Families of Infinite Translation Surfaces related to the Chamanara Surface

Zur Erlangung des akademischen Grades eines

DOKTORS DER NATURWISSENSCHAFTEN

von der KIT-Fakultät für Mathematik des
Karlsruher Instituts für Technologie (KIT)
genehmigte

DISSERTATION

von

Jan Kohlmüller

Tag der mündlichen Prüfung: 09. Februar 2022
Referenten: Prof. Dr. Frank Herrlich
Prof. Dr. Gabriela Weitze-Schmithüsen

Contents

1 Introduction

The Chamanara surface, first described by Reza Chamanara in [Cha04] is one of the most well-known examples of an infinite translation surface. Infinite translation surfaces is a relatively recent topic which has not been the focus of research up until the 21st century. Classical translation surfaces on the other hand have been studied extensively since the last quarter of the 20th century.

The most intuitive way of constructing a translation surface (in the classical sense) is by taking a finite number of disjoint polygons and gluing parallel edges of the same length and different orientation together via translations. By doing this for all edges we obtain a 2-dimensional oriented real manifold. Furthermore, for all charts which do not contain any of the former vertices, the transition functions are just translations. The points correlating to the former vertices are called *singularities* and each of them is a cone point of the surface with cone angle $2k\pi$ for some $k \in \mathbb{N}$. This k is also called the *order* of the singularity. This leads to another way of defining a translation surface, namely as a 2-dimensional oriented compact connected manifold for which the transition functions are translations apart from a finite number of cone singularities. There is also a third common way of defining a translation surface: A translation surface can be defined as a pair (X, ω), where X is a connected compact Riemann surface and ω is a nonzero holomorphic differential on X.

Translation surfaces possess many different interesting properties and are therefore related to various different fields. Nevertheless, in this work, we are mainly interested in the moduli spaces of translation surfaces. In the classical situation, these appear naturally as a vector bundle over the space $\mathcal{M}_g$ of Riemann surfaces of genus g and are stratified by the order of the singularities. The strata then have a complex orbifold-structure, which goes back to Veech ([Vee86]). This structure is realized by local coordinate functions called the period coordinates. In the last 20 years, several important results have been found concerning the structure of the strata. One of these was the research done by Kontsevich and Zorich in [KZ03], which showed that almost all of the strata are connected. Surely the most groundbreaking research was the one done by Eskin, Mirzakhani, and Mohammadi in [EMM15] and Eskin and Mirzakhani in [EM18], which led to Mirzakhani winning the Fields Medal in 2014. In these articles, they established a description of the orbit closures of the action of $\mathrm{GL}_2^+(\mathbb{R})$ on a stratum.

That being said, very little is known about the structure of the space of all infinite translation surfaces. Similar to the classical case an infinite translation surface can be obtained by gluing infinitely many polygons together instead of finitely many. Another method of constructing infinite translation surfaces is by allowing the polygons to have infinitely many edges. Again an infinite translation surface can be defined as a 2-dimensional oriented connected manifold X for which the transition functions are translations. In contrast to the finite case, X is not necessarily compact. In addition, the singularities are not part of the manifold X anymore but are rather found in the metric completion $\overline{X}$. Because we focus mainly on infinite translation surfaces we just

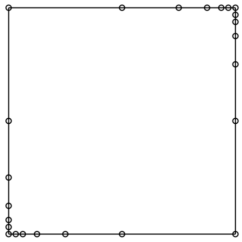

Figure 1: The polygon used to construct the Chamanara surface.

call them *translation surfaces* instead of "infinite translation surfaces" and call the classical translation surfaces *finite translation surfaces* instead. This is supported by the fact that each finite translation surface (after removing the singularities) is also an infinite translation surface.

The space of all infinite translation surfaces lacks the structure of the moduli spaces of finite translation surfaces. There have been different attempts at overcoming this issue either by the means of Veech groups or covering spaces (see e.g. [Tre14], [HT19] or [DHL14]). In [Hoo13] and [Hoo18] Patrick Hooper uses a different approach by using immersions and embeddings of translation surfaces to define a topology on the space $\mathcal{M}$ of all translation surfaces (finite and infinite ones). This topology is called the *immersive topology* and it will be our main tool to make statements about the structure of families of infinite translation surfaces. In addition, building up on this topology, we will define a finer topology on $\mathcal{M}$.

One recent attempt at investigating the structure of the space of all translation surfaces has been done by Rodrigo Treviño in [Tre18] building up on his work done with Kathryn Lindsey in [LT16]. There they introduced a way of constructing a translation surface out of a bi-infinite Bratteli diagram. This allows a topology defined on the set of Bratteli diagrams to be used on the set of translation surfaces constructed out of these Bratteli diagrams. Our approach follows a similar spirit. Instead of Bratteli diagrams we use point-symmetric polygons with infinitely many edges (called *PSI-polygons*) as a way to define coordinates on the set of translation surfaces constructed out of these polygons. In particular, we show that these coordinates agree with the immersive topology introduced in [Hoo13] and [Hoo18].

The basic idea of this procedure stems from the Chamanara surface mentioned in the beginning. This surface is constructed by taking the polygon depicted in Figure 1 and gluing diagonally opposing edges together. This leads to an infinite translation surface with infinite genus and one singularity. In [Cha04] the author does not only describe one surface but a whole family of surfaces. The polygon depicted in Figure 1 is constructed by dividing each side in half and then again dividing one of the resulting parts in half and so forth. This can be generalized by replacing the factor $\frac{1}{2}$ with another factor $\alpha \in (0,1)$ to receive a translation surface Ch_α. This results in the 1-parameter family introduced in [Cha04]. One starting idea will be to use this α as a coordinate function for this family. We then generalize this approach further to define similar coordinate functions on much larger families.

We will see that these coordinate functions agree with the immersive topology, i.e. they are homeomorphisms from these families to a suitable coordinate space. But as it turns out the immersive topology is not sufficient for coordinate functions of more complex families. The main problem is that this topology is too coarse to perceive certain properties. For example, the genus is not preserved under limits of sequences. The same holds for the area of surfaces: There are sequences of translation surfaces of infinite area which converge to a surface of finite area in the immersive topology. We will address this issue by introducing a slightly stronger topology on the space of all translation surfaces which we will therefore call the *strong immersive topology*. We will see that this topology is sufficient to resolve these problems. We also show that this topology is not too strong for our case, i.e. the coordinate functions are still homeomorphisms with respect to the strong immersive topology.

The structure of this thesis is as follows. In Section 2 we recall the basic definitions used in the later sections. In addition, we give a short summary of the structure of the moduli spaces of finite translation surfaces and introduce some examples of families of finite translation surfaces as a motivation for our further procedure.
In Section 3 we state the definition of the immersive topology as defined in [Hoo13] and [Hoo18] and repeat some basic facts from these works, which we use in later sections.
In Section 4 we start by introducing our basic tool, point-symmetric infinite polygons (PSI-polygons), and construct translation surfaces out of these polygons. This is done as a generalization of the Chamanara surface. Afterward, we investigate some basic properties and define convergence for these polygons. We then conclude that the convergence of PSI-polygons implies the convergence of the corresponding translation surfaces in the immersive topology in Theorem 4.18:

Theorem 4.18. *If $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ converges to $(P, (x_m)_{m\in\mathbb{N}})$, then $(\widehat{P}^{(n)})_{n\in\mathbb{N}}$ converges to $\widehat{P}$ in the immersive topology.*

In Section 5 we then establish several families of translation surfaces which are constructed out of PSI-polygons. Among other things we show in Theorem 5.1 that α as introduced above is, in fact, suitable as a coordinate:

Theorem 5.1. *The map*

$$\Psi\colon (0,1) \to \mathcal{M},\ \alpha \mapsto \mathrm{Ch}_\alpha$$

is a homeomorphism onto its image.

In Theorem 5.1 $\mathcal{M}$ is equipped with the immersive topology but Theorem 6.16 will show that the statement still holds for the strong immersive topology. In the other main statements of Section 5, namely Theorem 5.9 and Theorem 5.16, we introduce coordinates

for two much larger families. The first is an infinite-dimensional real family, the second a finite-dimensional complex family.
At last, in Section 6, we introduce the strong immersive topology on the space of all translation surfaces $\mathcal{M}$. This topology can be seen as a fusion of the immersive topology and the Gromov-Hausdorff topology for metric spaces. We investigate some basic properties of this topology and highlight some cases for which this topology has some advantages over the immersive topology. In Theorem 6.16 we show that an analogue of Theorem 4.18 still holds for the strong immersive topology. Finally, in Theorem 6.23 we establish an infinite-dimensional complex family.

2 Basics of translation surfaces

2.1 Fundamental definitions

This section will be devoted to repeating some basic facts about translation surfaces. We will only give a very brief introduction, mainly to fix notation and to give some motivation for the ideas used in the next sections. For a more detailed and beginner-friendly introduction see e.g. [Mas22] or [Zor06]. First, let us clarify how a translation surface is defined in our case. Most works regarding translation surfaces define a translation surface as the construct we will call "finite translation surface". But in our case, the term translation surface also includes those surfaces often referred to as "infinite translation surfaces". Nevertheless, we will also give a short introduction into the theory of finite translation surfaces in Section 2.2.

Definition 2.1. A *translation surface* is a pair $(X, \mathcal{A})$, where X is a connected oriented 2-dimensional manifold and $\mathcal{A}$ is a maximal translation atlas, i.e. a maximal atlas where all transition functions are translations.

These objects are often called infinite translation surfaces. Please remark that we do not require X to be a closed manifold. For example, every open path-connected subset of $\mathbb{C}$ is a translation surface. Often we will only write X for a translation surface if $\mathcal{A}$ is clear from the context. Throughout this work, we will often use $\mathbb{R}^2$ and $\mathbb{C}$ interchangeably (by identifying the point $(a, b) \in \mathbb{R}^2$ with the point $a + ib \in \mathbb{C}$) when the multiplication on $\mathbb{C}$ is not needed.

For a metric space X we will denote by $\overline{X}$ the metric completion of X. For a translation surface $(X, \mathcal{A})$, the translation atlas $\mathcal{A}$ gives rise to a unique metric on X.

Definition 2.2. Let $(X, \mathcal{A})$ be a translation surface. A *singularity* of X is a point in $\overline{X} \setminus X$.

In [Cha04] the author describes a family of translation surfaces. Each of these surfaces is often called *Chamanara surface* or sometimes *baker's map surface* and will be described in the following example.

Example 2.3 (Chamanara surface)**.** Let $P = [0, 1]^2$ be the closed unit square and $\alpha \in (0, 1)$. We divide the upper side of this square into segments

$$\begin{aligned} x_1 &= [0, 1 - (1-\alpha)^1] \times \{1\}, \\ x_2 &= [1 - (1-\alpha)^1, 1 - (1-\alpha)^2] \times \{1\}, \\ x_3 &= [1 - (1-\alpha)^2, 1 - (1-\alpha)^3] \times \{1\}, \\ &\vdots \end{aligned}$$

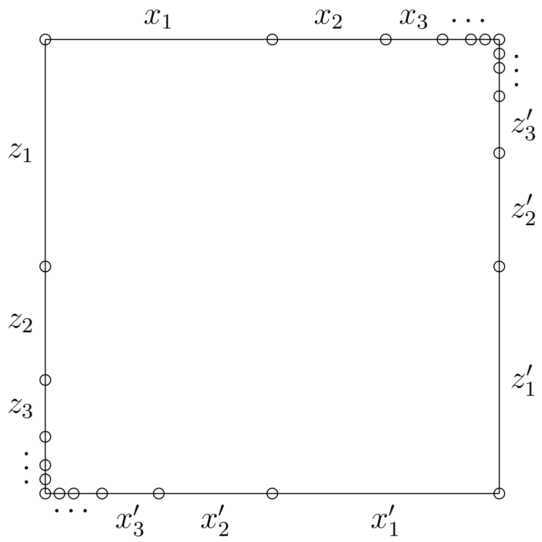

Figure 2: The "standard" Chamanara surface $\mathrm{Ch}_{\frac{1}{2}}$. The small circles correspond to the endpoints of the segments or, equivalently, to the singularity of the surface. In this work, we will always mark such points by circles.

We then divide the other sides in the same manner, see Figure 2 for a picture with $\alpha = \frac{1}{2}$. In the end, we remove the endpoints of each of these segments and glue each side x_i to the diagonally opposing side x'_i and each side z_i to the diagonally opposing side z'_i via a translation. The resulting surface is an infinite translation surface. Let us now look at the endpoints of these segments. These are not part of the surface anymore but they still lie in the metric completion. So they correspond to the singularities of this surface. Through the identification of the segments these points get identified into four points, two on the horizontal and two on the vertical sides. But these points are actually the same because they get infinitely close to each other at the upper right corner. Therefore this translation surface has only one singularity but infinite genus. In this work, we will denote this surface by Ch_α.

The Chamanara surface has been a recurring example in many works that deal with infinite translation surfaces with finite area (see e.g. [Ran16], [HR16] or [HT19]). In this work, we will use the Chamanara surface as some kind of starting point for an attempt to better understand the moduli space of translation surfaces.

For technical reasons, we will mostly regard pointed translation surfaces, i.e. translation surfaces with a marked point $o_X \in X$. We will denote such a translation surface by $(X, \mathcal{A}, o_X)$ or sometimes only by (X, o_X) or even only by X.

Definition 2.4. Two translation surfaces $(X, \mathcal{A})$ and $(Y, \mathcal{B})$ are isomorphic if there is a homeomorphism $\varphi\colon X \to Y$ such that for each $\psi \in \mathcal{A}$ it holds that $\psi \circ \varphi^{-1} \in \mathcal{B}$. If those translation surfaces are pointed with basepoints o_X and o_Y, we also require that

$\varphi(o_X) = o_Y$. Such a homeomorphism will be called a *translation isomorphism* or often just an isomorphism.

Definition 2.5. The *moduli space $\mathcal{M}$ of all pointed translation surfaces* is the set of all isomorphism classes of pointed translation surfaces.

For now, $\mathcal{M}$ is only a set. The notion of $\mathcal{M}$ as a space will be justified in Section 3 when we define a topology on $\mathcal{M}$.

In most cases, we will not differentiate between a translation surface X and the isomorphism class of X and therefore regard X as an element of $\mathcal{M}$.

Definition 2.6. The *total space of translation surfaces $\mathcal{E}$* is the set of all isomorphism classes of two-pointed translation surfaces, i.e. translation surfaces with two (not necessarily distinct) marked points.

The space $\mathcal{E}$ will be seen mostly as the space of pairs of a translation surface X together with a point $x \in X$. One of the marked points of these two-pointed translation surfaces corresponds to the basepoint, the other one to the point x. Or to be more precise:

Remark 2.7. *There is a natural projection*

$$\pi\colon \mathcal{E} \to \mathcal{M},\ (X, \mathcal{A}, o_X, x) \mapsto (X, \mathcal{A}, o_X)$$

and each fiber is as a set canonically bijective to a translation surface. Via this bijection, each fiber can then be endowed with a translation surface structure. As before we will often write (X, x) instead of $(X, \mathcal{A}, o_X, x)$.

As a last point in this section, we want to introduce the concept of a developing map which we will use on several occasions in the next sections. This concept exists in the more general case of (G, X)-structures, but we will only state the adaption to our case. For the more general language see [Thu97][§3.4].

Definition 2.8. Let $(X, \mathcal{A})$ be a simply connected translation surface. The *developing map* is the unique map $\mathrm{dev}_X\colon X \to \mathbb{C}$ with $\mathrm{dev}_X(o_X) = 0$ and such that for each $\varphi \in \mathcal{A}$, dev_X differs from φ on the domain of φ only by a translation.

The uniqueness of this map follows from analytic continuation. Basically, we start at the basepoint and a chart of a neighborhood of this point, so the image of dev_X is clear for every point of this neighborhood. We then continue with every chart whose domain intersects this neighborhood. Because X is connected this can be continued for all of X. Because the developing map can only be defined for simply connected translation surfaces it is often helpful to consider the universal cover of X.

Remark 2.9. *Let $(X, \mathcal{A})$ be a translation surface and C be a covering space of X together with the covering map $p\colon C \to X$. Then C has a natural structure as a translation surface via the atlas $\{\varphi \circ p \mid \varphi \in \mathcal{A}\}$. In particular, the universal cover of X is a simply connected translation surface. We will denote this translation surface by $\widetilde{X}$.*

We will mostly denote the points of the universal cover by $\widetilde{x}$ and paths whose image lies in $\widetilde{X}$ by $\widetilde{\gamma}$ and so on, to better distinguish whether an object belongs to X or $\widetilde{X}$.

For a pointed translation surface $(X, \mathcal{A}, o_X)$ the universal cover $\widetilde{X}$ becomes a pointed translation surface via the choice of a basepoint $o_{\widetilde{X}}$. This will always be done in a way such that $p(o_{\widetilde{X}}) = o_X$, where p is the universal covering map. This choice of a basepoint for $\widetilde{X}$ is then unique up to isomorphism.

2.2 Finite translation surfaces

Finite translation surfaces will not be the main focus of this work. Nevertheless, we will give a short summary of finite translation surfaces in this section. Our focus will be on the structure of the moduli space of finite translation surfaces. This will give us some starting points and motivation for finding coordinates for appropriate families of infinite translation surfaces. Let us start by defining what a finite translation surface is.

Definition 2.10. Let $(X, \mathcal{A})$ be a translation surface. If X has only finitely many singularities and $\overline{X}$ is a compact surface, then $\overline{X}$ is called a *finite translation surface*.

Often we will also call X a finite translation surface instead of $\overline{X}$. Sometimes we will call a translation surface an infinite translation surface if we want to emphasize that it is not a finite translation surface.

The above is not the only common way of defining finite translation surfaces. In fact, there are 3 common ways of defining finite translation surfaces as they appear for example in [Mas06]. The most constructive is via gluings of polygons, namely the following:

Remark 2.11. *Let $\mathcal{P}$ be a finite set of disjoint polygons (seen as subsets of $\mathbb{C}$), $E(\mathcal{P})$ be the set consisting of all edges of polygons in $\mathcal{P}$ and $\varphi\colon E(\mathcal{P}) \to E(\mathcal{P})$ a pairing, such that for all $x \in E(\mathcal{P})$ the edges x and $\varphi(x)$ differ only by a translation. In addition, the polygon that x belongs to and the polygon $\varphi(x)$ belongs to should be on different sides of these edges. Then, by gluing each $x \in E(\mathcal{P})$ to $\varphi(x)$ by the above translations, we get a finite translation surface.*

In addition, similar as we did in Example 2.3 when we remove the vertices from these polygons the resulting surface is a translation surface as defined in Definition 2.1. So we can assume that the singularities of these translation surfaces correspond to the vertices of the polygons. This also shows that the angle around such a singularity is a multiple of 2π.

The third way of defining finite translation surfaces is via abelian differentials. For a Riemann surface X we denote by $\Omega(X)$ the set of non-zero abelian differentials.

Definition 2.12. A *finite translation surface* is a pair (X, ω) where X is a compact, connected Riemann surface and $\omega \in \Omega(X)$.

Please note that this X corresponds to $\overline{X}$ in Definition 2.10. One can show that these three definitions are in fact equivalent. For a proof see e.g. [Mas22] or [Mas06]. In Definition 2.12 the singularities correspond to the zeroes of the abelian differential. The angle around such a singularity is then exactly $2\pi \cdot (k+1)$, where k is the multiplicity of this zero. We will call such a singularity a *singularity of order* k.

Note that an abelian differential ω allows us to make sense of the notion $\int_\gamma \omega$ for any path γ on X. The value of $\int_\gamma \omega$ is then just a complex number.

Now let X be a finite translation surface and g be the genus of X. Furthermore let $k_1, \ldots, k_n$ be the orders of the singularities of X. Then, analogue to the Gauß-Bonnet-formula it holds that

$$2g - 2 = \sum_{i=1}^{n} k_i.$$

This fact leads to a stratification of the moduli space. For this let us first take a look at the moduli spaces of Riemann surfaces. Let $\mathcal{M}_g$ be the set of all isomorphism classes of all Riemann surfaces of genus g. Then $\mathcal{M}_g$ has a well-established topology and is a $3g-3$-dimensional analytic space (see e.g. [IT92]). Furthermore, the moduli space consisting of pairs (Riemann surface, abelian differential) forms a vector bundle over $\mathcal{M}_g$ and corresponds to the moduli space $\Omega\mathcal{M}_g$ of translation surfaces of genus g. This fact yields a natural topology on this space[1].

This space is then stratified according to the number and multiplicity of the zeroes of the abelian differential. For this let $k_1, \ldots, k_n$ be a partition of $2g-2$. Then we define the stratum $\mathcal{H}(k_1, \ldots, k_n)$ to be the set of all translation surfaces where the orders of the singularities are exactly $k_1, \ldots, k_n$. From now on let $\kappa = (k_1, \ldots, k_n)$ so we can simply write $\mathcal{H}(\kappa)$ instead of $\mathcal{H}(k_1, \ldots, k_n)$. In most cases, only the topology of one stratum is considered and not the topology of all of $\Omega\mathcal{M}_g$. This is done because the topology of $\mathcal{H}(\kappa)$ has a nice structure through the period coordinates:

Definition 2.13. Let $(X, \omega) \in \mathcal{H}(\kappa)$ be a finite translation surface as defined in Definition 2.12. In addition let Σ be the set of singularities of (X, ω) and $\{\gamma_1, \ldots, \gamma_{2g+n-1}\}$

[1] There is also another method of defining a topology on this space via the developing map. This approach is found e.g. in [FM14] or [Yoc10].

be a basis of the relative homology group $H_1(X, \Sigma; \mathbb{Z})$. The *period coordinates* of (X, ω) are the complex numbers

$$\left(\int_{\gamma_i} \omega \right)_{i=1}^{2g+n-1} \in \mathbb{C}^{2g+n-1}.$$

The period coordinates provide $\mathcal{H}(\kappa)$ with the structure of a $(2g + n - 1)$-dimensional orbifold. To see this, it is necessary to have some way of "carrying over" the basis of the homology to nearby translation surfaces, but we will not investigate this any further. A reader interested and not familiar with this topic may find a proof of the fact that $\mathcal{H}(\kappa)$ is in fact a $(2g + n - 1)$-dimensional orbifold in [Wri15], [FM14] or [Yoc10]. Note that the Gauß-Bonnet-formula shows that $\dim_{\mathbb{C}}(\mathcal{H}(\kappa)) = \sum_{i=1}^{n} k_i + n + 1$.

One thing to note is that for most translation surfaces (X, ω), there is a neighborhood of (X, ω) in $\mathcal{H}(\kappa)$ on which the period coordinates act as an actual chart, i.e. a homeomorphism from this neighborhood to an open subset of $\mathbb{C}^{2g+n-1}$. So for the most part, $\mathcal{H}(\kappa)$ has the structure of a manifold instead of an orbifold.

Before we end this section, let us look at an example of finite translation surfaces and see some possible period coordinates for this example. In [CGL06] the authors construct a family of finite translation surfaces Y_{nm} in a very similar way to the Chamanara surface by dividing the horizontal lines into n edges and the vertical lines into m edges. The next example will be a generalization of this construction and will show that there is a canonical way to create a translation surface out of a point-symmetric polygon. This idea will be heavily used in this work and Section 4 will generalize this example to the infinite case.

Example 2.14. Let P be a point-symmetric polygon. Let o_P be the point of symmetry and $\Phi \colon P \to P$ be the reflection on o_P. For each edge x of P we denote the mirrored edge by $x' := \Phi(x)$. Applying the point-symmetry and then the reflection on the midpoint of x' we obtain a translation that maps x to x'. Therefore we can glue each edge x to the mirrored edge x' and get a finite translation surface $\widehat{P}$.

Let N be the number of edges of P and $M = \frac{N}{2}$. The singularities of such a translation surface correspond to the vertices of the polygon and one can see that all vertices are glued into a single singularity if M is even and into two different singularities if M is odd. So it holds that

$$\widehat{P} \in \begin{cases} \mathcal{H}(M-2), & \text{if } M \text{ is even,} \\ \mathcal{H}(\frac{M-3}{2}, \frac{M-3}{2}), & \text{if } M \text{ is odd.} \end{cases}$$

In each case it holds that $\dim_{\mathbb{C}}(\mathcal{H}(M-2)) = M = \dim_{\mathbb{C}}(\mathcal{H}(\frac{M-3}{2}, \frac{M-3}{2}))$, so $\widehat{P}$ lies in a stratum of dimension M. For an edge x_n we denote by $s(x_n)$ the start- and by $t(x_n)$ the endpoint of x_n. For each pair of edges consisting of one edge x_n and the mirrored edge x'_n we may choose a path γ_n in $\widehat{P}$ from $s(x_n)$ to $t(x_n)$ with image x_n. Then

$\{\gamma_n \mid n \in \mathbb{N},\ n \leq M\}$ is a basis of the homology $H_1(\widehat{P}, \Sigma; \mathbb{Z})$ and the corresponding period coordinates are

$$(t(x_1) - s(x_1), \ldots, t(x_M) - s(x_M)) \in \mathbb{C}^M.$$

As we have seen, due to the period coordinates, the strata have a very nice structure that has been well studied in the past. Unfortunately the same cannot be said about infinite translation surfaces. In fact, not very much is known about the structure of the moduli space of all translation surfaces and no analogue to the period coordinates is known for this space until today. In this work, we will try to find at least some structures in this space by looking at some families of infinite translation surfaces and providing them with some reasonable coordinates, i.e. homeomorphisms from open subsets of these families to a suitable Banach space. The following section will provide some examples to explain the idea and motivation behind this.

2.3 A family of finite translation surfaces

Example 2.15. We construct a sequence of finite translation surfaces. The construction will be very similar to Figure 2. For $n \in \mathbb{N}$ let X_n be the following (finite) translation surface:
As in Figure 2, we take $P = [0,1]^2$, $\alpha = \frac{1}{2}$, and we use the same subdivision of the upper side to construct the edges $x_1, \ldots, x_n$. But instead of continuing this splitting indefinitely, we define x_{n+1} to be the remainder of the upper side. We repeat this process for each side of P and glue opposite edges together. See Figure 3 for an example for $n = 1$ and $n = 2$.

For $n \to \infty$ these finite translation surfaces become more and more similar to the Chamanara surface. So it is plausible, that the sequence given by these translation surfaces converges to $\mathrm{Ch}_{\frac{1}{2}}$ in some sense. But this argument is rather intuitive and not very precise. The main problem here is, that the genus of these surfaces gets larger for greater n and the topology on the moduli space is only given on each $\mathcal{M}_g$ and not for the whole space. In addition, $\mathrm{Ch}_{\frac{1}{2}}$ is an infinite translation surface of infinite genus. So it is not contained in any $\mathcal{M}_g$ at all. The next section will be devoted to fixing this problem by introducing a topology on the moduli space of all translation surfaces.

But before we do that, let us look at a variation of the above sequence.

Example 2.16. Let $(\lambda_n)_{n\in\mathbb{N}}, (\mu_n)_{n\in\mathbb{N}}$ be sequences in $\mathbb{R}_+$ such that

$$\sum_{i=1}^{\infty} \lambda_i = \sum_{i=1}^{\infty} \mu_i = 1.$$

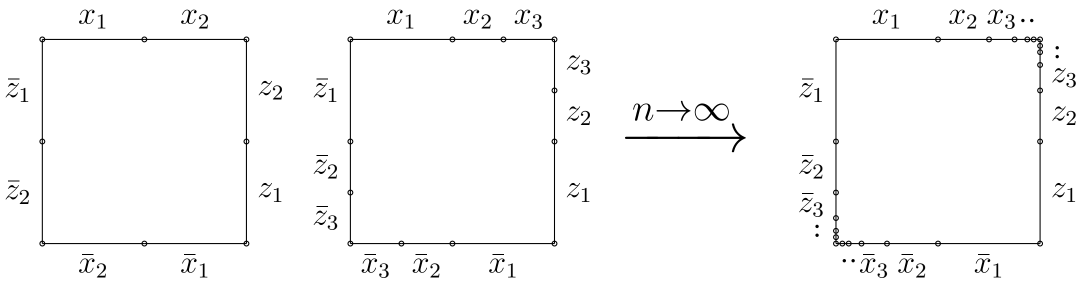

Figure 3: The translation surfaces X_n for $n = 1$ and $n = 2$. Each such translation surface is constructed from a polygon with $4 \cdot (n+1)$ edges. For $n \to \infty$ these surfaces become more and more like $\mathrm{Ch}_{\frac{1}{2}}$.

We construct the following translation surface: Let $P = [0,1]^2$. We again divide the upper side into the segments

$$x_1 = [0, \lambda_1] \times \{1\},\ x_2 = [\lambda_1, \lambda_1 + \lambda_2] \times \{1\},\ \ldots .$$

We divide the lower side in the same manner such that the result is point-symmetric for the point $\left(\frac{1}{2}, \frac{1}{2}\right)$. Next, we repeat this process for the vertical sides but with μ_i instead of λ_i. In the end, we glue each edge to its corresponding edge on the other side to get a translation surface X. So Ch_α is a special case of this construction with $\lambda_i = \mu_i = (1-\alpha)^{i-1}\alpha$.

As in Example 2.15, we may construct a sequence of finite translation surfaces X_n by doing the above construction up until x_n and defining x_{n+1} to be the rest of the upper side. We again repeat this for the vertical sides and glue the edges together accordingly. As in Example 2.14, we can choose a basis of the relative homology of X_n in such a way that the image of the period coordinates of X_n is

$$\left(\lambda_1, \ldots, \lambda_n, 1 - \sum_{i=1}^{n} \lambda_i, \mu_1 i, \ldots, \mu_n i, \left(1 - \sum_{i=1}^{n} \mu_i\right) \cdot i\right) \in \mathbb{C}^{2n+2}.$$

As before the sequence X_n converges to X in some sense. Therefore

$$(\lambda_1, \mu_1 i, \lambda_2, \mu_2 i, \ldots) \in \mathbb{C}^\infty$$

can be seen as reasonable coordinates for the translation surface X. The following sections will be devoted to creating a formal background for these coordinates. In Theorem 5.9 we will show that the above idea gives rise to a homeomorphism from the space of all translation surfaces that can be constructed in the above way to the space c_0 of all sequences converging to zero, equipped with the supremum norm.

3 Immersive topology

The goal of this section is to establish a topology on the set $\mathcal{M}$ of isomorphism classes of all translation surfaces. As mentioned in the previous section, the moduli spaces of finite translation surfaces have been the target of much research in the past. This cannot be said about infinite translation surfaces, but there are some results concerning this issue.

In [Bow13] a very similar approach is used as the one presented in Section 2.3. There the family of the so-called Arnoux-Yoccoz surfaces is constructed. These surfaces go back to the work done by Arnoux and Yoccoz in [AY81] and [Arn88]. For each $g \in \mathbb{N}$ one can construct such a finite translation surface (X_g, ω_g) of genus g. In [Bow13] a limit surface X_∞ is constructed, which is an infinite translation surface. In addition, for each $g \in \mathbb{N}$, a compact subsurface with boundary K_g of X_∞ and a piecewise-affine embedding $\iota_g \colon K_g \to X_g$ is given. It is also shown that $\iota_g^*(|\omega_g|)$ converges to $|\omega_\infty|$ on compact subsets of X_∞ as $g \to \infty$, where $|\omega_g|$ denotes the metric induced by ω_g on X_g. This is done to make clear that X_∞ can be seen as a limit space of the sequence $(X_g)_{g \in \mathbb{N}}$.

In [Hoo18] and [Hoo13] the author uses a very similar approach. While in [Bow13] the main goal is to investigate the above sequence, the goal of these works is to define a topology on all of $\mathcal{M}$. This approach uses the idea of embedding subsurfaces, too. The resulting topology is called the immersive topology and is the central tool used for most of this work. Therefore we now define this topology and review some properties of it. The following section mostly repeats facts stated in [Hoo18] and [Hoo13]. For a proof of these facts or a more detailed study please refer to these works. Let us first begin by defining what an immersion is.

Definition 3.1 ([Hoo18, §3.1]). Let $(X, \mathcal{A}, o_X)$ and $(Y, \mathcal{B}, o_Y)$ be pointed translation surfaces and $U \subset X$, $V \subset Y$ be path-connected subsets that contain the basepoint. An *immersion* from U to V is a continuous map $\iota \colon U \to V$ with $\iota(o_X) = o_Y$ which acts as a translation in local coordinates.

For the rest of this section, every subset of a translation surface is always path-connected and contains the basepoint.

We have already seen one prime example of an immersion, namely the following:

Example 3.2. For each simply connected translation surface X, the developing map is an immersion from X to $\mathbb{C}$.

For the same reasons as for the developing map, we can deduce uniqueness of immersions:

Proposition 3.3 ([Hoo18, Proposition 7]). *If there exists an immersion, then it is unique.*

We are mainly interested in the existence of such an immersion, so we write $U \rightsquigarrow V$ if there exists an immersion from U to V and $U \not\rightsquigarrow V$ if there exists no such immersion.

Graphically, the existence of an immersion just means that the first subset "fits" into the second one. This graphical explanation will be concretized by the following examples.

Example 3.4. If U and V are subsets of the same translation surface, then $U \rightsquigarrow V$ if and only if $U \subseteq V$.

This example also holds in a more general case. Let X, Y be simply connected translation surfaces and U and V be subsets as in Definition 3.1. If in addition $\mathrm{dev}_V := \mathrm{dev}_Y|_V$ is injective, then $U \rightsquigarrow V$ if and only if $\mathrm{dev}_U(U) \subseteq \mathrm{dev}_V(V)$. This is easy to see, because if $\iota\colon U \to V$ is an immersion, then $\mathrm{dev}_V \circ \iota\colon U \to \mathbb{C}$ is an immersion as well, so $\mathrm{dev}_V \circ \iota = \mathrm{dev}_U$ with Proposition 3.3 and therefore $\mathrm{dev}_U(U) \subseteq \mathrm{dev}_V(V)$. On the other hand if $\mathrm{dev}_U(U) \subseteq \mathrm{dev}_V(V)$ then dev_V^{-1} is also an immersion, so $\mathrm{dev}_V^{-1} \circ \mathrm{dev}_U\colon U \to V$ is an immersion.

Remark 3.5. *Immersions can also be examined at the level of paths: If $U \rightsquigarrow X$ and $U \rightsquigarrow Y$, then each path with image in U can be seen as a path in X or a path in Y.*

In most cases, the non-existence of an immersion can be graphically explained by the set "hitting" a singularity. The most simple such case is the following:

Example 3.6. We look at the two translation surfaces $(\mathbb{C}, 0)$ and $(\mathbb{C} \setminus \{x\}, 0)$ for some $x \in \mathbb{C} \setminus \{0\}$. Then $\mathbb{C} \not\rightsquigarrow \mathbb{C} \setminus \{x\}$ because the immersion would have to map x to x.

Definition 3.7. An injective immersion will be called an *embedding* and the existence of an embedding will be marked by $U \hookrightarrow V$.

Of course not every immersion is an embedding:

Example 3.8. Let T be the torus with an arbitrary marked point. Then $\mathbb{C} \rightsquigarrow T$, but of course $\mathbb{C} \hookrightarrow T$ does not hold.

