On the Coarse Geometry of Infinite Regular Translation Surfaces

Zur Erlangung des akademischen Grades eines

DOKTORS DER NATURWISSENSCHAFTEN

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

DISSERTATION

von

Christoph Karg, M.Sc.
aus Ludwigshafen a. Rh.

Tag der mündlichen Prüfung: 4. März 2020

Referenten: Prof. Dr. Gabriela Weitze-Schmithüsen
Prof. Dr. Frank Herrlich
Dr. Daniele Alessandrini

Acknowledgements

This PhD thesis was written during my time as a research assistant both at the Institute for Algebra and Geometry at Karlsruhe Institute of Technology and the Institute for Mathematics at the University of Heidelberg, financed in parts by the German Research Foundation DFG (RTG 2229 - "Asymptotic Invariants and Limits of Groups and Spaces") and the Institute for Algebra and Geometry.

As a wanderer between three worlds my sincere thanks go to my advisors in Saarbrücken, Karlsruhe and Heidelberg. First and foremost I warmly thank my main supervisor *Gabi Weitze-Schmithüsen* for making my time as her PhD student so pleasant and worthwhile. It was her visible motivation and the joyful atmosphere in her lectures that sparked my interest in geometric group theory and translation surfaces. During her supervision of my PhD thesis I appreciated her sense of humour, her encouragement to pursue my own ideas and her suggestions to attend various conferences. I want to thank my supervisor *Frank Herrlich* for his patient support and friendly advice, and for always finding me a room to work in. Also, I thank *Daniele Alessandrini* for his patient explanations and for sharing his wealth of knowledge about a variety of mathematical fields. Last but not least I want to thank *Anna Wienhard* for making possible my travels to conferences, without which this thesis would not have been possible. Through this, she gave me the opportunity to participate in an international network of mathematicians and gain important insights into academic life.

Moreover, I thank *Kasra Rafi* for inviting me to the University of Toronto for a two-month stay which permitted me to make new experiences, in- and outside of mathematics.

Shy as I am when it comes to talking about mathematics I thank *Daniele Alessandrini*, *Anja Randecker*, *Federica Fanoni* and *Beatrice Pozetti* a lot for "forcing" me to formulate my thoughts and listening patiently to my attempts. And I owe thanks to *Andy Sanders* for bringing groove and colour into the often gray academic rooms.

Furhermore, I want to thank all members of the Institute for Algebra and Geometry for making everyday life at work so pleasant with their various jokes, word plays and riddles.

The author acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 281869850 (RTG 2229).

Contents

CHAPTER 1

Introduction

The present thesis deals with translation surfaces which form a modern field of research compared to the history of mathematics. Since their appearance in [FK36] in the early 20th century where Fox and Kershner studied billiard trajectories on polygonal tables a lot of progress has been made in understanding the geometry of translation surfaces. In the course of this research many connections to various mathematical fields have been discovered: From the first study of polygonal billiards evolved a much broader interest in dynamical systems on translation surfaces, and furthermore relations to Teichmüller theory, algebraic geometry and - as it is the case for this thesis - geometric group theory. For a wide overview to translation surfaces and their applications see [Zor06], [HS06] and [Möl09]. The relation between billiards and translation surfaces is studied in [KMS86], [KS00] and [MT02]. In [Yoc10] translation surfaces are used to study the dynamics of interval exchange maps. In [Möl13] translation surfaces are studied from the perspective of algebraic geometry. And for an introduction to the Teichmüller theory of translation surfaces we refer to [Vee86], [Vee89] and [Wri15].

Apart from their importance in the mentioned research fields what makes translation surfaces so appealing is their very graphic nature. By definition, they are surfaces constructed from a finite or infinite collection of Euclidean polygons which are glued along parallel sides of the same length. This makes the geometry of translation surfaces to most parts quite approachable as polygons in the Euclidean plane are very familiar mathematical objects. Depending on whether the collection of polygons is finite or infinite we speak of *finite* and *infinite* translation surfaces. The study of finite translation surfaces has progressed quite far in the past years. And although there still are many open problems concerning finite translation surfaces, the interest in all types of infinite translation surfaces has grown ever more. This thesis attempts a contribution to shed light on the still mostly dark universe of infinite translation surfaces.

When trying to describe the global geometry of a translation surface we must first observe its local behaviour which is ambivalent in the following sense. On one hand a translation surface, apart from a set of *singular points*, is locally isometric to open subsets of the Euclidean plane. As this is a very familiar geometry we can use common tools and notions from Euclidean geometry as “straight lines”, “orientation”, “parallelity” and “directions” as “north” or “left”. This then allows us to locally study important objects as geodesics, metric discs and compact sets. On the other hand open disks around singular points may look very different from Euclidean open disks which makes the whole geometry of translation surfaces all the more interesting. As the curvature of a translation surface is concentrated in the singularities we have that the total angle sum around such a singularity is a natural multiple of 2π or even infinite, compared to regular points of the surface having an angle sum equal to 2π. This fact has interesting consequences: While there is exactly one geodesic ray starting in a regular point in any direction of S^1, there usually are more than one geodesic rays starting in a singularity in this same direction, see Figure 1.1.

Compared to infinite translation surfaces the local geometry of finite translation surfaces is rather easy to comprehend. They are compact surfaces having only finitely many singularities. Each singularity is *conical*, i.e. its total angle sum is finite and of the form $2\pi k$ for a natural k which is called the *multiplicity* of the singularity. Each disc of sufficiently small radius around such a conical singularity admits a translation covering of degree k onto a Euclidean disc of same radius and which is ramified over its center, see Figure 1.1. However, for infinite translation surfaces more types of singularities can appear. Firstly, there are ∞-*angle* singularities which, similar to conical singularities, admit a translation covering map of infinite degree from any sufficiently small disc around them onto a corresponding Euclidean disc. This existence of such a covering map is the reason why we classify conical and ∞-angle singularities as *tame* singularities: Although discs around such singularities are not isometric to common Euclidean discs they behave well enough for us to illustrate them using Euclidean discs glued together along parallel slits. Apart from these tame singularities a whole variety of *wild*

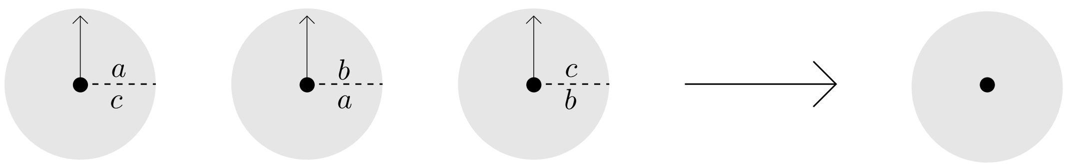

Figure 1.1: A disc around a singularity of multiplicity $k = 3$ together with a 3-sheeted ramified covering onto a Euclidean disc. Note that there are three distinct geodesics starting in that singularity in north direction.

singularities may occur on infinite translation surfaces. As the name suggests wild singularities are those singularities whose neighborhoods cannot be described as concisely as the ones of tame singularities. There is a multitude of infinite translation surfaces whose wild singularities exhibit different and astounding properties. See [BV13], [Ran16] and [Ran18] for a detailed introduction to wild translation surfaces.

This variety of singularity types complicates the study on the whole family of infinite translation surfaces. There rarely are insightful methods that can be applied to all types of infinite translation surfaces at once. A very reasonable restriction which often is included in the definition of translation surfaces is to require the discreteness of the subset of singular points. Hereby one discards surfaces such as open discs of the Euclidean plane which are translation surfaces by the general definition but do not have as interesting properties as translation surfaces with discrete singularities. But still in that case, often one has to restrict to special subfamilies and apply tools specifically designed for that kind of surface. Such subfamilies are for example infinite translation surfaces of finite area as the *Chamanara surface* ([Cha04]), translation surfaces constructed from rectangles with varying dimensions as the *stack of boxes* example ([Bow12] and [Ran16]), or translation surfaces obtained from inserting slits into the Eudlidean plan and gluing them back together in various ways as in [CRW19] and [Ran16]. Another class of translation surfaces - the one which we are interested in - consists of those surfaces constructed from infinitely many copies of one base polygon which are glued together in a regular way.

Considering this, it is clear that if we want to explore a notion as broad as "the geometry" of infinite translation surfaces we have to restrict ourselves to a suitable subfamily as well. But which? Usually, when it comes to exploring a great number of complicated objects it is very fruitful to begin with those objects which exhibit symmetries. A symmetrical object has the advantage that it can be understood as a whole often by studying a smaller, more approachable part whose copies can be arranged to form the original object. If, as in our case, this object is an infinite translation surface it makes sense to describe its symmetries using a suitable group action on the surface. And the role of the smaller, more approachable part should be played by a finite translation surface which we understand fairly well compared to the infinite ones. This observation suggests that we should define our subfamily of infinite translation surfaces which we want to study in detail as follows: A *regular* translation surface X as we will define it more precisely in Section 3.1 shall be a translation surface which permits a regular translation covering onto a finite translation surface X_0 and may be ramified at most over the singularities of the base surface X_0. We speak of G-regular translation surfaces when we require the deck transformation group of the corresponding covering to be isomorphic to G. A great advantage of this definition is that we can describe the symmetries of X with the group G, which allows us to use group theoretical methods in order to study the translation surface. Furthermore it allows us to deduce many properties of the infinite surface X from the base surface X_0 which is finite and hence easier to understand. However, this benefit has its price: As we will see all regular translation surfaces are tame. This means that with our approach a large part of the universe of infinite translation surfaces, namely the wild ones, is left unexplored.

Having found a reasonable family of translation surfaces to focus on it is now necessary to clarify what is meant by studying their "global geometry". Among the many ways to describe the geometry of Riemannian surfaces one common approach is to describe its *geodesics*, i.e. paths of shortest length between two arbitrary points and - closely related - to compute the distance between any two points of that space. Although geodesic segments in translation surfaces consist of common Euclidean line segments, in practice the exact computation of distances proves to be quite difficult. The reason is that the singularities lead to sudden directional changes of the geodesic which in general are hard to follow, see Figure 1.2. A

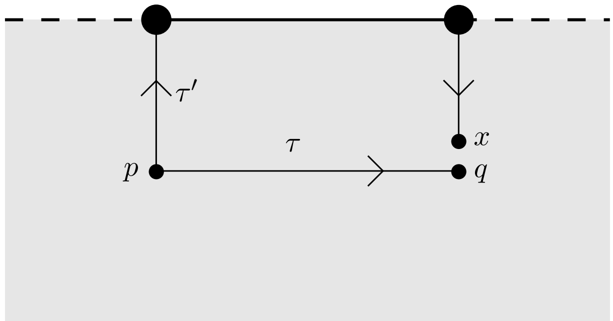

Figure 1.2: The Euclidean line segment τ between p and q is a geodesic segment. If we consider the geodesic segment τ' between p and the point x which may be arbitrarily close to q then the singularity acts as a shortcut.

good way to avoid this problem of computing exact distances is to study the *coarse geometry* of translation surfaces which only gives reasonable lower and upper bounds for the distance between two arbitrary points. This is usually done by finding a map, called *quasi-isometry*, from the translation surface to a metric space whose geometry is more familiar. In this way distances in the translation surface can be related to distances in the familiar metric space which gives us information about the rough global shape of the translation surface. Illustratively spoken two quasi-isometric spaces roughly have a similar shape “when looked at from a large distance”. For example the inclusion $\mathbb{Z} \hookrightarrow \mathbb{R}$ is a quasi-isometry while it is not possible to find a quasi-isometry from $\mathbb{R}$ to $\mathbb{R}^2$: The line $\mathbb{R}$ is simply too “thin” to behave metrically similar to the plane, see figure 1.3. When studying the coarse geometry of translation surfaces all finite translation surfaces become suddenly indistinguishable: Indeed, all bounded metric spaces are quasi-isometric to a one point set and therefore uninteresting from this perspective.

There are metric spaces whose coarse geometry can be described by means of group theory. For this, we usually consider finitely generated groups and turn a group itself into a metric space using the *word metric* with respect to a finite generating set. An important result in geometric group theory is the Švarc-Milnor Lemma, see [BH99], Proposition 8.19. It states in particular that if a finitely generated group G acts properly discontinuously on a proper metric space X such that the corresponding quotient space X/G is compact, then the metric space is quasi-isometric to the group equipped with the word metric corresponding to any finite generating set. As an example, the *Cayley graph* $\mathrm{Cay}(G, S)$ is a metric graph constructed from a group G, whose elements represent the vertices, and a finite generating set S. The G-left multiplication on this graph satisfies the conditions of the Švarc-Milnor Lemma and therefore both the graph and the group are quasi-isometric.

Similar to Cayley graphs we can construct any G-regular translation surface X from the finitely generated group G and a finite base translation surface. And as for Cayley graphs we have a close relation between the coarse geometry of the regular translation surface and its

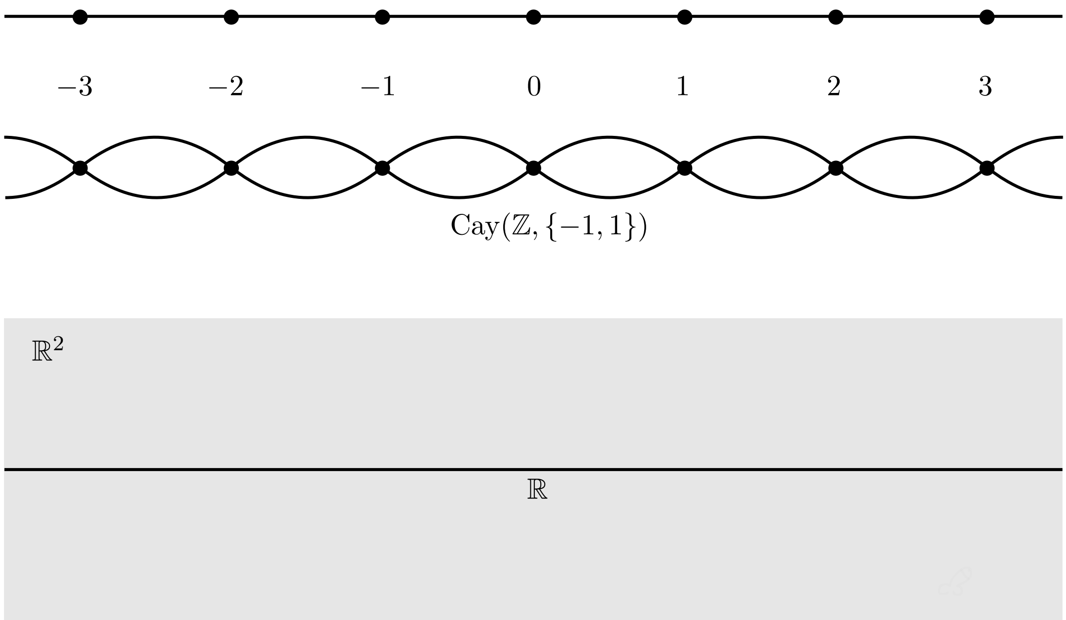

Figure 1.3: The Euclidean line $\mathbb{R}$ is quasi-isometric to $\mathbb{Z}$ and to the Cayley graph $\mathrm{Cay}(\mathbb{Z}, \{-1, 1\})$. However, $\mathbb{R}$ is not quasi-isometric to the Euclidean plane $\mathbb{R}^2$, as the plane contains points which are arbitrarily far from the line.

symmetry group. However, in general a G-regular translation surface will not be quasi-isometric to the group G. Illustratively speaking, the reason is that possible ∞-angle singularities of X lead to "shortcuts" in the geometry of the surface X which are not present in the group structure and hence geometry of G.

In this thesis we present methods to determine the coarse geometry of infinite regular translation surfaces. As an example we introduce *monodromy elements* in the group G that describe those "shortcuts" in the G-regular surface X. It turns out that in many cases the coarse geometry of X is determined by quotients of the group G, however this does not hold in general. For the general situation we construct a graph from the group G and the monodromy elements and prove that it is quasi-isometric to X.

This thesis is structured as follows. In Chapter 2 we present all the necessary definitions and tools that are needed in order to study infinite regular translation surfaces. General translation surfaces and affine maps are introduced in Section 2.1. An important example of such affine maps are translation coverings. In Section 2.2 we define singularities of translation surfaces and discuss their types and properties. In particular we are interested in tame singularities. We recall the metric on translation surfaces and study geodesics in more detail in Section 2.3. An important family of translation surfaces, the finite ones, is presented in Section 2.4. They are of high interest to us as many properties of infinite regular translation surfaces can be studied using properties of finite translation surfaces. In Section 2.5 we introduce singular loops in finite translation surfaces which are a useful tool in the study of regular translation surfaces. As we are interested in regular coverings of finite translation surfaces Section 2.6 provides all the necessary background from covering theory. Also, in order to study the coarse geometry of infinite regular translation surfaces we define quasi-isometries in Section 2.7, and give a short overview of the necessary results in geometric group theory. Furthermore, in Section 2.8 we give the definition of an end of a topological space and recall some properties of ends of groups.

In Chapter 3 we study the main objects of our interest: regular translation surfaces. As we will see regular translation surfaces that are infinite are highly symmetric and behave relatively well compared to other generic infinite translation surfaces. We define finite and infinite regular translation surfaces in Section 3.1. They permit a regular translation covering onto a finite translation surface, their base surface. We show that this implies that regular translation surfaces are always tame. Using the constructive characterization of translation surfaces we describe in Section 3.2 how to construct a regular translation surface from a finite translation surface and a finitely generated group. In Section 3.3 we prove that every regular translation surface can be constructed in this fashion. This turns out to be very helpful since we can describe quite complicated regular translation surfaces X by studying the more approachable data, namely the finite base surface and the deck transformation group G of the corresponding regular covering, by means of Euclidean geometry and group theory. In that case we speak of G-regular translation surfaces. A first use of this method is presented in Section 3.4. Here we use singular loops in order to describe the singularities of regular translation surfaces. In particular we characterize ∞-angle singularities using the order of certain elements in the corresponding group G. We then describe the whole set of singularities of a regular translation surface X by suitable cosets of G. This is a first indication that geometry and group theory are closely related for regular translation surfaces. As a first result we prove the following remarkable relation between geometry and group theory.

Theorem 1.1. *Let X be an infinite regular translation surface having at least one ∞-angle singularity. Then X has finitely many ∞-angle singularities if and only if G is virtually $\mathbb{Z}$.*

In Section 3.5 we study quotient spaces of regular translation surfaces using intermediate translation coverings which correspond to subgroups of G. Finally, in Section 3.6 we present important examples for infinite regular translation surfaces: The 2- and 3-staircase and the AB-surface. Although the 2- and 3-staircase are both $\mathbb{Z}$-regular and look highly similar from the constructive perspective we prove that their coarse geometry is very different: The 2-staircase is bounded whereas the 3-staircase is quasi-isometric to $\mathbb{Z}$. While in these two examples the quasi-isometry class seems to be closely related to the finitely generated deck transformation group G, the AB-surface exhibits a more complicated coarse geometry. Indeed we show that this surface is quasi-isometric to the countably infinite regular tree T_∞.

This observation leads us to the natural question whether two regular translation surfaces are quasi-isometric or not. More precisely, we formulate the question as follows:

Question 1.2. *Given a G-regular translation surface X, can we explicitly describe a graph which is quasi-isometric to X?*

Our attempt to answer this question is presented in Chapter 4. Firstly, we focus on special cases of G-regular translation surfaces before we answer this question for general regular translation surfaces. The simplest case is dealt with in Section 4.1. Here we consider G-regular translation surfaces having only conical singularities. Using the Švarc-Milnor-Lemma we get the following result.

Theorem 1.3. *Let X be a G-regular translation surface having only conical singularities. Then X is quasi-isometric to G.*

In Section 4.2 we consider G-regular translation surfaces where G is a boundedly generated group. In this case we can show that under moderate conditions a G-regular translation surface is quasi-isometric to a quotient group G/U for a normal subgroup U of G which depends on the covering. In particular this condition holds when G is abelian.

Theorem 1.4. *Let X be a G-regular translation surface where G is abelian. Then X is quasi-isometric to a quotient group of G.*

The study of the coarse geometry of general translation surfaces happens in Section 4.3. After a detailed observation on the length of geodesic segments in regular translation surfaces we formulate the main result of Section 4.

Theorem 1.5. *Let X be a G-regular translation surface. Then X is quasi-isometric to* $\mathrm{Cay}(G, T^\infty)$*, the Cayley graph of G with respect to an infinite generating system T^∞ of G.*

We see that the answer to the previously formulated question is "Yes". However, Theorem 1.5 does not tell us whether two given regular translation surfaces are quasi-isometric. It just reduces this problem to finding a quasi-isometry between corresponding intricate, locally infinite graphs, which in itself is not easy to solve. We conclude the study of the coarse geometry of regular translation surfaces in Section 4.4 where we describe the ends of a G-regular translation surface X. Theorem 1.6 tells us that each end of X comes from an end of G, i.e. of its Cayley graph with respect to a finite generating system.

Theorem 1.6. *There is a surjective map*

$$\mathrm{Ends}(G) \twoheadrightarrow \mathrm{Ends}(X).$$

As a side result of this theorem we obtain the following statement: Any $\mathbb{Z}$-regular translation surface containing at least one ∞-angle singularity is bounded and has exactly one end. When studying G-regular translation surface with G having infinitely many ends, its space of ends might however still be very complicated to describe as shows an example where G is the free group on two generators.

In Chapter 5 we present two applications of our results. We generalize the AB-surface to a family of regular translation surfaces having as deck transformation group the free group on n generators for $n \geq 2$. In the case $n = 2$ we obtain the original AB-surface. Using Theorem 1.5 we are able to prove that all translation surfaces of this family are quasi-isometric to T_∞.
As a second application we consider the Teichmüller space of translation structures on a closed topological surface. Using Theorem 1.5 we can show that the quasi-isometry class of a regular

translation surface does not depend on the choice of translation structure on the corresponding base surface. Illustratively speaking, this means that a small variation of the shape of the base polygon does not change the coarse geometry of the corresponding regular translation surface.

CHAPTER 2

Background

2.1 Translation surfaces

There are three equivalent possibilities to define translation surfaces. The first uses Euclidean polygons that are glued along parallel edges, the second uses translation structures and the third definition uses the notion of abelian differentials on a surface. For our purpose we will only need the first two and throughout this paper we mostly work with the first and most illustrative definition. For this reason we will put our focus on the first one. For a precise description of the other definitions and a detailed proof that all three are equivalent we refer to the book in progress by Valdez and Delecroix, see [DV], Chapter 1.

Consider an at most countable family $\mathcal{P}$ of polygons in the Euclidean plane. Here, a polygon is a simply connected and compact set whose boundary is a closed curve consisting of finitely many straight line segments, its edges. Let $E(\mathcal{P})$ be the set of all the edges in $\mathcal{P}$. Fix an orientation of the plane which induces an orientation on each polygon in $\mathcal{P}$ and hence on the edges of each polygon in $\mathcal{P}$. Suppose there is a map $gl : E(\mathcal{P}) \to E(\mathcal{P})$, called *gluing map*, where each edge $e \in E(\mathcal{P})$ is paired to a unique distinct edge $gl(e) \in E(\mathcal{P})$ such that $gl(e)$ and e differ by a translation and have opposite orientation. We also say that e and $gl(e)$ are *paired* edges. From the family $\mathcal{P}$ and the gluing map gl we construct a topological space $\overline{X}$ as follows: Consider the disjoint union $\bigsqcup_{P\in\mathcal{P}} P$ of all polygons and identify points on the edges using the translation map given by gl. The resulting quotient space is

$$\overline{X} := \bigsqcup_{P\in\mathcal{P}} P/ \sim_{gl},$$

obtained from 'gluing' all polygons in $\mathcal{P}$ along the paired edges. We have a natural quotient map $\pi : \bigsqcup_{P \in \mathcal{P}} P \to \overline{X}$ which is 1-to-1 in the interior of each polygon and 2-to-1 on the edges without their endpoints. A vertex v of a polygon $P \in \mathcal{P}$ is *of finite degree* if $\pi^{-1}(\pi(v))$ is finite and *of infinite degree* otherwise.

Definition 2.1. Let $\mathcal{P}$ and $gl : E(\mathcal{P}) \to E(\mathcal{P})$ as before. Let X be $\overline{X}$ without all vertices of infinite degree. If X is connected we call it *the (constructive) translation surface obtained from the family of polygons* $\mathcal{P}$. If the collection $\mathcal{P}$ is finite we call X a *finite* otherwise an *infinite translation surface*. Note that if X is a finite translation surface then $X = \overline{X}$. The set of all vertices of the family $\mathcal{P}$ after gluing is called the *singularities* of X and is denoted by $\mathrm{Sing}(X)$. We then denote by $X^* := \overline{X} \backslash \mathrm{Sing}(X)$ the punctured translation surface X without all vertices of the family $\mathcal{P}$. We equip X^* with the flat metric d obtained by extending the local Euclidean metric on each polygon in $\mathcal{P}$ and extend it to the metric d on $\overline{X}$.

Remark. In order to avoid confusion in the notation we remark here that $\overline{X}$ does not stand for the completion of the space X. In general $\overline{X}$ is not complete, for example if $\overline{X}$ consists of an infinite family of polygons whose diameter converges to zero. However, in the case of regular translation surfaces it indeed turns out that $\overline{X}$ is the metric completion of X which is why we stick with this notation.

It might not be clear at first why we do not define the whole quotient space $\overline{X}$ to be the translation surface, especially since in our definition above we allow singularities of X to lie 'outside' of X. However, while $X \subseteq \overline{X}$ indeed is a topological surface the constructed space $\overline{X}$ in general is not. Consider for example a vertex of infinite degree whose total angle sum is infinite, as shown on the right-hand side of Figure 2.2. Then the closed disc of any radius $r > 0$ centered in this vertex is not compact and hence $\overline{X}$ is not a locally compact space and in particular not a surface. We will see later that both spaces, X and $\overline{X}$, have their advantages. On the one hand X is a surface and hence is 'behaving well' as a topological space. On the other hand in the case of regular translation surfaces $\overline{X}$ is a complete metric space which allows us to study geodesics and gain information about the geometry of the space.

Consider any surface S equipped with a flat metric. We say that a point $\sigma \in S$ is a *conical point of angle* $2\pi\alpha$ for the flat metric if there exists an open neighborhood U of σ and a real number $\alpha > 0$ such that $U \backslash \{\sigma\}$ is isometric to $\mathbb{C} \backslash \{0\}$ equipped with the metric $(dr)^2 + (\alpha r d\theta)^2$ with respect to the polar coordinates (r, θ) of the Euclidean plane. As the name suggests it follows from this definition that a disc of sufficiently small radius ε around σ has the circumference $2\pi\alpha\varepsilon$. For the second definition let S be a connected topological surface. A *translation atlas* on S is a set of maps $\mathcal{T} := \{\varphi_i : U_i \to \mathbb{C}\}$ where $(U_i)_{i \in \mathbb{N}}$ is an open cover of S, each φ_i is a homeomorphism from U_i to $\varphi_i(U_i)$ and for each i, j the transition map $\varphi_j \circ \varphi_i^{-1} : \varphi_i(U \cap V) \to \varphi_j(U \cap V)$ locally is a translation in $\mathbb{C}$. The surface S together with such a translation atlas is as well naturally endowed with a flat metric simply by pulling back the Euclidean metric in $\mathbb{C}$.

Definition 2.2. Equivalently, a *(geometric) translation surface* is a tuple $X = (S, \Sigma, \mathcal{T})$ consisting of a connected topological surface S, a discrete subset $\Sigma \subset S$ and a maximal translation atlas $\mathcal{T}$ on $S\backslash\Sigma$ such that every $\sigma \in \Sigma$ is a conical point for the induced flat metric. Here, the subset Σ corresponds to the set of vertices of finite degree in the definition of constructive translation surfaces.

Let $X = (S, \Sigma, \mathcal{T})$ and $X' = (S', \Sigma', \mathcal{T}')$ be two (geometric) translation surfaces. An *affine map* between X and X' is a homeomorphism $f : S \to S'$ satisfying $f(\Sigma) = \Sigma'$ and which is of the form $x \mapsto Ax + b$ in local coordinates on $S\backslash\Sigma$, where $A \in \mathrm{GL}_2(\mathbb{R})$ and $b \in \mathbb{R}^2$. Note that the differential $Df = A$ of f is independent from the choice of charts. If Df is the unit matrix we call f a *translation* between X and X'. When the two translation surfaces are given as a glued collection of polygons then X and X' being equivalent means that we can obtain X' by cutting X along straight lines and gluing back along paired edges. In Figure 2.1 we see an affine map which is locally given by the linear transformation where $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$. A

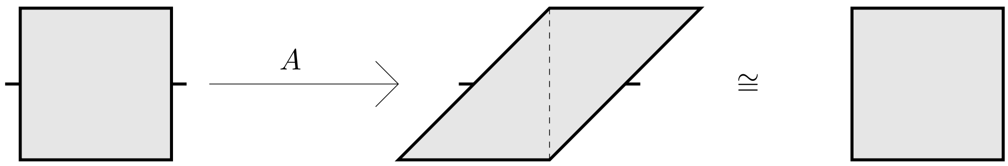

Figure 2.1: An affine map between the flat square torus and a translation surface obtained from a paralellogram which is equivalent to it.

translation covering from X onto X' is a topological covering map $p : S \to S'$ that is of the form $x \mapsto x + b$ in local coordinates, where $b \in \mathbb{R}^2$. A translation covering is called *cyclic* if its deck transformation group is cyclic, and hence isomorphic to $\mathbb{Z}$ or $\mathbb{Z}/d\mathbb{Z}$ for $d \in \mathbb{N}$. The number d is called the *degree* of the covering.

2.2 Singularities

Let X be a translation surface obtained from a family of polygons $\mathcal{P}$. As we have seen we can extend the natural flat metric on X^* to a metric on the space $\overline{X}$. In this way it makes sense to talk about open and closed discs of a given radius around any point in $\overline{X}$. Discs of a sufficiently small radius ε around any point in X^* are isometric to Euclidean ε-discs. However, discs around singularities behave differently. A singularity $\sigma \in \mathrm{Sing}(X)$ is called *tame* if there exists an $\varepsilon > 0$ and a cyclic translation covering

$$p : D_\varepsilon(\sigma)\backslash\{\sigma\} \to D_\varepsilon(0)\backslash\{0\}$$

from the punctured open ε-disc in $\overline{X}$ around σ to the punctured open ε-disc in $\mathbb{R}^2$. Otherwise the singularity is called a *wild* singularity. Similarly, we call the translation surface *tame* if each of its singularities is tame, and we call it *wild* in all other cases. A detailed classification of singularities in infinite translation surfaces can be found in [BV13] and [Ran16].

Let σ be a tame singularity. The corresponding cyclic translation covering has finite degree d if and only if σ is a conical point of angle $2\pi d$. We therefore call σ a *conical singularity* and define its *multiplicity* to be $k_\sigma := d - 1$. The reason for this definition is that conical points of angle $2\pi d$ are exactly the zeros of abelian differentials on X of order $d-1$, see [Str84], Chapter III, for a detailed explanation. In terms of the flat metric this means that the total angle sum around a conical singularity with multiplicity $d-1$ is exactly $2\pi d$. If the corresponding translation covering is infinitely cyclic we say that σ is an *infinite angle singularity* or short ∞*-angle singularity.*

Example. One classic example for a finite translation surface is given by a glued polygon as shown on the left-hand side in Figure 2.2. The resulting surface has genus 2 and all vertices of the polygon are identified to one singularity. Note that the total angle sum around this singularity is 6π. Hence it has one conical singularity with multiplicity $k = 2$.

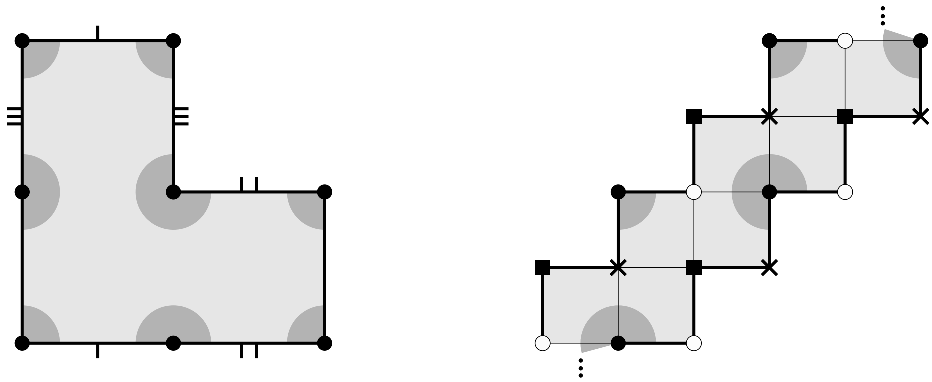

Figure 2.2: Two examples for tame translation surfaces.

On the right-hand side we see an important example, the *infinite* 2*-staircase.* It consists of infinitely many glued squares where opposite sides are identified using a horizontal resp. vertical translation. After identficiation there remain four distinct singularities which are all ∞-angle. Hence the resulting translation surface X is the complete quotient space without those four singularities. In particular, we see here that any closed disc in $\overline{X}$ around a singularity is non-compact.

For a closer study of the infinite staircase and its dynamical properties we refer to [HHW13].

2.3 GEODESICS

In this section we present the geometric properties of translation surfaces and are particularly interested in the behaviour of geodesics in such surfaces. For a detailed examination of the geometry of flat metrics, see [Dan10].

