New York Journal of Mathematics

New York J. Math. **20** (2014) 399–429.

Cutting sequences on translation surfaces

Diana Davis

Abstract. We analyze the cutting sequences associated to geodesic flow on a large class of translation surfaces, including Bouw–Möller surfaces. We give a combinatorial rule that relates a cutting sequence corresponding to a given trajectory, to the cutting sequence corresponding to the image of that trajectory under the parabolic element of the Veech group. This extends previous work for regular polygon surfaces to a larger class of translation surfaces. We find that the combinatorial rule is the same as for regular polygon surfaces in about half of the cases, and different in the other half.

Contents

1. Introduction

We are interested in *translation surfaces*, which are created by identifying opposite congruent edges of polygons. This creates a flat cone surface, a compact surface having a locally Euclidean metric away from finitely many

Received February 20, 2014.

2010 *Mathematics Subject Classification.* 37E35.

Key words and phrases. Translation surfaces; Veech surfaces; geodesics; symbolic dynamics.

ISSN 1076-9803/2014

points; these points are conical singularities whose angles are integer multiples of 2π. Such surfaces often contain distinguished graphs: The surfaces are built by identifying the edges of a union of polygons, and the edges of the polygons form the graph.

The simplest example of a translation surface is the flat torus, where the left and right edges of a square are identified, and the top and bottom edges are identified, in the usual way. The torus has been extensively studied and is well understood [6], [7]. Recently, John Smillie and Corinna Ulcigrai studied the regular octagon surface, where the four pairs of parallel edges are identified [8], [9]. Subsequent work investigated the double pentagon surface, made by gluing the five pairs of parallel edges of a pair of regular pentagons [2], [3].

A translation surface always has a group of locally affine automorphisms, and this group has a natural homomorphism to a subgroup of $\mathrm{SL}_2(\mathbf{R})$. The surface is called *Veech* if the image of this homomorphism is a lattice. Even when the surface is not Veech, it has some interesting automorphisms.

We are interested in cutting sequences on translation surfaces. To generate a cutting sequence, one chooses a direction on a surface, and considers the straight-line flow in that direction. One records the labels of the edges hit by an infinite geodesic, which gives us a bi-infinite sequence of edge labels, called the *cutting sequence*. The purpose of this paper is to describe the effect of the parabolic element of the Veech group on cutting sequences on several families of translation surfaces.

One question of interest is to characterize all possible cutting sequences for a given translation surface. Applying the parabolic element of the Veech group can help us to do this, by generating new valid cutting sequences from a known cutting sequence. Cutting sequences are also interesting because on the square torus, they give a continued fraction expansion for the slope of the associated trajectory [7]. John Smillie and Corinna Ulcigrai found a construction for the regular octagon surface that gives something analogous to a continued fraction sequence ([8], Theorem 2.3.1).

When the image of an affine automorphism is a parabolic element of $\mathrm{SL}_2(\mathbf{R})$, the element is called a *shear*. We consider the composition of a vertical reflection with this shear, an action that we call the *flip-shear*. Our main result characterizes the relationship between a cutting sequence corresponding to a trajectory, and the cutting sequence corresponding to the image of that trajectory under the flip-shear.

In general, the shear performs a twist in each cylinder of a surface's cylinder decomposition (Definition 2.1). Each cylinder must be twisted a whole number of times for the shear to be an automorphism, so the magnitude of the twist must be an integer multiple of the modulus of each cylinder. In particular, all the cylinders must have commensurable moduli.

Square-tiled surfaces are the simplest examples of translation surfaces with commensurable moduli. It turns out (Lemmas 5.3 and 5.4) that the

cylinders of regular polygon surfaces also have commensurable moduli. John Smillie and Corinna Ulcigrai gave a rule for the effect of the flip-shear on cutting sequences on the regular octagon surface, and on all regular $2n$-gon surfaces [8], [9]. Using the same methods, we showed that the same rule holds for all double regular n-gon surfaces for odd n [3].

Irene Bouw and Martin Möller [1] gave algebraic models for a family of translation surfaces, which we now call *Bouw–Möller surfaces*, whose Veech groups are (m, n, ∞)-triangle groups for most m and n. This generalized earlier work of William Veech [10] and Clayton Ward [11].

Motivated by the goal of understanding Bouw and Möller's work, Pat Hooper [5] and independently Ronen Mukamel defined a doubly-indexed family of translation surfaces by more geometric means. In [5], Hooper described several inequivalent ways of presenting these surfaces, both in terms of grid graphs and in terms of polygon gluings. In particular, he gave a polygon decomposition for the (m, n)-indexed surface consisting of m polygons that each have $2n$ edges, which we study in §6−§7. In each of these surfaces, all the cylinders have the same modulus. Hooper proved that his surfaces are usually the same as those constructed by Bouw and Möller, with possible exceptions when $m = n$ or when m and n are both even. Later, Alex Wright [12] showed that the Bouw–Möller and Hooper surfaces are in fact the same for all m and n. In this paper, we give a rule (Theorem 7.1) for the effect of the flip-shear on cutting sequences corresponding to certain trajectories on Bouw–Möller surfaces.

Results presented in this paper. We characterize the effect of the flip-shear on certain trajectories on all *perfect translation surfaces* with *common modulus M* (Definition 2.18). These include many families of surfaces, such as regular polygon surfaces, Bouw–Möller surfaces, and many rectilinear surfaces.

The results for regular polygon surfaces were already known [3], [8], but the results for Bouw–Möller surfaces and rectilinear surfaces are new. Roughly speaking, given a cutting sequence corresponding to a geodesic trajectory on a perfect translation surface, and the new cutting sequence corresponding to image of that trajectory under the flip-shear, to get from the original sequence to the new sequence, we keep the sandwiched letters, and also keep certain other letters. The kept letters are sandwiched in about half of the cases (see Table 1).

In §2, we give the general definitions. In §3, we discuss the action of the flip-shear in detail. In §4, we prove the main result. In the remainder of the paper, we discuss how the result applies to the regular polygon surfaces (§5) and Bouw–Möller surfaces (§6−§7). For a discussion of rectilinear surfaces, including certain L-shaped tables, see ([4], Chapter 8).

Acknowledgements. Many thanks to my advisor, Richard Schwartz, generally for his excellent guidance throughout my time in graduate school, and

specifically for the multitude of insightful suggestions he made for improving this paper (a subset of my thesis). Corinna Ulcigrai, John Smillie and I spent an enlightening three days in Bristol, UK in February 2012. We had extensive discussions about many aspects of the Bouw–Möller project, and I benefited from their insights. Sergei Tabachnikov and Pat Hooper shared their questions and insights with me over the course of several years. Ronen Mukamel taught me about cylinder decompositions and moduli.

2. Definitions

2.1. Translation surfaces.

Definition 2.1. Let S be a translation surface. A *cylinder decomposition* of S is a partition of S into finitely many flat parallel cylinders. (A *flat* cylinder does not contain a singular point in its interior.) The interiors of these cylinders are pairwise disjoint, and their union is S. Figure 1 shows a cylinder decomposition for the regular octagon surface.

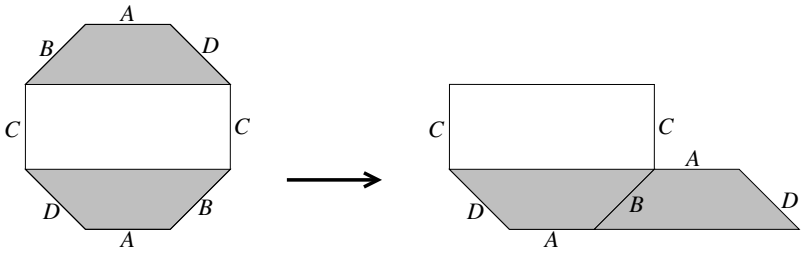

FIGURE 1. A cylinder decomposition for the regular octagon surface

Definition 2.2. A polygon is *level* if it never happens that there is a horizontal line that both contains a vertex of the polygon and crosses an edge of P at an interior point.

Definition 2.3. A polygon is *vertically symmetric* if it has a line of vertical symmetry. This means that we may translate the polygon in the plane so that the map $(x, y) \mapsto (-x, y)$ is a symmetry of the polygon.

Definition 2.4. A polygon is *special* if it is convex, level and vertically symmetric.

A translation surface has a *special decomposition* if it can be obtained by gluing together pairs of parallel edges of a finite union of special polygons. We require that the gluing is compatible with the vertical reflection symmetries of the polygons: Suppose that edge e_1 of polygon P_1 is glued to edge e_2 of polygon P_2. Let I_1 and I_2, respectively, be the vertical reflections preserving P_1 and P_2. Then $I_1(e_1)$ is glued to $I_2(e_2)$.

If a given translation surface has a special decomposition, we call it a *special translation surface*. Whenever we deal with a special translation surface, we will assume that it is given to us by way of a special decomposition.

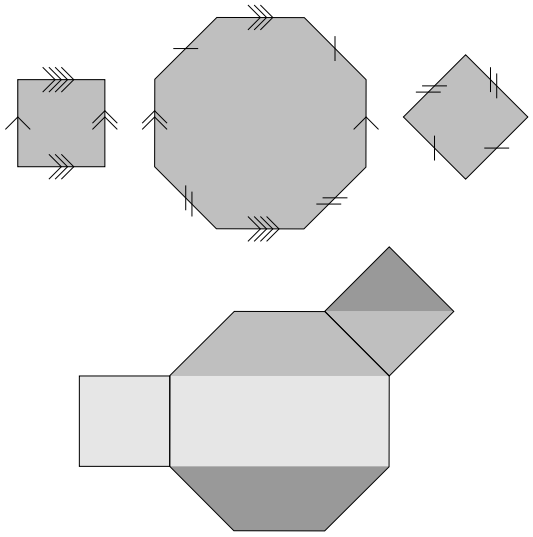

FIGURE 2. (a) A special decomposition, and (b) the horizontal cylinder decomposition, of the $(3, 4)$ Bouw–Möller surface

Figure 2 (a) shows a special decomposition of the $(3, 4)$ Bouw–Möller surface. It turns out that all Bouw–Möller surfaces have special decompositions (Lemma 6.5).