As a last point before we are able to define the immersive topology, we need to define the set PC. For a translation surface (X, o_X) the set $\mathrm{PC}(X)$ is the set of all path-connected subsets containing the basepoint. Now let (Y, o_Y) be another translation surface. $A \in \mathrm{PC}(X)$ and $B \in \mathrm{PC}(Y)$ are *isomorphic* if there is a bijective immersion $A \rightsquigarrow B$. The set PC is then just the set of isomorphism classes in $\bigcup_{X \in \widetilde{\mathcal{M}}} \mathrm{PC}(X)$.

Definition 3.9 ([Hoo13, §3.2])**.** The *immersive topology* on $\mathcal{M}$ is the coarsest topology such that all sets of the following forms are open:

(i) The set $\mathcal{M}_{\rightsquigarrow}(D) = \{X \in \mathcal{M} \mid D \rightsquigarrow X\}$ for every $D \in \mathrm{PC}$ homeomorphic to a closed disc.

(ii) The set $\mathcal{M}_{\not\rightsquigarrow}(U) = \{X \in \mathcal{M} \mid U \not\rightsquigarrow X\}$ for every $U \in \mathrm{PC}$ homeomorphic to an open disc.

(iii) The set $\mathcal{M}_{+}(D, U) = \{X \in \mathcal{M} \mid \exists \iota \colon D \rightsquigarrow X \text{ with } o_X \in \iota(U)\}$ for every $D \in \mathrm{PC}$ homeomorphic to a closed disc and every open $U \subseteq D^{\circ}$.

(iv) The set $\mathcal{M}_{-}(D, K) = \{X \in \mathcal{M} \mid \exists \iota \colon D \rightsquigarrow X \text{ with } o_X \notin \iota(K)\}$for every $D \in \mathrm{PC}$ homeomorphic to a closed disc and every closed $K \subseteq D$.

In this definition, we used implicitly that the existence of an immersion does not change under translation isomorphism. The notion of embedding also gives us a different kind of open sets which can be constructed from the other sets.

Theorem 3.10 ([Hoo13, Theorem 5])**.** *If $D \in \mathrm{PC}$ is homeomorphic to a closed disc, then the set*

$$\mathcal{M}_{\hookrightarrow}(D) \coloneqq \{X \in \mathcal{M} \mid D \hookrightarrow X\}$$

is open.

The following theorem shows that the immersive topology has some nice topological properties:

Theorem 3.11 ([Hoo13, Theorem 8])**.** *The immersive topology is Hausdorff and second countable.*

This tells us that it is very useful to consider convergence of sequences in $\mathcal{M}$. Due to the Hausdorff-property, every sequence can have at most one limit and the second countability provides that convergence of sequences can be used to determine if sets are closed or if maps are continuous.

In the following, we establish some criteria for the convergence of sequences in the immersive topology. For this let us first have a look at the convergence of sequences of simply connected translation surfaces. We denote by $\widetilde{\mathcal{M}} \subseteq \mathcal{M}$ the set of all isomorphism classes of simply connected pointed translation surfaces.

Theorem 3.12 ([Hoo18, §7])**.** *Let $X \in \widetilde{\mathcal{M}}$ and $(X_n)_{n \in \mathbb{N}}$ be a sequence in $\widetilde{\mathcal{M}}$. Then X_n converges to X if and only if the two following statements hold:*

(a) If $D \subseteq X$ is homeomorphic to a closed disc, then $D \rightsquigarrow X_n$ for almost all $n \in \mathbb{N}$.

(b) Let $Q \in \widetilde{\mathcal{M}}$. If $Q \hookrightarrow X_n$ holds for infinitely many $n \in \mathbb{N}$, then $Q \rightsquigarrow X$.

Before we can generalize this theorem for all translation surfaces in $\mathcal{M}$, we need to have a look at the space $\mathcal{E}$. We can expand the immersive topology to $\mathcal{E}$ and also establish some criteria for the convergence of sequences in this space.

Definition 3.13 ([Hoo13, §3.3]). The *immersive topology* on $\mathcal{E}$ is the coarsest topology such that the projection $\pi\colon \mathcal{E} \to \mathcal{M}$ is continuous and such that the set

$$\mathcal{E}_+(D, U) = \{(X, x) \in \mathcal{E} \mid \exists \iota\colon D \rightsquigarrow X \text{ such that } x \in \iota(U)\}$$

is open for every $D \in \mathrm{PC}$ homeomorphic to a closed disc and every open subset $U \subseteq D^\circ$.

Remark 3.14. *The topology on $\mathcal{E}$ is natural in the following sense: As mentioned in Remark 2.7, every translation surface X can be seen as a subset of $\mathcal{E}$. If we restrict the topology of $\mathcal{E}$ to the subspace topology on $X \subseteq \mathcal{E}$, we obtain the usual topology on X.*

Theorem 3.15 ([Hoo18, Proposition 43]). *Let X_n be a sequence converging to X in $\widetilde{\mathcal{M}}$. Then the sequence $(X_n, x_n)_{n\in\mathbb{N}}$ converges to (X, x) in $\mathcal{E}$ if and only if there is a compact subset $K \subseteq X$ with $x \in K$ and an $N \in \mathbb{N}$ such that there is an immersion $\iota_n\colon K \rightsquigarrow X_n$ for every $n > N$ and such that $d(x_n, \iota_n(x)) \to 0$ for $n \to \infty$.*

This allows us to make a statement about the convergence of non simply connected translation surfaces by looking at the universal cover of this surface.

Theorem 3.16 ([Hoo13, Theorem 17]). *Let $(X, o_X) \in \mathcal{M}$ and $(X_n, o_{X_n})_{n\in\mathbb{N}}$ be a sequence in $\mathcal{M}$. In addition let $p\colon \widetilde{X} \to X$ and $p_n\colon \widetilde{X}_n \to X_n$ be the universal covering maps. Then (X_n, o_{X_n}) converges to (X, o_X) if and only if the following statements hold:*

(a) $(\widetilde{X}_n, o_{\widetilde{X}_n})$ converges to $(\widetilde{X}, o_{\widetilde{X}})$.

(b) For every $\widetilde{o} \in p^{-1}(o_X)$, there is a sequence $(\widetilde{o}_n \in p_n^{-1}(o_{X_n}))$ converging to $\widetilde{o}$ in $\mathcal{E}$.

(c) For every subsequence (X_{n_k}) of (X_n) and every sequence of points $(\widetilde{o}_{n_k} \in p_{n_k}^{-1}(o_{X_{n_k}}))$ which converges to some $\widetilde{o} \in \mathcal{E}$ we have $\widetilde{o} \in p^{-1}(o_X)$.

4 Generalizations of the Chamanara surface

4.1 PSI-polygons

As we have seen in Example 2.14, there is a canonical way to construct a translation surface out of a point-symmetric polygon. This construction is very similar to that of the Chamanara surface. The main problem is that the corresponding point-symmetric polygon needed to create the Chamanara surface would have infinitely many edges. Therefore, it is not a polygon in the usual sense. This problem can be fixed by allowing polygons to have infinitely many edges. In this section, we want to give this approach a formal framework. Furthermore, in Section 4.2 we establish some properties of the space of translation surfaces constructed from such polygons.

We start by giving a formal definition.

Definition 4.1. A *line segment* is a set $x \subseteq \mathbb{C}$ of the form $x = \{u + tv \mid t \in [0,1]\}$ for some $u \in \mathbb{C}$, $v \in \mathbb{C}^\times$. We denote the *start* and *end* of x by $s(x) := u$ and $t(x) := u + v$. In addition, the *midpoint* $u + \frac{1}{2}v$ will be denoted by $m(x)$ and the *length* $|v|$ will be denoted by $l(x)$.

Definition 4.2. A *point-symmetric infinite polygon* (or short *PSI-polygon*) is a pair $(P, (x_n)_{n\in\mathbb{N}})$, where $P \subseteq \mathbb{C}$ is a point-symmetric subset homeomorphic to a closed disc and $(x_n)_{n\in\mathbb{N}} \subseteq P$ is a sequence of line segments $(x_n)_{n\in\mathbb{N}} \subseteq P$ such that the following conditions hold:

(a) $\partial P = \overline{\bigcup_{n\in\mathbb{N}} x_n}$,

(b) $\partial P \setminus \bigcup_{n\in\mathbb{N}} x_n$ is finite,

(c) $|x_i \cap x_j| \leq 1$ if $i \neq j$,

(d) for each $i \in \mathbb{N}$ there is exactly one $j \in \mathbb{N}$ with $x_i \cap x_j = \{t(x_i)\} = \{s(x_j)\}$ and

(e) for each $i \in \mathbb{N}$ there is exactly one $j \in \mathbb{N}$ with $x_i \cap x_j = \{s(x_i)\} = \{t(x_j)\}$.

Furthermore, if Φ_P denotes the reflection on the point of symmetry, then for each x_n there should be an $m \in \mathbb{N}$ with $\Phi_P(x_n) = x_m$ and $\Phi_P(s(x_n)) = s(x_m)$.

We will denote a PSI-polygon often just by P if $(x_n)_{n\in\mathbb{N}}$ is clear from the context. As for a finite polygon, the line segments x_n will be called the *edges* of P, and the startpoints of these segments will be called the vertices of P.

Convention 4.3. For a PSI-polygon $(P, (x_n)_{n\in\mathbb{N}})$ the point of symmetry will always be called o_P, Φ_P will always be the reflection on o_P. For each edge x_n of P the mirrored edge $\Phi_P(x_n)$ will also be called x'_n.
In most cases, we will only include half of the edges in the sequence $(x_n)_{n\in\mathbb{N}}$. This will be done in such a way that the sequence $x_1, x'_1, x_2, x'_2, \dots$ contains all the edges of P. So the sequence $(x_n)_{n\in\mathbb{N}}$ still contains all the necessary information. In addition, we mostly define $(x_n)_{n\in\mathbb{N}}$ in such a way that the closure of the union of these edges is connected.

The prime example of a PSI-polygon is of course the one depicted in Figure 2 which we used to construct the Chamanara surface. Our next goal will be to generalize the construction from Figure 2 to construct a translation surface out of an arbitrary PSI-polygon. But before we start, let us first take a look at one major difference to finite polygons, namely that not every point of ∂P must be contained in one edge. We will now show that the missing points correspond to the accumulation points of edges of P. So for example in Figure 2 every point of ∂P is contained in one of the edges except the points $(0, 0)$ and $(1, 1)$.

Proposition 4.4. *Let $(P, (x_n)_{n\in\mathbb{N}})$ be a PSI-polygon. Then the following statements hold:*

(a) *Every point in $\partial P \setminus \bigcup_{n\in\mathbb{N}} x_n$ is an accumulation point of $(s(x_n))_{n\in\mathbb{N}}$.*

(b) *Every accumulation point of $(s(x_n))_{n\in\mathbb{N}}$ is a vertex or lies in $\partial P \setminus \bigcup_{n\in\mathbb{N}} x_n$.*

Proof. (a) Let $y \in \partial P \setminus \bigcup_{n\in\mathbb{N}} x_n$. Because $\partial P = \overline{\bigcup_{n\in\mathbb{N}} x_n}$ there exists a sequence $(y_n)_{n\in\mathbb{N}}$ in $\bigcup_{n\in\mathbb{N}} x_n$ which converges to y. If there was an $N \in \mathbb{N}$ and an $i \in \mathbb{N}$ such that $y_n \in x_i$ for all $n \geq N$ then it would also hold that $y \in x_i$ because x_i is closed.
So there exists a subsequence $(y_{i_n})_{n\in\mathbb{N}}$ converging to y such that each y_{i_n} lies in a different x_j. Let $(j_n)_{n\in\mathbb{N}}$ be the sequence in $\mathbb{N}$ such that $y_{i_n} \in x_{j_n}$. Because P is homeomorphic to a disc, ∂P is homeomorphic to a circle. Then the sequence $(s(x_{j_n}))_{n\in\mathbb{N}}$ lies between the sequences $(y_{i_n})_{n\in\mathbb{N}}$ and $(y_{i_{n-1}})_{n\in\mathbb{N}}$ on this circle. Therefore it converges to y.

(b) So let $y \in x_m \setminus \{s(x_m), t(x_m)\}$ for an $m \in \mathbb{N}$. Again ∂P is homeomorphic to a circle and therefore $(\partial P \setminus x_m) \cup \{s(x_m), t(x_m)\}$ is closed. For each sequence $(y_n)_{n\in\mathbb{N}}$ in ∂P converging to y there has to be an $N \in \mathbb{N}$ such that $y_n \in x_m \setminus \{s(x_m), t(x_m)\}$ for $n \geq N$. Therefore only finitely many elements of such a sequence can be vertices of P, so y is not an accumulation point of these vertices.

□

Proposition 4.4 suggests that a point in ∂P can be a vertex as well as an accumulation point of vertices. In fact, an example of this behavior can be easily constructed by replacing all the vertical edges of Figure 2 with only one vertical edge of length 1.

We will now start the construction of a translation surface out of a PSI-polygon $(P, (x_n)_{n\in\mathbb{N}})$. In Example 2.14 this was already done for the finite case. Of course, in the infinite case, the vertices of P again correspond to the singularities of the translation surfaces. But in addition, the accumulation points which were examined in Proposition 4.4 also correspond to singularities. This holds for a similar reason as for the Chamanara surface.

We begin by defining the set

$$P^* := P \setminus \Sigma_P$$

where Σ_P is defined as

$$\Sigma_P := \overline{\{s(x_n) \mid n \in \mathbb{N}\}}. \tag{1}$$

We can then glue $x_n \cap P^*$ to $x_n' \cap P^*$ via a translation. The resulting topological space will be called $\widehat{P}$.

Proposition 4.5. *$\widehat{P}$ is a translation surface.*

Proof. For each point $\widehat{x} \in \widehat{P}$ corresponding to a point $x \in P^\circ$, there is an open neighborhood $\widehat{U}$ corresponding to an open set $U \subseteq P^\circ$. So the inclusion of U in $\mathbb{C}$ gives a natural chart.
As seen in Proposition 4.4, every point in $x \in \partial P \cap P^*$ belongs to $x_n \setminus \{s(x_n), t(x_n)\}$ for some $n \in \mathbb{N}$. So let $\widehat{x} \in \widehat{P}$ be the corresponding point in the surface. Then $\widehat{x}$ does not only correspond to the point $x \in x_n$, but also to the point $x' = \Phi(x) \in x_n'$. Then there are two open neighborhoods U and U' of x and x' in P homeomorphic to two open subsets of the upper halfplane. Choosing suitable subsets we can assume that $U, U' \subseteq P^*$ and $U' = \Phi(U)$. On the set $\widehat{U}$ corresponding to the set $U \cup U'$ there is a natural chart given by the inclusion of U in $\mathbb{C}$ and the translation which sends x_i' to x_i. $\square$

Convention 4.6. For a PSI-polygon $(P, (x_n)_{n\in\mathbb{N}})$ the translation surface constructed from P will always be called $\widehat{P}$. As in the above proof, for a point $x \in P$, the corresponding point in $\widehat{P}$ will be called $\widehat{x} \in \widehat{P}$ and vice versa. The same will be done for subsets of P or $\widehat{P}$.

Convention 4.7. We will always use $\widehat{o}_P$ as the basepoint of $\widehat{P}$.

As we have seen the Chamanara surface has only one singularity because each vertex of the PSI-polygon gets identified to one point in the corresponding translation surface. In fact, this is always true for all $\widehat{P}$ constructed out of a PSI-polygon P:

Proposition 4.8. *Let P be a PSI-polygon. Then $\widehat{P}$ has exactly one singularity.*

Proof. This follows essentially for the same reasons as it holds for the Chamanara surface. Let us first assume that no vertex (i.e. endpoint of an edge) of P is an accumulation point of vertices of P. Denote by $\{y_1, y_2, \ldots, y_n\}$ the accumulation points of the vertices of P such that on the part of ∂P between y_i and y_{i+1} there lies no other such accumulation point. In addition, let x_i denote the edges between y_1 and y_2 such that x_i is adjacent to x_{i-1} and x_{i+1} and such that $s(x_i) = t(x_{i+1})$. But through the identification in $\overline{\widehat{P}}$ it holds that

$$s(\widehat{x}_i) = t(\widehat{x'}_i) = s(\widehat{x'}_{i-1}) = t(\widehat{x}_{i-1}).$$

So again as for the Chamanara surface, every second vertex gets identified. This also holds for each part of ∂P between each y_i and y_{i+1} as well as for the part between y_n and y_1. Consider the sequence $(s(x_{2n}))_{n\in\mathbb{N}}$, which is identified to one singularity in $\widehat{P}$. It converges to y_2 for $n \to \infty$ and to y_1 for $n \to -\infty$ (or vice versa). So the two points corresponding to the sequences $(s(x_{2n}))_{n\in\mathbb{N}}$ and $(s(x_{2n+1}))_{n\in\mathbb{N}}$ get identified into one singularity by this convergence. For a similar reason the two points (corresponding to the two similarly defined sequences) between y_2 and y_3 also get identified with this singularity. Inductively, all these points get identified into one singularity.

Let us now return to the case, where a vertex is an accumulation point. W.l.o.g. let y_2 be a vertex, such that $(s(x_n))_{n\in\mathbb{N}}$ converges to y_2 for $n \to \infty$. Due to this convergence y_2 and $\Phi_P(y_2)$ get identified. Denote by x the edge such that $s(x) = y_2$. But then

$$\widehat{y_2} = \widehat{\Phi_P(y_2)} = \widehat{\Phi_P(s(x))} = s(\widehat{x'}) = t(\widehat{x})$$

holds in $\overline{\widehat{P}}$, so every vertex between y_2 and y_3 gets identified with y_2. Then the argument can be continued. □

As we have seen in Section 2 and 3, it is often useful to look at the universal cover of a translation surface. For a PSI-polygon $(P, (x_n)_{n\in\mathbb{N}})$, the universal cover of $\widehat{P}$ can always be constructed in the following way: Let S be the set consisting of half of the edges of P as noted in Convention 4.3 and F be the free group generated by S. We now take a look at the space $P^* \times F$. For each $g \in F$, the set $P^\circ \times \{g\}$ admits a natural structure of a translation surface, coming from P°. We then get a translation structure on all of $P^* \times F$ by identifying each $x_i \times \{g\}$ with $x'_i \times \{x_i \circ g\}$. The resulting translation surface will be denoted by $\widetilde{P}$. One example can be seen in Figure 4. We will always use (o_P, e) as the basepoint of $\widetilde{P}$, where e is the neutral element of F.

Proposition 4.9. *$\widetilde{P}$ is the universal cover of $\widehat{P}$.*

Proof. The map $p\colon \widetilde{P} \to \widehat{P}$, $(y, g) \mapsto \widehat{y}$ is continuous and surjective. A point $\widehat{y} \in \widehat{x}_i$ has a neighborhood of the form $\widehat{U} \cup \widehat{\Phi(U)}$ with $\widehat{U} \cap \widehat{\Phi(U)} \subseteq \widehat{x}_i$. Then $p^{-1}(\widehat{U} \cup \widehat{\Phi(U)})$ is the disjoint union of the sets $U \times \{g\} \cup \Phi(U) \times \{x_i \circ g\}$ over all $g \in F$. For each point $y \in P^\circ$ there is a neighborhood $U \subseteq P^\circ$ such that $p^{-1}(\widehat{U})$ is the disjoint union of the sets $U \times \{g\}$ for all $g \in F$. In addition, p respects the translation structure. Therefore

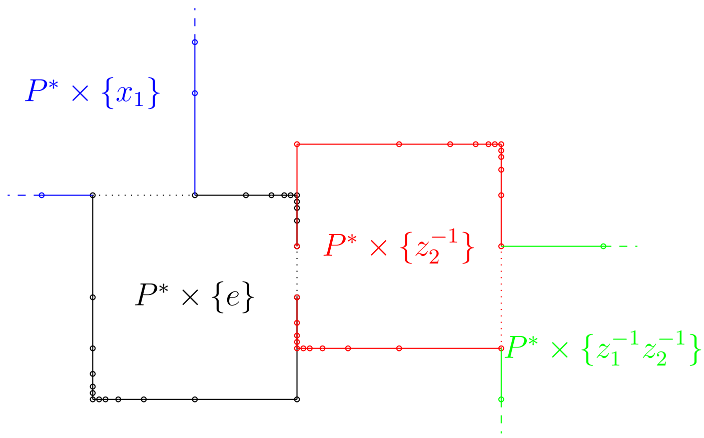

Figure 4: Part of the universal cover of $\mathrm{Ch}_{\frac{1}{2}}$. The PSI-polygon P is the one depicted in Figure 2.

$\widetilde{P}$ is a covering space of $\widehat{P}$.
Let us now look at the set $S \subseteq P$ consisting of the segments connecting o_P to the mid-points of the edges x_n. Then $S \times F \subseteq \widetilde{P}$ is homotopy equivalent to $\widetilde{P}$ and homeomorphic to a tree. Therefore $\widetilde{P}$ is simply connected. $\square$

The natural projection $\pi\colon P^* \to \widehat{P}$ is automatically continuous. Let X be a topological space. Then for every continuous map $f\colon X \to P^*$ the map $\widehat{f} \coloneqq \pi \circ f$ is also continuous. The converse is not always true. Nevertheless for a path $\widehat{\gamma}\colon [0,1] \to \widehat{P}$ a corresponding map $\gamma\colon [0,1] \to P^*$ can be constructed. This is just the intuitive approach of drawing the path $\widehat{\gamma}$ in the polygon P where it splits into pieces whenever it crosses an edge. This is formalized in the following remark:

Remark 4.10. *Let $(P, (x_n)_{n\in\mathbb{N}})$ be a PSI-polygon and let $\widehat{\gamma}\colon [0,1] \to \widehat{P}$ be a path. Then there is a map $\gamma\colon [0,1] \to P$ such that $\widehat{\gamma(a)} = \widehat{\gamma}(a)$ for all $a \in [0,1]$. Such a map can be constructed in the following way:*
We choose one of the up to two points of P corresponding to $\widehat{\gamma}(0)$ as $\gamma(0)$. Let $\widetilde{\gamma}\colon [0,1] \to \widetilde{P}$ be the lift of $\widehat{\gamma}$ which starts at $(\gamma(0), e)$. Let $a_1 \in [0,1]$ denote the maximum for which $\widetilde{\gamma}([0,a_1]) \subseteq P^* \times \{e\}$. For $a \in [0,a_1]$ we then set $\gamma(a) = \pi_1(\widetilde{\gamma}(a))$ where π_1 is the projection to the first coordinate. This can be repeated for each segment of $\widetilde{\gamma}$.

Then the map $\gamma\colon [0,1] \to P$ is piecewise continuous. The discontinuities of γ are those points where it meets or rather crosses one of the edges of P.

4.2 Convergence of PSI-polygons

We now want to investigate some conditions for the convergence of such translation surfaces constructed out of PSI-polygons. An intuitive thought might be that such translation surfaces converge if and only if the vertices of the corresponding PSI-polygons converge. Unfortunately, this is not the case. One problem that might occur is that for two different PSI-polygons P_1 and P_2 the corresponding translation surfaces $\widehat{P}_1$ and $\widehat{P}_2$ can be isomorphic. This problem can be solved but requires some work which will mainly be done in Lemma 5.15.

Another problem occurs because the immersive topology is rather weak and there are converging sequences of such translation surfaces, where the vertices of the corresponding PSI-polygons do not converge, even if each of these translation surfaces can be constructed from a unique PSI-polygon. We will return to this problem in Section 6.

On the upside, the converse is always true. If the vertices of a sequence of PSI-polygons converge, then the corresponding translation surfaces also converge. This is formalized in Theorem 4.18. The rest of this section is devoted to proving Theorem 4.18. Please note that by convergence of vertices we mean that they converge uniformly. This is formalized in the following definition.

Definition 4.11. Let $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ be a sequence of PSI-polygons and $(P, (x_m)_{m\in\mathbb{N}})$ be a PSI-polygon in such a way that the connection between the edges coincides, i.e. for all $m_1, m_2, n \in \mathbb{N}$ it holds that

$$s(x_{m_1}) = t(x_{m_2}) \Leftrightarrow s(x_{m_1}^{(n)}) = t(x_{m_2}^{(n)}).$$

In addition, the points of symmetry all coincide, i.e. $o_P = o_{P^{(n)}}\ \forall n \in \mathbb{N}$.

Then we say that $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ *converges* to $(P, (x_m)_{m\in\mathbb{N}})$ if for each $\varepsilon \in \mathbb{R}_+$ there exists an $N_\varepsilon \in \mathbb{N}$ such that

$$s(x_m^{(n)}) \in B_\varepsilon(s(x_m)) \quad \forall m \in \mathbb{N}, \forall n \geq N_\varepsilon$$

holds.

We will always denote sequences of PSI-polygons by $P^{(n)}$ instead of P_n in an effort not to overload the bottom index. Similar to before we will always denote points of $P^{(n)}$ by $x^{(n)}$, the same for edges, subsets, paths, etc.

Convention 4.12. For a sequence of PSI-polygons $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ converging to a PSI-polygon $(P, (x_m)_{m\in\mathbb{N}})$ and $\varepsilon \in \mathbb{R}_+$ we will denote by N_ε a number as defined in Definition 4.11, i.e. a number such that

$$s(x_m^{(n)}) \in B_\varepsilon(s(x_m)) \quad \forall m \in \mathbb{N}, \forall n \geq N_\varepsilon$$

holds.

Remark 4.13. *Let $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ be a sequence of PSI-polygons converging to $(P, (x_m)_{m\in\mathbb{N}})$. In addition let F be the free group used for the construction of the corresponding universal cover of $\widehat{P}$. Formally identifying x_i with $x_i^{(n)}$ we can use F for the construction of $\widetilde{P}$ as well as for the construction of each $\widetilde{P}^{(n)}$.*

As we have seen in Proposition 4.4 a PSI-polygon P has two kinds of points which correspond to the singularity of $\widehat{P}$. These points are either vertices of P or accumulation points of vertices. As it turns out convergence of PSI-polygons also implies convergence of the latter kind of points.

Proposition 4.14. *Let $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ be a sequence of PSI-polygons converging to $(P, (x_m)_{m\in\mathbb{N}})$. For each point $y \in \partial P$ which is an accumulation point of the vertices of P and for each $\varepsilon \in \mathbb{R}_+$, there exists for each $n \geq N_\varepsilon$ a point $y^{(n)} \in \partial P^{(n)}$ which is an accumulation point of the vertices of $P^{(n)}$ and such that*

$$y^{(n)} \in \overline{B_\varepsilon(y)}$$

holds. In addition the number of accumulation points of vertices of P and $P^{(n)}$ are the same.

Proof. Let $y \in \partial P$ be an accumulation point of the vertices of P. Then there is a sequence of vertices $(s(x_{n_k}))_{k\in\mathbb{N}}$ converging to y. W.l.o.g. we can assume that each x_{n_k} and $x_{n_{k+1}}$ are adjacent. Then $x_{n_k}^{(n)}$ and $x_{n_{k+1}}^{(n)}$ are also adjacent. Let us assume that the sequence $(s(x_{n_k}^{(n)}))_{k\in\mathbb{N}}$ does not converge. But then it is also not a Cauchy sequence and therefore the length of $\bigcup_{k=1}^{\infty} x_{n_k}$ would be infinite. This is a contradiction to $\partial P^{(n)}$ being homeomorphic to a circle.

So let $y^{(n)}$ denote the limit of $(s(x_{n_k}^{(n)}))_{k\in\mathbb{N}}$. Then $d(y^{(n)}, y) \leq \varepsilon$ or else we can find a large enough $k \in \mathbb{N}$ such that $d(s(x_{n_k}^{(n)}), s(x_{n_k})) \geq \varepsilon$.

This also shows, that the number of accumulation points of vertices of $P^{(n)}$ is greater or equal than the number of accumulation points of vertices of P. To show the opposite let $y^{(n)} \in \partial P^{(n)}$ be an accumulation point of the vertices of $P^{(n)}$. Then there is a sequence of vertices $(s(x_{m_k}^{(n)}))_{k\in\mathbb{N}}$ converging to $y^{(n)}$. With the same arguments as in the beginning of the proof, it follows, that the corresponding sequence of vertices in P also converges. □

One problem with the convergence of PSI-polygons is that even for very small ε there are always vertices of P closer together than ε. So their position in $P^{(n)}$ for $n > N_\varepsilon$ relative to one another is rather arbitrary. But at least for all edges x_i of a certain minimum length there is a certain set A_i around $m(x_i)$ (illustrated in Figure 5) which cannot intersect any other edge of $P^{(n)}$.

Lemma 4.15. *Let $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ be a sequence of PSI-polygons converging to $(P, (x_m)_{m\in\mathbb{N}})$. Then for each edge x_i of P and each $\delta \in \mathbb{R}_+$ with $\delta < \frac{1}{2}l(x_i)$, there exists an $\varepsilon \in \mathbb{R}_+$ ($\varepsilon < \delta$) such that for $n > N_\varepsilon$ each other edge of $P^{(n)}$ has distance greater than ε to the set*

$$A_i := \{v \in x_i \mid d(v, s(x_i)) \geq \delta \text{ and } d(v, t(x_i)) \geq \delta\}.$$

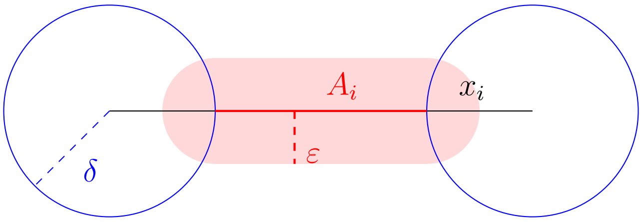

Figure 5: An illustration of the set A_i from Lemma 4.15. No edge of $P^{(n)}$ except $x_i^{(n)}$ can be in the area marked in red.

Proof. Similar to the proof of Proposition 4.4, the set A_i cannot contain an accumulation point of vertices of P and therefore also not an accumulation point of $\partial P \setminus A_i$. But then there is a $\xi \in \mathbb{R}_+$ such that each point in $\partial P \setminus A_i$ has distance greater than ξ to A_i. Now let $\varepsilon = \frac{\xi}{2}$ and $n > N_\varepsilon$. Then each edge $x_j^{(n)}$ of $P^{(n)}$ lies completely in $B_\varepsilon(x_j)$. But then each edge of $P^{(n)}$ (except $x_i^{(n)}$) still has distance greater than ε to each point in A_i. $\square$

Lemma 4.15 allows for some variation. For example this can also be done the other way around, i.e. for each edge $x_i^{(n)}$ of $P^{(n)}$ each edge of P has distance more than ε from the corresponding $A_i^{(n)}$ for large enough n.
Lemma 4.15 can also be adjusted to work for every edge longer than 2δ, because there are only finitely many of these edges.
Sometimes, we cite Lemma 4.15, when one of these variations is needed.

We will now show two statements which will be used in the proof of Theorem 4.18. For this, let $\widehat{\Sigma}_P$ denote the set of singularities of $\widehat{P}$ (consisting of only one element, see Proposition 4.8) and $\widetilde{\Sigma}_P$ the set of singularities of $\widetilde{P}$. As we have seen before, these sets correlate to the set Σ_P of vertices of P and their accumulation points defined in (1).

Lemma 4.16. *Let $(P, (x_n)_{n\in\mathbb{N}})$ be a PSI-polygon and $K \subseteq \widetilde{P}$ compact. Then there is an $\varepsilon \in \mathbb{R}_+$ such that*

$$d(x, y) > \varepsilon \quad \forall x \in K, y \in \widetilde{\Sigma}_P.$$

Proof. Let us first assume that for each $\varepsilon \in \mathbb{R}_+$ there is an $y \in \widetilde{\Sigma}_P$ and $x \in K$ such that $d(x, y) \leq \varepsilon$.
Then there is also a sequence $(x_n)_{n\in\mathbb{N}}$ in K such that

$$\forall \varepsilon \in \mathbb{R}_+, \exists N \in \mathbb{N}, \forall n \geq N \colon d(x_n, \widetilde{\Sigma}_P) \leq \varepsilon.$$

Because K is compact, $(x_n)_{n\in\mathbb{N}}$ has an accumulation point $x \in K$. Then for each $\varepsilon \in \mathbb{R}_+$ there is $y \in \widetilde{\Sigma}_P$ such that $d(x, y) \leq \varepsilon$. Therefore x is either a point in $\widetilde{\Sigma}_P$ or is an accumulation point of $\widetilde{\Sigma}_P$. But because of the definition of Σ_P (see (1)) these accumulation points also lie in $\widetilde{\Sigma}_P$. So it holds that $x \in \widetilde{\Sigma}_P$, which is a contradiction to $K \subseteq \widetilde{P}$. $\square$

One of the problems that we will encounter is that by shifting one vertex $s(x_i)$ of P [2] the copies $P^* \times \{x_i^k\}$ get shifted more and more if $k \to \infty$. This will cause some problems in the following proof. But at least the shift of $P^* \times \{g\}$ is bounded in correlation with the number of occurrences of x_i in g:

Lemma 4.17. *Let $\varepsilon \in \mathbb{R}_+$. If $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ converges to $(P, (x_m)_{m\in\mathbb{N}})$, then for all $g \in F$*

$$d(\mathrm{dev}_{\widetilde{P}}(o_P, g), \mathrm{dev}_{\widetilde{P}^{(n)}}(o_{P^{(n)}}, g)) < 4\varepsilon|g| \quad \forall n > N_\varepsilon$$

holds, where $|g|$ denotes the length of the word $g \in F$.

Proof. Let $\widetilde{\gamma}$ be a path in $\widetilde{P}$ from (o_P, e) to (o_P, x_i) and $\widetilde{\gamma}^{(n)}$ be a path in $\widetilde{P}^{(n)}$ from $(o_{P^{(n)}}, e)$ to $(o_{P^{(n)}}, x_i)$. We denote by $\widetilde{z}$ $(\widetilde{z}^{(n)})$ the intersection of the image of $\widetilde{\gamma}$ $(\widetilde{\gamma}^{(n)})$ with $x_i \times \{e\}$ $(x_i^{(n)} \times \{e\})$. This situation is depicted in Figure 6. W.l.o.g. we can assume that $d(\widetilde{z}, (s(x_i), e)) = d(\widetilde{z}^{(n)}, (s(x_i^{(n)}), e))$. The point z (or $z^{(n)}$) is an equivalence class consisting of two points (z_1, e) and (z_2, x_i) $((z_1^{(n)}, e)$ and $(z_2^{(n)}, x_i))$. Integrating along the path $\widetilde{\gamma}$ yields a vector $v_{\widetilde{\gamma}} \in \mathbb{C}$ which is equal to $\mathrm{dev}_{\widetilde{P}}(\widetilde{\gamma}(1)) - \mathrm{dev}_{\widetilde{P}}(\widetilde{\gamma}(0))$. Thus $v_{\widetilde{\gamma}} = (z_1 - o_P) + (o_P - z_2)$ and the same holds for $\widetilde{\gamma}^{(n)}$.