Let X be a translation surface and consider the space $\overline{X}$ with all singularities included. A *geodesic* in $\overline{X}$ is a path γ from an interval into $\overline{X}$ that is locally isometric. Given two points $x, y \in \overline{X}$ a *geodesic arc* between x and y is a geodesic $\gamma : [0, L] \to \overline{X}$ satisfying $L = d(x, y)$. We call its image $[x, y] := \gamma([0, L]) \subset \overline{X}$ a *geodesic segment* between x and y. A geodesic segment $[x, y]$ satisfying $[x, y] \cap \operatorname{Sing}(X) = \{x, y\}$ is called a *saddle connection.*

Let $z \in \overline{X}$ be a point on a geodesic segment which is not an endpoint. If $z \in X^*$ then the geodesic segment in a small neighborhood of z is a Euclidean line segment since X^* is locally isometric to the Euclidean plane. If $z \in \mathrm{Sing}(X)$ is a tame singularity then in a small neighborhood of z the geodesic segment consists of two Euclidean line segments starting in z such that the angle at z between both line segments is not less than π. In particular if z is a conical singularity of multiplicity 0, i.e. with the regular angle sum 2π, the geodesic segment through z locally is just a Euclidean line. However if z has multiplicity ≥ 1 or is ∞-angle it is possible to have a geodesic segment having an angle greater than π at a singularity. In particular, given a geodesic ending in such a singularity, there are uncountably many ways to extend this geodesic as long as the angle remains $\geq \pi$. In Figure 2.3 we see an example of a geodesic entering a singularity with multiplicity $k = 2$ and three possible ways of extending this geodesic.

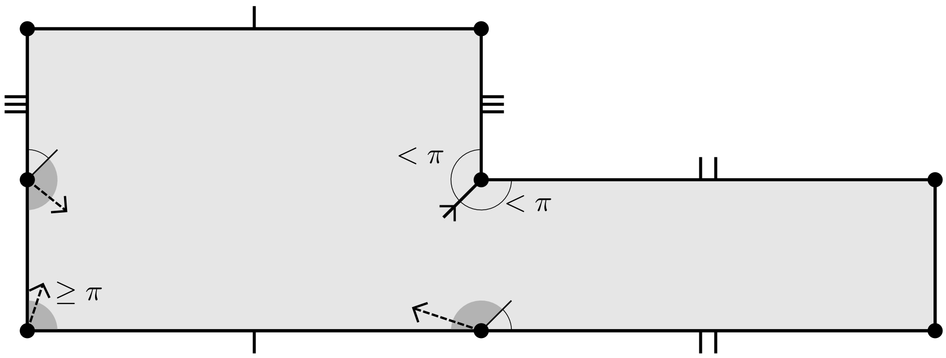

Figure 2.3: Three of infinitely many possible geodesics crossing a conical singularity.

The following lemma presents a sufficient condition for recognizing tame translation surfaces. Using this we will prove in Section 3.1 that all regular translation surfaces are tame.

Lemma 2.3. *Let X be a translation surface such that the length of each saddle connection in $\overline{X}$ is uniformly bounded below by a positive number L. Then X is tame. Furthermore any geodesic segment between two singularities consists of finitely many saddle connections.*

Proof. As the length of any saddle connection in $\overline{X}$ is bounded below by $L > 0$ each disc of radius $L/3$ around a singularity σ does not contain a second singularity. Hence, we can for example choose $L/3$ as uniform radius around each singularity to construct a cyclic covering of the corresponding punctured disc in Euclidean space. This proves that X is tame. Now consider a geodesic segment $[\sigma, \sigma']$ between two singularities σ, σ'. It has the length $d := d(\sigma, \sigma')$. As the length of each saddle connection is bounded below by $L > 0$ the number N of saddle connections in the segment $[\sigma, \sigma']$ satisfies the relation $d \geq L \cdot N$, and is thus finite. □

2.4 Finite translation surfaces

Finite translation surfaces form one large class of tame translation surfaces. By definition a finite translation surface X is obtained from a finite collection $\mathcal{P}$ of Euclidean polygons together with a gluing map $gl : E(\mathcal{P}) \to E(\mathcal{P})$. Since $\mathcal{P}$ is finite there are only finitely many vertices and the angle sum around each vertex is finite. As a result $X = \overline{X}$ is a closed, tame translation surface with finitely many conical singularities $\sigma_1, \dots, \sigma_n$ having multiplicities $k_1, \dots, k_n$. For finite translation surfaces we often shortly write Σ instead of $\mathrm{Sing}(X)$ and we say that X *has singularity type* $(k_1, \dots, k_n)$. Using a variant of the Gauss-Bonnet formula we have the following relation between the multiplicities and the genus g of the surface X:

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

In the following it will be much more convenient for us to think of a finite translation surface X as only one Euclidean polygon $P \subset \mathbb{R}^2$ glued along paired parallel edges of same length and different orientation. It is always possible to find such a polygon as shows the *zippered rectangle construction* by Yoccoz, see [Yoc10]. In particular we can even choose this polygon P to only have right angles. This is not necessary in general but will be helpful in a later proof. Let $\pi : P \to X$ be the natural quotient map. The polygon P has $2m$ sides and corners where m denotes the number of paired edges. The $2m$ corners in P are identified by the gluing to form the n singularities $\sigma_1, \dots, \sigma_n$. Applying the Euler formula on the finite translation surface obtained by gluing P we obtain the following relation for m:

$$m = 2g + n - 1.$$

Label the $2m$ edges by $e_1^{\pm}, \dots, e_m^{\pm}$ such that paired edges are of the form (e_i^+, e_i^-). Fix a point $\tilde{x}_0$ in the interior of P and let $x_0 := \pi(\tilde{x}_0)$ be its image in X^*. It is a common topological fact that the surface X^* of genus g and with n punctures has a fundamental group isomorphic to the free group on $m = 2g + n - 1$ generators. In our current setting we can now choose those generators in a natural way. Namely, we let c_i be the homotopy class of a loop in X^* based in x_0 that crosses only the edge $\pi(e_i^+) = \pi(e_i^-)$ exactly once. And we orient the loop in such a way that its preimage in the polygon P enters the edge e_i^+ and exits e_i^-. Figure 2.4 shows an example of such a choice of loops. Hence, the fundamental group of X^* is given by

$$\pi_1(X^*, x_0) = \langle c_1, \dots, c_m \rangle \cong F_m,$$

and we equip the fundamental group with the word metric with respect to the generating set $\{c_1 \dots, c_m\}$. Note that the choice of generators c_i depends on the polygon P. For a formal definition of the word metric on groups see Section 2.7.

Given a finite translation surface X coming from a Euclidean polygon $P \subset \mathbb{C}$ we define two different notions of diameter. The *diameter of* X is given by

$$\mathrm{diam}(X) := \max_{x,y \in X} d(x, y),$$

where d is the natural flat metric on X. For the second definition we define the *polygonal metric* d_P on P as follows. For any two points $x, y \in P \subset \mathbb{C}$ let $d_P(x, y)$ be the length of the

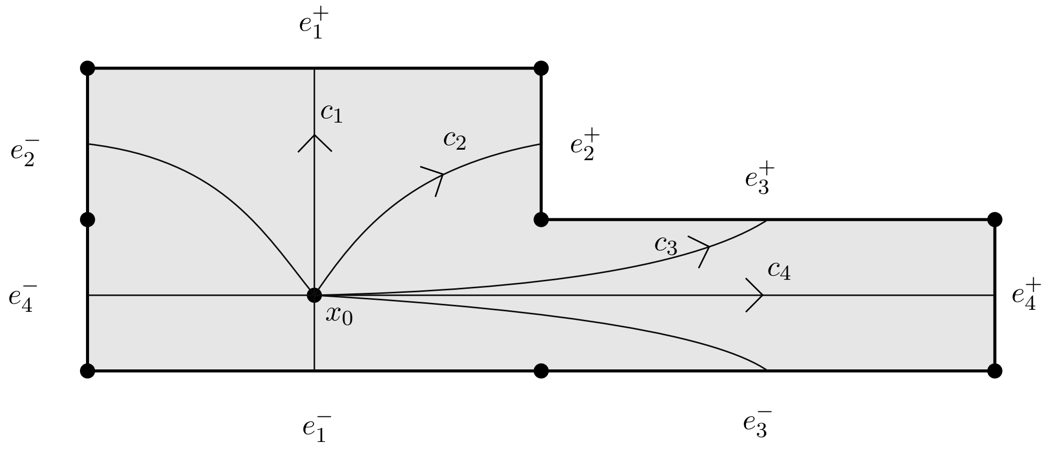

Figure 2.4: Four loops whose homotopy classes generate the free group $\pi_1(X^*, x_0)$.

shortest path between x and y which lies inside the polygon P. Note that if P is not convex then this shortest path may consist of several straight line segments as shown in figure 2.5 (a). We now define the *diameter of P* as

$$\operatorname{diam}(P) := \max_{x,y \in P} d_P(x, y).$$

In general both notions of diameter are not identical as shows the example in figure 2.5 (b).

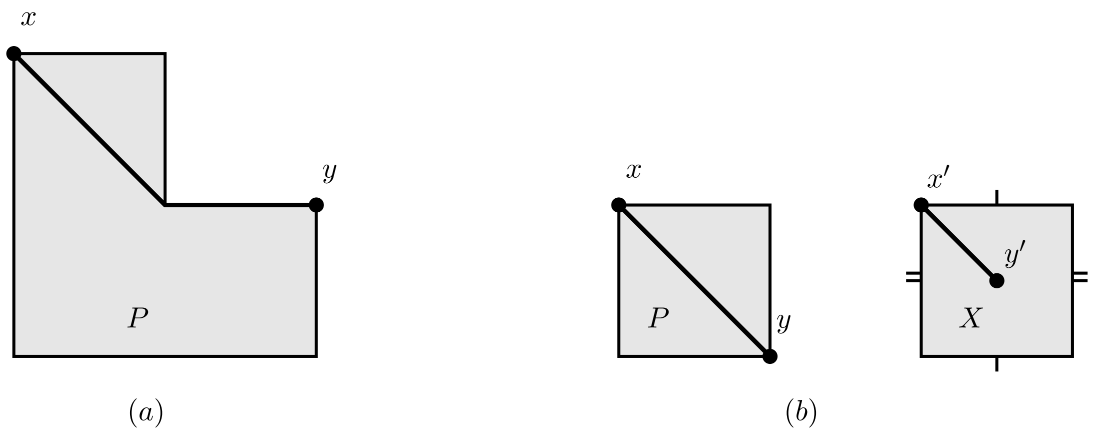

Figure 2.5: (a) A path in P realizing the distance $d_P(x, y)$. (b) Two paths in the polygon P resp. the corresponding translation surface X realizing the respective diameters.

2.5 Singular loops

For each singularity $\sigma_i \in \Sigma$ in the finite translation surface X we consider the homotopy class $r_i \in \pi_1(X^*, x_0)$ of a loop based in x_0 around σ_i. More precisely we choose this loop to be minimal with respect to its word length and say that r_i is a *singular loop* for σ_i. Note that this element r_i is not unique: If r_i has word length d_i then there are altogether $2d_i$ possible singular loops corresponding to the d_i cyclic conjugates of r_i and their inverses. By ignoring the inverses we thus have d_i possible choices of singular loops around σ_i all having word length d_i with respect to the generating system $\{c_1, \ldots, c_m\}$ of $\pi_1(X^*, x_0)$. By *cyclic conjugates* of a word $w = w_1 \ldots w_d$ we mean the words obtained by successively conjugating the word with $w_1^{-1}, w_2^{-1}, \ldots, w_{d-1}^{-1}$, i.e. the words $w_2 \ldots w_d w_1$, $w_3 \ldots w_d w_1 w_2$ etc. Illustratively the number d_i corresponds to the number of different 'corner sectors' crossed by a singular loop around a given singularity σ_i, see Figure 2.6. Or put differently, the number d_i is the number of corners

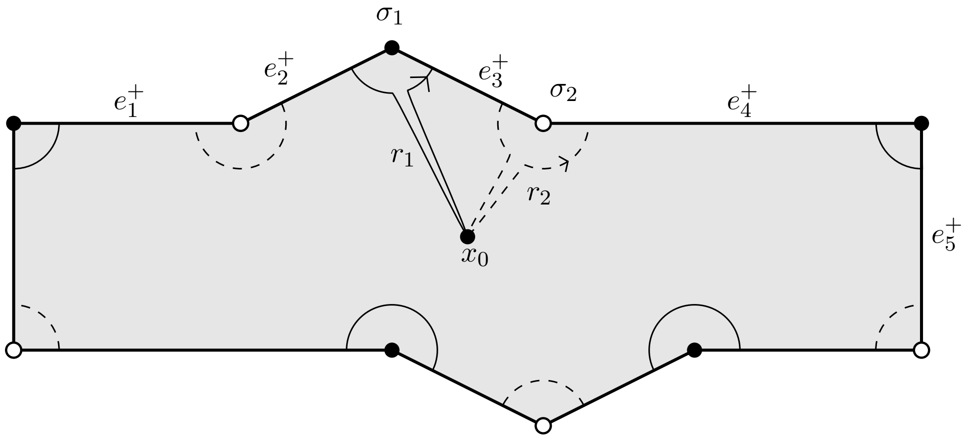

Figure 2.6: An example for singular loops around σ_1 and σ_2. Here $d = 5$ and $r_1 = c_3 c_4^{-1} c_5 c_1 c_2^{-1}$ and $r_2 = c_4 c_5^{-1} c_1^{-1} c_2 c_3^{-1}$.

in the polygon P which are identified with the singularity σ_i after gluing. This observation shows us that we can count the number $2m$ of corners in P by adding the number of 'corner sectors' for each singularity. In other words

$$\sum_{\sigma_i \in \Sigma} d_i = 2m.$$

In particular we can bound the word length d_i of those singular loops r_i by $2m$. If we now choose such singular loops $r_1, \ldots, r_n$ for each singularity we can describe the fundamental group $\pi_1(X, x_0)$ via the following group presentation, see [Sti93].

$$\pi_1(X, x_0) \cong \langle c_1, \ldots, c_m \mid r_1, \ldots, r_n \rangle.$$

2.6 REGULAR COVERINGS

We need some further general tools from covering theory, see [Hat02] or [Ful95]. Consider a topological covering $p : S \to S_0$ between two topological surfaces. The covering p is called *regular* if the deck transformation group

$$\mathrm{Deck}(S|S_0) := \{f \in \mathrm{Homeo}(S) \mid p \circ f = p\}$$

acts transitively on the fiber $p^{-1}(x)$ for any $x \in S_0$. We say that a regular covering is a *G-covering* if its deck transformation group is isomorphic to G. Two coverings $p : S \to S_0$ and $p' : S' \to S_0$ are *equivalent* if there is a homeomorphism $f : S \to S'$ such that $p' \circ f = p$. Note that in the case of p and p' being translation coverings this implies that f is a translation and hence both translation surfaces are equivalent.

Consider a regular topological covering $p : S \to S_0$ and fix a base point $\tilde{x}_0 \in S$ and $x_0 := p(\tilde{x}_0) \in S_0$. This defines a surjective group homomorphism, the *monodromy map*, as follows. Let

$$\mu : \pi_1(S_0, x_0) \to \mathrm{Deck}(S|S_0), \ c \mapsto f^{-1},$$

where f is the unique element in $\mathrm{Deck}(S|S_0)$ which maps $\tilde{x}_0$ to the endpoint of the lift $\tilde{c}$ of c starting in $\tilde{x}_0$. Note that this element is well defined as the endpoint lies in the fiber $p^{-1}(x_0)$ and $\mathrm{Deck}(S|S_0)$ acts transitively on this fiber. Note that it is necessary to take the inverse of f in order for μ to be a homomorphism. In the case that p is a G-covering with fixed isomorphism $\varphi : \mathrm{Deck}(S|S_0) \xrightarrow{\sim} G$ we often work with the alternative monodromy map $\varphi \circ \mu : \pi_1(S_0, x_0) \to G$. The following lemma shows that we can detect equivalent G-coverings by studying their monodromy maps.

Lemma 2.4. *Let $p : S \to S_0$ and $p' : S' \to S_0$ be two G-coverings with corresponding monodromy maps $\mu, \mu' : \pi_1(S_0, x_0) \twoheadrightarrow G$. The following are equivalent:*

(i) p and p' are equivalent coverings,

(ii) $\ker \mu = \ker \mu'$,

(iii) There is $\alpha \in \mathrm{Aut}(G)$ such that $\mu' = \alpha \circ \mu$.

Proof. We give a sketch of the proof here. For further details consider [Hat02].

(i)$\Rightarrow$ (ii) Let $\gamma \in \ker \mu$. In other words, each lift of γ in S is a closed path. By the equivalence of p and p' there is a homeomorphism $f : S \to S'$ satisfying $p' \circ f = p$. In particular, each lift of γ in S' is the image under f of a lift in S and hence a closed loop, too. This proves $\ker \mu \subseteq \ker \mu'$. The converse inclusion follows analoguously.

(ii)$\Rightarrow$ (i) Fix a universal covering $\tilde{S} \to S_0$ of S_0. The normal subgroup $U := \ker \mu = \ker \mu'$ in G acts freely and properly discontinuously on $\tilde{S}$ and hence induces a covering $q : U\backslash\tilde{S} \to S_0$. The quotient space by definition is homeomorphic to S, say via $f : S \to U\backslash\tilde{S}$. By definition of the covering q we have that $p = q \circ f$ and hence p and q are equivalent. The same holds for S' and p' and we get that p and p' are equivalent.

(ii)$\Leftrightarrow$ (iii) The direction (iii) $\Rightarrow$ (ii) is clear. Define the following map $\alpha : G \to G$ via $\alpha(g) := \mu'(\gamma)$ where γ is an arbitrary element in the preimage $\mu^{-1}(g) \subset \pi_1(S_0, x_0)$. This map is well-defined: If γ_1, γ_2 both lie in $\mu^{-1}(g)$ then $\gamma_1\gamma_2^{-1}$ lies in $\ker \mu$ which is $\ker \mu'$ by (ii). From this follows that $\mu'(\gamma_1) = \mu'(\gamma_2)$ and α is well-defined. Furthermore, it is not hard to check that α is a group homomorphism.

□

2.7 Quasi-isometries and Cayley graphs

The definitions and results in this chapter are taken from [FM12], [BH99] and [Bri08]. The proofs for all statements in this section can be found in [BH99]. Let (X, d) and (X', d') be two metric spaces. A map $f : (X, d) \to (X', d')$ is a *quasi-isometric embedding* if there are constants $A \geq 1, B \geq 0$ such that we have for all points $x, y \in X$

$$\frac{1}{A} \cdot d(x, y) - B \leq d'(f(x), f(y)) \leq A \cdot d(x, y) + B.$$

The map f is called *quasi-surjective* if there is a constant $R \geq 0$ such that for any point $z \in X'$ there is a point $x \in X$ with $d'(f(x), z) \leq R$. A quasi-surjective quasi-isometric embedding as above is called an *(A,B)-quasi-isometry* and we also say that the spaces (X, d) and (X', d') are *quasi-isometric*. Note that this property is indeed an equivalence relation as it is always possible to construct a quasi-isometry $g : X' \to X$ from a given quasi-isometry $f : X \to X'$.

There is an interesting result combining both the geometry of a metric space and properties of a group acting on that space in a certain way. A metric space (X, d) is *proper* if closed balls of finite radius in X are compact. It is a *geodesic* metric space if there exists a geodesic arc between any two points in X. Given a group G acting on the metric space (X, d) we call the action *cocompact* if there is a compact set $K \subseteq X$ such that $X = G.K$, or equivalently if the quotient space $G \backslash X$ of the action is compact. The action is *properly discontinuous* if, for each compact $K \subseteq X$, the set $\{g \in G \mid g.K \cap K \neq \emptyset\}$ is finite.

Remark. Let X be a tame translation surface with the induced flat metric d. We have seen before that closed balls of finite radius around ∞-angle singularities are non-compact. Hence $(\overline{X}, d)$ is a proper metric space if and only if $\mathrm{Sing}(X)$ only contains conical singularities.

We can turn a finitely generated group G with finite generating set S into a proper metric space endowed with the so-called *word metric* d_S induced by S as follows. We first define the *word length* of an element $g \in G$ to be

$$|g|_S := \min\{k \in \mathbb{N}_0 \mid g = s_1 \dots s_k,\ s_i \in S \cup S^{-1}, i = 1, \dots, k\},$$

where S^{-1} denotes the set of inverses of S. Then the *word metric* for two elements $g, h \in G$ is

$$d_S(g, h) := |g^{-1}h|_S$$

and defines a metric which is invariant under G-left multiplication. The proof of the following lemma can be found in [BH99], Proposition 8.19.

Lemma 2.5 (Švarc-Milnor-Lemma)**.** *Let (X, d) be a proper geodesic metric space. If a group G acts properly discontinuously and cocompactly by isometries on X, then G is finitely generated. Furthermore, for any choice of finite generating set S and basepoint $x_0 \in X$ the map $(G, d_S) \to (X, d)$ $g \mapsto g.x_0$ is a quasi-isometry.*

One important class of examples for quasi-isometric spaces is given by Cayley graphs of finitely generated groups. Let G be a group generated by a finite set S. From this we construct a graph $\mathrm{Cay}(G,S)$, the *oriented Cayley graph* of G with respect to S, as follows. The vertices of $\mathrm{Cay}(G,S)$ are the elements of G. And for each generator $s \in S$ and vertex $g \in G$ there is an edge from the vertex g to gs. Note that the unoriented Cayley graph is always a connected graph since S is a generating set. In the literature one sometimes finds this graph to be unoriented and also the restrictions on S may vary. It is often required that S may not contain the neutral element of G or any self-inverse elements. However, for our purposes we do not restrict S as long as all its elements generate G. In particular, we even allow S to contain two or more identical elements. Figure 2.7 shows part of the Cayley graph of $\mathbb{Z}$ with generating set $S = \{0, -2, 3\}$.

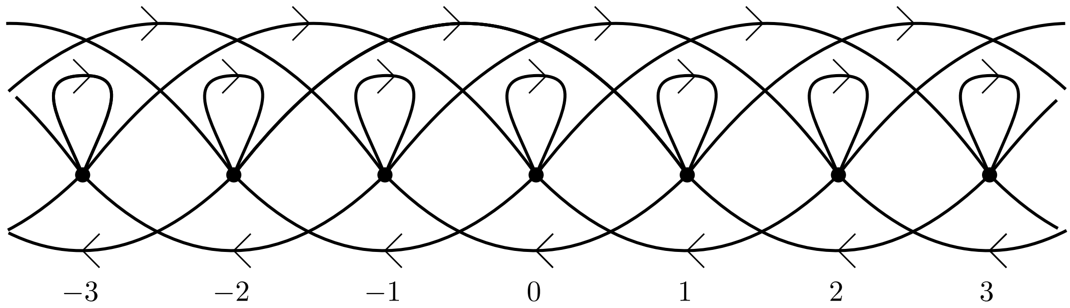

Figure 2.7: The Cayley graph $\mathrm{Cay}(\mathbb{Z}, \{0, -2, 3\})$.

The group G itself acts on $\mathrm{Cay}(G,S)$ by left-multiplication on the vertices: $(g,h) \mapsto gh$. In this way each element $g \in G$ induces an automorphism of $\mathrm{Cay}(G,S)$. Furthermore, we equip $\Gamma := \mathrm{Cay}(G,S)$ with the natural graph metric d_Γ where the distance between any two adjacent vertices is 1. With respect to this metric each automorphism induced by the G-action is in fact an isometry. One easily checks that the left-action of G on $\mathrm{Cay}(G,S)$ as metric spaces is properly discontinuous and cocompact. By the Švarc-Milnor Lemma 2.5 the injective map $G \to \mathrm{Cay}(G,S),\ g \mapsto g$ is a quasi-isometry. This proves the first statement of the following Lemma. For the second statement we refer to [FM12], Theorem 8.2.

Lemma 2.6. *Let G be a group generated by the finite subset S. Then (G, d_S) and $(\mathrm{Cay}(G,S), d_\Gamma)$ are quasi-isometric. Furthermore the quasi-isometry class of (G, d_S) does not depend on the finite generating set S.*

For this reason we will often omit the finite generating set S when we talk about the quasi-isometry class of G. And we often shortly say that some metric space X is quasi-isometric to G when it is clear that we are implicitly talking about word metrics corresponding to finite generating sets.

Remark. It is important to keep in mind that the previous lemma only holds for finite generating sets. While it is possible to define the Cayley graph with respect to an infinite generating set as we will do in Section 4.3, it turns out in general that this Cayley graph is not at all quasi-isometric to the group with word metric corresponding to a finite generating set. In an example at the beginning of Section 4.4 we consider the Cayley graph $\mathrm{Cay}(\mathbb{Z}, 2\mathbb{Z})$ and it turns out that this is a bounded metric space having only one end, in contrast to $\mathrm{Cay}(\mathbb{Z}, S)$ being unbounded and having two ends when S is a finite generating set.

There has been much work on the classification of quasi-isometry classes of metric spaces, and in particular of finitely generated groups. A helpful tool is to find objects associated to metric spaces that remain invariant under quasi-isometries. We call them *QI-invariants*. The usual approach in order to distinguish the quasi-isometry classes of two given metric spaces is to find a suitable QI-invariant and to show that it is different for those metric spaces. We give here some well-known facts about quasi-isometries, see [BH99], Chapter 8.

Lemma 2.7. *(i) Any two finite groups resp. bounded metric spaces are quasi-isometric.*

(ii) A group G is quasi-isometric to $\mathbb{Z}$ if and only if it contains $\mathbb{Z}$ as finite index subgroup.

(iii) For $n \geq 2$ all free groups F_n on n generators are quasi-isometric.

(iv) $\mathbb{Z}^m$ is quasi-isometric to $\mathbb{Z}^n$ if and only if $m = n$.

(v) If H is a finite index subgroup of G then H and G are quasi-isometric.

Some common QI-Invariants for finitely generated groups are e.g. finiteness, hyperbolicity, their growth rate and their space of ends. See the next section for the notion of ends.

2.8 The space of ends

The originial definiton of the space of ends can be found in [Fre31] and works with ascending sequences of compact subsets in a topological space. In this thesis we use the equivalent definition using proper rays which is introduced in [Hop44], §1, and which will turn out to be useful for our study of translation surfaces. All the following definitions and results can be found in [BH99].

Let X be a topological space. A *proper ray* in X is a map $r : [0, \infty) \to X$ such that the preimage of every compact set in X is again compact. In particular, for every compact $K \subseteq X$ there is $T \geq 0$ such that $r([T, \infty])$ does not intersect K. More illustratively, a proper ray is an infinite path in X that might cross a compact set many times but after a finite time it leaves this compact set and never reenters it. Two proper rays $r, r' : [0, \infty) \to X$ are *equivalent* if for every compact $K \subset X$ there is $T > 0$ such that the restricted rays $r|_{[T,\infty)}$ and $r'|_{[T,\infty)}$ lie in the same path component of $X \backslash K$. The equivalence class of a proper ray r, denoted by $\mathrm{end}(r)$, is called an *end* of X and we denote by $\mathrm{Ends}(X)$ the set of all ends of X.

Next, we equip $\mathrm{Ends}(X)$ with a topology by defining a notion of convergence as follows. A sequence of ends $\mathrm{end}(r_n)$ converges to an end $\mathrm{end}(r)$ for $n \to \infty$ if and only if for every compact set $K \subseteq X$ there exists a sequence of integers N_n such that $r_n[N_n, \infty)$ and $r[N_n, \infty)$ lie in the same path component of $X \backslash K$ whenever n is sufficiently large. The topology on $\mathrm{Ends}(X)$ is now defined by describing its closed sets, where $A \subseteq \mathrm{Ends}(X)$ is *closed* if it satisfies the following condition: if $\text{end}(r_n) \in A$ for all $n \in \mathbb{N}$ then $\text{end}(r_n) \to \text{end}(r)$ implies $\text{end}(r) \in A$. The following result, see [BH99], Proposition 8.29, shows that quasi-isometric spaces have homeomorphic spaces of ends.

Lemma 2.8. *Let X and X' be proper geodesic spaces. Then every quasi-isometry $f : X \to X'$ induces a homeomorphism $\text{end}(f) : \text{Ends}(X) \to \text{Ends}(X')$.*

Given a finitely generated group G with finite generating set S its Cayley graph $\Gamma := \text{Cay}(G, S)$ clearly is a geodesic space. And since S is a finite generating set, all vertices of Γ have finite valence and Γ is proper. We then define $\text{Ends}(G) := \text{Ends}(\Gamma)$. In this definition the space of ends $\text{Ends}(G)$ depends, a priori, on the choice of the generating set S. But the previous lemma together with Lemma 2.6 show that this is well-defined up to homeomorphisms.

One important result in geometric group theory about ends of finitely generated groups is the following, see [BH99], Theorem 8.32, and [Hop44], [Sta68] for details.

Theorem 2.9. *Let G be a finitely generated group. Then $\text{Ends}(G)$ has either $0, 1, 2$ ends or is a Cantor set. More precisely,*

(i) G has 0 ends if and only if it is finite,

(ii) G has 2 ends if and only if it contains $\mathbb{Z}$ as a subgroup of finite index,

(iii) G has infinitely many ends if and only if G can be expressed as a certain form of amalgamated free product or HNN extension.

Remark. These two results combined show us that for finitely generated groups the number of ends is a QI-invariant. However, this does not hold true for infinite generating sets. In Section 4.4 we will consider Cayley graphs of groups with infinite generating sets and see examples where such a Cayley graph for $\mathbb{Z}$ is bounded but does have one end, although $\mathbb{Z}$ when finitely generated has 2 ends.

Chapter 3

Regular Translation Surfaces

3.1 Definition

A translation surface X is called *regular* if there is a finite translation surface X_0 and a regular translation covering $p : X^* \to X_0^*$. The regular covering p can be extended to a continuous map $\overline{X} \to X_0$. We will denote this extension by p as well. Note that $p^{-1}(\mathrm{Sing}(X_0)) = \mathrm{Sing}(X)$. If the deck transformation group $\mathrm{Deck}(X^*|X_0^*)$ is isomorphic to some abstract group G we also say that X is a *G-regular* translation surface. We call the covered finite translation surface X_0 the *base surface of* X. Clearly, X is an infinite translation surface if and only if G is infinite.

Lemma 3.1. *Regular translation surfaces are tame.*

Proof. Let X be a regular translation surface and X_0 its base surface. We show that the minimal length of saddle connections in $\overline{X}$ is bounded below by a positive number. The claim then follows from Lemma 2.3. Since p is a local isometry each saddle connection in $\overline{X}$ is mapped to a saddle connection in X_0 of the same length. In particular the infimum of possible lengths of a saddle connection in X is exactly the infimum L of saddle connection lengths in X_0. Since X_0 is a finite translation surface L is a positive number. $\square$

3.2 Construction

We next describe an illustrative way of constructing G-regular translation surfaces and we prove that every regular translation surface can be constructed this way. This will turn out to be very useful in the later sections as we can consider regular translation surfaces as a glued collection of multiple copies of one and the same polygon.

We make use of the notions defined in Section 2.4. Fix a finite translation surface X_0 obtained by gluing paired parallel sides of a $2m$-sided polygon P. Let $\pi : P \to X_0$ be the corresponding quotient map. Fix further a finitely generated group G with m generators $s_1, \ldots, s_m$. By the universal property of free groups there is a unique homomorphism $\mu : F_m \cong \pi_1(X_0^*, M) \to G$ which maps $c_i^{-1} \mapsto s_i$. We now construct a translation surface X_μ as follows. Consider G copies of the polygon P, i.e. one copy $P \times \{g\}$ for each $g \in G$ and identify the sides of those polygons in a particular way. Namely we identify the e_i^+-side of each polygon $P \times \{g\}$ with the e_i^--side of the polygon $P \times \{gs_i\}$, $i = 1, \ldots, m$. As for the constructive definition of translation surfaces we define $\overline{X}_\mu$ to be $(P \times G)/_{\sim_\mu}$ where $\sim_\mu$ denotes the equivalence relation given by the gluing identification of sides, and X_μ is $\overline{X}_\mu$ without all vertices of infinite degree. Therefore points in $\overline{X}_\mu$ can be parametrized as equivalence classes of pairs (x, g) where $x \in P, g \in G$. We shortly write $[x, g] \in X_\mu$ for such an equivalence class. Note that the resulting space X_μ is connected since the group G is generated by $s_1, \ldots, s_m$. By construction X_μ is a translation surface and it is an infinite one if and only if G is infinite. Note that the definition of X_μ does not only depend on X_0 and μ but on the specific gluing polygon P we choose for X_0. However, we will show in the following section that different polygons for one and the same surface X_0 yield equivalent translation surfaces X_μ, hence we can omit P in the description of X_μ.

We still need to show that the constructed surface X_μ is a G-regular translation surface. In the way we constructed X_μ there is a natural left G-action by translations on X_μ given by $h.[x, g] := [x, hg]$ which when restricted to X_μ^* is clearly free and properly discontinuous. The latter follows from the fact that each element acts by translation of the copies of P and hence the distance between translated copies is bounded below. This way we get a covering map

$$p^* : X_\mu^* \to X_0^*, \ [x, g] \mapsto \pi(x),$$

which extends to a continuous map $p : \overline{X}_\mu \to X_0$. Furthermore, the quotient space of this action is exactly X_0^*. Hence the deck transformation group $\mathrm{Deck}(X_\mu^* | X_0^*)$ is isomorphic to G. Since G acts transitively on the fibre of each point $x \in X_0^*$ we have in particular that p restricted to X_μ^* is a regular covering. Note that the monodromy map corresponding to p^* is given by μ.