Lemma 2.5. *A special surface has a canonical cylinder decomposition, in which all the cylinders are horizontal.*

Proof. Consider the union of horizontal line segments connecting vertices of polygons in the special decomposition. These line segments piece together to give compact 1-manifolds without boundary, and hence circles. These circles are horizontal on the surface, and divide it into cylinders. This is clearly a cylinder partition. □

Figure 2 (b) shows how the special decomposition in Figure 2 (a) gives rise to a horizontal cylinder decomposition on the surface. Each cylinder is shaded a different color.

Definition 2.6. For us, an *isosceles trapezoid* is a vertically symmetric trapezoid or a vertically symmetric isosceles triangle. We think of the triangle as being a degenerate case of the trapezoid. The three main cases are a generic isosceles trapezoid, an isosceles triangle, and a rectangle.

Lemma 2.7. *Let S be a special translation surface. Assume that we have equipped S with its canonical cylinder decomposition. Each nonempty intersection of a cylinder with a special polygon is an isosceles trapezoid. Hence, each cylinder in the cylinder decomposition has a partition into isosceles trapezoids.*

Proof. Referring to the proof of Lemma 2.5, Figure 2 (b) shows how each region between two horizontal segments is an isosceles trapezoid (a rectangle, isosceles triangle or isosceles trapezoid). The result of the lemma is clear from this fact. □

Definition 2.8. A *level automorphism* of a translation surface is a locally affine automorphism that preserves the horizontal direction.

Figure 3 shows two examples of level automorphisms on the regular octagon surface. One is known as a *flip* and the other is known as a *shear*. The flip has order 2, and the shear has infinite order. The shear fixes each of the drawn horizontal line segments pointwise; it performs a Dehn twist in each horizontal cylinder. If R is the flip and S is the shear, then the middle picture in Figure 3 uses the notation $A' = R(A)$, etc. and the right picture uses the notation $A' = S(A)$, etc.

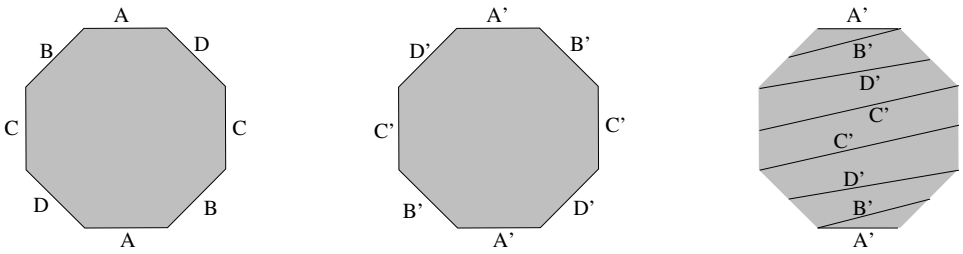

FIGURE 3. The octagon surface, and its images under the *flip* and the *shear*

Lemma 2.9. *Let S be a special translation surface. Then S admits an order-2 level automorphism.*

Proof. One can separately reflect each polygon in the special decomposition through its line of symmetry. The compatibility condition for the gluings guarantees that these separate maps fit together to give a global automorphism of S. □

Definition 2.10. Given a special translation surface, we call its order-2 affine element from Lemma 2.9 *the flip* and denote it by R.

We can use the existence of the flip R to give more information about the nature of the trapezoid decompositions of the cylinders in the horizontal cylinder decomposition.

Lemma 2.11. *Suppose that S is a special translation surface. Let $\mathscr{C}$ be one of the cylinders in the horizontal cylinder decomposition of S. Then $\mathscr{C}$ is partitioned into one or two trapezoids. In the case where $\mathscr{C}$ is partitioned into one trapezoid, this trapezoid is a rectangle.*

Proof. For the first claim, note that the flip R preserves each cylinder in the cylinder decomposition. The restriction $R|_{\mathscr{C}}$ is an orientation-reversing isometry of $\mathscr{C}$. Any such isometry fixes exactly two vertical line segments in $\mathscr{C}$. These two segments are diametrically opposed. At the same time, by construction, R fixes the vertical segment of symmetry of each trapezoid in the decomposition. Thus, there can be at most two such trapezoids.

For the second claim, suppose that $\mathscr{C}$ is decomposed into a single trapezoid. This means that opposite sides of the trapezoid are glued together. Thus, by definition, they are parallel, so the trapezoid is a rectangle. □

So far, we have exhibited one level automorphism on a special translation surface. We will exhibit two more. First, we need several more definitions.

Definition 2.12. The *modulus* of a horizontal cylinder is the ratio $\frac{W}{H}$. Here H is the height of the cylinder and W is the length of the horizontal translation needed to identify the opposite edges of a fundamental domain for the cylinder.

Definition 2.13. Let S be a special translation surface equipped with its canonical horizontal cylinder decomposition. We call the cylinder decomposition *perfect* if there is some number M such that the following is true of each cylinder $\mathscr{C}$ of the decomposition:

- $\mathscr{C}$ has modulus M and is decomposed into two isosceles trapezoids; or
- $\mathscr{C}$ has modulus $M/2$ and is decomposed into a single rectangle.

We call M the *common modulus* of the decomposition. If the first case occurs, we call $\mathscr{C}$ *typical*. If the second option occurs, we call $\mathscr{C}$ *exceptional*.

Remark 2.14. It might seem more natural to require that, in the exceptional case, the rectangle have modulus M. However, when we take the double cover of the cylinder in this case, we get a cylinder of modulus M that is decomposed into two trapezoids (rectangles). This turns out to be more natural for our purposes.

Lemma 2.15. *Let S be a special translation surface with a perfect horizontal cylinder decomposition of common modulus M. Then S admits an infinite-order level automorphism that fixes the tops and bottoms of the cylinders pointwise, and has derivative*

$$\begin{bmatrix} 1 & M \\ 0 & 1 \end{bmatrix}.$$

Proof. In each cylinder, there is a unique affine automorphism that has the desired effect and has the desired derivative. These automorphisms piece together to give the global automorphisms of the surface. □

Definition 2.16. Let S be a special translation surface whose horizontal cylinder decomposition is perfect and has common modulus M. We call the infinite-order level automorphism from Lemma 2.15 *the shear*, and we denote it by S_M.

Definition 2.17. Let S be a special translation surface whose horizontal cylinder decomposition is perfect and has common modulus M. (This guarantees the existence of R and S_M.) Then we define

$$V_M = S_M \circ R.$$

We call V_M the *flip-shear*. Corresponding to these elements, we have the composition of derivatives

$$V_M' = S_M' \circ R'.$$

In terms of matrices, this is

$$\begin{bmatrix} 1 & -M \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & M \\ 0 & 1 \end{bmatrix} \circ \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Definition 2.18. A *perfect translation surface* is a special translation surface whose cylinder decomposition is perfect.

A perfect translation surface admits the level automorphisms R, S_M and V_M, where M is the common modulus of the cylinder decomposition. In the next section, we will analyze how these elements act on certain distinguished arcs in a perfect translation surface.

Remark 2.19. All regular polygon surfaces (with a horizontal edge), all Bouw–Möller surfaces, and certain square-tiled surfaces are perfect translation surfaces. For regular polygon surfaces, the common modulus is $M = 2\cot \pi/n$ (Lemmas 5.3, 5.4). For the (m, n) Bouw–Möller surface, the common modulus is $M = 2\cot \pi/n + 2\frac{\cos \pi/m}{\sin \pi/n}$ (see Lemma 6.6 or [4], Lemma 6.6).

2.2. Trajectories and cutting sequences. The purpose of this paper is to describe the action of V_M on cutting sequences associated to geodesic trajectories on perfect translation surfaces. We analyze V_M because V_M, rotations, and reflections generate all the symmetries of the surface. The actions of rotations and reflections are easy to describe, but the action of V_M takes more work.

Since $V_M = S_M \circ R$, it is also true that S_M, rotations and reflections generate the group of symmetries of the surface. So instead of describing the action of V_M, we could choose to describe the action of S_M. However, the action of V_M turns out to be more elegant than that of S_M: The rule

for the effect of S_M is essentially "apply the rule given in this paper for V_M, and then permute the edge labels in a certain way." So in this paper we give the rule for V_M, and from there it is easy to deduce the effect of S_M.

If we know something about a certain trajectory, we can apply the symmetries of the surface to learn something about other trajectories on the surface. Ideally, we would be able to analyze all the trajectories in some set, and also be able to use the symmetries of the surface to transform all possible trajectories into the set we can analyze, thereby analyzing all possible trajectories. This is possible for regular polygon surfaces, because we can analyze trajectories with $\theta \in [0, \pi/n)$ and then apply the n-fold rotational symmetry of the surface to rotate any trajectory into that sector. This is not possible for Bouw–Möller surfaces using our methods, because the (m, n) Bouw–Möller surface has n-fold rotational symmetry, and yet we can only analyze trajectories with $\theta \in [0, \theta_S)$ (see Definition 3.2), where θ_S is at most $\pi/(2n)$. So we cannot analyze trajectories with $\theta \in [\pi/(2n), \pi/n)$, which is a limitation of the method. We hope to exploit the surfaces' other symmetries in the future to analyze all possible trajectories.

The main Theorem 4.2 proves the result for perfect translation surfaces in general, and the rest of the paper applies the result to several families of perfect translation surfaces: regular polygon surfaces (§5) and Bouw–Möller surfaces (§6–§7). For an application to rectilinear surfaces, see ([4], Chapter 8).

We will now define some notation for geodesic trajectories and their associated cutting sequences.

Definition 2.20. Given a geodesic trajectory on a perfect translation surface, the *cutting sequence* associated to the trajectory is the bi-infinite sequence of polygon edges that the trajectory crosses. (By convention, if a trajectory hits a vertex, it stops. In that case, the associated cutting sequence is not bi-infinite.) The *derived trajectory* is the image of the trajectory under the flip-shear V_M. The *derived sequence* is the cutting sequence associated to the derived trajectory.