In addition, it holds that $d(z_1, z_1^{(n)}) < 2\varepsilon$ and $d(z_2, z_2^{(n)}) < 2\varepsilon$ and therefore we have

$$\begin{aligned} d(\mathrm{dev}_{\widetilde{P}}(o_P, x_i), \mathrm{dev}_{\widetilde{P}^{(n)}}(o_{P^{(n)}}, x_i)) \;&=\; d(v_{\widetilde{\gamma}}, v_{\widetilde{\gamma}^{(n)}}) \\ = d(z_1 - o_P + o_P - z_2, z_1^{(n)} - o_{P^{(n)}} + o_{P^{(n)}} - z_2^{(n)}) \;&=\; d(z_1 - z_2, z_1^{(n)} - z_2^{(n)}) \\ \leq d(z_1, z_1^{(n)}) + d(z_2, z_2^{(n)}) \;&<\; 4\varepsilon. \end{aligned}$$

Inductively the claim follows. □

Another way to phrase Lemma 4.17 is, that if the two endpoints of one edge of P are shifted less than ε, then the corresponding translation surface gains or loses some area of diameter smaller than 2ε. This is depicted in Figure 7.

Formally, $\mathrm{dev}_{\widetilde{P}}$ cannot be extended to the metric completion $\overline{\widetilde{P}}$ because $\overline{\widetilde{P}}$ fails to be a translation surface. But we can still make sense of the notion $\mathrm{dev}_{\widetilde{P}}(s(x_i), g)$ in the following sense: For each $g \in F$, the set $\mathrm{dev}_{\widetilde{P}}(P^* \times \{g\})$ is a translate of P^*. Therefore the point $\mathrm{dev}_{\widetilde{P}}(s(x_i), g)$ can be defined to be the starting point of the segment $\mathrm{dev}_{\widetilde{P}}(x_i \times \{g\})$. For each singularity of $\widetilde{P}$ corresponding to an accumulation point $y \in \partial P$ there is a subsequence $(s(x_{i_n}))_{n\in\mathbb{N}}$ converging to y. We can proceed in a similar manner and define the point $\mathrm{dev}_{\widetilde{P}}(y, g)$ to be the limit of $(\mathrm{dev}_{\widetilde{P}}(s(x_{i_n}), g))_{n\in\mathbb{N}}$.

[2] i.e. by taking another PSI-polygon $P^{(1)}$ which has the same vertices as P, except the vertex $s(x_i)$ which is at another position.

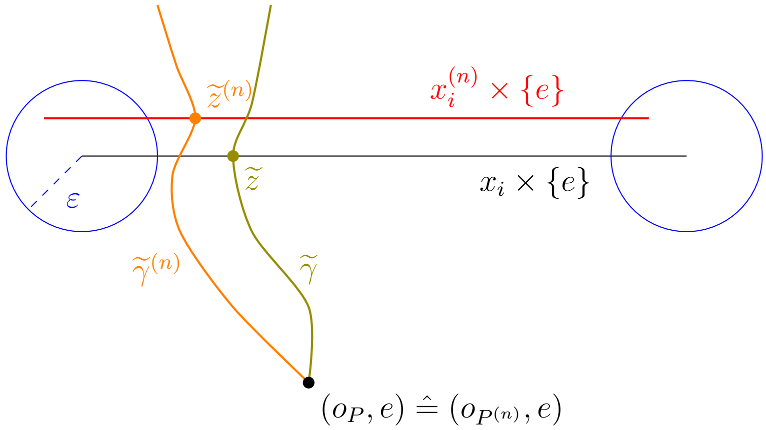

Figure 6: The situation in Lemma 4.17. Here the corresponding vector $z_1^{(n)} - o_{P^{(n)}}$ is slightly longer than $z_1 - o_P$, but the difference is smaller than ε.

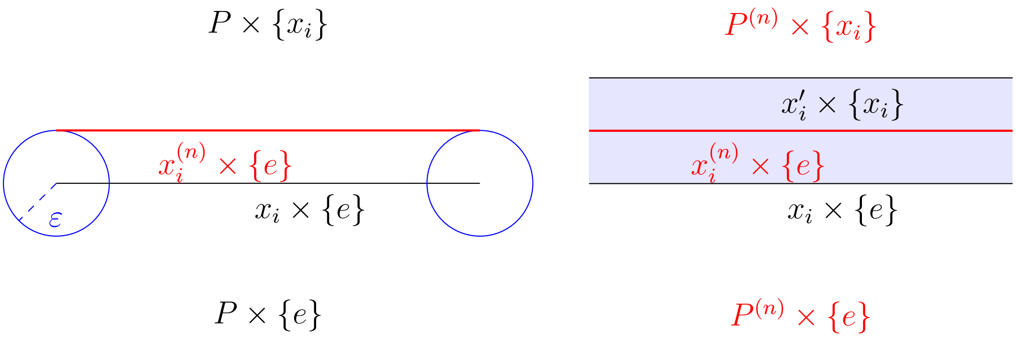

(a) The situation in $\widetilde{P}$.

(b) The corresponding situation in $\widetilde{P}^{(n)}$ if $x_i^{(n)}$ is shifted ε up compared to x_i.

Figure 7: Shifting x_i an ε farther away from o_P creates a new area with diameter 2ε marked in blue.

We now want to examine the reasons why a set cannot be immersed in a translation surface $\widetilde{P}$. As mentioned in Section 3, this is mostly due to the set "hitting" a singularity. We now want to make this notion more precise and investigate the situation for translation surfaces of the form $\widetilde{P}$. Let P be a PSI-polygon, $X \in \mathcal{M}$ and let $Q \subseteq X$ be a simply connected subset containing the basepoint.

We now take a look at the subset $V \subseteq Q$ such that $\mathrm{dev}_Q(V) \subseteq \mathrm{dev}_{\widetilde{P}}(P^* \times \{e\})$. Then let U_e be the connected component of V containing the basepoint. Then, as noted after Example 3.4, the map $\mathrm{dev}_{P\times\{e\}}^{-1} \circ \mathrm{dev}_{U_e}$ is a well defined immersion and therefore $U_e \rightsquigarrow \widetilde{P}$.

Now let $\overline{U_e} \subseteq Q$ be the set defined similarly as U_e, but with $\mathrm{dev}_Q(\overline{U_e})$ being a subset of $\mathrm{dev}_{\widetilde{P}}(P^* \times \{e\})$ together with the image of the corresponding singularities as defined above. If $U_e \neq \overline{U_e}$ then $Q \not\rightsquigarrow \widetilde{P}$: We can find a sequence in U_e converging to some point $x \in \overline{U_e} \setminus U_e$ such that $\mathrm{dev}_Q(x) = \mathrm{dev}_{\widetilde{P}}(s(x_i), e)$ for some $i \in \mathbb{N}$ (or $\mathrm{dev}_{\widetilde{P}}(y, e)$ for some accumulation point y of the vertices of P). This sequence then diverges in $\widetilde{P}$ because it converges to a point in $\overline{\widetilde{P}} \setminus \widetilde{P}$. Therefore the immersion from U_e to $\widetilde{P}$ cannot be extended to an immersion from $\overline{U_e}$ and especially not to an immersion from Q.

If instead $U_e = \overline{U_e}$ and there is no x_i with $\mathrm{dev}_Q(U_e) \cap \mathrm{dev}_{\widetilde{P}}(x_i \times \{e\}) \neq \emptyset$ then $U_e = Q$ and therefore Q can be immersed in $\widetilde{P}$. If instead there is x_i with $\mathrm{dev}_Q(U_e) \cap \mathrm{dev}_{\widetilde{P}}(x_i \times \{e\}) \neq \emptyset$ we can repeat this procedure for $P^* \times \{x_i\}$. We can do this by taking a point $b \in U_e$ such that $\mathrm{dev}_Q(b) \in \mathrm{dev}_{\widetilde{P}}(x_i \times \{e\})$ and using b as the new basepoint for this procedure on $P^* \times \{x_i\}$. If we can repeat this for all $h \in F$ then $Q \rightsquigarrow \widetilde{P}$ follows because then each point in Q lies in one of these U_h.

In summary, we can conclude that if $Q \not\rightsquigarrow \widetilde{P}$ there is an $h \in F$ such that there is $x \in \overline{U_h} \setminus U_h$. We can choose this h to be minimal, i.e $U_{h'} = \overline{U_{h'}}$ holds for all suffixes h' of h. Also there is x_i (or y) with $\mathrm{dev}_{\widetilde{P}}(s(x_i), h) = \mathrm{dev}_Q(x)$. Then the pair

$$(x, (s(x_i), h)) \in Q \times (P \times F)$$

or

$$(x, (y_i, h))$$

will be called a *critical point*.

For each such singularity of $\widetilde{P}$ corresponding to a point $(s(x_i), h)$ this singularity also corresponds to $(t(x_i'), x_i \circ h)$ and to $(s(x_j'), x_j \circ h)$ where x_j is the edge of P with $t(x_j) = s(x_i)$. This can be repeated for these points as well. The upshot is that if we look at the subset of F consisting of all h' such that $(s(x_i), h)$ also belongs to $P \times \{h'\}$ then this set has a shortest element g. In addition if $(x, (s(x_i), h))$ is a critical point then $h \in F$ is such a shortest element. This is true because $U_{h'} = \overline{U_{h'}}$ for all suffixes h' of h and each shorter element in this set has to be a suffix of h.

We are now ready to prove the aforementioned theorem.

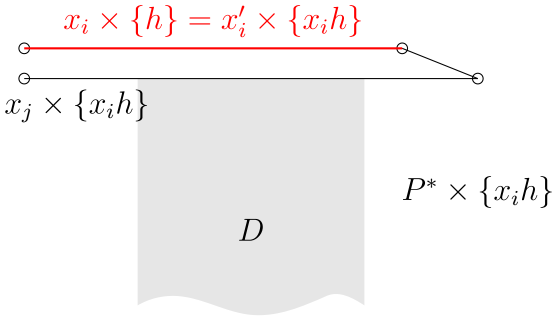

Figure 8: There might be an edge x_i which is so close to x_j that they switch position in $P^{(n)}$. The distance of D and $x_i \times \{h\}$ is relatively large in $P^* \times \{x_i h\}$ because it is taken by going around the drawn edges. But Lemma 4.15 shows that this behavior cannot happen for large enough n.

Theorem 4.18. *If $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ converges to $(P, (x_m)_{m\in\mathbb{N}})$, then $(\widehat{P}^{(n)})_{n\in\mathbb{N}}$ converges to $\widehat{P}$ in the immersive topology.*

Proof. We will use Theorem 3.16 to show the convergence in the immersive topology. To do this we first use Theorem 3.12 to show that $\widetilde{P}^{(n)}$ converges to $\widetilde{P}$. So let $D \subseteq \widetilde{P}$ be homeomorphic to a closed disc with $o_{\widetilde{P}} \in D$. We now have to show $D \rightsquigarrow \widetilde{P}^{(n)}$ for almost all $n \in \mathbb{N}$. With Lemma 4.16 it follows that there exists $\varepsilon \in \mathbb{R}_+$ such that $d(x, y) > \varepsilon \;\forall x \in D, y \in \widetilde{\Sigma}$. We will only focus on singularities corresponding to a point of the form $(s(x_i), h)$. For singularities of the form (y, h), where y is an accumulation point of the vertices of P the arguments are analogous.

Now let F be the free group as before. We denote by $G \subseteq F$ the set of all $g \in F$ with $D \cap (P^* \times \{g\}) \neq \emptyset$. Then there is $\varepsilon' \in \mathbb{R}_+$ such that $d(D, P^* \times \{h\}) > \varepsilon'$ for all $h \in F \setminus G$. This can be shown similar to Lemma 4.16: Let us assume that this statement is not true. Then there has to be a single point $x \in D \cap (P^* \times \{g\})$ such that for any $\delta \in \mathbb{R}_+$ there is a segment $x_i \times \{g\}$ closer than δ to x (as in Lemma 4.16 x can be obtained as the accumulation point of such a sequence $(x_n)_{n\in\mathbb{N}}$). But this is not possible because $\operatorname{dev}_{\widetilde{P}}(\partial P \times \{g\})$ is compact.

Note that this distance is measured in $\widetilde{P}$ and not in P. So it can still be possible that for one $x_i \times \{h\}$ the set $\operatorname{dev}_{\widetilde{P}}(x_i \times \{h\})$ is relatively close to $\operatorname{dev}_{\widetilde{P}}(D \cap P^* \times \{x_i h\})$ where the distance is measured in $\mathbb{C}$. This is possible because $x_i \times \{h\}$ may lie behind an edge $x_j \times \{x_i h\}$ that D goes through, compare Figure 8. But this implies that x_i' and x_j are relatively close together in P. This is exactly the behavior we get under control by applying Lemma 4.15 to each of these edges: D goes only through a finite number of edges $x_j \times \{g\}$, so there is $\delta \in \mathbb{R}_+$ such that D has distance greater than δ to each of the endpoints of each of the $x_j \times \{g\}$. Then we can choose ε'' in such a way that each x_i has distance greater than ε'' to such an A_j. Therefore the distance between $\operatorname{dev}_{\widetilde{P}}(x_i \times \{h\})$ and $\operatorname{dev}_{\widetilde{P}}(D \cap P^* \times \{x_i h\})$ is also greater than ε''.

Now let $\xi := \min(\varepsilon, \varepsilon', \varepsilon'')$, let $g \in F$ be a longest word in G and let $0 < \delta < \frac{1}{8|g|}\xi$. We now want to show that $D \rightsquigarrow \widetilde{P}_n$ for all $n > N_\delta$. Let us assume the opposite, so let $n > N_\delta$ and $D \not\rightsquigarrow P_n$. Then there is a critical point $(x, \left(s(x_i^{(n)}), h\right))$.

Let us assume for a moment that there is exactly one longest element $g \in G$ and each other element in G is a suffix of g. Let $U_h \subseteq D$ be as before for the critical point $(x, \left(s(x_i^{(n)}), h\right))$. Then there is a path γ in $\overline{U_h}$ from $o_{\widetilde{P}}$ to x such that $\gamma([0,1)) \subseteq U_h$. If we look at the embedding $\gamma^{(n)}$ of this path in $\widetilde{P}^{(n)}$ we see that if it crosses one edge $x_j^{(n)} \times \{e\}$ of $P^{(n)} \times \{e\}$ this edge has to be the last letter of g. This holds because every edge of $P^{(n)} \times \{e\}$ can be shifted at maximum δ. But this is smaller than ξ which is the distance of U_h to these edges (except $x_j^{(n)} \times \{e\}$, but this poses no problem due to Lemma 4.15 as explained above). Then, similar to the proof of Lemma 4.17 the first point of the path in $P^{(n)} \times \{x_j^{(n)}\}$ is shifted at maximum 2δ in relation to the first point of the path in $P \times \{x_j\}$. This is again the situation depicted in Figure 7.

We denote by $b \in [0,1)$ the first point such that $\gamma^{(n)}(b) \in P^{(n)} \times \{x_j\}$ and by c the last such point. Then for each $k \in \mathbb{N}$ the vertex $(s(x_k^{(n)}), x_j)$ still has distance more than $\xi - 3\delta$ from each point in $\gamma^{(n)}([b,c])$. This holds because each $x_k \times \{x_j\}$ with $x_k \times \{x_j\} \cap D = \emptyset$[3] has distance more than ξ from each point in $\mathrm{Im}(\gamma)$.

This can be continued inductively to show that each vertex $(s(x_k^{(n)}), h)$ still has distance more than

$$\xi - 3\delta|g| > \frac{1}{2}\xi$$

from the last part of $\gamma^{(n)}$. This is a contradiction to $(x, \left(s(x_i^{(n)}), h\right))$ being a critical point.

If G is not of this form then we can use that G consists of only finitely many elements because D is compact. Therefore we can repeat the above process for each such element in G that is not a suffix of a longer word. This leads to finitely many of the above constructed ξ and we can take the smallest such ξ. Then the rest follows as before.

Now let $Q \in \widetilde{\mathcal{M}}$ with $Q \hookrightarrow \widetilde{P}^{(n)}$ for infinitely many $n \in \mathbb{N}$, so we can see Q as a subsurface of these $\widetilde{P}^{(n)}$. We assume that $Q \not\rightsquigarrow \widetilde{P}$. Then again there is a critical point $(x, (s(x_i), h))$ and a corresponding $U_h \subseteq Q$. Like before we can choose a path γ in $\overline{U_h}$ from $o_{\widetilde{P}}$ to x such that $\gamma([0,1)) \subseteq U_h$. Then, because Q is open, there is $\xi \in \mathbb{R}_+$ such that $B_\xi(\mathrm{Im}(\gamma)) \subseteq Q$. But for similar reasons as before for large enough $n \in \mathbb{N}$ there has to be a singularity of $\widetilde{P}^{(n)}$ closer than ξ to $\gamma^{(n)}(1) \in Q \subseteq \widetilde{P}^{(n)}$[4]. This is a contradiction

[3] Only $x_j' \times \{x_j\}$ and up to one other edge do not fulfill this condition. Again, Lemma 4.15 shows that there cannot be an edge lying closely behind one of those edges.

[4] This singularity does not necessarily correspond to $(s(x_i^{(n)}), h)$ because another edge of $P^{(n)} \times \{h\}$ may lie before $(s(x_i^{(n)}), h)$ (similar to the case displayed in Figure 8). But then Lemma 4.15 shows that n can be chosen in such a way that this edge is sufficiently small, so that one of its endpoints is closer than ξ to $\gamma^{(n)}(1)$.

to $Q \hookrightarrow \widetilde{P}^{(n)}$.

We now use Theorem 3.16 to show that $\widehat{P}^{(n)}$ converges to $\widehat{P}$. Up to now we have shown (a). Let us show (b) and (c). The points in $p^{-1}(\widehat{o}_P)$ are exactly the points of the form (o_P, g). By taking a shortest path γ from (o_P, e) to (o_P, g) we get an $\varepsilon \in \mathbb{R}_+$ such that $d(\mathrm{Im}(\gamma), \widetilde{\Sigma}) > \varepsilon$ (again because $\mathrm{Im}(\gamma)$ is compact). The set $D := \{x \in \widetilde{P} \mid d(x, \mathrm{Im}(\gamma)) \leq \varepsilon\}$ is compact and simply connected. For all but finitely many $\widetilde{P}^{(n)}$ there exists an immersion $\iota^{(n)} \colon D \rightsquigarrow \widetilde{P}^{(n)}$. Then similar as before we see that $\iota^{(n)}(o_P, g)$ lies in $P^{(n)} \times \{g\}$ for sufficiently large n and with Lemma 4.17 (b) follows. Now let $K \subseteq \widetilde{P}$ be a compact path-connected subset containing the basepoint and $\iota^{(n)} \colon K \rightsquigarrow \widetilde{P}^{(n)}$ be the corresponding immersions. If $(o_{P^{(n)}}, h_n) \in \widetilde{P}^{(n)}$ converges to $x \in K$ then (again similar as before) there is an $h \in F$ such that $h_n = h$ for sufficiently large n . Therefore $x = (o_P, h)$ and (c) follows. $\square$

5 Families in $\mathcal{M}$

In this section, we want to use the tools introduced in Section 4 to establish coordinates on some families of infinite translation surfaces. These families consist of translation surfaces arising from PSI-polygons as established in Proposition 4.5. So let $(P, (x_n)_{n\in\mathbb{N}})$ be a PSI-polygon. When we compare it to the finite case considered in Example 2.14, the vector $(t(x_n) - s(x_n))_{n\in\mathbb{N}} \in \mathbb{C}^\infty$ seems like a suitable coordinate for $\widehat{P}$ (using Convention 4.3). However, when we take a closer look at it this causes several problems. The main problem will be the one mentioned in Section 4.2, namely that the convergence of the corresponding translation surfaces does not automatically imply the convergence of the PSI-polygons. Nevertheless, the above coordinates will be our guiding thought in the following section. Our main approach will be to look at some families of such surfaces such that the above implication holds. We will start by looking at some small families in Section 5.1 and then continue with two bigger families in Section 5.2 and Section 5.3.

5.1 Some simple families

As noted before the main idea of this approach of finding coordinates for infinite translation surfaces comes from looking at generalizations of the Chamanara surface. In this subsection, we want to look at some families which are still pretty similar to the Chamanara surface.

The first such family is the one introduced in [Cha04]. There the author describes a surface Ch_α for each $\alpha \in (0,1)$ as mentioned in Example 2.3. One very straightforward approach would be to use this α as a coordinate for such a translation surface. The following theorem will show that this approach is indeed viable.

Theorem 5.1. *The map*

$$\Psi\colon (0,1) \to \mathcal{M},\ \alpha \mapsto \mathrm{Ch}_\alpha$$

is a homeomorphism onto its image.

Proof. For $\alpha \in (0,1)$ let P_α be the PSI-polygon constructed in Figure 2, i.e. the PSI-polygon the Chamanara surface Ch_α is constructed from.

- Injective: Let $\alpha_1 \neq \alpha_2 \in (0,1)$.
  If $\mathrm{Ch}_{\alpha_1} = \mathrm{Ch}_{\alpha_2}$ holds in $\mathcal{M}$, there exists a translation isomorphism $\Phi\colon \mathrm{Ch}_{\alpha_1} \to \mathrm{Ch}_{\alpha_2}$. Because Φ has to send the basepoint of Ch_{α_1} to the basepoint of Ch_{α_2} it also sends every point of Ch_{α_1} corresponding to a point in $P^\circ_{\alpha_1}$ to the point of Ch_{α_2} corresponding to the same point. Now assume w.l.o.g. that $\alpha_2 < \alpha_1$. Then Φ has to map $\widehat{(\alpha_2, 1)} \in \mathrm{Ch}_{\alpha_1}$ to $\widehat{(\alpha_2, 1)}$, but $\widehat{(\alpha_2, 1)} \notin \mathrm{Ch}_{\alpha_2}$. ↯

- Continuous: Let $(\alpha_n)_{n\in\mathbb{N}}$ be a sequence in $(0,1)$ which converges to $\alpha \in (0,1)$. Then P_{α_n} converges to P_α and therefore it follows with Theorem 4.18 that Ch_{α_n} converges to Ch_α.
- Open: Let Ch_{α_n} be a sequence in $\mathrm{Im}(\Psi)$ which converges to $\mathrm{Ch}_\alpha \in \mathrm{Im}(\Psi)$. Then the universal covers $\widetilde{\mathrm{Ch}_{\alpha_n}}$ also converge to $\widetilde{\mathrm{Ch}_\alpha}$. We now assume that α_n does not converge to α. Then there exists $\varepsilon \in \mathbb{R}_+$ such that for all $N \in \mathbb{N}$ there is $n \in \mathbb{N}$, $n > N$ with $|\alpha - \alpha_n| > \varepsilon$. We can safely assume that $\varepsilon < \alpha$.
  In particular there are infinitely many α_n that meet this condition. Let
  $$D = \Big([\delta, 1-\delta]^2 \cup [\delta, \alpha - \frac{\varepsilon}{2}] \times [\delta, 1]\Big) \times \{e\} \subseteq \widetilde{\mathrm{Ch}_\alpha}$$
  for some sufficiently small $\delta \in \mathbb{R}_+$. Then D is homeomorphic to a closed disc and with Theorem 3.12(a) it follows that $D \rightsquigarrow \mathrm{Ch}_{\alpha_n}$ for almost all $n \in \mathbb{N}$. So $\alpha_n > \alpha - \frac{\varepsilon}{2}$ for almost all $n \in \mathbb{N}$. But still $|\alpha - \alpha_n| > \varepsilon$ holds for infinitely many $n \in \mathbb{N}$ and therefore also $\alpha_n > \alpha + \varepsilon$ holds for infinitely many α_n. But then the set $Q = (0,1)^2 \cup (0, \alpha+\varepsilon) \times (1-\delta, 1+\delta) \subseteq \mathbb{R}^2$ is a surface in $\widetilde{\mathcal{M}}$ and can be embedded in each of these infinitely many $\widetilde{\mathrm{Ch}_{\alpha_n}}$ but not in $\widetilde{\mathrm{Ch}_\alpha}$ which is a contradiction to Theorem 3.12(b).

□

Remark 5.2. *The sequence* Ch_{α_n} *converges to* $(0,1)^2$ *if* α_n *converges to* 0 *and it converges to the punctured torus if* α_n *converges to* 1.

Another very straightforward approach is to vary the shape of the PSI-polygon itself instead of its edges. This leads to two additional rather simple families.

Proposition 5.3. *(a) For* $\beta \in (0,\pi)$ *let* $(P_\beta, (x_n)_{n\in\mathbb{N}})$ *be the PSI-polygon such that* P_β *is the rhombus with side length* 1 *and with angle* β *at one corner. In addition, we divide the sides in the same way as they are divided in* $\mathrm{Ch}_{\frac{1}{2}}$*, with one accumulation point at the corner for which the angle is* β*. Then the map*

$$\Psi_a \colon (0,\pi) \to \mathcal{M},\ \beta \mapsto \widehat{P_\beta}$$

is a homeomorphism onto its image.

(b) For $l \in \mathbb{R}_+$*, let* $(P_l, (x_n)_{n\in\mathbb{N}})$ *be the PSI-polygon such that* P_l *is the square with side length* l*. Again we divide the sides in the same way as they are divided in* $\mathrm{Ch}_{\frac{1}{2}}$*. Then the map*

$$\Psi_b \colon \mathbb{R}_+ \to \mathcal{M},\ l \mapsto \widehat{P_l}$$

is a homeomorphism onto its image.

Sketch of Proof. The proof is very similar to the previous one, so we will only sketch it. For similar reasons as before these maps are injective and continuous.

To see that Ψ_a and Ψ_b are open we can take a sequence $\widehat{P}^{(n)}$ converging to $\widehat{P}$ in the image. Then, for all $\varepsilon \in \mathbb{R}_+$, the set $D_\varepsilon \subseteq P$ of all points with distance greater or equal to ε from ∂P can be embedded in $\widetilde{P}$ and therefore in almost all $\widetilde{P}^{(n)}$. This is sufficient to show that Ψ_a is open. For Ψ_b it is still possible that infinitely many of these $P^{(n)}$ are too large. But then if $P = P_l$ and ε sufficiently small we can take the set $Q_\varepsilon = P^\circ_{l+\varepsilon}$ which can be embedded in these infinitely many $\widetilde{P}^{(n)}$ but not in $\widetilde{P}$. □

At last, these three families can be combined, which is the contend of the following Theorem:

Theorem 5.4. *Let* $(\mathrm{Ch}_{\alpha,\beta,l}, (x_n)_{n\in\mathbb{N}})$ *be the PSI-polygon such that* $\mathrm{Ch}_{\alpha,\beta,l}$ *is the rhombus with side length* l *and with angle* β *at one corner. In addition we divide the sides in the same way as they are divided in* Ch_α *with one accumulation at the corner for which the angle is* β*. Then the map*

$$\Psi\colon (0,1) \times (0,\pi) \times \mathbb{R}_+ \to \mathcal{M},\ (\alpha,\beta,l) \mapsto \widehat{\mathrm{Ch}}_{\alpha,\beta,l}$$

is a homeomorphism onto its image.

Proof. This is essentially a combination of the previous statements in this section: Let $x \in (0,\pi) \times (0,1) \times \mathbb{R}_+$. A sequence in this space that does not converge to x cannot converge in all three coordinates. Then we can use the same sets as before for the coordinate for which this sequence does not converge. □

5.2 An infinite-dimensional real family

We will now return to the situation described in Example 2.16 and construct an infinite-dimensional family. In Example 2.16 we have seen a generalization of the Chamanara surface where the lengths of the edges were variable instead of being always in the same proportion. We have argued that these edges viewed as vectors in $\mathbb{C}$ make up suitable coordinates for such a translation surface. We will now pursue this approach. But instead of vectors in $\mathbb{C}$, we will take the lengths of these vectors as coordinates in $\mathbb{R}$ because in this case, such coordinates carry the same information. In fact, our image will lie in the following space:

Definition 5.5. By $c_0 = c_0(\mathbb{R})$ we will denote the real vector space of all null sequences in $\mathbb{R}$ together with the supremum norm $||(x_n)_{n\in\mathbb{N}}||_\infty = \sup\{|x_n| \mid n \in \mathbb{N}\}$.

Compared to Example 2.16 we will also need some additional technical restriction which will be motivated in the next example.

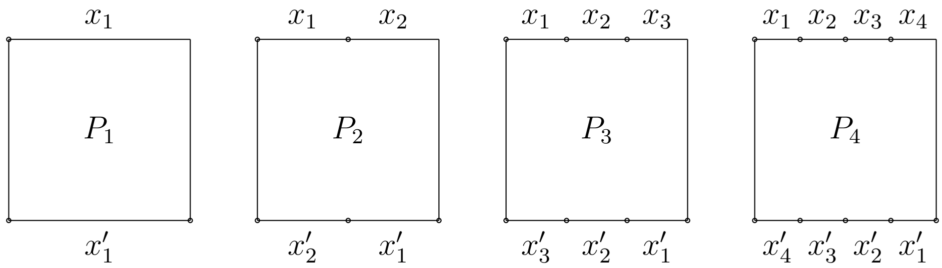

Figure 9: The polygons $P_1, \dots, P_4$. These polygons give rise to finite translation surfaces $\widehat{P}_n$ with genus $\lceil \frac{n}{2} \rceil$.

Example 5.6. We will construct a family of finite translation surfaces. For this let P_n be the (normal) polygon consisting of the set $[0,1]^2 \subseteq \mathbb{R}^2$ such that the upper and lower sides are split into n edges of equal length. The right and left side of these polygons are only one edge in each case. Figure 9 shows examples of these polygons for $n = 1, \dots, 4$. Then again as described in Example 2.14 (or similar to Proposition 4.5) we can construct finite translation surfaces $\widehat{P}_n$ out of these polygons. In Example 2.14 we have seen that

$$\widehat{P}_n \in \begin{cases} \mathcal{H}(n-1), & \text{if } n \text{ is odd,} \\ \mathcal{H}(\frac{n-2}{2}, \frac{n-2}{2}), & \text{if } n \text{ is even.} \end{cases}$$

If we apply the formula $2g - 2 = \sum_{i=1}^{m} k_i$ from Section 2.2 it follows that the genus g_n of $\widehat{P}_n$ satisfies

$$g_n = \begin{cases} \frac{n+1}{2}, & \text{if } n \text{ is odd,} \\ \frac{n}{2}, & \text{if } n \text{ is even.} \end{cases}$$

As in Figure 9 we label the upper sides by $x_1, x_2, \dots$ starting from left to right and the lower sides $x'_1, x'_2, \dots$ accordingly. Now for $m \in \mathbb{N}$, $m \geq n$ let $P_n^{(m)}$ be the polygon which we obtain from P_n by replacing the vertex $t(x_1)$ by an edge x_{ins} with length $\frac{1}{m}$ and midpoint at the point $t(x_1)$ and adding vertices accordingly. We do the same with x'_1. See Figure 10 for an example of this construction.

Then the sequence $(\widehat{P}_n^{(m)})_{m \in \mathbb{N}}$ converges to $\widehat{P}_n$ in the immersive topology. This can be seen in a very similar way to the proof of Theorem 4.18:

Part (a) of Theorem 3.12 follows in the same way as in Theorem 4.18. Now let $Q \in \widetilde{\mathcal{M}}$ such that Q can be embedded in infinitely many $\widetilde{P}_n^{(m)}$. Then this inclusion cannot cross a copy of x_{ins} in $\widetilde{P}_n^{(m)}$ if m is big enough, because this edge gets arbitrarily small for big m. Similar to Theorem 4.18 $Q \rightsquigarrow \widetilde{P}_n$ follows. The rest of Theorem 3.16 follows in a similar manner.

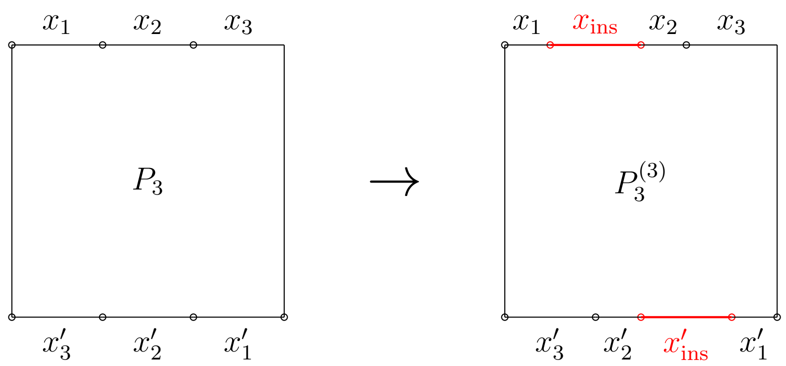

Figure 10: The polygons P_3 and $P_3^{(3)}$. The polygon $P_n^{(m)}$ is constructed from the polygon P_n by replacing the second vertex of the upper side by an edge of length $\frac{1}{m}$ and accordingly for the lower side.

This example is problematic for two reasons. The first is the insight that the genus is not preserved under limits in the immersive topology. We will return to this problem in Section 6. The second one concerns our planned way of assigning coordinates to the above-mentioned infinite translation surfaces.

For this, we observe that Example 5.6 can easily be expanded to the infinite case. Let us look for example at the polygons P_3 and $P_3^{(m)}$ and replace the rightmost edge of the upper side with infinitely many arbitrary edges which accumulate only in the upper right corner. We can do this in the same way for each $m \in \mathbb{N}$, i.e. we can take the same edges as replacement. We then split the leftmost edge of the lower side accordingly. The resulting PSI-polygons will be named Q_3 and $Q_3^{(m)}$ and the corresponding translation surfaces $\widehat{Q}_3^{(m)}$ converge to $\widehat{Q}_3$ for the same reasons as in Example 5.6. However, the coordinates mentioned before (ignoring the vertical edges) would be

$$\left(\frac{1}{3} - \frac{1}{2m}, \frac{1}{m}, \frac{1}{3} - \frac{1}{2m}, l_1, l_2, \dots\right) \in c_0$$

where the l_i are the lengths of these newly introduced edges. But for $m \to \infty$, this does not converge to

$$\left(\frac{1}{3}, \frac{1}{3}, l_1, l_2, \dots\right) \in c_0$$

which would be the coordinates of $\widehat{Q}_3$. This makes it necessary to introduce some conditions on these vertices in such a way that two vertices of such a PSI-polygon cannot be arbitrarily close. However, the vertices should still be able to accumulate at the lower left and upper right corners. The solution here is to forbid that these vertices are too close together when compared to the distance to an accumulation point. This is formulated in the following condition. Please note that this will be more general than necessary at the moment, but we will use this condition again later at which point the general case is needed.

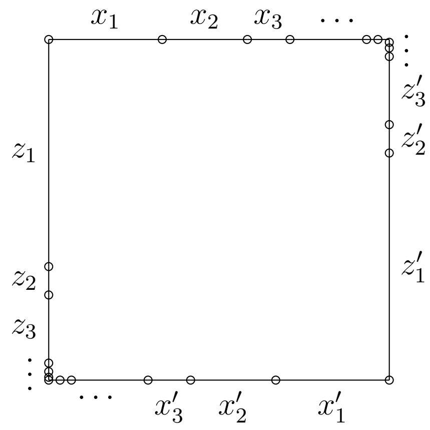

Figure 11: An example for a polygon in Re-Ch$_\delta$.

