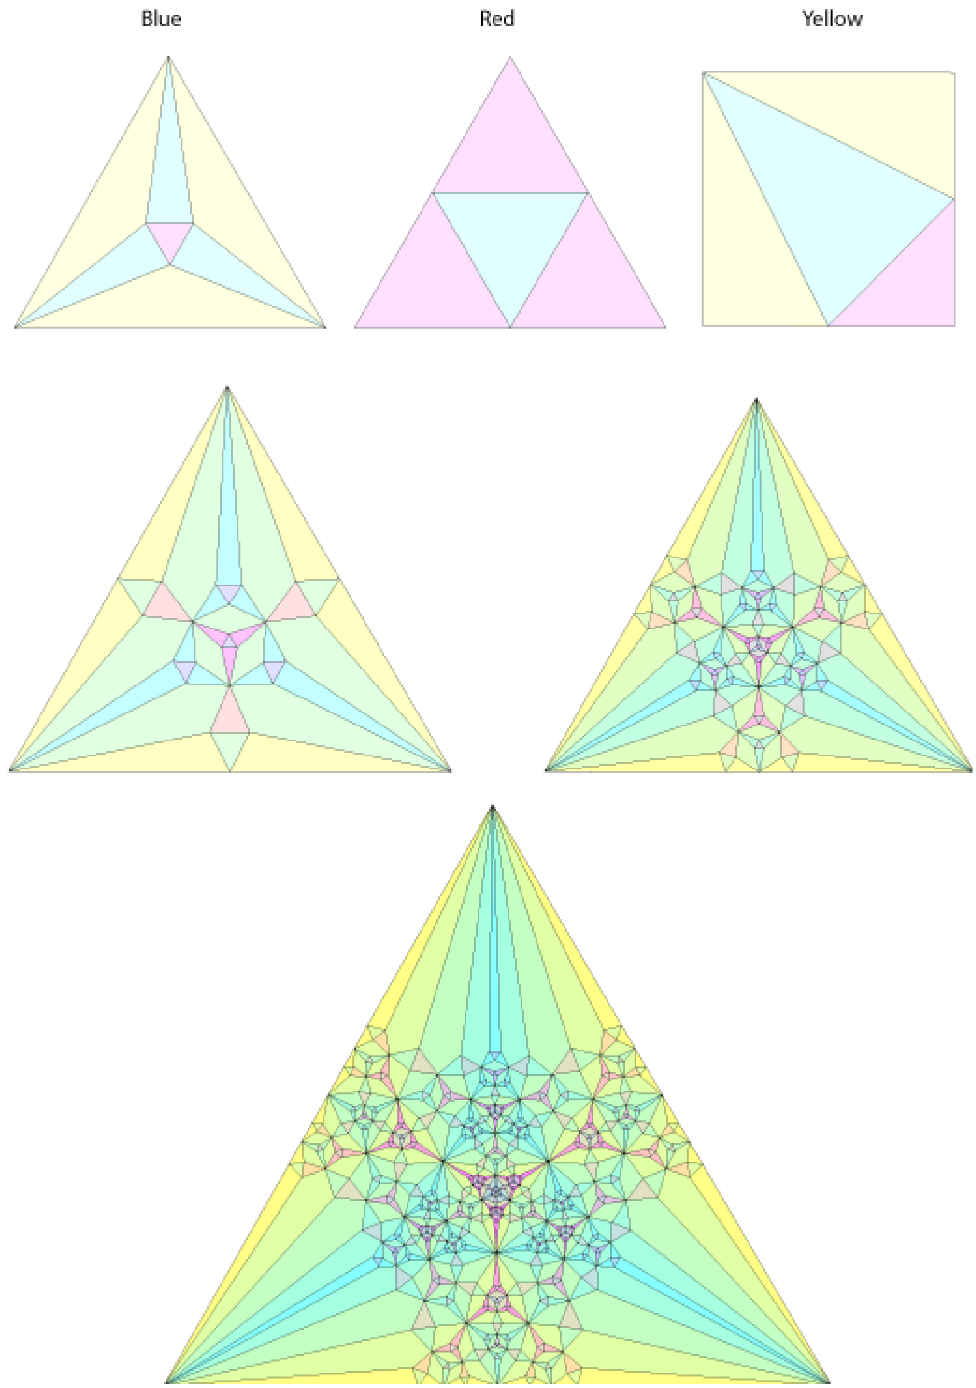
Blue
Red
Yellow

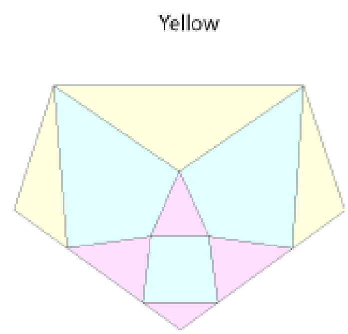
Blue
Red
Yellow

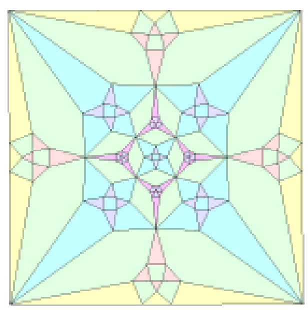

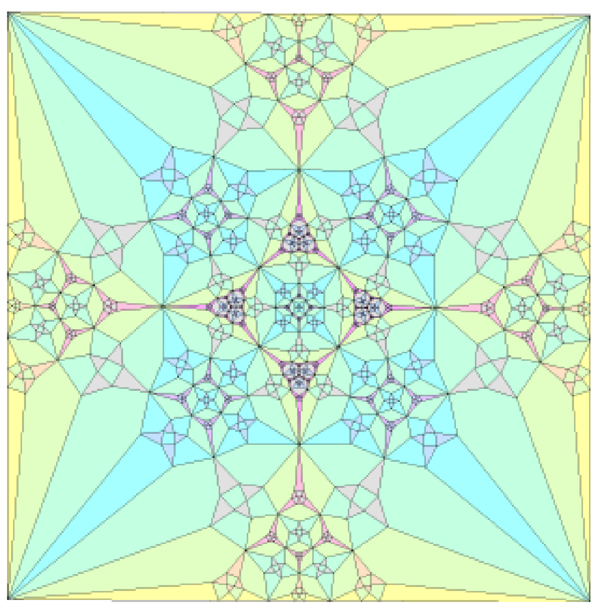

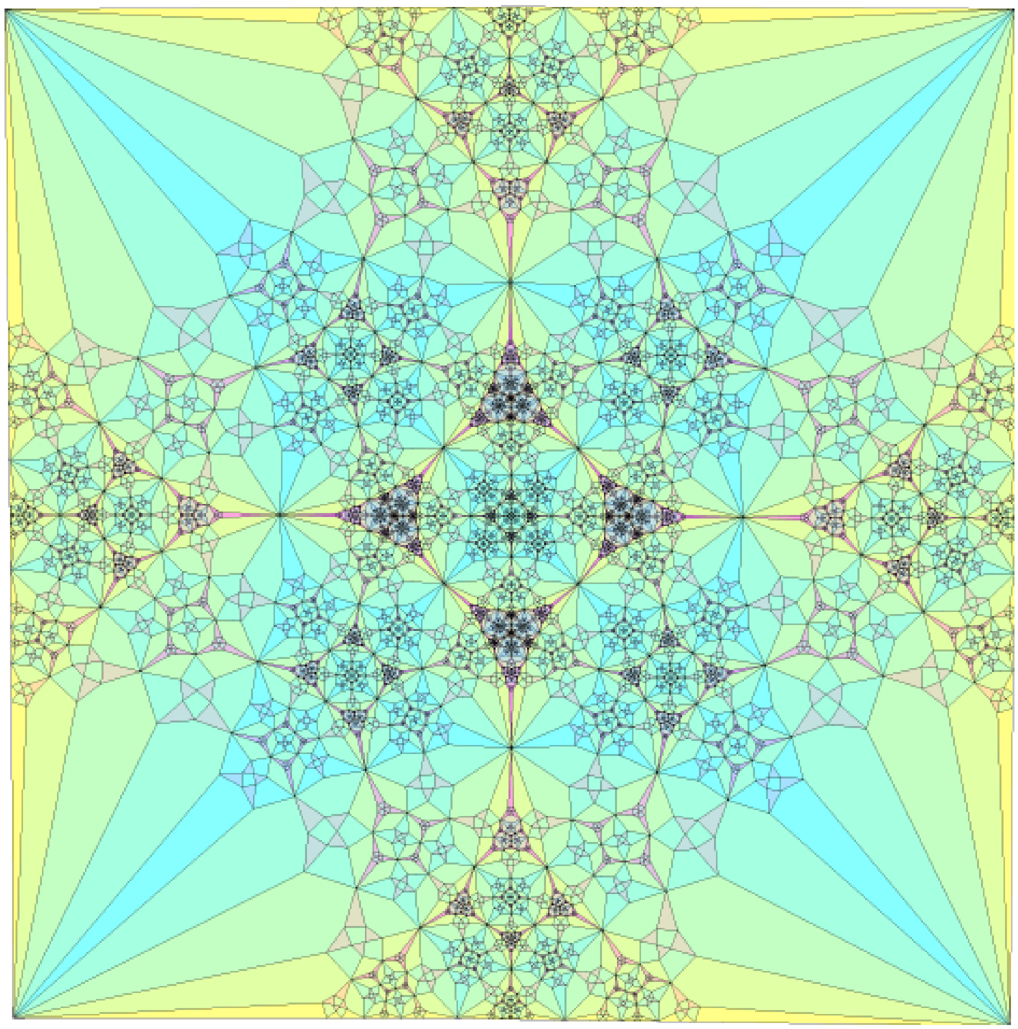

Blue

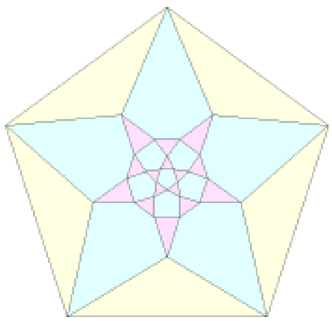

Red

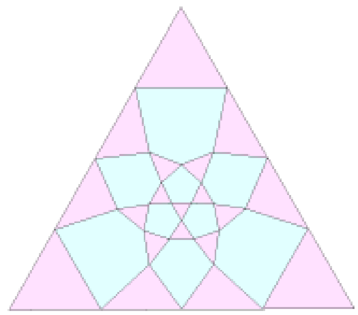

Yellow

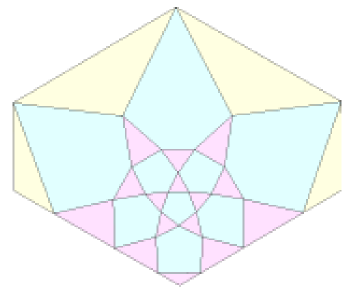

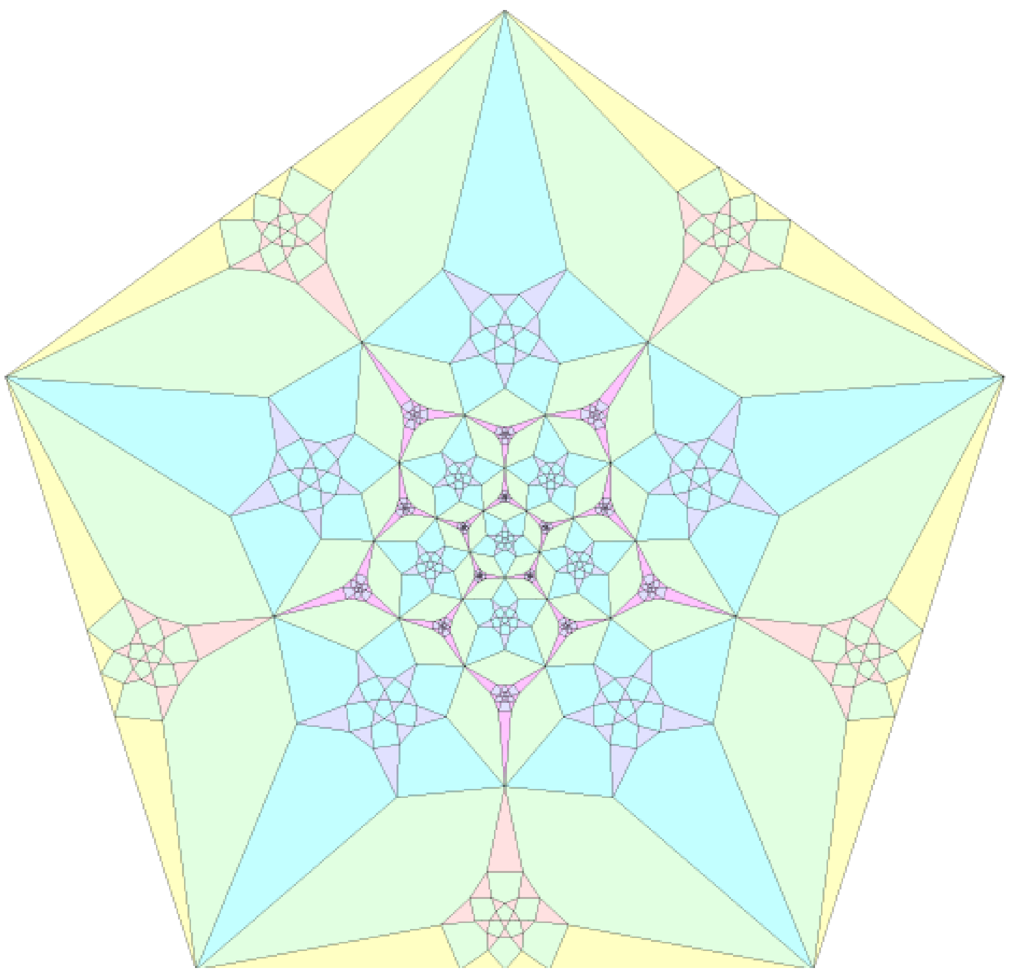

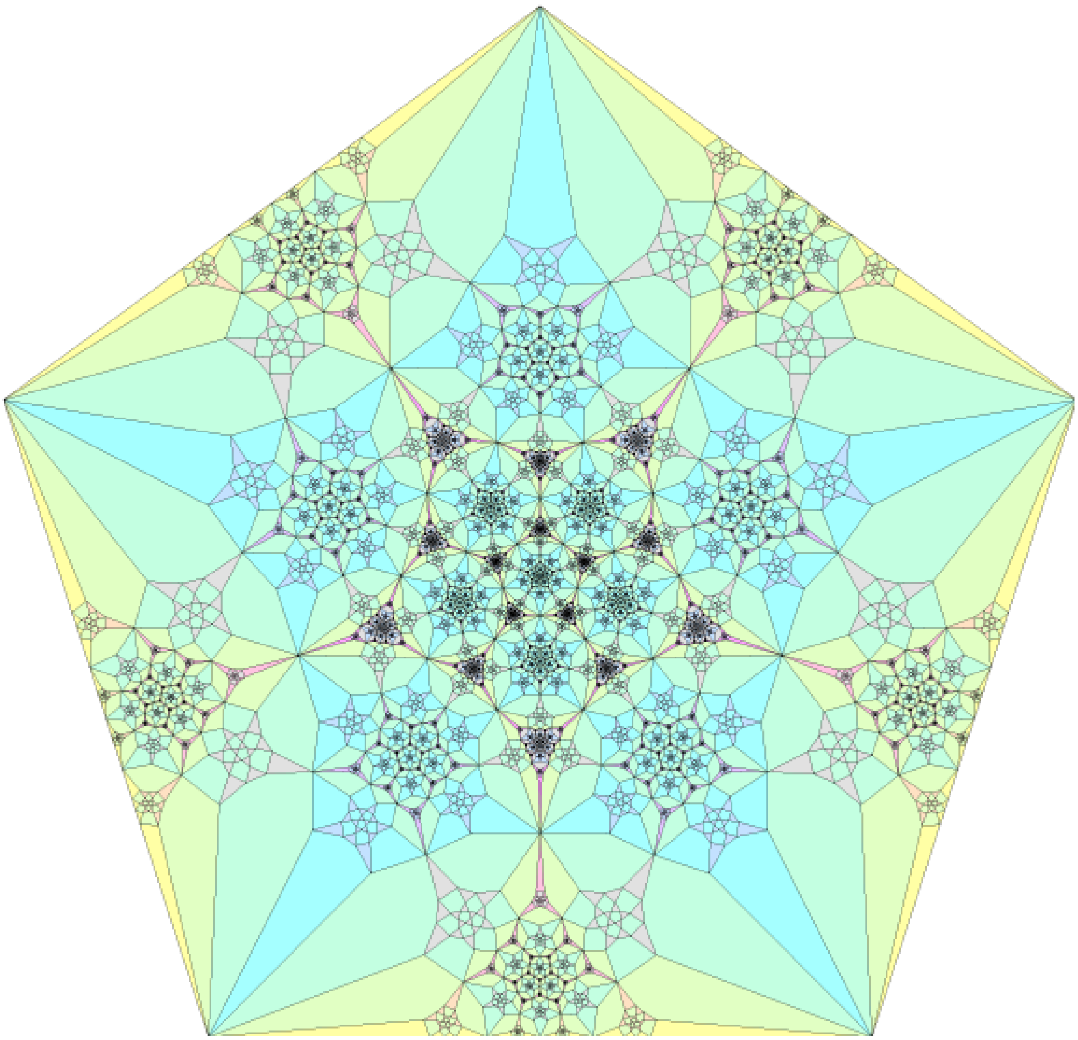

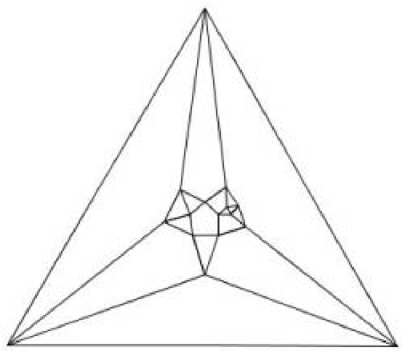

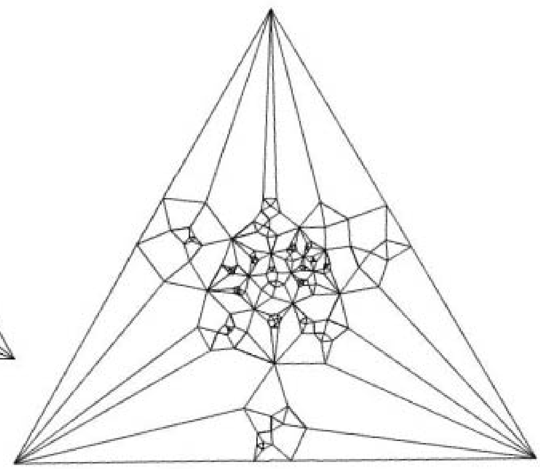

