Square Tilings of Translation Surfaces

Kevin Cong

January 16, 2022

Abstract

Translation surfaces are obtained by identifying opposite edges of a polygon with an even number of sides, paired together. We explore the question of tiling translation surfaces including the torus and the surfaces generated by the regular octagon with squares. Given any tiling, we identify its contacts graph, a triangulation formed by corresponding one vertex per square and drawing edges between vertices corresponding to adjacent squares. In particular, we prove that under certain conditions, there is exactly one torus tiling that has contacts graph a given torus triangulation. We then provide a method to approximately construct this tiling. We also show that the regular octagon translation surface cannot be tiled with squares. However, we give constructive tilings of translation surfaces corresponding to certain affine transformations of the octagon.

1 Introduction

Square tilings are an extremely intriguing aspect of mathematics. Not only are tilings aesthetically beautiful, but they also relate to many important mathematical concepts, such as packing and conformal theory. How does one effectively work with square tilings? In 1993, Oded Schramm [2] examined such tilings of a square using a combinatorial approach. His work was founded upon two important ideas: the contacts graph of a tiling and the *extremal metric* of such a graph.

Using these tools, Schramm proved the following striking theorem.

Theorem 1.1. *Given a triangulation T of a square $\mathcal{S}_1$ with vertices V_i, there exists a unique real number $h > 0$ and a tiling of a $\frac{1}{h}$ by h rectangle with squares Z_i such that Z_i and Z_j share an edge if and only if V_i and V_j are connected, with the condition that if V_i touches an edge of $\mathcal{S}_1$, then Z_i touches the corresponding edge of the rectangle.*

In this paper, we generalize several of Schramm's results concerning tori. We first show that Theorem 1.1 can be extended to cylinders and tori. In particular, for any triangulation of a cylinder or torus T, there exists a unique tiling of a cylinder or torus, respectively, with T as contacts graph, up to horizontal translations or vertical translations.

We then introduce *translation surfaces*, surfaces generated from identifying opposite edges in a polygon. We prove that the surface corresponding to the regular octagon cannot be tiled at all. However, in the same vein as the use of the variable h in Theorem 1.1, we show that there exist translation surfaces corresponding to an affine transformation of the regular octagon which are in fact tileable.

The remainder of the paper is structured as follows. Section 2 presents the preliminary definitions used in the remainder of the paper. Section 3 introduces results in the case of singly periodic or cylinder tilings. Section 4 gives results on doubly periodic or tori tilings, and Section 5 considers a generalization involving translation surfaces, constituting our main results. Finally, in Section 6, we present paths for further research.

2 Preliminary Definitions

We begin by introducing the key concepts used by Schramm: the *contacts graph* and the *extremal metric*.

The *contacts graph* of a tiling T which consists of finitely many squares $\{A_1, A_2, \dots A_n\}$ is the graph G that has vertex set V and edge set E, such that $V = \{V_1, V_2, \dots V_n\}$ and $\overline{V_i V_j} \in E$ if and only if A_i and A_j are adjacent in the tiling. As suggested by its name, this graph gives us a way of documenting the adjacencies of a tiling.

In order to understand the extremal metric, we must first understand the definitions of *metric extremal length.* Let G be an arbitrary simple graph.

Definition 1. *A metric m is any assigning of positive real numbers $m(V_i)$ to each vertex V_i.*

The *length* of any simple path $\gamma = V_1 V_2 ... V_n$ in G is then just $l(\gamma) = \sum_{i=1}^{n} m(V_i)$. By considering all possible paths in a graph from one set of vertices to another set for a given metric, we arrive at a notion for the length of a metric.

Definition 2. *Let V be the vertex set of G and S_1, S_2 be two disjoint subsets of V. Let m be a metric and $l(m)$ denote the infimum over all paths from S_1 to S_2 of the path lengths. Then $l(m)$ is the length of m on (G, S_1, S_2).*

Because this definition of length ignores scaling, we also present a normalized definition of length.

Definition 3. *Let V be the vertex set of G and S_1, S_2 be two disjoint subsets of V. The normalized length of (G, S_1, S_2) is $\frac{l(m)^2}{||m||^2}$, where $l(m)$ is the length and $||m|| = \sqrt{\sum_{v \in V} m_v^2}$ is the norm.*

Using this notion of normalization, we may define the *extremal length* of a graph and two vertex sets.

Definition 4. *Let V be the vertex set of G and S_1, S_2 be two disjoint subsets of V. The extremal length of the set (G, S_1, S_2) is the supremum over all metrics of the normalized length of each metric.*

Finally, we arrive at the definition of the *extremal metric*.

Definition 5. *An extremal metric is any metric whose normalized length equals the extremal length of (G, S_1, S_2).*

Schramm proved in [2] that the extremal metric is unique for any finite graph G. He then used the metric as follows. Let T be a triangulation of a square $\mathcal{S}_1$ with vertices V_i $(1 \leq i \leq n)$, let m be the extremal metric for (T, S_1, S_2), where S_1 consists of all vertices on the bottom edge of T and S_2 consists of all vertices on the top edge of T. Call a square tiling of a larger square $\mathcal{S}_2$ *admissible* if any vertex touching the bottom edge of $\mathcal{S}_1$ corresponds to a square touching the bottom edge of $\mathcal{S}_2$. Schramm's following theorem then relates the extremal metric to tilings.

Theorem 2.1. *[2] In any admissible tiling, the size of a square corresponding to $V \in T$ equals $m(V)$. Conversely, there exists a unique admissible tiling of $\mathcal{S}_2$ with contacts graph T.*

Note that Theorem 2.1 implies Theorem 1.1.

Schramm [2] also determined a method to approximately construct the tilings. We remark that because the extremal metric solves an extremal problem, it is difficult to determine its exact values. In the next sections, we generalize these results for singly and doubly periodic tilings. We begin with singly periodic tilings.

3 Singly Periodic Tilings

In the case of singly periodic tilings, many of Schramm's results for finite tilings still hold. Let T be a singly periodic triangulation invariant under the transformation $T + 1$, where $T + 1$ refers to the image of T translated right by 1. In what follows, we work only with such triangulations.

Call a periodic tiling Z of a horizontal strip bounded by $y = 0$ and $y = h$ with contacts graph T *admissible* if for any $V, V + k \in T$, their corresponding squares V', V'_k satisfy $V'_k = V' + k$, and if V intersects $y = 0$ if and only if V' intersects $y = 0$, and similarly for $y = 1$. In addition, for any vertex V of the T on either side (top or bottom) of the strip, the corresponding square V' touches the same side in Z. Finally, if there are two vertices V, W touching the same side of the strip such that one is to the left of the other, then V', W' satisfy the same left-right relation. We will work with admissible tilings only.