Remark. We can easily illustrate G-regular coverings constructed as before by drawing the Euclidean polygon P and labeling each edge with a generator of G. If two edges are paired we often omit the label of one of both edges since each label in G of an edge is simply the inverse of the label of its paired egde. See the following example.

Example. Consider the G-regular translation surface covering X_0 given by Figure 3.1 where $G = S_3$ is the symmetry group on 3 elements. The monodromy map μ is given as shown on the left-hand side. It assigns to each generator of $\pi_1(X_0^*, x_0)$ a transposition in S_3 such that all transpositions together generate the whole group. The preimages of the black resp. white singularities of the base surface are the three black resp. white singularities on the regular translation surface. An edge on the right-hand side is glued to the unique parallel edge having the same singularities as endpoints.

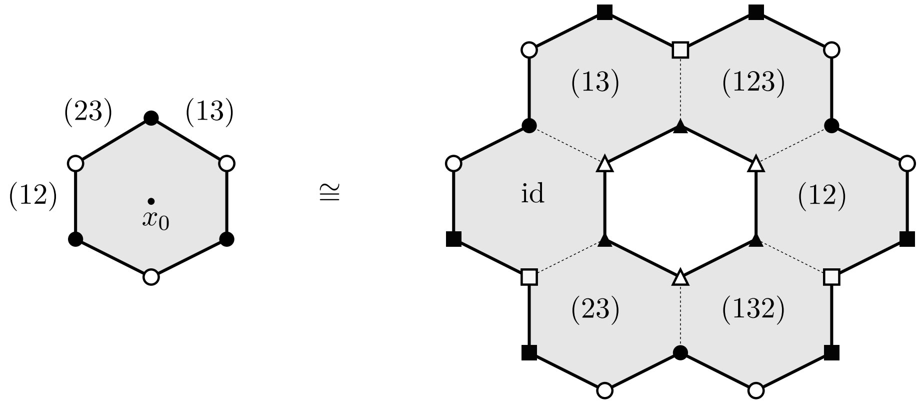

Figure 3.1: An S_3-regular translation surface.

3.3 Equivalence of regular translation surfaces

In this section we show that each regular translation surface X is equivalent to a constructed one of the form X_μ and that the equivalence class is independent of the choice of base polygon. For this reason we will from now on only consider G-regular translation surfaces of the form X_μ constructed from G copies of one Euclidean polygon.

Proposition 3.2. *(i) For each regular translation surface X there is a polygon P, a finitely generated group G and a monodromy map μ such that X is equivalent to the constructed translation surface $X_\mu = (P \times G)/\sim_\mu$ without the vertices of infinite degree.*

(ii) The equivalence class of X does not depend on the choice of base polygon P.

Proof. Let $p : X^* \to X_0^*$ be the regular translation covering associated to the regular translation surface X. Choose an arbitrary Euclidean polygon $P \subset \mathbb{R}^2$ with paired edges such that $P/_\sim$ is equivalent to X_0 where $\sim$ is the map gluing two paired edges together. In other words, there is a local translation $T : P/_\sim \to X_0$. Such a polygon can for example be found using the zippered rectangle construction presented in [Yoc10]. Fix a point $\tilde{x}_0$ lying in the interior of P and let $x_0 := T(\tilde{x}_0) \in X_0$. Let $G := \mathrm{Deck}(X^*|X_0^*)$ be the deck transformation group of the covering p. As shown in Section 2.6 the covering p induces the monodromy homomorphism $\mu : \pi_1(X_0^*, x_0) \to G$. This monodromy in return induces a labeling of the edges of P as follows: To each edge e of P is assigned a unique homotopy class of paths in $P/_\sim$ based in $\tilde{x}_0$ crossing only the edge e once. Let $\gamma_e \in \pi_1(X_0^*, x_0)$ be the image of this class under T and define $\mu(\gamma_e^{-1}) \in G$ to be the label of e. The inverse argument of μ comes from the fact that we defined μ as a homomorphism. Now construct the translation surface X_μ from $\overline{X}_\mu := (P \times G)/_{\sim_\mu}$ as explained above.

As we have already seen this construction induces a new covering $p' : X_\mu^* \to (P/_\sim)^*$ which as well is a G-regular translation covering. Let $q := T \circ p' : X_\mu^* \to X_0^*$ and consider the following commuting diagram. By the construction of X_μ we know that the corresponding monodromy maps μ_p and μ_q of those two translation coverings are identical. Lemma 2.4 then implies that p and q are equivalent as topological coverings and so there exists a homeomorphism

$$\begin{tikzcd}
X_\mu^* \arrow[r, dashed, "f"] \arrow[d, "p'"'] \arrow[dr, "q"] & X^* \arrow[d, "p"] \\
(P/\sim)^* \arrow[r, "T"'] & X_0^*
\end{tikzcd}$$

$f\ :\ X^* \to X_\mu^*$ such that $p = q \circ f$. The resulting map f locally is a translation since p and q are locally translations. The map f can be extended continuously to the metric completion $\overline{X} \to \overline{X}_\mu$ satisfying $f(\mathrm{Sing}(X)) = \mathrm{Sing}(X_\mu)$. Hence X and X_μ are equivalent translation surfaces which proves (i). Since the choice of P was arbitrary this also proves (ii). $\square$

Remark. With these results in mind we will from now on always assume that all G-regular translation surfaces are of the form $X_\mu = (P \times G)/_{\sim_\mu}$ for some polygon P on $2m$ sides and a surjective monodromy homomorphism $\mu\ :\ F_m \to G$.

Example. Consider the following two G-regular translation surfaces $X_\mu, X_{\mu'}$ where $G = F(a, b)$ is the free group on two generators, as shown in Figure 3.2. The two finite base translation

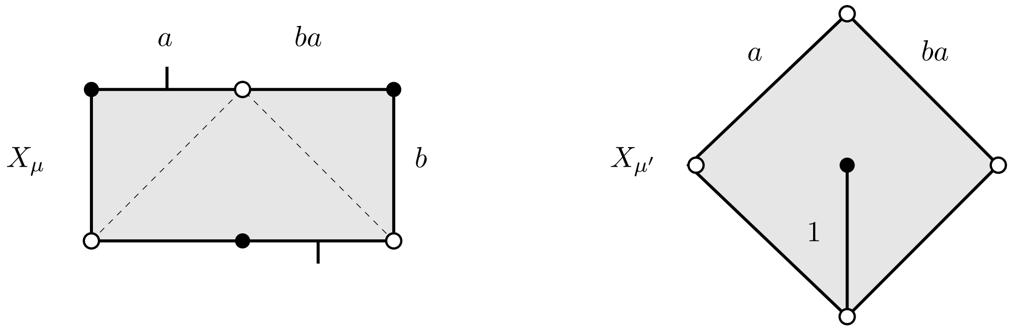

Figure 3.2: Two $F(a, b)$-regular translation surfaces with equivalent base surfaces

surfaces X_0 and X_0' are equivalent since we can cut the left-hand polygon along the dotted lines and glue the remaining parts together to form the right-hand polygon. We claim that both infinite G-regular translation surfaces are equivalent as well. In order to see this we choose generators c_1, c_2, c_3 of the fundamental group of the equivalent base surfaces and study the monodromy maps μ and μ' of the coverings, see Figure 3.3. The monodromy map μ of the left-hand covering is easily seen to be

$$\mu\ :\ c_1^{-1} \mapsto a,\ c_2^{-1} \mapsto ba,\ c_3^{-1} \mapsto b,$$

while following the paths of the right-hand covering results in a monodromy map

$$\mu'\ :\ c_1^{-1} \mapsto a,\ c_2^{-1} \mapsto a^{-1}ba^2,\ c_3^{-1} \mapsto a^{-1}ba.$$

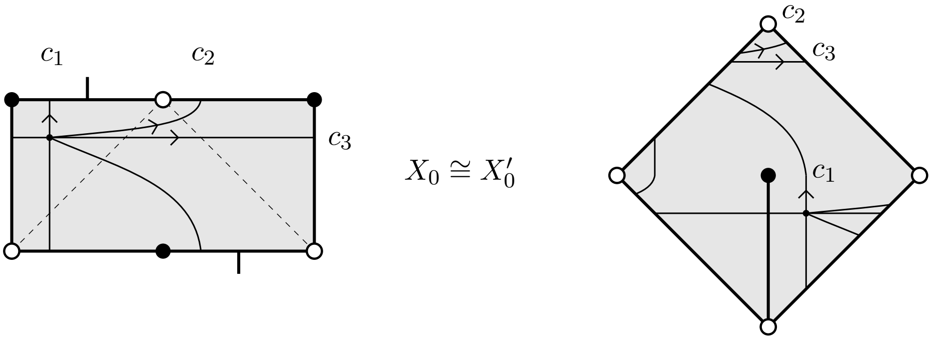

Figure 3.3: A choice of three generators of $\pi_1(X_0^*, x_0)$ in the equivalent base surfaces

Although both maps are not identical a short calculation shows that the inner automorphism α of $F(a, b)$

$$\alpha \ : \ a \mapsto a, \ b \mapsto a^{-1}ba,$$

satisfies $\mu' = \alpha \circ \mu$, and hence both monodromy maps are equivalent. Lemma 2.4 now implies that the corresponding translation coverings are equivalent es well and therefore the two regular translation surfaces X_μ and $X_{\mu'}$ are equivalent as well. An illustrative proof of this fact can also be seen in Figure 3.4, where a part of the regular translation surfaces is shown. Here, one can "see" the isometry between both translation surfaces mapping black resp. white

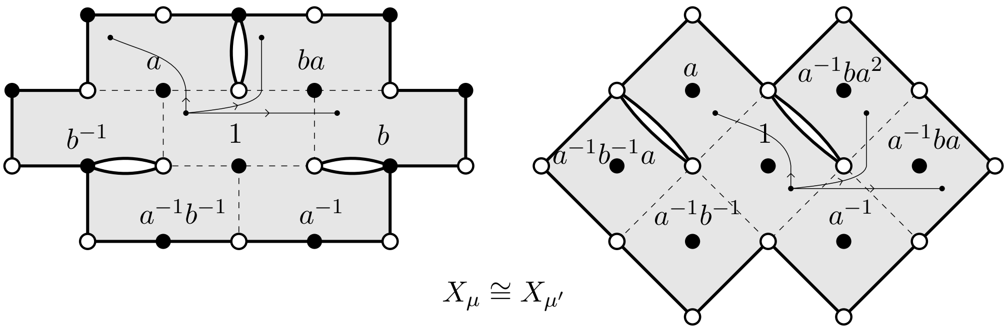

Figure 3.4: Both $F(a, b)$-regular translation surfaces turn out to be equivalent

singularities of X_μ onto black resp. white singularities of $X_{\mu'}$. Note that the white singularities are all ∞-angle singularities. That is why we draw those curved slits into the surfaces which should be thought of as straight slits.

3.4 Singularity types

A lot of data and properties of a G-regular surface X_μ can be extracted simply by studying the base surface X_0 and the monodromy map $\mu \ : \ \pi_1(X_0^*, x_0) \to G$. We show now that we can describe the types of singularities of X_μ using methods of group theory.

Consider first a finite translation surface X_0 with singularities $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ which is given by some simply connected polygon P with $2m$ sides. Let $\tilde{x}_0$ be a point in the interior of P and $x_0 \in X_0$ its corresponding image under the gluing map. As described in Section 2.4

let $c_1, \dots, c_m$ be the loops in X_0^* corresponding to the side pairs of P and which generate the free group $\pi_1(X_0^*, x_0)$ of rank m. For each singularity σ_i choose a singular loop $r_i \in \pi_1(X_0^*, x_0)$ around it as shown in Section 2.5. Now fix a surjective homomorphism $\mu : \pi_1(X_0^*, x_0) \to G$ such that $s_i := \mu(c_i^{-1})$ form a finite generating system S for the group G, $i = 1, \dots, m$. For the singular loop r_j around each singularity σ_j define $u_j := \mu(r_j^{-1}) \in G$ to be the *monodromy element* for σ_j, $j = 1, \dots, n$. Note again that a different choice of singular loop r' around σ_j results in a monodromy element u' which is conjugated to u_j in G or its inverse. Finally we define the set

$$\Sigma^\infty := \{\sigma_j \mid u_j \text{ has infinite order in } G,\ j = 1, \dots, n\} \subseteq \Sigma.$$

The set Σ^∞ is well defined as it does not depend on the choice of the singular loop r_j: From the observations made in Section 2.5 we know that since a different singular loop r' around σ_j results in an element u' conjugated to u_j in G or inverse to it, both elements u_j and u' have the same order $\in \mathbb{N} \cup \{\infty\}$. Note however that the subset Σ^∞ always depends on the homomorphism μ.

As we have seen before μ induces a G-regular translation surface X_μ and a regular covering which can be extended to $p : \overline{X}_\mu \to X_0$ such that $\operatorname{Sing}(X_\mu) = p^{-1}(\Sigma)$. Since X_μ is tame it only has conical or ∞-angle singularities. The following lemma tells us that we can recognize the ∞-angle singularities in X_μ as the preimages of Σ^∞ under p.

Lemma 3.3. *Let X_μ and Σ^∞ be as before.*

(i) $\sigma \in \Sigma^\infty$ if and only if each of its preimages is ∞-angle.

(ii) $\sigma \notin \Sigma^\infty$ if and only if each of its preimages is conical. More precisely, if σ has multiplicity k_σ then each of its preimages has multiplicity

$$(k_\sigma + 1)\operatorname{ord}(u_\sigma) - 1.$$

Proof. Let $\sigma \in \Sigma$ and $\tilde{\sigma} \in p^{-1}(\sigma)$ one of its preimages in $\overline{X}_\mu$. We show that $\sigma \in \Sigma \backslash \Sigma^\infty$ if and only if $\tilde{\sigma}$ is conical. On the polygon P fix a corner C which is mapped onto σ under the gluing map. By construction of X_μ we find a copy $P \times \{g\}$ of the polygon such that the corner after identification $[C, g]$ is exactly the preimage $\tilde{\sigma}$. Consider the unique singular loop $r \in \pi_1(X_0^*)$ around σ corresponding to the corner C and consider its lift $\tilde{r}$ in X_μ^* starting in the point $[x_0, g] \in X_\mu^*$. The endpoint of $\tilde{r}$ is then $[x_0, gu]$ where $u = \mu(r^{-1})$ since μ is the monodromy map of the covering. Similarly for every $k \in \mathbb{Z}$ the endpoint of the lift of r^k starting in $[x_0, g]$ is given by $[x_0, gu^k]$. See Figure 3.5 for an example. Obviously $\tilde{\sigma}$ is a conical singularity if and only if the lift of some power of r is a closed loop in X_μ^*, i.e. if and only if there is a $k \in \mathbb{Z}$ such that $gu^k = g$ or equivalently if and only if $u \in G$ has finite order. This last statement is equivalent to $\sigma \notin \Sigma^\infty$.

Let us calculate the multiplicity in the finite case. Since $\sigma \notin \Sigma^\infty$ we know that each of its preimages $\tilde{\sigma}$ is conical. Each lift of the singular loop r corresponds to a sector of angle sum

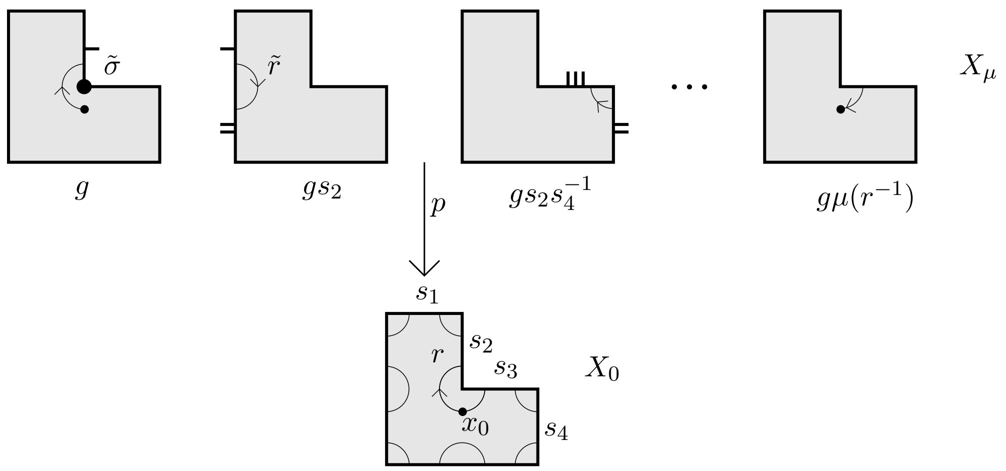

Figure 3.5: The lift $\tilde{r}$ of a singular loop starts in the copy g and ends in the copy $g\mu(r^{-1})$.

$2\pi(k_\sigma + 1)$ around $\tilde{\sigma}$. Since the lift of $r^{\mathrm{ord}(u)}$ is a closed simple loop around $\tilde{\sigma}$ and $\mathrm{ord}(u)$ is minimal with respect to this property we can add up the angle sums of the $\mathrm{ord}(u)$ many sectors and get a total angle sum of $2\pi(k_\sigma+1)\cdot \mathrm{ord}(u)$. Hence its multiplicity is $(k_\sigma+1)\mathrm{ord}(u)-1$. $\square$

We will now describe the singularity set $\mathrm{Sing}(X_\mu)$ in more detail and denote by $\mathrm{Sing}^\infty(X_\mu) := p^{-1}(\Sigma^\infty)$ the set of all ∞-angle singularities of X_μ. Remember that the points in $X_\mu = (P \times G)/_{\sim_\mu}$ are given by equivalence classes $[x, g]$ of pairs (x, g) where $x \in P$ and $g \in G$. Then $[x, g]$ is a singularity (possibly of multiplicity 0) if and only if x is a corner of the polygon P. Now for each $\sigma \in \Sigma$ of the n singularities fix one corner C_σ of P representing it and let $u_\sigma \in G$ be the monodromy element corresponding to the corner C_σ. Clearly, $\mathrm{Sing}(X_\mu) = \{[C_\sigma, g] \mid \sigma \in \Sigma, g \in G\}$.

Because of $[C_\sigma, g] = [C_\sigma, h]$ if and only if $h = gu_\sigma^m$ for some $m \in \mathbb{Z}$ we know that $[C_\sigma, g] = [C_\sigma, gu_\sigma^m]$ for all $m \in \mathbb{Z}$. Hence we can identify this singularity with the coset $g\langle u_\sigma\rangle$ of $G/\langle u_\sigma\rangle$. This observation proves the following Lemma.

Lemma 3.4. *There is a bijection*

$$\mathrm{Sing}(X_\mu) = \{[C_\sigma, g\langle u_\sigma\rangle] \mid \sigma \in \Sigma, g \in G\} \xrightarrow{\sim} \bigsqcup_{\sigma\in\Sigma} G/\langle u_\sigma\rangle,$$

and similarly

$$\mathrm{Sing}^\infty(X_\mu) \cong \bigsqcup_{\sigma\in\Sigma^\infty} G/\langle u_\sigma\rangle.$$

Lemma 3.4 permits us to make two interesting statements about the structure of regular translation surfaces. Firstly, for finite regular translation surfaces we can compute the genus and the singularity type from the data of the given base surface X_0 and the monodromy μ. Secondly, we can characterize infinite regular translation surfaces only having finitely many ∞-angle singularities.

Corollary 3.5. *Let X_μ be a finite G-regular translation surface with base surface X_0. Let $u_\sigma \in G$ be a monodromy element for each singularity $\sigma \in \Sigma$ of X_0. Then the genus of X_μ can be computed using the following formula:*

$$2 - 2g(X_\mu) = |G| \cdot \left(2 - 2g(X_0) - \sum_{\sigma\in\Sigma}\left(1 - \frac{1}{\operatorname{ord}(u_\sigma)}\right)\right).$$

Proof. Since G is a finite group all the elements u_σ have finite order and the set Σ^∞ is empty. Hence the continuous extension map $p : \overline{X}_\mu \to X_0$ yields a ramified covering which is possibly ramified over the n singularities of X_0. The ramification index e_σ for each singularity is exactly the order of u_σ. Let m_σ be the number of preimages in $p^{-1}(\sigma)$ and let $d = |G|$ be the finite degree of the unramified covering $p : X_\mu^* \to X_0^*$. Then we have $m_\sigma \cdot e_\sigma = d$ for all $\sigma \in \Sigma$. Consider a triangulation of X_0 where Σ is a subset of the vertices with F faces, E edges and V vertices. This induces a triangulation on X_μ via the covering map p. Then every face of dimension 1, 2 and any vertex of the triangulation in X_0 that is not a singularity has exactly d preimages in X_μ. Applying the Riemann-Hurwitz formula to the triangulation in X_μ yields

$$\begin{aligned}\chi(X_\mu) &= d\cdot F - d\cdot E + d\cdot(V-n) + \sum_{\sigma\in\Sigma} m_\sigma\\ &= d\cdot(F-E+V) - d\cdot n + \sum_{\sigma\in\Sigma}\frac{d}{e_\sigma}\\ &= d\cdot\chi(X_0) - d\cdot\sum_{\sigma\in\Sigma}\left(1-\frac{1}{e_\sigma}\right).\end{aligned}$$

Using the relation $\chi(X_\mu) = 2 - 2g(X_\mu)$ for the genus finishes the proof. □

Example. We can use the formula from Corollary 3.5 to compute the genus of the S_3-regular translation surface in Figure 3.1. The base surface is the flat torus and has genus $g(X_0) = 1$. The monodromy element corresponding to the singular loop starting at the black top corner is (12) and has order 2. The same holds for the white bottom corner. Hence we have

$$2 - 2g(X_\mu) = 6\cdot\left(2 - 2\cdot 1 - \left(1-\frac{1}{2}\right) - \left(1 - \frac{1}{2}\right)\right) = -6,$$

and X_μ has genus $g(X_\mu) = 4$.

Corollary 3.6. *Let X_μ be G-regular having at least one ∞-angle singularity. Then X_μ has only finitely many ∞-angle singularities if and only if G is virtually $\mathbb{Z}$, i.e. if G contains an infinite cyclic subgroup of finite index.*

Proof. By the Lemma 3.4 we have a bijection

$$\operatorname{Sing}^\infty(X_\mu) \cong \bigsqcup_{\sigma\in\Sigma^\infty} G/\langle u_\sigma\rangle,$$

and hence we have that $\operatorname{Sing}^\infty(X_\mu)$ is finite if and only if $\sum_{\sigma\in\Sigma^\infty}[G:\langle u_\sigma\rangle]$ is finite. Since Σ

is finite, this union is finite if and only if $[G : \langle u_\sigma \rangle]$ is finite for all $\sigma \in \Sigma^\infty$. Therefore, G is virtually cyclic. Note that we need the condition $\mathrm{Sing}^\infty(X_\mu) \neq \emptyset$, otherwise $\Sigma^\infty = \emptyset$ holds and we do not necessarily find an infinite cyclic subgroup of G. On the other hand it is quite easy to see that if G is virtually cyclic then any infinite cyclic subgroup of G is of finite index and hence $\sum_{\sigma\in\Sigma^\infty}[G : \langle u_\sigma\rangle]$ is finite: Let $U = \langle u \rangle$ be an infinite cyclic subgroup of G and $W = \langle w \rangle$ a cyclic subgroup of finite index in G. The finite index of W implies that there is an $m \neq 0$ such that $u^m \in W$, otherwise all cosets of the form $u^m W$ would be pairwise distinct, a contradiction to the finite index. In particular, the infinite cyclic group $\langle u^m \rangle$ is a subgroup of $U \cap W$. As U and W both are infinite cyclic subgroups of G it follows that $U \cap W$ must be an infinite cyclic subgroup of finite index in both W and U. The index formula

$$[G : U] \cdot [U : (U \cap W)] = [G : (U \cap W)] = [G : W] \cdot [W : (U \cap W)]$$

then implies that U has finite index in G. □

3.5 Quotient surfaces

Consider a G-regular translation surface X_μ on which we have the natural G-action by translations. Given a subgroup $U \leq G$ we can study the orbit space $U\backslash X_\mu$ of this action. Clearly, if $U = G$ then the quotient space is again the finite base surface X_0, possibly without singularities. The following lemma shows that the quotient space always is a translation surface, but not necessarily a regular one.

Lemma 3.7. *The quotient space $U\backslash X_\mu$ is a translation surface. More precisely,*

(i) $U\backslash X_\mu$ is a finite translation surface if and only if $U \leq G$ has finite index,

(ii) The intermediate covering $U\backslash X_\mu \to X_0$ is a regular translation covering if and only if $U \leq G$ is a normal subgroup.

Proof. By definition there is a regular covering $p : X_\mu^* \to X_0^*$ with $\mathrm{Deck}(X_\mu^*|X_0^*) = G$. As a subgroup of G we know that U as well acts freely and properly discontinuously from left via translations and hence the quotient map $\pi_U : X_\mu^* \to U\backslash X_\mu^*$ is a translation covering. Furthermore the subgroup U induces an intermediate covering $p_U : U\backslash X_\mu^* \to X_0^*$ satisfying $p = p_U \circ \pi_U$. Hence p_U is a translation covering onto X_0^* as well and $U\backslash X_\mu^*$ is a translation surface. This covering is regular if and only if U is a normal subgroup which implies (ii). Finally $U\backslash X_\mu^*$ is a finite translation surface if and only if the fiber $p_U^{-1}(x)$ is finite for any $x \in X_0^*$. Since there is a bijective correspondence between the fiber and the coset space $U\backslash G$ the finiteness of the fiber is equivalent to the finite index of U in G and (i) follows. □

We can explicitly construct the quotient surface $U\backslash X_\mu$ from the polygon collection of X_μ as follows. Since the U-left action identifies all copies of the form $P \times \{u\}$, $u \in U$, the quotient surface $U\backslash X_\mu$ consists of the collection $P \times (U\backslash G)$ of copies of P which are glued together similarly to the case of X_μ: Each edge e_i^+ of a polygon $P \times \{Ug\}$ is glued to the edge e_i^- of the polygon $P \times \{Ugs_i\}$. Again, removing identified vertices of P of infinite degree in the space $(P \times (U\backslash G))/_\sim$ results in a translation surface which is the quotient surface $U\backslash X_\mu$.

Example. Consider the previous example of the S_3-regular surface. We have the subgroup $U := \langle(12)\rangle \leq S_3$ of order two which is not normal in S_3. There are three cosets $U = \{\mathrm{id}, (12)\}, U(13) = \{(13), (132)\}$ and $U(23) = \{(23), (123)\}$. The quotient surface $U\backslash X_\mu$ is shown in Figure 3.6. Note that the quotient surface covers the base surface X_0 but not regularly: The lift of the loop γ in X_0 in $\tilde{x}$ is a closed path whereas the lift in $\tilde{y}$ is not closed.

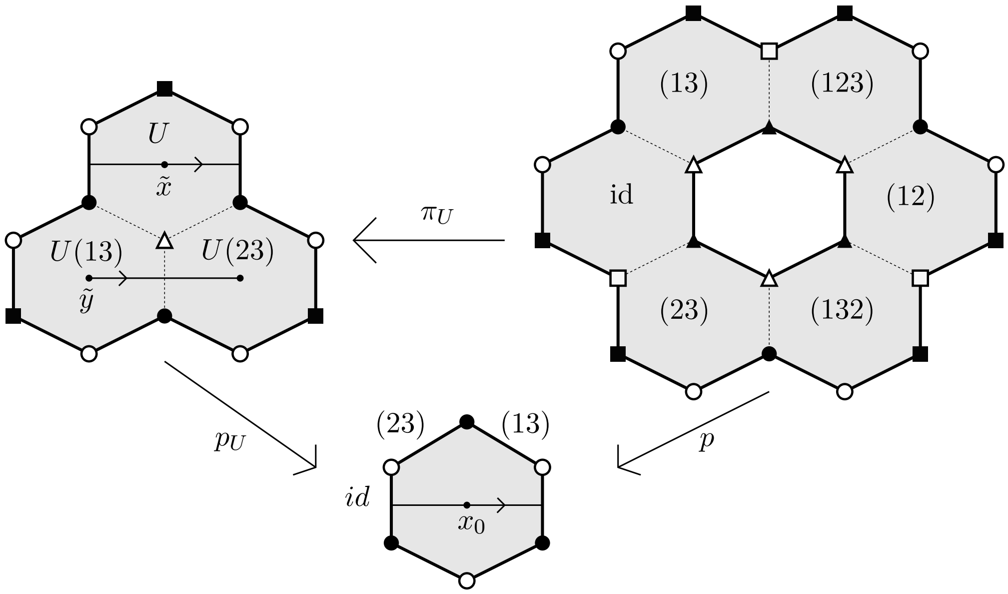

Figure 3.6: A non-regular quotient surface of the S_3-regular translation surface X_μ.

3.6 EXAMPLES AND THEIR QUASI-ISOMETRY CLASSES

In the following we present three examples of regular translation surfaces and their geometric properties. In particular we will determine their quasi-isometry class. As finite translation surfaces are all quasi-isometric to a point it is clear that only the coarse geometry of infinite regular translation surfaces will be of interest for us. The resulting observations will then be generalized in the next section.

Example (1). Figure 3.7 shows two examples of $\mathbb{Z}$-regular translation surfaces, the left one covering a torus, the right one covering a genus-2 surface. We call them the *infinite* 2- resp. 3-*staircase*, denoted by X_2 and X_3. Although both translation surfaces 'look highly similar' there are some substantial differences. Namely, as is shown in the figure the 2-staircase has exactly four ∞-singularities while the 3-staircase has infinitely many conical singularities, all of them having multiplicity $k = 2$. In particular, the completion of the 3-staircase is a proper metric space whereas the completion of the 2-staircase is not. One more important difference between those two similar translation surfaces is that they are not in the same quasi-isometry class:

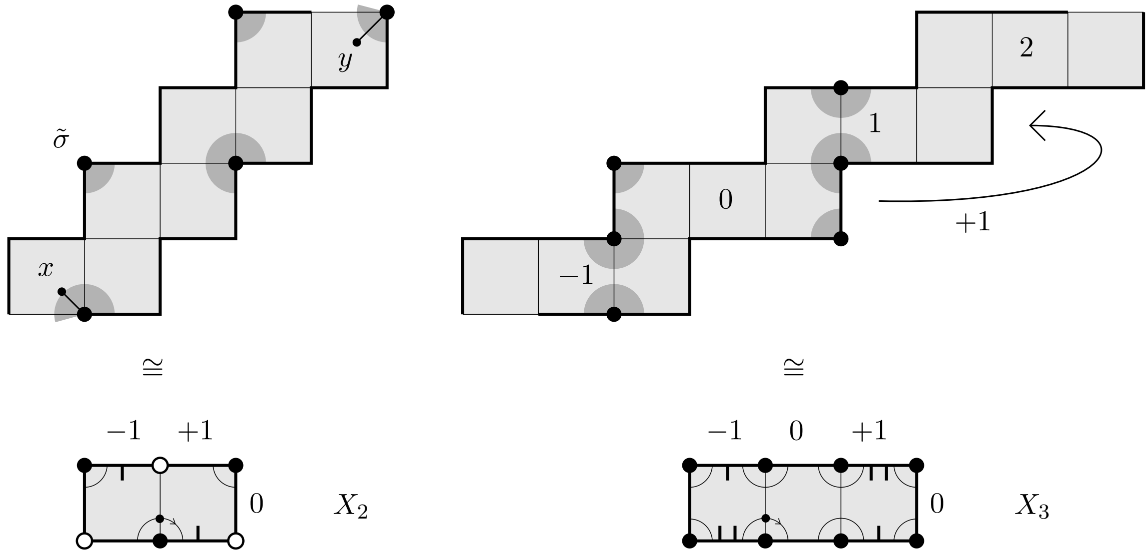

Figure 3.7: The 2-staircase has four ∞-angle singularities and is a bounded metric space. The 3-staircase only has conical singularities and is quasi-isometric to $\mathbb{Z}$.

Let us take a closer look at the geometry of the completions of these two surfaces. Although the complete 2-staircase $\overline{X}_2$ has infinite area it is a bounded metric space: Each copy of the base polygon P in $\overline{X}_2$ contains all four ∞-angle singularities. Let $\tilde\sigma$ be one of them. Then for any two points x, y in $\overline{X}_2$ we have

$$d(x,y) \leq d(x,\tilde\sigma) + d(\tilde\sigma, y) \leq 2 \cdot \mathrm{diam}(P),$$

since the distance $d(x,\tilde\sigma)$ is bounded by the largest Euclidean distance between any two points in the copy of P containing x and similarly for y.

In order to determine the quasi-isometry class of the 3-staircase X_3 note that the complete space $\overline{X}_3$ is a proper metric space as it only contains conical singularities. We have the natural $\mathbb{Z}$-action of the deck transformation group on $\overline{X}_3$ which is illustratively given by translating the copies of Q up- or downwards. This action is easily seen to be properly discontinuous and it is cocompact since the quotient space is the compact base translation surface. By the Švarc-Milnor-Lemma 2.5 it follows then that X_3 is quasi-isometric to $\mathbb{Z}$.