Lemma 2.21. *$V_M^{-1} = V_M$, i.e., V_M is an involution.*

Proof. By definition, $V_M = S_M \circ R$. So

$$V_M \circ V_M = S_M \circ R \circ S_M \circ R = S_M \circ (R \circ S_M \circ R) = S_M \circ (R \circ S_M \circ R^{-1}).$$

We have $R = R^{-1}$ because R is a reflection, and thus has order 2. Also, $R \circ S_M R^{-1} = S_M^{-1}$, because a reflection composed with a perpendicular shear composed with a reflection is a shear in the opposite direction. Hence,

$$V_M \circ V_M = S_M \circ (R \circ S_M \circ R^{-1}) = S_M \circ S_M^{-1} = \mathrm{Id}. \qquad \square$$

We use the following notations only in Lemma 2.22: Let τ be a geodesic trajectory on a perfect translation surface with common modulus M, and let $V_M\tau$ be the image of τ under V_M, i.e. the derived trajectory. Let

$L = \{A, B, \ldots\}$ be the set of polygon edges, and let $L' = \{A', B', \ldots\}$ be their transformed images. Let $c(\tau)$ be the cutting sequence associated to τ, i.e. the sequence of polygon edges that τ crosses, so $c(\tau) \in L^{\mathbb{Z}}$. Let $c'(\tau)$ be the sequence of transformed edges that τ crosses, so $c'(\tau) \in L'^{\mathbb{Z}}$. Let $c(V_M\tau)$ be the cutting sequence associated to $V_M\tau$, i.e. the *derived* sequence.

Lemma 2.22. $c(V_M\tau) = c'(\tau)$.

In words: Determining the derived cutting sequence for a given trajectory is equivalent to determining which sheared edges the trajectory crosses.

Proof. Consider $V_M^{-1}\tau$, the transformed trajectory, and the sequence of polygon edges $c(V_M\tau)$ that it crosses. When we apply V_M, it acts on both the trajectory and on the polygon edges, so the intersection points are preserved. By Lemma 2.21, the image of $V_M\tau$ under V_M is $V_M(V_M^{-1}\tau) = \tau$, and the images of the polygon edges under V_M are the transformed edges, so the intersection points are then the intersection between the trajectory τ and the transformed edges. Symbolically,

$$c(V_M\tau) = c(V_M^{-1}\tau) = V_M^{-1}\tau \cap L = V_M(V_M^{-1}\tau) \cap V_M(L) = \tau \cap L' = c'(\tau). \ \square$$

Given a trajectory, the derived sequence is the cutting sequence corresponding to the derived trajectory, which is the image of the original trajectory under V_M. The main Theorem 4.2, and its applications Theorems 5.6 and 7.1 to various families of surfaces, give a combinatorial rule for determining the derived sequence from the original cutting sequence. It is useful to introduce the following notation to keep track of various types of edge crossings, which we use in the statement and proof of the main Theorem 4.2.

Definition 2.23. On a perfect translation surface S, a trajectory with $\theta \in [0, \theta_S)$ cuts through an edge on the left side (which includes the bottom edge) of a polygon and then an edge on the right side (which includes the top edge) of a given polygon. If the left and right edges are on the same level, we call it type (0). If the right edge is a level above, we call it type (1). Each three-letter sequence in the cutting sequence corresponds to crossing two polygons, which give us four cases (see Figure 4): $(00), (10), (01)$, and (11).

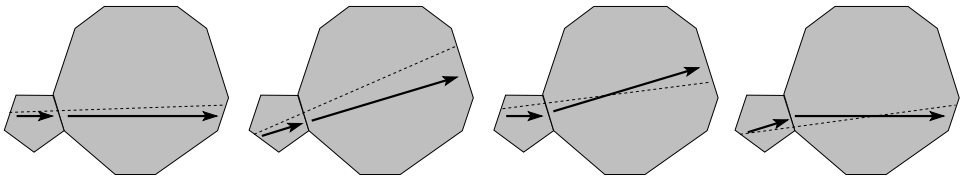

FIGURE 4. Examples of $(00), (11), (01)$, and (10) cases, with example trajectories (dotted) for each.

3. The action on edges

Definition 3.1. Let S be a perfect translation surface. A *gluing edge* of S is an arc in S that is the image of a pair of polygon edges under the identification map. Within each cylinder, the gluing edges are the non-horizontal edges of the trapezoids in the trapezoid partition. Figure 5 shows a typical cylinder with the gluing edges labeled A and B. In an exceptional cylinder that is decomposed into a single rectangle, there is just one gluing edge.

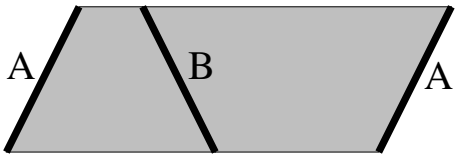

FIGURE 5. A typical cylinder with its gluing edges A and B

Definition 3.2. The *horizontal edges* are the tops and bottoms of the cylinders in the horizontal cylinder decomposition.

The *slanted edges* are the diagonals of positive slope in the trapezoids of the trapezoid decomposition. In terms of the special polygons, each slanted edge joins a vertex on the left side of a polygon to the vertex on the right that is one level above.

In a perfect translation surface S, we let θ_S be the smallest angle between any slanted edge and the positive horizontal. An example of slanted edges, their associated angles, and θ_S is in Figure 6.

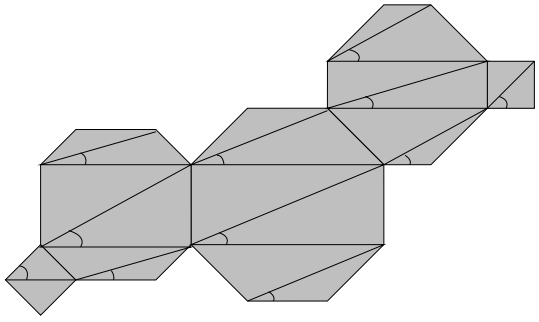

FIGURE 6. For the $(5, 4)$ Bouw–Möller surface shown, θ_S is the minimum of the marked angles.

The main Theorem 4.2 analyzes a geodesic on a perfect translation surface, where the geodesic's angle is between 0 and θ_S.

We are interested in the action of V_M on the edges within a cylinder. First, we will describe what happens for a typical cylinder. The degenerate

case (when one or more of the trapezoids degenerates to a triangle) is a limiting case of the typical case. The exceptional case (when the cylinder is one rectangle) is a special case of the typical case, using a double cover of the exceptional cylinder.

Figure 7 shows a portion of the universal cover of the cylinder $\mathscr{C}$. The universal cover itself is an infinite strip. The cylinder $\mathscr{C}$ is the quotient of the strip by the minimal label-preserving horizontal transformation, which is the generator of the deck transformation group.

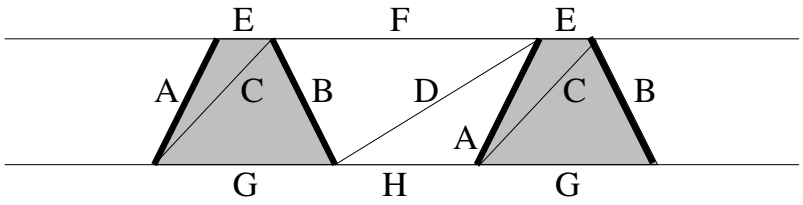

FIGURE 7. Part of the universal cover of a typical cylinder

We will persistently abuse terminology by saying things like "The edges labeled A and B in Figure 7 are the gluing edges," when we really mean that these edges are lifts to the universal cover of the gluing edges.

In Figure 7, $\mathscr{C}$ is decomposed into two trapezoids, one shaded and one unshaded:

- Edges A and B are the gluing edges.
- Edges C and D are the slanted edges.
- Edges E, F, G and H are the horizontal edges.

Figure 8 (a) shows the action of R on our cylinder. The action of R *does not* lift to an action on the universal cover, so here we are moving the objects in the cylinder by R and then lifting to the universal cover. In Figure 8, $a = R(A), b = R(B)$, and so on.

In the cylinder, R acts as a symmetry on each of the trapezoids, swapping the left and right boundaries of each. The effect from Figure 7 to Figure 8 (a) on the non-horizontal trapezoid boundaries is that the edge labels A and B have been swapped. R preserves each top and bottom boundary of each trapezoid. The effect from Figure 7 to Figure 8 (a) on the top and bottom boundaries is that the edge labels E, F, G and H are preserved.

Figure 8 (b) shows the action of V_M on $\mathscr{C}$, as depicted in the universal cover. As was the case with R, V_M does not lift to an action on the universal cover. Instead, there are two equivalent descriptions of what Figure 8 (b) shows:

- We let V_M act on A, B, C, D, E, F, G, H in the cylinder $\mathscr{C}$ (Figure 7), and then lift their images to the universal cover.
- We lift the action of S_M to the universal cover and let it act on the picture in Figure 8 (a). On the universal cover, S_M is a shear, fixing

the bottom boundary and shifting the top boundary to the right by a full modulus distance, indicated by the horizontal arrow.

Since $V_M = R \circ S_M$, these produce the same picture.

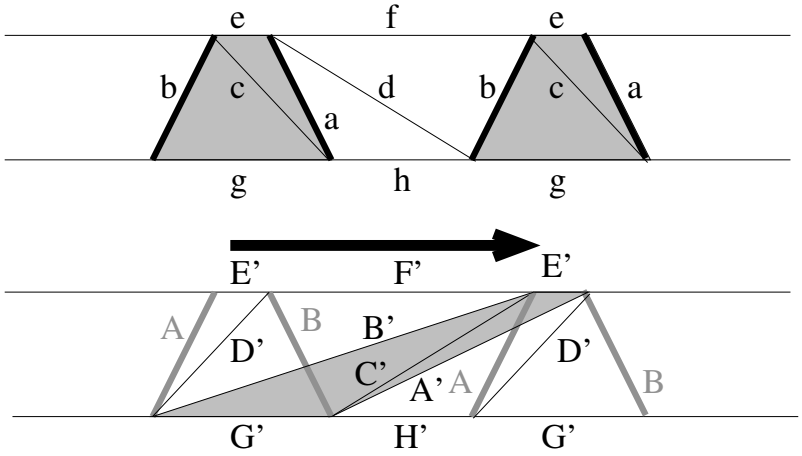

FIGURE 8. The action of (a) R and (b) V_M, depicted in the universal cover

In Figure 8 (b), we have used the notation

$$A' = S_M(a) = S_M \circ R(A) = V_M(A),$$

and so on.

The original gluing edges A and B are shown in a lighter color. This is important, because the proof of the main Theorem 4.2 analyzes the configuration of an edge with respect to its image under V_M. Here, we can observe that A' intersects A, and B' intersects B.

The shaded trapezoid in Figure 8 is (a lift of) $V_M(T)$, where T is the shaded trapezoid in the decomposition of $\mathscr{C}$. This trapezoid is not particularly interesting. More interesting are the *slanted trapezoids*:

Definition 3.3. We partition each cylinder into two *slanted trapezoids*, each of which is the region between two slanted edges.