In order to generalize Schramm's results, we will need to modify the definition of the extremal metric.

3.1 The Extremal Metric for Singly Periodic Tilings

For a singly periodic triangulation, consider the paths from any vertex V to $V+1$. Define length and normalized length in the same way as in Section 2 except with this set of paths. A valid metric is then any metric m where for any vertices $V_0, V_1 \in T$, if $V_1 = V_0 + k$, then $m(V_1) = m(V_0)$. We define the extremal length as the supremum over all valid metrics of all normalized lengths, and an extremal metric as one attaining this supremum.

The following theorem, an analogue to Theorem 3.1 in [2], shows that this is a functional definition.

Theorem 3.1. *Any singly periodic triangulation T has a unique extremal metric m, up to scaling.*

Proof. We follow the proof of Theorem 3.1 in [2]. Call a metric *admissible* if it has unnormalized length at least 1. Then the set S of admissible metrics is nonempty, closed, and convex. Now, the norm is a strictly convex function because $\lambda x^2 + (1-\lambda)y^2 > [\lambda x + (1-\lambda)y]^2 \iff (\lambda)(1-\lambda)(x-y)^2 \geq 0$, which is true. Therefore, there exists a unique element m of S with minimal norm. Then, m must have length 1, as otherwise scaling m down reduces its norm. Since any other metric can be scaled to one of length 1, its normalized length must be strictly less than that of m unless it is m itself. Hence m is an extremal metric, and all other extremal metrics can be scaled to m. This completes the proof. $\square$

Next, we will relate the extremal metric to the sizes of squares in a singly periodic tiling.

3.2 Extremal Metric and Square Sizes

As shown below, the sizes of the squares in any admissible tiling Z correspond to the values of the extremal metric on T.

Theorem 3.2. *Let Z be any admissible tiling with contacts graph T. Then, for any square $V' \in Z$ corresponding to a vertex $V \in T$, the side length of V' equals $m(V)$.*

Proof. Suppose the strip is bounded by $y = t_0$ and $y = t_1$ for $t_1 - t_0 = h > 0$. Then let S_{R_1} be the set of squares in Z which intersect $y = t_0$ and let S_{R_2} be the set of squares in Z which intersect $y = t_1$.

Let $l(m)$ be the length of m. Given a line $x = r$, suppose it intersects $V' \in S_{R_1}$ and $W' \in S_{R_2}$. Clearly, V', W' are unique. Then, consider the unique path γ from V' to W' traversing through only squares which intersect $x = r$. We define the function $f : \mathbb{R} \to \mathbb{R}$ such that $f(r) = l(\gamma)$.

Clearly $f(r) \geq l(m)$. The key is then, like in [2], to integrate this over a range $r \in [0, M]$, for some integer M. This yields:

$$M \cdot l(m) \leq \int_0^M f(r)\, dr. \tag{1}$$

In the integral on the right hand side of 1, every square $V' \in T_0$ corresponding to a vertex $V \in T_0$ and its translations $V' + m$ contributes a total of at most

$$M \cdot m(V)s(V), \tag{2}$$

where $s(V)$ is the side length of V'.

Combining inequalities (1) and (2) and using the Cauchy-Schwarz inequality,

$$M \cdot l(m) \leq M \sum_{V \in T_0} m(V)s(V) \leq M||m|| \cdot ||s||. \tag{3}$$

Here, s is the metric which assigns to each vertex of T to side length of its corresponding square in Z. Then by a density argument, $||s||^2 = h$. In addition, $l(s) = h$ because the shortest path must always go from $y = t_0$ to $y = t_1$ vertically. Manipulating inequality 3 yields

$$\frac{l(m)^2}{||m||^2} \leq \frac{1}{M}||s||^2 = \frac{l(s)^2}{||s||^2}.$$

Since m is the unique extremal metric on T, it follows that $s \equiv m$. The lemma is thus proved. $\square$

Having related tilings to extremal metrics, we now turn to nonuniqueness.

3.3 Nonuniquenesses

Unlike in square tilings, singly periodic tilings are not necessarily unique when given the contacts triangulation, even when we restrict ourselves to admissible tilings. However, the nonuniquenesses may be concisely characterized. The following theorem determines all possible nonuniquenesses.

Theorem 3.3. *Suppose Z and Z' are singly periodic admissible tilings with contacts graph T. Then, one of the following is true:*

(1) There exists a vector $\lambda' \in \mathbb{R}^2$ for which $Z' = Z + \lambda'$.

(2) There exists real numbers a and P and a partition of the strip into horizontal strips of the form $H_k = \{(x, y), x \in \mathbb{R}, y \in [a_k, a_{k+1}]\}$ ($k \in \mathbb{Z}$), such that $Z \cap H_k$ and $Z' \cap H_k$ differ only by a horizontal translation.

Proof. First, choose any vertex V corresponding to a square V' which touches the bottom of the strip. Then, there is a unique vertex W adjacent to V which also corresponds to a square touching the bottom of the strip, W', and right of V'. Continuing and using periodicity, we may determine the locations of all the squares touching the bottom row easily.

Next, given any two nonaligning topmost squares such that we have not placed any squares directly above any of them, observe that we can uniquely place a square which touches both of them. This is because there are at most two squares touching both of them, which we may determine through the triangulation. However, these squares must have different heights along the strip. Yet through our metric, remark that there is at least one minimal path going through any one square, and the sum of metrics of all squares between the bottommost one and this square in the past gives the height of the square. Therefore, we may uniquely determine which square is in which location.

Also, each time we place a square, we may immediately place all translates of the square. Therefore, because we place new squares (even up to translation) each time, we will eventually have placed all the squares.

The nonuniqueness occurs when after having placed some set of squares, we obtain exactly a smaller strip. Then, the remaining squares must tile some smaller strip. In addition, we can determine which squares touch the bottom and top of the strip via the initial conditions and the squares which have already been placed. We may therefore continuously repeat this process until all squares have been placed, obtaining uniqueness up to horizontally translating each strip.

Thus, the square placements are unique up to horizontally translating strips. This completes the proof of the theorem. □

We end this section with the following remark.

Remark 6. *The proof of Theorem 3.3 also provides a method to construct singly periodic tilings given the contacts triangulation.*

4 Doubly Periodic Tilings

We now examine doubly periodic tilings of the plane. Let T be a doubly periodic triangulation invariant under the transformations $T+1$ and $T+i$. In what follows of this section, we will work with only such triangulations.