Example (2). For the last example let $F(a,b)$ be the free group on two generators a and b. We consider the $F(a,b)$-regular surface $X_{(a,b)}$ as shown in Figure 3.8 where the base surface consists of four unit squares and has four singularities of multiplicity $k = 0$. For further observations we will call this surface the *AB-surface.* The two black singularities have the monodromy a resp. a^{-1} while the two white ones have monodromy b resp. b^{-1}, all of infinite order in $F(a,b)$. The resulting surface $X_{(a,b)}$ is a rather intricate but nonetheless flat surface having only ∞-angle singularities. More precisely, by Lemma 3.4 we have a bijection

$$\mathrm{Sing}(X_{(a,b)}) = \mathrm{Sing}^\infty(X_{(a,b)}) \cong F(a,b)/\langle a\rangle \sqcup F(a,b)/\langle a\rangle \sqcup F(a,b)/\langle b\rangle \sqcup F(a,b)/\langle b\rangle.$$

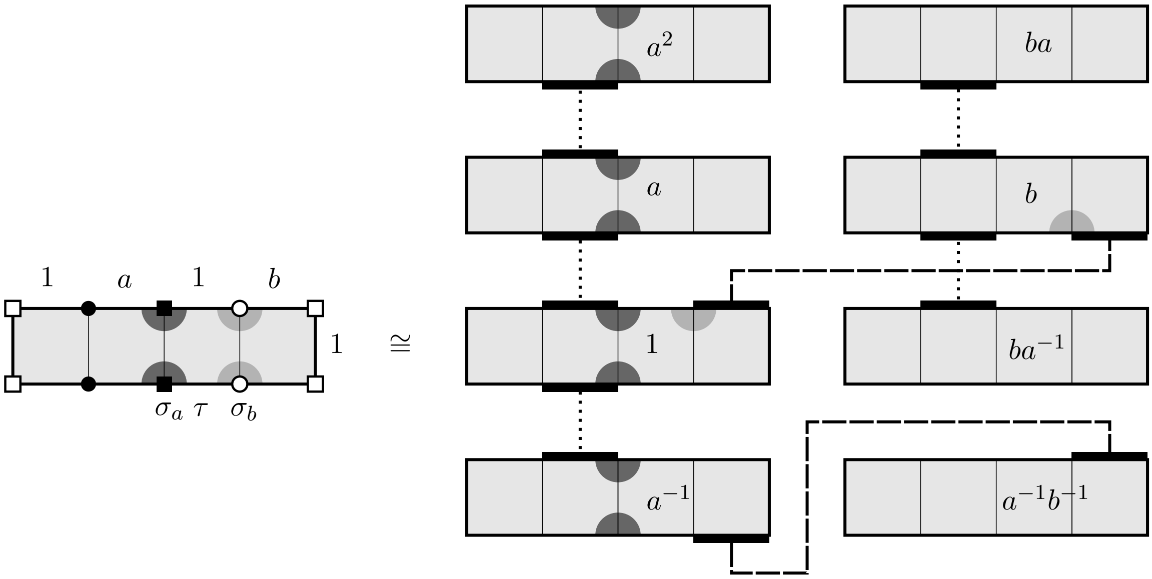

Figure 3.8: The $F(a,b)$-regular translation surface $X_{(a,b)}$

For example the black square singularity in $X_{(a,b)}$ corresponds to the coset $\langle a\rangle$ because it is contained in all copies $P \times \{a^m\}$ for $m \in \mathbb{Z}$. The following proposition determines the quasi-isometry class of $X_{(a,b)}$.

Proposition 3.8. *Let T_∞ be the regular tree of countably infinite valence. Then $X_{(a,b)}$ and T_∞ are quasi-isometric.*

Proof. We prove the claim in two steps. In step 1 we define a graph Δ and embed it into $\overline{X}_{(a,b)}$. We then show that this embedding is a quasi-isometry and therefore $\overline{X}_{(a,b)}$ as well as $X_{(a,b)}$ are quasi-isometric to Δ. In step 2 we prove that Δ is graph-isomorphic to T_∞. This concludes the proof.

Step 1. We define the *A-B-graph* Δ as follows. Let $G := F(a,b)$ and $A := \langle a\rangle$, $B := \langle b\rangle$ the relevant cyclic subgroups. The vertex set of Δ is defined to be

$$V(\Delta) := G/A \sqcup G/B$$

and two vertices gA and hB are adjacent if and only $gA \cap hB \neq \emptyset$. There are no further edges in between G/A and G/B which makes Δ a bipartite graph with vertex partitions G/A and G/B. For the embedding let σ_a resp. σ_b be the black square resp. white circle singularity of the base surface X_0 with corresponding monodromies a and b as shown in Figure 3.8. For each $g \in F(a,b)$ all the copies $(ga^m)_{m\in\mathbb{Z}}$ of P have exactly one preimage of σ_a in common, similar for the copies $(gb^m)_{m\in\mathbb{Z}}$. As observed before this means that we can describe preimages of σ_a and σ_b by the cosets G/A and G/B. More precisely, we have a bijection between the preimage $p^{-1}(\sigma_a)$ and G/A given by

$$G/A \longrightarrow p^{-1}(\sigma_a),\ gA \mapsto [C_a, g],$$

where C_a is a fixed corner of the polygon P corresponding to the singularity σ_a. An analogue statement holds for $p^{-1}(\sigma_b)$ and G/B.

Consider now the geodesic segment τ between σ_a and σ_b in X_0 which is simply the horizontal saddle-connection between them. We define $\Delta' := p^{-1}(\tau) \subset \overline{X}_{(a,b)}$ to be its preimage under the covering p and we show that Δ' is an embedding of the A-B-graph Δ defined above into $X_{(a,b)}$. The vertex set $V(\Delta')$ is $p^{-1}(\sigma_a) \sqcup p^{-1}(\sigma_b)$ and thus bijective to $G/A \sqcup G/B$. And two vertices $[C_a, gA]$ and $[C_b, hB]$ in Δ' are connected by an edge if and only if there is common copy $P \times \{k\}$ in $X_{(a,b)}$ such that both singularities lie in this copy. This is of course equivalent to $k \in gA \cap hB \neq \emptyset$. So both graphs are isomorphic.

It remains to show that this embedding is a quasi-isometry. It clearly is quasi-surjective since each point in $\overline{X}_{(a,b)}$ has distance at most $\operatorname{diam}(P)$ to a singularity $[\sigma_a, gA]$. And it is even an isometric embedding since the shortest path between two singularities in $V(\Delta')$ is simply a sequence of horizontal saddle connections of length 1 which are preimages of τ. An illustration of the situation is given in Figure 3.9.

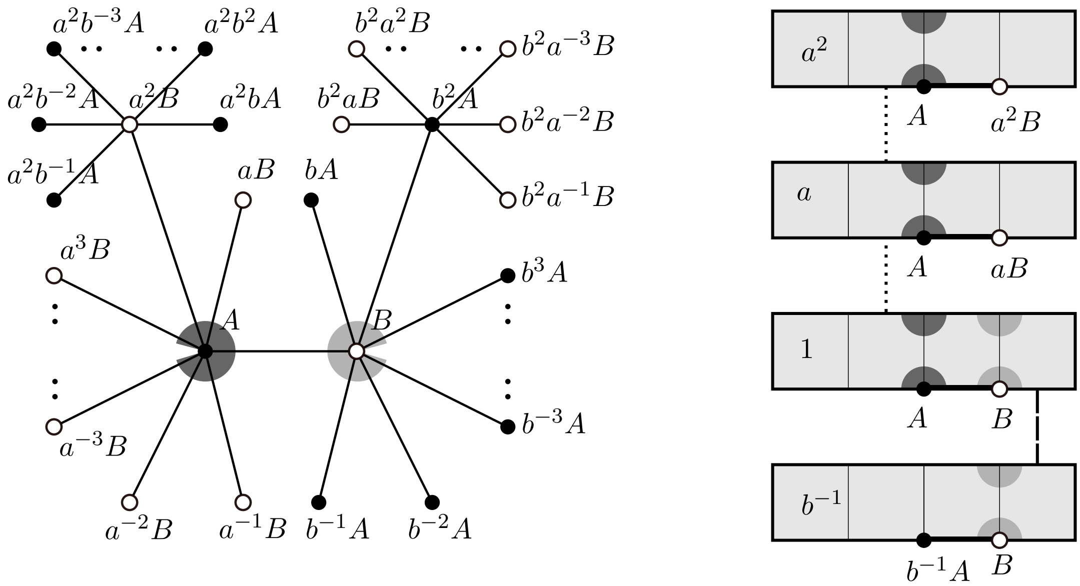

Figure 3.9: On the left-hand side is the graph Δ which is isomorphic to T_∞. On the right-hand side we see how Δ is isometrically embedded into $\overline{X}_{(a,b)}$. Each bold lower edge of a copy is an embedded edge of Δ. The dotted lines represent the gluings of the edges labelled with a resp. b.

Step 2. We have to show that Δ is a regular tree of countably infinite vertex degree. The vertex A has countably infinitely many neighbors, namely all vertices of the form $a^m B$, $m \in \mathbb{Z}$. Similar for the vertex B, see Figure 3.9. It is not hard to show that the left multiplication by elements of $G = F(a,b)$ is an action via graph automorphisms which is transitive when restricted to the partitions of $V(\Delta)$. Hence all vertices have infinite vertex degree.

It remains to prove that Δ is a tree. To each vertex of Δ we can assign a *syllable length* in $\mathbb{N}_0$ and a *syllable pattern* as follows. Let w.l.o.g. gA be a vertex in Δ, then there is a unique representant $g_0 \in gA$ of minimal word length. The syllable length $l(gA)$ of gA is then defined as

$$l(gA) := \min\{k \in \mathbb{N}_0 \mid g_0 = s_1 \dots s_k,\ s_i \in \{a^m, b^m \ : \ m \in \mathbb{Z}\}\}.$$

Note that all adjacent vertices except from the vertices A and B always have a syllable length that differs by 1, . For example, the syllable length of the vertex $a^3b^{-4}a^2b^{10}B$ is 3 since the minimal representant is $g_0 = a^3b^{-4}a^2$. Furthermore, this element g_0 is either of the form $a^*b^*a^*b^* \dots$ or $b^*a^*b^*a^* \dots$, the two possible syllable patterns.

Now assume Δ contains a closed edge path w.l.o.g. based in the vertex A (otherwise apply a suitable G-left multiplication). This path cannot contain the edge $\{A, B\}$ because this is the only edge where the syllable pattern changes. Hence, A is the only vertex of syllable length 0 in the closed edge path. Each path starting in A and not crossing B is a sequence of vertices such that the syllable length of the vertices increases by 1 in that sequence. So if we move along the closed edge path in one direction starting in A the syllable length must increase with each new vertex but on the other hand the syllable length has to be 0 again when the path returns to A, a contradiction. Hence Δ does not contain closed edge paths and must be a tree. □

Remark. By an edge contraction argument it is not hard to see that for all $n \geq 3$ the n-regular trees T_n are all quasi-isometric to each other. How about the infinite regular tree T_∞? In order to answer this question we need to define the following notion for metric spaces. A metric space (X, d) has the *finite packing property*, shortly denoted by FPP, if there is $\rho \geq 0$ such that for all radii $R \geq r > \rho$ each ball of radius R contains at most finitely many disjoint balls of radius r. The following lemma shows that having the FPP is a QI-invariant and therefore T_n is not quasi-isometric to T_∞. I thank Moishe Kohan for the proof idea, see [htt].

Lemma 3.9. *Let $f : (X, d) \to (Y, d')$ be a quasi-isometry. Then X has the FPP if and only if Y has the FPP. In particular, for all $n \in \mathbb{N}$ the tree T_n is not quasi-isometric to T_∞.*

Proof. It suffices to show that if X has the FPP then Y has the FPP. We need to make two observations. For this let $\alpha \geq 1, \beta \geq 0$ and $\delta \geq 0$ be the parameters of the quasi-isometry f, i.e. for all $x_1, x_2 \in X$ we have

$$\frac{1}{\alpha} d(x_1, x_2) - \beta \leq d'(f(x_1), f(x_2)) \leq \alpha d(x_1, x_2) + \beta,$$

and for each $y \in Y$ there is $x \in X$ such that $d'(f(x), y) \leq \delta$. In the following we denote by $B(x, R)$ the closed ball around $x \in X$ of radius $R \geq 0$.

Observation 1: For all $y \in Y, R \geq 0$ there is a point $x \in X$ and radius $\tilde{R} \geq 0$ such that

$$f^{-1}(B(y, R)) \subseteq B(x, \tilde{R}).$$

In order to prove this claim let $y \in Y, R \geq 0$. By the quasi-surjectivity there is $x \in X$ such that $d'(f(x), y) \leq \delta$. Now suppose that $z \in f^{-1}(B(y, R))$, i.e. $d'(f(z), y) \leq R$. Using first the quasi-isometry and then the triangle inequality we get

$$\begin{aligned} d(z,x) &\leq \alpha d'(f(z),f(x)) + \alpha\beta \\ &\leq \alpha\left(d'(f(z),y) + d'(y,f(x))\right) + \alpha\beta \\ &\leq \alpha R + \alpha\delta + \alpha\beta. \end{aligned}$$

Setting $\tilde{R} := \alpha R + \alpha\delta + \alpha\beta$ proves that $z \in B(x,\tilde{R})$.

Observation 2: Let $\rho \geq 0$ be such that $\frac{\rho-\delta-\beta}{\alpha} \geq 0$. Then for all $y \in Y, r > \rho$ there is a point $x \in X$ and radius $\tilde{r} > 0$ such that

$$f(B(x,\tilde{r})) \subseteq B(y,r).$$

In order to prove this claim let ρ be as mentioned and $y \in Y, r \geq \rho$. Define $\tilde{r} := \frac{r-\delta-\beta}{\alpha} > 0$ by assumption on ρ. Again by quasi-surjectivity there is $x \in X$ such that $d'(f(x),y) \leq \delta$. Now let $z \in f(B(x,\tilde{r}))$, i.e. there is $t \in X$ such that $f(t) = z$ and $d(t,x) \leq \tilde{r}$. Similar as for Observation 1 we get

$$\begin{aligned} d'(z,y) &\leq d'(f(t),f(x)) + d'(f(x),y) \\ &\leq \alpha d(t,x) + \beta + \delta \\ &\leq \alpha\tilde{r} + \beta + \delta \\ &= r, \end{aligned}$$

and therefore $z \in B(y,r)$.

Now let (X,d) have the FPP with corresponding parameter $\rho \geq 0$, and assume that (Y,d') does not. We lead this assumption to a contradiction, hence proving the first part of the lemma. Set $\rho_Y := \alpha\rho + \beta + \delta \geq 0$. Since by assumption Y does not have the FPP, there are radii $R \geq r > \rho_Y$ and a point $y \in Y$ such that the ball $B(y,R)$ does contain an infinite family of disjoint balls of radius r. Denote these balls by $B(y_i,r) \subset Y$ for $y_i \in Y$, $i \in I$ infinite. By Observation 1 the preimage $f^{-1}(B(y,R))$ is contained in a ball $B(x_0,\tilde{R})$ for some $x_0 \in X$. By Observation 2 we have $B(x_i,\tilde{r}) \subseteq f^{-1}(B(y_i,r))$ for some $x_i \in X$, i.e. each preimage of the pairwise disjoint balls $B(y_i,r)$ contains a ball of radius $\tilde{r} = \frac{r-\delta-\beta}{\alpha}$, satisfying $\tilde{r} > \rho$ since $r > \rho_Y$. Altogether we have an infinite family of pairwise disjoint balls $B(x_i,\tilde{r})$ contained in the ball $B(x_0,\tilde{R})$ such that $\tilde{r} > \rho$. This is a contradiction to X having the FPP.

For the last part of the lemma we simply need to show that T_n, $n \in \mathbb{N}$, has the FPP whereas T_∞ does not. Thus, both metric spaces cannot be quasi-isometric. As T_n is locally compact we can set $\rho = 0$ and have that any ball of radius R contains at most finitely many balls of radius $0 < r \leq R$. On the other hand since T_∞ is a regular tree of infinte valence we can find for all $\rho > 0$ a sufficiently large radius R such that a corresponding ball may contain infinitely many balls of sufficiently small radius $r > \rho$. □

Altogether we have seen that the $F(a,b)$-regular translation suface $X_{(a,b)}$ is quasi-isometric to the infinite regular tree T_∞. In particular, this observation tells us that there are regular translation surfaces that are not quasi-isometric to a Cayley graph of a finitely generated group G with respect to a finite generating system, since such a Cayley graph always has the FPP whereas T_∞ does not. In the following section we will show however that any regular translation surface is quasi-isometric to such a Cayley graph with respect to an *infinite* generating system.

Chapter 4

The Coarse Geometry of Translation Surfaces

4.1 Regular translation surfaces with only conical singularities

Firstly, we consider the "simplest" possible regular translation surfaces, namely those having no ∞-angle singularities or in other words exactly those regular translation surfaces that are also complete. It turns out that for those surfaces it is not difficult to determine their quasi-isometry class.

Theorem 4.1. *Let X_μ be a G-regular translation surface having only conical singularities. Then X_μ is quasi-isometric to G.*

Proof. The proof generalizes the example of the 3-staircase in Section 3.6, using the Švarc-Milnor Lemma. Since $X_\mu = \overline{X}_\mu$ only contains conical singularities it is a proper geodesic space. We have the natural G-action on X_μ by left multiplication given by $g.[x, h] = [x, gh]$ which corresponds to the translation of glued copies of P in X_μ. This action is cocompact since the quotient space $G\backslash X_\mu$ is isometric to the base surface X_0. For the properly discontinuous action note that compact subsets in X_μ are bounded, closed and contain only finitely many singularities as those lie discretely in X_μ. Since X_μ consists of copies of one and the same Euclidean polygon, any compact set $K \subset X_\mu$ must be contained in the union of finitely many of those copies. As each element $g \in G$ acts on X_μ by translating the copies of the polygon this implies that only finitely many translates $g.K$ intersect the set K. Hence the action is properly discontinuous. Lemma 2.5 then implies that X_μ and G are quasi-isometric. $\square$

Interestingly, the coarse geometric behaviour of regular translation surfaces changes as soon as they contain ∞-angle singularities. For example the 2-staircase is a $\mathbb{Z}$-regular translation surface containing four ∞-angle singularities. But as we have seen it is not quasi-isometric to $\mathbb{Z}$ since it is a bounded metric space.

Example. Figure 4.1 shows an example of a $\mathbb{Z}^2$-regular translation surface. The base surface X_0 has one singularity of multiplicity $k = 2$ and the corresponding monodromy element is the neutral element. Hence in the covering surface X_μ each singularity is again a conical singularity of multiplicity 2 as well. The previous theorem implies that X_μ is quasi-isometric to $\mathbb{Z}^2$.

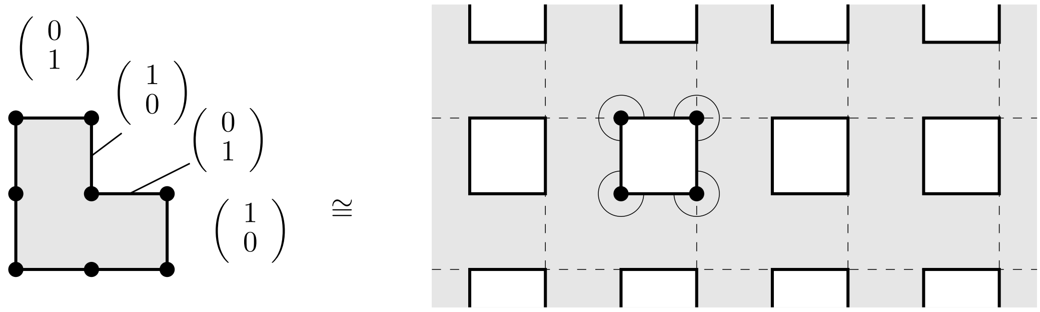

Figure 4.1: A $\mathbb{Z}^2$-regular translation surface with only conical singularities and hence quasi-isometric to $\mathbb{Z}^2$.

4.2 Boundedly generated subgroups

Now let us try to understand the case where X_μ does contain ∞-angle singularities. In particular we have $X_\mu \neq \overline{X}_\mu$. Note that since X_μ is obtained by removing a discrete subset of $\overline{X}_\mu$ the inclusion $X_\mu \hookrightarrow \overline{X}_\mu$ is a quasi-isometry. For that reason we do not distinguish between X_μ and $\overline{X}_\mu$ when it comes to determining the quasi-isometry type of a regular translation surface. Assuming that X_μ has an ∞-angle singularity, by Lemma 3.3 this is equivalent to the base surface X_0 having a singularity $\sigma \in \Sigma$ such that a corresponding monodromy element $u_\sigma \in G$ has infinite order in G. Geometrically this means that for any $g \in G$ all the infinitely many polygons $P \times \{gu_\sigma^m\}$, $m \in \mathbb{Z}$, intersect in a common corner which corresponds to a preimage of σ in $\overline{X}_\mu$. In other words, even if u_σ has large word length in G the copies $P \times \{g\}$ and $P \times \{gu_\sigma\}$ are very close in the metric space $\overline{X}_\mu$: Any two points lying in the union of these copies have maximal distance $2\mathrm{diam}(P)$. In order to see this choose for each of the two points a shortest path fully contained inside the copy of P from the point to the common corner. Then each of the two paths then has length at most $\mathrm{diam}(P)$.

The previous observation shows that it might be a good idea to identify elements g and gu_σ in G such that the word length between them becomes short as well. For this, let $T := \{u_\sigma \mid \sigma \in \Sigma\}$ for a choice of monodromy elements $u_\sigma \in G$ for each singularity σ in X_0. Then define $U := U_\mu \leq G$ to be the subgroup generated by these finitely many elements and call it the *monodromy subgroup of G generated by T*. Clearly, the monodromy subgroup U depends on the choice of T, but if the choice T of monodromy elements is clear from the context we will often speak of "the" monodromy subgroup of G. Now, with this definition, each subcollection of all polygons $P \times \{u_\sigma^m\}$, $m \in \mathbb{Z}$, in $\overline{X}_\mu$ forms a bounded subspace since they all have one corner in common. Having observed this it might seem likely that X_μ and the quotient surface $U\backslash X_\mu$ are quasi-isometric since, loosely spoken, we collapse an infinite but bounded subset $\bigsqcup_{\sigma\in\Sigma} P \times \{u_\sigma^m\}$ to just one copy of the base polygon. Unfortunately this idea is wrong in general. We will see later an example where both translation surfaces are not quasi-isometric. However, our clue turns out to be true in some special case, namely when U is a so-called boundedly generated subgroup.

Definition 4.2. Let (G, S) be a finitely generated group with finite generating set S. Recall that the *word length* of $g \in G$ with respect to S is given by

$$|g|_S = \min\{k \mid g = s_1 \dots s_k \text{ where } s_i \in S \cup S^{-1}\}.$$

A *syllable* of h is a maximal factor s^m of g where $s \in S \cup S^{-1}$ and $m \in \mathbb{Z}$. The *syllable length* of $g \in G$ with respect to S is defined as

$$|g|_S^{syl} = \min\{k \mid g = s_1^{m_1} \dots s_k^{m_k} \text{ where } s_i \in S \cup S^{-1} \text{ and } m_i \in \mathbb{Z}\}.$$

Note that we always have $|g|_S^{syl} \leq |g|_S$ for all $g \in G$. We say that G is *boundedly generated by* S if there is a constant $L \in \mathbb{N}$ such that for all $g \in G$ we have $|g|_S^{syl} \leq L$, and the group G is called *boundedly generated* if there is a finite subset $S \subseteq G$ such that G is boundedly generated by S. The syllable length induces a G-left invariant metric d_S^{syl} on G via $d_S^{syl}(g,h) = d_S^{syl}(1_G, g^{-1}h) := |g^{-1}h|_S^{syl}$. Here, the G-left invariance follows directly from the definition. For the proof that d_S^{syl} is a metric note that we have $|1_G|_S^{syl} = 0$, $|g|_S^{syl} = |g^{-1}|_S^{syl}$ and $|gh|_S^{syl} \leq |g|_S^{syl} + |h|_S^{syl}$ for all $g, h \in G$ as ending and starting syllables of g and h may cancel out.

Example. (i) Finite groups and finitely generated abelian groups are boundedly generated. The latter follows from the fact that we can write any element g in an abelian group as $g = s_1^{m_1} \dots s_n^{m_n}$ where $s_1, \dots, s_n$ are any generators of the group.

(ii) The Heisenberg group is defined as

$$H = \langle X, Y \rangle = \left\{ \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \middle| \; a, b, c \in \mathbb{Z} \right\} = \{Y^b Z^c X^a \mid a, b, c \in \mathbb{Z}\},$$

where

$$X = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \; Y = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \; Z = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Since every element in H can be written as a product $Y^b Z^c X^a$, $a, b, c \in \mathbb{Z}$, the Heisenberg group H is boundedly generated by $\{X, Y, Z\}$ with maximal syllable length 3.

(iii) If $n \geq 2$ then the free group F_n is not boundedly generated. For example $F(a, b)$ is not boundedly generated by $S = \{a, b\}$ since the infinite sequence $w_i := (ab)^i$ satisfies $|w_i|_S^{syl} = 2i$ which is unbounded for $i \to \infty$. It is more technical to show that the free group F_n is not boundedly generated by any finite subset. However, a variety of proofs exists using geometric or algebraic methods or methods involving bounded cohomology. The claim follows for instance from results in [Gri80].

We need a lemma about the distance in G-regular translation surfaces before we can formulate the quasi-isometry statement.

Lemma 4.3. *Let $X_\mu = (P \times G)/_{\sim_\mu}$ be a G-regular translation surface and let U be the monodromy subgroup of G generated by a choice T of monodromy elements. Then for all $x, y \in P$ and for all $u \in U$ we have*

$$d([x, 1_G], [y, u]) \leq (|u|_T^{syl} + 1) \cdot \operatorname{diam}(P).$$

Proof. Fix corners $C_\sigma \in P$ in the polygon P respresenting the singularities of the base surface X_0 and let $T = \{u_\sigma \mid \sigma \in \Sigma\}$ be a choice of corresponding monodromy elements. It follows then from the description of $\operatorname{Sing}(X_\mu)$ in Lemma 3.4 that all the copies $P \times \{gu_\sigma^m\}$, $m \in \mathbb{Z}$, have the singularity $[C_\sigma, g] = [C_\sigma, gu_\sigma^m] \in \operatorname{Sing}(X_\mu)$ in common for each $g \in G$.

First observe that any two points $[x, g], [y, g]$ in $\overline{X}_\mu$ lying in the same copy $P \times \{g\}$, $g \in G$, satisfy

$$d([x, g], [y, g]) \leq d_P(x, y) \leq \operatorname{diam}(P),$$

since the distance between them in $\overline{X}_\mu$ is at most the polygonal distance d_P between x and y inside $P \times \{g\}$. For the definition of the polygonal distance, see Section 2.4.

Now let $u \in U$ and $k := |u|_T^{syl}$. Then u is of the form $u = u_{\sigma_1}^{m_1} \dots u_{\sigma_k}^{m_k}$ and we get

$$\begin{aligned}
d([x, 1_G], [y, u]) &\leq d([x, 1_G], [C_{\sigma_1}, 1_G]) + \sum_{i=1}^{k-1} d([C_{\sigma_i}, u_{\sigma_1}^{m_1} \dots u_{\sigma_{i-1}}^{m_{i-1}}], [C_{\sigma_{i+1}}, u_{\sigma_1}^{m_1} \dots u_{\sigma_i}^{m_i}]) \\
&\quad + d([C_{\sigma_k}, u_{\sigma_1}^{m_1} \dots u_{\sigma_{k-1}}^{m_{k-1}}], [y, u]) \\
&= d([x, 1_G], [C_{\sigma_1}, 1_G]) + \sum_{i=1}^{k-1} d([C_{\sigma_i}, u_{\sigma_1}^{m_1} \dots u_{\sigma_i}^{m_i}], [C_{\sigma_{i+1}}, u_{\sigma_1}^{m_1} \dots u_{\sigma_i}^{m_i}]) \\
&\quad + d([C_{\sigma_k}, u], [y, u]) \\
&\leq d_P(x, C_{\sigma_1}) + \sum_{i=1}^{k-1} d_P(C_{\sigma_i}, C_{\sigma_{i+1}}) + d_P(C_{\sigma_k}, y) \\
&\leq (k + 1) \cdot \operatorname{diam}(P).
\end{aligned}$$

See Figure 4.2 for an illustration of the proof. Here the elements denote the copy of the corresponding polygon P in $\overline{X}_\mu$ and the spiraling movement shall illustrate the lift of a singular loop starting at a corner C_σ thus connecting two copies via one ∞-angle singularity. □

Remark. Before we state the result of this section we need to make an observation about lifts of paths in regular translation surfaces. For this consider a regular translation surface X_μ

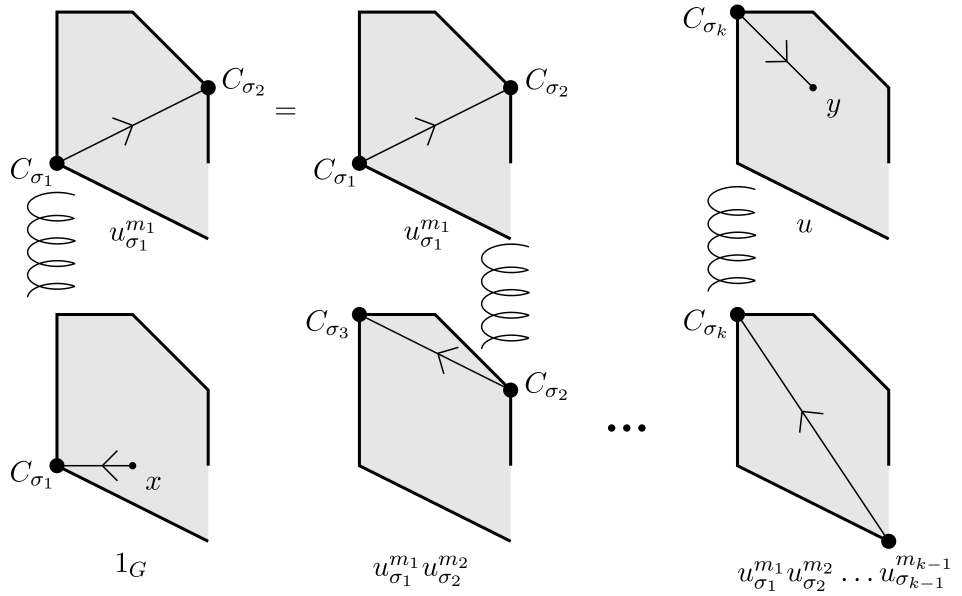

Figure 4.2: A path from $[x, 1_G]$ to $[y, u]$ crossing $|u|_T^{syl} + 1$ copies of the polygon P.

with finite base surface X_0 and the corresponding continuous map $p : \overline{X}_\mu \to X_0$. As the restriction $X_\mu^* \to X_0^*$ is an unramified cover any path γ inside X_0^*, say starting in the point x_0, can be lifted to a unique path in X_μ^* starting in $\tilde{x}_0$ for any preimage $\tilde{x}_0 \in p^{-1}(x_0)$. Of course, the same argument holds when we consider ending points instead of starting points. In general this unique lifting property of paths is not given when considering ramified coverings. Similarly, in the case of regular translation coverings there usually are many possible lifts of a path in X_0 which contains singularities. In the following we present a method how to choose such a lift from these options and thus make the lift unique.

For this let γ be a path in X_0 and assume it crosses only one conical singularity σ. The case when γ crosses several singularities can be dealt with inductively. As this will be the case for the proofs to come we may assume furthermore that γ is a concatenation of straight line segments. Denote by α the angle in counterclockwise direction between the ingoing segment γ_1 and the outgoing segment γ_2 of γ at σ. We now give a construction of a unique lift $\tilde{\gamma}$: There is a unique lift $\tilde{\gamma}_1$ of the ingoing segment γ_1 starting in a preimage $\tilde{x}_0$ of x_0, the starting point of γ. Its endpoint is a singularity $\tilde{\sigma} \in p^{-1}(\sigma)$. If $\tilde{\sigma}$ is a conical singularity of angle sum $2\pi d$ then there are d possible lifts of γ_2 starting in $\tilde{\sigma}$, and if it is an ∞-angle singularity there are even infinitely many possible lifts. However, in both cases there is only exactly one lift $\tilde{\gamma}_2$ starting in $\tilde{\sigma}$ which forms the same angle α with $\tilde{\gamma}_1$ at $\tilde{\sigma}$. The concatenation $\tilde{\gamma} := \tilde{\gamma}_1 * \tilde{\gamma}_2$ hence yields a lift of γ and since the angle is fixed this lift is unique.

Theorem 4.4. *Let X_μ be a G-regular translation surface. Fix a choice T of monodromy elements and let $U \leq G$ be the corresponding monodromy subgroup generated by T. Suppose U is boundedly generated by T. Then X_μ and its quotient surface $U \backslash X_\mu$ are quasi-isometric. In addition,*

(i) If U has finite index in G, then X_μ is a bounded metric space

(ii) If U is normal in G, then X_μ is quasi-isometric to the quotient group G/U.