Definition 5.7. Let $\delta \in \mathbb{R}_+$. We say that a PSI-polygon P fulfills the *δ-condition* if for every two non-adjacent edges x_i and x_j it holds that

$$d(x_i, x_j) > \delta \cdot \max\left(\min_{y \in Y}(d(x_i, y)), \min_{y \in Y}(d(x_j, y))\right)$$

where Y denotes the set of all accumulation points of vertices of P.

We are now ready to formally define our aforementioned family.

Definition 5.8. For $\delta \in (0, \frac{1}{2})$ we define Re-Ch$_\delta$ to be the set of PSI-polygons fulfilling the δ-condition such that for each such PSI-polygon $(P, (x_n)_{n\in\mathbb{N}})$ the set P is exactly $[-1, 1]^2 \subseteq \mathbb{R}^2$ and the points $s(x_n)$ have exactly two accumulation points positioned at $(-1, -1)$ and $(1, 1) \in \mathbb{R}^2$. $\widehat{\text{Re-Ch}_\delta} \subseteq \mathcal{M}$ is the set of isomorphism classes of translation surfaces arising from polygons in Re-Ch$_\delta$.

Figure 11 shows an example of such a PSI-polygon. We will always denote the edges on $[-1, 1] \times \{1\}$ by $x_1, x_2, \ldots$ starting from the left and the edges on $\{-1\} \times [-1, 1]$ by $z_1, z_2 \ldots$ starting from the top. The point y will always denote the accumulation point $(1, 1)$ and the endpoints of these edges will always be the ones closer to y.

Please note that compared to Example 2.16, we changed the set to be $[-1, 1]^2$ instead of $[0, 1]^2$ so that the point of symmetry is now $(0, 0)$.

Now each translation surface in $\widehat{\text{Re-Ch}_\delta}$ is constructed from a unique PSI-polygon in Re-Ch$_\delta$ (see the following proof). This enables us to use this polygon to define infinite-dimensional coordinates for this translation surface. This is the content of the following theorem.

Theorem 5.9. *The map*

$$\Psi\colon \widehat{\text{Re-Ch}_\delta} \to c_0^2,\ \widehat{P} \mapsto ((d(t(x_n),y))_{n\in\mathbb{N}}, (d(t(z_n'),y))_{n\in\mathbb{N}})$$

is a homeomorphism onto its image.

Proof.

- well-defined: Because the sequence $(t(x_n))_{n\in\mathbb{N}}$ converges to y, the sequence $(d(t(x_n),y))_{n\in\mathbb{N}}$ does converge to 0. In addition, similar to Theorem 5.1, no two such PSI-polygons give rise to isomorphic translation surfaces. So each $\widehat{P} \in \widehat{\text{Re-Ch}_\delta}$ comes from a unique $P \in \text{Re-Ch}_\delta$.
- injective: Such a PSI-polygon is uniquely given by the above coordinates so the map is injective.
- open: This follows directly from Theorem 4.18.
- continuous: Let $P^{(n)}, P \in \text{Re-Ch}_\delta$ be PSI-polygons such that $\widehat{P}^{(n)}$ converges to $\widehat{P}$. Assume that $\Psi(\widehat{P}^{(n)})$ does not converge to $\Psi(\widehat{P})$. Then there is $\varepsilon \in \mathbb{R}_+$ such that for infinitely many $n \in \mathbb{N}$ an $i_n \in \mathbb{N}$ exists such that $t(x_{i_n}^{(n)}) \notin B_\varepsilon(t(x_{i_n}))$ or $t(z_{i_n}^{(n)}) \notin B_\varepsilon(t(z_{i_n}))$. W.l.o.g. we can assume that infinitely many of these vertices are of the form $t(x_{i_n}^{(n)})$. Let us first assume that there is $i \in \mathbb{N}$ such that $t(x_i^{(n)}) \notin B_\varepsilon(t(x_i))$ for infinitely many $n \in \mathbb{N}$. Then these points $t(x_i^{(n)})$ have to accumulate at a point $(a,1) \in [-1,1] \times \{1\}$.

  Case 1: $a < \pi_1(t(x_i))$[5]. If there is no $j \in \{1,\dots,i-1\}$ with $t(x_j) = (a,1)$, then there is $\xi \in \mathbb{R}_+$ such that $B_\xi((a,1))$ does not contain any vertex of P. We can then take the two line segments l_1 from $(0,0)$ to $(a,0)$ and l_2 from $(a,0)$ to $(a,1)$ and the set

  $$D \coloneqq \overline{B_\xi(l_1 \cup l_2)} = \{x \in \mathbb{R}^2 \mid d(x,l_1) \le \xi \text{ or } d(x,l_2) \le \xi\}.$$

  Then D is homeomorphic to a disc and can be immersed in $\widehat{P}$ (and $\widetilde{P}$) but not in these infinitely many $\widehat{P}^{(n)}$ (or $\widetilde{P}^{(n)}$), which is a contradiction to Theorem 3.16. If there is $j \in \mathbb{N}$ such that $t(x_j) = (a,1)$, we can deduce from the δ-condition that the points $t(x_j^{(n)})$ cannot accumulate at the point $(a,1)$ because the points $t(x_i^{(n)})$ already accumulate at this point. Therefore the points $t(x_j^{(n)})$ have to accumulate at a point to the left of $(a,1)$ and we can repeat the argument with this point. Because there are only finitely many vertices on $[-1,1] \times \{1\}$ to the left of $t(x_i)$ there has to be one $k \in \mathbb{N}$ such that $t(x_k^{(n)})$ accumulates at a point which is not a vertex of P.

  Case 2: $a > \pi_1(t(x_i))$. Similar as for case 1, we can deduce that there is $k \in \{1,\dots,i\}$ and $\xi \in \mathbb{R}_+$ such that for infinitely many $n \in \mathbb{N}$ no vertex of $P^{(n)}$ lies in $B_\xi(t(x_k))$. Now we can choose the two line segments l_1 from 0 to $(\pi_1(t(x_k)),0)$

[5] By $\pi_1\colon \mathbb{R}^2 \to \mathbb{R}$ we denote the projection to the first coordinate.

and l_2 from $(\pi_1(t(x_k)), 0)$ to $t(x_k)$. Then the set $Q := B_\xi(l_1 \cup l_2)$ can be immersed in infinitely many $\widetilde{P}^{(n)}$ but not in $\widetilde{P}$, which is again a contradiction to Theorem 3.16.

If such an i does not exist we instead look at the set $V := [-1, 1-\varepsilon] \times \{1\}$. Then only finitely many vertices of P lie in V. Let $k \in \mathbb{N}$ be the greatest number such that $t(x_k) \in V$. According to the requirements, for infinitely many $n \in \mathbb{N}$ there is $i_n \in \mathbb{N}$ such that $t(x_{i_n}^{(n)}) \notin B_\varepsilon(t(x_{i_n}))$. But this also means that $t(x_{i_n}^{(n)}) \in V$. Because infinitely many of these i_n are greater than $k+1$ and because for those it holds that $\pi_1(t(x_{k+1}^{(n)})) < \pi_1(t(x_{i_n}^{(n)}))$ we can conclude that $t(x_{k+1}^{(n)}) \in V$. But then $t(x_{k+1}^{(n)})$ cannot accumulate at $t(x_{k+1})$ and we can use the same argument as for case 1.

□

Remark 5.10. *Similar to Theorem 5.4, we can also vary the lengths of the sides and the angles between the sides to get a family with coordinates in* $c_0^2 \times (0, \pi) \times \mathbb{R}_+$.

5.3 A finite-dimensional family of infinite translation surfaces

We now want to introduce a different kind of family. In the last subsection, we introduced an infinite-dimensional real family. In contrast, we will now construct a finite-dimensional complex family of infinite translation surfaces. The idea behind this is simply to fix all but finitely many edges of a PSI-polygon and use these finitely many edges as coordinates. In our case, the fixed edges will be exactly like the edges of the Chamanara surface. The following definition will make this more precise:

Definition 5.11. Let $n \in \mathbb{N}$ and Fi-Ch_n be the set of PSI-polygons P of the following form:

- $o_P = 0 \in \mathbb{C}$.
- There are exactly 2 accumulation points y and y' of the vertices of P. In addition $d(y, y') > 2$ holds.
- The two parts of ∂P adjacent to y with length 1 are the same (up to translation) as the upper and right side of the Chamanara surface $\text{Ch}_{\frac{1}{2}}$ and the interior of P lies down left of these two sides. These parts will be called the *converging sides*. In addition, there is no other edge of P cutting through the triangle spanned by the two converging sides.
- Apart from these two parts P has exactly $2n$ edges, called *finite edges*. The corresponding start- and endpoints of these edges will be called the *finite vertices*.

$\widehat{\text{Fi-Ch}_n} \subseteq \mathcal{M}$ is the set of isomorphism classes of translation surfaces arising from polygons in Fi-Ch_n.

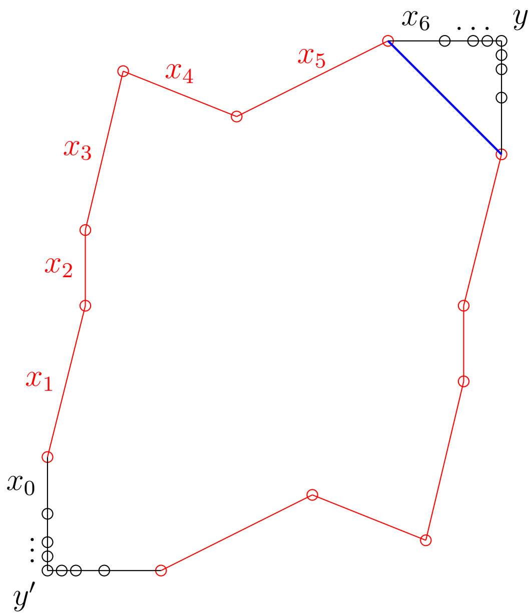

Figure 12: A PSI-polygon in Fi-Ch_5. The red edges and vertices are the finite edges and vertices. The black parts are the converging sides and no edge can lie in the triangle consisting of the blue line and the two adjacent converging sides.

Figure 12 shows an example of such a polygon. Of course the set $\widehat{\text{Fi-Ch}_0}$ would only consist of the Chamanara surface $\text{Ch}_{\frac{1}{2}}$.

Convention 5.12. Let $n \in \mathbb{N}$ and $P \in \text{Fi-Ch}_n$. We will denote the finite edges on one side by $x_1, \ldots, x_n$ such that each edge x_i is adjacent to x_{i-1} and x_{i+1} for $i \in \{2, \ldots, n-1\}$. The corresponding edges on the other side will be denoted by x_i'. We will continue this labeling onto the converging sides, so the edge in the converging side adjacent to x_n is labeled x_{n+1} and the edge adjacent to x_1 is labeled x_0.

Our basic idea now is to use the finite edges as coordinates and get a homeomorphism to $\mathbb{C}^n$. This is, in principle, not very difficult. However, a general problem with our way of creating coordinates appears and has to be tackled. The following example will demonstrate this problem:

Example 5.13. Let $P^{(1)}$ be the PSI-polygon sketched in Figure 13 and $P^{(2)}$ be the PSI-polygon sketched in Figure 14. Then both $P^{(1)}$ and $P^{(2)}$ lie in Fi-Ch_4. But $\widehat{P}^{(1)}$ and $\widehat{P}^{(2)}$ are translation isomorphic via the translation isomorphism Φ sketched in Figure 15.

This behavior causes some problems. $\widehat{\text{Fi-Ch}_4}$ consists of isomorphism classes of translation surfaces. Hence the classes that $\widehat{P}^{(1)}$ and $\widehat{P}^{(2)}$ belong to are actually the same. Therefore our usual procedure is not sufficient to assign unique coordinates to this class. This is not a very surprising behavior when compared to the finite case mentioned in

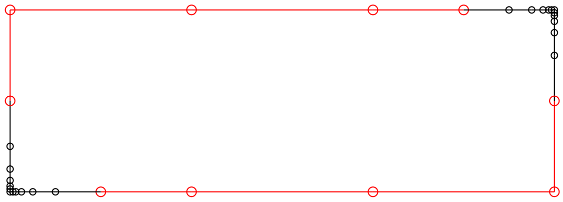

Figure 13: The PSI-polygon $P^{(1)} \in \text{Fi-Ch}_4$. Again the finite edges and vertices are marked in red.

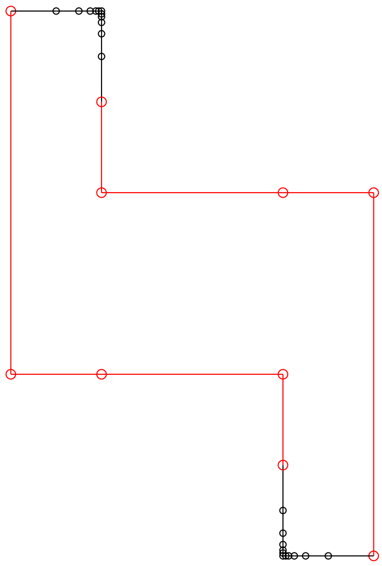

Figure 14: The PSI-polygon $P^{(2)} \in \text{Fi-Ch}_4$.

Section 2.2. There, for each isomorphism class of finite translation surfaces, we get a chart for some neighborhood of this class in the moduli space by choosing a basis of the relative homology (ignoring orbifold issues). In the present case, for the isomorphism class of $\widehat{P}$ with $P \in \text{Fi-Ch}_n$, the choice of the PSI-polygon P corresponds to selecting such a basis. Because the families we have considered up to now were rather restricted and a PSI-polygon has many restricting properties as well, the generating PSI-polygon was unique for these families. However Example 5.13 shows that these restrictions on PSI-polygons are not enough to ensure uniqueness in a more general case. But the solution to this problem can be found analogously to the finite case by restricting ourselves to local coordinates instead of global ones. In fact, we will see that for each PSI-polygon $P \in \text{Fi-Ch}_n$ we get an open neighborhood of the isomorphism class of $\widehat{P}$ such that P gives rise to unique coordinates from this neighborhood to an open subset of $\mathbb{C}^n$. This will give $\widehat{\text{Fi-Ch}_n}$ the structure of an n-dimensional complex manifold. We start by introducing these open neighborhoods.

Definition 5.14. Let $n \in \mathbb{N}$, $\varepsilon \in \mathbb{R}_+$ and $P \in \text{Fi-Ch}_n$. Then the set $B_\varepsilon(P) \subseteq \text{Fi-Ch}_n$

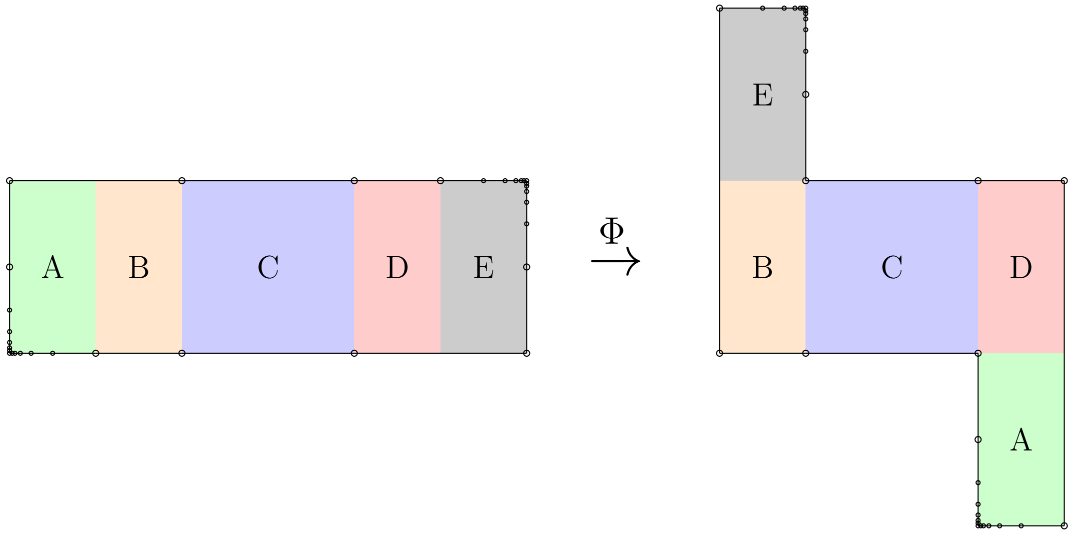

Figure 15: The map Φ is obtained by mapping the areas A-E of $\widehat{P}^{(1)}$ on the left to the corresponding areas of $\widehat{P}^{(2)}$ on the right via translation. Then Φ maintains the connections between these areas and is therefore a translation isomorphism.

consists of all PSI-polygons $P^{(1)} \in \text{Fi-Ch}_n$ such that for each finite edge $x_i^{(1)}$ of $P^{(1)}$

$$d(s(x_i), s(x_i^{(1)})) < \varepsilon$$

holds, where the distance is measured in $\mathbb{C}$. Similar to before $\widehat{B_\varepsilon(P)} \subseteq \mathcal{M}$ denotes the set of isomorphism classes of translation surfaces arising from polygons in $B_\varepsilon(P)$.

Note that for the last finite vertex $t(x_n)$ it also holds that

$$d(t(x_n), t(x_n^{(1)})) < \varepsilon$$

because for each $P \in \text{Fi-Ch}_n$

$$t(x_n) = s(x_1') + (-1, 1) = \Phi_P(s(x_1)) + (-1, 1)$$

holds.

Of course $\widehat{B_\varepsilon(P)}$ is a neighborhood of $\widehat{P}$ in $\widehat{\text{Fi-Ch}_n}$, but to be of any use, this neighborhood has to be small enough, so ε has to be sufficiently small. For $P \in \text{Fi-Ch}_n$ let $\delta_P \in \mathbb{R}_+$ such that

$$\delta_P < d(x_i, \partial P \setminus (x_i \cup x_{i-1} \cup x_{i+1}))\ \forall i \in \{1, \dots, n\}.$$

From now on, let $\varepsilon \in \mathbb{R}_+$ such that $\varepsilon < \frac{1}{16}\delta_P$ and $\varepsilon < \frac{1}{2^{n+5}}$. In addition, ε should fulfill another condition. For better readability we will state this condition not until it is needed in the proof.

This guarantees that each of the $2n$ finite edges has length greater than 16ε. Our next goal is to show that such a set $\widehat{B_\varepsilon(P)}$ is open in the immersive topology:

Lemma 5.15. *Let $n \in \mathbb{N}$, $P \in$ Fi-Ch$_n$ and $\varepsilon \in \mathbb{R}_+$ that fulfills the above mentioned conditions for P. Then $\widehat{B_\varepsilon(P)}$ is open in the immersive topology.*

Proof. We will give an explicit description of $\widehat{B_\varepsilon(P)}$ as a finite intersection of open sets. Unfortunately, this requires some work and therefore this proof will be very long. To accommodate this, we subdivide the proof into 4 parts:

In part 1 we construct V as an intersection of some sets which are open in the immersive topology.
In part 2 we show that $\widehat{B_\varepsilon(P)}$ is a subset of V.
In part 3 we start to show that V is a subset of $\widehat{B_\varepsilon(P)}$.
In part 4 we tackle one case omitted from part 3, namely the case mentioned in Example 5.13: It is possible that there are two PSI-polygons $P^{(1)}, P^{(2)} \in$ Fi-Ch$_n$ such that $\widehat{P}^{(1)} \cong \widehat{P}^{(2)}$. But this means that $\widehat{P}^{(1)} \in \widehat{B_\varepsilon(P)}$ does not automatically imply $P^{(1)} \in B_\varepsilon(P)$. This can happen, if $P^{(2)}$ lies in $B_\varepsilon(P)$ instead (only one of them can lie in $B_\varepsilon(P)$, compare Theorem 5.16). So in part 4, we show that for each translation surface $X \in V$ there is at minimum one $P \in B_\varepsilon(P)$ such that $\widehat{P} \cong X$. We want to emphasize that this is the most important part of the proof, because it shows that picking an open neighborhood $\widehat{B_\varepsilon(P)}$ actually corresponds to choosing one of the generating PSI-polygons for each translation surface $X \in \widehat{B_\varepsilon(P)}$.

Part 1: In this part, we will introduce the open sets which are used later. At the end of part 1 we will define the set

$$V := \mathcal{M}_{\hookrightarrow}(D_1) \cap \mathcal{M}_{\rightsquigarrow}(D_2) \cap \bigcap_{i=-n}^{2n+1} \mathcal{M}_+(K_i, U_i) \cap \widehat{\text{Fi-Ch}_n}$$

which is open in $\widehat{\text{Fi-Ch}_n}$ by definition (compare Section 3). Before we begin to define the individual sets, we want to summarize their usage in part 3. For this let $P^{(1)} \in$ Fi-Ch$_n$ such that $\widehat{P}^{(1)} \in V$.

- For the set $\mathcal{M}_{\hookrightarrow}(D_1)$ we will see that $D_1 \subsetneq (P^{(1)})^\circ$ when seen as subsets of $\mathbb{C}$. This allows us, to use D_1 as some kind of starting point to arrange the edges of $P^{(1)}$ around it. For this to be valid, we have to ignore the issues tackled in part 4 for a moment.
- The set $\mathcal{M}_+(K_i, U_i)$ allows us, to know the positions of the midpoints of the edges of $P^{(1)}$ relative to D_1. This will be illustrated by "arrows" going from one side of D_1 to the opposite side (compare Figure 22). In particular, this allows us to know the position of the converging sides relative to D_1.

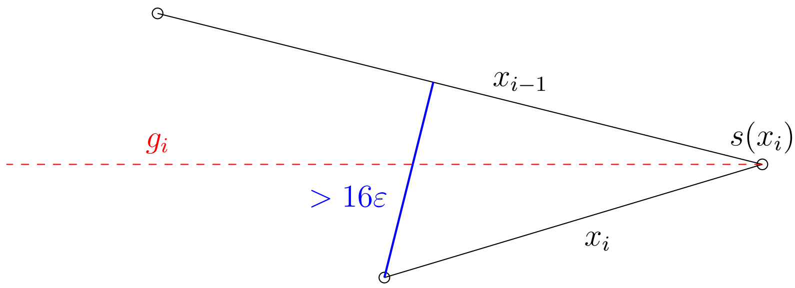

Figure 16: The endpoint of the shorter of two adjacent edges has distance greater than 16ε from the other edge, because it also belongs to the next edge, in this case to the edge x_{i+1}.

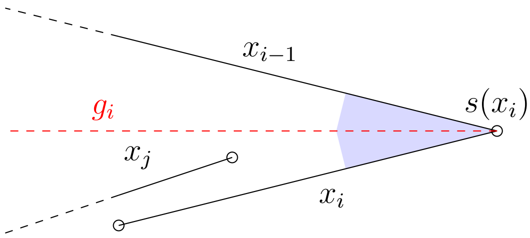

Figure 17: Every point in the area marked in blue is closer than ε to x_{i-1} or x_i. The points on g_i left of this area either lie in D_1, or there is another edge x_j closer than ε to these points. But then x_j is too close to x_i or x_{i-1}.

- By knowing the position of the midpoint and one endpoint of an edge up to ε[6], we know the position of the other endpoint only up to 3ε. The set $\mathcal{M}_{\rightsquigarrow}(D_2)$ will be used to overcome this issue.

So, after much preparation, let us start with defining the sets. For this let

$$D_1 \coloneqq \{x \in P \mid d(x, \partial P) \geq \varepsilon\} \subseteq P.$$

Then D_1 is compact and contains the basepoint. In addition D_1 is simply connected. This can be seen by looking at the path that stays exactly ε away from ∂P. Because ε is small enough this path cannot cross itself.
For the same reasons, the set

$$D_1' \coloneqq \{x \in P \mid d(x, \partial P) \geq 3\varepsilon\} \subseteq P$$

is also simply connected, compact and contains the basepoint.

The next set requires a bit more work. For each of the edges with index i between 1 and $n+1$, we take a bisector g_i of the edges x_{i-1} and x_i at the point $s(x_i)$ (see Figure 16). Then g_i has to meet D_1 at some point. This can be seen in the following way:
Assume g_i meets ∂P before it meets D_1. The endpoint of the shorter of the two edges x_i and x_{i-1} has distance greater than 16ε from the other edge (see Figure 16). So there

[6]That means, we know a point $x \in \mathbb{C}$ such that the midpoint (or endpoint) lies in $B_\varepsilon(x) \subseteq \mathbb{C}$.

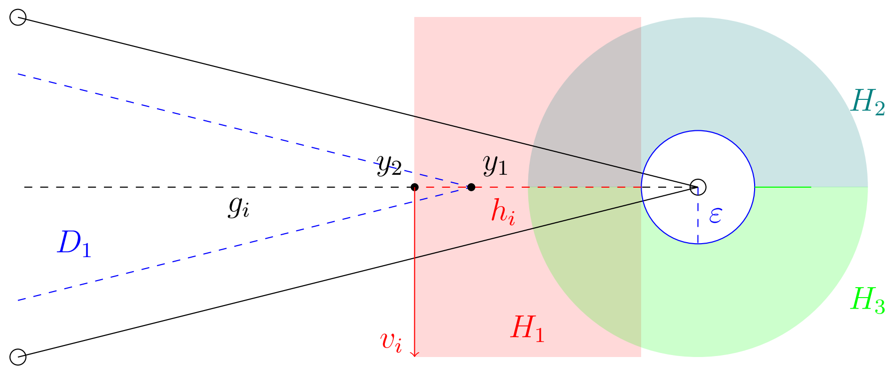

Figure 18: A sketch of the sets H_1, H_2 and H_3. They can be constructed for each g_i and together with D_1 they yield the set D_2.

has to be a point of g_i in the interior of P, with distance greater than ε from these two edges. Because this distance grows continuously, there is also such a point at which this distance is smaller than 2ε. But because this point does not lie in D_1, there has to be another edge x_j closer than ε to this point (see Figure 17). This shows that x_j is closer than 4ε to the edges x_i and x_{i-1}, which is a contradiction. In addition, we have seen that the first point of g_i that lies in D_1 has distance greater than 3ε from each other edge x_j.

Now let $C_i \coloneqq \mathbb{C} \setminus \{s(x_i)\}$. Then P° can be embedded in the universal cover $\widetilde{C}_i$ and therefore, we can view D_1 and g_i as subsets of $\widetilde{C}_i$. We now take the following subset of $\widetilde{C}_i$ as shown in Figure 18: Let $v_i \in \mathbb{C}$ be a vector orthogonal to g_i, with length exactly 3ε. Let y_1 be the first point at which g_i meets D_1. Let y_2 be the point on g_i that has distance ε from y_1 and which is further away from $s(x_i)$ than y_1. By h_i we denote the part of g_i between the first point of distance ε to $s(x_i)$ and y_2. We now take the set

$$H_1 \coloneqq \{u_1 + u_2 \in \widetilde{C}_i \mid u_1 \in h_i, u_2 = a \cdot v_i \text{ with } a \in [-1, 1]\}$$

and combine it with the set H_2, obtained by taking the subset of g_i of points of distance greater or equal than ε and smaller or equal than 3ε to $s(x_i)$ and turning it clockwise for 180 degrees. We then combine it with the set H_3 which is obtained by turning this set counter-clockwise instead. We want to show, that the union of these sets, together with D_1 can be immersed in $\widetilde{P}$.

But first, we want to specify the additional condition for ε. For this, we consider each of the $n+1$ finite vertices. ε should then be small enough, that the red line α in Figure 19 is shorter than $\delta_P - 4\varepsilon$. To be more precise: There should exist a point y_3 on g_i, such that the blue line orthogonal to g_i meets ∂P at distance greater than 3ε from y_3 and such that $d(y_3, s(x_i)) < \delta_P - 4\varepsilon$. In addition, this condition should also hold for each of the finite vertices in each $P^{(1)} \in B_\varepsilon(P)$ and the corresponding value $\delta_{P^{(1)}}$. This is a

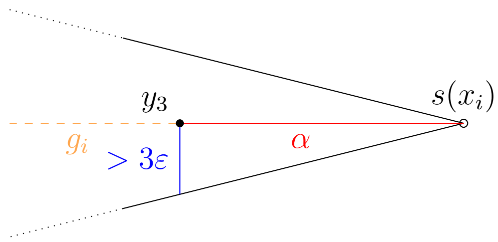

Figure 19: The additional condition on ε states, that ε should be small enough such that the red line α is smaller than $\delta_P - 4\varepsilon$ for each relevant g_i. This is possible, because the blue line gets shorter for smaller ε and therefore also the red line gets shorter. But δ_P is fixed for P.

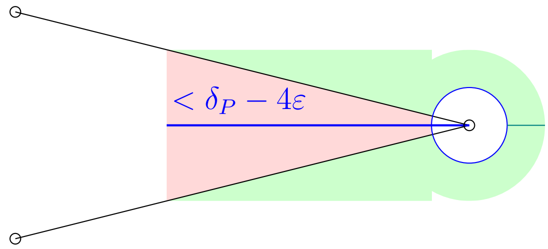

Figure 20: One part of the set D_2 immersed in $\widetilde{P}$. Because of the previous condition for ε, the blue line is short enough, such that each point in the green area is closer than δ_P to each edge. This does also hold for each $P^{(1)} \in B_\varepsilon(P)$.

viable condition, because $\delta_{P^{(1)}}$ can only be 2ε smaller than δ_P and each angle at a vertex in $P^{(1)}$ has to be at least half of the corresponding angle in P (because the distance marked in blue in Figure 16 has to be greater than 16ε).

Next, we want to show, that these sets can be immersed in $\widetilde{P}$. D_1 can of course be immersed. We now take a look at Figure 20. Of course, the inner red part can be immersed in $\widetilde{P}$. The additional condition on ε (compare Figure 19) shows that H_1 collides at earliest $\delta_P - 4\varepsilon$ away from $s(x_i)$ with one of the two edges which are adjacent to $s(x_i)$. Therefore each point in the outer green part is at maximum δ_P away from one of these edges. These sets can therefore intersect a finite number of edges but they cannot contain a vertex:
Such a vertex has to be one of the endpoints of one of the edges it has crossed because every other vertex is too far away. But this is impossible because the intersection of these sets with such an edge has distance greater than ε from the endpoints.

Therefore, these sets can be immersed in $\widetilde{P}$. We now define D_2 to be the union of the images of these sets under this immersion for all of the $n+1$ finite vertices. Note that we have already shown that D_2 can be immersed in $\widehat{P}$ and in each $\widehat{P}^{(1)}$ for each $P^{(1)} \in B_\varepsilon(P)$.

Now, we define for each $i \in \{-n, \dots, 2n+1\}$ the set K_i as the subset of $\widetilde{P}$ consisting of the union of the following sets as depicted in Figure 21:

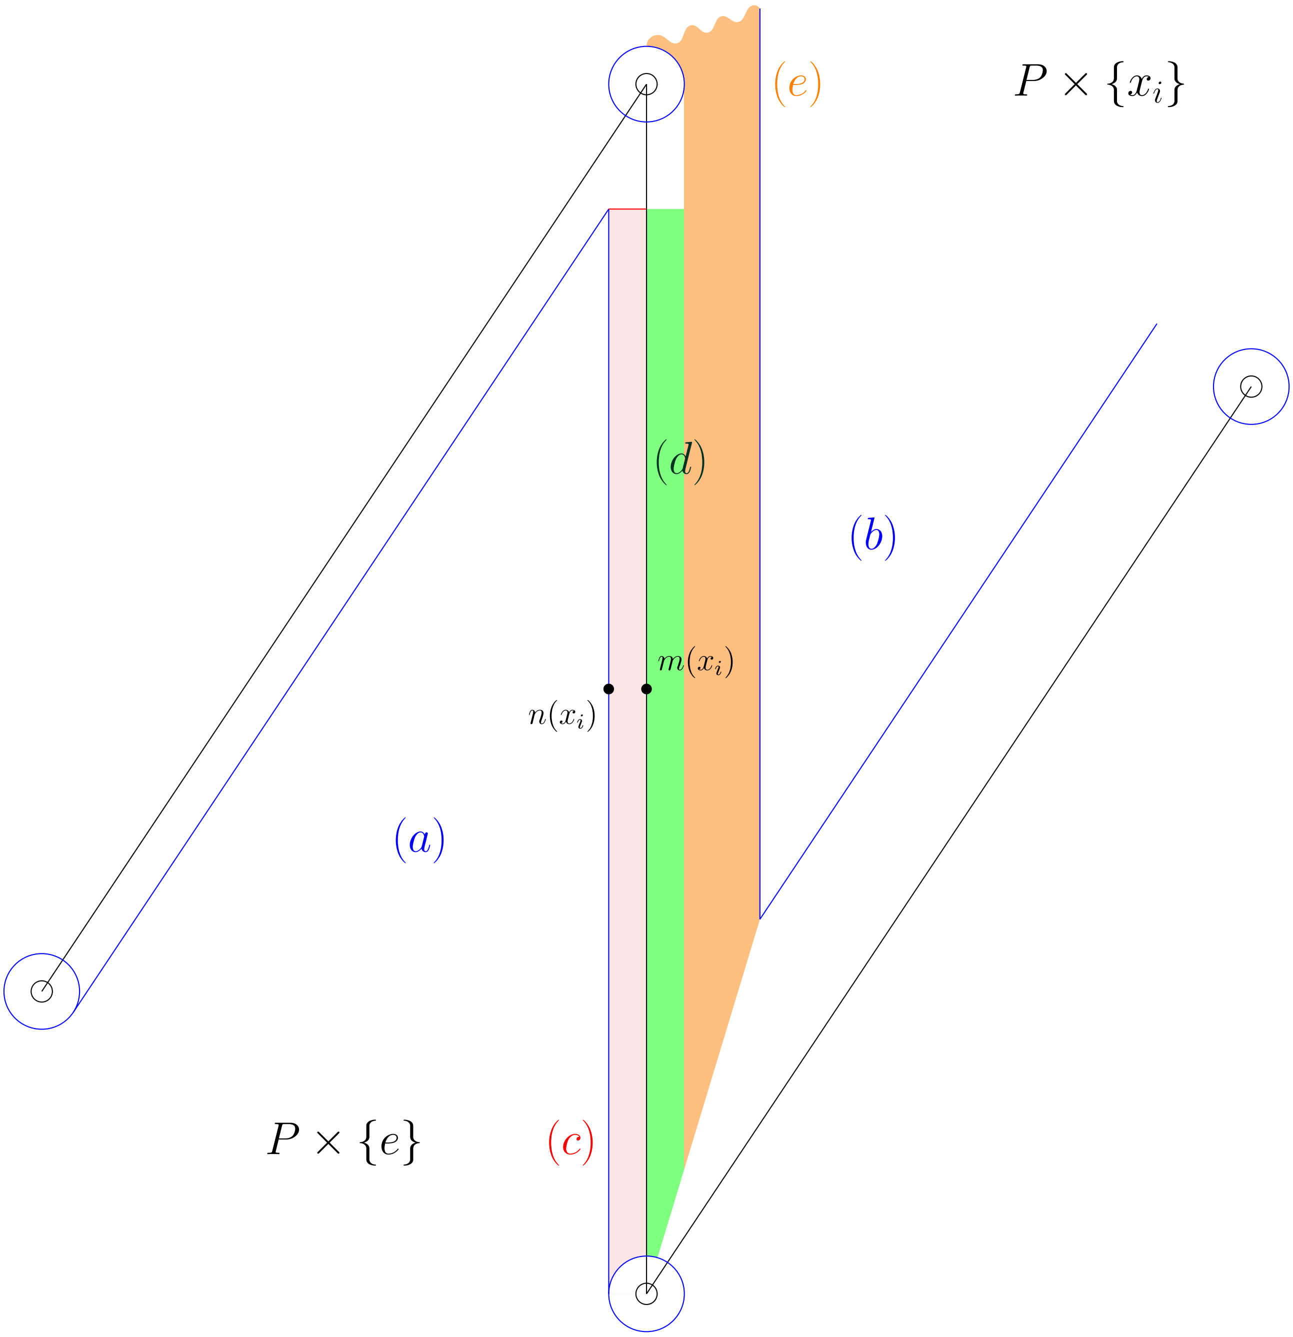

Figure 21: A sketch of the sets (a)-(e) used to create K_i. Because it should be possible to immerse K_i in each $\widehat{P^{(1)}} \in \widehat{B_\varepsilon(P)}$, the set (b) must have distance at least 3ε from δP (compare Lemma 4.17), which makes this construction necessary.