Call a doubly periodic tiling $Z = \{A_1, \dots A_n\}$ *admissible* if its contacts graph is a admissible triangulation T, and there is a vector $\lambda_Z = [u, v]$, $v > 0$, such that whenever two vertices V_i, V_j of T satisfy $V_j = V_i + [m, n]$, the corresponding squares V_i', V_j' satisfy $V_j' = V_i' + m + n\lambda$. We will restrict our focus to admissible tilings.

We then have the following theorem of existence.

Theorem 4.1. *[3] Let T be an admissible triangulation. Then there exists a admissible tiling Z with contacts graph T.*

Proof. The theorem follows from Theorem 6.5 in [3], by rounding the corners of the squares to make them convex disks.

□

4.1 The Extremal Metric for Doubly Periodic Tilings

We now adapt the idea of an extremal metric for doubly periodic tilings.

Let T_0 be a subset of vertices of T such that every vertex in T can be obtained by a translation by $[m, n]$ ($m, n \in \mathbb{Z}$) of some vertex in T_0. Let T_0 contain $|T_0|$ vertices, and let S_0 be the set of all finite paths in T from any vertex $V \in T_0$ to $V + 1$.

Then, the *extremal metric* of a triangulation T is defined in the same way as for square and singly periodic tilings, except the set of paths is S_0. First we will show that this definition is functional.

Theorem 4.2. *Any admissible doubly periodic triangulation T has a unique extremal metric m.*

Proof. The proof is the exact same as that of Theorem 3.1, as all the statements in the proof still hold. Thus this theorem is true as well. □

Now we analyze the relation between the extremal metric and the sizes of squares in admissible tilings. In fact, similar theorems hold for doubly periodic tilings as square and singly periodic tilings.

4.2 Extremal Metric and Square Sizes

The following theorem then shows that the sizes of the squares in any admissible tiling Z equal the values of the extremal metric on corresponding vertices of T:

Theorem 4.3. *Let T be a admissible triangulation with extremal metric m, and let Z be any admissible tiling with contacts graph T. Then, for any square $V' \in Z$ corresponding to a vertex $V \in T$, the side length of V' equals $m(V)$.*

Proof. The proof is somewhat similar to that of Theorem 3.2.

Let $\lambda = [u, h]$. Consider a large M by N rectangle $\mathcal{R}$ in the plane with bottom-left corner the origin, where M is an integer multiple of h and N is an integer, and let S_R be the set of squares in Z which intersect the left edge of $\mathcal{R}$.

Let $l(m)$ be the length of m. Given a line $y = r$ intersecting $\mathcal{R}$, suppose it intersects $V' \in S_R$. Clearly, V' is unique. Then, consider the unique path γ from V' to $V' + N$ traversing through only squares which intersect $y = r$. We define the function $f : R- > R$ such that $f(r) = l(\gamma)$.

Observe that γ can be written as the union of N paths going from $V' + i$ to $V' + i + 1$ for $0 \leq i \leq N - 1$. Since each of these paths has length at least $l(m)$, we must have

$$N \cdot l(m) \leq l(\gamma_t). \tag{4}$$

The key is then to integrate inequality (4) over the range of $r \in [0, M]$. We obtain:

$$MN \cdot l(m) \leq \int_0^M f(r)\, dr. \tag{5}$$

In the integral on the right hand side of (5), every square $V' \in T_0$ corresponding to a vertex $V \in T_0$ and its translations $V' + m + n\lambda$ contributes a total of at most

$$\frac{1}{h} MN \cdot m(V') \cdot s(V') + \frac{1}{h}(2M + 2N + 4), \tag{6}$$

where $s(V')$ is the side length of V'. This is because given V', at most $\frac{1}{h}(M+2)(N+2)$ of its translations, including itself, can intersect $\mathcal{R}$, and the side length of any square is clearly at most 1.

Combining inequalities (5) and (6) and using the Cauchy-Schwarz inequality,

$$MN \cdot l(m) \leq \sum_{V' \in T_0} \frac{1}{h} MN \cdot m(V') \cdot s(V') + \frac{1}{h}(2M + 2N + 4)$$
$$\implies MN \cdot s(m) \leq \frac{1}{h} MN ||m|| \cdot ||s|| + \frac{1}{h}(2M + 2N + 4)|T_0|. \quad (7)$$

Here, s is the metric which assigns to each vertex of T to the side length of its corresponding square in Z. Then by a density argument, $||s||^2 = h$. Therefore, manipulating (7) yields

$$\frac{l(m)^2}{||m||^2} \leq (\frac{1}{\sqrt{h}} + \frac{1}{h} \frac{2M + 2N + 4}{MN} |T_0|)^2. \quad (8)$$

Observe that $l(s) = 1$. Indeed, we may take some $y = r$ and some square V' intersecting the line. Then, the set of squares between V' and $V'+1$ intersecting $y = r$ forms a path with length 1, as Z is admissible. In addition, no path from V' to $V' + 1$ may have length less than 1, because the sum of side lengths in such a path is at least the difference in the x coordinates of V' and $V' + 1$.

Therefore, setting $M = hN$ and taking the limit of inequality (8) as N approaches infinity yields

$$\frac{l(m)^2}{||m||^2} \leq \frac{1}{h} = \frac{l(s)^2}{||s||^2}.$$

Since m is the unique extremal metric on T, it follows that $s \equiv m$. The lemma is thus proved. □

4.3 Nonuniquenesses

For doubly periodic tilings, the possible nonuniqueness is similar to that for singly-periodic tilings, except that we may have either horizontal or vertical strips. These possible nonuniquenesses are given in the theorem below.

Theorem 4.4. *Let T be a admissible triangulation. Suppose Z and Z' are admissible tilings with contacts graph T. Then, one of the following is true:*

(1) There exists a vector $\lambda' \in \mathbb{R}^2$ for which $Z' = Z + \lambda'$.

(2) There exist real numbers a and P, a vector λ', and a partition of the plane into horizontal strips of the form $H_k = \{(x, y), x \in \mathbb{R}, y \in [a_k, a_{k+1}]\}$ ($k \in \mathbb{Z}$), such that $Z \cap H_k$ and $Z' + \lambda' \cap H_k$ differ only by a horizontal translation.

(3) There exist real numbers a and P, a vector λ', and a partition of the plane into vertical strips of the form $C_k = \{(x, y), y \in \mathbb{R}, x \in [a_k, a_{k+1}]\}$ ($k \in \mathbb{Z}$), such that $Z \cap C_k$ and $Z' + \lambda' \cap C_k$ differ only by a vertical translation.

In order to prove this theorem, we will need several preliminary lemmas.

