

Chapter 2

Basic topological and geometrical aspects

In the first part of this chapter we explain the invariants that classify all orientable surfaces up to homeomorphism. In the second part we discuss infinite-type translation surfaces that arise as covering spaces of translation surfaces. In the last two sections we study singularities.

2.1 Topological classification of surfaces

The topological classification theorem for orientable surfaces with finitely-generated fundamental group and empty boundary states that any such surface is determined up to homeomorphism by a pair of non-negative integers (g, n) corresponding to the genus and the number of punctures, respectively [Ful95]. The topological classification of orientable surfaces with infinitely-generated fundamental group and empty boundary was done by I. Richards [Ric63] and is based on the work of Raymond and Freudenthal's theory of *ends* of topological spaces. Roughly speaking, to a sufficiently nice non-compact topological space X one can associate a topological invariant called the space of ends which encodes the different forms on which a point in X can escape to infinity. In the particular case when X is a surface, there are two kinds of ends: those that are accumulated by genus and those which have genus zero (a.k.a. *planar ends*). Ends together with genus define the topological invariants needed to classify orientable surfaces with infinitely-generated fundamental group and empty boundary.

2.1.1 Ends of topological spaces

For the purposes of this text we use two definitions of the space of ends of a topological space, which in our context are equivalent. One uses nested sequences of open sets and the other proper rays. The space of ends was originally introduced by H. Freudenthal [Fre31]. For a modern approach (in English) and a more detailed discussion about these two definitions we refer to [DK03].

Definition 2.1.1. [Fre31] Let X be a locally-compact, locally-connected, connected-Hausdorff space and $U_1 \supseteq U_2 \supseteq \dots$ be an infinite sequence of non-empty connected open subsets of X such that for each $i \in \mathbb{N}$ the boundary ∂U_i of U_i is compact and $\bigcap_{i \in \mathbb{N}} \overline{U_i} = \emptyset$. Two such sequences $U_1 \supseteq U_2 \supseteq \dots$ and $U'_1 \supseteq U'_2 \supseteq \dots$ are said to be equivalent if for every $i \in \mathbb{N}$ there exist j such that $U_i \supseteq U'_j$ and viceversa, that is, for every $i \in \mathbb{N}$ there exist j such that $U'_i \supseteq U_j$. The corresponding equivalence classes are called *ends* of X and we denote by $\text{Ends}(X)$ the set of all ends of X .

For every non-empty open subset U of X with compact boundary let

$$U^* := \{[U_1 \supseteq U_2 \supseteq \dots] \in \text{Ends}(X) \mid U_j \subset U \text{ for some } j\}. \quad (2.1)$$

We endow $\text{Ends}(X)$ with the topology generated by all sets U^* . The collection formed by all open sets of X and all sets of the form $U \cup U^*$ is a base for a topology of $X' := X \cup \text{Ends}(X)$.

Theorem 2.1.2: [Ray60]

Let $X' = X \cup \text{Ends}(X)$ be the topological space defined above. Then,

1. X' is Hausdorff, connected and locally-connected.
2. $\text{Ends}(X)$ is closed and has no interior points in X' .
3. $\text{Ends}(X)$ is totally-disconnected.
4. X' is compact.
5. If V is any open connected set in X' , then $V \setminus \text{Ends}(X)$ is connected.

We stress that in his work Raymond defines the space of ends as an inverse limit obtained by considering complements of nested sequences of relatively-compact subsets of X . His approach is equivalent to the one presented here. In summary, the space X' is a compactification of X obtained by adding a point for each way one can escape to infinity in X . Even though Theorem 2.1.2 works for a large class of spaces, we will be using it in two simple contexts: surfaces and infinite (but locally-finite) graphs.

Definition 2.1.3. Let X be a path-connected, locally-compact space. A proper continuous map $r : [0, \infty) \rightarrow X$ is called *proper ray* in X . Two proper rays r_1, r_2 are equivalent if for every compact $K \subset X$ there exists an $n \in \mathbb{N}$ such that $r_1([n, \infty))$ and $r_2([n, \infty))$ are contained in the same path-connected component of $X \setminus K$. The equivalence class of a proper ray r is denoted by $\text{end}(r)$ and:

$$\{\text{end}(r) : r \text{ is a proper ray in } X\}, \quad (2.2)$$

is also called the space of ends of X .

A sequence of ends $\text{end}(r_n)$ converges to $\text{end}(r)$ if for every compact $K \subset X$ there exists a sequence of integers M_n such that $r_n[M_n, \infty)$ and $r[M_n, \infty)$ lie in the same path component of $X \setminus K$ whenever n is sufficiently large.

Exercise 2.1.4

Prove that definitions 2.1.1 and 2.1.3 lead to homeomorphic spaces of ends when X is a locally-compact path-connected Hausdorff space.

Exercise 2.1.5

Let $\beta_n = \{e_i\}_{i=1}^n$ be the standard set of \mathbb{Z}^n . For example, $\beta_1 = \{1\}$ and $\beta_2 = \{(1,0), (0,1)\}$. Prove that the Cayley graph of \mathbb{Z}^n w.r.t. β_n has one end if $n \geq 2$ and two ends if $n = 1$. Denote by F_n the free group generated by $n \geq 2$ elements $\{a_1, \dots, a_n\}$ and Γ_n the corresponding Cayley graph. Prove that all spaces $\text{Ends}(F_n)$, $n \geq 2$ are homeomorphic to the Cantor set.