(a) $D_1 \times \{e\}$,

(b) $D_1' \times \{x_i\}$,

(c) The subset of $P \times \{e\}$ enclosed by D_1, x_i, $B_\varepsilon(s(x_i) \times \{e\})$ and $B_\varepsilon(t(x_i) \times \{e\})$. If one of these balls does not touch D_1, we instead use the line which is closest to this ball, orthogonal to x_i and which meets D_1 at distance exactly ε from x_i.

(d) The subset of $P \times \{x_i\}$ that we get by reflecting the set (c) along $x_i \times \{e\}$. From this set, we take only those points that are closer (or equal) to x_i than to any other edge.

(e) $\{x \in P \times \{x_i\} \mid \varepsilon \leq d(x, x_i \times \{e\}) \leq 3\varepsilon \text{ and } d(x, x_i) \leq d(x, x_j) \text{ for all } j \neq i\}$.

Then each of the sets (a)-(e) is compact and simply connected. Because the sets (c)-(e) are adjacent to exactly 2 other of these sets and the sets (a) and (b) are adjacent to exactly one other set (see Figure 21) the union K_i is also simply connected and compact. We then take

$$U_i \coloneqq B_{2\varepsilon}(o_P) \times \{x_i\}$$

which is an open subset of K_i.

We now define the set

$$V \coloneqq \mathcal{M}_{\hookrightarrow}(D_1) \cap \mathcal{M}_{\rightsquigarrow}(D_2) \cap \bigcap_{i=-n}^{2n+1} \mathcal{M}_+(K_i, U_i) \cap \widehat{\text{Fi-Ch}_n}.$$

We now want to show that $V = \widehat{B_\varepsilon(P)}$.

Part 2: First, we show that $\widehat{B_\varepsilon(P)} \subseteq V$. For this, let $P^{(1)}$ be a PSI-polygon such that $\widehat{P}^{(1)} \in \widehat{B_\varepsilon(P)}$. Then the embedding of D_1 lies in the interior of $P^{(1)}$, therefore $\widehat{P}^{(1)} \in \mathcal{M}_{\hookrightarrow}(D_1)$. We already showed $\widehat{P}^{(1)} \in \mathcal{M}_{\rightsquigarrow}(D_2)$ at the time we defined it.

At last, we want to show that K_i can be immersed in $\widehat{P}^{(1)}$. For that, we see that by moving each of the endpoints of x_i at maximum ε, the point $m(x_i)$ can also not be moved more than ε. Because $\text{dev}(o_P, x_i) = 2m(x_i)$, it holds that

$$d(\text{dev}(o_P, x_i), \text{dev}(o_{P^{(1)}}, x_i)) < 2\varepsilon.$$

Therefore part (b) of K_i can be immersed in $\widehat{P}^{(1)}$ and $\widehat{o}_{P^{(1)}}$ lies in the image of U_i under this immersion. That is, if no other edge cuts through the parts (c)-(e). This holds, because each edge except x_i has distance more than 8ε to $m(x_i)$.
Let us look at this argument in more detail: Let $n(x_i)$ denote the point in P, which is exactly ε away from x_i and $m(x_i)$. Because ε is small enough, $n(x_i) \in D_1$ and because $m(x_i)$ has distance greater than 8ε from each other edge the point $n(x_i)$ also has distance greater than 7ε from each other edge. Therefore each other edge of $P^{(1)}$ has distance more than 6ε from $n(x_i)$ ($n(x_i) \in D_1$ lies in the interior of $P^{(1)}$). The same holds for the

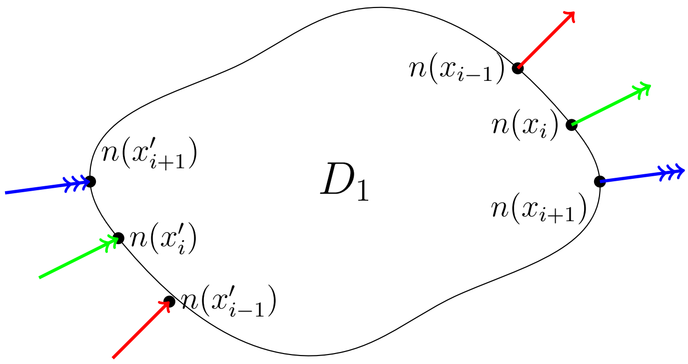

Figure 22: The set $\mathcal{M}_+(K_i, U_i)$ leads to a connection from $n(x_i) \in D_1$ to $n(x'_i) = n(x_i)' \in D_1$ via a straight line segment. This connection will be represented by arrows.

reflected point $n(x'_i) = (n(x_i))'$. But $d(n(x'_i), n((x_i^{(1)})')) < \varepsilon$. This shows that $n((x_i^{(1)})')$ lies either in the (d) or (e) part of K_i[7] and each other edge of $P^{(1)}$ has distance more than 5ε from this point.

Part 3: We will now show that $V \subseteq \widehat{B_\varepsilon(P)}$. So let $P^{(1)} \in$ Fi-Ch$_n$ such that $\widehat{P^{(1)}} \in V$. Because $P^{(1)} \in \mathcal{M}_{\hookrightarrow}(D_1)$ we can embed D_1 in $\widehat{P}^{(1)}$. This does not mean, that no edge of $\widehat{P}^{(1)}$ intersects the image of D_1. This behavior can be seen in Example 5.13, where the interior of $P^{(1)}$ can be embedded in $\widehat{P}^{(2)}$, but crosses two edges. But for the moment, we will mostly depict it as though no edges intersect the image of D_1, but the following arguments still hold if this is not true. Because $D_1 \hookrightarrow \widehat{P}^{(1)}$ we can see D_1 as a subset of $\widehat{P}^{(1)}$.

Let us first take a look at one of the K_i. There has to be an immersion $\iota_i \colon K_i \to \widehat{P}^{(1)}$ with $o_{P^{(1)}} \in \iota_i(U_i)$. Therefore $\iota_i(D'_1 \times \{x_i\})$ is a slightly smaller and slightly shifted (at maximum 2ε) copy of $\iota_i(D_1 \times \{e\}) = D_1$. Therefore $\iota_i(D'_1 \times \{x_i\}) \subseteq D_1$ and $\iota_i|_{D'_1 \times \{x_i\}}$ is injective. We define $H_i \subseteq K_i$ to be the subset corresponding to the (c)-(e) part of the definition of K_i. We now take a look at the point $n(x'_i) \in D_1$. It follows, that $n(x'_i) \in \iota_i(H_i)$ because a 2ε-ball around $(n(x'_i), x_i)$ lies completely in H_i. In addition, there is a line segment, which lies completely in $\iota_i(H_i)$ and which connects $n(x_i)$ to $n(x'_i)$. This is true for each such i and we will represent this connection by arrows. We can repeat this for each such i to get arrows from D_1 to D_1 like they are shown in Figure 22.

To summarize, each K_i yields an arrow from $n(x_i) \in D_1$ to $n(x'_i) \in D_1$. These arrows actually represent a line, lying in K_i, from $(n(x_i), e)$ to a point $a \in B_{2\varepsilon}(n(x'_i), x_i)$ with $\iota_i(a) = n(x'_i)$. We will say, that such an arrow corresponds to the edge x_i, if it was constructed out of the set K_i.

[7] Or, to be more precise, in the image of the (d) or (e) parts under the map which immerses $K_i \backslash D'_1 \times \{x_i\}$ into $\widehat{P}^{(1)}$. The existence of such an immersion follows (similar as for the set D_2) from the fact that each vertex, which is not an endpoint of x_i, is sufficiently far away from x_i.

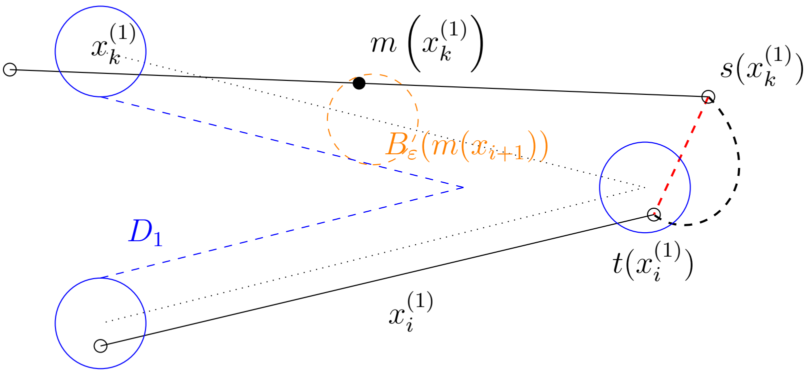

Figure 23: One endpoint of the edge $x_k^{(1)}$, responsible for the next arrow, has to coincide with one endpoint of $x_i^{(1)}$ or else there will be edges (represented by the dashed black line) beyond the red dashed line. These edges cannot be responsible for an arrow, which is a contradiction.

In $\widehat{P}^{(1)}$ this line segment goes through a finite number of edges, which can be written as a word $x_{n_0}^{(1)} \dots x_{n_m}^{(1)}$. Because of the symmetry, the line segment, which connects $n(x_i')$ to $n(x_i)$ has to go through the mirrored versions of these edges. Therefore $x_{n_i}^{(1)} = (x_{n_{m-i}}^{(1)})'$. But when we look at this line segment embedded in $P^{(1)}$ as mentioned in Remark 4.10, going through an edge $x_{n_i}^{(1)}$ corresponds to subtracting $2(m(x_{n_i}^{(1)}) - o_{P^{(1)}})$. This shows, that m is even and only the edge $x_{n_{(\frac{m}{2})}}$ is responsible for the corresponding arrow, i.e. for the shift (relative to D_1) which is noted by this arrow. In addition, if an edge of $P^{(1)}$ is responsible for an arrow, the midpoint of this edge has to be at maximum ε away from $m(x_i)$ with x_i being the edge, this arrow corresponds to. This also shows, that each edge can be responsible for at maximum one arrow.

We now look at the arrows corresponding to the edges of the converging sides. Because there are $n+1$ of them, at least one of the edges of the converging sides of $P^{(1)}$ has to be responsible for one of these arrows. But then, because of the structure of these converging sides, the word consists of only one letter. In addition, it has to have the same index as the edge, this arrow corresponds to or else the other converging side of $P^{(1)}$ would cut through D_1 or be too far away from it. Therefore each of the arrows corresponding to the converging sides is realized by only one edge and it always has the right index. This also shows that each other edge has to be responsible for exactly one arrow.

The rest of the claim will be shown by some form of induction. Starting from the converging sides, we assume that up to a specific arrow (corresponding to x_i), for each of these arrows (corresponding to x_j) the corresponding word consists of only the letter $x_j^{(1)}$. In addition, it should hold that

$$d\left(s(x_j), s(x_j^{(1)})\right) < \varepsilon \quad \text{and} \quad d\left(t(x_j), t(x_j^{(1)})\right) < \varepsilon.$$

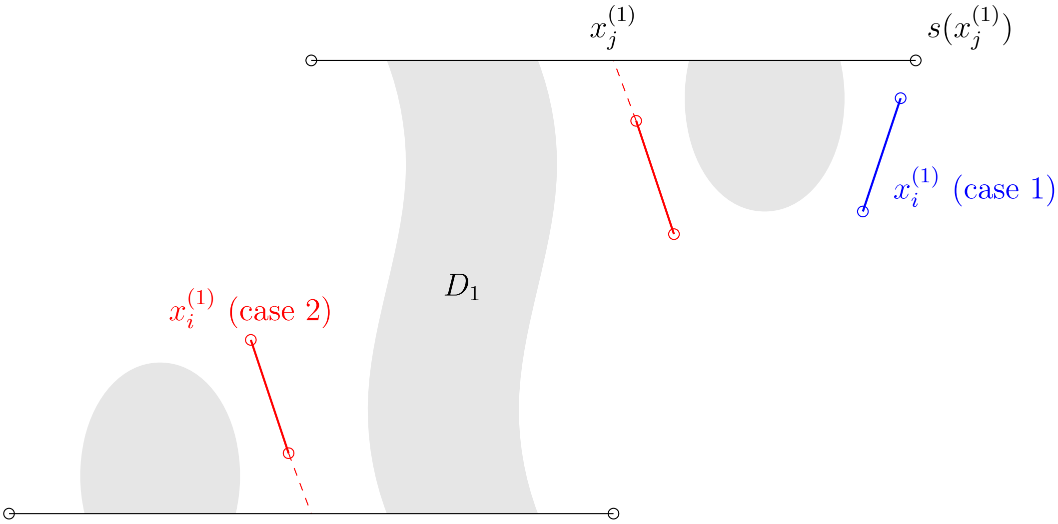

Figure 24: The situation, if D_1 goes through one of the edges of $P^{(1)}$ exactly twice. In the first case $x_i^{(1)}$ is of the form of the blue line, in the second case of the red line. Note, that in case 1 the upper end of the blue edge is actually the same as $s(x_j^{(1)})$.

The following situation is also depicted in Figure 23: We look at the next arrow (corresponding to x_{i+1}). Then let $x_k^{(1)}$ be the edge of $P^{(1)}$ responsible for this arrow. This means that $d(m(x_k^{(1)}), m(x_{i+1})) < \varepsilon$. Let us now, for a moment, return to the issue of D_1 going through edges of $P^{(1)}$. When we look at D_1 as a subset of $P^{(1)}$ this cuts the set into several components.

Let us first assume, that the arrows corresponding to x_i and x_{i+1} completely lie in the same component [8]. Then one of the endpoints of $x_k^{(1)}$ (w.l.o.g. $s(x_k^{(1)})$) has to be the same as the endpoint of $x_i^{(1)}$ (w.l.o.g. $t(x_i^{(1)})$) or else there have to be some edges of $P^{(1)}$ connecting these two vertices (marked by the black dashed line in Figure 23). But the area bound by these edges can only be reached from D_1 through a region lying between these two arrows (marked by the red dashed line in Figure 23). But this means, that these edges cannot be responsible for any arrow, which would result in more than n finite edges in $P^{(1)}$.

Together with the condition for the midpoints it follows that $d(t(x_k^{(1)}), t(x_{i+1})) < 3\varepsilon$. But in fact, this distance cannot be greater than ε or else D_2 cannot be immersed in $\widehat{P}^{(1)}$. This is true because either H_2 or H_3 (from the definition of D_2) only goes through edges, that are fixed by the inductive condition and each of these edges is far enough from any other edge (except the adjacent edges of course). This shows the induction for this case.

[8] By completely we mean, that the whole area of D_1 that is adjacent to H_i and H_{i+1} lies in the same component.

Part 4: At last, we show that this inductive argument can even be continued if D_1 goes through one of the edges $x_j^{(1)}$ of $P^{(1)}$. Then, this edge also has to be responsible for an arrow and again $m(x_j^{(1)})$ has to be at maximum ε away from the midpoint of the edge of P corresponding to that arrow. Let us first assume that D_1 goes through that edge exactly twice. Then the situation is as depicted in Figure 24 and there are two possibilities.

Case 1: $x_j^{(1)}$ is not responsible for the next arrow, i.e. the arrow corresponding to x_{i+1}. This is the case, if $x_i^{(1)}$ is the blue line in Figure 24. For a similar reason as before, one endpoint of $x_j^{(1)}$ has to be the same as one endpoint of $x_i^{(1)}$ (compare Figure 23). So $s(x_j^{(1)})$ and $t(x_i^{(1)})$ are actually the same. But this just means, that the inductive argument can be continued on the other side of $x_j^{(1)}$.

Case 2: $x_j^{(1)}$ is responsible for the next arrow. This is the case, if $x_i^{(1)}$ is the red line in Figure 24. But then the inductive argument can also be continued, because the next vertex[9] is relative to D_1 at the correct position. To see this, we can look at one of the red dashed lines. We see that the position, relative to D_1 of the next vertex is 2 times this line from the previous vertex. Together with the fact that $d(m(x_j^{(1)}), m(x_{i+1})) < \varepsilon$ we get that the distance of this next vertex to $t(x_{i+1})$ is smaller than 3ε. Again with D_2 we get, that this distance is smaller than ε.

So this inductive argument tells us only the position of these vertices relative to D_1 embedded in $P^{(1)}$. Neither does it tell us their real position nor the connections between these vertices. But before we tackle this problem, we take a look at the case, where D_1 goes through that edge more than twice and show that nothing of relevance can happen there.

For each area in $P^{(1)}$ between two such connected components, that does not contain $m(x_j^{(1)})$, there has to be a vertex near $x_j^{(1)}$ or else the next arrow would also be realized by $x_j^{(1)}$. We can assume that starting from such an area, the next arrow starts after going through $x_j^{(1)}$, or else the fact that D_1 goes through $x_j^{(1)}$ one more time has no relevance for the above argumentation (Figure 25 depicts this situation).

But then, on the other side of $x_j^{(1)}$, there also has to be a vertex, so the situation is as depicted in Figure 26. However, the blue edge has to be connected to the blue point by a series of edges not already present in Figure 26 (see Figure 27). But then the part of D_1, lying between the spikes in the middle cannot be connected to the rest of D_1. This argument is shown in Figure 28 for the case, that there is an odd number of these red spikes. The blue parts are only connected to each other, so D_1 cannot leave the area between them. But this means, that D_1 can go through $x_j^{(1)}$ at maximum three times.

[9]That means the vertex that lies between the arrow for which $x_j^{(1)}$ is responsible for and the next arrow alongside D_1.

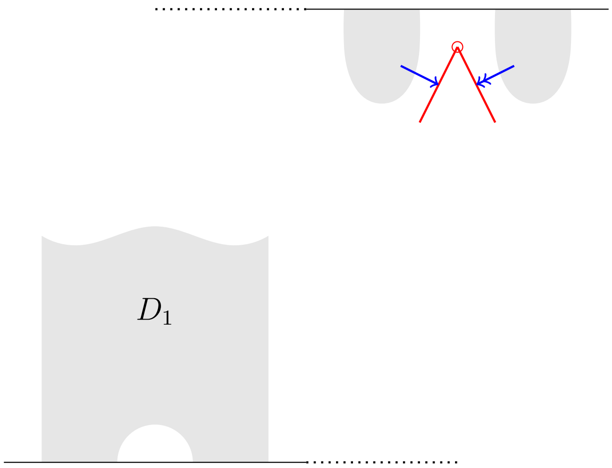

Figure 25: The situation, if there is no other arrow on the other side of $x_j^{(1)}$. In this case, there can only be a small hole on the other side and we can count this as only one crossing.

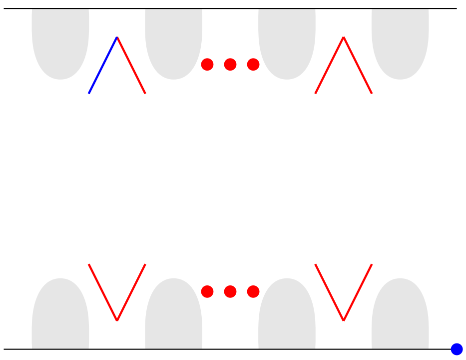

Figure 26: Between each two parts of D_1 there has to be a vertex relatively close to $x_i^{(1)}$, except for the part close to $m(x_i^{(1)})$. Because each vertex is the endpoint of exactly two edges we get these displayed spikes.

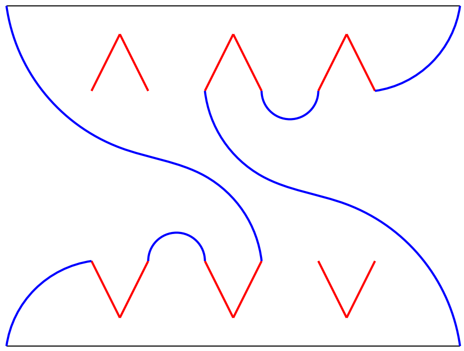

Figure 27: The blue edge and blue vertex in Figure 26 have to be connected by a series of edges not already present in Figure 26. If this would not be the case, the situation will always be similar to the one displayed here. But then there is a spike in the interior of $P^{(1)}$. Note that the blue lines only describe that there are some edges joining these vertices, not how many there are.

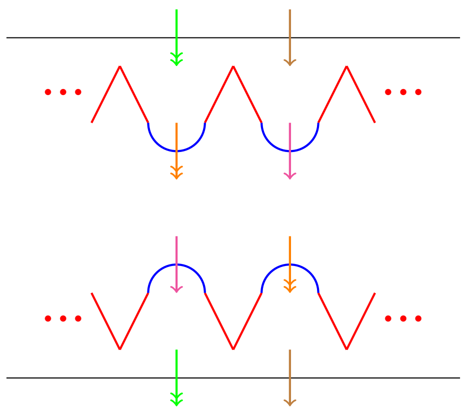

Figure 28: If there are more than 3 crossings, some part of the interior cannot be connected to the rest. Displayed here is the case for an odd number of spikes. For an even number, the same holds, but the disconnected area lies between only two spikes (instead of three as displayed here).

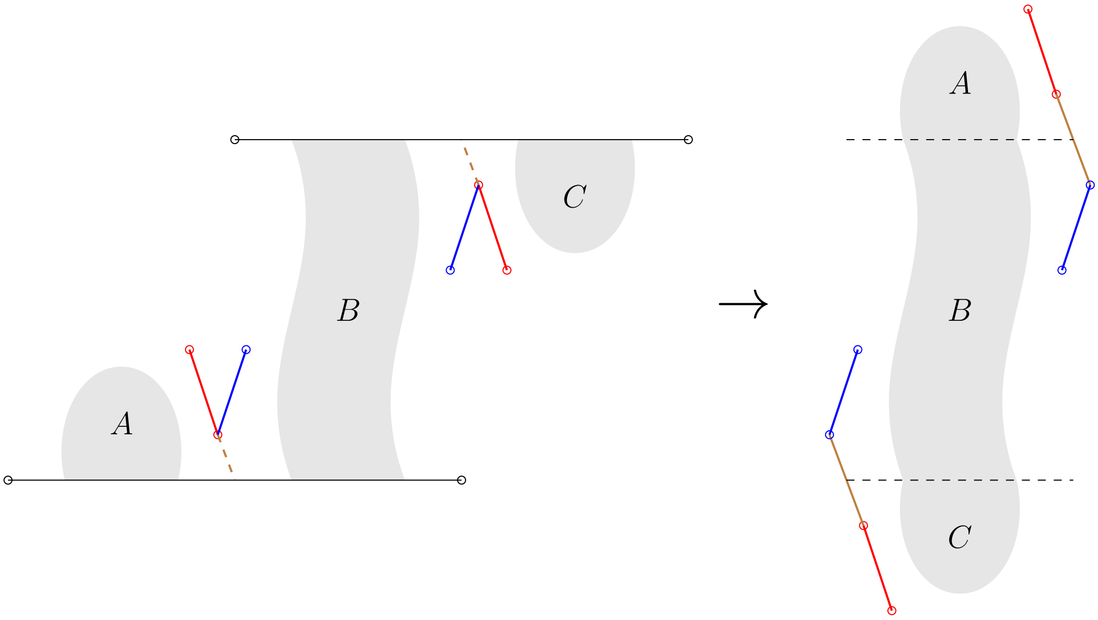

Figure 29: The isomorphism from $\widehat{P}^{(1)}$ to another $\widehat{P}^{(2)}$ for which D_1 goes through none of the edges $x_i^{(2)}$. This transformation replaces the black edge with the brown one. So the number of edges stays the same, but the order of these edges changes.

Because the point of symmetry has to be in the interior of $P^{(1)}$, it follows that D_1 cannot go through $x_j^{(1)}$ exactly three times.

At last, we show, that there is another PSI-polygon $P^{(2)} \in \text{Fi-Ch}_n$, such that D_1 lies in the interior of $P^{(2)}$ and such that $\widehat{P}^{(1)}$ and $\widehat{P}^{(2)}$ are isomorphic as translation surfaces. We construct $P^{(2)}$ by taking D_1 and arranging vertices around it at the position we know from $P^{(1)}$. Then each of these vertices is at maximum ε away from one vertex of P and we connect these vertices of $P^{(2)}$ the same way they are connected around P. Then we get $P^{(2)} \in B_\varepsilon(P)$. It remains to show, that the corresponding translation surfaces are isomorphic. The isomorphism is clear on $D_1 \subseteq P^{(1)}$ and the parts of $P^{(1)}$ between D_1 and $\partial P^{(1)}$, that belong to one of the arrows, for which D_1 does not go through the corresponding edge. For the other edges, the isomorphism is sketched in Figure 29. This concludes the proof. $\square$

We can now reap what we sowed and show that $\widehat{\text{Fi-Ch}_n}$ is a complex manifold.

Theorem 5.16. *Let $P \in \text{Fi-Ch}_n$ and $\varepsilon \in \mathbb{R}_+$ be as above. Then the map*

$$\Psi\colon \widehat{B_\varepsilon(P)} \to \mathbb{C}^n,\ \widehat{P}^{(1)} \mapsto \begin{pmatrix} s(x_1^{(1)}) \\ \vdots \\ s(x_n^{(1)}) \end{pmatrix}$$

is a homeomorphism onto its open image and therefore $\widehat{\text{Fi-Ch}_n}$ *is an* n*-dimensional complex manifold.*

Proof.

- well-defined: We have to check, that there are no two PSI-polygons in $B_\varepsilon(P)$, such that the translation surfaces arising from these polygons are isomorphic. But the set D_1 as defined above lies in the interior of each such polygon. So let $P^{(1)}$ and $P^{(2)}$ be two different PSI-polygons in $B_\varepsilon(P)$. Then there is $i \in \{1, \dots, n\}$ such that $s(x_i^{(1)}) \neq s(x_i^{(2)})$. W.l.o.g let $d(D_1, s(x_i^{(1)})) \leq d(D_1, s(x_i^{(2)}))$. But then we can fix a line segment in $\mathbb{C}$ going from the point in D_1 closest to $s(x_i^{(1)})$ to $s(x_i^{(1)})$. Again for similar reasons as in the last proof, D_1 together with this line segment can be immersed in $\widehat{P}^{(2)}$. But these sets cannot be immersed in $\widehat{P}^{(1)}$ which shows that these translation surfaces are not isomorphic.
- injective: As for Re-Ch$_\delta$ a PSI-polygon is uniquely given by the above coordinates so the map is injective.
- continuous: Let $P^{(i)}, P^{(\infty)} \in B_\varepsilon(P)$ be PSI-polygons, such that $(\widehat{P}^{(i)})_{i\in\mathbb{N}}$ converges to $\widehat{P}^{(\infty)}$. If we assume that $(\Psi(\widehat{P}^{(i)}))_{i\in\mathbb{N}}$ does not converge to $\Psi(\widehat{P}^{(\infty)})$ there is $j \in \{1, \dots, n\}$ such that $(s(x_j^{(i)}))_{i\in\mathbb{N}}$ does not converge to $s(x_j^{(\infty)})$. But then there is $\xi \in \mathbb{R}_+$ such that $s(x_j^{(i)}) \notin B_\xi(s(x_j^{(\infty)}))$ for infinitely many $i \in \mathbb{N}$. Of course $\xi < \varepsilon$. Similar to before, we can create a set D_2 but for the polygon $P^{(\infty)}$ instead of P and with H_2 and H_3 ranging between ξ and ε instead of ε and 2ε. As before, D_2 can be immersed in $\widehat{P}^{(\infty)}$, but not in infinitely many $\widehat{P}^{(i)}$, which is a contradiction.
- open: Again this follows directly from Theorem 4.18.
- image open: The image is exactly $\bigtimes_{i=1}^{n} B_\varepsilon(s(x_i))$ which is open in $\mathbb{C}^n$.

$\square$

Remark 5.17. *Similar to Theorem 5.4 and Remark 5.10, we can combine this family with the previous families. For some of these combinations, it might be necessary to fix a path which lies in the interior of each such polygon (compare Section 6.5).*

6 Strong immersive topology

6.1 Motivation

In this last section, we ultimately want to construct an infinite-dimensional complex family in Section 6.5. To do this, we first construct a topology which is slightly stronger than the immersive topology and which we will call the strong immersive topology. In Section 6.3 and Section 6.4 we will then explore some properties of this topology.

Before we start to define the topology we want to give a motivation for why a stronger topology on $\mathcal{M}$ might be more suitable in some cases. For this, we give a few examples of sequences in $\mathcal{M}$ which converge in the immersive topology.

Example 6.1. Let P again be the PSI-polygon from Figure 2 which we used to construct the Chamanara surface. Let $P^{(n)}$ be the same PSI-polygon but with $t(x_n)$ shifted to the point $(\frac{1}{2}, \frac{4}{3})$ and $t(x_n')$ accordingly. See Figure 30 for the first three PSI-polygons in this sequence. Of course, the interior of each $P^{(n)} \subseteq \mathbb{C}$ is also a translation surface as an open subset of $\mathbb{C}$. This sequence then converges to the unit square $(0,1)^2$ in the immersive topology. This can be easily seen by applying Theorem 3.12: Each subset of $(0,1)^2$ can be immersed in each of the larger sets $(P^{(n)})^\circ$. For the second part of Theorem 3.12 let $Q \in \widetilde{\mathcal{M}}$ be a simply connected translation surface which can be embedded in infinitely many $P^{(n)}$. Then the embedding cannot contain a part which is a subset of one of the additional "spikes", i.e. the parts of $P^{(n)} \setminus (0,1)^2$. This holds because the set $x_n \setminus \{s(x_n), t(x_n)\} \subseteq (0,1) \times \{1\}$ is only contained in the interior of $P^{(n)}$ and $P^{(n-1)}$ and not in the interior of any other $P^{(m)}$. Therefore the image of this embedding of Q is a subset of $(0,1)^2$ and can therefore also be embedded in this limit surface.

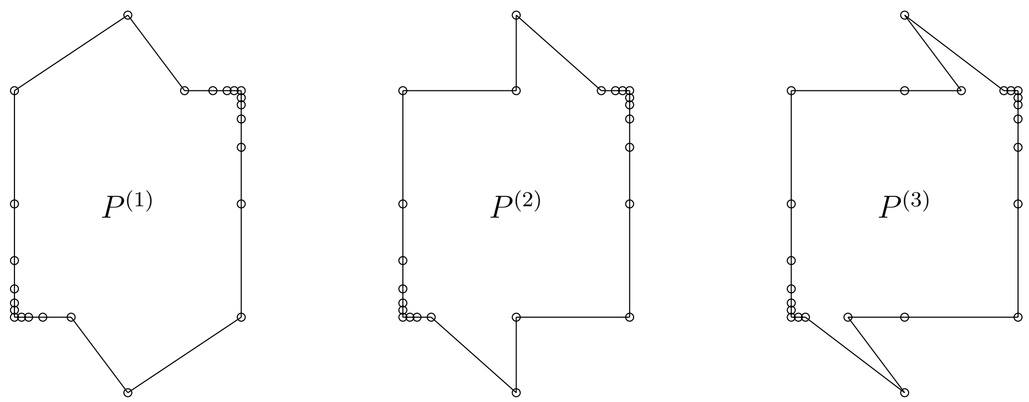

Figure 30: The first 3 elements of the sequence $(P^{(n)})_{n\in\mathbb{N}}$ converging to $(0,1)^2$. This sequence converges because for each Q, which can be embedded in infinitely many $P^{(n)}$, the image of the embedding is a subset of $(0,1)^2$.

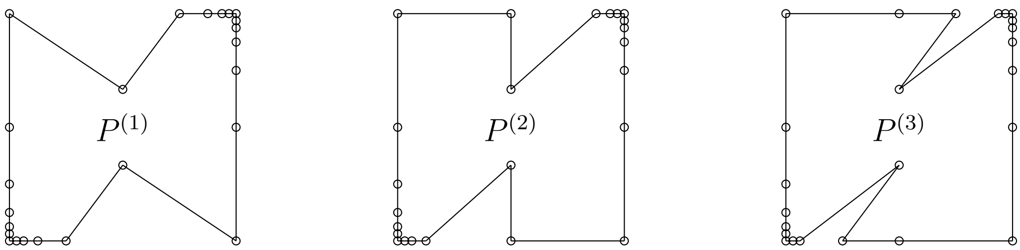

Figure 31: In contrary to Example 6.1, the sequence $(P^{(n)})_{n\in\mathbb{N}}$ of which the first 3 elements are depicted here does not converge to P from Figure 2. This is true because each compact $D \subseteq P$ containing $(\frac{1}{2}, \frac{2}{3})$ cannot be immersed in any $P^{(n)}$.

This is a general behavior of the immersive topology. A graphical intuition of this is that the topology cannot observe something being added to a translation surface, if the connection to this added part is arbitrarily small. But the topology is sensitive if something is deleted instead. For example, if the vertices in Example 6.1 are shifted to the point $(\frac{1}{2}, \frac{2}{3})$ instead of $(\frac{1}{2}, \frac{4}{3})$ (see Figure 31) the sequence does not converge anymore, because $D = [\frac{1}{10}, \frac{9}{10}]^2$ can be immersed in $(0,1)^2$ but not in any of those $P^{(n)}$. The general reasoning here is that this sequence instead converges to the polygon in Figure 32a where the lines from $(0,0)$ to $(\frac{1}{2}, \frac{1}{3})$ and from $(\frac{1}{2}, \frac{2}{3})$ to $(1,1)$ are removed from the unit square. For the sequence from Example 6.1 the intuitive limit of this sequence would be the one depicted in Figure 32b, where the line from $(\frac{1}{2}, -\frac{1}{3})$ to $(0,0)$ and the line from $(1,1)$ to $(\frac{1}{2}, \frac{4}{3})$ is added to the unit square. But this set is not a translation surface anymore. In the immersive topology this sequence then converges to the next similar set which is a translation surface. This describes the main difference we want to achieve by introducing a stronger topology, namely that sequences of this type do not converge at all.