First, pick any vertex V of our triangulation T and let square A correspond to V. Place A randomly on the plane. Hereafter we will use for some square V' and constant r the notation $V' + r$ to denote $V' + [r, 0]$.

Call a vertex V_i of T *admissible* if there exists real numbers r, s and a path γ such that $V + r, V_i, V + s$ are in γ in that order, and γ has length $s - r$.

From now on, assume that neither (2) nor (3) occur. Then we have the following lemma.

Lemma 4.5. *All vertices of T are admissible.*

Proof. Suppose vertex W is not admissible, and let it correspond to $W' \in Z$. Since A is fixed, let A intersect the horizontal line $y = k_0$.

Without loss of generality, assume that W' is above $y = k_0$. Define a sequence of squares $V'_1, V'_2, \ldots V'_n, \ldots$ as follows: Let V'_1 be the lowest square which shares a vertical edge with W' and is to the right of W'. Given V'_n, let V'_{n+1} denote the lowest square to the right of and sharing a vertical edge with V'_n and to the right of V'_n.

Similarly, define $U'_1, \ldots U'_n$ to the left of W' such that U'_1 is the lowest square sharing a vertical edge with and to the right of W', and U'_{i+1} is the lowest square sharing a vertical edge with and to the right of U'_i.

Clearly it is impossible that V'_i is entirely above $y = k_0$ and V'_{i+1} is entirely below $y = k_0$. The same thing holds for the U'_i. There are now two possibilities:

(i) Both sequences eventually intersect $y = k_0$, say at U'_j and V'_k. Then, we may choose r, s appropriately such that $A + r$ is left of U'_j and $A + s$ is right of U'_k.

Define a sequence of squares intersecting $y = k_0$: $W'_0 = A + r$ and W'_{i+1} is adjacent to W'_i. Suppose that $U'_j = W'_x, V'_k = W'_y$, and $A + s = W'_z$.

Consider the path

$$\gamma = W'_0 W'_1 \ldots W'_x U'_j U'_{j-1} \ldots U'_1 W' V'_1 \ldots V'_{k-1} W'_y W'_{y+1} \ldots W'_z.$$

It is easy to see that the length of the γ is exactly $s - r$, because it is the sum of the side lengths of the squares in γ. In addition, consider any other path γ' taking $A + r$ to $A + s$. By projecting all the squares in the path onto the line $y = 0$, we obtain a set of intervals which cover a larger interval of length $s - r$. Since the extremal metric corresponds to the square side lengths, we see that the length of the path γ' is at least $s - r$.

Therefore, γ is the shortest path from $A + r$ to $A + s$, and since $W' \in \gamma$, W must be admissible.

(ii) At least one sequence does not intersect $y = k_0$. Without loss of generality, suppose V'_i does not. Set $f(V'_i)$ to be the y-coordinate of the bottom-left corner of V'_i. Clearly, the sequence $f(V'_i)$ is nonincreasing but bounded below. Therefore, it converges to some real t, and so for any sufficiently small ϵ, there exists $I \in \mathbb{Z}$ such that for all $i \geq I$, $y = t + \epsilon$ intersects the interior of V'_i.

Suppose there is any square S' whose interior has nonempty intersection with the line $y = t$. Then, observe that we may choose $\epsilon > \min\{\{m(V), V \in T\}, 0\}$ such that $y = t + \epsilon$ intersects the interior of S' as well. Then it is clear that all squares of the sequence V_i' for $i > I$ must intersect $y = t + \epsilon$ as otherwise their side length is no greater than ϵ, contradicting our definition of ϵ.

Choose sufficiently large l such that $f(S' + l) > f(V_{I+1})$. From the above remarks, $S' + l = V_j'$ for some j. But then $f(V_j')$ is less than t, contradicting the definition of t.

Thus, all squares of Z are either above $y = t$ or below $y = t$. Suppose $\lambda = [p, q]$. Then split the complex plane into strips $H_k = \{(x, y), x \in \mathbb{R}, y \in [t + kq, t + (k+1)q]\}$. It is clear that $Z \cap H_k$ is well defined because Z has no squares which intersect more than one strip.

Then, let the set of vertices which are admissible be S. By the above work, for any non-admissible vertex U corresponding to square U' above A, we may associate a line $y = g(U)$ with $g(U) < f(U')$, which no square intersects. By symmetry, for any non-admissible vertex U corresponding to square U' below A, we may associate a line $y = g(U)$ with $g(U) > f(U')$ which no squares intersect.

Consider $y = t$ and $y = t - q$. Clearly all admissible vertices are in said strip, since no vertex outside the strip shares a vertical edge with any edge inside the strip.

Now, if any square in the strip U is not admissible, then we may reduce the height of the strip by replacing one of the strip bounds by $y = g(U)$. It is clear that this reduces the strip height by at least the size of the smallest square. Therefore, repeating this process, we eventually obtain a strip with the entire set of squares corresponding to admissible vertices contained inside, and no other squares.

Then, the set of squares touching the boundaries of the strip is uniquely determined because it is simply the set of admissible vertices which have a non-admissible neighbor. Thus, using Theorem 3.3, since A is fixed and obviously admissible by definition, the strip is also fixed (i.e. same for Z and Z'), since the set of admissible vertices is independent of the tiling. Suppose this strip is bounded by $y = a_0$ and $y = a_1$, $a_0 < a_1$. Then first note that k is uniquely determined by the norm of the metric, since the total area of all the squares must be k by a density argument. Then by translating the squares between $y = a_0$ and $y = a_1$, the strip between $y = a_0 + k$ and $y = a_1 + k$ can also be uniquely covered by squares. Then, we may place the squares bounded by $y = a_1$ and $y = a_0 + k$ using Theorem 3.3, up to translation. Finally, we may shift the squares between $y = a_0$ and $y = a_0 + k$ to yield the entire tiling.

Note that this is done independently of the tiling (i.e. only with the metric), hence Z and Z' have both been split into horizontal strips on which they differ by a constant, i.e. (2) holds, contradiction.

The lemma is therefore proved in both cases. $\square$

Let the *x-coordinate* $x(S')$ of a square S' be the x-coordinate of the bottom left corner of the square.

Next, we prove that:

Lemma 4.6. *The x-coordinate of every square is determined uniquely by T.*

Proof. Let V' denote any square. By Lemma 4.5, V' is admissible. Thus, there exists real numbers r, s and a path $A+r, V_1', V_2', \ldots, V_n', V', U_1', U_2', \ldots, U_m', A+s$ with length $s-r$.