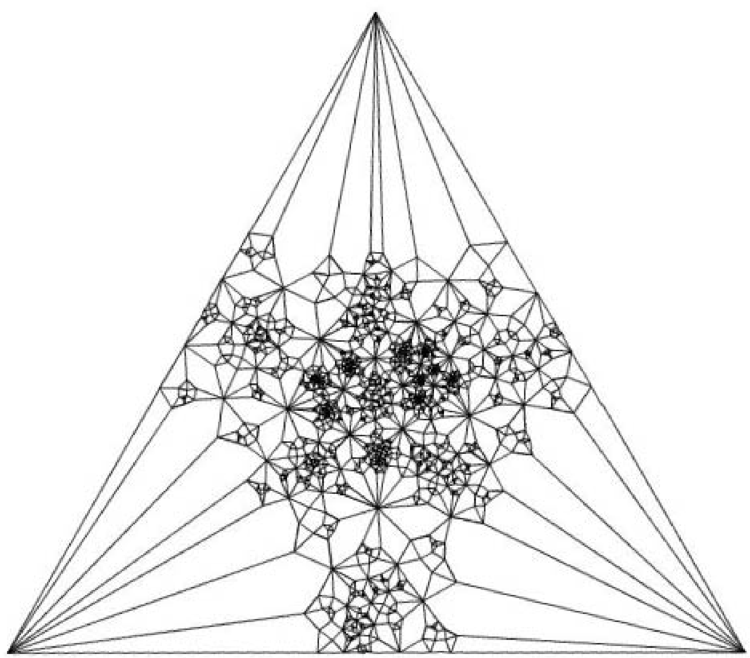

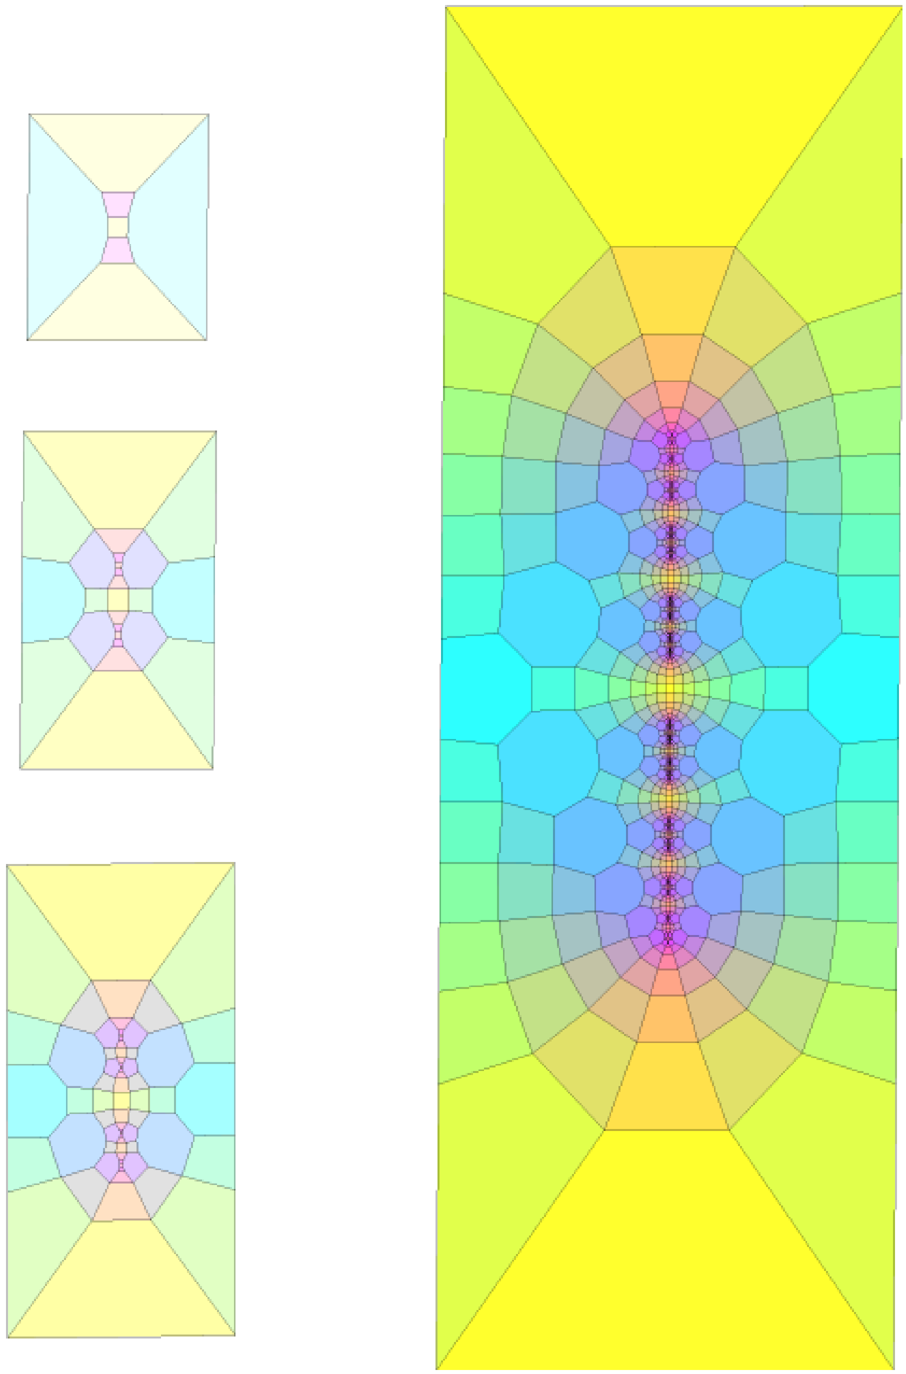

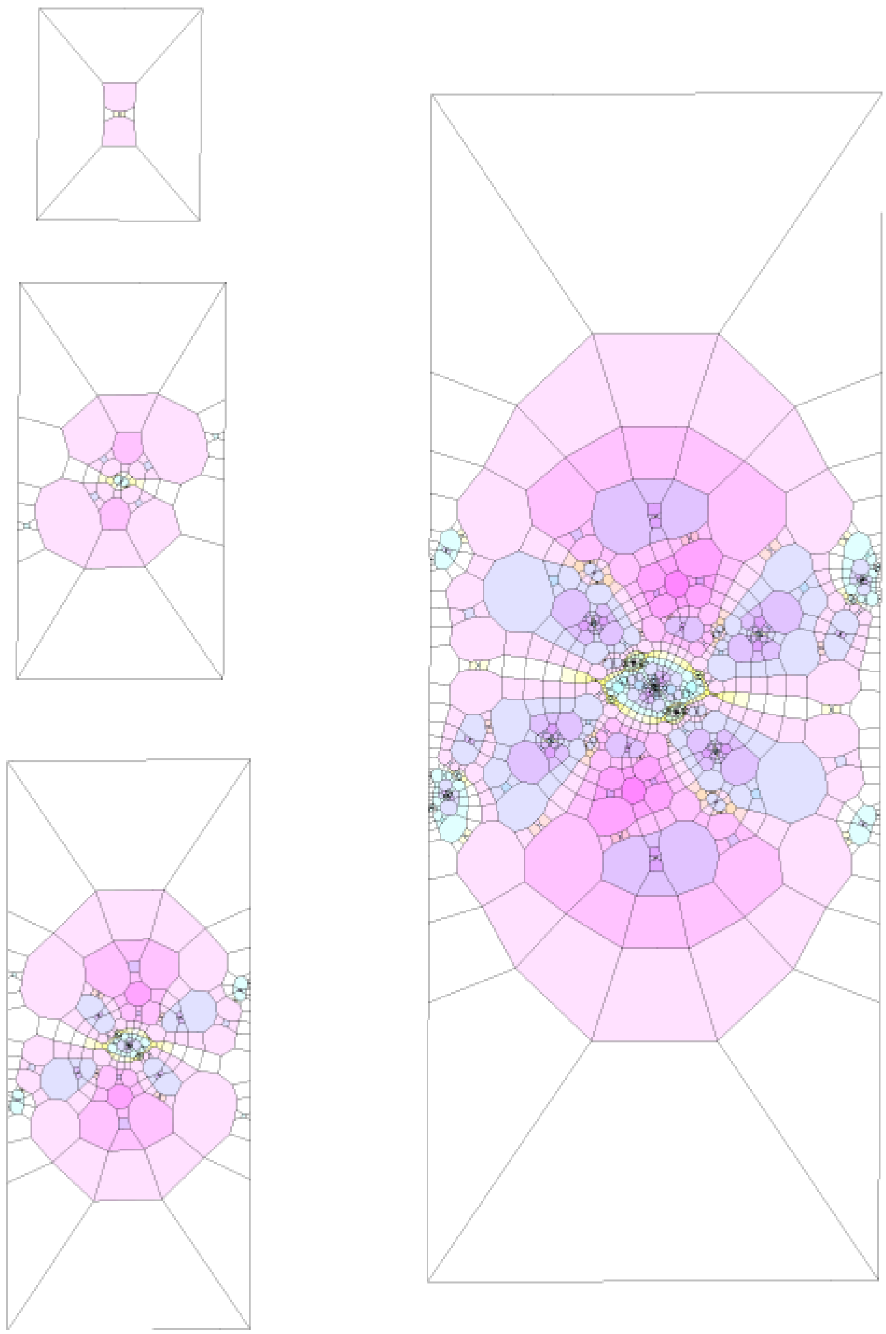

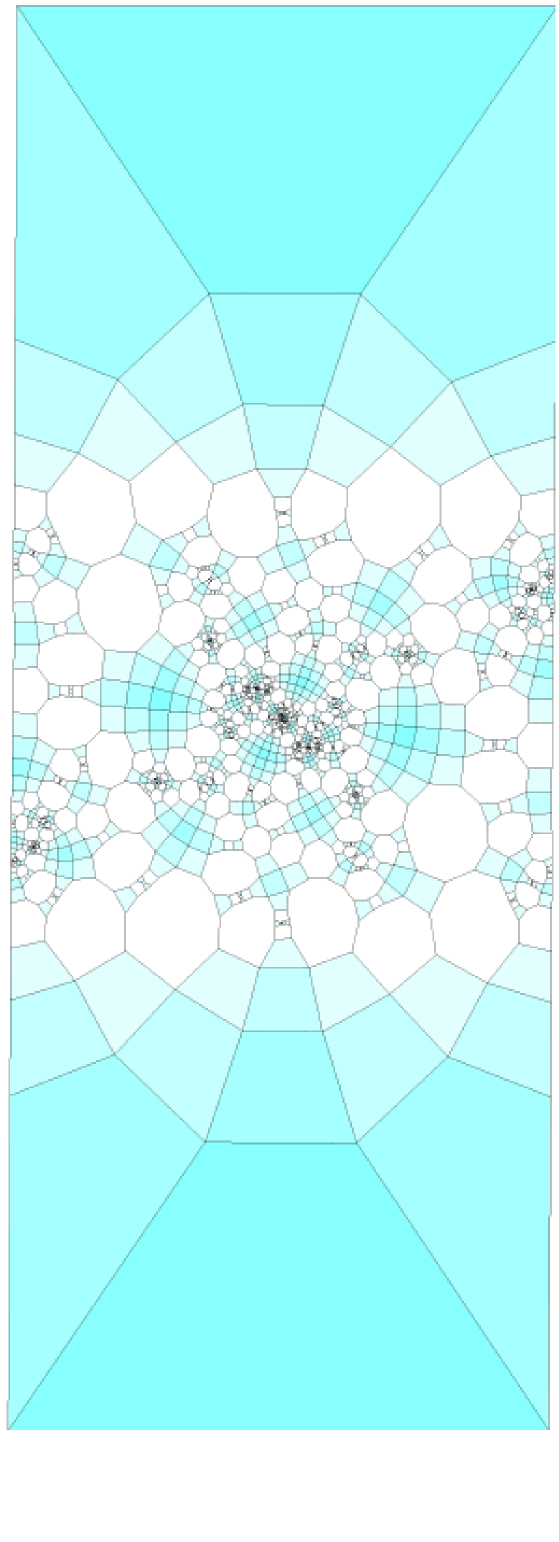

14 Answers to Exercises

Exercise. Draw the first stage of the construction of the universal cover of the Borromean rings. Pick one loaded pair and draw the second stage of the construction restricted to those two regions, just as we did above for the trefoil.

Solution. Look at the last part of Figure 61. This is the projection we begin with. Inside each triangle, draw