Proof. We consider the surjective translation covering

$$\pi_U \ : \ \overline{X}_\mu \longrightarrow U\backslash\overline{X}_\mu, \ [x,g] \mapsto [x, Ug]$$

and prove that it is a quasi-isometric embedding with respect to the metrics d resp. d_U on $\overline{X}_\mu$ resp. $U\backslash\overline{X}_\mu$. Let w.l.o.g. $[x, 1_G]$ and $[y, g]$ be two points in $\overline{X}_\mu$ where $x, y \in P$ and $g \in G$. The quotient map π_U sends them to $[x, U]$ and $[y, Ug]$ in $U\backslash\overline{X}_\mu$. Let γ be a geodesic arc from $[x, U]$ to $[y, Ug]$ of length $d_U([x, U], [y, Ug])$. Now lift γ to the unique path $\tilde{\gamma}$ in $\overline{X}_\mu$ ending in $[y, g]$. By the previous remark this lift is well defined. The starting point of $\tilde{\gamma}$ then is of the form $[x, u_0]$ for some element $u_0 \in U$ and we have

$$d([x, u_0], [y, g]) \leq \text{length}(\tilde{\gamma})$$

The situation is illustrated in Figure 4.3. Since π_U is a translation covering both paths $\tilde{\gamma}$ and γ have the same length and we have

$$\begin{aligned} d([x, 1_G], [y, g]) &\leq d([x, 1_G], [x, u_0]) + d([x, u_0], [y, g]) \\ &\leq d([x, 1_G], [x, u_0]) + \text{length}(\tilde{\gamma}) \\ &= d([x, 1_G], [x, u_0]) + \text{length}(\gamma) \\ &= d([x, 1_G], [x, u_0]) + d_U([x, U], [y, Ug]). \end{aligned}$$

By Lemma 4.3 we know that $d([x, 1_G], [x, u_0]) \leq (|u_0|_T^{syl} + 1) \cdot \text{diam}(P)$. If U is boundedly generated this last expression is bounded by some constant $L' := (L + 1) \cdot \text{diam}(P)$ not depending on x, y and u_0. Hence

$$d([x, 1_G], [y, g]) \leq d_U([x, U], [y, Ug]) + L'.$$

On the other hand, consider a geodesic arc $\tilde{\gamma}$ between $[x, 1_G]$ and $[y, g]$ of length $d([x, 1_G], [y, g])$ in $\overline{X}_\mu$. Project this down to a path $\gamma := \pi_U(\tilde{\gamma})$ in $U\backslash\overline{X}_\mu$ between $[x, U]$ and $[y, Ug]$ of same length. However, the length of γ is surely at least the distance between $[x, U]$ and $[y, Ug]$ and we get

$$d([x, 1_G], [y, g]) = \text{length}(\tilde{\gamma}) = \text{length}(\gamma) \geq d_U([x, U], [y, Ug])$$

Both inequalities imply that π_U is a quasi-isometry between $\overline{X}_\mu$ and $U\backslash\overline{X}_\mu$. Hence the translation surfaces X_μ and $U\backslash X_\mu$ are as well quasi-isometric.

We now prove the implications (i) and (ii). The statement (i) follows directly from Lemma 3.7 (i) since $U\backslash G$ is finite. For the proof of (ii) let U be a normal subgroup of G, so $U\backslash G = G/U$ is a group. From Lemma 3.7 (ii) it follows that the quotient surface $U\backslash X_\mu$ is a G/U-regular translation surface. We claim that $U\backslash X_\mu$ only has conical singularities. Then by Theorem 4.1 there is a quasi-isometry between $U\backslash X_\mu$ and $U\backslash G$ and therefore X_μ and $U\backslash G$ are quasi-

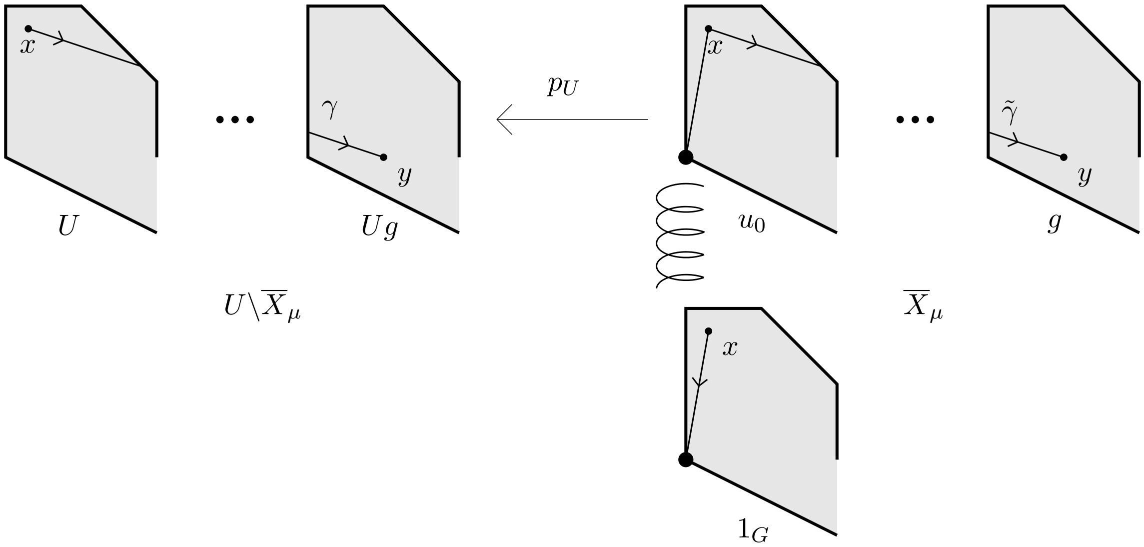

Figure 4.3: A path from $[x, 1_G]$ to $[y, g]$ containing a subpath $\tilde{\gamma}$ which is a lift of a geodesic segment from $[x, U]$ to $[y, Ug]$.

isometric. In order to prove our claim recall that the quotient surface $U\backslash X_\mu$ is of the form $(P \times (U\backslash G))/_{\sim_{\mu'}}$, where $\mu' = \pi \circ \mu \ : \ \pi_1(X_0^*, x_0) \to U\backslash G$ is the projected monodromy map and π the canonical quotient map from G to $U\backslash G$. For each singularity σ in the base surface X_0 choose a singular loop r_σ and let $u_\sigma = \mu(r_\sigma) \in U$ be the corresponding monodromy element with respect to the covering $X_\mu^* \to X_0^*$. Then the monodromy element of σ with respect to the intermediate covering $p_U \ : \ U\backslash X_\mu^* \to X_0^*$ is $\mu'(r_\sigma) = \pi \circ \mu(r_\sigma) = \pi(u_\sigma) = U$, which has order 1 in $U\backslash G$. Hence the set Σ^∞ with respect to $X_{\mu'}$ is empty. Lemma 3.3 finally implies that $U\backslash X_\mu$ only has conical singularities - having exactly the same multiplicity as their images in X_0 - and the claim is proved. $\square$

Remark. In Theorem 4.4 the condition that $U \leq G$ is boundedly generated is indeed necessary. A counterexample is given by the AB-surface $X_{(a,b)}$ presented in Section 3.6. As we have shown, the finite base surface X_0 has four singularities all having as monodromy elements a resp. b. Hence the monodromy subgroup U generated by a and b is the whole deck transformation group $G = F(a, b)$. Note that $F(a, b)$ is not boundedly generated by $\{a, b\}$. The quotient surface is $U\backslash X_\mu = G\backslash X_\mu = X_0^*$ and thus a bounded metric space whereas Proposition 3.8 implies that X_μ is quasi-isometric to T_∞, the ∞-regular tree. Hence in that case X_μ and $U\backslash X_\mu$ are not quasi-isometric.

Corollary 4.5. *Let X_μ be a G-regular translation surface where G is an abelian group. Let U be a monodromy subgroup of G. Then X_μ is quasi-isometric to G/U.*

Proof. As a subgroup of an abelian group U is normal in G and again abelian. In particular, it is boundedly generated by any generating set. The claim then follows from Theorem 4.4 (ii). $\square$

We are now able to characterize the quasi-isometry class of regular translation surfaces having at least one but at most finitely many ∞-angle singularities.

Theorem 4.6. *Let X_μ be a G-regular translation surface containing at least one ∞-angle singularity. Then X_μ has only finitely many ∞-angle singularities if and only if G is quasi-isometric to $\mathbb{Z}$. In this case X_μ is a bounded metric space.*

Proof. The main argument for the proof of the equivalence was given in the proof of Corollary 3.6. It implies that X_μ has finitely many ∞-angle singularities if and only if G is virtually $\mathbb{Z}$. The last statement is equivalent to G being quasi-isometric to $\mathbb{Z}$, by Lemma 2.7 (ii).

It remains to show that under the mentioned conditions X_μ is a bounded space. For this let X_μ contain only finitely many ∞-angle singularities, so G is virtually $\mathbb{Z}$. Fix a choice $T = \{u_\sigma \mid \sigma \in \Sigma\}$ of monodromy elements and let $U \leq G$ be the corresponding monodromy subgroup generated by T. Since $\mathrm{Sing}^\infty(X_\mu) \neq \emptyset$ at least one monodromy element $u \in T$ must have infinite order and so U contains the infinite cyclic subgroup $\langle u \rangle$. Since G is virtually $\mathbb{Z}$ any infinite cyclic subgroup, in particular $\langle u \rangle$, is of finite index in G and therefore U also has finite index in G. See the proof of Lemma 3.6 for details.

We still need to prove in a last step that U is boundedly generated by T, then Theorem 4.4 (i) implies that X_μ is a bounded metric space which proves the corollary. Since $\langle u \rangle$ has finite index in U there are finitely many cosets $\langle u \rangle, h_1\langle u \rangle, \dots, h_k\langle u \rangle$ in $U/\langle u \rangle$ for representants $h_1, \dots, h_k \in U$. Let $L := \max_{i=1,\dots,k} |h_i|_T$ be the maximum word length of these representants $h_1, \dots, h_k$ and define $h_0 := 1_G$. Now every element $g \in U$ must lie in such a coset and hence is of the form $g = h_i u^m$ with $i = 0, \dots, k$ and $m \in \mathbb{Z}$. So for the syllable length of g we get

$$|g|_T^{syl} \leq |h_i|_T^{syl} + |u^m|_T^{syl} \leq L + 1,$$

which only depends on the choice of representants but not on g. Hence U is boundedly generated by T. $\square$

Example. Consider the Euclidean plane with horizontal slits between integral points and glued together as illustrated in Figure 4.4. At first glance this space X_μ 'looks like $\mathbb{R}^2$ up to horizontal

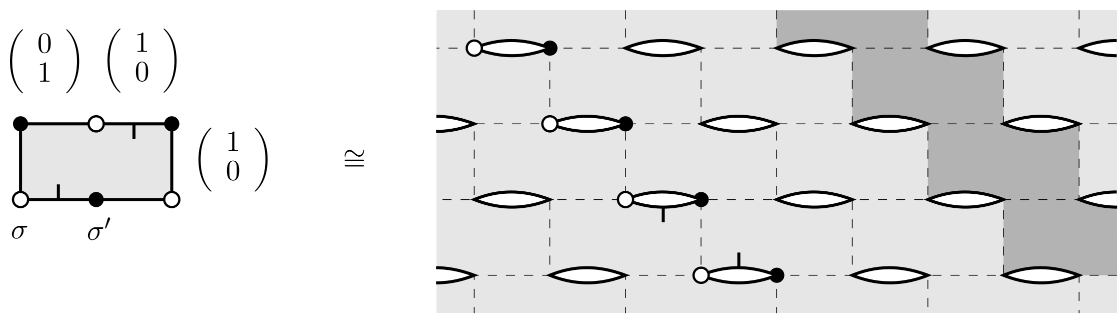

Figure 4.4: A $\mathbb{Z}^2$-regular translation surface being quasi-isometric to $\mathbb{Z}$.

slits' and one might think that it is therefore quasi-isometric to $\mathbb{R}^2$. But this first impression is wrong. In fact, the space X_μ is a $\mathbb{Z}^2$-regular translation surface covering the base surface X_0

as shown. The base surface X_0 has two singularities σ, σ' and computing the corresponding monodromy elements $u, u' \in \mathbb{Z}^2$ yields $u = u' = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. Hence the monodromy subgroup is

$$U = \langle \begin{pmatrix} 0 \\ 1 \end{pmatrix} \rangle \cong \mathbb{Z},$$

and Corollary 4.5 implies that X_μ is quasi-isometric to $G/U \cong \mathbb{Z}$. This fact can be illustrated by taking a closer look at the singularities of X_μ. Each such ∞-angle singularity is contained in a copy of the 2-staircase as shown in the light gray subspace of X_μ. Each such copy is a bounded space. Therefore, loosely spoken X_μ consists of $\mathbb{Z}$ copies of this bounded metric space. Or in other words, from the two possible dimensions in $\mathbb{R}^2$, i.e. possible "ways to stretch out towards infinity", one dimension is contracted to something bounded which turns X_μ geometrically into a strip stretching out towards infinity in only one direction.

Example. We have seen in Theorem 4.6 that if G has $\mathbb{Z}$ as finite index subgroup, then any G-regular translation surface X_μ with an ∞-angle singularity is bounded. However, the converse statement is not true. To see this, consider the $\mathbb{Z}^2$-regular translation surface X_μ as shown in Figure 4.5. The finite base surface is a torus consisting of the shown three squares

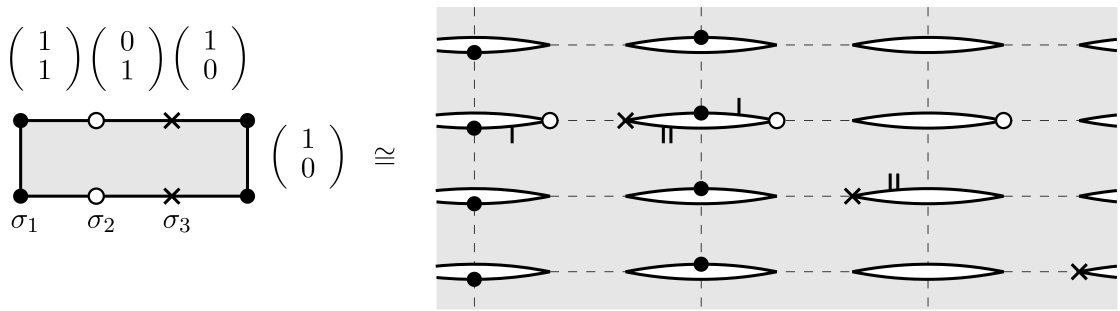

Figure 4.5: A $\mathbb{Z}^2$-regular translation surface that is bounded.

where opposite sides are glued together. Here, the monodromy elements corresponding to the singularities $\sigma_1, \sigma_2, \sigma_3$ in X_0, when computed in clockwise direction, are

$$u_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \; u_2 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \; u_3 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

It follows that the monodromy subgroup is

$$U = \langle \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \rangle = \mathbb{Z}^2,$$

and Corollary 4.5 implies that X_μ is quasi-isometric to $\mathbb{Z}^2/U \cong \{0\}$ and therefore a bounded metric space. This also makes sense illustratively since starting from one copy one can stretch to infinity vertically (black singularity), horizontally (white singularity) and diagonally (crossshaped singularity) while always being in bounded distance from the starting point. However, $\mathbb{Z}^2$ is of course not virtually cyclic or equivalently, there are infinitely many singularities since we have the bijection

$$\begin{aligned}\mathrm{Sing}^\infty(X_\mu) &\cong \left(\mathbb{Z}^2/u_1\mathbb{Z}\right) \sqcup \left(\mathbb{Z}^2/u_2\mathbb{Z}\right) \sqcup \left(\mathbb{Z}^2/u_3\mathbb{Z}\right)\\ &\cong \mathbb{Z} \sqcup \mathbb{Z} \sqcup \mathbb{Z}.\end{aligned}$$

Example. In the last and more abstract example we present a H-regular translation surface X_μ as shown in Figure 4.6. Here

$$H = \langle X, Y \rangle = \left\{ \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; a, b, c \in \mathbb{Z} \right\} = \{ Y^b Z^c X^a \mid a, b, c \in \mathbb{Z} \}$$

is the Heisenberg group, where

$$X = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \; Y = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \; Z = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The group H has the group presentation

$$\langle X, Y, Z \mid XZ = ZX,\; YZ = ZY,\; XY = ZYX \rangle.$$

X_μ :

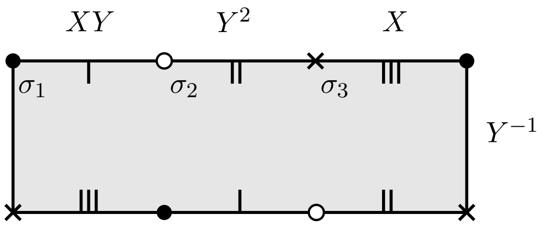

Figure 4.6: An H-regular translation surface that is quasi-isometric to $\mathbb{Z}$.

We claim now that X_μ is quasi-isometric to $\mathbb{Z}$. However, as H is not abelian we may not make use of Corollary 4.5. Consider the three top corners of the polygon associated to the singularities $\sigma_1, \sigma_2, \sigma_3$ as shown. Choosing singular loops in counterclockwise direction starting in these corners yields corresponding monodromy elements

$$u_1 = XYX^{-1}Y^{-1} = Z,\; u_2 = Y^2(XY)^{-1} = YX^{-1},\; u_3 = X(Y^{-1})^{-1}Y^{-2} = XY^{-1}.$$

Hence the monodromy subgroup U generated by $T = \{Z, YX^{-1}, XY^{-1}\}$ is

$$U = \langle Z, YX^{-1} \rangle.$$

Consider the surjective group homomorphism

$$\Phi \ : \ H \to \mathbb{Z}, \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \mapsto a+b,$$

then its kernel is given by

$$\ker \Phi = \left\{ \begin{pmatrix} 1 & -a & c \\ 0 & 1 & a \\ 0 & 0 & 1 \end{pmatrix} \,\middle|\, a, c \in \mathbb{Z} \right\} = \{ Y^a Z^c X^{-a} \mid a, c \in \mathbb{Z} \}.$$

Note that, using the commutation relations in H, we have

$$(YX^{-1})^a = YX^{-1}YX^{-1} \ldots YX^{-1} = Z^{a'} Y^a X^{-a},$$

for some integer $a' \in \mathbb{Z}$. From this observation follows that $U = \ker \Phi$ and hence is a normal subgroup of H. Furthermore, U is boundedly generated by T: Since any element $u \in U$ is of the form $Y^a Z^c X^{-a} = Z^c Y^a X^{-a} = Z^{c'}(YX^{-1})^a$, for $c' := c - a'$, we have that u has syllable length $|u|_T^{syl} = |Z^{c'}(YX^{-1})^a|_T^{syl} \leq 2$.

Altogether, we have that U is a normal subgroup of H that is boundedly generated by T. Theorem 4.4 now implies that X_μ is quasi-isometric to $H/U \cong \Phi(H) = \mathbb{Z}$. Figure 4.7 shows a part of X_μ and we see that in this example it is much more difficult to guess the quasi-isometry class of X_μ as the group structure of H is much more complex.

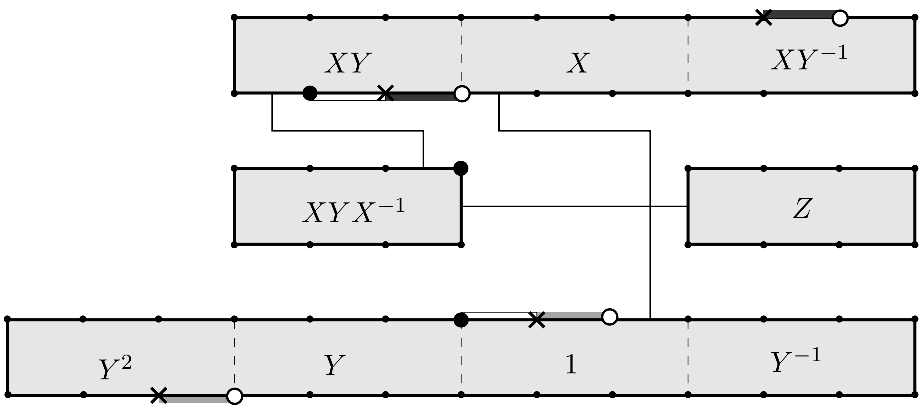

Figure 4.7: The H-regular translation surface X_μ turns out to be quasi-isometric to $\mathbb{Z}$.

4.3 Quasi-isometry classes for general regular translation surfaces

So far we have studied some examples of regular translation surfaces and their geometric properties. For some of the examples we could generalize our observations and we were able to prove that they belonged to a class of regular translation surfaces which are quasi-isometric to finitely generated groups. But not all examples had a quasi-isometry type that was as easy to describe, for example the $F(a,b)$-regular translation surface $X_{(a,b)}$ which is quasi-isometric to the locally infinite regular tree T_∞. In the following we will take a closer look at general regular translation surfaces and describe a locally infinite graph, in fact a Cayley graph with respect to an infinite generating set, which the surface is quasi-isometric to.

Before we state the result we need to understand how geodesic segments in a regular translation surface behave. In particular we only consider geodesic segments between two singularities. Lemma 2.3 then implies for regular translation surfaces that any geodesic segment between two singularities is a concatenation of finitely many saddle connections. It is thus necessary to understand the role of saddle connections in a regular translation surface.

Let X_0 be a finite translation surface obtained from a Euclidean polygon P with $2m$ paired sides and let $\pi : P \to X_0$ be the natural gluing map. For any saddle connection τ in X_0 its preimage $\pi^{-1}(\tau)$ has finitely many connected components in P which are all straight line segments in P parallel to the direction of τ. Denote by $N(\tau) := N(\tau, P) \in \mathbb{N}$ this number of segments. For a translation covering $p : X_\mu^* \to X_0^*$ this number can also be interpreted as follows: Given a saddle connection $\tilde{\tau}$ in $\overline{X}_\mu$ the number $N(p(\tilde{\tau}), P)$ is exactly the number of copies of P in $\overline{X}_\mu$ which are crossed by $\tilde{\tau}$, see Figure 4.8. The following lemma states that for right-angled polygons we can closely relate the number $N(\tau, P)$ to the length of τ.

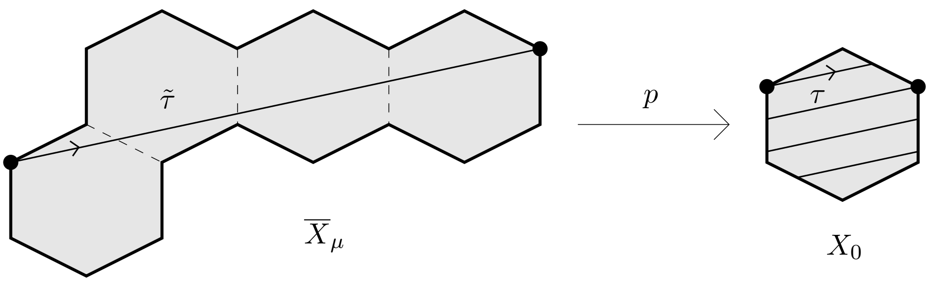

Figure 4.8: The saddle connection $\tilde{\tau}$ in $\overline{X}_\mu$ crosses $N(\tau, P) = 4$ copies of the hexagon P. Its image τ under p has $N(\tau, P) = 4$ connected components inside the polygon P.

Lemma 4.7. *Let P be a right-angled polygon for X_0 and τ a saddle connection in X_0. Then there are constants $A \geq 1, B \geq 0$ independent of τ such that*

$$\frac{1}{A} \cdot N(\tau, P) - B \leq \operatorname{length}(\tau) \leq A \cdot N(\tau, P) + B.$$

Proof. We first prove the right inequality. Let $k = N(\tau) \in \mathbb{N}$, i.e. the preimage of τ in P consists of k parallel straight line segments $t_1, \dots, t_k$ inside of P. Their lengths are clearly bounded above by $\operatorname{diam}(P)$. Hence

$$\operatorname{length}(\tau) = \sum_{i=1}^{k} \operatorname{length}(t_i) \leq \operatorname{diam}(P) \cdot k = \operatorname{diam}(P) \cdot N(\tau).$$

Now let us prove the left inequality. As P is right angled we can define its *minimum height*

h_{min} resp. *minimum width* w_{min} as follows. By assumption P only has horizontal and vertical sides, their number being p resp. q. Let $y_1, \ldots, y_p$ be the y-coordinates of the p horizontal sides and $x_1, \ldots, x_q$ the x-coordinates of the q vertical sides of P. Here, $p + q = 2m$. Define

$$h_{min} := \min\{|y_i - y_j| \ : \ y_i \neq y_j,\ i, j = 1, \ldots, p\} > 0$$

and

$$w_{min} := \min\{|x_i - x_j| \ : \ x_i \neq x_j,\ i, j = 1, \ldots, q\} > 0$$

Let $L := \text{length}(\tau) > 0$ be the length of the saddle connection τ whose preimage in P consists of $N := N(\tau)$ straight line segments. The set

$$\{x \in \tau \mid \pi^{-1}(x) \in \partial P\}$$

then consists of $N+1$ or N points on the saddle connection, depending on τ being a closed saddle connection or not. More precisely, we denote the points in this set by $\sigma_0, x_1, x_2, \ldots, \sigma_N$ where σ_0 is the starting singularity of τ, the $x_1, \ldots, x_{N-1}$ correspond to crossings with horizontal or vertical sides of P and σ_N is the ending singularity of τ. Note that if τ is closed we have $\sigma_0 = \sigma_N$. See Figure 4.9 for an example.

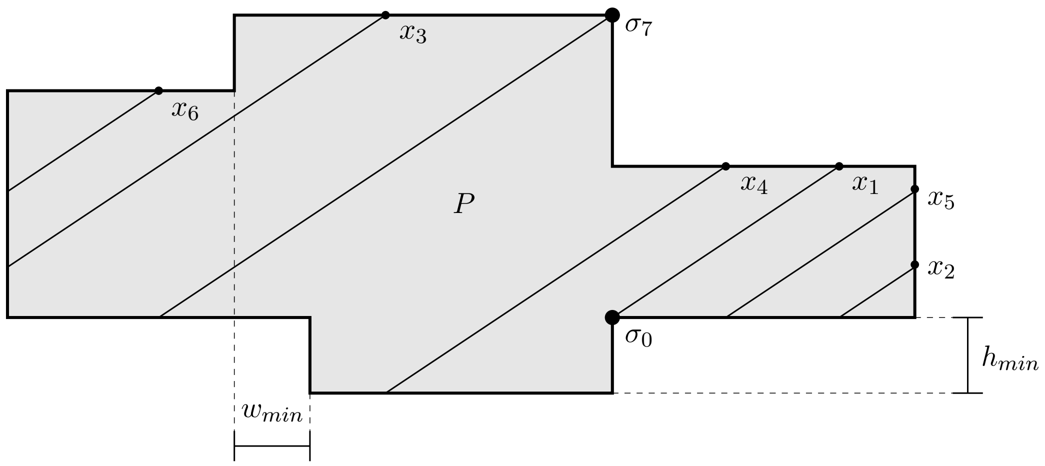

Figure 4.9: A saddle connection τ inside the polygon P. Here, $N(\tau, P) = 7$ and we have $a = 4$ horizontal and $b = 2$ vertical edge crossings.

Say, τ crosses a horizontal and b vertical edges, i.e. $a + b = N - 1$. Note that the case $a = 0$ resp. $b = 0$ is possible, for example if τ is a vertical resp. horizontal saddle connection. Now let $x_{i_1}, \ldots, x_{i_a}$ be the a horizontal edge crossings and consider the segments $[\sigma_0, x_{i_1}], [x_{i_1}, x_{i_2}], \ldots [x_{i_a}, \sigma_N]$ forming a partition of τ into $a+1$ segments. In our example, the horizontal edge crossings are given by the points x_1, x_3, x_4, x_6. By definition of h_{min} we have that each of these $a+1$ segment has length at least h_{min} and this implies that the number a of crossed horizontal edges satisfies

$$(a + 1) \cdot h_{min} \leq L.$$

Analoguously, for the number b of crossed vertical edges we get

$$(b+1) \cdot w_{min} \leq L.$$

Altogether this implies

$$N = a + b + 1 \leq L\left(\frac{1}{h_{min}} + \frac{1}{w_{min}}\right) - 1,$$

and therefore

$$N(\tau) \leq \text{length}(\tau)\left(\frac{1}{h_{min}} + \frac{1}{w_{min}}\right) - 1.$$

Note that the constants h_{min} and w_{min} only depend on the shape of P but not on the choice of τ. $\square$

Let us return to our original problem, determining the quasi-isometry type of a G-regular translation surface X_μ which is determined by a base surface X_0 and a surjective monodromy map $\mu : \pi_1(X_0^*, x_0) \to G$. As before we denote by s_i the images in G of the m edge loop generators c_i of $\pi_1(X_0^*, x_0)$ and set $S := \{s_1, \ldots, s_m\}$. The idea is to construct a G-invariant graph Γ which X_μ is quasi-isometric to. More precisely, we want to construct Γ as a Cayley graph of G with an infinite generating system S_μ^∞.

As before, the vertex set should naturally be G where each element $g \in G$ represents the point $[x_0, g]$ inside the polygon $P \times \{g\}$ in $\overline{X}_\mu$. Then, we put edges between two elements in G whenever the two corresponding copies are "close" in $\overline{X}_\mu$. Now what does "close" mean in this context? Firstly, it is clear that two copies sharing a common side are 'close'. Hence each element of S should be contained in the generating set S_μ^∞. Secondly, we have already observed that for each singularity $\sigma \in \Sigma$ of X_0 we can choose a monodromy element $u_\sigma \in G$ such that the copy $P \times \{1_G\}$ shares a singularity in $\overline{X}_\mu$ with all the copies $P \times \{u_\sigma^k\}$, $k \in \mathbb{Z}$. But any other monodromy element associated to σ gives us copies that are as well close to the copy $P \times \{1_G\}$. Therefore it makes sense to include all elements u^k in S_μ^∞ where u is a monodromy element of a singularity and $k \in \mathbb{Z}$. The following definition gives the formal setup.

Definition 4.8. Let $s_i := \mu(c_i) \in G$, $i = 1 \ldots, m$, be the images of the edge loops and set $S := \{s_1, \ldots, s_m\}$. In Section 2.5 we defined the singular loops around a singularity in X_0: For $\sigma \in \Sigma$ there are $d_\sigma \in \mathbb{N}$ cyclically conjugated singular loops $r_{\sigma,1}, \ldots, r_{\sigma,d_\sigma} \in \pi_1(X_0^*, x_0)$ around σ plus their inverses such that the number d_σ satisfies

$$\sum_{\sigma \in \Sigma} d_\sigma = 2m,$$

and in particular

$$d_\sigma \leq 2m \text{ for all } \sigma \in \Sigma. \tag{4.1}$$

Let T be the set containing all corresponding monodromy elements $u_{\sigma,i} := \mu(r_{\sigma,i}) \in G$ for

$\sigma \in \Sigma$, $i = 1, \dots, d_\sigma$. Then T contains exactly $2m$ elements counted with possible multiplicities of elements. If w is an element of the free group $\pi_1(X_0^*, x_0)$ and of the form $w = pq$, where the product pq already is a reduced word, then p is called a *prefix of* w. Similarly, a *prefix* of a monodromy element $u = \mu(r) \in T$ is an element $v = \mu(p) \in G$ where $p \in \pi_1(X_0^*, x_0)$ is a prefix of the corresponding singular loop r. We now define the *monodromy generating set* to be

$$S_\mu^\infty := S \cup \left\{ u^k \in G \;\middle|\; u \in T, k \in \mathbb{Z} \right\}.$$

With this infinite generating set $S_\mu^\infty \subset G$ we can now define a locally infinite Cayley graph similarly to Cayley graphs with finite generating sets. As for general graphs, we turn $\mathrm{Cay}(G, S_\mu^\infty)$ into a metric space by setting the length of each edge to 1. Furthermore, the generating set S_μ^∞ induces a word metric on G, denoted by d_∞. For simplicity we define $|g|_\infty := d_\infty(1_G, g)$. Note that both metrics are invariant under G-left multiplication. Similarly as for finitely generated Cayley graphs one can show: With respect to these metrics the inclusion $G \hookrightarrow \mathrm{Cay}(G, S_\mu^\infty)$ is a quasi-surjective isometric embedding and hence a quasi-isometry.

Remark. By the previous definition we have for all $g \in G$

$$|g|_\infty \leq |g|_S \text{ and } |g|_\infty \leq |g|_T^{syl},$$

since $S \subseteq S_\mu^\infty$ and $u^k \in S_\mu^\infty$ for each $u \in T$. Let r be a singular loop around $\sigma \in \Sigma$ which, as observed in Section 2.5, has word length d_σ with respect to the generating edge loops $\{c_1, \dots, c_m\}$. Moreover, let $u \in G$ its corresponding monodromy element and v a prefix of u. Then (4.1) implies

$$|v|_\infty \leq |v|_S \leq |u|_S \leq d_\sigma \leq 2m. \tag{4.2}$$

As the following theorem shows this such defined generating set is the right candidate in order to find a suitable quasi-isometry between X_μ and $\mathrm{Cay}(G, S_\mu^\infty)$.

Theorem 4.9. *Let X_μ be a G-regular translation surface. Fix a corner C in the base polygon P. Then the map Φ_C defined by*

$$\Phi_C \;:\; (G, d_\infty) \to (\overline{X}_\mu, d),\; g \mapsto [C, g],$$

is a quasi-isometry. In particular, X_μ is quasi-isometric to $\mathrm{Cay}(G, S_\mu^\infty)$.