Figure 9 shows a typical cylinder, depicted in the universal cover. It has two gluing edges (A and B), so it has two slanted edges (C and D), so it has two slanted trapezoids ($CFDG$ and $DECH$). This is a copy of Figure 7, with different trapezoids shaded.

We will use slanted trapezoids extensively in proving the main theorem in §4. In particular, because we can partition a perfect translation surface into cylinders, and each cylinder into slanted trapezoids, we can partition any perfect translation surface into slanted trapezoids. We will use this partition to partition a geodesic trajectory into finite segments, whose behavior we can easily analyze.

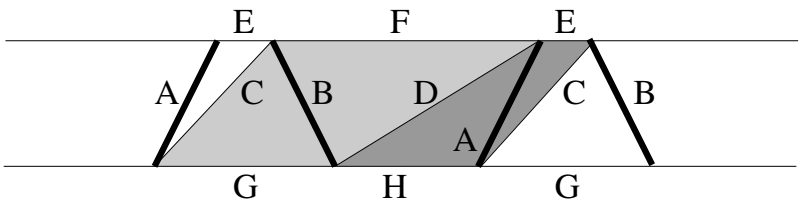

FIGURE 9. A typical cylinder, partitioned into two slanted trapezoids (shaded), depicted in the universal cover

The action is similar in a typical cylinder where one or both of the trapezoids degenerates to a triangle, or where one of the trapezoids degenerates to a slanted edge.

In an exceptional cylinder, there is only one gluing edge A, rather than two gluing edges A and B, and the gluing edge is vertical. There is only one slanted edge C. We use two copies of the rectangle (one shaded and one unshaded, as in Figure 8), which form a double cover of the cylinder, and then the rest of the construction is the same. For expanded details, see ([4], Chapter 3).

4. The main result

Recall (Definition 2.20) that the *cutting sequence* associated to a geodesic trajectory on a translation surface is the bi-infinite sequence of polygon edges that the trajectory crosses. The *derived trajectory* is the image of the trajectory under the flip-shear V_M, and the *derived sequence* is the cutting sequence associated to the derived trajectory. The purpose of this paper is to determine the relationship between the original cutting sequence and the derived sequence for perfect translation surfaces. Theorem 4.2 gives the relationship when $\theta \in [0, \theta_S)$: The derived sequence is a subsequence of the original cutting sequence, obtained by keeping certain letters and removing the rest. It turns out that to determine if a given letter survives into the derived sequence, we only need to look at the letters that precede and follow it, i.e. at the three-letter *word* for which it is the middle letter:

Definition 4.1. Recalling Definition 2.23, a *(11) word* is a three-letter sequence corresponding to a trajectory cutting through three edges that form a (11) case. A *(00) word*, a *(01) word* and a *(10) word* are defined analogously.

Theorem 4.2 (Main Theorem)**.** *Let S be a perfect translation surface with common modulus M.*

Consider a geodesic trajectory on the surface whose angle θ satisfies $0 \leq \theta < \theta_S$, and its associated cutting sequence. Mark a letter in the cutting sequence if:

- *It corresponds to a gluing edge, and it is the middle letter of a* (00) *or* (11) *word, or*

- *it corresponds to a horizontal edge.*

Then the derived sequence consists precisely of the marked letters, read from left to right.

Outline of Proof. Lemma 2.22 shows that determining the derived cutting sequence for a given trajectory is equivalent to determining which sheared edges the trajectory crosses. To determine which sheared edges a trajectory crosses, we partition the trajectory into disjoint segments whose union is the entire trajectory: We partition the translation surface into finitely many disjoint slanted trapezoids (Definition 3.3), and each finite segment of the trajectory that we analyze is the part of the trajectory that lies in a particular slanted trapezoid. Then we analyze what happens on each of those segments, and then pool the information to determine what happens for the entire infinite trajectory.

To determine the derived sequence from the original cutting sequence, the key question is whether a line segment that passes through edges $E_1E_2E_3$ must cross edge $E_2' = V_M(E_2)$, or not. This is a topological question, based on the location of edge E_2' in relation to the edges E_1, E_2, E_3. The location of the segment E_2' depends on the configuration of the cylinder for which the edge E_2 is a gluing edge: whether it is typical, degenerate or exceptional. Lemma 4.3 and its Corollary 4.4 are simple topological observations about configurations of line segments in quadrilaterals and trapezoids. Lemmas 4.5-4.8 apply these observations to each of the disjoint segments of the trajectory.

The local analysis allows us to locally determine the effect of V_M on each segment, and thus the effect on its associated three-letter word in the cutting sequence, telling us if the middle letter survives into the derived sequence. We then move across the sequence from left to right, looking at the middle letter in each overlapping three-letter word and marking those indicated in Theorem 4.2. After we have done this, the derived sequence is the marked letters, read from left to right.

Lemma 4.3. *Consider a convex quadrilateral with its two diagonals, and a trajectory that intersects the quadrilateral (Figure* 10 *(a)). The trajectory crosses both diagonals if and only if the edges it crosses are non-adjacent. In the adjacent case, a diagonal is crossed if and only if it shares the common vertex of the two edges.*

Proof. Trivial. □

Corollary 4.4. Let the the *top, left, bottom* and *right* edges be as in Figure 10 (b), and define the *positive* and *negative* diagonals according to whether their slope is positive or negative when the vertices of the trapezoid are horizontally shifted to transform the trapezoid into a rectangle. Then on a trapezoid, a trajectory passing from the left edge to the top edge, the left

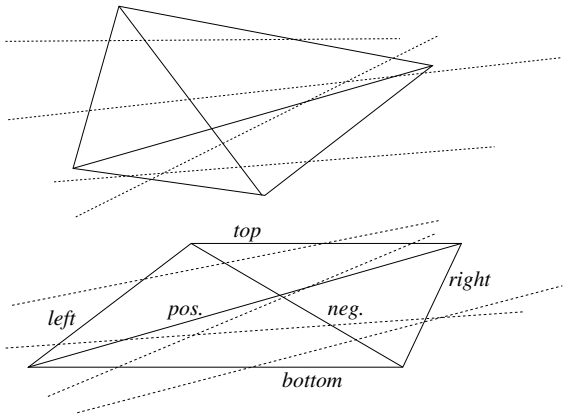

FIGURE 10. (a) A convex quadrilateral, and (b) a trapezoid, with trajectories intersecting their diagonals

edge to the right edge, the bottom edge to the top edge, or the bottom edge to the right edge crosses the negative diagonal. Of these, only trajectories passing from the left edge to the right edge, or from the bottom edge to the top edge, cross the positive diagonal.

The key question, to determine the derived sequence from the original cutting sequence, is whether a trajectory that crosses edges $E_1E_2E_3$ crosses $E_2' = V_M(E_2)$. Lemmas 4.5 and 4.6 establish that this happens if and only if the word $E_1E_2E_3$ is in a (00) or (11) case. Each lemma proves one direction.

Lemma 4.5. *Assume that E_2 is the gluing edge for a typical cylinder. Consider the word $E_1E_2E_3$ and the corresponding portion of the trajectory that lies between the intersections with edges E_1 and E_3. Consider the shorter segment s that is the intersection of that segment with the slanted trapezoid that has E_2 as a negative diagonal. If s crosses E_2', then $E_1E_2E_3$ is a* (00) *or* (11) *word.*

Proof. Refer to Figure 11. The hypothesis states that s crosses the shaded trapezoid's positive diagonal B'. By Corollary 4.4, there are two ways for this to occur: Either s passes from the left edge C to the right edge D, or from the bottom edge G to the top edge F. From this, we wish to recover the information about which of the gluing edges, which are the bold edges I, K, A, B, L, N, are hit by the extension of s to a trajectory. This will determine which case the word is in.

Suppose that s passes from the left edge to the right edge of the shaded trapezoid, so from edge C to D. The angle restriction $\theta \in [0, \theta_S)$ requires that the trajectory is less steep than all of the slanted edges. So a trajectory intersecting edge C must hit edge A just before it hits C, because the angle restriction prevents its coming from edge L. By Corollary 4.4, a trajectory intersecting edge C and then edge D must intersect edge B in between.

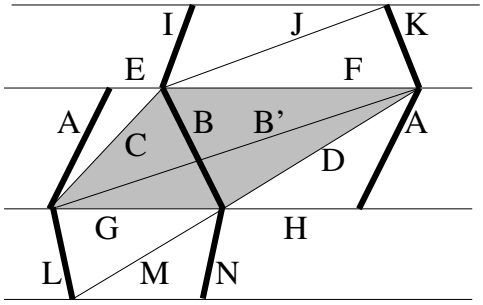

FIGURE 11. A typical cylinder with a slanted trapezoid shaded, and parts of two adjacent cylinders

Similarly, because of the angle restriction, a trajectory intersecting edge D must then intersect edge A, because it is not steep enough to intersect edge K. So a trajectory passing from edge C to D must cut through gluing edges ABA. This is a (00) case.

Suppose that s passes from the bottom edge to the top edge of the shaded trapezoid, so from edge G to edge F. Because the trajectory must be less steep than slanted edge M, it must intersect L just before edge G. By Corollary 4.4, between its intersections with edges G and F, it must pass through edge B. Again, because the trajectory must be less steep than the slanted edge J, it must intersect K just after F. So a trajectory passing from edge G to F must cut through gluing edges LBK. This is a (11) case. □

Lemma 4.6. *Assume that E_2 is the gluing edge for a typical cylinder. Consider the word $E_1E_2E_3$ and the corresponding portion of the trajectory that lies between the intersections with edges E_1 and E_3. Consider the shorter segment s that is the intersection of that segment with the slanted trapezoid that has E_2 as a negative diagonal. If $E_1E_2E_3$ is a* (00) *or* (11) *word, then s crosses E_2'.*

Proof. Suppose that $E_1E_2E_3$ is a (00) word. Referring to Figure 11, this means that the corresponding trajectory passes from edge A to B to A. In so doing, s crosses edge C and then D. This means that s goes from the left edge of the slanted trapezoid to the right edge, so by Corollary 4.4, s intersects the positive diagonal B' of the slanted trapezoid, which is E_2' as desired.