Project each square on the path to the line $y=0$. Since the extremal metric corresponds to the side length of the square, we obtain a sequence of intervals covering an interval of length $s-r$, with total length equalling $s-r$. Therefore, any two consecutive intervals intersect at their endpoints only.

This implies that the x-coordinate of V' equals $x(A+r)+m(A+r)+m(V_1')+m(V_2')+\ldots+m(V_n')$, which is uniquely determined by T, proving the lemma. $\square$

Now, define $y(S')$ to be the y-coordinate of the bottom-left corner of a square S'.

The following two lemmas will pinpoint the y-coordinate of each square.

Lemma 4.7. *Suppose that the squares S_0' and S_1' are fixed and share a horizontal edge, and that both squares intersect the line $x=r$. Then for any other square W' whose interior intersects $x=r$, the value $y(W')$ is determined uniquely.*

Proof. Recall that we have fixed the location of square A. Suppose the interior of A intersects the line $x=r$. By Lemma 4.6, the x-coordinate of every square is determined uniquely by T, hence the set S of squares which have nonempty intersection with $x=r$ is fixed.

Inductively place a sequence of squares S_i' whose locations are all fixed as follows: S_1' has already been placed. Given the placement of S_n', let S_{n+1}' be a square which is adjacent to S_n' and in the set S, but not adjacent to or equivalent to S_{n-1}'. Then by Lemma 4.6, the x-coordinate of S_{n+1}' is determined uniquely. In addition, by definition, the bottom edge of S_{n+1}' must lie on the same horizontal line as the top edge of S_n', hence $y(S_{n+1}') = y(S_n') + m(S_n)$. Since S_n' has been placed, $y(S_n')$ is fixed, hence $y(S_{n+1}')$ is fixed as well.

Thus, the location of S_{n+1}' is fixed and the square may be placed uniquely.

Finally, since $x=r$ intersects the interior of W', there must exist k such that $W' = S_k'$. Hence the y-coordinate of S_k' is determined uniquely. $\square$

Lemma 4.8. *Let $N>0$ be a fixed real number. Then there exists a real number m and a finite set of squares S', whose locations are uniquely determined by the location of A and the triangulation T, such that for every $r \in [m, m+N]$, the line $x=r$ intersects at least two squares in S' which share a horizontal edge.*

Proof. Set m such that $m+N = x(A)$. Set $A = A_1$. Using Lemma 4.7, we may uniquely place a square B_1 which shares a horizontal side with A_1.

We aim to inductively construct a sequence of quadruples $(A_i, B_i, x = r_i)$, such that each pair shares a horizontal edge, both squares intersect $x=k$ for any $k \in [r_i, r_{i-1}]$, r_i is the maximum of $x(A_i)$ and $x(B_i)$, and the sequence $\{r_i\}$ is strictly decreasing.

Define (A_1, B_1) as above, and simply let r_1 be the maximum of $x(A_1)$ and $x(B_1)$. Given A_k, B_k, using Lemma 4.6, we may construct a sequence of squares

$S_1 = A_k, S_2 = B_k, S_3, \ldots$ such that S_i and S_{i+1} share a horizontal edge and each S_i intersects $x = r_k$. We have the following two cases:

(i) If there exists i such that $x(S_i) < r_k$, take the smallest such i. Then, let A_{k+1} be S_i, and let B_{k+1} be the unique square adjacent to S_i and S_{i-1} and to the left of S_{i-1}. Clearly B_{k+1} is uniquely determined. Let the maximum of $x(B_{k+1})$ and $x(A_{k+1})$ be r_{k+1}. By definition, A_{k+1} and B_{k+1} must both intersect the line $x = x(r_{k+1})$. Also, $r_{k+1} < r_k$ and both A_{k+1} and B_{k+1} intersect r_{k+1}. In addition, by the minimal choice of i and the fact that $r_k = x(S_{i-1})$, we see that both squares intersect $x = r_k$.

(ii) If there does not exist such i, then since A_k and B_k intersect $x = r_k$, we must have $x(S_i) = r_k$ for all i. Then observe that we may partition the plane into strips bounded by $x = r_k$ and $x = r_k + l + 1$. By Lemma 4.6, the x-coordinate of every square is fixed, hence there is a unique way to partition the set of vertices in T into sets R_t such that R_t is the subset that corresponds to the set of squares within the strip bounded by $x = r_k + t, r_k + t + 1$. In addition, Lemma 4.6 implies that the set of squares intersecting the boundary of any strip is uniquely determined.

Therefore, by the vertical analogue to Theorem 3.3, each of these strips may be tiled uniquely up to vertical translation, hence we obtain the case (3) in the statement of Theorem 4.4. This contradicts our assumption, hence (ii) may not occur.

This process yields an inductive construction for the triples $(A_i, B_i, x = r_i)$. Set l to be the minimum square side length. Then, it is easy to see that $r_{i+1} \leq \max(r_i - l, r_{i-1} - l)$, since by the construction method (i) above, if $r_{k+1} = x(B_{k+1})$, then $r_{k+1} = r_k - m(B_{k+1}) \leq r_k - l$, and otherwise $r_{k+2} = x(B_{k+2})$.

Therefore the sequence r_i is not bounded, since $r_k \leq r_{k-2} - l$ and $l > 0$. Then set S to include all pairs A_i and B_i up to some sufficiently large index k, where $r_k < m$. From the construction of (A_i, B_i), we see that all $k \in [m, m+N]$ satisfy that $k \in [r_j, r_{j-1}]$ for some j and therefore $x = k$ intersects A_j and B_j, which share a horizontal edge.

Since A_i and B_i are determined uniquely by our process above, this completes the proof of the lemma. □

Now, we are ready to prove Theorem 4.4.

Proof. Take $N > 1+4\delta$, where δ is the maximum of the lengths of all the minimal paths between a vertex V and $V + i$ in T. Consider a set S' and real number m as described in the statement of Lemma 4.8. Then, for every $r \in [m, m + N]$, there are two squares intersecting $x = r$ which share a horizontal edge. By Lemma 4.6, the set of squares intersecting $x = r$ is uniquely determined, and by Lemmas 4.6 and 4.7, their locations are fixed as well.

Since $N > 1 + 4\delta$, there must be some two squares V', V'' which intersect the vertical strip bounded by $x = m$ and $x = m + N$, which correspond to vertices $V, V + [0, 1]$ in T. By the above comments, the positions of V' and V'' are uniquely determined. Therefore, $\lambda = [x(V'') - x(V'), y(V'') - y(V')]$ is fixed as well.

In addition, for every vertex $V \in T_0$, since $N > 1$, there exists a real number k such that $m < x(V+k) < m+N$. Then by Lemmas 4.7 and 4.8, the square V'_k corresponding to $V+k$ can be uniquely placed. Then the square corresponding to V can be obtained by shifting V'_k left by k.