This example can also be varied in different ways. One variation can be achieved by exchanging the upper half of the spike in Example 6.1 by a copy of the open upper half-plane. This leads to a sequence of translation surfaces of infinite area converging to a translation surface of area 1.

The other variation can be achieved by looking at the corresponding translation surfaces $\widehat{P}^{(n)}$ and $\widehat{P}$ constructed out of the PSI-polygons introduced in Example 6.1 instead of the PSI-polygons themselves. Similarly to Theorem 4.18, we see that this sequence also converges. The reasoning here is very similar to the one in Example 6.1: Each subset $D \subseteq \widetilde{P}$ homeomorphic to a closed disc only goes through a finite number of edges and for almost all $P^{(n)}$ these edges are the same as in P. The same follows for each path γ in any Q which can be embedded in infinitely many $\widetilde{P}^{(n)}$ (compare the proof of Theorem 4.18). But when we try to assign complex coordinates in $c_o(\mathbb{C})^2$ to these polygons in a similar way as we did in Theorem 5.9 for real coordinates, the image of this sequence does not converge anymore. So this coordinate function would not be continuous. In Lemma 6.22 we will see that (with the help of some additional restrictions on the PSI-polygons) this problem will be fixed by the stronger topology.

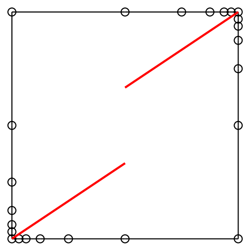

(a) The limit of the sequence which is depicted in Figure 31. The red lines depict points removed from the square.

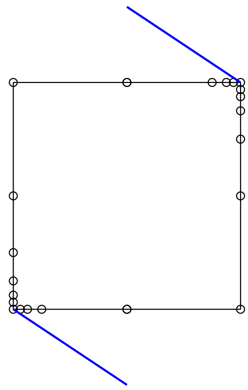

(b) An intuitive limit of the sequence which is depicted in Figure 30. The blue lines depict points added to the square.

Figure 32: A depiction of the limits of the two previous sequences. The sequence depicted in Figure 31 does not converge to the square because it converges to the surface shown on the left. The picture on the right shows an intuitive limit of the sequence which is depicted in Figure 30. But this set is not a translation surface.

Before we define the topology we want to recall Example 5.6. There we have seen that there are sequences of finite translation surfaces in $\mathcal{M}$ of genus g, converging to a finite translation surface of genus $g-1$ in the immersive topology. This is also an indication that a slightly stronger topology is desirable in which this behavior is not possible anymore. We will not investigate this phenomenon that much, because most of the translation surfaces which we consider have infinite genus. But the general idea of the strong immersive topology can be derived from this example which we will examine further in the next section.

6.2 Definition

The basic idea of the strong immersive topology comes from the following observation: Let us again take a look at the sequence $(\widehat{P}_2^{(m)})_{m\in\mathbb{N}}$ as defined in Example 5.6. This converges to $\widehat{P}_2$ in the immersive topology. As mentioned the genus of $\widehat{P}_2^{(m)}$ is 2 for each $m \in \mathbb{N}$ and the genus of $\widehat{P}_2$ is 1. But we can observe that these translation surfaces are very different metric spaces, although they are very similar from the viewpoint of immersions. To see this let $\varepsilon \in \mathbb{R}_+$ be sufficiently small (e.g. smaller than $\frac{1}{100}$). We can then look at the points $a = (\frac{1}{2}, 1-\varepsilon)$ and $b = (\frac{1}{2}, \varepsilon)$. The distance of these points is 2ε in $\widehat{P}_2^{(m)}$ because they are connected via the additional edge x_{ins}. But the distance in $\widehat{P}_2$ is much larger because $(\frac{1}{2}, 1)$ and $(\frac{1}{2}, 0)$ correspond to two different singularities of $\widehat{P}^{(2)}$.

So the distance of a and b is slightly larger than the distance of these singularities which is $\frac{1}{2}$. This is independent of m. The main idea will be to also enforce convergence as metric spaces. One topology often used for sets of metric spaces is the Gromov-Hausdorff topology which is based on the Gromov-Hausdorff metric[10]. But this metric is mostly used for compact (or at least for locally compact) metric spaces. Instead, we will use a slight variation of the topology induced by this metric and additionally adapt it in such a way, that it is compatible with the translation structure. To do this the following definitions will come to help.

Remark 6.2. *Let X be a translation surface and $\gamma\colon [0,1] \to X$ a path. Then there is a natural corresponding path $\overset{\bullet}{\gamma}\colon [0,1] \to \mathbb{C}$ with $\overset{\bullet}{\gamma}(0) = 0$. This path can be obtained in different ways. One way is by analytic continuation. Another one is by looking at the path $\operatorname{dev}_{\widetilde{X}}(\widetilde{\gamma})$ where $\widetilde{\gamma}$ is a lift of γ and translating this path such that it starts at 0. If $\operatorname{Im}(\gamma)$ is contained in a closed simply-connected set D there is also a third way, namely by taking $\gamma(0)$ as the basepoint of X and defining $\overset{\bullet}{\gamma}$ as the composition of γ with the unique immersion of D into $\mathbb{C}$.*
In the case of a geodesic path, this construction agrees with the definition of the holonomy vector.

Definition 6.3. Let X and Y be translation surfaces, $\gamma_1\colon [0,1] \to X$ and $\gamma_2\colon [0,1] \to Y$ be paths in X and Y and $\varepsilon \in \mathbb{R}_+$. Then γ_1 and γ_2 will be called *ε-similar* if

$$d\left(\overset{\bullet}{\gamma}_1(a), \overset{\bullet}{\gamma}_2(a)\right) < \varepsilon \quad \forall a \in [0,1]$$

and if

$$|l(\gamma_1) - l(\gamma_2)| < \varepsilon.$$

For the Gromov-Hausdorff metric ε-approximations (sometimes called ε-isometries as in [BBI01]) are often used. For this metric, the existence of such an ε-approximation between two metric spaces implies that the Gromov-Hausdorff distance of these spaces is smaller than 2ε. We will take a similar approach but define a more strict form of ε-approximations, which is compatible with the translation structure.

Definition 6.4. Let $(X, o_X), (Y, o_Y)$ be translation surfaces and $\varepsilon \in \mathbb{R}_+$. A map $\varphi\colon X \to Y$ is called *ε-translation approximation* or often just *ε-approximation* if all of the following conditions hold:

(a) $\varphi(o_X) = o_Y$.

(b) $\operatorname{Im}(\varphi)$ is ε-dense in Y, i.e. for each $y \in Y$ there is an $x \in X$ with $d(\varphi(x), y) < \varepsilon$.

(c) For each $x_1, x_2 \in X$ and each shortest path $\gamma_1\colon [0,1] \to X$ from x_1 to x_2 there is a path $\gamma_2\colon [0,1] \to Y$ from $\varphi(x_1)$ to $\varphi(x_2)$ which is ε-similar to γ_1.

[10]For a definition and further information see e.g. [BBI01].

(d) For each $x_1, x_2 \in X$ and each shortest path $\gamma_2\colon [0,1] \to Y$ from $\varphi(x_1)$ to $\varphi(x_2)$ there is a path $\gamma_1\colon [0,1] \to X$ from x_1 to x_2 which is ε-similar to γ_2.

Because the conditions (c) and (d) hold for all shortest paths, they imply that for each $x_1, x_2 \in X$ it holds that

$$|d(x_1, x_2) - d(\varphi(x_1), \varphi(x_2))| < \varepsilon.$$

If X and Y are translation surfaces and $\Phi\colon X \to Y$ is a translation isomorphism then Φ is also an ε-approximation for each $\varepsilon \in \mathbb{R}_+$. In addition, if Z is another translation surface and $\varphi\colon Y \to Z$ is an ε-approximation for some $\varepsilon \in \mathbb{R}_+$ then $\varphi \circ \Phi\colon X \to Z$ is an ε-approximation as well. Therefore for two isomorphism classes $X, Y \in \mathcal{M}$ the existence of an ε-approximation is independent of a chosen representative. As before we will mostly make no difference between a translation surface and its isomorphism class and just identify $X, Y \in \mathcal{M}$ with one of its representatives.

At last, this allows us to define the strong immersive topology:

Definition 6.5. The *strong immersive topology* on $\mathcal{M}$ is the coarsest topology on $\mathcal{M}$, which is finer than the immersive topology and such that all sets of the form

$$\Delta_\varepsilon(X) \coloneqq \{Y \in \mathcal{M} \mid \exists \xi\text{-approximation } \varphi\colon X \to Y \text{ for a } \xi \in \mathbb{R}_+, \xi < \varepsilon\}$$

for $X \in \mathcal{M}$ and $\varepsilon \in \mathbb{R}_+$ are open.

6.3 Some technical statements

Before we can start examining and using the strong immersive topology we have to prove some rather technical prerequisites about ε-similar paths and ε-approximations. This section is dedicated to this task. The statements we prove in this section will all be used at later stages of Section 6.

A reader not interested in those technical details may as well skip to Section 6.4 and look up those statements when they are needed.

At first, we show that being ε-similar behaves well under concatenation of paths.

Lemma 6.6. *Let $X^{(1)}$ be a translation surface and $\gamma_1^{(1)}, \gamma_2^{(1)}\colon [0,1] \to X^{(1)}$ be two paths such that $\gamma_1^{(1)}(1) = \gamma_2^{(1)}(0)$. In addition, let $X^{(2)}$ be another translation surface and $\gamma_1^{(2)}, \gamma_2^{(2)}\colon [0,1] \to X^{(2)}$ be two paths such that $\gamma_1^{(2)}(1) = \gamma_2^{(2)}(0)$.*

If there are $\varepsilon_1, \varepsilon_2 \in \mathbb{R}_+$ such that $\gamma_1^{(1)}$ and $\gamma_1^{(2)}$ are ε_1-similar and such that $\gamma_2^{(1)}$ and $\gamma_2^{(2)}$ are ε_2-similar then $\gamma_1^{(1)} * \gamma_2^{(1)}$ and $\gamma_1^{(2)} * \gamma_2^{(2)}$ are $(\varepsilon_1 + \varepsilon_2)$-similar, where $*$ denotes the concatenation of paths.

Proof. Because $\gamma_1^{(1)}$ and $\gamma_1^{(2)}$ are ε_1-similar it holds that

$$d(\dot{\gamma}_1^{(1)}(1), \dot{\gamma}_1^{(2)}(1)) < \varepsilon_1.$$

Then $(\gamma_1^{(1)} \overset{\bullet}{*} \gamma_2^{(1)})$ is obtained by translating $\dot{\gamma}_2^{(1)}$ such that it starts at $\dot{\gamma}_1^{(1)}(1)$. This is done by adding the vector $\dot{\gamma}_1^{(1)}(1) \in \mathbb{C}$. The same holds for $(\gamma_1^{(2)} \overset{\bullet}{*} \gamma_2^{(2)})$.

But then for $a \in [\frac{1}{2}, 1]$

$$\begin{aligned}
&d((\gamma_1^{(1)} \overset{\bullet}{*} \gamma_2^{(1)})(a), (\gamma_1^{(2)} \overset{\bullet}{*} \gamma_2^{(2)})(a)) \\
&= d(\dot{\gamma}_2^{(1)}(2a-1) + \dot{\gamma}_1^{(1)}(1), \dot{\gamma}_2^{(2)}(2a-1) + \dot{\gamma}_1^{(2)}(1))) \\
&\leq d(\dot{\gamma}_2^{(1)}(2a-1), \dot{\gamma}_2^{(2)}(2a-1)) + d(\dot{\gamma}_1^{(1)}(1), \dot{\gamma}_1^{(2)}(1)) < \varepsilon_1 + \varepsilon_2.
\end{aligned}$$

Also

$$\begin{aligned}
&|l(\gamma_1^{(1)} * \gamma_2^{(1)}) - l(\gamma_1^{(2)} * \gamma_2^{(2)})| \\
&= |l(\gamma_1^{(1)}) + l(\gamma_2^{(1)}) - l(\gamma_1^{(2)}) - l(\gamma_2^{(2)})| \\
&\leq |l(\gamma_1^{(1)}) - l(\gamma_1^{(2)})| + |l(\gamma_2^{(1)}) - l(\gamma_2^{(2)})| < \varepsilon_1 + \varepsilon_2.
\end{aligned}$$

□

In addition the composition of two approximations is again an approximation:

Lemma 6.7. *Let $X^{(1)}, X^{(2)}, X^{(3)} \in \mathcal{M}$ be translation surfaces, $\varepsilon^{(1)}, \varepsilon^{(2)} \in \mathbb{R}_+$ and $\varphi^{(1)}\colon X^{(1)} \to X^{(2)}$ an $\varepsilon^{(1)}$-approximation and $\varphi^{(2)}\colon X^{(2)} \to X^{(3)}$ an $\varepsilon^{(2)}$-approximation. Then $\varphi^{(2)} \circ \varphi^{(1)}\colon X^{(1)} \to X^{(3)}$ is an $(\varepsilon^{(1)} + 2\varepsilon^{(2)})$-approximation.*

Proof. By definition we have $\varphi^{(2)} \circ \varphi^{(1)}(o_{X^{(1)}}) = o_{X^{(3)}}$.

If $\gamma^{(1)}\colon [0,1] \to X^{(1)}$ and $\gamma^{(2)}\colon [0,1] \to X^{(2)}$ are $\varepsilon^{(1)}$-similar and $\gamma^{(2)}\colon [0,1] \to X^{(2)}$ and $\gamma^{(3)}\colon [0,1] \to X^{(3)}$ are $\varepsilon^{(2)}$-similar, then $\gamma^{(1)}\colon [0,1] \to X^{(1)}$ and $\gamma^{(3)}\colon [0,1] \to X^{(3)}$ are $(\varepsilon^{(1)} + \varepsilon^{(2)})$-similar. This shows the conditions (c) and (d) of Definition 6.4.

Now let $x^{(3)} \in X^{(3)}$. Then there is $x^{(2)} \in X^{(2)}$ with $d(\varphi^{(2)}(x^{(2)}), x^{(3)}) < \varepsilon^{(2)}$. In addition there is $x^{(1)} \in X^{(1)}$ with $d(\varphi^{(1)}(x^{(1)}), x^{(2)}) < \varepsilon^{(1)}$. Then

$$\begin{aligned}
&d(\varphi^{(2)} \circ \varphi^{(1)}(x^{(1)}), x^{(3)}) \\
&\leq d(\varphi^{(2)}(\varphi^{(1)}(x^{(1)})), \varphi^{(2)}(x^{(2)})) + d(\varphi^{(2)}(x^{(2)}), x^{(3)}) \\
&< (\varepsilon^{(2)} + d(\varphi^{(1)}(x^{(1)}), x^{(2)})) + \varepsilon^{(2)} \\
&< 2\varepsilon^{(2)} + \varepsilon^{(1)},
\end{aligned}$$

so condition (b) holds. □

In the definition of an ε-approximation, it is only required that such an ε-similar path exists for each shortest path. This can be expanded to include non-shortest paths by subdividing such a path into short subpaths which are ε-similar to shortest paths. This is the content of the next lemma and is only considered in the context of PSI-polygons.

Lemma 6.8. *Let $P^{(1)}$ and $P^{(2)}$ be PSI-polygons, $\varphi\colon \widehat{P}^{(1)} \to \widehat{P}^{(2)}$ a ξ-approximation for some $\xi \in \mathbb{R}_+$, $\gamma^{(1)}\colon [0,1] \to (P^{(1)})^\circ$ a path of length l and $\delta \in \mathbb{R}_+$ such that $d(\partial P^{(1)}, \operatorname{Im}(\gamma^{(1)})) > \delta$. Then there is a path $\gamma^{(2)}\colon [0,1] \to \widehat{P}^{(2)}$ from $\varphi(\widehat{\gamma}^{(1)}(0))$ to $\varphi(\widehat{\gamma}^{(1)}(1))$ which is $\left(\left\lceil \frac{l}{\delta} \right\rceil \cdot \xi + \delta\right)$-similar to $\widehat{\gamma}^{(1)}$.*

Proof. Let $a_1 \in [0,1]$ be such that $\gamma^{(1)}([0,a_1])$ has length δ. Then $d(\widehat{\gamma^{(1)}(0)}, \widehat{\gamma^{(1)}(a_1)}) \leq \delta$ and the shortest path connecting these two points lies in the interior of $P^{(1)}$. We call this shortest path $\gamma^{(1)}_{\mathrm{sh},1}$. Then $\gamma^{(1)}_{\mathrm{sh},1}$ is δ-similar to $\gamma^{(1)}([0,a_1])$ and in addition $\dot{\gamma}^{(1)}_{\mathrm{sh},1}(1) = \dot{\gamma}^{(1)}(a_1)$ because they both lie in $(P^{(1)})^\circ$ which is simply connected. We now subdivide $\gamma^{(1)}$ into finitely many segments of length smaller than δ. Similar to Lemma 6.6 we get a path $\gamma^{(1)}_{\mathrm{sh}}\colon [0,1] \to \widehat{P}^{(1)}$ which is δ-similar[11] to $\gamma^{(1)}$ and which goes through all of the start- and endpoints of these segments and consecutively connects these via shortest paths.

Let $a_1 \in [0,1]$ be such that $\gamma^{(1)}(a_1) = \gamma^{(1)}_{\mathrm{sh}}(a_1)$ is the endpoint of the first segment. Then there is a path $\gamma^{(2)}_{\mathrm{sh},1}\colon [0,1] \to \widehat{P}^{(2)}$ from $\varphi(\gamma^{(1)}_{\mathrm{sh}}(0))$ to $\varphi(\gamma^{(1)}_{\mathrm{sh}}(a_1))$ which is ξ-similar to $\gamma^{(1)}_{\mathrm{sh}}|_{[0,a_1]}$.

Now let $a_2 \in [0,1]$ be such that $\gamma^{(1)}(a_2)$ is the endpoint of the second segment. Then a ξ-similar path $\gamma^{(2)}_{\mathrm{sh},2}$ can again be found. With Lemma 6.6 it follows inductively that there is a path $\gamma^{(2)}_{\mathrm{sh}}$ which is $(\xi \cdot m)$-similar to $\gamma^{(1)}_{\mathrm{sh}}$, where m is the number of these segments. But because we can choose the above subdivision in a way that each such segment (except the last one) has length δ there are $\left\lceil \frac{l}{\delta} \right\rceil$ many of them. □

Lemma 6.9. *Let $P^{(1)}$, $P^{(2)}$ be PSI-polygons, $\varepsilon \in \mathbb{R}_+$, $\gamma^{(1)}\colon [0,1] \to (P^{(1)})^\circ$ a path of length l and $\delta \in \mathbb{R}_+$, $\delta < \varepsilon$ such that $d(\partial P^{(1)}, \operatorname{Im}(\gamma^{(1)})) > \delta$. If there is a ξ-approximation $\varphi\colon \widehat{P}^{(1)} \to \widehat{P}^{(2)}$ for*

$$\xi \leq \frac{\delta\varepsilon - \delta^2}{2l}$$

then there is a path $\gamma^{(2)}\colon [0,1] \to \widehat{P}^{(2)}$ from $\varphi(\widehat{\gamma}^{(1)}(0))$ to $\varphi(\widehat{\gamma}^{(1)}(1))$ which is ε-similar to $\widehat{\gamma}^{(1)}$.

Proof. This is a direct consequence of Lemma 6.8. □

[11] $d(\dot{\gamma}^{(1)}_1(1), \dot{\gamma}^{(2)}_1(1))$ in Lemma 6.6 is actually zero in this case.

In the proof of Lemma 5.15, we have defined the set D_1 of all points of distance more than ε to ∂P. For later use, we now fix a more general definition.

Definition 6.10. Let P be a PSI-polygon and $\varepsilon \in \mathbb{R}_+$. Then

$$\mathrm{IP}_\varepsilon \coloneqq \{x \in P \mid d(x, \partial P) \geq \varepsilon\}.$$

Similar to our previous approach for a PSI-polygon $P^{(n)}$ we denote the corresponding set by $\mathrm{IP}_\varepsilon^{(n)}$.

This set is mostly used because an appropriate approximation has to map it into a slightly larger version of itself which is the content of the next lemma:

Lemma 6.11. *Let $P^{(1)}$ and $P^{(2)}$ be two PSI-polygons, $\varepsilon \in \mathbb{R}_+$ and*

$$l \coloneqq \sup\{d_{\mathrm{IP}_\varepsilon^{(1)}}(x, o_P) \mid x \in \mathrm{IP}_\varepsilon^{(1)}\},$$

where $d_{\mathrm{IP}_\varepsilon^{(1)}}$ denotes the metric of $\mathrm{IP}_\varepsilon^{(1)}$, i.e. the distance measured by the length of shortest paths lying in $\mathrm{IP}_\varepsilon^{(1)}$.

In addition let $\varphi\colon \widehat{P}^{(1)} \to \widehat{P}^{(2)}$ be an $\frac{\varepsilon}{8l}$-approximation. If $\mathrm{IP}_{2\varepsilon}^{(1)}$ is connected and $\mathrm{IP}_\varepsilon^{(1)}$ is also a subset of $(P^{(2)})^\circ$, then $\varphi(\widehat{\mathrm{IP}_{2\varepsilon}}^{(1)}) \subseteq \widehat{\mathrm{IP}_\varepsilon}^{(1)}$[12]*.*

Proof. This follows directly from Lemma 6.9 by setting $\delta = \frac{\varepsilon}{2}$ and by taking a shortest path for each $x \in \mathrm{IP}_{2\varepsilon}^{(1)}$ from $o_{P^{(1)}}$ to x which lies in $\mathrm{IP}_{2\varepsilon}^{(1)}$. $\square$

As a final observation in this section, we see that every point in P gets arbitrarily close to IP_ε if we choose ε to be sufficiently small.

Lemma 6.12. *Let P be a PSI-polygon and $\alpha \in \mathbb{R}_+$. Then there is $\beta \in \mathbb{R}_+$ such that $P \subseteq B_\alpha(\mathrm{IP}_\beta)$.*

Proof. Let us assume the opposite. Then there is a sequence $(a_n)_{n\in\mathbb{N}}$ in ∂P with $d(a_n, \mathrm{IP}_{\frac{1}{n}}) > \alpha$ for all $n \in \mathbb{N}$. Because ∂P is compact it accumulates at a point $a \in \partial P$. But then there is a point $x \in P^\circ$ with $d(x, a) < \alpha$. For this point x there is some $N \in \mathbb{N}$ such that $d(\partial P, x) > \frac{1}{N}$. This is a contradiction to $d(a_n, \mathrm{IP}_{\frac{1}{n}}) > \alpha$ for all $n \in \mathbb{N}$ because a is an accumulation point of a_n. $\square$

[12] The set $\widehat{\mathrm{IP}_\varepsilon}^{(1)}$ is the set of points in $\widehat{P}^{(1)}$ corresponding to points in $\mathrm{IP}_\varepsilon^{(1)}$. But because $\mathrm{IP}_\varepsilon^{(1)} \subseteq (P^{(2)})^\circ$, $\widehat{\mathrm{IP}_\varepsilon}^{(1)}$ can also be seen as a subset of $\widehat{P}^{(2)}$.

6.4 Properties of the strong immersive topology

Before we start using this topology we want to make sure that it behaves nicely. For this let us first return to our examples and show that the defined sequences do not converge anymore.

Example 6.13. Let $(P^{(n)})_{n\in\mathbb{N}}$ and P be the PSI-polygons defined in Example 6.1. Then (again viewed as a translation surface) $(P^{(n)})_{n\in\mathbb{N}}$ does not converge to P in the strong immersive topology[13]. This holds because a point in the middle of such a spike is more than $\frac{1}{6}$ away from each point in $(0,1)^2 \subseteq P^{(n)}$. But a small variation of Lemma 6.11 show that no point in such a spike lies in the image of an ε-approximation from P to $P^{(n)}$ (for ε sufficiently small).

This argument still holds for the sequence $(\widehat{P}^{(n)})_{n\in\mathbb{N}}$ and the translation surface $\widehat{P}$ and shows that it does not converge as well. This behavior will be examined in a more general case in Lemma 6.22.

Example 6.14. We return to the sequence $(\widehat{P}_2^{(m)})_{m\in\mathbb{N}}$ and the translation surface $\widehat{P}_2 \in \mathcal{M}$ defined in Example 5.6. As we have seen, this sequence converges to $\widehat{P}_2$ in the immersive topology. But it does not converge in the strong immersive topology. This follows again from Lemma 6.11: Recall that the points a and b introduced at the beginning of Section 6.2 have distance larger than $\frac{1}{2}$ in $\widehat{P}_2$. More precisely, for a shortest path γ connecting these two points, $\dot{\gamma}$ leaves $B_{\frac{1}{4}}(0)$. But for a sufficiently small ε, an ε-approximation has to map these points to two points which are very close to a and b. But these points then have arbitrarily small distance in $\widehat{P}_2^{(m)}$.

Again as mentioned in Section 3, to make the convergence of sequences a useful criterion a topological space has to be at least first countable. This was the case for the immersive topology which is even second countable. First countability still holds for the strong immersive topology which is shown by the following proposition:

Proposition 6.15. *The strong immersive topology on $\mathcal{M}$ is Hausdorff and first countable.*

Proof. The strong immersive topology is still Hausdorff because it is finer than the immersive topology which is Hausdorff.

For $\varepsilon \in \mathbb{R}_+$ an ε-approximation is also a δ-approximation for each $\delta \in \mathbb{R}_+$, $\delta \geq \varepsilon$. Therefore for a translation surface $X \in \mathcal{M}$ it holds that $\Delta_\varepsilon(X) \subseteq \Delta_\delta(X)$.

[13] Therefore this sequence does not converge at all, because the immersive topology is Hausdorff and therefore each sequence can have at most one limit.

Now let $\mathfrak{B}$ be a countable basis of the immersive topology. We claim that

$$B := \left\{\Delta_{\frac{1}{n}}(X) \cap U \mid n \in \mathbb{N}, U \in \mathfrak{B}\right\}$$

is a countable neighborhood basis of X. That B is countable follows directly from the fact that $\mathfrak{B}$ is countable.

To show that it is a neighborhood basis let $Y \in \mathcal{M}$ and $\varepsilon \in \mathbb{R}_+$ such that $X \in \Delta_\varepsilon(Y)$. Then there is a ξ-approximation $\varphi_1 \colon Y \to X$ with $\xi < \varepsilon$. Now let $n \in \mathbb{N}$ such that $\frac{1}{n} < \frac{\varepsilon - \xi}{2}$. In addition let $Z \in \Delta_{\frac{1}{n}}(X)$ and $\varphi_2 \colon X \to Z$ be a δ-approximation for some $\delta < \frac{1}{n}$. Then Lemma 6.7 shows that $\varphi_2 \circ \varphi_1 \colon Y \to Z$ is a $(\xi + 2\delta)$-approximation. But because

$$\xi + 2\delta < \xi + \frac{2}{n} < \xi + (\varepsilon - \xi) = \varepsilon$$

it holds that $Z \in \Delta_\varepsilon(Y)$. Therefore $\Delta_{\frac{1}{n}}(X) \subseteq \Delta_\varepsilon(Y)$ holds.

We have defined the strong immersive topology by giving a subbasis. So each open set can be written as a union of finite intersections of elements of this subbasis. This shows that each open set containing X contains an open subset of the form

$$O := \bigcap_{i=0}^{m} \Delta_{\varepsilon_i}(Y_i) \cap U$$

for some $m \in \mathbb{N}$, $\varepsilon_i \in \mathbb{R}_+$, $Y_i \in \mathcal{M}$ and $U \subseteq \mathcal{M}$ open in the immersive topology. In addition O can be chosen in such a way that $X \in O$. But then the above argument tells us that for each $\Delta_{\varepsilon_i}(Y_i)$ there is $\Delta_{\frac{1}{n}}(X)$ with $\Delta_{\frac{1}{n}}(X) \subseteq \Delta_{\varepsilon_i}(Y_i)$ so the intersection of these $\Delta_{\frac{1}{n}}(X)$ together with a suitable element of $\mathfrak{B}$ is a subset of O and lies in B. This shows that B is a countable neighborhood basis of X. $\square$

Finally, we will return to our main focus of translation surfaces constructed out of PSI-polygons. Of course, the coordinate functions introduced in Section 5 will still be continuous when we use the strong immersive topology on $\mathcal{M}$. In Section 5 we always used Theorem 4.18 to show that they are open. So in the best case, this Theorem still holds if we use the strong immersive topology instead. As it turns out, this is indeed the case and is a clear indication that the strong immersive topology is not too strong for our case.

Theorem 6.16. *If $(P^{(n)}, (x_m^{(n)})_{m\in\mathbb{N}})_{n\in\mathbb{N}}$ converges to $(P, (x_m)_{m\in\mathbb{N}})$, then $(\widehat{P}^{(n)})_{n\in\mathbb{N}}$ converges to $\widehat{P}$ in the strong immersive topology.*

Proof. Theorem 4.18 already shows that the sequence converges in the immersive topology. It remains to show that for each $\Delta_{\varepsilon_1}(\widehat{P}')$ with $\widehat{P} \in \Delta_{\varepsilon_1}(\widehat{P}')$ it follows that almost all $\widehat{P}^{(n)}$ lie in $\Delta_{\varepsilon_1}(\widehat{P}')$. From Lemma 6.7 it follows that there is $\varepsilon_2 \in \mathbb{R}_+$ such that

$\Delta_{\varepsilon_2}(\widehat{P}) \subseteq \Delta_{\varepsilon_1}(\widehat{P'})$ (compare the proof of Proposition 6.15). In addition for $\xi \in \mathbb{R}_+$ with $\xi < \varepsilon_2$ the set

$$\{X \in \mathcal{M} \mid \exists \xi\text{-approximation } \varphi \colon \widehat{P} \to X\}$$

is a subset of $\Delta_{\varepsilon_2}(\widehat{P})$. So it is sufficient to show that for any $\xi \in \mathbb{R}_+$ there exists an $\varepsilon \in R_+$ such that there is a ξ-approximation from $\widehat{P}$ to $\widehat{P}^{(n)}$ for each $n > N_\varepsilon$.

So let $\xi \in \mathbb{R}_+$ and let $\delta \in \mathbb{R}_+$ such that $\delta < \frac{\xi}{80}$ and there is no edge in P with length exactly 2δ. Then we choose $\varepsilon \in \mathbb{R}_+$ to be sufficiently small, i.e. small enough such that

$$8\frac{\varepsilon \cdot l \cdot m}{l_{\inf}} + 36\delta + 32\varepsilon < \xi$$

holds. For better readability, we will define the constants l, m and $l_{\inf}$ at a later point. Please note that δ and ε are much smaller than actually needed. But by making it smaller we have to put less thought into getting exact values for similar paths later in the proof.

Additionally, ε should be small enough that

$$2\varepsilon < |2\delta - l(x_i)|$$

holds for each edge x_i of P. This is possible because there is only a finite number of edges of P which are longer than δ.

Now let $D \in \mathbb{R}_+$ such that $P \subseteq B_{\frac{D}{2}}(o_P) \subseteq \mathbb{C}$. Then ε should also fulfill

$$\varepsilon < \frac{\delta^2}{4D},$$

which is equivalent to

$$D\left(1 - \frac{\delta - 4\varepsilon}{\delta}\right) < \delta.$$

Lastly, ε should be small enough such that Lemma 4.15 still holds for 3ε (for each long enough edge of P as well as $P^{(n)}$), i.e. the ε which we use here is smaller than one third of the ε arising from Lemma 4.15. In the following $n \in \mathbb{N}$ is always chosen such that $n > N_\varepsilon$.

The map $\widehat{\varphi} \colon \widehat{P} \to \widehat{P}^{(n)}$ is given through definition of a map $\varphi \colon P \to P^{(n)}$ defined as follows (see also Figure 33):

(1) On $\mathrm{IP}_{4\varepsilon}$, φ is the identity.

(2) Each edge x_i is mapped to $x_i^{(n)}$ by the affine map which maps $s(x_i)$ to $s(x_i^{(n)})$ and $t(x_i)$ to $t(x_i^{(n)})$.

(3) For each remaining point we chose a closest point x in ∂P and map this point to $\varphi(x)$ as defined in (2).

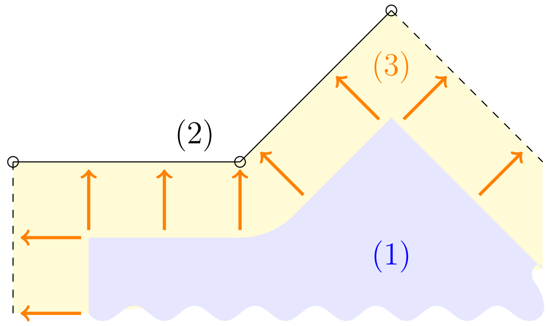

Figure 33: A sketch of the subdivision of $\widehat{P}$ used to define φ. The blue part (1) is mapped via the identity to $\widehat{P}^{(n)}$. The edges marked in black (2) are mapped to the corresponding edges of $\widehat{P}^{(n)}$ and stretched accordingly. The orange area (3) is mapped via the arrows to the edges.

Then $\varphi(o_P) = o_{P^{(n)}}$. In addition $\mathrm{Im}(\varphi) = \mathrm{IP}_{4\varepsilon} \cup \partial P^{(n)}$. But in $P^{(n)}$ each point either lies in $\mathrm{IP}_{4\varepsilon}$ or has distance smaller than 5ε to $\partial P^{(n)}$, so the image is 5ε-dense in $P^{(n)}$. So $\widehat{\varphi}$ fulfills the condition (a) and (b) of Definition 6.4.

We now show that $\widehat{\varphi}$ fulfills the conditions (c) and (d). We begin showing this for two points $q_1, q_2 \in \mathrm{IP}_{4\varepsilon}$. Let $\widehat{\gamma}\colon [0,1] \to \widehat{P}$ be a shortest path from q_1 to q_2. Again there is a corresponding map $\gamma\colon [0,1] \to P$ (compare Remark 4.10). We now want to find a ξ-similar path from q_1 to q_2 in $\widehat{P}^{(n)}$. This may seem trivial at first but is in fact rather complicated, because γ may cross any number of edges of P. In addition, the positioning of these edges can be very different in $P^{(n)}$ if the length of these edges is smaller than ε. This makes it necessary to distinguish two cases: In the first, we look at the case if γ crosses only edges of a certain length. Then γ can cross ∂P only a finite number of times. In the second one, we look at the case if γ comes close to a vertex. In this case, we can use the corresponding singularity of $\widehat{P}^{(n)}$ as a shortcut. At this point, we want to call to mind that Proposition 4.8 shows that $\widehat{P}$ has only one singularity, which we will call y.

Case 1: $d(\widehat{\gamma}(a), y) > \delta$ for all $a \in [0,1]$.
Then $\widehat{\gamma}$ can only cross a finite number of edges, namely those of length greater than 2δ. Each edge in $P^{(n)}$ can at maximum be 2ε longer or shorter than the corresponding edge in P. The above condition for ε implies that $l(x_i)+2\varepsilon < 2\delta$ for each edge x_i of P shorter than 2δ and $l(x_i) - 2\varepsilon > 2\delta$ for each edge x_i of P longer than 2δ. This shows that the edges of $P^{(n)}$ which are longer than 2δ are exactly the same as the edges of P longer than 2δ.