2.1.2 Topological classification of orientable surfaces

By *surface* we mean a connected topological manifold S of real dimension two. We stress that in this text all surfaces are required to be second-countable topological spaces, *i.e.* to have a countable basis for their topology. In particular surfaces are metrizable. Henceforth all the surfaces considered in this text are orientable and, unless explicitly stated, will have empty boundary. A simple closed *curve* in S is a continuous injective map $\alpha : \mathbb{S}^1 \rightarrow S$, where $\mathbb{S}^1 \subset \mathbb{C}$ is the unit circle. A curve is said to be *essential* if it is not isotopic to the boundary curve of a neighbourhood of a puncture of S nor to a point. We often abuse language and use the term curve to refer to the map α , its image and its isotopy class in S . A simple curve α is *separating* if $S \setminus \alpha$ has two connected components.

Definition 2.1.6. A set of disjoint isotopy classes of curves $\{a_i\}_{i \in I} \subset S$ is called multicurve. The genus of S is the maximal cardinality of a multicurve $\{a_i\}_{i \in I}$ for which there exist representatives $\{\alpha_i\}_{i \in I}$ such that $S \setminus \cup_{i \in I} \alpha_i$ is connected. If such cardinality is infinite we say that S has infinite genus. An end $[U_1 \supseteq U_2 \supseteq \dots] \in \text{Ends}(S)$ is called planar if there exists an $i \in \mathbb{N}$ such that U_i has genus zero. We define $\text{Ends}_\infty(S) \subset \text{Ends}(S)$ as the subspace of all ends which are not planar, and we refer to them as *ends accumulated by genus* or *ends of infinite genus*.

Remark that a surface S has genus zero if every simple essential curve in S is a separating curve. We have chosen the notation $\text{Ends}_\infty(S)$, for if $[U_1 \supseteq U_2 \supseteq \dots] \in \text{Ends}_\infty(S)$ then every U_i is an infinite genus surface. In other words, non-planar ends always have infinite genus. It follows from the definitions that $\text{Ends}_\infty(S)$ forms a closed subset of $\text{Ends}(S)$.

Theorem 2.1.7: Topological classification of orientable surfaces.

Two orientable surfaces S and S' with empty boundary are homeomorphic if and only if they have the same genus $g \in \{\mathbb{Z}^{\geq 0} \cup \infty\}$, and both $\text{Ends}_\infty(S) \subset \text{Ends}(S)$ and $\text{Ends}_\infty(S') \subset \text{Ends}(S')$ are homeomorphic as nested topological spaces, that is, there exists a homeomorphism $h : \text{Ends}(S) \rightarrow \text{Ends}(S)'$ whose restriction $h| : \text{Ends}_\infty(S) \rightarrow \text{Ends}_\infty(S')$ is a homeomorphism as well.

This result was first announced by Kerékjártó, but a complete proof was given by Richards [Ric63]. It tells us that the complete topological invariant of an arbitrary orientable topological surfaces is the genus plus a couple of nested topological spaces characterising the ways a point in S has to escape to infinity and, among these ways, which of them carry infinite genus. Remark that if S is a surface of genus g having n punctures, then $\text{Ends}_\infty(S) = \emptyset$ and $\text{Ends}(S) = \{1, \dots, n\}$ endowed with the discrete topology.

From Theorem 2.1.2 we deduce that both $\text{Ends}_\infty(S)$ and $\text{Ends}(S)$ are closed, totally-disconnected and Hausdorff. Hence the couple $\text{Ends}_\infty(S) \subset \text{Ends}(S)$ is homeomorphic to a couple of nested closed subsets of $C' \subset C$ of the standard Cantor set obtained by removing middle thirds from an interval. On the other hand, the following result tells us that every such pair defines a homeomorphism class of orientable surfaces.

Theorem 2.1.8: [Ric63]

Let $C' \subset C$ be a nested pair of closed subsets of the Cantor set. Then there exist a surface S such that $\text{Ends}_\infty(S) \subset \text{Ends}(S)$ is homeomorphic to $C' \subset C$ as nested pair of topological spaces.

Sketch of proof. Consider C as a subspace of points in the sphere \mathbb{S}^2 . Then $\text{Ends}(S = \mathbb{S}^2 \setminus C)$ is homeomorphic to C . An easy way to see this is to imagine that points in C are pulled away to infinity. If $U_1 \supseteq U_2 \supseteq \dots$ represents a point $c' \in C'$ then, by gluing a sufficiently small torus into a sufficiently small neighborhood of a point $p \in U_j \setminus U_{j+1}$ for every $j \gg 1$, we produce the desired surface S . \square

From the preceding theorem we deduce that there are uncountably many different topological types of surfaces whose fundamental group is not finitely-generated. Some of these surfaces seem to appear more naturally than others, and hence have names of their own. The following nomenclature can be attributed to A. Phillips and D. Sullivan [PS81] or E. Ghys [Ghy95].