Thus, every vertex $V \in T_0$ corresponds to a square V' which may be uniquely placed. Thus, Z and Z' must be the same tiling, proving the theorem. $\square$

4.4 Constructing the Tiling

The proof of Theorem 4.4 also provides us with the following method for constructing any given tiling from the contacts triangulation T and the extremal metric m.

First fix any square A corresponding to some $V \in T$.

Let the minimum value of the metric on T be r. Then, for each vertex V in T_0, take any path from V to $V+1$, and set it's length as $g(V)$. Then any path with more than $\frac{g(V)}{r}$ vertices cannot be the shortest path from V to $V+1$. Hence, we may take all paths with at most $\frac{g(V)}{r}$ vertices starting from V and in finite time deduce the length of the shortest path from V to $V+1$.

Repeating this process for each $V \in T_0$, we may deduce the length of the metric.

For the next steps, we borrow the notation of x-coordinate and y-coordinate as well as the functions $x(V')$ and $y(V')$ from the proof of Theorem 4.4.

For any vertex $W \neq X$, we claim that we can determine in finite time whether there exists $k, l > 0$ and a path from X to $X+l$ passing through $W+k$ with length l. Call W *X-accessible* if this path exists. Note that any such path will consist of vertically adjacent squares.

Indeed, follow in the method given by the proof of Lemma 4.5. As usual, let V'_0 be the square corresponding to V_0 for any vertex V_0. Take the sequences U'_i and V'_i defined there. Set $|T_0|$ to be the number of vertices in T_0. Note that in the first $|T_0|+1$ squares, there must be two squares which are translations of each other. If they are horizontal translations, then we obtain case (2) of Theorem 4.4, because W is not X-accessible. Otherwise, their y-coordinates must differ by at least k. Take any path γ from W to X, and observe that the difference in the y-coordinates of W and X is at most $l(\gamma)+m(W)+m(X)$. Thus, continuing the above process, we will eventually arrive at some V'_j such that V'_j and X' both intersect $y = r$ for some r, where $j \leq \frac{l(\gamma)+m(W)+m(V)}{|T_0|+1}$, or we will enter case (2) of Theorem 4.4. If the former occurs, we may take consecutive squares intersecting $y = r$ until we arrive at horizontal translation of X'. Note that this will happen within $\frac{1}{r}$ squares. Doing the same for the U_i and translating appropriately shows that we only need to consider paths from X' of length at most $3 + 2\frac{l(\gamma)+m(W)+m(V)}{|T_0|+1} + \frac{1}{r}$, of which there are finitely many to check.

Then, check for each $W \in T_0$ whether W is V-accessible. Then we may similarly check whether $W + [0, 1]$ and $W + 1$ are W-accessible. If any of our checks yields a negative answer, then due to Lemma 4.5, we are in case (2) of Theorem 4.4.

If in any of the above we obtain that we are in case (2) of Theorem 4.4. Then, take a sequence of vertices $V, V_1, V_2, \ldots, V_m$ such that no two consecutive vertices correspond to squares sharing a vertical edge and each V_i is V-accessible, but V_m has neighbors which are not V-accessible. This is possible because using a similar method as above, U and W share a vertical edge if and only if there is a path from U to $U + 1$ of length 1 which passes through W. We may suppose that the y-coordinates of the V_i form an increasing sequence.

Then, we may determine the height of one line bounding the horizontal strip which includes A. Next, we may take another sequence $V, W_1, W_2, \ldots, W_n$ satisfying equivalent conditions, but with W_1 and V_1 disconnected in the graph with vertex set the neighbors of V, and edges between two X, Y if $\overline{XY}$ is an edge of T and neither X, Y share a vertical edge with V. It is clear that the y-coordinates of this sequence are decreasing. The construction implies that this yields the bounds of the minimal horizontal strip consisting A, which must contain only V-accessible vertices from the proof of 3.5.

Now, take a sequence of vertices from V_m such that their adjacent corresponding squares share a vertical edge, their x-coordinates form a monotonically increasing sequence, and each vertex has a neighbor who is not V-accessible. Because a tiling must exist from Theorem 4.1, such a sequence will eventually reach $V_m + 1$. From this, we obtain by translation the entire set of vertices corresponding to squares touching the top edge of the strip. Similarly, we may deduce the set of vertices corresponding to squares touching the bottom edge of the strip.

Then, by taking lines $x = t + \frac{kr}{2}$, $k \in \mathbb{Z}$, $t \in R$ which intersect X_1, X_2 on the top and bottom edges of the strip, we may find all V-accessible vertices up to horizontal translation through taking all vertices on some path of minimal length from X_1 to X_2. The choice of $x = t+kr$ and the fact that each square has length at least r means that every V-accessible vertex corresponds to a square intersecting at least one such line, and is thus included in some minimal path.

Then using the remark to Theorem 3.3 we may construct the tiling of our strip containing A, bounded by $y = t_0$ and $y = t_1$ ($t_0 < t_1$). Suppose $\lambda = [p, q]$. Remark that q is determined by the metric, as shown in the proof of Theorem 4.4. Then we may tile by translation any strip $y = t_0 + kq, y = t_1 + kq$, $k \in \mathbb{Z}$. Finally, we may determine all squares adjacent to the borders $y = t_1$ and $y = t_0 + q$, hence by the remark to Theorem 3.3 we may construct the tiling of this strip and its vertical translations as well.

This then produces our tiling in this case.

Otherwise, using the above remarks, we may for each $W+[m, n]$ create a path from V to W and then to $W+1, W+2, \ldots, W+m, W+[m, 1], \ldots, W+[m, n]$ with adjacent corresponding squares being vertically adjacent and the x-coordinates monotonically increasing. Then we may extract the x-coordinate of W using the same method as in the proof of Lemma 4.6.

Then, as in the proof of Theorem 4.4, we may obtain the positions of any square in a fixed vertical strip bounded by $x = m$ and $x = m+N$, for $N > 1+4\delta$, where δ is the maximum of the lengths of all the minimal paths between V and $V + i$, for $V \in T_0$.

Combining this with the proofs of Lemmas 4.7 and 4.8, we may determine the position of any square. Note that this process is finite. Indeed, all steps in the proof of Lemma 4.7 are finite, because to determine the S_i' in the proof can be done by looking at all neighbors of S_{i-1}'.

In the proof of Lemma 4.8, borrowing the notation, the only possibly infinite step is in determining whether (i) or (ii) holds. But we may continue to determine the locations of S_i until we have at least $|T_0| + 1$ consecutive S_i with the same x-coordinate, where $|T_0|$ denotes the number of vertices of T_0, or have found that we are in case (i).