Proof. The second assertion follows from the first one and the previous observation. The map Φ_C is quasi-surjective since any point in $\overline{X}_\mu$ lies in a copy $P \times \{g\}$ and thus has distance at most $\mathrm{diam}(P)$ from $\Phi_C(g) = [C, g]$. It remains to show that Φ_C is a quasi-isometric embedding, i.e. that there are constants $A' \geq 1$, $B' \geq 0$ only depending on P and μ such that for all $g, h \in G$ we have

$$\frac{1}{A'} \cdot d_\infty(h, g) - B' \leq d([C, h], [C, g]) \leq A' \cdot d_\infty(h, g) + B'.$$

Since both metrics are G-left invariant we can w.l.o.g. assume that $h = 1_G$. We proceed in three steps: In step 1 we prove the right-hand inequality. In step 2 we study saddle connections in more detail and finally in step 3 we prove the left-hand inequality.

Step 1:
Let $g \in G$ with $|g|_\infty = k$, i.e. we can write g as $g = g_1 \dots g_k$ where each g_i lies in S_μ^∞. Set $g_0 := 1_G$. Then, using the G-left invariance in the first equality, we get

$$\begin{aligned} d([C,1_G],[C,g]) &\leq \sum_{i=0}^{k-1} d([C,g_0\dots g_i],[C,g_0\dots g_{i+1}]) \\ &= \sum_{i=0}^{k-1} d([C,1_G],[C,g_{i+1}]) \\ &\leq \sum_{i=0}^{k-1} 2\cdot \operatorname{diam}(P) \\ &= 2\cdot\operatorname{diam}(P)\cdot k \\ &= 2\cdot\operatorname{diam}(P)\cdot |g|_\infty, \end{aligned}$$

where in the last inequality we used the fact that for $g_i \in S_\mu^\infty$ the copies $P \times \{1_G\}$ and $P \times \{g_i\}$ always share either a side if $g_i \in S$ or a corner if $g_i = u^k$ for some $u \in T$ by definition of S_μ^∞. Hence the distance is bounded by $2 \cdot \operatorname{diam}(P)$. See Figure 4.10 for an illustration.

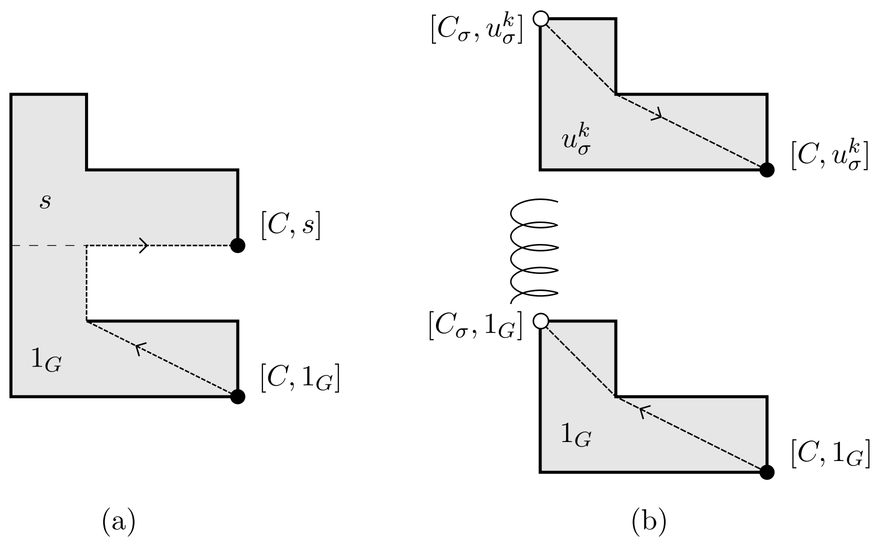

Figure 4.10: (a) Two copies sharing a side. (b) Two copies sharing a corner which is a singularity.

Step 2:
Let τ be an oriented saddle connection in $\overline{X}_\mu$. Let $h^{start} \in G$ be the copy of P it crosses first and $h^{end} \in G$ the one it crosses last. The starting point $\tilde{\sigma}$ of τ lies in $\operatorname{Sing}(X_\mu)$ and thus is of the form $[C', h]$ for some corner C' of P and $h \in G$. Let $u \in T$ be its corresponding monodromy element. Then, by cyclically moving around $\tilde{\sigma}$ starting from copy h, we see that h^{start} can be written as $hu^a v$, where $a \in \mathbb{Z}$ and v is a prefix of u. See Figure 4.11. Hence,

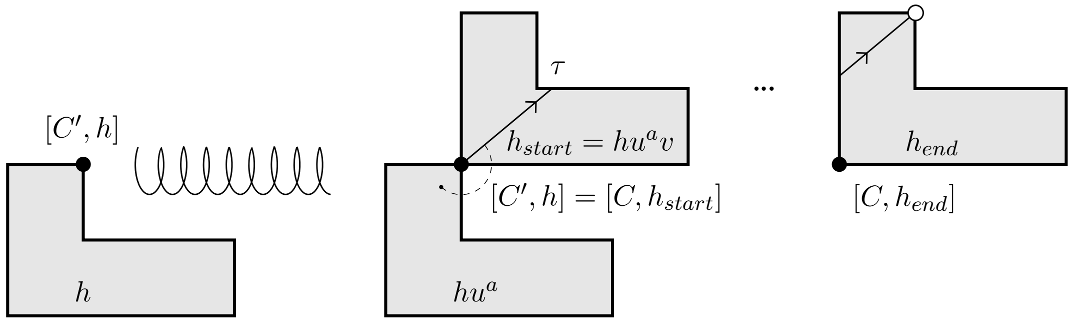

Figure 4.11: Given a saddle connection τ starting in the singularity $[C', h]$, we see that we can express the first crossed copy by τ with a term $hu^a v$, where v is a prefix of u.

using (4.2) we have

$$|h^{-1}h^{start}|_\infty \leq |u^a|_\infty + |v|_\infty = 1 + |v|_\infty \leq 1 + 2m.$$

By definition, the saddle connection τ crosses $N := N(\tau, P)$ copies of P. Hence there are generators $s_1, \ldots, s_{N-1} \in S$ such that $h^{end} = h^{start} s_1 \ldots s_{N-1}$. In case τ crosses only one copy, we have $N = 1$ and hence $h^{end} = h^{start}$. Note that

$$|s_1 \ldots s_{N-1}|_\infty \leq |s_1 \ldots s_{N-1}|_S \leq N - 1.$$

Step 3:
Proof of the left inequality. Let $g \in G$ be arbitrary and let c be a geodesic segment from $[C, 1_G]$ to $[C, g]$ in $\overline{X}_\mu$ realizing the length $d([C, 1_G], [C, g])$. Then c is a concatenation of finitely many saddle connections $\tau_1, \ldots, \tau_s$ and we know

$$d([C, 1_G], [C, g]) = \text{length}(c) = \sum_{i=1}^{s} \text{length}(\tau_i).$$

Note that since the covering $X_\mu^* \to X_0^*$ is locally isometric, the length of each saddle connection in $\overline{X}_\mu$ is bounded by below by the length $\text{minsc}(X_0) > 0$ of the shortest saddle connection in X_0. This constant only depends on the geometry of X_0. Hence $d([C, 1_G], [C, g]) \geq s \cdot \text{minsc}(X_0)$ or equivalently

$$s \leq \text{minsc}(X_0)^{-1} \cdot d([C, 1_G], [C, g]). \tag{4.3}$$

Now for each $i = 1, \ldots, s-1$ the saddle connection τ_{i+1} has starting and ending copies h_{i+1}^{start} and h_{i+1}^{end} such that the ending copy h_i^{end} of τ_i and h_{i+1}^{start} have a singularity of the form $[C_i', h_i^{end}]$ in common, see Figure 4.12. From step 2 it follows that $h_{i+1}^{start} = h_i^{end} u_i^{a_i} v_i$ for a corresponding monodromy element $u_i \in T$, $a_i \in \mathbb{Z}$ and v_i a prefix of u_i. And it also follows that we can write $h_i^{end} = h_i^{start} g_i$ for some $g_i \in G$ satisfying $|g_i|_\infty \leq N(\tau_i, P) - 1$, for all $i = 1, \ldots, s$. For the first and last saddle connection, note that the copy h_1^{start} has a singularity with the copy 1_G in common and similarly the copy h_s^{end} with the copy g. Hence we can write

$$h_1^{start} = u_0^{a_0} v_0 \text{ and } g = h_s^{end} u_s^{a_s} v_s,$$

for appropriate elements $u_0, u_s \in T$, prefixes v_0, v_s and $a_0, a_s \in \mathbb{Z}$. Altogether we can therefore

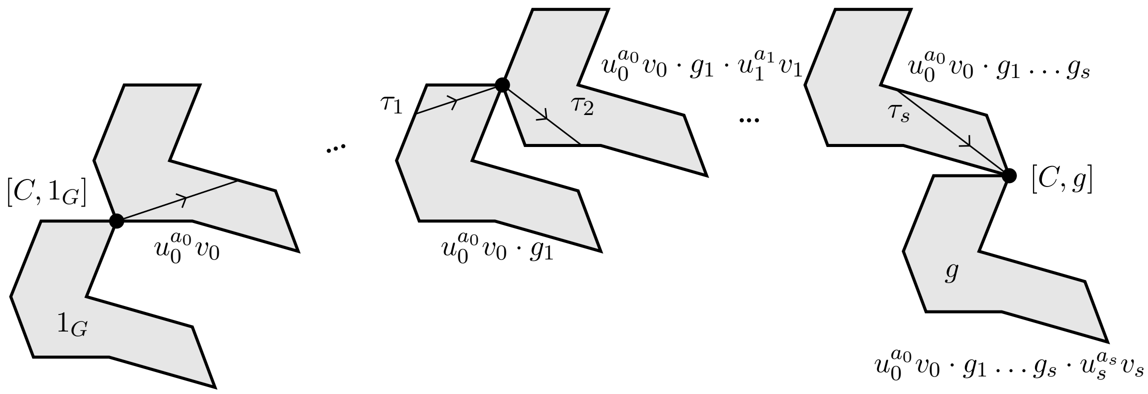

Figure 4.12: A geodesic segment between $[C, 1_G]$ and $[C, g]$ consisting of s saddle connections. Each saddle connection τ_i contributes to a factor g_i. Each singularity contributes to a factor $u_i^{a_i} v_i$.

write g as a product

$$g = u_0^{a_0} v_0 \cdot g_1 \cdot u_1^{a_1} v_1 \cdot g_2 \cdot \ldots g_s \cdot u_s^{a_s} v_s.$$

And using the inequalities from step 2 we get

$$\begin{aligned} |g|_\infty &\le \sum_{i=0}^{s} |u_i^{a_i} v_i|_\infty + \sum_{i=1}^{s} |g_i|_\infty \\ &\le (s+1) \cdot (1 + 2m) + \sum_{i=1}^{s} (N(\tau_i, P) - 1) \\ &= 2m \cdot s + 2m + 1 + \sum_{i=1}^{s} N(\tau_i, P) \end{aligned}$$

Now, Lemma 4.7 provides us with constants $A \geq 1$, $B \geq 0$ such that for all $i = 1, \ldots, s$ we have

$$N(\tau_i, P) \leq A \cdot \text{length}(\tau_i) + B.$$

This yields

$$\begin{aligned} |g|_\infty &\le 2m \cdot s + 2m + 1 + \sum_{i=1}^{s} (A \cdot \text{length}(\tau_i) + B) \\ &= (2m + B)s + 2m + 1 + A \cdot \sum_{i=1}^{s} \text{length}(\tau_i) \\ &= (2m + B)s + 2m + 1 + A \cdot d([C, 1_G], [C, g]). \end{aligned}$$

We now use inequality (4.3) for s and get

$$\begin{aligned} d_\infty(1_G, g) &= |g|_\infty \\ &\leq (2m+B)\mathrm{minsc}(X_0)^{-1} \cdot d([C,1_G],[C,g]) + 2m + 1 + A \cdot d([C,1_G],[C,g]) \\ &= d([C,1_G],[C,g]) \cdot \left(\frac{2m+B}{\mathrm{minsc}(X_0)} + A\right) + 2m + 1 + A. \end{aligned}$$

Note that all appearing values A, B, m and $\mathrm{minsc}(X_0)$ only depend on X_0 and the choice of polygon P for X_0, and hence are independent from the points $\Phi_C(1_G)$ and $\Phi_C(g)$. This concludes our proof. □

Remark. Using this theorem we are able to give an alternative proof of Theorem 4.1: If X_μ only has conical singularities, then Lemma 3.3 implies that all monodromy elements in T have finite order. In particular the set $S_\mu^\infty = S \cup \{u^k | u \in T, k \in \mathbb{Z}\}$ is a finite generating set for G and our theorem implies that X_μ is quasi-isometric to the finitely generated Cayley graph $\mathrm{Cay}(G, S_\mu^\infty)$ and hence to the finitely generated group G.

The proofs of Lemma 4.7 and Theorem 4.9 give us concrete bounds for the metrics in some special cases, for example regular coverings of origamis. An *origami* or *square-tiled surface* O is a connected finite translation surface consisting of a glued family of Euclidean unit squares. For example, we can realize an origami from a right-angled Euclidean polygon having integer side lengths. We call it a *d-origami* if it consists of d squares. For an introduction to origamis, see [Sch04], and their relation to Teichmüller space, see [Her12].

Corollary 4.10. *Let X_μ be a G-regular translation surface covering a d-origami O which is obtained from a right-angled Euclidean polygon having integer side lengths. Let $|\cdot|_\infty$ be the word metric on G induced by μ as defined in Definition 4.8. Then for all $g \in G$ we have the inequalities:*

$$\frac{1}{4d} \cdot d([C,1_G],[C,g]) \leq |g|_\infty \leq (4d+2) \cdot d([C,1_G],[C,g]) + 4d + 3.$$

In particular, if X_μ is a G-regular square-tiled surface, i.e. if O is simply the square torus, then

$$\frac{1}{4} \cdot d([C,1_G],[C,g]) \leq |g|_\infty \leq 6 \cdot d([C,1_G],[C,g]) + 7.$$

Proof. Given an Origami O with d squares and a right-angled polygon P for it, one easily computes the values

$$m \leq d+1,\ h_{min} \geq 1,\ w_{min} \geq 1,\ \mathrm{minsc}(O) \geq 1.$$

Note that the inequalities may occur when the base surface does not contain singularities lying in the inner of P. Furthermore as the diameter of one unit square is $\sqrt{2}$ we get $\mathrm{diam}(P) \leq \sqrt{2}d \leq 2d$. We will from now on use the larger bound $2d$ for a simpler presentation of the bounds. However, one may insert the bound $\sqrt{2}d$ for a closer approximation. The proof of Lemma 4.7 implies that

$$\text{length}(\tau) \leq d\sqrt{2} \cdot N(\tau, P) \leq 2d \cdot N(\tau, P),$$

giving us constants $A = 2d$ and $B = 0$. Inserting all these values into the inequalities of Theorem 4.9 proves the first assertion. Setting $d = 1$ proves the last one. $\square$

Example. We return to our example of the AB-surface $X_{(a,b)}$ presented in Section 3.6. As observed before, all singular loops have either monodromy a or b, hence the infinite generating set S_μ^∞ is of the form

$$S_\mu^\infty = \{a^k, b^k | k \in \mathbb{Z}\}.$$

From Proposition 3.8 and Theorem 4.9 it follows then that the following three metric spaces lie in the same quasi-isometry class:

$$X_{(a,b)} \sim_{QI} T_\infty \sim_{QI} Cay\left(F(a,b), \{a^k, b^k | k \in \mathbb{Z}\}\right).$$

Figure 4.13 illustrates the locally infinite Cayley graph $Cay\left(F(a,b), S_\mu^\infty\right)$. As an example, the distance between the vertices 1 and $a^3b^{-4}ab^2$ is 4.

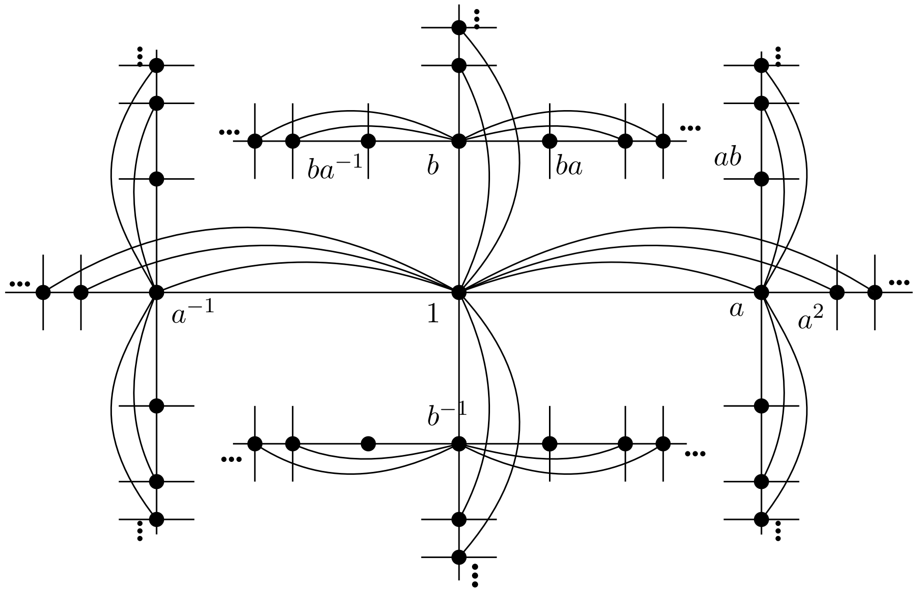

Figure 4.13: The locally infinite Cayley graph of $F(a,b)$ with respect to the infinite generating set $\{a^k, b^k | k \in \mathbb{Z}\}$. It is quasi-isometric to the infinite regular tree T_∞.

4.4 The space of ends of regular translation surfaces

In this section we study the ends of a G-regular translation surface X_μ and relate the number of ends with the number of ends in G. By the definition given in Section 2.8 an end is an equivalence class of a proper ray $r : [0, \infty) \to X_\mu$. For this purpose it is necessary to characterize compact subsets in X_μ.

Lemma 4.11. *Let $K \subseteq X_\mu$ where X_μ is of the form $(P \times G)\backslash_{\sim_\mu}$. Then K is compact if and only if K is closed in $\overline{X}_\mu$ and intersects only finitely many copies of P.*

Proof. If K is compact then it is also closed. Note that since K lies in X_μ, by defintion of X_μ it only contains conical singularities but no ∞-angle singularities. Assume that it intersects infinitely many copies of P. In particular, we find a sequence $(g_i)_{i\in\mathbb{N}}$ of pairwise distinct elements in G such that K intersects the interior of each copy $P \times \{g_i\}$. For each such intersection choose a point $x_i \in X_\mu$ contained in it. Then the sequence $(x_i)_{i\in\mathbb{N}} \subset K$ is an infinite sequence lying discretely in X_μ, a contradiction to K being compact.

In return let K be closed in $\overline{X}_\mu$ and intersect finitely many copies of P. Let $(x_i)_{i\in\mathbb{N}}$ be any sequence in K. The finiteness condition implies that there must be at least one copy $P \times \{g\}$ containing an infinite subsequence of $(x_i)_{i\in\mathbb{N}}$. This copy is compact as it is isometric to a Euclidean polygon, hence the infinite subsequence has an accumulation point x in $P \times \{g\}$ and therefore the whole sequence $(x_i)_{i\in\mathbb{N}}$ has the accumulation point x in K. The latter is equivalent to K being compact. □

Example. Let us compute the number of ends in the example of the 2- resp. 3-staircase X_2 resp. X_3. Using Lemma 4.11 any compact set $K \subset X_2$ intersects finitely many copies of the base polygon. Therefore there is a minimal distance $\varepsilon > 0$ between K and any ∞-angle singularity, see Figure 4.14 (a). This implies that the complement $X_2 \backslash K$ is connected: Any two points in $X_2 \backslash K$ can be connected via a path which spirals around an ∞-angle singularity at a distance less than ε if necessary. This shows that the bounded translation surface X_2 has one end.

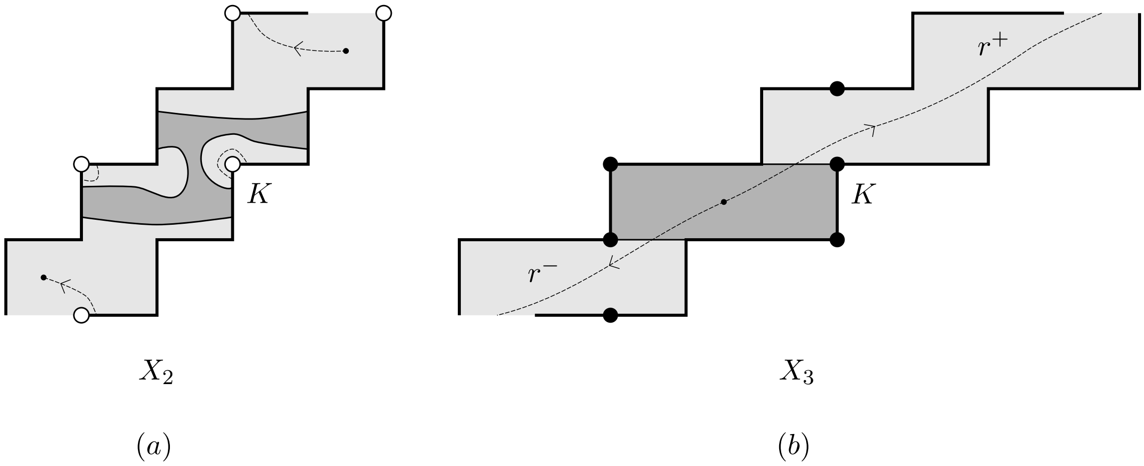

Figure 4.14: (a) The 2-staircase has one end. (2) The 3-staircase has two ends given by the equivalence classes of the rays r^+ and r^-.

Remark. The example of the 2-staircase is very insightful for the study of Cayley graphs with respect to infinite generating sets. Computing the monodromy subgroup of $\mathbb{Z}$ in the 2-staircase example results in $U = 2\mathbb{Z}$. Therefore, by Theorem 4.5 the surface X_2 is quasi-isometric to $\mathbb{Z}/2\mathbb{Z}$ and hence bounded. On the other hand Theorem 4.9 implies that X_2 is quasi-isometric to the infinitely generated Cayley graph $\mathrm{Cay}(\mathbb{Z}, 2\mathbb{Z})$ which consequently is bounded as well. The discussion in the previous example shows that $\mathrm{Cay}(\mathbb{Z}, 2\mathbb{Z})$ as well has exactly one end. This follows from the fact that any compact set K of $\mathrm{Cay}(\mathbb{Z}, 2\mathbb{Z})$ only contains finitely many vertices which therefore lie in an interval of the form $[-n, n]$, $n \in \mathbb{N}$. But then the vertices $-(n+1)$ and $n+1$ are adjacent and hence $\mathrm{Cay}(\mathbb{Z}, 2\mathbb{Z})\backslash K$ has one connected component and thus one end.

Altogether we see that for a finitely generated group the locally infinite Cayley graphs, i.e. those with respect to infinite generating systems, in general behave very different from locally finite ones. As observed for $G = \mathbb{Z}$ in general they are not quasi-isometric to each other and do not have the same number of ends.

On the other hand, consider the subset $K \subset X_3$ which consists of one copy of the closed base polygon inside X_3. Since all singularities are conical this copy indeed is compact. The resulting complement $X_3\backslash K$ has two unbounded path components as there is no path between points "above" and "below" the deleted copy, see Figure 4.14 (b). Each such unbounded path component corresponds to one end of X_3 and therefore X_3 has two ends.

Definition 4.12. Let $\tilde{\sigma}$ be an ∞-angle singularity of X_μ. Fix any point $\tilde{x} \in X_\mu$. We define three proper rays

$$r_{\tilde{\sigma}}, r_{\tilde{\sigma}}^{+}, r_{\tilde{\sigma}}^{-} \;:\; [0, \infty) \to X_\mu$$

as follows. Let $r_{\tilde{\sigma}}$ be a ray with $r_{\tilde{\sigma}}(0) = \tilde{x}$ which ends in the singularity $\tilde{\sigma}$, i.e. the closure of $r_{\tilde{\sigma}}([0, \infty))$ in $\overline{X}_\mu$ contains the additional point $\tilde{\sigma}$. And let $r_{\tilde{\sigma}}^{+}$ resp. $r_{\tilde{\sigma}}^{-}$ be rays with $r_{\tilde{\sigma}}^{+}(0) = r_{\tilde{\sigma}}^{-}(0) = \tilde{x}$ that spiral sufficiently close around $\tilde{\sigma}$ in counterclockwise resp. clockwise direction. Note that $r_{\tilde{\sigma}}$ is a proper ray: Each compact set $K \subset X_\mu$ intersects only finitely many copies of the base polygon and hence there is a minimal distance $\varepsilon = \varepsilon(K) > 0$ between K and any ∞-angle singularity, see the example of the 2-staircase. This means that for each compact K there is a time $T > 0$ such that $r_{\tilde{\sigma}}(T)$ has distance less than ε from $\tilde{\sigma}$ and hence the ray $r_{\tilde{\sigma}}$ leaves the compact K. See Figure 4.15 for an illustration of the three rays.

As in the 2-staircase example we see that all three rays are equivalent since for each compact set $K \subset X_\mu$ there is an $\varepsilon > 0$ such that the ε-disc around $\tilde{\sigma}$ does not intersect K. And since K is contained in the union of only finitely many copies of the base polygon, each spiraling ray leaves K after some finite time.

We denote the common end of these three rays by $\mathrm{end}(\tilde{\sigma}) \in \mathrm{Ends}(X_\mu)$ and say that this is the *end of the singularity* $\tilde{\sigma}$. Notice that distinct ∞-angle singularities do not necessarily induce distinct ends of singularities. As an example, the 2-staircase has four ∞-angle singularities that all induce one and the same end of X_2.

If we recall the construction pattern of the metric space $\overline{X}_\mu$ as gluing G copies of the original polygon P along parallel edges we might recognize similarities with the way edges were defined in the Cayley graph $\Gamma := \mathrm{Cay}(G, S)$ in Section 2.7. Firstly fix a point x_0 in X_0^* and embed the wedge product of m circles into X_0 using the m closed edge paths based in x_0 whose homotopy class is given by the generators c_i of $\pi_1(X_0^*, x_0)$. This wedge product is the embedding of

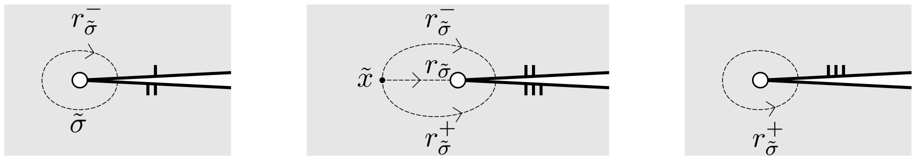

Figure 4.15: Three equivalent proper rays in X_μ whose equivalence class is the end $\mathrm{end}(\tilde{\sigma})$ of the ∞-angle singularity $\tilde{\sigma}$.

a graph on one vertex and m loop edges into X_0. Define $\tilde{\Gamma} \subset X_\mu$ as the preimage of that embedded graph under the G-covering p. It has vertex set $\{x_0\} \times G \cong G$ and a path between $[x_0, g]$ and $[x_0, h]$ whenever g and h are connected by an edge in Γ. Note that this embedding of Γ into X_μ is quasi-surjective since any point in X_μ is at most $\mathrm{diam}(P) < \infty$ away from some vertex $[x_0, g]$. This embedding, denoted by $\iota : \Gamma \to X_\mu$, will turn out helpful for describing ends of X_μ. See Figure 4.16 for an example.

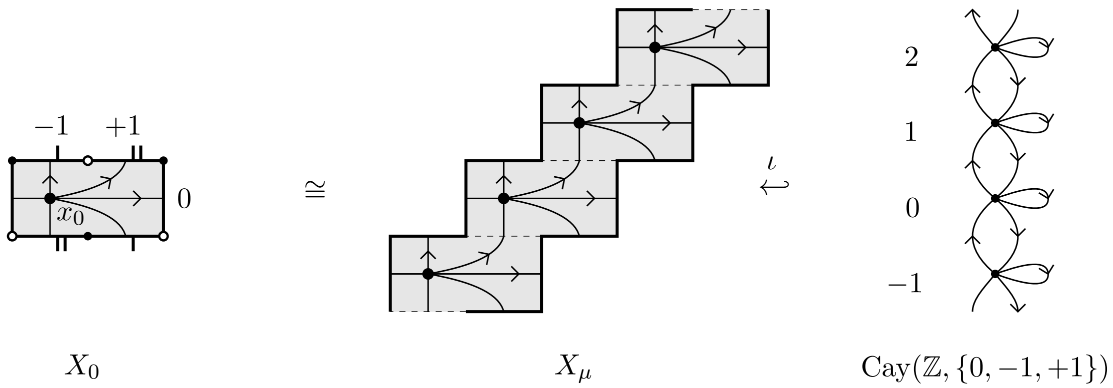

Figure 4.16: Embedding a Cayley graph of $\mathbb{Z}$ into the $\mathbb{Z}$-regular 2-staircase X_2.

Theorem 4.13. *The embedding ι induces a surjective map*

$$\begin{aligned} \iota_E \ : \ \mathrm{Ends}(\Gamma) &\twoheadrightarrow \mathrm{Ends}(X_\mu), \\ \mathrm{end}(r) &\mapsto \mathrm{end}(\iota(r)). \end{aligned}$$

Remark. In other words, Theorem 4.13 states that all ends of a G-regular translation surface X_μ can be described by $\mathrm{Ends}(G)$. However the map ι_E is not injective in general. For example if X_μ does contain an ∞-angle singularity the two corresponding clockwise and counterclockwise rays around it yield one and the same end in $\mathrm{Ends}(X_\mu)$.

Proof. We firstly prove that the map ι_E is well-defined. Let $r, r' : [0, \infty) \to \Gamma$ be two proper rays with $\mathrm{end}(r) = \mathrm{end}(r')$. We need to show that $\mathrm{end}(\iota(r)) = \mathrm{end}(\iota(r'))$. Note that ι is a proper map, hence the rays $\iota(r)$ and $\iota(r')$ are again proper rays. For this, let $K \subset X_\mu$ be a compact subset. Define $K_0 := \iota^{-1}(K)$ which is a compact subset of Γ since ι is proper. Since r and r' are equivalent there is $T > 0$ such that $r([T, \infty))$ and $r'([T, \infty))$ lie in the same path component of $\Gamma\backslash K_0$. Applying ι this implies that $\iota(r)([T,\infty))$ and $\iota(r')([T,\infty))$ lie in the same path component of $\iota(\Gamma)\backslash K$, and in particular of $X_\mu\backslash K$. Hence $\iota(r)$ and $\iota(r')$ are equivalent.

We now prove the surjectivity of ι_E. Consider $\mathrm{end}(r) \in \mathrm{Ends}(X_\mu)$ for a proper ray $r : [0,\infty) \to X_\mu$. We distinguish two cases.

Case 1: $\forall g \in G, T > 0\ \exists t > T : r(t) \notin P \times \{g\}$, i.e. the ray exits each polygon it enters. We show in two steps that we can homotope r into the image $\iota(r_0)$ of a proper ray r_0 in Γ.
Step 1: Suppose r contains a conical singularity $\tilde{\sigma}$ of X_μ. Then there is a copy $P \times \{g\}$ where r enters the singularity and a copy $P \times \{h\}$ where it leaves it. We can then homotope r locally around $\tilde{\sigma}$ into r' such that r' forms the segment of a circle of sufficiently small radius, see Figure 4.17 (a).
Step 2: Suppose r does not contain any singulariy of X_μ, otherwise apply Step 1. This means that whenever r crosses a copy $P \times \{g\}$ it must intersect an entering and an exiting edge of P. More precisely, let $r(t_0) < r(t_1) \in (\partial P) \times \{g\}$ be the entering and exiting point of r. Then $r([t_0,t_1])$ is a path lying inside the copy $P \times \{g\}$. Since P is simply connected we can locally homotope $r|_{[t_0,t_1]}$ into a segment $\iota(r_0)|_{[t_0',t_1']}$ of the image of a ray r_0 in Γ. See Figure 4.17 (b) for an illustration of the homotopy. Altogether we can combine this to a homotopy between the rays r and $\iota(r_0)$. Since both rays always "travel together", i.e. cross exactly the same copies of P they cannot be separated by an compact set $K \subset X_\mu$. It follows that both rays are equivalent and we have

$$\mathrm{end}(r) = \mathrm{end}(\iota(r_0)) = \iota_E(\mathrm{end}(r_0)).$$

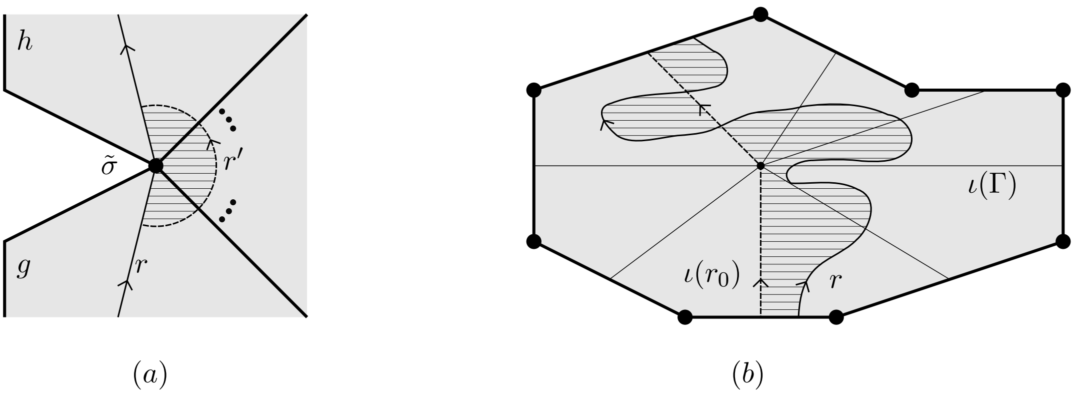

Figure 4.17: (a) A ray r crossing a conical singularity can be homotoped to r' only crossing edges of P. (b) A ray r only crossing edges of P can be homotoped to a ray of the form $\iota(r_0)$ where r_0 is a ray in Γ.