Now suppose that $E_1E_2E_3$ is a (11) word. In Figure 11, this means that the corresponding trajectory passes from edge L to B to K. In so doing, s crosses edge G and then F. This means it goes from the bottom edge of the slanted trapezoid to the top edge, so by Corollary 4.4, s intersects the positive diagonal B' of the slanted trapezoid, which is E_2' as desired. □

Most of the cylinders in the surfaces we discuss in this paper are the union of two trapezoids, covered in Lemma 4.7. Sometimes, as on each end of a

Bouw–Möller surface, one of the trapezoids is degenerate and is a triangle, and other times, as in a double $(2n+1)$-gon, both of the trapezoids in the cylinder decomposition degenerate to triangles. We consider these to be special cases of two trapezoids, so we cover these as part of Lemma 4.7. Sometimes, as in a regular $4n$-gon, the middle cylinder is a rectangle, but a double cover (two rectangles) is two trapezoids, so we cover this in Lemma 4.7 as well. Finally, horizontal edges are not the gluing edges for any cylinder, so we consider them in Lemma 4.8.

Lemma 4.7. *Let A be a gluing edge of a typical cylinder of a perfect translation surface of common modulus M. Consider a trajectory with $\theta \in [0, \theta_S)$ that crosses A, and let s be the segment that is the intersection of the trajectory with the slanted trapezoid that has A as its negative diagonal. Then s crosses $A' = V_M(A)$ if and only if the corresponding A in the cutting sequence occurs as the middle letter of a* (11) *or* (00) *word.*

Proof. This is a direct corollary of Lemmas 4.5 and 4.6.

Because this is the key result of the paper, we will draw several pictures to carefully work through why this is true.

There are three cases of a typical cylinders: A cylinder that is the union of two isosceles trapezoids, and then the two limiting cases, where one of the trapezoids is an isosceles triangle, or where both trapezoids are isosceles triangles. We will treat the general case first, and then give further arguments for the limiting cases.

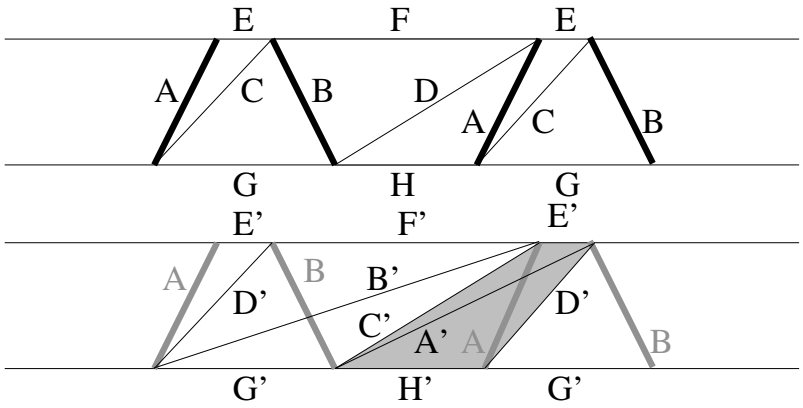

FIGURE 12. A typical cylinder, and the result of applying V_M, depicted in the universal cover. The slanted trapezoid with A, A' as diagonals is shaded

Consider a cylinder $\mathscr{C}$ of modulus M that is the union of two isosceles trapezoids glued along edges A and B. The top half of Figure 12 shows a portion of the universal cover of the cylinder $\mathscr{C}$. This is a copy of Figure 7. The bottom half of Figure 12 shows the action of V_M on $\mathscr{C}$, again depicted

in the universal cover. This is a copy of Figure 8 (b), except that we have shaded a different trapezoid.

In a perfect translation surface, each cylinder is partitioned into two isosceles trapezoids (in a typical case) or one rectangle (in an exceptional case), divided by the gluing edges. Now, we use a different partition, into slanted trapezoids, where the two trapezoids are divided by the slanted edges.

Consider a trajectory with $\theta \in [0, \theta_S)$ crossing the shaded slanted trapezoid in Figure 12. By Corollary 4.4, every such trajectory crosses A. By Lemmas 4.5–4.6, the trajectory crosses A' if and only if it is a (00) or (11) case.

The same analysis applies to the other slanted trapezoid in the cylinder (shaded in Figure 13): Every trajectory with $\theta \in [0, \theta_S)$ that crosses the shaded slanted trapezoid crosses B, but only those in the (00) or (11) cases cross B'.

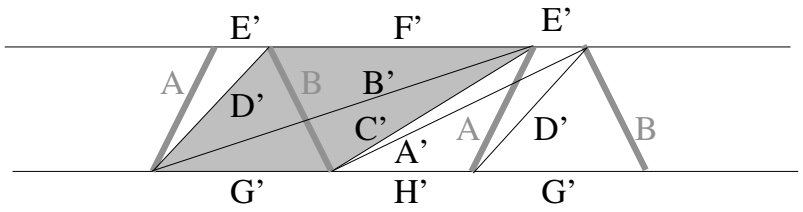

FIGURE 13. A typical cylinder, and the result of applying V_M, depicted in the universal cover. The slanted trapezoid with B, B' as diagonals is shaded

To spell things out in detail, we will discuss the limiting cases where one or two of the trapezoids are triangles.

If one of the trapezoids degenerates to a triangle, this means that edge E' (without loss of generality) is contracted to a point, and edge D' is not needed.

Consider edge A. The top edge has been contracted to a point, so only the cases (00) and (10) are possible. A trajectory in the (10) case crosses edge H' and then edge A. Such a trajectory clearly cannot cross edge A'. A trajectory in the (00) case crosses BAB, so by Corollary 4.4, it crosses A'. So indeed, the trajectory crosses A' if and only if it is in a (00) or (11) case, because the (11) case is not possible. For edge B, the analysis is the same as in the two-trapezoid case.

If both of the trapezoids degenerate to triangles, this means that both edges E' and H' contract to a point. Edges A, A', C', D' all coincide, so edges C' and D' are not needed and edge A' coincides with edge A.

Thus, trivially, a trajectory crosses A' if and only if it crosses A. The only way to cross edge A is in the (00) word BAB, so again, the trajectory crosses A' if and only if it is in a (00) or (11) case. For edge B (looking

at the slanted trapezoid whose diagonals are B and B'), the analysis is the same as in the two-trapezoid case.

Finally, we will consider an exceptional cylinder. We take a double cover of the exceptional cylinder, so in Figure 13 this means that gluing edges A and B have the same label (say, A) and are vertical. Similarly, each pair of edge labels C and D, E and F, and G and H is given the first label (C, E, G). We consider the slanted trapezoid $CE'CG'$, with positive diagonal A' and negative diagonal A, and then the analysis is the same as in the typical case.

For an expanded discussion of the degenerate and exceptional cases, see ([4], Chapter 3 and Lemma 4.8). □

Lemma 4.8. *A trajectory crosses a horizontal edge A if and only if it crosses its image $A' = V_M(A)$.*

Proof. The horizontal shear S_M is an isometry on horizontal segments, and R is a reflection, so $V_M = S_M \circ R$ is an isometry on horizontal segments. This means that $A = A'$, so the result is trivial. □

Lemma 4.9. *The derived sequence consists of the edge labels of the "primed" edges crossed by the trajectory, in the same order.*

Proof. We partitioned the infinite trajectory into countably many finite segments, each of which is the intersection of the trajectory with a slanted trapezoid. Lemma 2.22 tells us that the derived sequence consists of the crossed edge labels. These are a subset of the full sequence of edge labels, in the same order, so the order does not change. □

Now we may prove the theorem:

Proof of Theorem 4.2. By Lemma 2.22, determining the derived cutting sequence for a given trajectory is equivalent to determining which sheared edges the trajectory crosses. Lemmas 4.7–4.8 prove that the trajectory crosses the sheared edges as stated in the theorem. Lemma 4.9 tells us that their order does not change. This proves the result. □

5. Regular polygon surfaces

We will prove a corollary of Theorem 4.2, that the combinatorial rule for obtaining the derived sequence from the original cutting sequence for regular polygon surfaces is to keep only the *sandwiched* edge labels:

Definition 5.1. A *sandwiched* label is one that is preceded and followed by the same label, as A is in the sequence $\ldots CBABDAC \ldots$.

This result was already known for even polygon surfaces [8] and double odd polygon surfaces [3] but this is a new method of proof.

Definition 5.2. A *regular n-gon surface* is a double regular n-gon for n odd, and may be a single or double regular n-gon for n even. For a single n-gon, identify opposite parallel edges to obtain a translation surface (this

is only possible when n is even). For a double n-gon, take two copies of the regular n-gon, one of which is a reflection of the other, and identify opposite parallel edges. We assume that an edge of the polygon is horizontal.

John Smillie and Corinna Ulcigrai analyzed the even cases and showed, among many other results for these surfaces, that the combinatorial rule for obtaining the derived sequence from the original cutting sequence for regular $2n$-gon surfaces is to keep only the sandwiched edge labels ([8], Proposition 4.1.1). They used "transition diagrams," which we explore further in Section 7.2. We subsequently showed that Smillie and Ulcigrai's transition diagram method extends to double odd regular polygons as well ([3], Theorem 6.4). Note that Smillie and Ulcigrai proved many results for the regular octagon surface, and for even regular polygon surfaces in general, and we only extend a few of their results to other surfaces.

The two rules in Theorem 4.2 are reduced to this one simple rule ("keep only the sandwiched edge labels") in the case of regular polygon surfaces. The key insight is that because of the surfaces' symmetry, sandwiching occurs exactly in the (00) and (11) cases.

In Lemmas 5.3 (double) and 5.4 (single), we prove that regular n-gon surfaces are perfect, with common modulus $2\cot \pi/n$, so that we may apply the Main Theorem 4.2 to them. We will need the following trigonometric identities:

$$1 + 2\cos\theta + \ldots + 2\cos k\theta + \cos(k+1)\theta = 2\cot(\theta/2)\sin((k+1)\theta); \tag{1}$$

$$1 + 2\cos\theta + \ldots + 2\cos k\theta = \frac{\sin((k+1/2)\theta)}{\sin(\theta/2)}. \tag{2}$$

Equation (1) is proven in ([3], Lemma 2.1). To prove Equation (2), write the sum as $\sum_{j=-k}^{k} \cos j\theta$ and then rewrite $\cos j\theta = \frac{e^{ij\theta}+e^{-ij\theta}}{2}$ and simplify the expression. For details, see [4].