In the former, it is easy to see that we are in case (3) of Theorem 4.4. We may also find some $p, q \in \mathbb{Z}$ for which $p\lambda + q = 0$. Therefore, we may take T_1 to be the union of all sets $T_0 + r\lambda + s$, $0 \leq r < p$ and $0 \leq s \leq q$. Then we may use T_1 as the set which when translated horizontally and vertically yields the whole plane. In this way, we have altered our λ so that it is of the form $[0, t]$. Then it is easy to see that in this case we may use the same method as in the proof of Theorem 4.5 to determine the y-coordinates of each square. Combining yields the positions of every square in T_1 and knowing the translations, this completes the construction.

In the latter, we continue with the process outlined in the proof to completely determine the positions of any square.

Using the above procedure, we may determine the position of squares $V', V'+\lambda$ corresponding to vertices $V, V + i$ in T_0. This thus gives us the value of λ. Then for any $V \in T_0$, we may identify the position of the corresponding square V'.

Since any vertex in T may be written as $V + m + ni$ for some $V \in T_0$, the position of the corresponding square is $V' + m + n\lambda$, where V' corresponds to V. Since we know λ and the position of V', this yields the position of every square, completing the construction in all cases.

We are therefore able to construct the tiling if we can deduce the metric from the contacts graph. While this is difficult due to the definition of the extremal metric (it is unclear how to quickly solve the complex extremal problem), we have this following remark, with which we end this section.

Remark 7. *We may approximate the values of the extremal metric in a naive way. Indeed, first take any metric of norm 1 and evaluate the length, L. Then, take all possible metrics m such that $m(V)$ for any vertex V is an integer multiple of r, where r is some small rational number. Taking the metric that gives the maximal length over all of these yields a likely admissible approximation for the extremal metric.*

5 Tilings of Translation Surfaces

The doubly periodic tilings explored in Section 4 can be equivalently thought of as tilings of a torus, as identifying the opposite sides of a parallelogram results in a torus. This motivates us to consider tiling *translation surfaces*, a natural generalization of the torus.

5.1 Translation Surfaces: Definition and Examples

Roughly speaking, a *translation surface* is equivalent to a collection of polygons. To formalize this concept, there are many equivalent definitions, of which we use the third given in [4].

Definition 8. *Let $\mathcal{P}_i$, $1 \leq i \leq n$ be polygons with an even number of sides such that there exists a pairing of the sides satisfying the condition that any two paired sides are parallel and of equal length. Then identifying opposite edges in each $\mathcal{P}_i$ and taking the union of the polygons yields a translation surface.*

Remark 9. *The torus is a translation surface, because it can be obtained through taking any parallelogram as the polygon $\mathcal{P}_1$.*

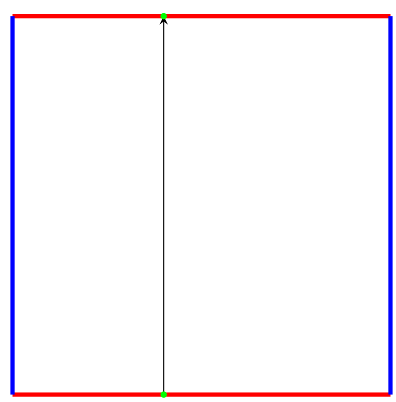

The torus as a translation surface; note that the two green points are identified.

We now analyze a more complex translation surface, namely, the octagonal translation surface.

5.2 The Octagonal Translation Surface

First we give the formal definition of this surface.

Definition 10. *The octagonal translation surface is the translation surface $\mathcal{H}$ obtained by identifying opposite edges of a regular octagon.*

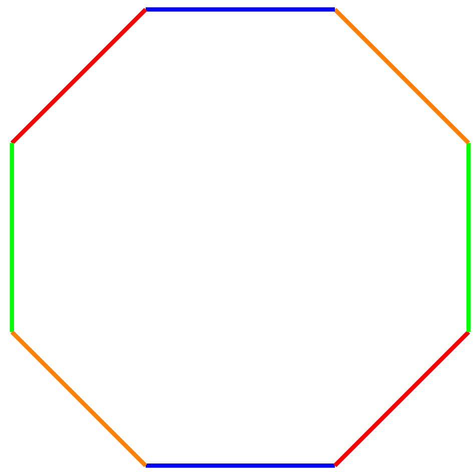

The octagonal translation surface. Opposite sides are identified with the same color.

One natural question, analogous to the existence of tori tilings, is whether this surface may be tiled with squares. We first need the concept of Hamel bases.

Definition 11. *A Hamel basis of the real numbers is a countably infinite set* $\{x_r\}$*,* $r \in R$*, such that each real number* x *can be uniquely written as a sum*

$$x = \sum_{s \in S} x_s$$

for a finite set $S \subset R$*.*

The following theorem, whose proof uses ideas from Freiling and Rinne [1], resolves this question in the negative.

Theorem 5.1. *There does not exist a square tiling of* $\mathcal{H}$*.*

Proof. Assume that $\mathcal{H}$ has side length $\sqrt{2}$. Let f be any additive function, and define the f-area of an isosceles triangle with legs a, a to be $\frac{f(a)^2}{2}$ and the f-area of any rectangle with sides a, b to be $f(a)f(b)$. Suppose that a figure $\mathcal{F}$ can be cut into a union of pieces $\bigcup \mathcal{F}_i$. Let the f-area of any figure be the sum of the f-areas of $\mathcal{F}_i$. Because f is additive, the f-area is well-defined and it is invariant under cutting and rearranging pieces.

Suppose that $\mathcal{H}$ can be tiled by squares. Then although f does not necessarily take on positive values, the f-area of $\mathcal{H}$ must be nonnegative because it each square has area $f(x)^2 \geq 0$. On the other hand, as shown in Figure 1, we may rearrange two isosceles right triangles of side length 1 to change $\mathcal{H}$ into the union of two 1 by $1 + \sqrt{2}$ rectangles and one $\sqrt{2}$ by $2 + \sqrt{2}$ rectangle.

Figure 1: Rearranging two isosceles triangles.