Definition 2.1.9. An infinite-genus surface S with only one end is called a *Loch Ness monster*. An infinite-genus surface with two ends, each accumulated by genus, is called a *Jacob's Ladder*. If a surface S has only planar ends and $\text{Ends}(S)$ is homeomorphic to the Cantor set, then S is called a *Cantor tree*. On the other hand, if S has no planar ends and $\text{Ends}(S)$ is homeomorphic to the Cantor set, then S is called a *blooming Cantor tree*. These are illustrated in figure 2.1.

It is not difficult to endow each infinite-type topological surface with a translation surface structure. Indeed, given that surfaces under consideration are second countable, they always admit a Riemann surface structure and hence a holomorphic 1-form which is not identically zero. However, as we see in what follows, most of the examples of infinite-type surfaces that we discussed in the previous chapter are homeomorphic to the Loch Ness monster.

Example 2.1.10. Recall that in example 1.2.2 we defined for every $\alpha \in (0, 1)$ a surface B_α called baker's surface. We claim that for every parameter, B_α is a Loch Ness monster. It is sufficient to show this is true for $\alpha = \frac{1}{2}$, case in which $B_{\frac{1}{2}}$ has area 1. Consider figure 2.2 and for each $n \in \mathbb{N}$ the simple closed curve c_n formed by the union of the segments I_n and J_n of slope -1 given by: I_n joins the midpoint of A_n in the lower side of the square to the midpoint of B_n in the right side, and J_n joint the midpoint of B_n in the left side to the midpoint of A_n in the top. All curves in $\{c_n\}_{n \in \mathbb{N}}$ are non-isotopic. Moreover, from the figure one can see that $B_{\frac{1}{2}} \setminus \{c_1, \dots, c_n\}$ is connected for any $n \in \mathbb{N}$, therefore $B_{\frac{1}{2}}$ has infinite genus. To show that $B_{\frac{1}{2}}$ is one-ended we use the following lemma, whose proof is left as an exercise to the reader:

Lemma 2.1.11. *Let S be any surface. The space of ends $\text{Ends}(S)$ has only one element if and only if for every compact $K \subset S$ there exists a compact subset K' of S containing K and such that $S \setminus K'$ is connected.*

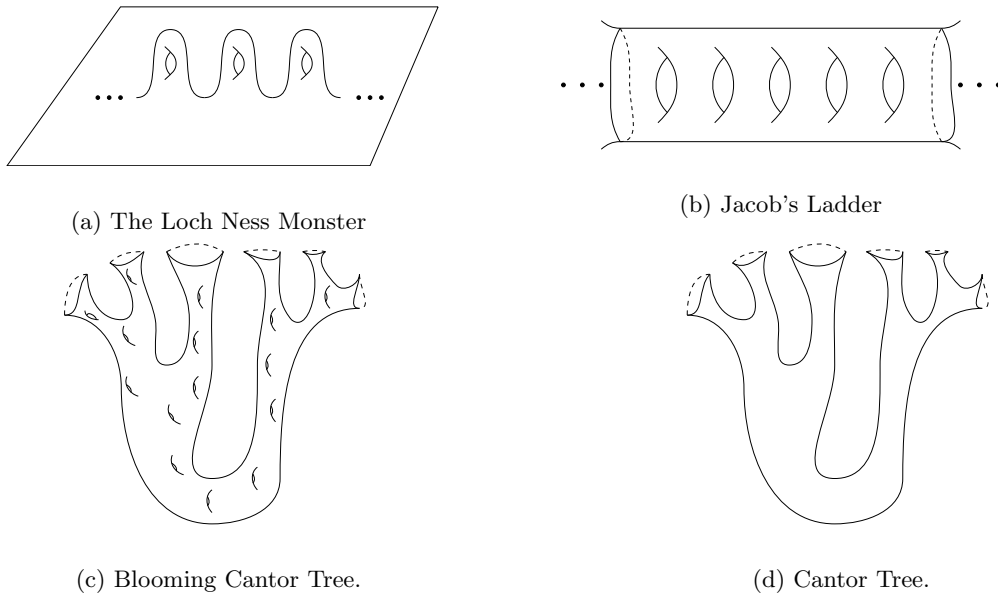


Figure 2.1: Famous infinite-type surfaces

Every compact subset $K \in B_{\frac{1}{2}}$ is contained in the complement of a small neighbourhood U of the points in the square to which the translation surface structure cannot be extended. These points are the extrema of the segments labeled by $A_n, B_n, n \in \mathbb{N}$ in the figure plus the corners a, b, c and d . If we denote by K' this complement then by definition it is a compact set. Any point in $B_{\frac{1}{2}} \setminus K'$ can be joined by an arc to the extremity of one segment in $\{A_n, B_n\}_{n \in \mathbb{N}}$. Given that the extremities of these segments accumulate to the corners b and d we conclude that one can connected any two points in the complement of K' through an arc.




Figure 2.2: $B_{\frac{1}{2}} \setminus \{c_1, \dots, c_n\}$ is connected for any $n \in \mathbb{N}$, hence it has infinite genus.

Very similar arguments are used in the next section to prove that the infinite staircase and any infinite-type surface comming from the billiard in generic triangle are homeomorphic to the Loch-Ness monster.