Lemma 5.3. *A double regular n-gon surface is perfect, with common modulus $2\cot \pi/n$.*

Proof. We will prove that each cylinder of a double regular n-gon surface (a) is the union of two isosceles trapezoids or two isosceles triangles, and (b) has modulus $2\cot \pi/n$.

(a) The interior of the polygon that lies between horizontal levels k and $k+1$ is an isosceles trapezoid for $k = 0, \ldots, n-2$ and is an isosceles triangle for $k = n-1$ when n is odd. A cylinder consists of gluing together two congruent copies of such pieces, one from each polygon.

(b) Assume that the edges of a regular polygon have length 1. The exterior angle of the polygon is $\theta = 2\pi/n$. In a regular polygon, the horizontal segment connecting the vertices at level 0 is 1 (the horizontal edge), at level 1 is $1 + 2\cos\theta$, and at level $k \neq 0$ in general is $1 + 2(\cos\theta + \ldots + \cos(k\theta))$ (see Figure 14).

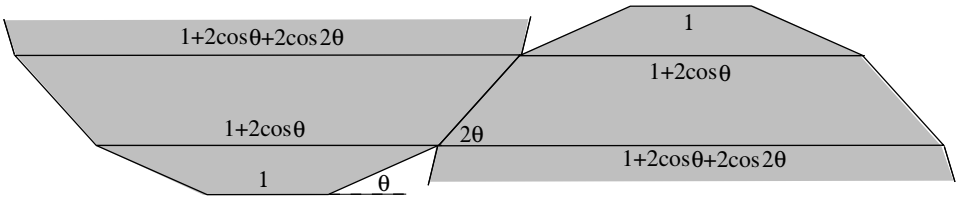

FIGURE 14. Two regular polygons, with opposite parallel edges identified

When we glue two regular polygons together to make a surface, two copies of the trapezoid or triangle between level k and level $k+1$ are glued together to make the kth cylinder. So the width of the kth cylinder is

$$1 + 2(\cos\theta + \ldots + \cos(k\theta)) \quad + \quad 1 + 2(\cos\theta + \ldots + \cos((k+1)\theta))$$
$$= 2\,(1 + 2\cos\theta + \ldots + 2\cos k\theta + \cos(k+1)\theta) = 2\cot(\theta/2)\sin((k+1)\theta),$$

by applying Equation (1).

The vertical distance between level k and level $k+1$ is $\sin(k+1)\theta$.

The modulus of the kth cylinder is then

$$\frac{2\cot(\theta/2)\sin(k+1)\theta}{\sin(k+1)\theta} = 2\cot\theta/2 = 2\cot\pi/n. \qquad \square$$

Lemma 5.4. *A regular $2n$-gon surface is perfect, with common modulus $2\cot\pi/n$.*

Proof. We will prove that each cylinder of a single regular $2n$-gon surface is the union of two isosceles trapezoids (typical), or is one rectangle (exceptional). In the typical case, the cylinder has modulus $2\cot\pi/n$. In the exceptional case, the cylinder has modulus $\cot\pi/n$.

The region that lies between a pair of non-horizontal, non-vertical edges of an even-gon is an isosceles trapezoid. It is not a triangle, because $2n$-gons have top and bottom horizontal edges. For the cylinders that are the union of two trapezoids, the calculation is the same as in Lemma 5.3, so the modulus is $2\cot\pi/n$.

A $4m$-gon with horizontal edges also has two vertical edges, which are identified. The cylinder glued along the vertical edge consists of a single rectangle. The rectangle's horizontal edges are between (using the same notation as in Lemma 5.3) levels $m-2$ and $m-1$, so the width is

$$1 + 2(\cos\theta + \ldots + \cos m\theta)$$

and its height is 1, so its modulus is just the width:

$$1 + 2\cos\theta + \ldots + 2\cos(m\theta) = \frac{\sin((m+1/2)\theta)}{\sin(\theta/2)},$$

by applying Equation (2).

Applying the angle-addition formula, we can rewrite this as

$$\frac{\sin(m\theta)\cos(\theta/2) + \cos(m\theta)\sin(\theta/2)}{\sin(\theta/2)} = \cot\theta/2 = \cot\pi/n,$$

since $m\theta = m(2\pi/(4m)) = \pi/2$, so $\sin(m\theta) = 1, \cos(m\theta) = 0$. This is a single rectangle of modulus $\frac{1}{2}(2\cot\pi/n)$, so the surface is perfect. □

Lemma 5.5. *For a regular n-gon surface, $\theta_S = \pi/n$.*

Proof. A regular polygon is cyclic. The angle of each slanted edge can be taken as the angle between two of the polygon's diagonals. Each of these pairs of diagonals spans a chord of the circle. The Central Angle Theorem says that this angle is half that of the central angle spanned by the same chord. The result follows. □

Theorem 5.6. *Consider a regular n-gon surface, and let $M = 2\cot\pi/n$. Consider a geodesic trajectory on the surface whose angle θ satisfies $\theta \in [0, \pi/n)$, and its associated cutting sequence. Then the derived sequence consists precisely of the sandwiched letters, read from left to right.*

To streamline the proof, we introduce the following terminology:

Definition 5.7. If a given letter in the cutting sequence survives into the derived sequence, we say that it is *kept*.

Theorem 4.2 says that an gluing edge label is kept if and only if it is the middle letter of a (00) or (11) word. We must show that this happens for regular polygon surfaces if and only if the edge label is sandwiched. We will show this for gluing edges and for horizontal edges. This will prove the result.

Proof. By Lemmas 5.3 and 5.4, regular polygon surfaces are perfect with common modulus M, and by Lemma 5.5, $\theta_S = \pi/n$. So the hypotheses of Theorem 4.2 are satisfied, and we may apply it.

Suppose that a given gluing edge label is sandwiched. Then it occurs in the middle of a three-label sequence $E_1E_2E_1$ for some edge labels E_1, E_2 (not necessarily distinct). If E_1 and E_2 are at the same level, then the transition E_1E_2 is type (0), and the transition E_2E_1 is also case (0), so it is case (00) and by Theorem 4.2, E_2 is kept. Similarly, if E_1 and E_2 are on different levels, then both transitions are type (1), so it is case (11) and by Theorem 4.2, E_2 is kept.

Suppose that a given gluing edge label is not sandwiched. Then it occurs in the middle of a three-label sequence $E_1E_2E_3$ for some edges E_1, E_2, E_3 not necessarily all distinct, but $E_1 \neq E_3$. Because regular polygon surfaces are perfect, each cylinder has at most two gluing edges, so at least two of the edges must be in different levels. If E_1, E_3 are in the same level and E_2 is in a different level, then $E_1 = E_3$ by symmetry, so E_2 is sandwiched, a contradiction. If E_1, E_2 are in the same level and E_3 is in a different level,

or if E_2, E_3 are in the same level and E_1 is in a different level, these are cases (01) and (10), respectively, so by Theorem 4.2, E_2 is not kept.

Theorem 4.2 says that a horizontal edge label is always kept. By symmetry, a horizontal edge label is always sandwiched in regular polygon surfaces. So a horizontal edge label is kept if and only if it is sandwiched. □

6. Introduction to Bouw–Möller surfaces

In Lemmas 5.3 and 5.4, we showed that for the n-indexed family of regular n-gon surfaces, the cylinder moduli are equal, or in a $2:1$ ratio. It turns out that there is an (m,n)-indexed family of surfaces whose cylinder moduli are all equal, called *Bouw–Möller surfaces*.

As discussed in the introduction, after Irene Bouw and Martin Möller gave algebraic models for their Bouw–Möller surfaces [1], Pat Hooper gave a polygon decomposition for them [5], which we will use here. The polygon decomposition gives a family of translation surfaces created by identifying opposite parallel edges of a collection of m "semi-regular" polygons that each have $2n$ edges.

In Definitions 6.1 and 6.4, we define Bouw–Möller surfaces. In Lemmas 6.2 and 6.6, we show that they are perfect translation surfaces, so we can apply the main Theorem 4.2, which we do in Chapter 7.

The language of Definitions 6.1 and 6.4 is taken directly from [5].

Definition 6.1. The (a,b) semi-regular $2n$-gon has edge vectors given by:

$$\mathbf{v}_i = \begin{cases} a\,[\cos\frac{i\pi}{n}, \sin\frac{i\pi}{n}] & \text{if } i \text{ is even} \\ b\,[\cos\frac{i\pi}{n}, \sin\frac{i\pi}{n}] & \text{if } i \text{ is odd} \end{cases}$$

for $i = 0, \ldots, 2n-1$. Denote this $2n$-gon by $P_n(a,b)$. The edges whose edge vectors are $\mathbf{v_i}$ for i even are called *even edges*. The remaining edges are called *odd edges*. We restrict to the case where at least one of a or b is nonzero. If a or b is zero, $P_n(a,b)$ degenerates to a regular n-gon.

Lemma 6.2. *Semi-regular polygons are special.*

Proof. Recall Definition 2.4, that a polygon is *special* if it is convex, level and vertically symmetric.

Semi-regular polygons are convex because the exterior angle increases monotonically from 0 to $(2n-1)/n$.

To show that they are level, we must show that vector $\mathbf{v}_i$ has the same y-value as vector $\mathbf{v}_{2n-i}$. To show that they are vertically symmetric, we will show that the region of a semi-regular polygon between two levels is an isosceles trapezoid or triangle. To do this, we will show that the x-values of $\mathbf{v}_i$ and $\mathbf{v}_{2n-i}$ have the same magnitude. If i is even,

$$\mathbf{v}_{2n-i} = a\left[\cos\frac{(2n-i)\pi}{n}, \sin\frac{(2n-i)\pi}{n}\right] = a\left[-\cos\left(\frac{i\pi}{n}\right), \sin\left(\frac{i\pi}{n}\right)\right].$$

Since $\mathbf{v}_i = a\left[\cos\left(\frac{i\pi}{n}\right), \sin\left(\frac{i\pi}{n}\right)\right]$, this proves both results.

For i odd, a is replaced with b. □

Corollary 6.3. Each Bouw–Möller surface admits the Veech element R, whose derivative is a vertical flip. This follows from the vertical symmetry.