Case 2: $\exists g \in G, T > 0 : r([T,\infty))$ lies inside one copy $P \times \{g\}$, i.e. the ray r ends up in a polygon and never leaves it. Then, as r is proper, it must converge to a corner C of $P \times \{g\}$ which necessarily represents an ∞-angle singularity in X_μ. More precisely we can construct compact sets $K \subset P \times \{g\}$ as illustrated in Figure 4.18, and we see that r must end in one path component of $(P \times \{g\}) \backslash K$ and thus converge to a corner of the polygon. Hence the point $\tilde{\sigma} = [C, g]$ is an ∞-angle singularity and we have $\mathrm{end}(r) = \mathrm{end}(\tilde{\sigma})$ is

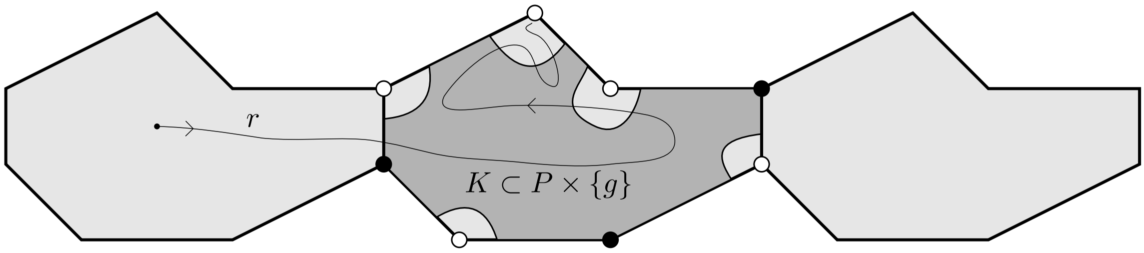

Figure 4.18: As r remains inside K for a finite time, it must end up converging to a white ∞-angle singularity. Black singularities represent conical singularities.

the end of $\tilde{\sigma}$. Consider a proper ray r^- that spirals clockwise around $\tilde{\sigma}$ as constructed in Definition 4.12. It follows that the ray r is equivalent to r^-. Since $\tilde{\sigma}$ is an ∞-angle singularity this clockwise ray exits each polygon it enters. Hence we can apply Case 1 to r^- and we see that r^- is equivalent to a proper ray of the form $\iota(r_0)$ for a ray r_0 in Γ. Hence

$$\mathrm{end}(r) = \mathrm{end}(r^-) = \mathrm{end}(\iota(r_0)) = \iota_E(\mathrm{end}(r_0)).$$

□

Corollary 4.14. *Let X_μ be a G-regular translation surface having only conical singularities. Then* $\mathrm{Ends}(G)$ *and* $\mathrm{Ends}(X_\mu)$ *are homeomorphic.*

Proof. Since X_μ does not have ∞-angle singularities it is a proper geodesic space and Theorem 4.1 implies that there is a quasi-isometry between X_μ and G with respect to a finite generating system. By Lemma 2.8, this quasi-isometry induces a homeomorphism between the corresponding spaces of ends. □

Corollary 4.15. *Let X_μ be a G-regular translation surface having at least one ∞-angle singularity. If G has finitely many ends, then X_μ has exactly one end.*

Proof. Since X_μ has at least one ∞-angle singularity this implies that G is an infinite group and thus has $1, 2$ or infinitely many ends, see Lemma 2.9. If it has finitely many ends it therefore must have 1 or 2 ends. If G has one end there is nothing to prove since ι_E is a surjection and X_μ an unbounded space.

Now suppose that G has two ends. By assumption X_μ has an ∞-angle singularity $\tilde{\sigma}$ which by Lemma 3.4 is of the form $[C, g\langle u\rangle]$ for a corner C of the base polygon, a monodromy element u for the conical singularity of the base surface corresponding to C and some element $g \in G$. Moreover, by Lemma 3.3 the monodromy element u has infinite order in G. Now consider the clockwise resp. counterclockwise proper rays $r^+_{\tilde{\sigma}}$ resp. $r^-_{\tilde{\sigma}}$ around $\tilde{\sigma}$ together yielding one end $\mathrm{end}(\tilde{\sigma})$ as described in Definition 4.12. We can assume that both rays start in the copy g of the base polygon P. Then $r^+_{\tilde{\sigma}}$ resp. $r^-_{\tilde{\sigma}}$ crosses the copies $P \times \{gu^m\}$ resp. $P \times \{gu^{-m}\}$ for $m \in \mathbb{N}_0$. Applying Case 1 from the proof of Theorem 4.13 shows that the rays $r^\pm_{\tilde{\sigma}}$ can be homotoped into rays $\iota(r^{\pm})$ which are images of proper rays $r^{\pm}$ in $\Gamma := \mathrm{Cay}(G, S)$ for a choice of finite generating system S for G. It follows that the ray r^+ resp. r^- contains all vertices of the form gu^m resp. gu^{-m} for $m \in \mathbb{N}_0$ and therefore r^+ and r^- induce two distinct ends in G. As G only has two ends by assumption and both ends are mapped to the one end $\mathrm{end}(\tilde{\sigma})$ of X_μ under the surjective map ι_E it then follows that X_μ can only have one end. □

Summarizing our observations from Theorem 4.6 and Corollary 4.15 so far in the case of $\mathbb{Z}$-regular translation surfaces we get the following result.

Corollary 4.16. *Let X be a $\mathbb{Z}$-regular translation surface. If X contains at least one ∞-angle singularity, then it is a bounded metric space having finitely many ∞-angle singularities and exactly one end. If X only has conical singularities, then it is quasi-isometric to $\mathbb{Z}$ and has two ends.*

Note that the converse statement of Corollary 4.15 is false in general. The following example presents a G-regular translation surface X_μ where G has infinitely many ends but X_μ only one.

Example. Consider the $F(a, b)$-regular translation surface X_μ as shown in Figure 4.19. It is a common fact that the free group $F(a, b)$ has infinitely many ends. However, we prove that X_μ is a simply connected non-compact surface and hence homeomorphic to an open disc. The open disc has exactly one end corresponding to a ray inside it converging to the boundary. This implies that X_μ also has one end.

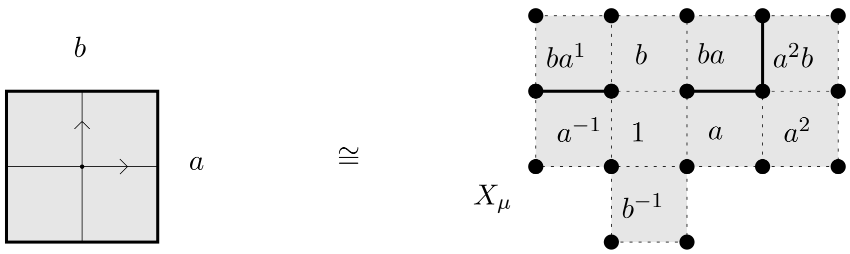

Figure 4.19: A G-regular translation surface having only one end while $G = F(a, b)$ has infinitely many ends.

Proof that X_μ is simply connected: Similar to our previous examples X_μ is a non-compact surface only having ∞-angle singularities. We have the natural embedding

$$\iota \ : \ \mathrm{Cay}(F(a, b), \{a, b\}) =: \Gamma \to X_\mu$$

which in this case is a homotopy equivalence. Its homotopy inverse is the map $\rho_\iota \ : \ X_\mu \twoheadrightarrow \Gamma$ illustrated in Figure 4.20, which projects any point of X_μ onto $\iota(\Gamma)$. Hence Γ and X_μ are homotopy equivalent. More precisely, $\iota(\Gamma) \subset X_\mu$ is a deformation retract. This homotopy equivalence induces an isomorphism between fundamental groups $\pi_1(X_\mu) \cong \pi_1(\Gamma)$, the last one being trivial since the Cayley graph Γ is isomorphic to the simply connected 4-regular tree T_4. Therefore, X_μ is as well simply connected.

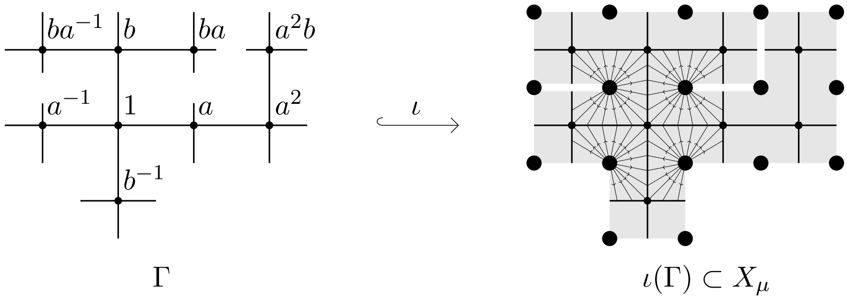

Figure 4.20: The map ρ_ι which projects any point of X_μ to $\iota(\Gamma)$ and via ι^{-1} to Γ is a homotopy equivalence with homotopy inverse ι.

We end this discussion with a last example which shows that when $|\mathrm{Ends}(G)| = \infty$, the space of ends $\mathrm{Ends}(X_\mu)$ may be more complex to describe: We again consider the AB-surface $X_{(a,b)}$ already introduced in Section 3.6. Into $X_{(a,b)}$ we embed the Cayley graph $\Gamma := \mathrm{Cay}(F(a,b), \{a, b, 1, 1, 1\})$ as described before and denote this embedding by $\iota : \Gamma \to X_{(a,b)}$. Note that Γ is isomorphic to a 4-regular tree with three loops of unit length attached to each vertex in $F(a,b)$. Therefore, each end of Γ is the equivalence class of a unique geodesic ray starting in 1. Moreover, each such ray corresponds to a unique infinite sequence of words in $F(a,b)$ whose word length increases by 1, namely these are exactly the vertices of Γ the ray is crossing. As an example the horizontal ray in Γ starting in 1 and moving to the right along a-edges corresponds to the sequence $(a^m)_{m\in\mathbb{N}_0}$. For simplicity, we shortly identify the ends of Γ with its geodesic rays starting in 1. This way we call an end (or ray) of Γ *horizontal* if it is of the form $(ga^{\pm m})_{m\in\mathbb{N}_0}$ and *vertical* if it is of the form $(gb^{\pm m})_{m\in\mathbb{N}_0}$ for $g \in G$, and *alternating* otherwise. Two horizontal resp. vertical ends (or rays) are *opposite* if they are of the form $(ga^m)_{m\in\mathbb{N}_0}$ and $(ga^{-m})_{m\in\mathbb{N}_0}$ resp. $(gb^m)_{m\in\mathbb{N}_0}$ and $(gb^{-m})_{m\in\mathbb{N}_0}$. See Figure 4.21 for an example.

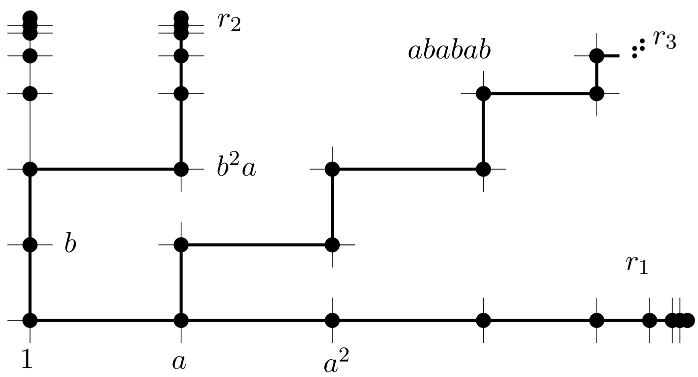

Figure 4.21: Three possible ends in Γ. The ray $r_1 = (a^m)_{m\in\mathbb{N}_0}$ is horizontal, $r_2 = (b^2ab^m)_{m\in\mathbb{N}_0}$ is vertical and $r_3 = ((ab)^m)_{m\in\mathbb{N}_0}$ is alternating.

Corollary 4.17. *The AB-surface $X_{(a,b)}$ has infinitely many ends. More precisely, we have a bijection* $\mathrm{Ends}(X_{(a,b)}) \cong \mathrm{Ends}(F(a,b))/_\sim$ *where each horizontal end $(ga^m)_{m\in\mathbb{N}_0}$ is identified with its opposite $(ga^{-m})_{m\in\mathbb{N}_0}$ and each vertical end $(gb^m)_{m\in\mathbb{N}_0}$ is identified with its opposite*

$(gb^{-m})_{m\in\mathbb{N}_0}$.

Proof. By Theorem 4.13 we have a surjective map

$$\iota_E \ : \ \mathrm{Ends}(\Gamma) \twoheadrightarrow \mathrm{Ends}(X_{(a,b)}), \ \mathrm{end}(r) \mapsto \mathrm{end}(\iota(r)).$$

Let us describe the ends of singularities of $X_{(a,b)}$: By Lemma 3.4 we have the bijection

$$\mathrm{Sing}^{\infty}(X_{(a,b)}) \cong F(a,b)/A \sqcup F(a,b)/A \sqcup F(a,b)/B \sqcup F(a,b)/B,$$

where $A := \langle a\rangle$ and $B := \langle b\rangle$. Therefore all singularities of the form gA induce one end in $\mathrm{Ends}(X_{(a,b)})$ denoted by $\mathrm{end}(gA)$. The image under ι of a horizontal ray $(ga^m)_{m\in\mathbb{N}_0}$ resp. its opposite $(ga^{-m})_{m\in\mathbb{N}_0}$ are equivalent to the clockwise resp. counterclockwise rays in $X_{(a,b)}$ that form the end $\mathrm{end}(gA)$ of the singularity gA. This follows from the fact that all these rays cross exactly the same copies $P\times\{ga^m\}$ resp. $P\times\{ga^{-m}\}$ for $m\in\mathbb{N}_0$, and hence cannot be separated by any compact set. In other words, the two opposite ends $(ga^m)_{m\in\mathbb{N}_0}$ and $(ga^{-m})_{m\in\mathbb{N}_0}$ are mapped to one end $\mathrm{end}(gA)$ under ι_E. An analogue statement holds for ends $\mathrm{end}(gB)$ of singularities gB. This observation implies that the map ι_E descends to a well-defined surjective map

$$\iota_E^* \ : \ \mathrm{Ends}(\Gamma)/_\sim \twoheadrightarrow \mathrm{Ends}(X_{(a,b)}).$$

We prove that this map is also injective by showing that distinct ends in $\mathrm{Ends}(\Gamma)/_\sim$ are mapped to distinct ends in $X_{(a,b)}$. Consider two distinct ends of Γ that are not opposite. They correspond to two geodesic rays r, r' in Γ starting in 1. Let $v\in F(a,b)$ be the last vertex both rays have in common. W.l.o.g. $v=1$, otherwise apply a left multiplication by v^{-1} on Γ. Consider the bi-infinite path $R : \mathbb{R}\to\Gamma$ whose image is the union of the images of r and r'. Since r and r' are not opposite rays this new path R is not of the form $(a^m)_{m\in\mathbb{Z}}$ or $(b^m)_{m\in\mathbb{Z}}$. Therefore there is a vertex f in the path R which is adjacent to a vertex in R of the form $fa^{\pm1}$ and to a vertex in R of the form $fb^{\pm1}$. Geometrically spoken, the path R "has a corner in f" in the Cayley graph Γ. Now define the compact set $K_f\subset X_{(a,b)}$ as shown in Figure 4.22 which is a subset of the copy $P\times\{f\}$. Note that $X_{(a,b)}\backslash K_f$ has two path components, one containing the singularities fA and one containing fB. Since r' and r only have the vertex 1 in common, one of the components contains $r'([T,\infty])$ while the other component contains $r[(T,\infty)]$ for sufficiently large $T>0$. Thus we have that K_f separates the embedded rays $\iota(r)$ and $\iota(r')$. This implies that the ends $\mathrm{end}(\iota(r)) = \iota_E(\mathrm{end}(r))$ and $\mathrm{end}(\iota(r')) = \iota_E(\mathrm{end}(r'))$ are distinct. □

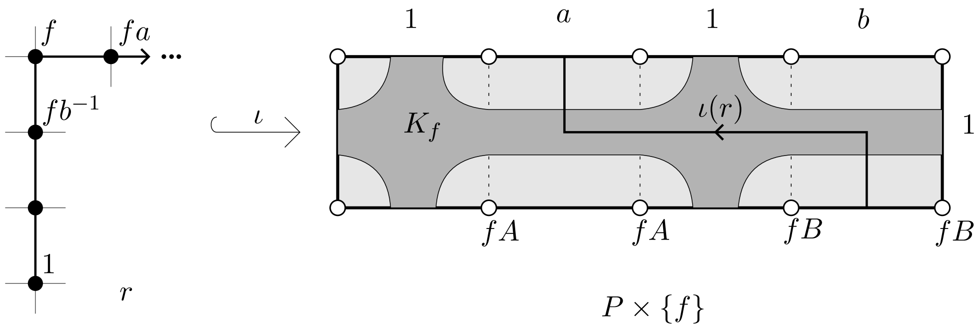

Figure 4.22: The ray r in Γ has a corner at the vertex f. Hence $\iota(r)$ must cross the compact set K_f in X_μ. So, K_f separates r and r'.

CHAPTER 5

Applications of Theorem 4.9

5.1 A generalization of the surface $X_{(a,b)}$

Theorem 4.9 allows us to compute the quasi-isometry class of a whole family of regular translation surfaces. Namely, the study of the AB-surface $X_{(a,b)}$ can be extended to the family of regular translation surfaces which are obtained in a similar fashion but having as deck transformation group any free group on $n \geq 2$ generators. After defining this family properly we will show that all these surfaces obtained this way are quasi-isometric to the infinite regular tree T_∞. This argument furthermore gives us an alternative proof of Proposition 3.8. In this section we make use of common notions in graph theory. For a detailed introduction to graph theory, see [Die18].

Definition 5.1. Let $n \geq 2$ and consider $F_n := F(a_1, \ldots, a_n)$, the free group on n generators $a_1, \ldots, a_n$. Similar to the definition of $X_{(a,b)}$ we define the F_n-regular translation surface $X_{F_n} := X_{(a_1,\ldots,a_n)}$ as shown in Figure 5.1. Furthermore, we define the graph $T_{n,\infty}$ which is the

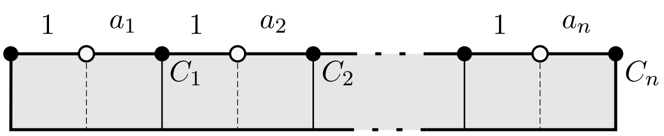

Figure 5.1: The $F(a_1, \ldots, a_n)$-regular translation surface X_{F_n}.

bipartite (n, ∞)*-regular tree*, i.e. it is a tree whose vertices either have valence n or ∞ and such that any vertex of valence n is only adjacent to vertices of valence ∞ and vice versa. An illustration of $T_{5,\infty}$ is given in Figure 5.2.

As we will see in this section, all three spaces X_{F_n}, $T_{n,\infty}$ and T_∞ are quasi-isometric. In

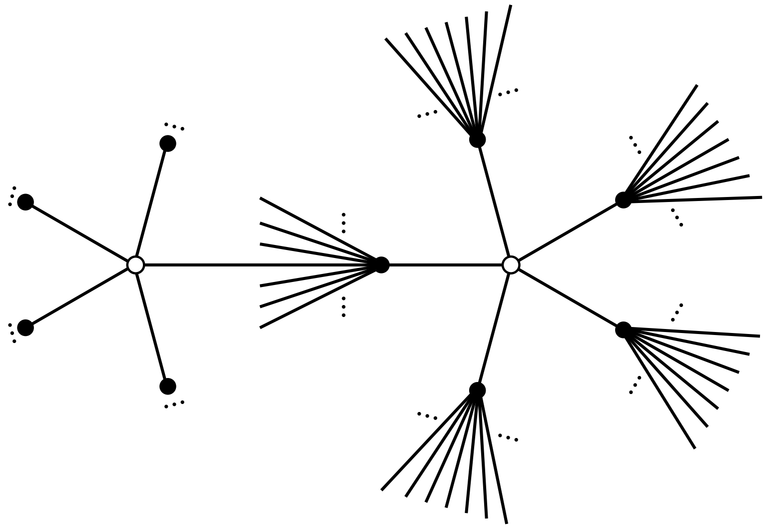

Figure 5.2: The bipartite $(5, \infty)$-regular tree $T_{5,\infty}$. Vertices of infinite valence are marked by black dots.

particular, for all $n, m \geq 2$ we have that the regular translation surfaces X_{F_n} and X_{F_m} are quasi-isometric. Before we can prove this we need to study the coarse geometry of infinite graphs in more detail. Let $\Gamma = (V, E)$ be an infinite graph. A *block cover* $\mathcal{B} = \{B_i \mid i \in I\}$ on Γ is a collection of complete subgraphs B_i of size ≥ 2, called *blocks*, of Γ such that each edge of Γ is contained in at least one block B_i. The *star graph* $\mathcal{S} := \mathcal{S}(\Gamma, \mathcal{B})$ of Γ with respect to $\mathcal{B}$ is a graph constructed as follows: The vertex set of $\mathcal{S}$ is $V(\Gamma) \sqcup \mathcal{B}$ and there are only edges between vertices $v \in V$ and $B \in \mathcal{B}$ if and only if v lies in the block B. This makes $\mathcal{S}$ a bipartite graph with vertex partitions $V(\Gamma)$ and $\mathcal{B}$. The name "star graph" comes from the fact that when constructing $\mathcal{S}$ from Γ each block B is replaced by a star graph having B as central vertex, see Figure 5.3 for an example.

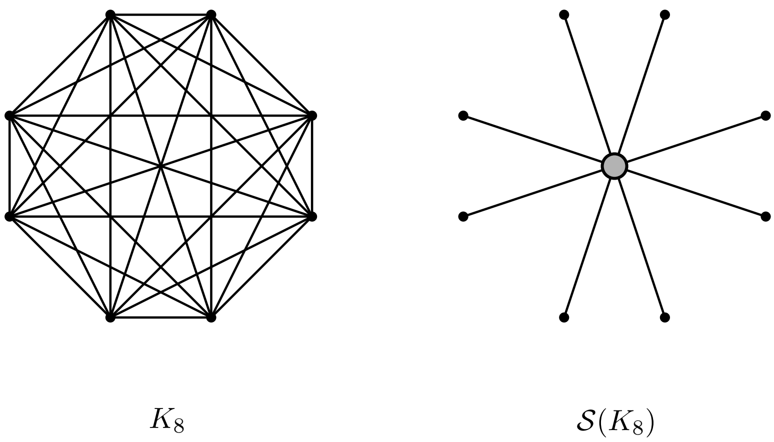

Figure 5.3: The star graph of the complete graph K_8 on 8 vertices with respect to the block $\mathcal{B} = \{K_8\}$.

Note that restricting the natural graph metric of a graph $\Gamma = (V, E)$ to the vertex set V induces a metric on V. In particular with respect to this induced metric on V the inclusion

$V \hookrightarrow \Gamma$ is a quasi-surjective isometric embedding, and hence a quasi-isometry. This is useful for us since instead of studying embeddings $\Gamma \hookrightarrow X_\mu$ we can study the simpler restrictions onto V. Also, the following lemma implies that a graph Γ and its star graph $\mathcal{S}(\Gamma, \mathcal{B})$ are always quasi-isometric.

Lemma 5.2. *Let $\mathcal{B} = \{B_i \mid i \in I\}$ be a block cover of Γ. Then the inclusion*

$$\Phi \ : \ V(\Gamma) \to V(\mathcal{S}) = V(\Gamma) \sqcup \mathcal{B}, \qquad v \mapsto v$$

is a quasi-isometry satisfying for all $v, v' \in V(\Gamma)$:

$$d_{\mathcal{S}}(\Phi(v), \Phi(v')) = 2 \cdot d(v, v'),$$

where d resp. $d_{\mathcal{S}}$ is the natural graph metric on Γ resp. $\mathcal{S}$.

Proof. The map Φ is quasi-surjective: Each vertex in $\mathcal{S}$ is either of the form $v \in V(\Gamma)$ or $B \in \mathcal{B}$. In the first case, we clearly have $d_{\mathcal{S}}(v, \Phi(v)) = 0$ and in the latter case choose any vertex $w \in V(\Gamma)$ lying inside B. Then $d_{\mathcal{S}}(B, \Phi(w)) = 1$, by definition of the star graph.

Let $v, v' \in V(\Gamma)$, $v \neq v'$, and $k := d(v, v') \in \mathbb{N}$. Choose a geodesic in Γ between v and v', which is a path consisting of a sequence $v = v_0, v_1, \ldots, v_k = v'$ of pairwise distinct vertices. Since $\mathcal{B}$ is a block cover each edge $\{v_{i-1}, v_i\}$ lies inside a block B_i for $i = 1, \ldots, k$. Hence in $\mathcal{S}$ we have that B_i is adjacent to both v_{i-1} and v_i. We can now form a path in $\mathcal{S}$ from $\Phi(v) = v$ to $\Phi(v') = v'$ given by the vertex sequence $v, B_1, v_1, B_2, \ldots, B_k, v_k = v'$ having length $2k$, see Figure 5.4. This proves the inequality

$$d_{\mathcal{S}}(\Phi(v), \Phi(v')) \leq 2 \cdot d(v, v').$$

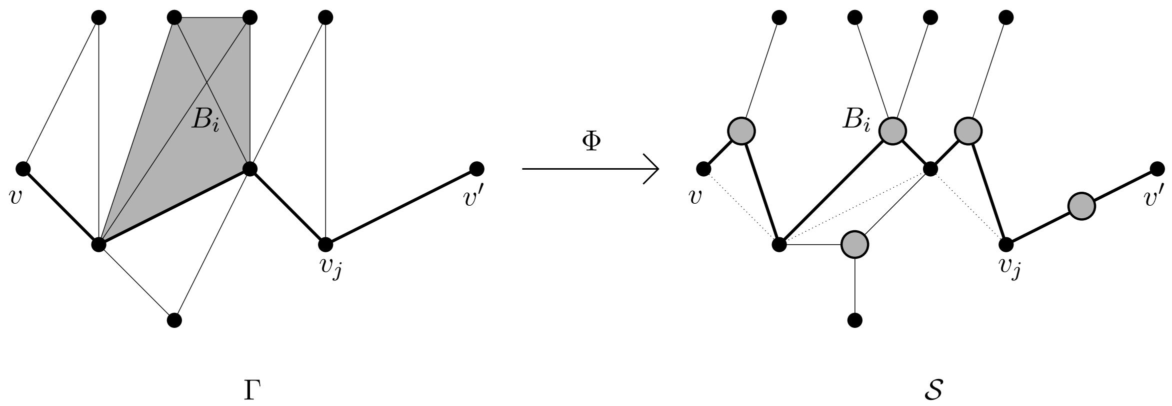

Figure 5.4: A geodesic path in Γ of length 4 and its image under Φ in the star graph $\mathcal{S}(\Gamma)$ having length 8. Here all maximal complete subgraphs of Γ yield the blocks of $\mathcal{B}$.

We claim that this path $v, B_1, v_1, B_2, \ldots, B_k, v'$ is a geodesic segment in $\mathcal{S}$ from v to

v'. This proves that the inequality is indeed an equality. Assume there was a shorter path from v to v' inside $\mathcal{S}$, say of length $2j < 2k$. This path then must be of the form $v = v_0', B_1', v_1', \ldots, B_j', v_j' = v'$. By definition of the edges in $\mathcal{S}$ we see that for all $i = 1, \ldots, j$ the vertices v_{i-1}' and v_i' lie in the block B_i' and thus are adjacent in Γ. However, this implies that $v, v_1', \ldots, v_j' = v'$ is a path in Γ from v to v' of length $j < k = d(v, v')$, a contradiction. $\square$

We are now able to prove that all regular translation surfaces X_{F_n} and bipartite trees $T_{n,\infty}$ are quasi-isometric to T_∞.

Theorem 5.3. *Let $n \geq 2$. Then:*

(i) The surface X_{F_n} is quasi-isometric to $T_{n,\infty}$,

(ii) The trees $T_{n,\infty}$ and T_∞ are quasi-isometric.

In particular, for all $n \geq 2$ the surfaces X_{F_n} are quasi-isometric to T_∞.

Proof. (i) The base surface has $2n$ singularities and the corresponding monodromy elements for each singularity are given by $a_1^\pm, a_2^\pm, \ldots, a_n^\pm$. By Lemma 3.4 and similarly to the AB-surface we can describe the singularities of X_{F_n} by

$$\mathrm{Sing}(X_{F_n}) = \mathrm{Sing}^\infty(X_{F_n}) \cong \bigsqcup_{i=1}^{n} \left(F_n/A_i \sqcup F_n/A_i\right),$$

where $A_i := \langle a_i \rangle \leq F_n$ is the infinite cyclic subgroup generated by a_i. Define $\sigma_{gA_i} := [C_i, g]$, where C_i is a corner of the base polygon, see Figure 5.1. In other words σ_{gA_i} is the unique ∞-angle singularity which lies in the preimage of the conical singularity corresponding to C_i and which is contained in all copies ga_i^m, $m \in \mathbb{Z}$, of the base polygon.

We define a graph Δ, analogously as for the surface $X_{(a,b)}$, as follows: Let the vertex set of Δ be $V := \bigsqcup_{i=1}^n F_n/A_i$ and put an edge between gA_i and hA_j if and only if $i \neq j$ and $gA_i \cap hA_j \neq \emptyset$. This makes Δ an F_n-left invariant n-partite graph with partitions $F_n/A_1, \ldots, F_n/A_n$.

Step 1:
We prove that Δ is quasi-isometric to $\mathrm{Cay}(F_n, \{a_1^{m_1}, \ldots, a_n^{m_n} \mid m_i \in \mathbb{Z}\})$. Theorem 4.9 implies then that Δ and X_{F_n} are quasi-isometric. Note that $\mathrm{Cay}(F_n, \{a_1^{m_1}, \ldots, a_n^{m_n} \mid m \in \mathbb{Z}\}$ is isometric to the group F_n equipped with the syllable metric d_T^{syl} with respect to the generating set $T = \{a_1, \ldots, a_n\}$. Consider the following map

$$\Phi \;:\; (F_n, d_T^{syl}) \to (V(\Delta), d_\Delta),\; g \mapsto gA_1.$$

The map is quasi-surjective: Given a vertex hA_i in Δ we have that hA_1 and hA_i are either equal or adjacent vertices since $h \in hA_1 \cap hA_i$. Hence $d_\Delta(\Phi(h), hA_i) \leq 1$.

Let us prove the first of two inequalities for a quasi-isometric embedding. Let $g \in F_n$ with $|g|_T^{syl} = k \in \mathbb{N}$. Then g is of the form $g = a_{i_1}^{m_1} a_{i_2}^{m_2} \ldots a_{i_k}^{m_k}$ with $m_1, \ldots, m_k \in \mathbb{Z}\backslash\{0\}$

and $a_{i_r} \neq a_{i_{r+1}}$. In Δ we construct a path from $A_1 = \Phi(1)$ to $gA_1 = \Phi(g)$ given by the following sequence of adjacent vertices

$$A_1 \sim A_{i_1} \sim a_{i_1}^{m_1} A_{i_2} \sim a_{i_1}^{m_1} a_{i_2}^{m_2} A_{i_3} \sim \cdots \sim \prod_{j=1}^{k-1} a_{i_j}^{m_j} A_{i_k} = gA_{i_k} \sim gA_1.$$

Note that the cases $i_1 = 1$ and $i_k = 1$ are possible. In these cases the corresponding adjacencies at the beginning and end of the path have to be replaced by equalities. Hence this path has length at most $k+1$ and we have

$$d_\Delta(\Phi(1), \Phi(g)) \leq k + 1 = d_T^{syl}(1, g) + 1,$$

and the F_n-left invariance of the metrics proves the corresponding inequality for $d_\Delta(\Phi(h), \Phi(g))$ for any $g, h \in F_n$.