another copy of this exact same projection. To subdivide the outside face, take this projection and enlarge it and rotate it by π so the center triangle overlaps with the original triangle, and paste it on.

The replacement rule of a loaded pair is as follows:

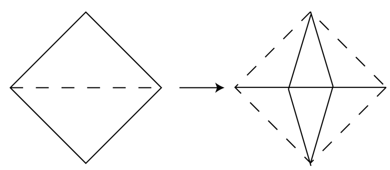

Exercise. Finish the proof by showing that the only alternating link projection with 4 edges and 2 vertices is the Hopf link. Show that the only homogeneous alternating link projection with 12 edges and 6 vertices is the Borromean rings.

Solution. Now, a link with 2 vertices and 4 edges where every region has two edges must be the Hopf link; each edge must go from one vertex to the other, since we aren't allowed singular regions with one edge. So following an edge from one vertex to another, we see it must return by another edge and form a loop. Since there are two pairs of edges, there are two loops, intersecting in two vertices, so we get the Hopf link.

In the case $n = 3$, we have 6 vertices and 8 faces, all of which are triangles. Draw a triangle (see Figure 61). Since every link projection has four edges at each vertex, draw them in, too, on the outside. Now, none of these new edges can be the

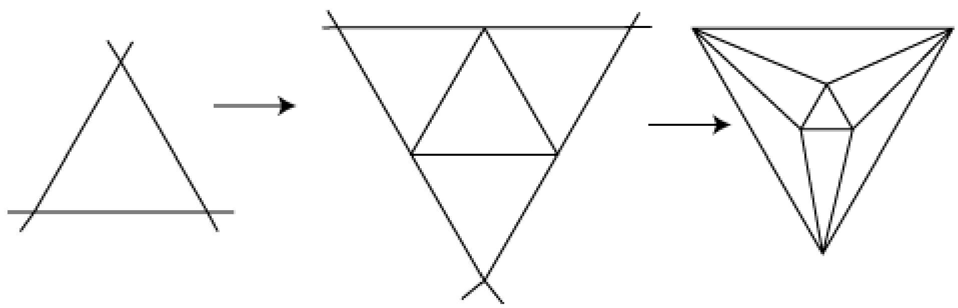

Figure 61: The only link where every region is a triangle is the Borromean rings

same, since that would create bigons. So on each edge of the original triangle, we have three distinct edges connecting in two vertices, so they must connect in a third vertex. Again, each of the three vertices thus created must be distinct, or we would create bigons. Looking at each pair of these vertices, we see that they are connected by two edges, which have a vertex in the middle with four edges already attached to it. Thus, we have three vertices that are connected already, so we must add an edge to the two outside ones and form another triangle. Three triangles are formed in this manner, and by this point, all vertices have the correct number of edges coming out and we have 6 vertices, 12 edges, and 8 faces, all of which are triangles. Since this is an alternating link, it is uniquely determined (up to reflection) by its graph. Thus, it is the Borromean rings.

Note that this is really the 2-d replacement rule process described in Section 9. Since every region is a triangle, the T's and P's are the same. We start with 3 T faces (of the original triangle). Gluing on three triangles gives three DT pairs; gluing on three more triangles leaves three white circles, so we glue a single triangle on to all three open faces and complete the graph.

Exercise. Find all the tiles used in the subdivision rule described above for the Hopf. Draw out the first three or four stages.