Definition 6.4. The (m, n) Bouw–Möller surface is made by identifying the edges of m semi-regular polygons $P(0), \ldots, P(m-1)$. Define $P(k)$ by

$$P(k) = \begin{cases} P_n\left(\sin\frac{(k+1)\pi}{m}, \sin\frac{k\pi}{m}\right) & \text{if } m \text{ is odd} \\ P_n\left(\sin\frac{k\pi}{m}, \sin\frac{(k+1)\pi}{m}\right) & \text{if } m \text{ is even and } k \text{ is even} \\ P_n\left(\sin\frac{(k+1)\pi}{m}, \sin\frac{k\pi}{m}\right) & \text{if } m \text{ is even and } k \text{ is odd.} \end{cases}$$

We form a surface by identifying the edges of the polygon in pairs. For k odd, we identify the even edges of $P(k)$ with the opposite edge of $P(k+1)$, and identify the odd edges of $P(k)$ with the opposite edge of $P(k-1)$. The cases in the definition of $P(k)$ are chosen so that this gluing makes sense.

The (m, n) Bouw–Möller surface is a collection of m polygons, where the first and mth are regular n-gons, and the $m-2$ middle polygons are semi-regular $2n$-gons (see Figure 15). If m is odd, the central polygon is regular, and if m is even, the two central polygons are regular.

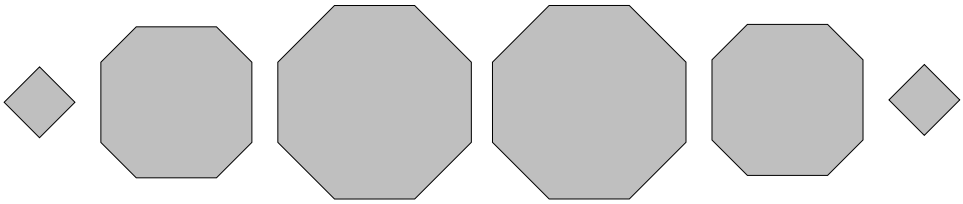

FIGURE 15. The Bouw–Möller surface $M_{(6,4)}$. Edges of the square $P(0)$ on the far left are glued to oppositely-oriented parallel edges of the semi-regular octagon $P(1)$. The remaining edges of $P(1)$ are glued to the regular octagon $P(2)$, and so on.

Bouw–Möller surfaces are translation surfaces because each surface is a collection of polygons with opposite parallel edges identified. Each $2n$-gon has n of its edges glued to the polygon on its left, and the other n edges glued to the polygon on its right. Each of the n-gons on the ends is only glued to one other polygon.

Corollary 6.5 (to Lemma 6.2)**.** Bouw–Möller surfaces are special.

Proof. Lemma 6.2 shows that semi-regular polygons are special, and Definition 6.4 gives a decomposition of a Bouw–Möller surface into semi-regular polygons. □

Lemma 6.6. *Bouw–Möller surfaces are perfect.*

Proof. Recall (Definition 2.18) that a perfect translation surface is a special translation surface that has a perfect cylinder decomposition with some common modulus M. Corollary 6.5 shows that Bouw–Möller surfaces are special. Showing that Bouw–Möller surfaces have a perfect decomposition requires that:

(a) Every cylinder is the union of two isosceles trapezoids.
(b) Every cylinder has the same modulus.

(a) We showed in Lemma 6.5 that Bouw–Möller surfaces are special, so the only way a cylinder could fail to be the union of two isosceles trapezoids is if it consisted of a single rectangle with just one gluing edge. This occurs when a vertical edge of a polygon is identified with an opposite parallel edge of the same polygon. By the gluing rules in Definition 6.4, each polygon edge is glued to an edge of another polygon, not an edge of the same polygon, so this cannot occur, and every cylinder is the union of two isosceles trapezoids.

(b) In fact, the edge lengths of Bouw–Möller surfaces are chosen precisely so that the modulus is the same; see [1], Theorem 8.1. The calculation is a elementary but lengthy trigonometric result: First, we calculate the length of the horizontal segment connecting the two vertices of polygon $P(k)$ at each level, and then add two of these together to find the width of each cylinder. A simple calculation gives the height of the cylinder, and then the modulus of the cylinder is the ratio of width to height. It turns out that the modulus of every cylinder is $2\cot\pi/n + 2\frac{\cos\pi/m}{\sin\pi/n}$. For details, see ([4], Lemma 6.6). □

7. Results for Bouw–Möller surfaces

In Lemma 6.6, we showed that Bouw–Möller surfaces are are perfect. This tells us that we may apply the main Theorem 4.2 to them, which is what we do in this chapter. In § 7.1, we prove Theorem 7.1, which describes how to determine the derived sequence from a cutting sequence on a Bouw–Möller surface, and we work through an example. In § 7.2, we describe the *transition diagrams* for Bouw–Möller surfaces, extending some of John Smillie and Corinna Ulcigrai's ideas for regular polygons ([8], § 2.1.2) to these surfaces.

7.1. Cutting sequences on Bouw–Möller surfaces.

Theorem 7.1. *Let $M = 2\cot\pi/n + 2\frac{\cos\pi/m}{\sin\pi/n}$, the modulus of the cylinders of the (m, n) Bouw–Möller surface S. Consider a geodesic trajectory with $0 \leq \theta < \theta_S$, and its associated cutting sequence. Mark a letter in the cutting sequence if:*

- *It corresponds to a gluing edge, and it is the middle letter of a* (00) *or* (11) *word, or*
- *it corresponds to a horizontal edge.*

Then the derived sequence consists precisely of the marked letters, read from left to right.

Proof. By Lemma 6.6, Bouw–Möller surfaces are are perfect with common modulus M, so we may apply Theorem 4.2 to them, and the same result holds. □

Example 7.2. We will determine the derivation rule for cutting sequences on the $(3, 4)$ Bouw–Möller surface, shown in Figure 16. (The ordering of the edge numbers may seem arbitrary, but it is chosen to simplify the *transition diagrams* in Section 7.2.)

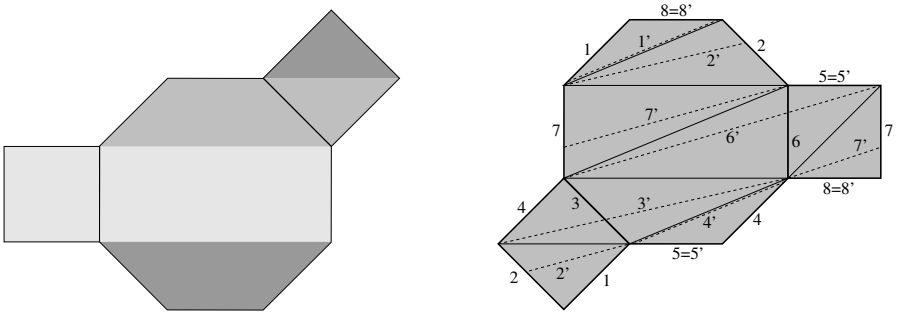

FIGURE 16. The $(3, 4)$ Bouw–Möller surface (a) showing the three cylinders, and (b) with edge labels, sheared edges (dotted lines) and auxiliary edges (thin)

We will consider the gluing edges one cylinder at a time, and then the horizontal edges.

First, we consider the cylinder with gluing edges 1 and 2. By Theorem 7.1, 1 and 2 are kept if and only if they are the middle numbers of a (00) or (11) word.

The (00) word with 1 as the middle number is 212. The (00) and (11) cases with 2 as the middle number are 121 and 723. The middle numbers are removed in the remaining (01) and (10) cases, which are 218, 123 and 721.

By the same reasoning on the cylinder with gluing edges 3 and 4, 3 is kept in the (00) and (11) cases 434 and 236, and removed in the other (01) and (10) cases 436 and 234. 4 is kept in the (00) case 343 and removed in the (10) case 543.

The gluing edges of the central cylinder are 6 and 7. The (00) cases are 767 and 676. The (11) cases are 365 and 872. By Theorem 7.1, the middle numbers are kept in these cases. The (01) cases are 765 and 672. The (10) cases are 367 and 876. By Theorem 7.1, the middle numbers are removed in these cases.

By Theorem 7.1, horizontal edges 5 and 8 are always kept. Possible three-number transitions with middle numbers 5 or 8 are 654 and 187, so the middle number is kept in these cases.

The numbers that are kept and removed are summarized in Table 1. Of the 12 three-letter transitions where the middle letter is kept, in the six (00) cases the middle letter is sandwiched, and in the six (11) cases, it is not.

kept (00)	kept (11)	removed (01), (10)
212		218
121	723	123, 721
434	236	436, 234
343		543
	654	
767	365	765, 367
676	872	672, 876
	187	

TABLE 1. All possible three-number transitions for the $(3, 4)$ Bouw–Möller surface, indicating when the middle number is kept. The table is sorted by middle number

7.2. Transition diagrams for Bouw–Möller surfaces. To determine the derived sequence from a cutting sequences on a regular $2n$-gon surface, John Smillie and Corinna Ulcigrai used *transition diagrams* ([8], § 2.1.2). In this section, we will describe the form of transition diagrams for Bouw–Möller surfaces. Using the transition diagrams may yield a new proof for Theorem 7.1, more similar to the proofs used in [3], [8] and [9].

Definition 7.3. A *transition diagram* for a translation surface S is a directed graph whose vertices are edge labels, and whose arrows connect two edge labels if a trajectory with $\theta \in [0, \theta_S)$ can intersect the first edge and then the second edge.

Example 7.4. We construct the transition diagram for the regular octagon surface (Figure 17). For the regular octagon, $\theta_S = \pi/8$, so we consider trajectories with $\theta \in [0, \pi/8)$. A trajectory crossing horizontal edge 1 must next cross edge 2, so we connect 1 to 2 with an arrow. A trajectory crossing edge 2 may next cross edge 1 or edge 3, so we connect 2 to 1 and to 3. The rest of the transition diagram is constructed similarly.

Example 7.5. We construct the transition diagram for the $(3, 4)$ Bouw–Möller surface (Figure 18). For this surface, $\theta_S = \pi/8$, so we again consider trajectories with $\theta \in [0, \pi/8)$. A trajectory crossing edge 1 then crosses edge 2 or 8, so we connect these with arrows. A trajectory crossing edge 2 then crosses edge 1 or 3, so we connect these with arrows. The rest of the diagram is constructed similarly.