Exercise 2.1.12

Let $\varepsilon > 0$ be small (say less than $\frac{1}{10}$) and consider the orbit in \mathbb{R}^2 of the square $[0, \varepsilon]^2$ by the translation group $(x, y) \rightarrow (x + n, y + m)$, $(n, m) \in \mathbb{Z} \times \mathbb{Z}$. Consider the infinite polygon given by the complement of this orbit, and the translation surface M obtained from identifying pairs of opposite sides in each square. Prove that M is a Loch-Ness monster. Use the translation $(x, y) \rightarrow (x, y + 1)$ to produce a quotient of M homeomorphic to Jacob's ladder.

Exercise 2.1.13

Construct a translation surface homeomorphic to (A) the Cantor tree and (B) the Cantor blooming tree by gluing polygons as in definition 1.1.1 in the preceding chapter.

2.2 Covering spaces

A simple way to produce many examples of infinite-type translation surfaces is to consider coverings of finite-type translation surfaces where the deck transformation group is infinite. As we see in the next chapters, coverings are particularly relevant for it is possible to prove non-trivial statements about their Veech groups or the dynamics of their translation flow.

Definition 2.2.1. Let M be a translation surface, $\Sigma \subset M$ a discrete subset containing the conical singularities of M and define $M^0 := M \setminus \Sigma$. Consider a non-ramified covering map $p: \widetilde{M}^0 \rightarrow M^0$ for which the covering surface \widetilde{M}^0 has the translation surface structure defined by the pullback of the translation surface structure of M^0 via p . Let \widetilde{M} be the translation surface obtained by adding to \widetilde{M}^0 all conical singularities in $\text{Sing}(\widetilde{M}^0)$ and regular points to which the projection p can be extended continuously into a map $\widetilde{M} \rightarrow M$. We call $p: \widetilde{M} \rightarrow M$ a *translation covering*.

It is important to remark that translation covering is an abuse of language because the map $p: \widetilde{M} \rightarrow M$ is not necessarily surjective. This is illustrated in the case of the infinite staircase in figure 2.3.

In analytic terms, if we think of M^0 as a pair (X, ω) , then $\widetilde{M}^0 = (\widetilde{X}^0, \eta)$, where \widetilde{X}^0 is the natural Riemann surface structure inherited from M^0 via p and $\eta = p^*\omega$. By definition all points in \widetilde{M}^0 are regular and the deck transformation group $\text{Deck}(p)$ acts on \widetilde{M}^0 by translations (see definition 1.1.26). We say that the translation covering is *finite* if the fibers of the covering map are finite.

Exercise 2.2.2

Prove that if $p: \widetilde{M} \rightarrow M$ is a finite translation covering over a finite-type translation surface M , then \widetilde{M} is always a translation surface.

G-coverings. Let $p: \widetilde{M} \rightarrow M$ be a translation covering. Recall that the fibers of the covering map $p: \widetilde{M}^0 \rightarrow M^0$ are naturally identified with the cosets of the subgroup $p_*\pi_1(\widetilde{M}^0) < \pi_1(M^0)$ and that any covering of M^0 is determined (up to covering isomorphism) by (the conjugacy class of) a subgroup of $\pi_1(M^0)$. Given that the induced map $p_*: \pi_1(\widetilde{M}^0) \rightarrow \pi_1(M^0)$ is always injective, we will also write $\pi_1(\widetilde{M}^0) < \pi_1(M^0)$ when there is no ambiguity.

Most of the translation coverings we study in this text are normal coverings, that is to say $\pi_1(\widetilde{M}^0)$ is a normal subgroup of $\pi_1(M^0)$, hence the deck transformation group $\text{Deck}(p)$ of the translation covering can be identified with $G = \pi_1(M^0)/\pi_1(\widetilde{M}^0)$ and acts transitively on the fibers of the covering. *In this situation* $p: \widetilde{M}^0 \rightarrow M^0$ is called a *G-covering*, because the group G also acts freely and properly discontinuously on \widetilde{M}^0 , $M^0 = \widetilde{M}^0/G$ and p is just the projection $\widetilde{M}^0 \rightarrow \widetilde{M}^0/G$. For more details on *G-coverings* we refer the reader to the classical book of W. Fulton [Ful95].

Example 2.2.3 (\mathbb{Z} -coverings and the infinite staircase). Let M and Σ as before. A \mathbb{Z} -covering of $M^0 = M \setminus \Sigma$ is given by the kernel of a surjective morphism $f: \pi_1(M^0) \rightarrow \mathbb{Z}$. Since \mathbb{Z} is abelian, f factorizes through a map $\phi: H_1(M^0; \mathbb{Z}) \rightarrow \mathbb{Z}$, i.e. $\phi \in H^1(M^0; \mathbb{Z})$. On the other hand, the algebraic intersection form

$$\iota: H_1(M^0; \mathbb{Z}) \times H_1(M, \Sigma; \mathbb{Z}) \rightarrow \mathbb{Z}$$

is non-degenerate and hence produces a natural identification $H^1(M^0; \mathbb{Z}) \simeq H_1(M, \Sigma; \mathbb{Z})$. In particular, the map ϕ is represented by a unique element c of $H_1(M, \Sigma; \mathbb{Z})$, that is for every, $v \in H_1(M^0, \mathbb{Z})$ we have $\phi(v) = \iota(v, c)$.