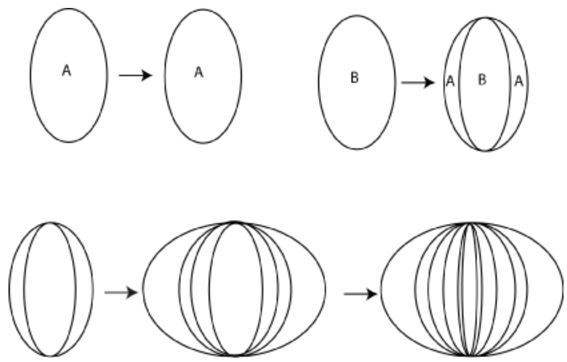

Figure 62: On the top are the tiles. The bottom shows the subdivisions of the link, where we start with two A tiles and two B tiles (one B tile is the outside face).

Solution. See Figure 62.

Exercise. Find the subdivision rule given the Borromean rings by the above theorem.

Solution. The subdivision rule for the Borromean rings is the one given in the Introduction (Section 1).

Exercise. Dispersed links must be prime. Complete the above proof by showing that the alternating projection obtained is prime.

Solution. If it is not prime, then there is some region that shares two edges with another region. It can't be two of the original regions, because then the old link would not have been prime. It can't be one original edge region and one of the R_i, since then R and the original region would have made the old link composite. Finally, R' and R_i share only one edge for each i. Thus, since the diagram is prime, the link is prime. It is clearly still non-split, since our new edges are connected to the old edges.

Exercise. Prove that the Hopf link cannot have a subdivision rule with only one tile type.

Solution. Assume there is onlye one tile type. We can make the edge subregions of a non-loaded bigon either both move out, both stay, or one move and one stay. The only one that would give a well-defined subdivision rule in the first stage of subdivision is if one stayed and one moved. So the subdivision rule would be to divide the bigon in half. But in loaded bigons in the second stage, the single new edge would have to move to the left or to the right to cover the stationary edges in non-loaded regions, and then the subdivision rule would be that the bigon doesn't divide at all. Thus, we have a contradiction.

Exercise. Show that this link can be obtained by inscribing and/or circumscribing four pentagons to a two-braid knot with ten vertices.

Solution. Take a ten-vertex two-braid and inscribe a pentagon in it by touching every other interior edge. Then inscribe a pentagon inside of that pentagon by connecting all the midpoints of the edges. Draw two more pentagons by doing the same thing to the exterior of all the bigons that have not been modified yet, and you have the big platonic link.

Note that it can also be obtained by adding a single closed loop to the (5,5)-radial link.

Exercise. Draw the first few stages of the subdivision rule for a two-braid link according to both methods described in this thesis. What similarities and differences do you notice?

Solution. If we begin with the standard projection of the link having p bigons, the two subdivisions are the same, except the one from Section 10 has two lines for every one line of the alternating link method.

References

[1] C. C. Adams. *The Knot Book*. American Mathematical Society, 2004.

[2] C. C. Adams. Noncompact fuchsian and quasi-fuchsian surfaces in hyperbolic 3manifolds. *Algebraic and Geometric Topology*, 7:565–582, 2007.

[3] C. Ashley. *Ashley Book of Knots*. Doubleday, 1944.

[4] M. Brazil. Calculating growth functions for groups using automata. In W. Bosma and A. J. Van Der Poorten, editors, *Computational Algebra and Number Theory*, pages 1–18. Springer-Verlag, 2004.

[5] J. W. Cannon. The combinatorial riemann mapping theorem. *Acta Mathematica*, 173(2):155–234, 1994.

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

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

[8] J. W. Cannon and E. L. Swenson. Recognizing constant curvature discrete groups in dimension 3. *Transactions of the American Mathematical Society*, 350(2):809–849, 1998.

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

[10] J.W. Cannon and Ph. Wagreich. Growth functions of surface groups. *Math. Ann.*, 293:239–257, 1992.

[11] D. Dunham. Hyperbolic celtic knot patterns. In *Conference Report: Bridges 2000: Mathematical Connections in Art, Music and Science*, pages 13–22, Winfield, Kansas, USA, July .

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

[13] P.J. Lu and et al. Decagonal and quasi-crystalline tilings in medieval islamic architecture. *Science*, 315:1106–1110, 2007.

[14] W. Menasco. Polyhedra representation of link complements. In S. J. Lomonaco, editor, *Low-dimensional Topology*, volume 20 of *Contemporary Mathematics*, pages 305–325. American Mathematical Society, 1983.

[15] W. Menasco. Closed incompressible surfaces in alternating knot and link complements. *Topology*, 23:37–44, 1984.

[16] V. V. Prasolov and A. B. Sossinsky. *Knots, Links, Braids and 3-Manifolds*. American Mathematical Society, 1991.

[17] P. Prusinkiewicz and A. Lindenmayer. *The Algorithmic Beauty of Plants*. Springer-Verlag, 1996.

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

[19] W. P. Thurston. *Three-Dimensional Geometry and Topology*, volume 1. Princeton University Press, 1997.