Arranging the edge labels in numerical order in a straight line, as we did in Example 7.4, gives us the transition diagram shown at the top right of

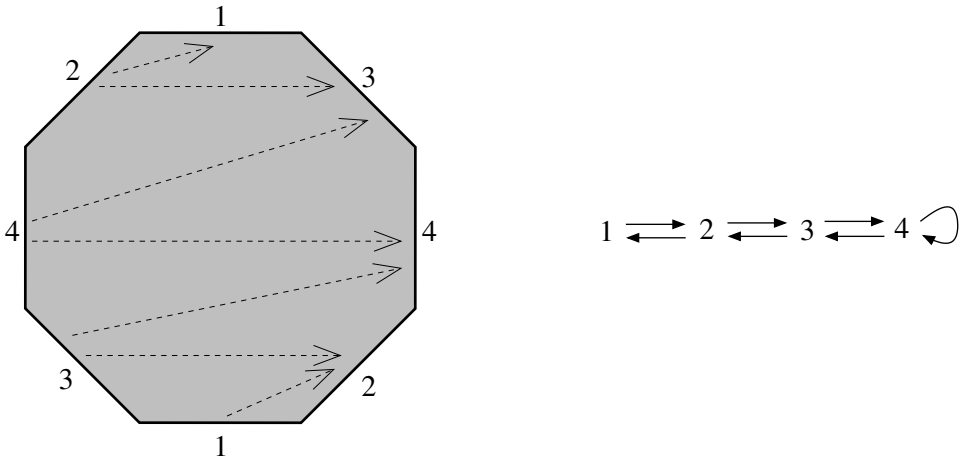

FIGURE 17. The regular octagon surface, and its transition diagram

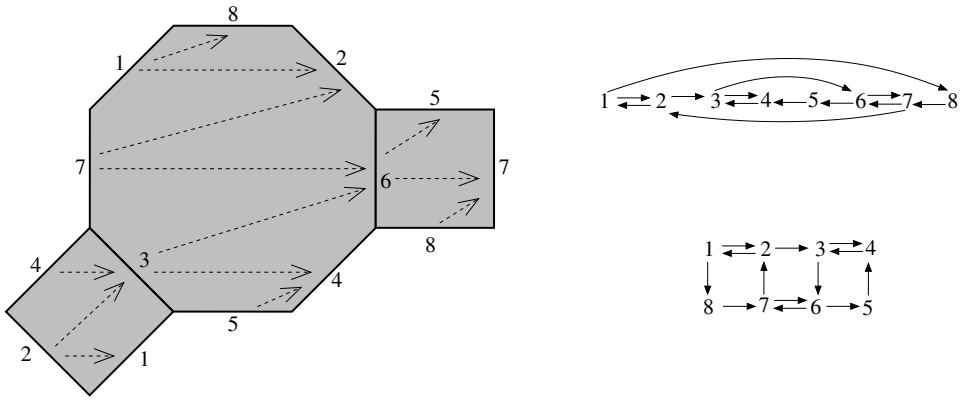

FIGURE 18. The $(3, 4)$ Bouw–Möller surface, and its transition diagram in two forms

Figure 18. We can rearrange the transition diagram into a rectangle, as shown at the bottom right of Figure 18.

It turns out that the rectangular form of the transition diagram in Figure 18 extends, and holds for Bouw–Möller surfaces in general. Figure 19 shows the polygon decomposition and the transition diagram for the $m = 6, n = 5$ Bouw–Möller surface.

There are n numbers on each row, snaking back and forth, with $m - 1$ rows in all. The edge labels in the kth row correspond to the edge labels that appear for the first time in $P(k-1)$. For example, in the $(3, 4)$ Bouw–Möller

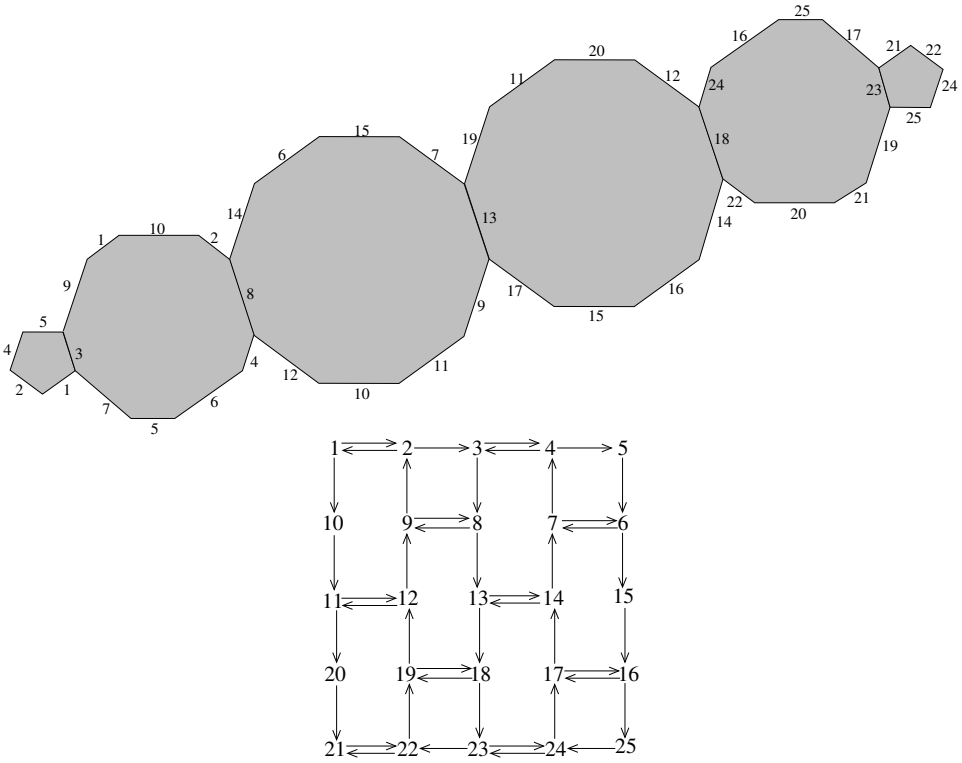

FIGURE 19. The $(6, 5)$ Bouw–Möller surface and its transition diagram

surface, $1, 2, 3, 4$ appear in the first tilted square $P(0)$, $5, 6, 7, 8$ appear in the octagon $P(1)$, and no new edge labels appear in the final square $P(2)$.

The transition diagram for the (m, n) Bouw–Möller surface, with $m \geq 2, n \geq 3$, always follows the pattern of Figure 19, with the upper-left corner the same and the diagram expanding (or contracting) down and to the right. Showing that the transition diagrams always have this form is tedious, so we omit the proof.

References

[1] BOUW, IRENE I.; MÖLLER, MARTIN. Teichmüller curves, triangle groups, and Lyapunov exponents. *Ann. of Math.* (2) **172** (2010), no. 1, 139–185. MR2680418 (2012b:32020), Zbl 1203.37049, arXiv:math/0511738, doi: 10.4007/annals.2010.172.139.

[2] DAVIS, DIANA; FUCHS, DMITRY; TABACHNIKOV, SERGE. Periodic trajectories in the regular pentagon. *Mosc. Math. J.* **11** (2011), no. 3, 439–461, 629. MR2894424 (2012k:37074), Zbl 1276.37033, arXiv:1102.1005.

[3] DAVIS, DIANA. Cutting sequences, regular polygons, and the Veech group. *Geom. Dedicata* **162** (2013), 231–261. MR3009542, Zbl 06133885, doi: 10.1007/s10711-012-9724-2.

[4] Davis, Diana. Cutting sequences on translation surfaces. Ph.D. thesis, *Brown University*, 2013. http://math.brown.edu/∼diana/.

[5] Hooper, W. Patrick. Grid graphs and lattice surfaces. *Int. Math. Res. Not. IMRN* **2013** no. 12, 2657–2698. MR3071661, doi: 10.1093/imrn/rns124.

[6] Lothaire, M. Algebraic combinatorics on words. Encyclopedia of Mathematics and its Applications, 90. *Cambridge University Press, Cambridge*, 2002. xiv+504 pp. ISBN: 0-521-81220-8. MR1905123 (2003i:68115), Zbl 1001.68093, doi: 10.1007/s00233-004-0146-9.

[7] Series, Caroline. The geometry of Markoff numbers. *Math. Intelligencer* **7** (1985), no. 3, 20–29. MR0795536 (86j:11069), Zbl 0566.10024, doi: 10.1007/BF03025802.

[8] Smillie, John; Ulcigrai, Corinna. Beyond Sturmian sequences: coding linear trajectories in the regular octagon. *Proc. Lond. Math. Soc.* (3) **102** (2011), no. 2, 291–340. MR2769116 (2011m:37012), Zbl 1230.37021, doi: 10.1112/plms/pdq018.

[9] Smillie, John; Ulcigrai, Corinna. Geodesic flow on the Teichmüller disk of the regular octagon, cutting sequences and octagon continued fractions maps. *Dynamical numbers - interplay between dynamical systems and number theory*, 29–65. Contemp. Math., 532. *Amer. Math. Soc., Providence, RI* 2010. MR2762132 (2012h:37075), Zbl 1222.37012, arXiv:1004.2265.

[10] Veech, W. A. Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards. *Invent. Math.* **97** (1989), no. 3, 553–583. MR1005006 (91h:58083a), Zbl 0676.32006, doi: 10.1007/BF01388890.

[11] Ward, Clayton C. Calculation of Fuchsian groups associated to billiards in a rational triangle. *Ergodic Theory Dynam. Systems* **18** (1998), no. 4, 1019–1042. MR1645350 (2000b:30065), Zbl 0915.58059, doi: 10.1017/S0143385798117479.

[12] Wright, Alex. Schwarz triangle mappings and Teichmüller curves: the Veech–Ward–Bouw–Möller curves. *Geom. Funct. Anal.* **23** (2013), no. 2, 776–809. MR3053761, Zbl 1267.30099, arXiv:1203.2685, doi: 10.1007/s00039-013-0221-z.

[13] Zemlyakov, A.N.; Katok, A. B. Topological transitivity of billiards in polygons. *Mat. Zametki* **18** (1975), no. 2, 291–300. MR0399423 (53 #3267), Zbl 0323.58012, doi: 10.1007/BF01818045.

Department of Mathematics, Northwestern University, 2033 Sheridan Road, Evanston, IL 60208-2370
diana@math.northwestern.edu

This paper is available via http://nyjm.albany.edu/j/2014/20-23.html.