Exercise 2.2.4

Show that an element ϕ of $H^1(M^0; \mathbb{Z})$ is surjective if and only if it is represented by a primitive class c in $H_1(M, \Sigma; \mathbb{Z})$ (i.e. there is no $n \in \mathbb{Z} \setminus \{0\}$ such that $c = nc'$ with $c' \in H_1(M, \Sigma; \mathbb{Z})$).

The preceding exercise implies that every primitive element $c \in H_1(M, \Sigma; \mathbb{Z}) \simeq H^1(M^0; \mathbb{Z})$ determines a \mathbb{Z} -covering of M . In the following lines we explain in detail how the infinite staircase introduced in example 1.2.3 can be seen as a \mathbb{Z} -covering. Consider the rectangle P on the right hand side of Figure 2.3 formed by two unit squares. If we identify using translations parallel sides labeled with the same letter the result is a flat torus M with two marked points $\Sigma \subset M$ coming from the corners of the aforementioned squares. Denote by $\{A, B, C\}$ the oriented basis of $H_1(M, \Sigma; \mathbb{Z})$ defined by the identification of the sides $\{A^\pm, B^\pm, C^\pm\}$ respectively (in the figure the orientation of these is illustrated). The \mathbb{Z} -covering $p: \widetilde{M} \rightarrow M$ defined by the cycle $c = B - A$ can be concretely constructed from the polygon P and c as follows. Let $\{\phi_A, \phi_B, \phi_C\}$ and $\{\gamma_A, \gamma_B, \gamma_C\}$ be the basis of $H^1(M^0; \mathbb{Z}) = \text{Hom}(H_1(M^0; \mathbb{Z}), \mathbb{Z})$ defined by $\{A, B, C\}$ and its dual in $H_1(M^0; \mathbb{Z})$, respectively. If we denote by ϕ_c the cohomology class determined by c , then \widetilde{M}^0 is given by performing the following identifications (using translations) in the infinite family $P \times \mathbb{Z}$: identify the side (A^+, n) with $(A^-, n + \phi_c(\gamma_A)) = (A^-, n - 1)$, (B^+, n) with $(B^-, n + \phi_c(\gamma_B)) = (B^-, n + 1)$ and (C^+, n) with $(C^-, n + \phi_c(\gamma_C)) = (C^-, n)$. The covering space $p: \widetilde{M}^0 \rightarrow M^0$ is depicted on the left hand side of Figure 2.3 and is exactly the *infinite staircase* introduced in Section 1.2.3. It is important to remark that the metric completion of \widetilde{M}^0 is not a surface, because in order to obtain it we have to add four infinite angle singularities, none of which has a compact neighbourhood. Therefore in this case the extension of $p: \widetilde{M}^0 \rightarrow M^0$ to the metric completion is not a map between translation surfaces.




Figure 2.3: The staircase as a \mathbb{Z} -cover of a torus.

We now illustrate how to generalize the above construction to arbitrary \mathbb{Z}^d -coverings, for $d \geq 2$. Let P be a polygon with an even number of parallel sides of the same length $\{\xi_i^+, \xi_i^-\}_{i=1}^N$ and M the compact translation surface obtained by identifying ξ_i^+ with ξ_i^- using a translation for all $i = 1, \dots, N$. Denote by ξ_1, \dots, ξ_N the identified segments in M and $\Sigma \subset M$ the finite set of points defined by their extremities, i.e. corresponding to the vertices of P . The set $\{\xi_1, \dots, \xi_N\}$ is a basis for the relative homology $H_1(M, \Sigma; \mathbb{Z}) \simeq \mathbb{Z}^N$. Let $\{\phi_{\xi_1}, \dots, \phi_{\xi_N}\}$ be the basis of $H^1(M^0; \mathbb{Z})$ determined by $\{\xi_1, \dots, \xi_N\}$ and $\{\gamma_1, \dots, \gamma_N\}$ its dual in $H_1(M^0; \mathbb{Z})$.

Now let $\mathcal{C} = \{c_1, \dots, c_d\}$ be $d \geq 2$ linearly independent cycles in $H_1(M, \Sigma; \mathbb{Z})$ and $\{\phi_{c_1}, \dots, \phi_{c_d}\}$ the elements these determine in $H^1(M^0, \mathbb{Z})$. We construct a \mathbb{Z}^d -covering $p: \widetilde{M}^0 \rightarrow M^0$ as follows: by definition the set \mathcal{C} spans a rank d sublattice $\Lambda(\mathcal{C})$ of $H_1(S, \Sigma; \mathbb{Z})$, i.e. isomorphic to \mathbb{Z}^d . Let P^0 be P deprived of

its vertices and in $P^0 \times \Lambda(\mathcal{C}) \simeq P^0 \times \mathbb{Z}^d$ we identify (using translations) for each $i = 1, \dots, d$ and $\hat{n} \in \mathbb{Z}^d$ the side (ξ_i^+, \hat{n}) in $P^0 \times \{\hat{n}\}$ with the side $(\xi_i^-, \hat{n} + (\phi_{c_1}(\gamma_i), \dots, \phi_{c_d}(\gamma_i)))$ in $P^0 \times \{\hat{n} + (\phi_{c_1}(\gamma_i), \dots, \phi_{c_d}(\gamma_i))\}$. The quotient $\widetilde{M}^0 := P^0 \times \Lambda(\mathcal{C}) / \sim$ is the \mathbb{Z}^d -covering of M^0 defined by the cycles in \mathcal{C} .