We now look at the set Γ of all paths in P and in each $P^{(n)}$ (for each $n > N_\varepsilon$) which go from one of the edges longer than 2δ to another one of those edges and which stay more

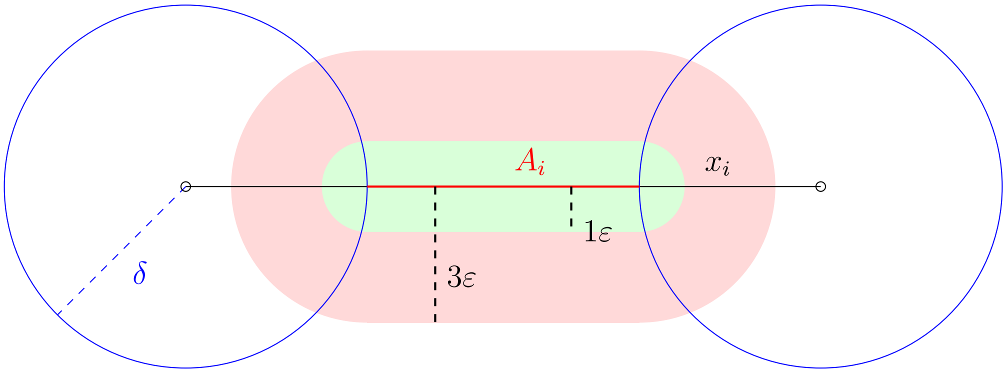

Figure 34: Each path which starts at $x_i^{(n)}$ at a point of distance greater than δ to the endpoints of $x_i^{(n)}$, has to start in the green area. This holds because $n > N_\varepsilon$. In addition because Lemma 4.15 holds for 3ε, no edge of $P^{(n)}$ can intersect the red area. So each path from such a point on $x_i^{(n)}$ to another edge $x_j^{(n)}$ has to be longer than 2ε.

than δ away[14] from each vertex. Then the set

$$\{l(\gamma) \mid \gamma \in \Gamma\}$$

is bounded from below by 2ε. This holds because each such path in $P^{(n)}$ which goes from $x_i^{(n)}$ to $x_j^{(n)}$ starts at maximum ε away from the set A_i of Lemma 4.15 (belonging to the edge x_i of P). But Lemma 4.15 shows that $x_j^{(n)}$ has distance greater than 3ε from this set. See Figure 34 for an illustration of this situation. We denote by l_{inf} the infimum of these lengths. Then $\widehat{\gamma}$ can cross an edge x_i at maximum $\frac{l(x_i)}{l_{\text{inf}}}$ times or else it would be shorter to go along this edge instead of going through P° to x_i'. Together it follows that $\widehat{\gamma}$ can cross $\widehat{\partial P}$[15] a maximum of

$$\frac{l \cdot m}{l_{\text{inf}}}$$

times, where m is the number of edges of P of length greater than 2δ and l is the length of the longest edge of P plus 2ε.

Next, we look at one of the segments of γ, i.e. a part of $\widehat{\gamma}$ going from one edge x_i to another edge x_j which crosses no edge[16]. We call this segment γ_{seg}. Our next goal is to construct a 4ε-similar path to γ_{seg} in $P^{(n)}$. The following steps are also depicted in Figure 35. No edge of P except x_i and x_j intersects the set

$$B_{2\varepsilon}(\text{Im}(\gamma_{\text{seg}})) \subseteq \mathbb{C}.$$

[14] With the distance measured in P, or $P^{(n)}$ respectively.

[15] By $\widehat{\partial P}$ we denote the subset of $\widehat{P}$ corresponding to the edges of P.

[16] The start- or endpoint of this part can also be the start- or endpoint of $\widehat{\gamma}$ instead of a point on an edge, but then the following arguments still hold.

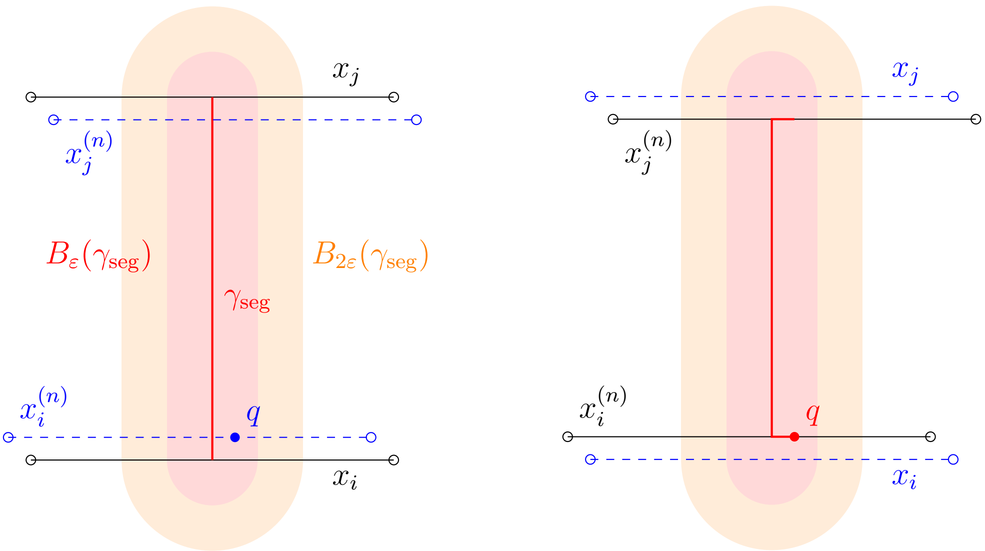

(a) The path γ_{seg} in P. The blue dashed lines represent the positions of the corresponding edges in $P^{(n)}$.

(b) A 4ε-similar path to γ_{seg} in $P^{(n)}$. The blue dashed lines represent the positions of the corresponding edges in P.

Figure 35: For each shortest path γ_{seg} connecting two edges of P, which has distance greater than δ to each vertex, a 4ε-similar path in $P^{(n)}$ can be constructed. No edge of P except the two displayed ones can intersect the orange area ($B_{2\varepsilon}(\gamma_{\text{seg}})$). Therefore no edge of $P^{(n)}$ except the two displayed ones can intersect the red area ($B_\varepsilon(\gamma_{\text{seg}})$).

This holds because of Lemma 4.15 and the fact that $2\varepsilon < \delta$: No vertex can lie in $B_{2\varepsilon}(\mathrm{Im}(\gamma_{\text{seg}}))$ and no edge can intersect γ_{seg}. In addition because of Lemma 4.15 no other edge can intersect $B_{3\varepsilon}(\gamma_{\text{seg}}(0))$ and $B_{3\varepsilon}(\gamma_{\text{seg}}(1))$, so no edge can intersect $B_{2\varepsilon}(\mathrm{Im}(\gamma_{\text{seg}}))$ because γ_{seg} is a straight line segment.

But then no other edge of $P^{(n)}$ intersects the set

$$B_\varepsilon(\mathrm{Im}(\gamma_{\text{seg}})) \subseteq \mathbb{C},$$

because each edge cannot be shifted more than ε compared to P. Now let q be a point on $x_i^{(n)}$ with distance smaller than ε to $\gamma_{\text{seg}}(0)$. Then there is a 4ε-similar path in $P^{(n)}$ from q to a point in $x_j^{(n)}$ with distance smaller than ε to $\gamma_{\text{seg}}(1)$. This path is found by going from q to the nearest point of $\mathrm{Im}(\gamma_{\text{seg}}) \cap P^{(n)}$ then following γ_{seg} as long as possible and then going to the nearest point on $x_j^{(n)}$, compare Figure 35.

By concatenating all of these parts we get a path in $\widehat{P}^{(n)}$ which is

$$\frac{4\varepsilon \cdot l \cdot m}{l_{\inf}} < \xi$$

similar to $\widehat{\gamma}$ (compare Lemma 6.6). For the same reasons, for each shortest path $\widehat{\gamma}^{(n)}$ in $\widehat{P}^{(n)}$ connecting $\widehat{\varphi}(q_1)$ and $\widehat{\varphi}(q_2)$ such that the image of $\widehat{\gamma}^{(n)}$ has distance more than δ from the singularity of $\widehat{P}^{(n)}$, a ξ-similar path in $\widehat{P}$ connecting these two points can be constructed.

Case 2: $d(\widehat{\gamma}(a), y) \leq \delta$ for an $a \in [0,1]$.
Then there is an edge x_i, such that $d(\gamma(a), s(x_i)) \leq \delta$, but this distance is measured in $\overline{\widehat{P}}$, i.e. there is a path of length smaller than δ in $\overline{\widehat{P}}$ such that the corresponding map in P connects $\gamma(a)$ and $s(x_i)$[17].

Let a_1 denote the first such a and a_2 the last. Then, because $\widehat{P}$ has only one singularity it follows that $d(\widehat{\gamma}(a_1), \widehat{\gamma}(a_2)) \leq 2\delta$ and therefore $\widehat{\gamma}([a_1, a_2])$ is shorter than 2δ. We can then construct a path in $\overline{\widehat{P}}$ which is the same as $\widehat{\gamma}$ until it reaches $\widehat{\gamma}(a_1)$ and then goes directly (i.e. on a shortest path) to the singularity, from where it goes directly to $\widehat{\gamma}(a_2)$. After that, it is again the same as $\widehat{\gamma}$. This is not a viable path in $\widehat{P}$ because it has the singularity in its image. But it can easily be adjusted to an arbitrarily similar path, which is also a path in $\widehat{P}$. This path $\widehat{\gamma}_1$ is then 4δ-similar to $\widehat{\gamma}$.

We can then assume that $\widehat{\gamma}_1$ does not cross an edge x_j at a point with distance smaller than δ to one of the endpoints of x_j. Because if it did, this intersection point has to be closer than δ to $s(x_i)$, because a_1 is the first such a. Then we can change $\widehat{\gamma}_1$ to go along this edge to the singularity. By doing this twice (once for the part before $\widehat{\gamma}_1$ hits the singularity and once for the part after) we get a new path for which the assumption is true and which is still 8δ-similar to $\widehat{\gamma}$.

[17]Instead of a vertex $s(x_i)$, this could also be an accumulation point of vertices. But the argument stays the same, because there are vertices arbitrarily close to this accumulation point.

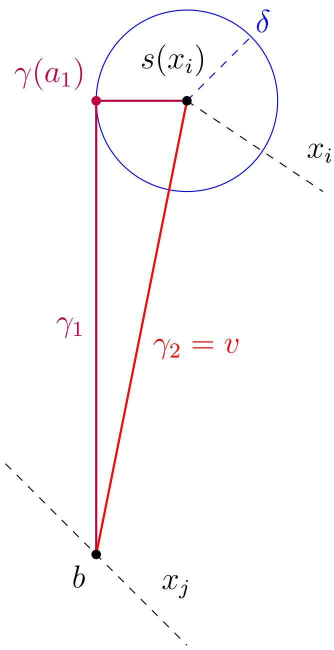

Figure 36: A depiction of the situation if γ_1 goes through $\gamma(a_1)$ after b. Then a similar path γ_2 can be constructed, such that the part between b and $s(x_i)$ is a straight line segment, called v.

Let now b denote the first point in $\mathrm{Im}(\gamma_1)$ such that the part of γ_1 between b and $s(x_i)$ is connected, i.e. does not go through an edge. So either $b = \gamma(0)$ or b lies on an edge x_j and has distance more than δ to the endpoints of that edge. If γ_1 goes through b after it has reached $\gamma(a_1)$ then the part of γ_1 between b and $s(x_i)$ is a straight line segment, because this part is a shortest path in P connecting these two points and there is no other vertex closer to b than $s(x_i)$. We then set $\gamma_2 = \gamma_1$.

So let us now take a look at the case if γ_1 goes through b before it has reached $\gamma(a_1)$. This case is also depicted in Figure 36. Then the part of γ_1 between b and $\gamma(a_1)$ is also a straight line segment with distance more than δ to each singularity. So there is a straight line segment connecting b and $s(x_i)$ and we denote by γ_2 the path which is the same as γ_1, but with the part between b and $s(x_i)$ replaced by this line segment. Then γ_2 is 12δ-similar to γ. In the following the straight line segment connecting b and $s(x_i)$ will be called v. In addition by g we denote the line which contains v. Please note that for each $x \in v$ with

$$d(x, s(x_i)) > \delta > D\left(1 - \frac{\delta - 4\varepsilon}{\delta}\right)$$

it still holds that $d(x, s(x_k)) > 4\varepsilon$ for each edge x_k (this is just Thales' theorem, see Figure 37).

Up to this point, we have only modified our path in P. Let us now look at the situation in $P^{(n)}$. Because of the assumption made in the second last paragraph, we can repeat

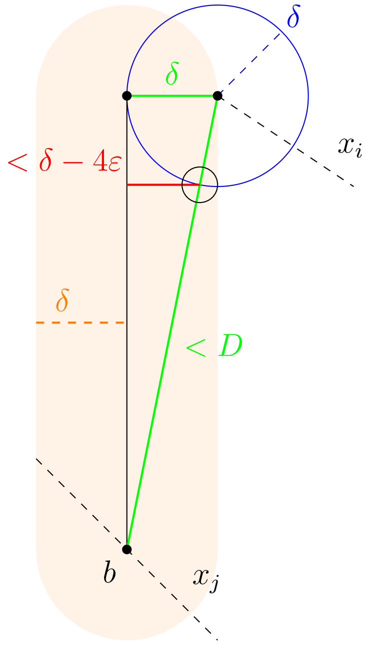

Figure 37: Because the long green line is shorter than D, Thales's theorem implies that the red horizontal line is shorter than $\delta - 4\varepsilon$. Therefore the 4ε-ball around x (painted in black) lies completely in the orange area which contains no vertex of P.

the construction as in case 1 up to the last edge x_j that $\widehat{\gamma}_2$ crosses before it reaches the singularity. Therefore we get a

$$4\frac{\varepsilon \cdot l \cdot m}{l_{\text{inf}}} + 12\delta$$

similar path up to a point $b^{(n)} \in x_j^{(n)}$ (or $b^{(n)} = \gamma(0)$ if no such edge exists) with distance less than ε from b.

We now look at the straight line segment $v^{(n)}$ in $\mathbb{C}$ from $b^{(n)}$ to $s(x_i^{(n)})$. If $v^{(n)}$ meets no edge of $P^{(n)}$, then going along $v^{(n)}$ leads to a

$$4\frac{\varepsilon \cdot l \cdot m}{l_{\text{inf}}} + 12\delta + 2\varepsilon$$

similar path in $P^{(n)}$ to $s(x_i^{(n)})$.

We now look at the case where $v^{(n)}$ hits an edge $x_k^{(n)}$. If one endpoint of $x_k^{(n)}$ lies in $B_{2\varepsilon}(v)$ then this endpoint has to be at maximum δ away from $s(x_i)$ as we have seen in Figure 37. We can then go to this endpoint instead of $s(x_i^{(n)})$ and get an

$$\frac{\varepsilon \cdot l \cdot m}{l_{\text{inf}}} + 18\delta + 2\varepsilon$$

similar path in $P^{(n)}$ instead.

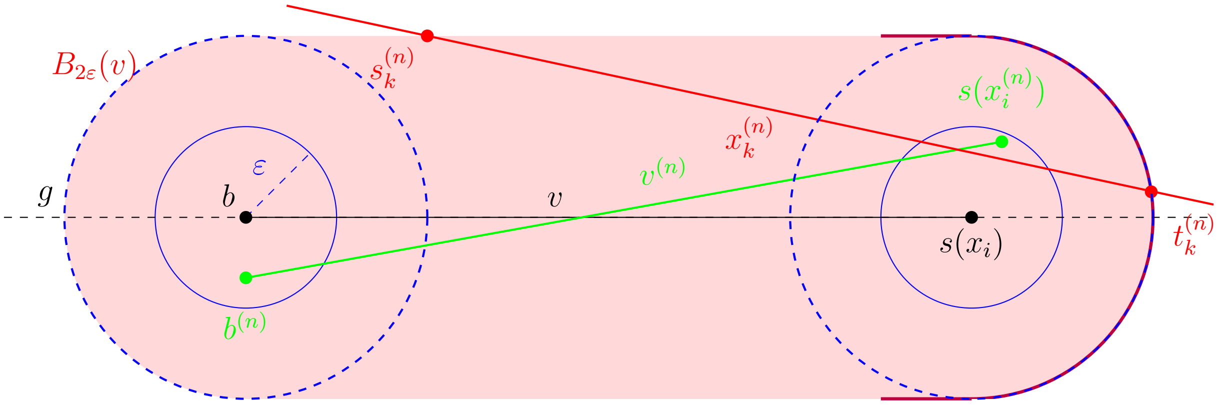

Figure 38: A depiction of the case, where an edge $x_k^{(n)}$ intersects $v^{(n)}$, has no endpoint in $B_{2\varepsilon}(v)$ and the endpoints of $x_k^{(n)}$ are more than δ away from the intersection point. Then because x_k does not intersect v one of the points $s_k^{(n)}$ and $t_k^{(n)}$ has to be on the purple line and therefore less than 3ε away from $s(x_i)$.

So let us look at the case if no endpoint of $x_k^{(n)}$ lies in $B_{2\varepsilon}(v)$. Let us take a closer look at how such an edge can be positioned. The following considerations are also depicted in Figure 38. With Lemma 4.15 (or if $b = \gamma(0)$ just with $b \in \mathrm{IP}_{4\varepsilon}$) it follows that no edge of $P^{(n)}$ can intersect $B_{2\varepsilon}(b)$. Denote by $s_k^{(n)}$ and $t_k^{(n)}$ the first and last point at which $x_k^{(n)}$ hits $B_{2\varepsilon}(v)$. The edge x_k does not intersect v. This leads to two options for the positioning of the points $s_k^{(n)}$ and $t_k^{(n)}$. One option is that they are both at a position at maximum ε away from two points which are at the same side of g. The other one is that they both are at a position at maximum ε away from two points which border $B_{2\varepsilon}(s(x_i))$. Then at least one of these points has to be on the purple line in Figure 38.

But then this point is closer than 3ε to $s(x_i)$ and Lemma 4.15 tells us that it is closer than δ to one of the endpoints of $x_k^{(n)}$[18]. So again, as before we can go to this endpoint of $x_k^{(n)}$ instead of $s(x_i^{(n)})$ and get a

$$4\frac{\varepsilon \cdot l \cdot m}{l_{\inf}} + 16\delta + 8\varepsilon$$

similar path.

The same procedure can be done for the part of γ_2 from the singularity to the endpoint. All in all we get a

$$8\frac{\varepsilon \cdot l \cdot m}{l_{\inf}} + 32\delta + 16\varepsilon < \xi$$

similar path in $P^{(n)}$. These steps can also be done for a shortest path in $P^{(n)}$ to construct a similar path in P.

[18] No point of an edge $x_k^{(n)}$ with distance more than δ to $s(x_k^{(n)})$ and $t(x_k^{(n)})$ can lie in $B_{3\varepsilon}(s(x_i))$.

This shows that conditions (c) and (d) of a ξ-approximation are fulfilled for two points of the form (1) of the above subdivision, i.e. for two points in $\mathrm{IP}_{4\varepsilon}$.

If one of these points is of the form (2), then these arguments still hold. If one of these points (w.l.o.g. the startpoint) is of the form (3), then we can extend γ to start at the closest point x in ∂P instead and repeat the same construction as before. This makes the path up to 4ε longer so the resulting path is still

$$8\frac{\varepsilon \cdot l \cdot m}{l_{\inf}} + 36\delta + 24\varepsilon < \xi$$

similar. So if both points are of the form (3) we still get a

$$8\frac{\varepsilon \cdot l \cdot m}{l_{\inf}} + 36\delta + 32\varepsilon < \xi$$

similar path. All in all, we have shown that the conditions (c) and (d) of a ξ-approximation are fulfilled by $\widehat{\varphi}$. □

Remark 6.17. *Theorem 6.16 tells us, that the strong immersive topology is not that much stronger in the case of translation surfaces coming from PSI-polygons. But the additional restrictions imposed through the strong immersive topology are a lot stricter when we take translation surfaces of infinite area into account. Most prominently the convergence of a sequence of translation surfaces in the immersive topology implies the convergence of their universal covers. This is not the case for the strong immersive topology.*

In particular, if $(P^{(n)})_{n\in\mathbb{N}}$ is a sequence of PSI-polygons converging to another PSI-polygon P this does not imply that $(\widetilde{P}^{(n)})_{n\in\mathbb{N}}$ converges to $\widetilde{P}$ in the strong immersive topology because a very small shift of one edge x_i can shift the copy $P\times\{x_i^m\}$ an arbitrarily high amount if m is large enough.

6.5 An infinite-dimensional complex family

At last, we want to use the strong immersive topology to construct even larger families of infinite translation surfaces. Unfortunately, because of the behavior discovered in Example 5.13, the converse of Theorem 6.16 is still not true. But we can help ourselves by looking only at such PSI-polygons for which the convergence of the vertices to the accumulation points still happens in a somewhat orderly manner. That means, that the projection of these vertices to an appropriate 1-dimensional affine subspace strictly converges to the accumulation point. This leads to an infinite-dimensional complex family.

For this section let $n \in \mathbb{N}$, $y \in \mathbb{R}^2 \setminus B_1(0)$, $\delta \in (0, \frac{1}{16})$ and $v_1, v_2 \in \mathbb{R}^2$ be linearly independent with length 1. In addition $\gamma\colon [0,1] \to \mathbb{R}^2$ is a path with $\gamma(0) = 0$ and $y = \gamma(1)$.

Convention 6.18. For the 1-dimensional affine space $A_1 = y + \langle v_1 \rangle$ we denote by $\pi_{v_1} \colon \mathbb{R}^2 \to A_1$ the orthogonal projection onto this space. Moreover there is a natural ordering on A_1 with $y + v_1 > y$.

Definition 6.19. $\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$ is the set of PSI-polygons P of the following form (compare Figure 39):

- $o_P = 0 \in \mathbb{C}$.
- P fulfills the δ-condition (compare Definition 5.7).
- $\gamma([0, 1))$ lies in the interior of P.
- There are exactly 2 accumulation points of the vertices of P and they are positioned at $y = \gamma(1)$ and y'.
- Each $x \in \mathbb{R}^2 \setminus \{y\}$ with $d(x, y) < 1$ which lies (mathematically positive) between v_1 and v_2 also lies in P. Or to be more precise, if x fulfills

  $$\det(x - y \;\; v_1) \geq 0 \text{ and } \det(x - y \;\; v_2) \leq 0,$$

  then $x \in P$.
- $\overline{B_1(y)} \cap \partial P$ is connected and therefore $(\overline{B_1(y)} \cap \partial P) \setminus \{y\}$ has two connected components. One of these components, together with the one edge which lies partially in this component and partially outside of $\overline{B_1(y)}$ will be called a *converging side*. The two converging sides at y will be denoted by c_1 and c_2 and it should hold that

  $$\det(x - y \;\; v_1) < 0 \text{ and } \pi_{v_1}(x) > y$$

  for all $x \in c_1$ and

  $$\det(x - y \;\; v_2) > 0 \text{ and } \pi_{v_2}(x) > y$$

  for all $x \in c_2$.
- Let x_i and x_{i+1} be two adjacent edges in c_1 such that x_{i+1} is closer to y (with the distance measured on ∂P) than x_i. Then it should hold that

  $$\pi_{v_1}(s(x_i)) > \pi_{v_1}(s(x_{i+1})).$$

  The same should hold for all edges in c_2.
- P has exactly $2n$ edges which do not lie in any c_i (or c_i'). These are called *finite edges*. The corresponding start- and endpoints of these edges will be called the *finite vertices*.

$\widehat{\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)} \subseteq \mathcal{M}$ is the set of isomorphism classes of translation surfaces arising from polygons in $\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$.

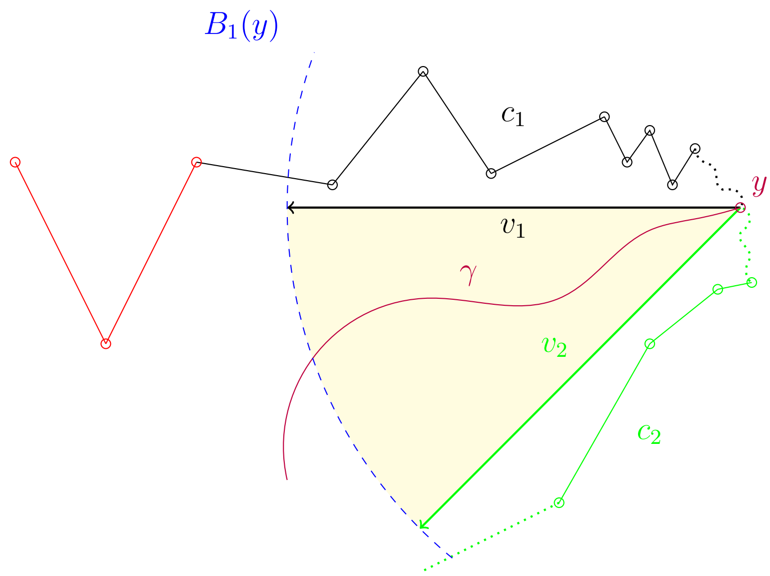

Figure 39: An illustration of the conditions applied to a PSI-polygon in $\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$. The red lines correspond to the finite edges, the black lines to the edges in c_1, and the green lines to the edges in c_2. The yellow area has to lie in the inside of such a PSI-polygon.

We often write v instead of v_1 (v_2) if the index is clear from the context. So for example the map π_{v_1} (π_{v_2}) can be written just as π_v.

We have already seen many examples of PSI-polygons lying in such a $\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$. In fact, all of the PSI-polygons we have considered in Section 5 lie in an appropriate $\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$.

From now on, for $P \in \text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$ let E be the set of all finite edges of P together with the first 4 non-finite edges of each converging side. For the following definition the edges of each such PSI-polygon are labeled in the same way as described in Convention 5.12. The following situation is also illustrated in Figure 40.

Definition 6.20. Let $\varepsilon \in \mathbb{R}_+$ and $P \in \text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$. Then the set $B_\varepsilon(P) \subseteq \text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$ consists of PSI-polygons $P^{(1)}$ such that for each edge $x_i \in E$ it holds that

$$d(s(x_i), s(x_i^{(1)})) < \varepsilon$$

and

$$d(t(x_i), t(x_i^{(1)})) < \varepsilon,$$

where the distance is measured in $\mathbb{C}$.

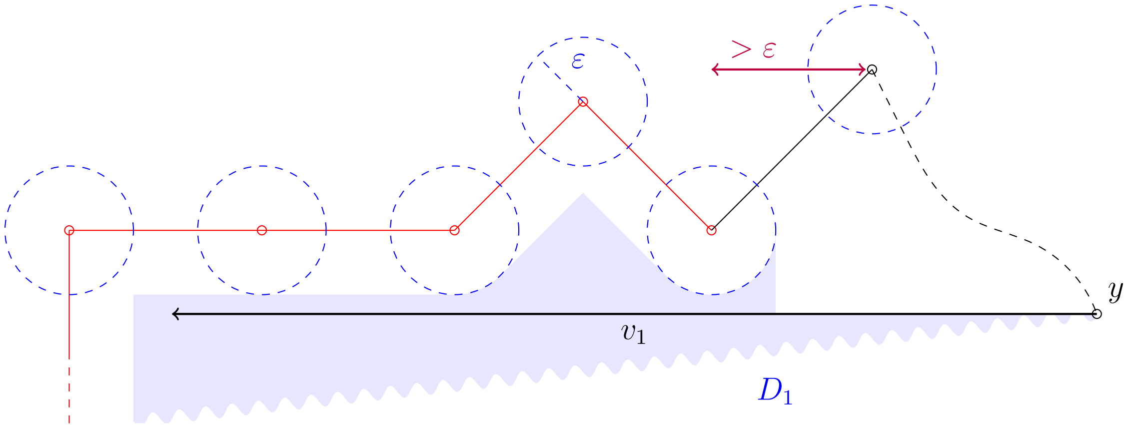

Figure 40: An illustration of the additional condition imposed on $B_\varepsilon(P)$. Here the distance marked by the purple arrow must be greater than ε. The edges which lie in E are marked in red. In addition, the set D_1 used in the proof of Lemma 6.21 is also depicted as the blue area.

In addition for those edges $x_i \notin E$ which are adjacent to an edge in E it should hold that

$$d(\pi_v(s(x_i)), \pi_v(t(x_i^{(1)}))) > \varepsilon$$

where the start of these edges lies farther away from the corresponding y or y'. Again the set $\widehat{B_\varepsilon(P)} \subseteq \mathcal{M}$ is the set of isomorphism classes of translation surfaces arising from polygons in $B_\varepsilon(P)$.

Again ε has to be sufficiently small to be of any use. So let $P \in \text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$ and $\delta_P \in \mathbb{R}_+$ be such that

$$\delta_P < d(x_i, \partial P \setminus (x_i \cup x_{i-1} \cup x_{i+1})) \; \forall x_i \in E$$

where x_{i-1} and x_{i+1} are again the edges adjacent to x_i. From now on let $\varepsilon \in \mathbb{R}_+$ be such that $\varepsilon < \frac{1}{16}\delta_P$.

In addition, for each non-finite edge $x_i \in E$ it should hold that

$$d(\pi_v(s(x_i)), \pi_v(t(x_i))) > 16\varepsilon$$

and

$$d(\pi_v(s(x_i)), s(x_i)) > 16\varepsilon.$$

Additionally for each edge $x_i \notin E$ which is adjacent to an edge in E it should hold that

$$d(\pi_v(s(x_i)), \pi_v(t(x_i))) > 2\varepsilon.$$

The last condition guarantees that $P \in B_\varepsilon(P)$.

ε should also fulfill

$$\varepsilon < \delta d(x_i, y)$$

for all $x_i \in E$. At last, ε should fulfill the additional condition stated in the proof of Lemma 5.15 (at the bottom of page 43) for each edge in E.

Now an analogue for Lemma 5.15 can be proven for $\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$.

Lemma 6.21. *Let $P \in \text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$ and $\varepsilon \in \mathbb{R}_+$ that fulfills the above conditions for P. Then $\widehat{B_\varepsilon(P)}$ is open in the immersive topology and therefore also open in the strong immersive topology.*

Proof. This follows for the same reasons as in Lemma 5.15 but with a slight variation for the converging sides.

First, we again define D_1. It will be similar as in Lemma 5.15, but with some points removed. Namely for c_1 we remove all those points lying between $A_1 = y + \langle v_1 \rangle$, ∂P and the line orthogonal to A_1 which goes through $\pi_{v_1}(s(x_i)) - \varepsilon v_1$ for the edge $x_i \notin E$ of c_1 which is adjacent to an edge in E. This set is also depicted in Figure 40. We repeat this removal for each converging side. Now for a point $x \in D_1$ which lies between A_1 and c_1 it holds that $x - \varepsilon v_1 \in P$ and $x + \varepsilon v_1 \in P$. But then the rectangle spanned by the line segment from $\pi_{v_1}(x - \varepsilon v_1)$ to $x - \varepsilon v_1$ and the line segment from $x - \varepsilon v_1$ to $x + \varepsilon v_1$ lies completely in P. This shows that D_1 is simply connected.

The set D_2 will be adjusted accordingly and in addition, for each vertex only one of the sets H_2 and H_3 is used, namely the one in the direction, in which more edges of E lie. But to accommodate the missing set, we will make the present set larger by turning the line 360 degree instead of 180.

Now for all non-finite edges $x_i \in E$ we replace the K_i from Lemma 5.15 with the following sets which are illustrated in Figure 41 and which will be used again later: Let $s(x_i)$ be farther away from y than $t(x_i)$. We define $H_i \subseteq \mathbb{C}$ to be the rectangle spanned by the line segment from $\pi_{v_1}(s(x_i)) - \varepsilon v_1$ to $\pi_{v_1}(t(x_i)) + \varepsilon v_1$ (the purple arrow in Figure 41) and the line segment from $\pi_{v_1}(m(x_i))$ to $\pi_{v_1}(m(x_i)) + 2 \cdot (m(x_i) - \pi_{v_1}(m(x_i)))$ (the green arrow in Figure 41). Then $D_1 \cup H_i$ can be embedded in $\widetilde{P}$, so let $K_i \subseteq \widetilde{P}$ be the set obtained by the union of this embedded set and $D_1' \times \{x_i\}$ (again compare Lemma 5.15 and remove points as above) which is then connected (compare the requirements for ε) and compact. Again we take

$$U_i := B_\varepsilon(o_P) \times \{x_i\}$$

which is an open subset of K_i.

Similar to Lemma 5.15 we define

$$V := \mathcal{M}_{\hookrightarrow}(D_1) \cap \mathcal{M}_{\rightsquigarrow}(D_2) \cap \bigcap_{x_i \in E} \mathcal{M}_+(K_i, U_i) \cap \widehat{\text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)}.$$

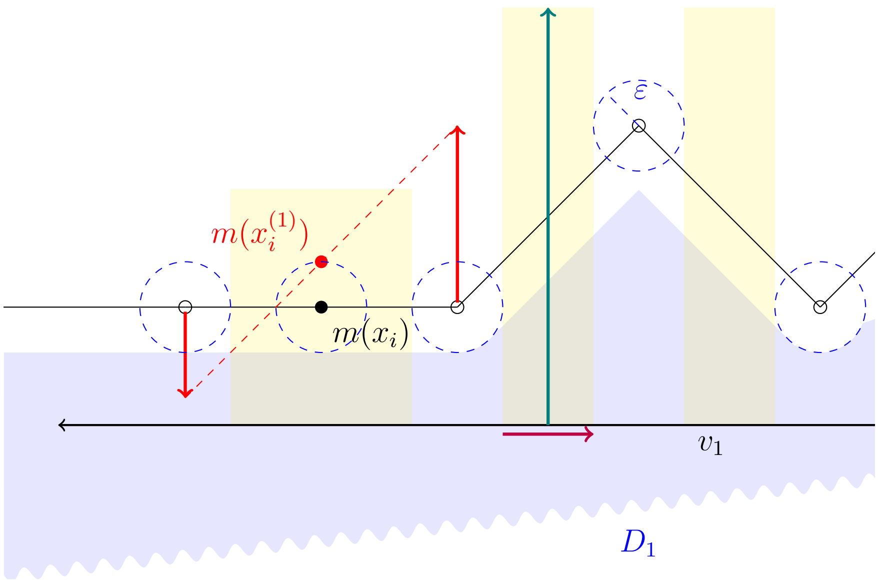

Figure 41: A sketch of the construction of the new K_i in the proof of Lemma 6.21. The yellow areas correspond to the rectangles named H_i. The effects of these new K_i are also illustrated: If in any $P^{(1)}$ one vertex is too far up compared to P the next vertex is too far down. This is depicted by the red arrows.