For the last inequality let $g \in F_n$ such that $d_\Delta(\Phi(1), \Phi(g)) = d_\Delta(A_1, gA_1) = k \in \mathbb{N}$. This means there is a geodesic in Δ from $A_1 =: A_{i_0}$ to gA_1 of length k which is given by a sequence of adjacent vertices of the form

$$A_{i_0} \sim v_1 A_{i_1} \sim v_2 A_{i_2} \sim \cdots \sim v_k A_{i_k} = gA_1,$$

where $v_1, \ldots, v_k \in F_n$ are minimal representatives of the cosets. Note that since Δ is n-partite subsequent indices must be distinct, i.e. we have $1 = i_0 \neq i_1 \neq i_2 \neq \ldots \neq i_k = 1$. Consider now the first edge of the geodesic. Since A_{i_0} and $v_1 A_{i_1}$ are adjacent they have an element of F_n in common. It follows that the minimal representant v_1 must be of the form $v_1 = a_{i_0}^{m_0}$ for some $m_0 \in \mathbb{Z}$. Similarly, since $v_1 A_{i_1} = a_{i_0}^{m_0} A_{i_1}$ and $v_2 A_{i_2}$ are adjacent it follows that v_2 is of the form $v_2 = a_{i_0}^{m_0} a_{i_1}^{m_1}$ for some $m_1 \in \mathbb{Z}$. Inductively, it follows that v_k is of the form $v_k = a_{i_0}^{m_0} a_{i_1}^{m_1} \ldots a_{i_{k-1}}^{m_{k-1}}$ for integer exponents. Since $v_k A_{i_k} = gA_1$ it follows that g is of the form $g = a_{i_0}^{m_0} a_{i_1}^{m_1} \ldots a_{i_{k-1}}^{m_{k-1}} a_1^{m_k}$ and thus $|g|_T^{syl} \leq k+1$. Therefore, we have

$$d_T^{syl}(1, g) - 1 \leq d_\Delta(\Phi(1), \Phi(g)),$$

which proves step 1.

Step 2:
In order to finish the proof of (i) we need to show that Δ is quasi-isometric to $T_{n,\infty}$. For each $g \in F_n$ define the induced subgraph B_g of Δ on n vertices $gA_1, \ldots, gA_n$. Since $g \in \bigcap_{i=1}^n gA_i \neq \emptyset$, this subgraph B_g is a complete subgraph of Δ. We claim that the collection $\mathcal{B} := \{B_g \mid g \in F_n\}$ is a block cover on Δ: Consider any edge with vertices gA_i and hA_j, $i \neq j$. Then by definition the intersection $gA_i \cap hA_j$ contains an element, say $v \in F_n$. It follows that $vA_i = gA_i \in B_v$ and $vA_j = hA_j \in B_v$. Since B_v is a complete subgraph the whole edge $\{gA_i, hA_j\}$ is then contained in B_v which proves the claim. We compute the star graph $\mathcal{S} := \mathcal{S}(\Delta, \mathcal{B})$. Its vertex set is given by $V(\Delta) \sqcup \mathcal{B}$ which we identify with $\bigsqcup_{i=1}^n F_n/A_i \sqcup F_n$. By definition of the star graph there are no edges of the

form $\{gA_i, hA_j\}$ and $\{v, w\}$ where $g, h, v, w \in F_n$. The edges are precisely of the form $\{gA_i, v\}$ where $v \in gA_i$, see Figure 5.5 for an example when $n = 4$. In other words, $\mathcal{S}$ consists of stars with center $g \in F_n$ and neighbors gA_i, $i = 1, \ldots, n$.

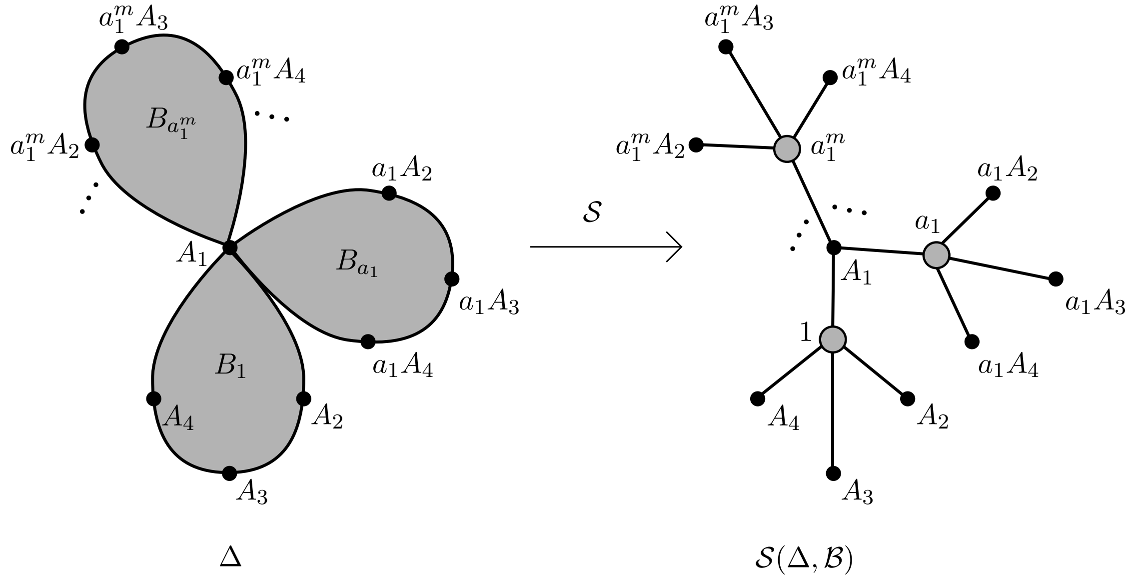

Figure 5.5: The neighborhood of A_1 in Δ consists of infinitely many blocks $B_{a_1^m}$, $m \in \mathbb{Z}$. In the star graph $\mathcal{S}$ the vertex A_1 has valence ∞ and each block vertex has valence 4. The star graph $\mathcal{S}$ is isomorphic to $T_{4,\infty}$.

We claim that $\mathcal{S}$ is isomorphic to $T_{n,\infty}$, where F_n forms the vertex set of valence n and $\bigsqcup_{i=1}^n F_n/A_i$ the vertex set of infinite valence. Clearly, each vertex $v \in F_n$ is only adjacent to $vA_1, \ldots, vA_n$ and thus has valence n. On the other hand each vertex gA_i is only adjacent to all the vertices ga_i^m, $m \in \mathbb{Z}\backslash\{0\}$ and thus has infinite valence. Since there are only edges in $\mathcal{S}$ between vertices having distinct valence we have that $\mathcal{S}$ is bipartite. We need to show that $\mathcal{S}$ is a tree, this then concludes the proof. The proof that $\mathcal{S}$ is acyclic is very similar to step 2 in the proof of Lemma 3.8. Assume that $\mathcal{S}$ contains a closed path. W.l.o.g. this path contains the vertex $1 \in F_n$. Starting from 1 the vertex sequence of the path then is of the form

$$1 \sim A_{i_1} \sim a_{i_1}^{m_1} \sim a_{i_1}^{m_1} A_{i_2} \sim a_{i_1}^{m_1} a_{i_2}^{m_2} \sim a_{i_1}^{m_1} a_{i_2}^{m_2} A_{i_3} \sim \ldots$$

for integers $m_1, m_2, m_3, \cdots \in \mathbb{Z}\backslash\{0\}$, see Figure 5.5. Every second vertex is an element in F_n whose syllable length with respect to the generating set $\{a_1, \ldots, a_n\}$ is increasing by 1. This is a contradiction, since in a closed path this sequence of syllable lengths must again decrease in order to return to 1 having syllable length 0.

(ii) We prove the statement using a contraction argument. For this denote the two vertex partitions of $T_{n,\infty}$ by V_n and V_∞ corresponding to the valency of the vertices. Fix a vertex $v_0 \in V_n$, called the root. The root v_0 now induces an orientation of the tree $T_{n,\infty}$ as follows: Given two adjacent vertices $v, w \in V(T_{n,\infty})$ we orient the edge from v to w

if and only if $d(v, v_0) < d(w, v_0)$. See Figure 5.6 for an illustration of the oriented tree $T_{4,\infty}$. This way, the root v_0 has n outgoing edges and each vertex $v \in V_n \backslash \{v_0\}$ has $n-1$ outgoing edges and one incoming edge. For $v \in V_n$ we call $w \in V_\infty$ a *lower neighbor* if v and w are adjacent and there is an edge from v to w. For any $v \in V_n$, $v \neq v_0$, we define S_v to be the subtree of $T_{n,\infty}$ induced by the vertices v and its $n-1$ lower neighbors, i.e. the smallest subtree containing these vertices. For v_0 choose some $n-1$ lower neighbors and let S_{v_0} be the corresponding induced subtree. Clearly for all $v \in V_n$ the subgraphs S_v are disjoint stars with center v and $n-1$ leaves.

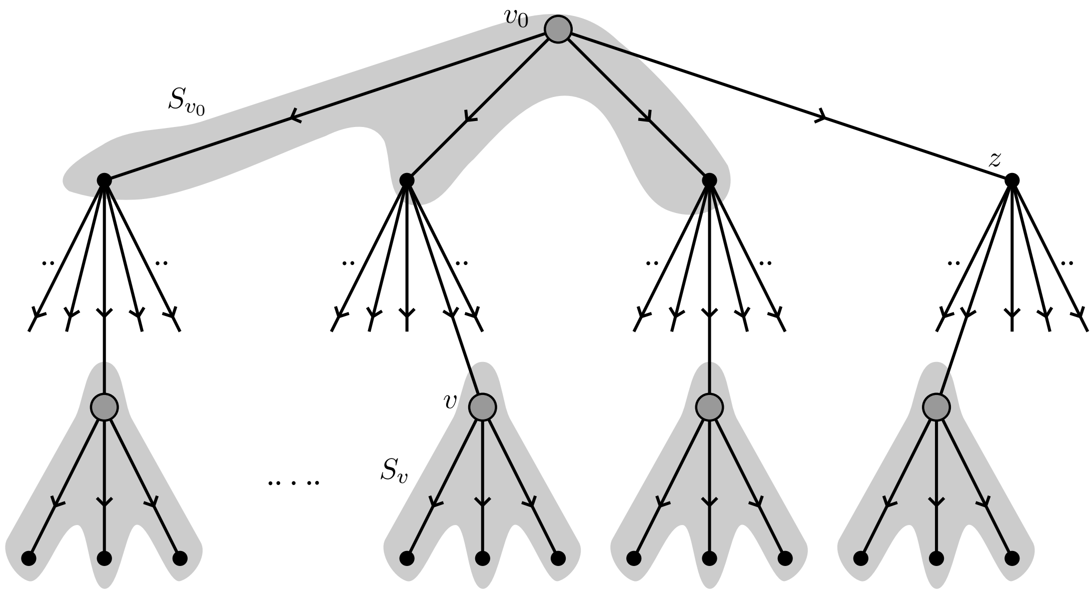

Figure 5.6: The oriented $(4, \infty)$-regular tree $T_{4,\infty}$ with root v_0. To each vertex of finite valence we construct the star subgraph S_v.

Define T to be the graph obtained from $T_{n,\infty}$ after contracting each star S_v, $v \in V_n$, to one distinct vertex and let $f : T_{n,\infty} \to T$ be the corresponding quotient map which clearly is surjective. In a first step we show that T is graph-isometric to T_∞ and then we prove that f is a quasi-isometry.

Step 1:
Since we contract each subtree S_v of the tree $T_{n,\infty}$ to one distinct vertex the resulting graph is a tree again. Hence it remains to prove that each vertex in T has countably infinite valence. Consider a vertex $z \in V(T)$. If z is not a contracted star $f(S_v)$ for some $v \in V_n$ then z has to be the unique neighbor of the root v_0 which is not contained in the subtree S_{v_0}. If $z = f(S_v)$ for some $v \in V_n$ then its neighbors in T are all vertices $f(S_w)$ where $w \in V_n$ with $d(v, w) = 2$, see figure 5.7. This shows that the set of neighbors of z is countably infinite as it is a finite union of countably infinitely many vertices. This concludes step 1.

Step 2:

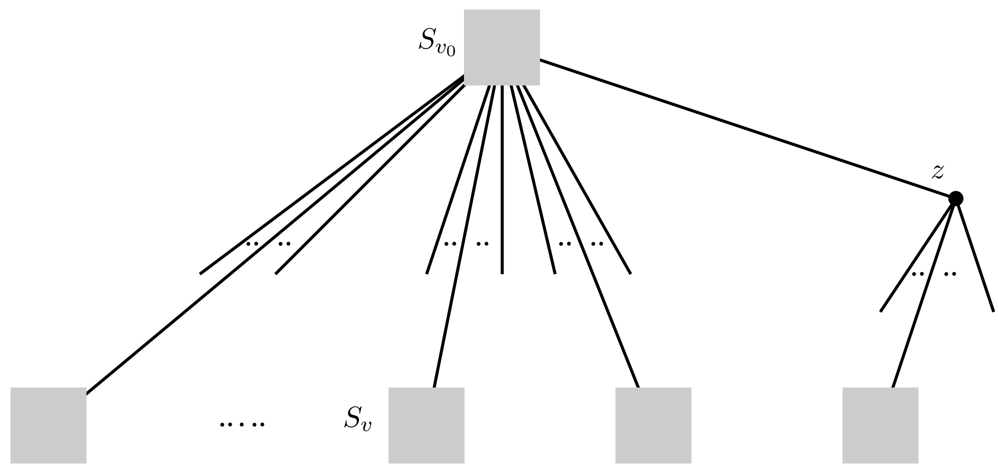

Figure 5.7: The image of $T_{4,\infty}$ under the contraction f. Each contracted star is a vertex having countably infinitely many neighbors.

Since f is a contraction of subgraphs we directly have

$$d_T(f(v), f(w)) \leq d(v, w),$$

for all vertices $v, w \in V(T_{n,\infty})$ where d resp. d_T denotes the natural graph metric on $T_{n,\infty}$ resp. T. Now let $v, w \in V(T_{n,\infty})$ and $k := d(v, w) \in \mathbb{N}_0$. Note that since the stars S_x, $x \in V_n$, are disjoint, any geodesic segment in $T_{n,\infty}$ has only connected subpaths of length at most 2 which are contracted under f. See Figure 5.8 for an example. In other words, at least every third edge of a geodesic segment must be an edge not contracted under f, which implies for the geodesic segment between v and w of length k

$$d_T(f(v), f(w)) \geq \frac{1}{3}(k-2),$$

and hence

$$d(v, w) \leq 3 \cdot d_T(f(v), f(w)) + 2.$$

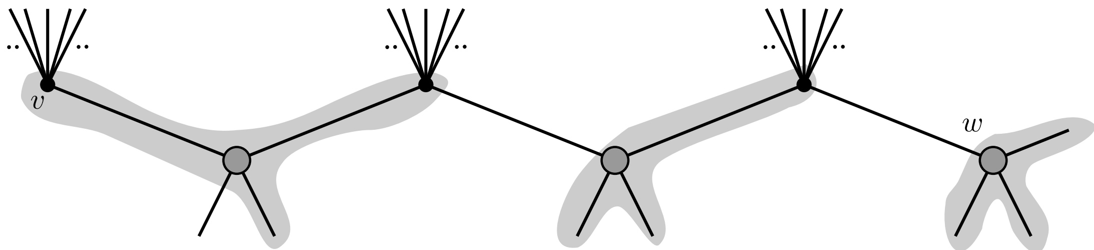

Figure 5.8: The vertices v and w have distance 5 in $T_{4,\infty}$. Since all star subgraphs are disjoint at least one third of all edges in the geodesic segment from v to w are not contracted under f. Here, the images $f(v)$ and $f(w)$ have distance 2.

□

Remark. The fact that $\overline{X}_{F_n}$ is quasi-isometric to $T_{n,\infty}$ can also be visualised using an embedding of the graph $T_{n,\infty}$ into the translation surface $\overline{X}_{F_n}$ as follows. In each copy $P \times \{g\}$ of the base polygon fix a point $[x, g]$ which we identify with $g \in F_n$. See Figure 5.9 for an illustration when $n = 4$. From this point consider the n straight line segments from $[x, g]$ to the singularities of $\overline{X}_{F_n}$ given by $gA_1, \ldots, gA_n$ as shown. These form the edges of the embedded graph. Note that this graph is bipartite and that each vertex g has valency n while each vertex that is a singularity has infinite valence. Hence this embedded graph is isomorphic to $T_{n,\infty}$. Theorem 5.3 tells us that this embedding indeed is a quasi-isometry.

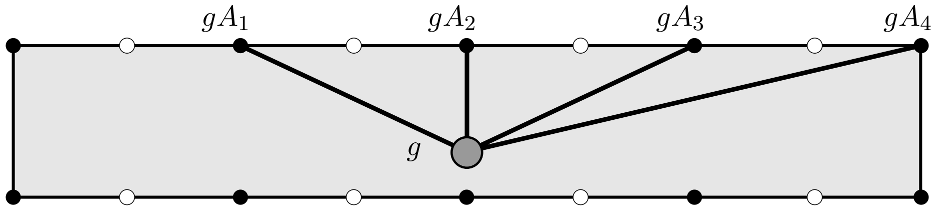

Figure 5.9: An embedding of the star graph $\mathcal{S}(\Delta)$ into $\overline{X}_{F_4}$. Each copy of the base polygon represents a vertex $g \in F_4$ which is adjacent to the four cosets $gA_1, \ldots, gA_4$ representing singularities in $\overline{X}_{F_4}$.

5.2 The Teichmüller space of translation structures

Another observation we can conclude from Theorem 4.9 deals with the Teichmüller space of translation structures. It is defined analoguously to the well studied Teichmüller space of hyperbolic structures on a topological surface, as explained in [IT92] and [FM14].

For this, we fix a closed surface S_0 together with a finite subset $\Sigma \subset S_0$. Furthermore also fix a regular topological covering $p : S \to S_0^*$ where $S_0^* := S_0 \backslash \Sigma$. We define an equivalence relation on the set of translation structures on S_0^* as follows. Two translation structures $\mathcal{T}, \mathcal{T}'$ are *equivalent* if there is a translation $f : (S_0, \Sigma, \mathcal{T}) \to (S_0, \Sigma, \mathcal{T}')$ between the corresponding finite translation surfaces which is isotopic to the identity on S_0. Then the *Teichmüller space of translation structures on* S_0, denoted by $\Omega T(S_0)$, is the space of equivalence classes of translation structures on S_0. Given such a translation structure $\mathcal{T}$ in a class of $\Omega T(S_0)$ this defines a finite translation surface $X_{0,\mathcal{T}} := (S_0, \Sigma, \mathcal{T})$ as in the geometric definition given in Section 2.1. We can now lift the translation structure $\mathcal{T}$ via p to a translation structure $\tilde{\mathcal{T}}$ on S. In this way the space $X_{\mathcal{T}}^* := (S, \mathcal{T})$ is equipped with a flat metric and hence is a possibly punctured translation surface. In particular the covering $p : X_{\mathcal{T}}^* \to X_{0,\mathcal{T}}$ is a translation covering. From this follows that the metric completion $\overline{X_{\mathcal{T}}^*}$ consists of $X_{\mathcal{T}}^*$ with additional conical and ∞-angle singularities. Let $X_{\mathcal{T}}$ be the translation surface obtained by removing all ∞-angle singularities from $\overline{X_{\mathcal{T}}^*}$. By definition this is a regular translation surface with base surface $X_{0,\mathcal{T}}$.

Corollary 5.4. *For all $\mathcal{T}, \mathcal{T}' \in \mathcal{T}(S_0)$ we have that the regular translation surfaces $X_{\mathcal{T}}$ and $X_{\mathcal{T}'}$ are quasi-isometric.*

Proof. Let $\mathcal{T} \in \Omega T(S_0)$. By Theorem 4.9 the resulting regular translation surface $X_{\mathcal{T}}$ is quasi-isometric to the Cayley graph $\Gamma := \mathrm{Cay}(G, S_\mu^\infty)$ where G is the deck transformation group of the regular covering p and μ is the corresponding monodromy map. Note that Γ only depends on the covering p and the choice of singular loops on S_0 which is purely topological data. In particular, the graph Γ does not depend on the choice $\mathcal{T}$ of translation structure on S_0. Hence we have for all $\mathcal{T}, \mathcal{T}' \in \Omega T(S_0)$ that

$$X_{\mathcal{T}} \sim_{QI} \Gamma \sim_{QI} X_{\mathcal{T}'}.$$

□

Remark. In other words Corollary 5.4 tells us that the quasi-isometry class of a regular translation surface does not depend on the choice of translation structure on the base surface. For example in Figure 5.10 we see two different translation structures on a genus 2 surface with one puncture. The resulting regular translation surfaces however are quasi-isometric.

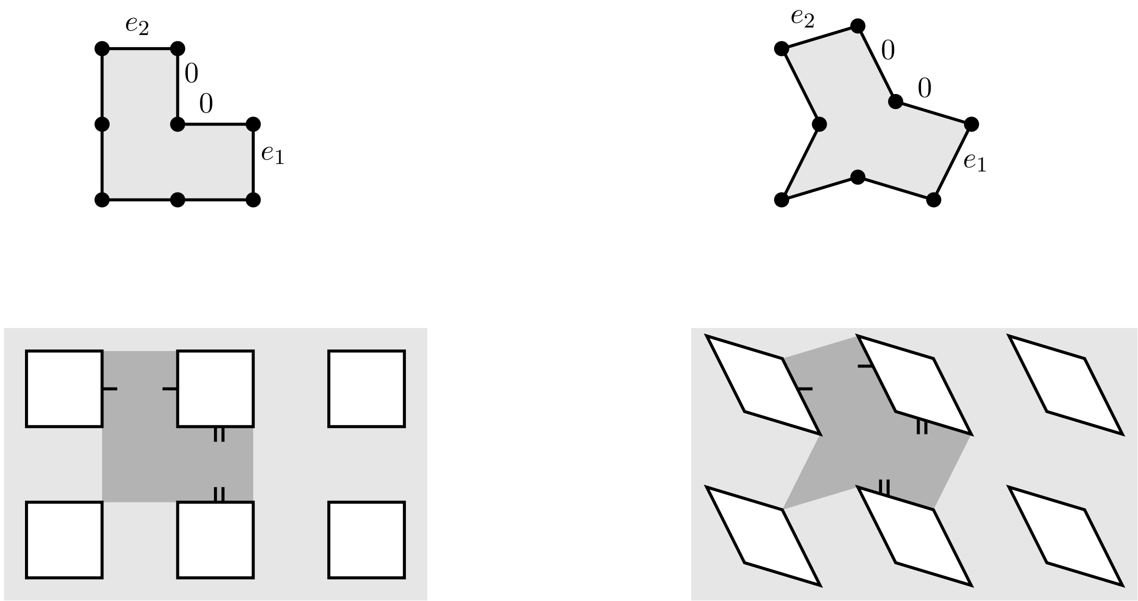

Figure 5.10: A $\mathbb{Z}^2$ covering of a genus 2 surface equipped with two different translation structures. The corresponding regular translation surfaces are quasi-isometric. Here, e_1 and e_2 are the standard generators of $\mathbb{Z}^2$.

Bibliography

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature.* Vol. 319. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 1999, pp. xxii+643. ISBN: 3-540-64324-9. DOI: `10.1007/978-3-662-12494-9`. URL: `https://doi.org/10.1007/978-3-662-12494-9`.

[Bow12] Joshua P. Bowman. "Finiteness conditions on translation surfaces". In: *Quasiconformal mappings, Riemann surfaces, and Teichmüller spaces.* Vol. 575. Contemp. Math. Amer. Math. Soc., Providence, RI, 2012, pp. 31–40. DOI: `10.1090/conm/575/11388`. URL: `https://doi.org/10.1090/conm/575/11388`.

[Bri08] Martin R. Bridson. "Geometric and combinatorial group theory". In: *ThePrinceton Companion to Mathematics.* Princeton University Press, 2008, pp. 431–448.

[BV13] Joshua P. Bowman and Ferrán Valdez. "Wild singularities of flat surfaces". In: *Israel J. Math.* 197.1 (2013), pp. 69–97. ISSN: 0021-2172. DOI: `10.1007/s11856-013-0022-y`. URL: `https://doi.org/10.1007/s11856-013-0022-y`.

[Cha04] R. Chamanara. "Affine automorphism groups of surfaces of infinite type". In: *In the tradition of Ahlfors and Bers, III.* Vol. 355. Contemp. Math. Amer. Math. Soc., Providence, RI, 2004, pp. 123–145. DOI: `10.1090/conm/355/06449`. URL: `https://doi.org/10.1090/conm/355/06449`.

[CRW19] Lucien Clavier, Anja Randecker, and Chenxi Wu. "Rotational component spaces for infinite-type translation surfaces". In: *Geom. Dedicata* 201 (2019), pp. 57–80. ISSN: 0046-5755. DOI: `10.1007/s10711-018-0381-y`. URL: `https://doi.org/10.1007/s10711-018-0381-y`.

[Dan10] Klaus Dankwart. "On the large-scale geometry of flat surfaces". PhD thesis. Rheinische Friedrich-Whilhelms-Universität Bonn, 2010.

[Die18] Reinhard Diestel. *Graph theory.* Fifth. Vol. 173. Graduate Texts in Mathematics. Paperback edition of [MR3644391]. Springer, Berlin, 2018, pp. xviii+428.

[DV] Vincent Delecroix and Ferrán Valdez. *Infinite Translation Surfaces in the Wild*. URL: https://www.labri.fr/perso/vdelecro/infinite-translation-surfaces-in-the-wild.html.

[FK36] Ralph H. Fox and Richard B. Kershner. "Concerning the transitive properties of geodesics on a rational polyhedron". In: *Duke Math. J.* 2.1 (1936), pp. 147–150. ISSN: 0012-7094. DOI: 10.1215/S0012-7094-36-00213-2. URL: https://doi.org/10.1215/S0012-7094-36-00213-2.

[FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Vol. 49. Princeton Mathematical Series. Princeton University Press, Princeton, NJ, 2012, pp. xiv+472. ISBN: 978-0-691-14794-9.

[FM14] Giovanni Forni and Carlos Matheus. "Introduction to Teichmüller theory and its applications to dynamics of interval exchange transformations, flows on surfaces and billiards". In: *J. Mod. Dyn.* 8.3-4 (2014), pp. 271–436. ISSN: 1930-5311. DOI: 10.3934/jmd.2014.8.271. URL: https://doi.org/10.3934/jmd.2014.8.271.

[Fre31] Hans Freudenthal. "Über die Enden topologischer Räume und Gruppen". In: *Math. Z.* 33.1 (1931), pp. 692–713. ISSN: 0025-5874. DOI: 10.1007/BF01174375. URL: https://doi.org/10.1007/BF01174375.

[Ful95] William Fulton. *Algebraic topology*. Vol. 153. Graduate Texts in Mathematics. A first course. Springer-Verlag, New York, 1995, pp. xviii+430. DOI: 10.1007/978-1-4612-4180-5. URL: https://doi.org/10.1007/978-1-4612-4180-5.

[Gri80] R. I. Grigorčuk. "On Burnside's problem on periodic groups". In: *Funktsional. Anal. i Prilozhen.* 14.1 (1980), pp. 53–54. ISSN: 0374-1990.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002, pp. xii+544.

[Her12] Frank Herrlich. "Introduction to origamis in Teichmüller space". In: *Strasbourg master class on geometry*. Vol. 18. IRMA Lect. Math. Theor. Phys. Eur. Math. Soc., Zürich, 2012, pp. 233–253. DOI: 10.4171/105-1/3. URL: https://doi.org/10.4171/105-1/3.

[HHW13] W. Patrick Hooper, Pascal Hubert, and Barak Weiss. "Dynamics on the infinite staircase". In: *Discrete Contin. Dyn. Syst.* 33.9 (2013), pp. 4341–4347. ISSN: 1078-0947. DOI: 10.3934/dcds.2013.33.4341. URL: https://doi.org/10.3934/dcds.2013.33.4341.

[Hop44] Heinz Hopf. "Enden offener Räume und unendliche diskontinuierliche Gruppen". In: *Comment. Math. Helv.* 16 (1944), pp. 81–100. ISSN: 0010-2571. DOI: 10.1007/BF02568567. URL: https://doi.org/10.1007/BF02568567.

[HS06] Pascal Hubert and Thomas A. Schmidt. "An introduction to Veech surfaces". In: *Handbook of dynamical systems. Vol. 1B*. Elsevier B. V., Amsterdam, 2006, pp. 501–526. DOI: 10.1016/S1874-575X(06)80031-7. URL: https://doi.org/10.1016/S1874-575X(06)80031-7.

[HS10] Pascal Hubert and Gabriela Schmithüsen. "Infinite translation surfaces with infinitely generated Veech groups". In: *J. Mod. Dyn.* 4.4 (2010), pp. 715–732. ISSN: 1930-5311. DOI: 10.3934/jmd.2010.4.715. URL: https://doi.org/10.3934/jmd.2010.4.715.

[htt] Moishe Kohan (https://math.stackexchange.com/users/84907/moishe-kohan). *Is the infinite regular tree T_∞ quasi-isometric to the n-regular tree T_n?* Mathematics Stack Exchange. URL:https://math.stackexchange.com/q/3066187 (version: 2019-01-08). eprint: https://math.stackexchange.com/q/3066187. URL: https://math.stackexchange.com/q/3066187.

[IT92] Y. Imayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Translated and revised from the Japanese by the authors. Springer-Verlag, Tokyo, 1992, pp. xiv+279. ISBN: 4-431-70088-9. DOI: 10.1007/978-4-431-68174-8. URL: https://doi.org/10.1007/978-4-431-68174-8.

[KMS86] Steven Kerckhoff, Howard Masur, and John Smillie. "Ergodicity of billiard flows and quadratic differentials". In: *Ann. of Math. (2)* 124.2 (1986), pp. 293–311. ISSN: 0003-486X. DOI: 10.2307/1971280. URL: https://doi.org/10.2307/1971280.

[KS00] Richard Kenyon and John Smillie. "Billiards on rational-angled triangles". In: *Comment. Math. Helv.* 75.1 (2000), pp. 65–108. ISSN: 0010-2571. DOI: 10.1007/s000140050113. URL: https://doi.org/10.1007/s000140050113.

[Möl09] Martin Möller. "Affine groups of flat surfaces". In: *Handbook of Teichmüller theory. Vol. II*. Vol. 13. IRMA Lect. Math. Theor. Phys. Eur. Math. Soc., Zürich, 2009, pp. 369–387. DOI: 10.4171/055-1/11. URL: https://doi.org/10.4171/055-1/11.

[Möl13] Martin Möller. "Teichmüller curves, mainly from the viewpoint of algebraic geometry". In: *Moduli spaces of Riemann surfaces*. Vol. 20. IAS/Park City Math. Ser. Amer. Math. Soc., Providence, RI, 2013, pp. 267–318.

[MT02] Howard Masur and Serge Tabachnikov. "Rational billiards and flat structures". In: *Handbook of dynamical systems, Vol. 1A*. North-Holland, Amsterdam, 2002, pp. 1015–1089. DOI: 10.1016/S1874-575X(02)80015-7. URL: https://doi.org/10.1016/S1874-575X(02)80015-7.

[Ran16] Anja Randecker. "Geometry and topology of wild translation surfaces". PhD thesis. Karlsruher Institut für Technologie (KIT), 2016. 151 pp. ISBN: 978-3-7315-0456-6. DOI: 10.5445/KSP/1000050964.

[Ran18] Anja Randecker. "Wild translation surfaces and infinite genus". In: *Algebr. Geom. Topol.* 18.5 (2018), pp. 2661–2699. ISSN: 1472-2747. DOI: 10.2140/agt.2018.18.2661. URL: https://doi.org/10.2140/agt.2018.18.2661.

[Sch04] Gabriela Schmithüsen. "An algorithm for finding the Veech group of an origami". In: *Experiment. Math.* 13.4 (2004), pp. 459–472. ISSN: 1058-6458. URL: http://projecteuclid.org/euclid.em/1109106438.

[Sta68] John R. Stallings. "On torsion-free groups with infinitely many ends". In: *Ann. of Math. (2)* 88 (1968), pp. 312–334. ISSN: 0003-486X. DOI: 10.2307/1970577. URL: https://doi.org/10.2307/1970577.

[Sti93] John Stillwell. *Classical topology and combinatorial group theory*. Second. Vol. 72. Graduate Texts in Mathematics. Springer-Verlag, New York, 1993, pp. xii+334. ISBN: 0-387-97970-0. DOI: 10.1007/978-1-4612-4372-4. URL: https://doi.org/10.1007/978-1-4612-4372-4.

[Str84] Kurt Strebel. *Quadratic differentials*. Vol. 5. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1984, pp. xii+184. ISBN: 3-540-13035-7. DOI: 10.1007/978-3-662-02414-0. URL: https://doi.org/10.1007/978-3-662-02414-0.

[Vee86] William A. Veech. "The Teichmüller geodesic flow". In: *Ann. of Math. (2)* 124.3 (1986), pp. 441–530. ISSN: 0003-486X. DOI: 10.2307/2007091. URL: https://doi.org/10.2307/2007091.

[Vee89] W. A. Veech. "Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards". In: *Invent. Math.* 97.3 (1989), pp. 553–583. ISSN: 0020-9910. DOI: 10.1007/BF01388890. URL: https://doi.org/10.1007/BF01388890.

[Wri15] Alex Wright. "Cylinder deformations in orbit closures of translation surfaces". In: *Geom. Topol.* 19.1 (2015), pp. 413–438. ISSN: 1465-3060. DOI: 10.2140/gt.2015.19.413. URL: https://doi.org/10.2140/gt.2015.19.413.

[Yoc10] Jean-Christophe Yoccoz. "Interval exchange maps and translation surfaces". In: *Homogeneous flows, moduli spaces and arithmetic*. Vol. 10. Clay Math. Proc. Amer. Math. Soc., Providence, RI, 2010, pp. 1–69.

[Zor06] Anton Zorich. "Flat surfaces". In: *Frontiers in number theory, physics, and geometry. I.* Springer, Berlin, 2006, pp. 437–583. DOI: 10.1007/978-3-540-31347-2_13. URL: https://doi.org/10.1007/978-3-540-31347-2_13.