Exercise 2.2.5: The infinite staircase is a Loch Ness monster.

Let $p: \widetilde{M}^0 \rightarrow M^0$ be the \mathbb{Z} -covering given by the infinite staircase as depicted in figure 2.3.

1. For each $k \in \mathbb{Z}$, let $V_k \subset \widetilde{M}^0$ be the maximal vertical cylinder containing the horizontal saddle connection defined by the side A_k^+ and γ_k its the core curve. Show that $S \setminus \bigcup_{k \in \mathbb{Z}} \gamma_k$ is connected and any to curves in $\bigcup_{k \in \mathbb{Z}} \gamma_k$ are non-isotopic.
2. If $K \subset \widetilde{M}^0$ is compact, show that there exist a neighbourhood U of $\text{Sing}(\widetilde{M}^0)$ whose complement contains K . Use that neighbourhood to construct a compact set K' containing K such that $\widetilde{M}^0 \setminus K'$ is connected. *Hint:* the set K must be contained in a finite union of steps of the staircase; then proceed as in the case of baker's surface explained in example 2.1.10.

Deduce from the preceding points that the infinite staircase is homeomorphic to the Loch Ness monster.

Remark 2.2.6. It is somehow surprising that the infinite staircase is a Loch Ness monster, for it actually looks like having two ends and moreover the deck transformation group of the covering $p: \widetilde{M}^0 \rightarrow M^0$ it defines is \mathbb{Z} , which is a group with two ends. The main issue here is that $\text{Deck}(p)$ does not act cocompactly (because of the existence of infinite cone angle singularities) and the infinite staircase is not quasi-isometric to the Cayley graph of \mathbb{Z} .

Exercise 2.2.7

Let M be a finite-type translation surface obtained from a polygon P as in the preceding paragraphs and suppose G is an infinite group for which there exist a surjective group morphism $f: \pi_1(M \setminus \Sigma) \rightarrow G$. Describe in concrete geometrical terms how to construct the G -covering $\widetilde{M} \rightarrow M$ defined by the kernel of f .

Example 2.2.8 (Origamis). The infinite staircase belongs to the class of translation surfaces called *square-tiled surfaces* or *origamis*. In what follows we explain several equivalent ways to define this kind of surfaces.

Definition 2.2.9. Let E be the flat torus \mathbb{C}/\mathbb{Z}^2 with a marked point at 0. An *square-tiled surface* (or *origami*) is a translation covering $p: M \rightarrow E$, ramified at most over 0. A square-tiled surface is infinite if the fibers of the corresponding cover are infinite.

Square-tiled surfaces are described by subgroups of the free group F_2 . Indeed, for every square-tiled surface $p: M \rightarrow E$ we can consider the non-ramified covering $p_1: M^0 \rightarrow E^0$, where $E^0 := E \setminus \{0\}$ and $M^0 = M \setminus p^{-1}(0)$. Hence, p determines a subgroup H of $F_2 = \pi_1(E^0)$ up to conjugacy. Reciprocally, every subgroup H of the free group on two generators F_2 determines a non-ramified translation covering of the punctured torus $p: M^0 \rightarrow E^0$, up to covering isomorphism. If one completes both M^0 and E^0 using the distance induced by the flat metric, the result is a square-tiled surface $p: M \rightarrow E$. For example, the commutator subgroup $[F_2, F_2]$ determines a covering $p: \mathbb{C} \setminus \mathbb{Z}^2 \rightarrow \mathbb{C}/\mathbb{Z}^2 \setminus \{0\}$ with deck transformation group isomorphic to \mathbb{Z}^2 . The completion of this covering w.r.t. the distance induced by the flat metric is just the universal cover $p: \mathbb{C} \rightarrow \mathbb{C}/\mathbb{Z}^2$.

Exercise 2.2.10

Describe a subgroup $H < F_2$ that determines the infinite staircase as a square-tiled surface. Is H in this case normal?

From the constructive point of view, a square-tiled surface is just a connected surface tiled by copies of the same square. More precisely, let $\{C_i\}_{i \in I}$ be a family of copies of the unit square where the set of indexes I is at most countable. Consider two bijections $r, t: I \rightarrow I$ (the gluing rules) and for each $i \in I$ glue using a translation:

- the right edge of C_i to the left edge of $C_{r(i)}$ and
- the top edge of C_i to the bottom edge of $C_{t(i)}$,

in such a way that the resulting topological space is connected. Remark that as in the case of the infinite staircase, if the vertices of the squares C_i are taken into consideration then the topological space that results from the gluings described above could be a non-locally compact space, for infinite angle singularities may appear.

Exercise 2.2.11

Describe a pair of bijections $r, t : \mathbb{Z} \rightarrow \mathbb{Z}$ for which the gluings in the family $\{C_k\}_{k \in \mathbb{Z}}$ produce the infinite staircase.