In addition, remark that 1 and $\sqrt{2}$ are linearly independent over $\mathbb{Q}$ as $\sqrt{2}$ is irrational. Thus, there exists a Hamel basis with these two as basis elements, and because all additive functions are completely determined by its value on the Hamel basis, we may set $x = f(1)$ and $y = f(\sqrt{2})$. Therefore, it suffices to find suitable x, y such that the f-area of $\mathcal{H}$ is negative, whereupon we arrive at a contradiction. However, the f-area of $\mathcal{H}$ equals $2x(x+y) + y(2x+y) = 2x^2 + 4xy + y^2$. Therefore, setting $x = 1$ and $y = -1$ results in an f-area of $-1 < 0$. This provides the desired contradiction, hence $\mathcal{H}$ cannot be tiled by squares. $\square$

Remark 12. *The use of Hamel bases in the above proof shows the motivation and the ideas behind this method. However, if desired, one may avoid this notion altogether because it suffices to assign values to $f(k)$ where k is one of finitely many side lengths that appear in our tiling and the rearrangement. Because 1 and $\sqrt{2}$ are linearly independent over $\mathbb{Q}$, we may assign $x = f(1)$ and $y = f(\sqrt{2})$, and then for any remaining elements assign them according to the additivity condition. Then we may proceed as above and complete the proof.*

However, recall that in the tori tilings, we allowed for tilings periodic under $T + 1$ and $T + \lambda$. Equivalently, we allowed the torus to undergo an affine transformation. Let us consider a similar idea for the octagonal translation surface. Specifically, we examine the regular octagon under a vertical stretch by a factor of l. Call the translation surface generated by this octagon $\mathcal{H}_l$.

Similar methods as in the proof of Theorem 5.1 show that for many values of l, $\mathcal{H}_l$ is not tileable.

Theorem 5.2. *There does not exist a square tiling of $\mathcal{H}_l$ if l is of the form $a + b\sqrt{2}$, where a, b are rational, and $|a| > |b\sqrt{2}|$.*

Proof. We follow the method used in the proof of Theorem 5.1. Maintain the same definition of f-area and also set $x = f(1)$ and $y = f(\sqrt{2})$.

Using the rearrangement and tiling given by 1, the f-area of the octagon is

$$2xf((1+\sqrt{2})l) + yf((2+\sqrt{2})l).$$

If $l \neq a + b\sqrt{2}$ for rational a, b, then the representation of l in the Hamel basis includes another real number r. Then 1, $\sqrt{2}$, $(1+\sqrt{2})l$, and $(2+\sqrt{2})l$ are linearly independent over $\mathbb{Q}$. Therefore we may choose x, y, $f((1+\sqrt{2})l)$, and $f((2+\sqrt{2})l)$ arbitrarily. Hence setting $x = y = 1$, $f((1+\sqrt{2})l) = 1$, and $f((2+\sqrt{2})l) = -100$ makes the f-area negative, implying as in the proof of 5.1 that no square tiling exists.

If $l = a + b\sqrt{2}$ for rational a, b, then the f-area may be rewritten as

$$2x((a+2b)x + (a+b)y) + y((2a+2b)x + (2b+a)y),$$

or

$$(2a+4b)x^2 + (4a+4b)xy + (2b+a)y^2.$$

There exist x, y such that this value is negative if $(4a+4b)^2 > 4(2a+4b)(a+2b) \iff 8a^2 > 16b^2$, or $|a| > |b\sqrt{2}|$. Therefore, $\mathcal{H}_l$ cannot be tiled if $|a| > |b\sqrt{2}|$. Combined with the previous case, this concludes the proof. □

However, there do exist certain values of l for which $\mathcal{H}_l$ is tileable.

Theorem 5.3. *There exist square tilings of $\mathcal{H}_l$ for $l = r(1+\sqrt{2})$ and $l = r\sqrt{2}$, where r is a rational number.*

Proof. Assume $r = 1$; for general r, the vertical stretch implies that if $\mathcal{H}_k$ can be tiled, then $\mathcal{H}_{rk}$ can be tiled by a by ar squares. Since r is rational, the result follows.

Now we refer to Figures 2 and 3 below. The former provides a tiling for $l = 1 + \sqrt{2}$ via three squares. This tiling uses the rectangle tiling obtained via the triangle rearrangement. These rectangles are all similar, so a vertical shift renders them as squares. The latter one provides a tiling for $l = \sqrt{2}$ via only two squares. □

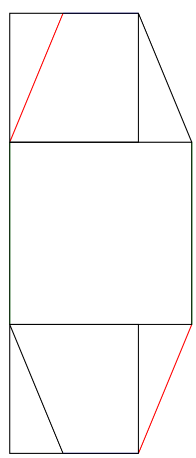

Figure 2: Three squares tiling the translation surface obtained from vertically stretching the octagonal translation surface by a factor of $l = 1 + \sqrt{2}$.

We may thus conclude that while the octagonal translation surface $\mathcal{H}$ may not be square-tiled, there exists some affine transformations of $\mathcal{H}$ which may

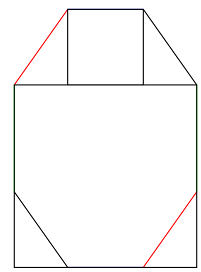

Figure 3: Two squares tiling the translation surface obtained from $l = \sqrt{2}$.

be square-tiled. This paves the way for further work on tiling this translation surface.

6 Further Research

There are several avenues of continuation from this work. Firstly, one may strive for a proof of Theorem 4.1 using the construction methods outlined in this paper. Unfortunately, the lack of a simple formula for the locations of the squares similar to that in the proof of Theorem 5.1 in [2].

In addition, our work on the octagonal translation surface $\mathcal{H}$ indicates a natural continuation of working with triangulations on an affine transform of $\mathcal{H}$. The most general (and a very intriguing) question one may ask is whether Theorems 4.1 and 4.4 apply in the case of the octagonal translation surface or translation surfaces in general. Results concerning this would generalize the theorems presented here and in [3] and [2].

7 Acknowledgements

I would like to thank my mentor, Professor Sergiy Merenkov, for suggesting this project and for his infinite guidance and support throughout, without whom this work would not have been possible. I also wish to thank Dr. Tanya Khovanova and Dr. Kent Vashaw for providing numerous invaluable edits, as well as Professors Slava Gerovitch and Pavel Etingof for giving me the wonderful opportunity to participate in the PRIMES-USA program where this research was conducted.

References

[1] C Freiling and D Rinne. Tiling a square with silimar rectangles. *Math. Res. Let.*, 1(5):547–558, 1994.

[2] Oded Schramm. Square tilings with prescribed combinatorics. *Isr. J. Math.*, 84(1):97–118, 1993.

[3] Oded Schramm. Conformal uniformization and packings. *Isr. J. Math.*, 1996.

[4] Alex Wright. Translation surfaces and their orbit closures: an introduction for a broad audience. *EMS Surv. Math. Sci.*, 2(1):63–108, 2015.