Now let $P^{(1)}$ be a PSI-polygon such that $\widehat{P}^{(1)} \in \widehat{B_\varepsilon(P)}$. Then the embedding of D_1 lies in the interior of $P^{(1)}$, therefore $\widehat{P}^{(1)} \in \mathcal{M}_{\hookrightarrow}(D_1)$. This follows because the first vertex which is not an endpoint of an edge of E lies in or above the area we removed from D_1. $\widehat{P}^{(1)} \in \mathcal{M}_{\hookrightarrow}(D_2)$ follows as in Lemma 5.15, together with the insight, that the sum of the interior angles at two neighboring vertices in a converging side is more than 180 degree. $\widehat{P}^{(1)} \in \mathcal{M}_{+}(K_i, U_i)$ follows in the same way as in Lemma 5.15.

Now let $P^{(1)} \in \text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$ such that $\widehat{P}^{(1)} \in V$. For the same reasons as in Lemma 5.15 we get arrows from D_1 to D_1 with the arrows corresponding to the edges of the converging sides facing orthogonal to A_i. In addition, because $\gamma([0,1))$ lies in the interior of P, the point y is at the correct position relative to D_1. But this also means, that for every arrow, which lies at least partially in $B_1(y)$, the edge of $P^{(1)}$ responsible for this arrow belongs to the same converging side this arrow corresponds to.

But then the endpoints of these edges in $P^{(1)}$ are at most ε away from the endpoints of the edges in P corresponding to these arrows. To see this, we first see, that the endpoints of these edges have to lie (in v-direction) between the H_i belonging to those arrows. In addition, in the direction orthogonal to v, these endpoints cannot be more than ε closer to $y + \langle v \rangle$ or they would collide with D_2. But this also means, that no such endpoint can be more than 3ε farther away from $y + \langle v \rangle$, because the position of the midpoint is given by the arrow and therefore the next vertex would be too close to $y + \langle v \rangle$ (compare again Figure 41). But then $\widehat{P} \in \mathcal{M}_{\rightsquigarrow}(D_2)$ shows that it can be no more than ε farther away from $y + \langle v \rangle$. Then D_2 also enforces that these endpoints lie in an ε-ball from the corresponding endpoints in P.

Because ε fulfills

$$\varepsilon < \delta d(x_i, y)$$

for all $x_i \in E$ the δ-condition implies that only one vertex can lie in each of these ε-balls.

The rest of the proof follows exactly as in Lemma 5.15. □

In Section 5 the main problem we faced was that the converse of Theorem 4.18 is not true in the general case. But as we have seen in Theorem 5.9 those translation surfaces behave nicely if the vertices are lined up in a 1-dimensional affine subspace. If this is not the case then problems can arise because one edge can cut through the interior of P (see Example 5.13). Another problem which can arise is that one edge goes too far away from P like in Example 6.1. In the following lemma, we show that both of these problems cannot happen in our case. The first, because we restrict ourselves to $B_\varepsilon(P)$ and the second because the strong immersive topology forbids such behavior.

Lemma 6.22. *Let* $\left(P^{(n)}\right)_{n\in\mathbb{N}}$ *be a sequence in* $\text{Co-Ch}_{n,\delta}(y, v_1, v_2)$ *and* $P \in \text{Co-Ch}_{n,\delta}(y, v_1, v_2)$ *such that* $(\widehat{P}^{(n)})_{n\in\mathbb{N}}$ *converges to* $\widehat{P}$ *in the strong immersive topology. Let* $c_i^{(n)}$ *be a converging side of* $P^{(n)}$ *and* $c^{(n)}$ *be the same converging side, but with the last edge (the edge which lies only partially in* $\overline{B_1(y)}$*) removed. Then for all* $\varepsilon \in \mathbb{R}_+$ *it holds that* $c^{(n)} \subseteq B_\varepsilon(\partial P)$ *for almost all* $n \in \mathbb{N}$.

Proof. W.l.o.g. let the converging side be $c_1^{(n)}$ and the edges which lie (again at least partially) in $c_1^{(n)}$ will be labeled $x_1^{(n)}, \dots$ again in such a way that $x_i^{(n)}$ and $x_{i+1}^{(n)}$ are adjacent and $x_{i+1}^{(n)}$ is closer to y than $x_i^{(n)}$ for all $i \in \mathbb{N}$. We label the edges of the corresponding converging side c_1 of P in the same way.

Let us assume that the statement is not true. Then for infinitely many n there is $i_n \in \mathbb{N} \setminus \{1\}$ such that $x_{i_n}^{(n)} \not\subseteq B_\varepsilon(\partial P)$. For each such n we always choose i_n to be the smallest such integer. Now let $\text{IP} \coloneqq \text{IP}_\varepsilon$ and

$$\text{OP} \coloneqq \{x \in \mathbb{C} \mid d(x, P) \geq \varepsilon\}.$$

Case 1: $x_{i_n}^{(n)} \cap \text{IP} \neq \emptyset$ for infinitely many $n \in \mathbb{N}$.
Let $x \in x_{i_n}^{(n)} \cap \text{IP}$. Then $B_\varepsilon(x) \subseteq P$ and therefore also $x - \varepsilon v_1 \in P$. But this implies $\pi_{v_1}(x) - \varepsilon v_1 \in P$, so $\pi_{v_1}(x) \geq y + \varepsilon v_1$.

We can now take $\mathcal{M}_+(K_j, U_j)$ as in the proof of Lemma 6.21 for all edges x_j of $c_1 \subseteq \partial P$ which fulfill $\pi_{v_1}(x_j) \not\subseteq B_\varepsilon(y)$. Of course, these are only finitely many edges. We get that for only finitely many $n \in \mathbb{N}$ an $s(x_j^{(n)})$ (for such an edge $x_j^{(n)}$) can be shifted more than ε towards A_1 or more than 3ε away from A_1, compared to the situation in P. A difference to the situation in Lemma 6.21 is that we cannot assume that only one vertex of $P^{(n)}$ lies in each area between two H_i. But if there are at least two vertices in such an area, then the edges between these vertices cannot intersect IP (for all but finitely many n). This holds because the converse would imply, in this situation, that there is a vertex in IP, but $\overline{\mathrm{IP}}$ can be immersed in all but finitely many $\widehat{P}^{(n)}$. All in all we get that such a point $x \in x_{i_n}^{(n)} \cap \mathrm{IP}_\varepsilon$ can only exist for finitely many $n \in \mathbb{N}$ which is a contradiction.

Case 2: $x_{i_n}^{(n)} \cap \mathrm{OP} \neq \emptyset$ for infinitely many $n \in \mathbb{N}$.

Let $\mu \in (0,1)$ be such that no point $b \in c_1$ with $d(b, A_1) > \frac{\varepsilon}{4}$ fulfills $\pi_{v_1}(b) \leq y + \mu v_1$.

For the following let $\xi \in \mathbb{R}_+$ with

$$\xi < \frac{\mu\delta^2\varepsilon}{32}.$$

Therefore it also holds that

$$\xi < \frac{\delta\varepsilon}{32}.$$

In addition, let $\beta' \in \mathbb{R}_+$ be the value given by Lemma 6.12 for $\alpha = \frac{\xi}{4}$ and $\beta = \frac{\beta'}{2}$. We will now show that there cannot be a $\frac{\beta}{8l}$-approximation from $\widehat{P}$ to any of those $\widehat{P}^{(n)}$, where l is defined as in Lemma 6.11 for IP_β.

Because case 1 can also be done for IP_β instead of $\mathrm{IP} = \mathrm{IP}_\varepsilon$ we can assume that no edge of a converging side (apart from the first edge) intersects IP_β. In addition Lemma 6.21 shows that there is $\varepsilon' \in \mathbb{R}_+$, $\varepsilon' < \beta$ such that $B_{\varepsilon'}(\widehat{P})$ is open. Then only finitely many $\widehat{P}^{(n)}$ do not lie in $B_{\varepsilon'}(\widehat{P})$. Therefore $\mathrm{IP}_\beta \subseteq P^{(n)}$ for infinitely many $n \in \mathbb{N}$. We can see $\widehat{\mathrm{IP}}_\beta$ as a subset of $\widehat{P}$ as well as $\widehat{P}^{(n)}$. Furthermore, β can be chosen in such a way that $\mathrm{IP}_{2\beta}$ is connected (compare the statement shown for D_1 at the beginning of the proof of Lemma 6.21).

Claim 1. For each such n there is a point $\widehat{x}^{(n)} \in \widehat{P}^{(n)}$ which has distance greater than ξ from each point in $\widehat{\mathrm{IP}_\beta} \subseteq \widehat{P}^{(n)}$.

Let us assume for a moment that claim 1 holds. Then only finitely many $\widehat{P}^{(n)}$ do not lie in $\Delta_{\frac{\beta}{8l}}(\widehat{P})$. A $\frac{\beta}{8l}$-approximation $\widehat{\varphi}^{(n)} \colon \widehat{P} \to \widehat{P}^{(n)}$ has to map at least one point $\widehat{x} \in \widehat{P}$ to a point in $B_{\frac{\beta}{8l}}(\widehat{x}^{(n)}) \subseteq \widehat{P}^{(n)}$.

Then $\widehat{x}$ cannot lie in $\widehat{\mathrm{IP}}_{2\beta}$ or else $\widehat{\varphi}^{(n)}(\widehat{x})$ would lie in $\widehat{\mathrm{IP}}_\beta$ (compare Lemma 6.11). But then

$$d(\widehat{x}^{(n)}, \widehat{\varphi}^{(n)}(\widehat{x})) > \xi > \frac{\beta}{8l},$$

which is a contradiction. If $\widehat{x}$ does not lie in $\widehat{\mathrm{IP}}_{2\beta}$ then there is a point $\widehat{a} \in \widehat{\mathrm{IP}}_{2\beta}$ with $d(\widehat{x}, \widehat{a}) < \alpha = \frac{\xi}{4}$ (compare the definition of β). But then $\widehat{\varphi}^{(n)}(\widehat{a}) \in \widehat{\mathrm{IP}}_\beta$, so

$$d(\widehat{\varphi}^{(n)}(\widehat{a}), \widehat{\varphi}^{(n)}(\widehat{x})) > d(\widehat{\varphi}^{(n)}(\widehat{a}), \widehat{x}^{(n)}) - \frac{\beta}{8l} > d(\widehat{\varphi}^{(n)}(\widehat{a}), \widehat{x}^{(n)}) - \frac{\xi}{4} > \frac{3\xi}{4}.$$

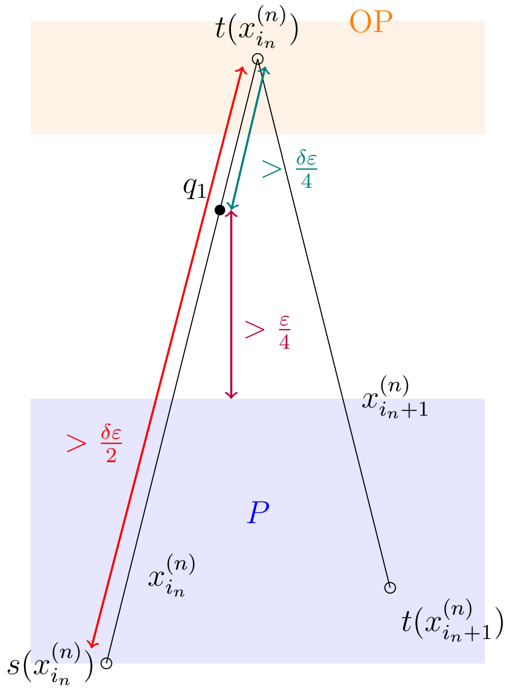

Figure 42: If at least one point of $x_{i_n}^{(n)}$ lies in OP then $x_{i_n}^{(n)}$ is longer than $\frac{\delta\varepsilon}{2}$. In addition there is a point $q_1 \in x_{i_n}^{(n)}$ with distance more than $\frac{\delta\varepsilon}{4}$ to each endpoint of $x_{i_n}^{(n)}$ and more than $\frac{\varepsilon}{4}$ to P.

But

$$|d(\widehat{\varphi}^{(n)}(\widehat{a}), \widehat{\varphi}^{(n)}(\widehat{x})) - d(\widehat{a}, \widehat{x})| > \frac{3\xi}{4} - \frac{\xi}{4} > \frac{\xi}{4}$$

and this is a contradiction to $\widehat{\varphi}^{(n)}$ being a $\frac{\beta}{8l}$-approximation and therefore also a $\frac{\xi}{4}$-approximation.

It remains to show that claim 1 holds, which follows essentially from the δ-condition. Nevertheless, it is a rather technical construction. The rest of this proof will be dedicated to proving claim 1.

Proof of Claim 1. For this let $x \in x_{i_n}^{(n)} \cap \mathrm{OP}$. Then

$$l(x_{i_n}^{(n)}) > \delta \cdot d(t(x_{i_n}^{(n)}), y)$$

and

$$l(x_{i_n}^{(n)}) > d(t(x_{i_n}^{(n)}), x) > \delta \cdot d(t(x_{i_n}^{(n)}), x).$$

So it follows that

$$l(x_{i_n}^{(n)}) > \max(\delta \cdot d(t(x_{i_n}^{(n)}), x), \delta \cdot d(t(x_{i_n}^{(n)}), y)).$$

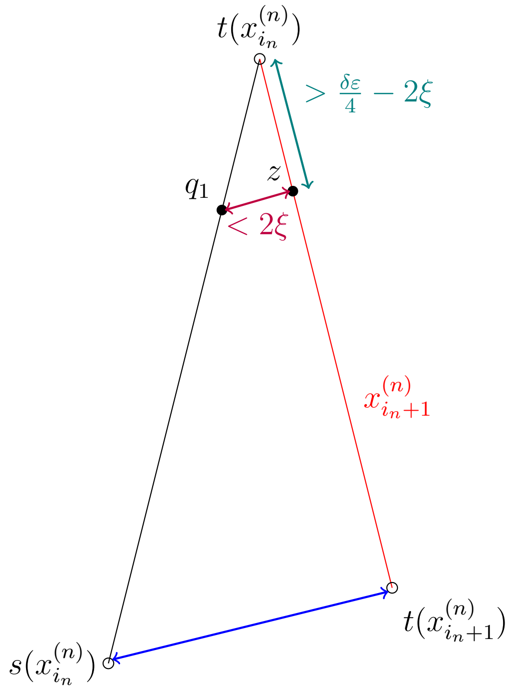

Figure 43: Thales' theorem implies that the length of the blue line is smaller than $\frac{l(x^{(n)}_{i_n+1})\cdot 2\xi}{\frac{\delta\varepsilon}{4}-2\xi}$. Therefore $L := d(t(x^{(n)}_{i_n+1}), x^{(n)}_{i_n})$ is also smaller than this number.

In addition $d(x, y) > \varepsilon$ because $x \in \mathrm{OP}$ and so the triangle-inequality implies $l(x^{(n)}_{i_n}) > \frac{\delta\varepsilon}{2}$. This, as well as the following steps, is also illustrated in Figure 42.

Then there exists a point $q_1 \in x^{(n)}_{i_n}$ which is at minimum $\frac{\delta\varepsilon}{4}$ away from $s(x^{(n)}_{i_n})$ and $t(x^{(n)}_{i_n})$ and still more than $\frac{\varepsilon}{4}$ away from P. If each other edge of $P^{(n)}$ is more than 2ξ away from q_1 (with the distance measured in $P^{(n)}$) the point $q_2 \in P^{(n)}$, which is exactly ξ away from $x^{(n)}_{i_n}$ and q_1 is still more than ξ away from P and from each other edge, so we can set $\widehat{x}^{(n)} := \widehat{q}_2$.

So let us now assume that another edge of $P^{(n)}$ is less than 2ξ away from q_1. Because q_1 has distance greater than $\frac{\varepsilon}{4}$ from y and $\xi < \frac{\delta\varepsilon}{4}$ the other edge, which is close to this point has to be $x^{(n)}_{i_n-1}$ or $x^{(n)}_{i_n+1}$. W.l.o.g. we assume it is $x^{(n)}_{i_n+1}$ and that $t(x^{(n)}_{i_n}) = s(x^{(n)}_{i_n+1})$ and $d(t(x^{(n)}_{i_n+1}), A_1) \geq d(s(x^{(n)}_{i_n}), A_1)$ hold. Let $z \in x^{(n)}_{i_n+1}$ be the closest point to q_1. Then

$$d(s(x^{(n)}_{i_n+1}), z) > \frac{\delta\varepsilon}{4} - 2\xi$$

and

$$d(q_1, z) < 2\xi.$$

Now it follows for $L := d(t(x_{i_n+1}^{(n)}), x_{i_n}^{(n)})$ (compare Figure 43):

$$L < \frac{l(x_{i_n+1}^{(n)}) \cdot 2\xi}{\frac{\delta\varepsilon}{4} - 2\xi} < \frac{l(x_{i_n+1}^{(n)}) \cdot 2\xi}{\frac{\delta\varepsilon}{4} - \frac{\delta\varepsilon}{8}} = \frac{l(x_{i_n+1}^{(n)}) \cdot 16\xi}{\delta\varepsilon} < \frac{l(x_{i_n+1}^{(n)})}{2}$$

and similar for the angle $\alpha_{i_n+1}^{(n)}$ at $s(x_{i_n+1}^{(n)})$ it follows that

$$\sin \alpha_{i_n+1}^{(n)} < \frac{16\xi}{\delta\varepsilon} < \frac{1}{2},$$

so $\alpha_{i_n+1}^{(n)} < \frac{\pi}{4}$. This implies that $t(x_{i_n+1}^{(n)})$ is closer to A_1 than $s(x_{i_n+1}^{(n)})$.

Now it follows from the δ-condition that

$$d(t(x_{i_n+1}^{(n)}), A_1) < d(t(x_{i_n+1}^{(n)}), y) < \frac{1}{\delta} L$$

and therefore the same holds for $s(x_{i_n}^{(n)})$. We now define $\pi_{v_1}^{\perp} : \mathbb{R}^2 \to \mathbb{R},\ x \mapsto x - \pi_{v_1}(x)$. Then

$$d(\pi_{v_1}^{\perp}(s(x_{i_n}^{(n)})), \pi_{v_1}^{\perp}(t(x_{i_n+1}^{(n)}))) < \frac{1}{\delta} L$$

and therefore

$$d(\pi_{v_1}^{\perp}(m(x_{i_n}^{(n)})), \pi_{v_1}^{\perp}(m(x_{i_n+1}^{(n)}))) < \frac{1}{\delta} L.$$

This is also depicted in Figure 44. According to the depiction in the previous figures we say that a point a_1 is higher (or lower) than another point a_2 if the distance of $\pi_{v_1}^{\perp}(a_1)$ to A_1 is greater (or smaller) than that of $\pi_{v_1}^{\perp}(a_2)$ to A_1. The height of a point can just be defined to be this distance.

Now let m be the point on the bisector g between these edges with

$$\pi_{v_1}^{\perp}(m) = \frac{1}{2} \pi_{v_1}^{\perp}(m(x_{i_n}^{(n)})) + \pi_{v_1}^{\perp}(m(x_{i_n+1}^{(n)}))$$

and a be the highest point of P which lies between $x_{i_n}^{(n)}$ and $x_{i_n+1}^{(n)}$ (see again Figure 44).

Let us first assume that the height of a is greater than $\frac{\varepsilon}{4}$. Similar to above it follows that

$$d(\pi_{v_1}(s(x_{i_n}^{(n)})), y) < d(s(x_{i_n}^{(n)}), y) < \frac{l(x_{i_n}^{(n)}) \cdot 16\xi}{\delta^2\varepsilon} < \frac{16\xi}{\delta^2\varepsilon}$$

because no edge of a converging side can be longer than 1. The definition of μ then leads to a contradiction because

$$\mu < d(\pi_{v_1}(a), y) < d(\pi_{v_1}(s(x_{i_n}^{(n)})), y) < \frac{16\xi}{\delta^2\varepsilon} < \mu.$$

So a is lower than $\frac{\varepsilon}{4}$. This shows that the height of m is $\frac{\varepsilon}{4}$ greater than the height of a, because the height of x is at least ε and therefore, the height of $t(x_{i_n}^{(n)})$ is also at least ε.

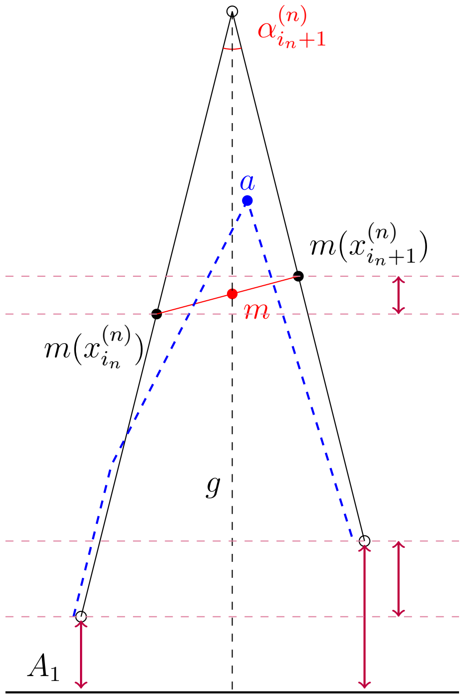

Figure 44: In this situation, each of the purple arrows is shorter than $\frac{L}{\delta}$. This follows because these edges are very close together, so the δ-condition implies that the endpoints are very close to y. This also shows that the highest point a of P between these edges cannot be that much left of y. Some possible edges of P are depicted by the blue dashed lines.

Let us now look at a path in $\widehat{P}^{(n)}$ which starts and ends at g, goes parallel to v_1 and which goes through $\widehat{x}_{i_n}^{(n)}$ and afterwards through $(\widehat{x}_{i_n+1}^{(n)})'$ (see Figure 45). It changes its height by a maximum of $2\frac{L}{\delta}$. But this path has a length of at minimum $\frac{L}{2}$. So this path changes its height only $\frac{4}{\delta}$ faster than its length. This shows that a path from $\widehat{m}$ has to be more than $\frac{\varepsilon\delta}{32}$ long to reach a point in P. This holds because m is more than $\frac{\varepsilon}{4}$ higher than a[19]. Because $\xi < \frac{\delta\varepsilon}{32}$ this shows that no path of length smaller than ξ starting from $\widehat{m}$ reaches $\widehat{\mathrm{IP}}_\beta$, so we can set $\widehat{x}^{(n)} := \widehat{m}$ and we have found such a point.

□

At last, we again create a coordinate function on each of these $\widehat{B_\varepsilon(P)}$. But let us first fix notation. So for the next Theorem, for $P^{(1)} \in \text{Co-Ch}_{n,\delta}(y, v_1, v_2)$ we denote by $x_1^{(1)}, \dots, x_n^{(1)}$ the finite edges on one side, i.e. n finite edges in such a way that the union of these edges is connected and by $(a_n^{(1)})_{n\in\mathbb{N}}$ and $(b_n^{(1)})_{n\in\mathbb{N}}$ all edges each of one converging side respectively. This should again be done in a way that the union of all of these edges

[19]The lower height might be reached before completing a full cycle, so we changed $\frac{\varepsilon\delta}{4\cdot 4}$ to $\frac{\varepsilon\delta}{32}$ just to be sure.

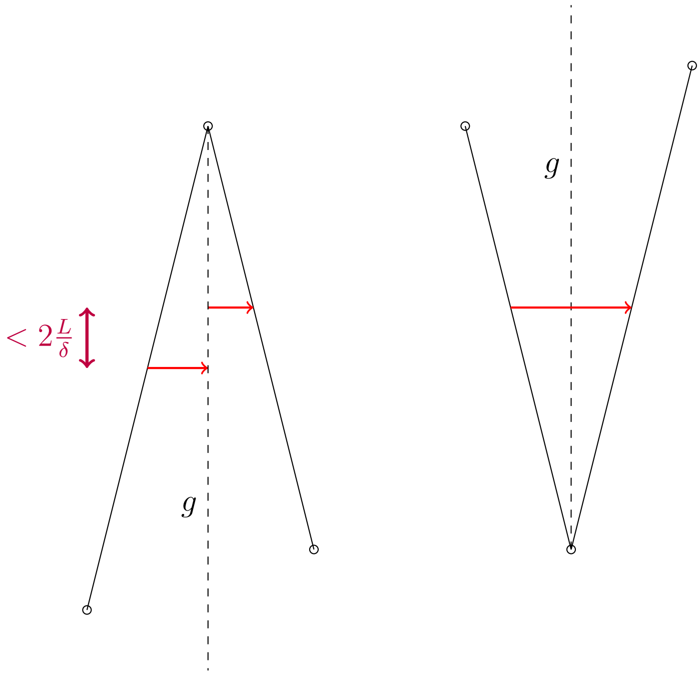

Figure 45: The path depicted by the red arrow has length of at least $\frac{L}{2}$, so it changes its height at maximum $\frac{4}{\delta}$ faster than its length.

is connected. The a_n and b_n are sorted in a way that a higher index indicates being closer to the corresponding accumulation point. At last $t(a_n^{(1)})$ $(t(b_n^{(1)}))$ should be closer to y (y') than $s(a_n^{(1)})$ $(s(b_n^{(1)}))$ for each $n \in \mathbb{N}$.

By $c_0(\mathbb{C})$ we denote the complex vector space of all null sequences in $\mathbb{C}$ together with the supremum norm $||(x_n)_{n\in\mathbb{N}}||_\infty = \sup\{|x_n| \mid n \in \mathbb{N}\}$.

Theorem 6.23. *Let $P \in \text{Co-Ch}_{n,\delta}(\gamma, v_1, v_2)$ and $\varepsilon \in \mathbb{R}_+$ be as above. Then the map*

$$\Psi\colon \widehat{B_\varepsilon(P)} \to \mathbb{C}^{n+1} \times c_0(\mathbb{C})^2,\ \widehat{P}^{(1)} \mapsto \left(\begin{pmatrix} s(x_1^{(1)}) \\ \vdots \\ s(x_n^{(1)}) \\ t(x_n^{(1)}) \end{pmatrix}, (t(a_n^{(1)}) - y)_{n\in\mathbb{N}}, (t(b_n^{(1)}) - y')_{n\in\mathbb{N}} \right)$$

is a homeomorphism onto its image.

Proof.
- well-defined: This follows in the same way as for Theorem 5.16 because no edge can cut through the area between the two converging sides.
- injective: This follows in the same way as before.
- continuous: Let $P^{(i)}, P^{(\infty)} \in B_\varepsilon(P)$ be PSI-polygons such that $(\widehat{P}^{(i)})_{i\in\mathbb{N}}$ converges to $\widehat{P}^{(\infty)}$. If we assume that $(\Psi(\widehat{P}^{(i)}))_{i\in\mathbb{N}}$ does not converge to $\Psi(\widehat{P}^{(\infty)})$ there is

either $j \in \{1, \dots, n\}$ such that $s(x_j^{(i)})$ does not converge to $s(x_j^{(\infty)})$ or the vertices which do not converge are located in the converging sides. In the first case, a contradiction is reached in the same way as in Theorem 5.16. In the second case, Lemma 6.22 allows us to reach a contradiction in a similar way to Theorem 5.9 by using π_v instead of the real coordinates and adjusting the set D to reach up to ∂P orthogonal from v.

- open: This follows directly from Theorem 6.16.

□

7 Reference of most commonly used objects and notations

For reference, we give a short list of the most commonly used objects/notations in this work in the order of appearance:

Object/Notation	Place to find
Chamanara surface	Example 2.3
$\mathcal{M}$	Definition 2.5
$\mathcal{E}$	Definition 2.6
Developing map dev_X	Definition 2.8
Immersion / $\rightsquigarrow$	Definition 3.1
Embedding / $\hookrightarrow$	Definition 3.7
$\mathcal{M}_{\rightsquigarrow}(D), \mathcal{M}_{\not\rightsquigarrow}(U), \mathcal{M}_{+}(D,U), \mathcal{M}_{-}(D,K)$	Definition 3.9
PSI-polygon P	Definition 4.2
Edge x_n	Definition 4.2
Vertex $s(x_n)/t(x_n)$	Definition 4.2
o_P	Convention 4.3
$x_n' = \Phi_P(x_n)$	Convention 4.3
$\widehat{P}$, $\widehat{x}$, $\widehat{\gamma}$	Convention 4.6
$\widetilde{P}$	Proposition 4.9
Convergence of PSI-polygons	Definition 4.11
N_ε	Definition 4.11
$\dot{\gamma}$	Remark 6.2
ε-similar	Definition 6.3
ε-approximation	Definition 6.4
$\Delta_\varepsilon(X)$	Definition 6.5
IP_ε	Definition 6.10

References

[Arn88] Pierre Arnoux. Un exemple de semi-conjugaison entre un échange d'intervalles et une translation sur le tore. *Bulletin de la Société Mathématique de France*, 116(4):489–500, 1988.

[AY81] Pierre Arnoux and Jean-Christophe Yoccoz. Construction de difféomorphismes pseudo-Anosov. *Comptes Rendus de l'Académie des Sciences. Série I*, 292, 1981.

[BBI01] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A Course in Metric Geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, 2001.

[Bow13] Joshua Bowman. The complete family of Arnoux-Yoccoz surfaces. *Geometriae Dedicata*, 164:113–130, 2013.

[CGL06] Reza Chamanara, Frederick Gardiner, and Nikola Lakic. A hyperelliptic realization of the horseshoe and baker maps. *Ergodic Theory and Dynamical Systems*, 26(6):1749 – 1768, 2006.

[Cha04] Reza Chamanara. Affine automorphism groups of surfaces of infinite type. In *In the tradition of Ahlfors and Bers, III*, volume 355 of *Contemporary Mathematics*, pages 123–145. American Mathematical Society, Providence, RI, 2004.

[DHL14] Vincent Delecroix, Pascal Hubert, and Samuel Lelièvre. Diffusion for the periodic wind-tree model. *Annales Scientifiques de l'École Normale Supérieure*, 47:1085–1110, 2014.

[EM18] Alex Eskin and Maryam Mirzakhani. Invariant and stationary measures for the $\mathrm{SL}(2,\mathbb{R})$ action on moduli space. *Publications mathématiques de l'IHÉS*, 127:95–324, 2018.

[EMM15] Alex Eskin, Maryam Mirzakhani, and Amir Mohammadi. Isolation, equidistribution, and orbit closures for the $\mathrm{SL}(2,\mathbb{R})$ action on moduli space. *Annals of Mathematics*, 182(2):673–721, 2015.

[FM14] Giovanni Forni and Carlos Matheus. Introduction to Teichmüller theory and its applications to dynamics of interval exchange transformations, flows on surfaces and billiards. *Journal of Modern Dynamics*, 8(3&4):271–436, 2014.

[Hoo13] W. Patrick Hooper. Immersions and the space of all translation structures. arXiv:1310.5193, 2013.

[Hoo18] W. Patrick Hooper. Immersions and translation structures I: The space of structures on the pointed disk. *Conformal Geometry and Dynamics*, 22:235–270, 2018.

[HR16] Frank Herrlich and Anja Randecker. Notes on the Veech group of the Chamanara surface. arXiv:1612.06877, 2016.

[HT19] W. Patrick Hooper and Rodrigo Treviño. Indiscriminate covers of infinite translation surfaces are innocent, not devious. *Ergodic Theory and Dynamical Systems*, 39(8):2071–2127, 2019.

[IT92] Yoichi Imayoshi and Masahiko Taniguchi. *An Introduction to Teichmüller Spaces*. Springer-Verlag, Tokyo, 1992.

[KZ03] Maxim Kontsevich and Anton Zorich. Connected components of the moduli spaces of Abelian differentials with prescribed singularities. *Inventiones Mathematicae*, 153(3):631–678, 2003.

[LT16] Kathryn Lindsey and Rodrigo Treviño. Infinite type flat surface models of ergodic systems. *Discrete and Continuous Dynamical Systems*, 36(10):5509–5553, 2016.

[Mas06] Howard Masur. Ergodic theory of translation surfaces. In *Handbook of Dynamical Systems*, volume 1B, page 527–547. Elsevier, Amsterdam, 2006.

[Mas22] Daniel Massart. A short introduction to translation surfaces, Veech surfaces, and Teichmüller dynamics. In A. Papadopoulos, editor, *Surveys in Geometry I*. Springer Nature, 2022. to appear.

[Ran16] Anja Randecker. *Geometry and topology of wild translation surfaces*. PhD thesis, Karlsruher Institut für Technologie (KIT), 2016.

[Thu97] William P. Thurston. *Three-Dimensional Geometry and Topology, Volume 1*. Princeton University Press, 1997.

[Tre14] Rodrigo Treviño. On the ergodicity of flat surfaces of finite area. *Geometric and Functional Analysis*, 24(1):360–386, 2014.

[Tre18] Rodrigo Treviño. Flat surfaces, Bratteli diagrams and unique ergodicity à la Masur. *Israel Journal of Mathematics*, 225(1):35–70, 2018.

[Vee86] William A. Veech. The Teichmüller geodesic flow. *Annals of Mathematics*, 124(3):441–530, 1986.

[Wri15] Alex Wright. Translation surfaces and their orbit closures: An introduction for a broad audience. *EMS Surveys in Mathematical Sciences*, 2(1):63–108, 2015.

[Yoc10] Jean-Christophe Yoccoz. Interval exchange maps and translation surfaces. In *Homogeneous flows, moduli spaces and arithmetic*, volume 10 of *Clay Mathematics Proceedings*, pages 1–69. American Mathematical Society, Providence, RI, 2010.

[Zor06] Anton Zorich. Flat surfaces. In *Frontiers In Number Theory Physics And Geometry I: On Random Matrices, Zeta Functions, and Dynamical Systems*, pages 439–585. Springer, Berlin, 2006.

Acknowledgement

Zum Abschluss möchte ich mich bei allen Menschen bedanken, die mich bei der Erstellung dieser Arbeit unterstützt haben.

Mein größter Dank gilt dabei Frank Herrlich, der, trotz seiner vielen anderen Verpflichtungen, immer Zeit für mich hatte, wenn ich ihn brauchte. Seine freundliche, humorvolle und offene Art waren stets eine Bereicherung, sei es beim Gang in die Mensa, der Diskussion über mathematische Themen oder dem Planen des nächsten Übungsblattes.
Außerdem möchte ich Gabriela Weitze-Schmithüsen dafür danken, dass sie sofort bereit war, Referentin für diese Arbeit zu sein.

Mein Dank gilt außerdem allen, die Teile meiner Arbeit gelesen haben: Leonid Grau für die viele Zeit, die er dort hinein investiert hat und die vielen hilfreichen Anmerkungen, die dabei herauskamen. Kevin Klinge für das Hinweisen auf viele unklare Stellen. Barbara Matthes, der in kürzester Zeit alle groben Fehler aufgefallen sind. Stephan Stock für das Hinzufügen zahlreicher Kommata und die Tipps, welche damit einhergingen. Und nicht zuletzt Vanessa Jeske für all die sprachlichen Verbesserungsvorschläge.

Zuletzt möchte ich noch allen Mitgliedern der Kaffeerunde danken, durch die sich die Arbeitsgruppe wie ein zweites Zuhause angefühlt hat und die bei allen Anliegen stets hilfsbereit zur Stelle waren.