Example 2.2.12 (Irrational billiards). In Section 1.2.1 we explained how the dynamics of a billiard ball on an Euclidean polygonal table P can be interpreted as the dynamics of the translation flow on a translation surface $M(P)$ obtained from P by unfolding. In the following lines we sketch the proof of theorem 1.2.3 for totally-irrational triangles. More precisely, we show that for every totally-irrational triangle the surface $M(P)$ is homeomorphic to the Loch Ness monster. The general case is detailed in [Val09b].

Let $\{\xi_1, \xi_2, \xi_3\}$ be the sides of a totally irrational triangle¹ P and $\mathbb{S}^2(P)$ be the Euclidean surface² obtained by identifying two copies of $P \setminus \text{Vertices}(P)$ along sides with the same labels using orientation preserving isometries. The Euclidean surface $\mathbb{S}^2(P)$ is homeomorphic to a three-punctured sphere. Remark that with respect to its natural Euclidean metric (inherited by from $P \subset \mathbb{R}^2$) the sectional curvature of $\mathbb{S}^2(P)$ at every point is zero but $\mathbb{S}^2(P)$ is not a translation surface.

Exercise 2.2.13

Let P be a totally-irrational triangle.

1. Show that there is a natural normal covering map $p : M(P) \rightarrow \mathbb{S}^2(P)$ coming from the unfolding construction whose deck transformation group is isomorphic to \mathbb{Z}^2 .
2. Show that $\text{Deck}(p)$ acts by isometries on $M(P)$.

Hence $p : M(P) \rightarrow \mathbb{S}^2(P)$ is a \mathbb{Z}^2 -covering but its base is *not* a translation surface. We include this example to illustrate how the discussion about G -coverings we had before can be extended to covering maps where the base is an Euclidean surface and the covering space a translation surface.

The surface $M(P)$ has infinite genus because, by choosing the appropriate subgroup of F_2 , one can find an intermediate covering

$$M(P) \rightarrow M_n \rightarrow \mathbb{S}^2(P)$$

where M_n is a finite-type positive genus surface. To prove that $M(P)$ has only one end we use lemma 2.1.11 and argue in a similar way as we did for baker's surface or the infinite staircase: if $K \subset M(P)$ is compact, then $p(K) \subset \mathbb{S}^2(P)$ is contained in K_ε , the complement of the union of a sufficiently small ε -neighbourhood of $\text{Vertices}(P)$ in $\mathbb{S}^2(P)$. Since K is compact, it is possible to find finitely many copies of K_ε in $p^{-1}(K_\varepsilon)$ covering K . The union of these finitely many copies of K_ε defines a compact set $K' \subset M(P)$. Think of $M(P)$ as $D \times \mathbb{Z}^2$, where D is a fundamental domain for the action of the deck transformation group (isomorphic to \mathbb{Z}^2) of $p : M(P) \rightarrow \mathbb{S}^2(P)$. Since K' is compact, it is contained in a ball of the form:

$$B_N := \{D \times \{(n, m)\} \mid n^2 + m^2 < N\}$$

for some large $N \in \mathbb{N}$. Remark that $M(P) \setminus B_N$ is connected, because \mathbb{Z}^2 is a group with only one end. Every point $z \in M(P) \setminus K'$ can be connected by a path to a point z' which projects to $p(z') \in \mathbb{S}^2(P) \setminus K_\varepsilon$. Given that P is a totally irrational triangle, it is possible to connect z' through an arc to a copy of D outside B_N , that is a fundamental of the form $D \times \{(n_0, m_0)\}$, for some $n_0, m_0 > N$. Given that the choice of z was arbitrary and $M(P) \setminus B_N$ is connected, one can connect any two points in $M(P) \setminus K'$ through an arc.

¹An Euclidean triangle P is said to be *totally-irrational* if it has the following property: if $\lambda_i \pi, \lambda_j \pi$ are interior angles of P and $n_i \lambda_i + n_j \lambda_j \in \mathbb{Z}$ for some $n_i, n_j \in \mathbb{Z}$, then $n_i = n_j = 0$. Remark that totally-irrational triangles are generic in the space of triangles.

²By Euclidean surface we mean a surface endowed with an atlas where transition functions are isometries of the Euclidean plane.

Remark 2.2.14. If $p: \tilde{S} \rightarrow S$ is a non-ramified normal covering and S is a compact surface, then $\text{Ends}(\tilde{S})$ is homeomorphic to the space of ends of the Cayley graph of $\text{Deck}(p)$, which we denote by $\Gamma(\text{Deck}(p))$. If in addition \tilde{S} has genus, then by using the transitive action of $\text{Deck}(p)$ on fibers we conclude that $\text{Ends}(\tilde{S}) = \text{Ends}_\infty(\tilde{S})$ and $\text{Ends}(\Gamma(\text{Deck}(p)))$ are homeomorphic. This means that the deck transformation group imposes some restrictions of the topology of a non-ramified covering. For example, given that a finitely-generated group has either 0,1,2 or infinitely many ends, there exist no non-ramified normal covering $\tilde{S} \rightarrow S$ (with S a compact surface) for which $\text{Ends}(\tilde{S})$ has $n \geq 3$ elements.

Question 2.2.15

Let P be a generalized irrational polygon (see definitions 1.2.4 and 1.2.5) which is not simply connected. Describe the topology of $S(P)$. Is it always homeomorphic to a Loch Ness monster?

Dilation surfaces. The infinite staircase forms part of a larger class of infinite-type translation surfaces called λ -staircases, which were defined in Section 1.2.6. Every λ -staircase admits an affine automorphism with a derivative of the form $r_+ \text{Id}$, where $r_+ > 1$ is the largest root of $x^2 - \lambda x + 1$ and the corresponding quotient is a compact surface whose transition maps are dilatations (see Figure 1.21). We remark that this is similar to the covering situation for irrational billiards $p: M(P) \rightarrow \mathbb{S}^2(P)$ that we described above. In this section we briefly discuss dilation surfaces and how they are related to translation surfaces.

A *dilation atlas* on a topological surface S is an atlas $\mathcal{T} = \{\phi_i: U_i \rightarrow \mathbb{C}\}$ for which the transition maps $\phi_j \circ \phi_i^{-1}: \phi_i(U_i \cap U_j) \rightarrow \phi_j(U_i \cap U_j)$ are dilatations, *i.e.* maps of the form $z \mapsto az + b$ with $a \in \mathbb{R}^*$.

Definition 2.2.16. Let S be a topological surface and $\Sigma \subset S$ a discrete set of points. A *dilation surface* is a pair (S, \mathcal{T}) where \mathcal{T} is a maximal dilation atlas on $S \setminus \Sigma$ that can be extended to a Riemann surface atlas on S .

In other words, if (S, \mathcal{T}) is a dilation surface then $S \setminus \Sigma$ admits a (G, X) -structure (in the sense of Thurston) for which $X = \mathbb{C}$ and G is the group of dilations in $\text{Aut}(\mathbb{C})$. In this sense, every point in $p \in \Sigma$ has a punctured neighbourhood (with fundamental group isomorphic to \mathbb{Z} , generated by a loop γ) whose holonomy representation is of the form $\gamma \rightarrow (z \rightarrow re^{2\frac{\pi}{k}i}z)$, for some $r \in \mathbb{R}^*$ and $k \in \mathbb{N}$. By the similarity with the translation surface case, we call say in this case that p has a singularity of angle $2\pi k$ at p . One has to be careful though for this does not mean necessarily that p is a conical singularity in the sense of the definition included in exercise 1.1.4, because dilation surfaces only accept Euclidean metrics when they are translation surfaces.

Translation coverings. Associated to the kernel of the holonomy of the (G, X) -structure of a dilation surface S we have a normal covering $\pi: \tilde{S} \rightarrow S$. By definition, \tilde{S} inherits an atlas from S where all transition functions are translations and hence we call it the *translation covering associated to the dilation surface* S . Given that \mathbb{R}^* is an abelian torsion free group, every non-trivial translation covering over a dilation surface S with finitely generated fundamental group has deck transformation group isomorphic to \mathbb{Z}^d , for some $d \geq 1$. The archetypical example of this kind of coverings are λ -staircases, as illustrated in figure 1.21.

2.3 Singularities and the space of linear approaches

In this section we explain how the natural extension of the tangent bundle of a translation surface M to its metric completion provides the necessary invariants to describe the singularities of M . Let us begin by recalling two indispensable notions.

Saddle connections and holonomy vectors. As we saw in the preceding chapter, the term *flow* is an abuse of language when referring to the translation flow F_θ^t on M for it might not be defined for all points in M for all $t \in \mathbb{R}$. We distinguish those orbits of F_θ^t which are not defined for all times in the future or in the past.

Definition 2.3.1. Let M be a translation surface, $z \in M$ and $I \subset \mathbb{R}$ the *maximal* domain of definition of $F_\theta^t(z)$.

- If $I = \mathbb{R}$, the orbit of z under F_θ^t is called a *regular geodesic*³,
- if I is unbounded but $I \neq \mathbb{R}$, this orbit is called a *separatrix*. Furthermore, if I is bounded below (respectively above) the orbit is called a forward separatrix (respec. backward), and
- if I is bounded, the orbit of z under F_θ^t called a *saddle connection* in direction θ .

If γ is the trace of a saddle connection in direction θ , we denote by v_γ the vector of length $|\gamma|$ and direction θ . This vector is called the *holonomy vector associated* to the saddle connection γ . The set formed by all holonomy vectors is denoted by $V_{hol}(M)$.

It is clear from the definition that saddle connections are geodesics whose extremities are points in $\text{Sing}(M)$ or marked points⁴, and that $v \in V_{hol}(M)$ implies that $-v \in V_{hol}(M)$. On the other hand, affine automorphisms send saddle connections to saddle connections and hence $\Gamma(M)$, the Veech group of M , acts on $V_{hol}(M)$. As we see in the next chapter, this action can be used to deduce properties of $\Gamma(M)$, for example it is a well known fact (see [MT02]) that for finite-type translation surfaces $V_{hol}(M)$ is an infinite discrete subset of \mathbb{R}^2 and hence $\Gamma(M)$ must be discrete in this context.

Exercise 2.3.2

Let P be a polygon and $M(P)$ the translation surface obtained by unfolding.

1. Show that the natural projection $\pi_P : M(P) \rightarrow P$ sends saddle connections to *generalized diagonals*, that is, billiard trajectories whose extremities are vertices of P . Is π_P restricted to a saddle connection a bijection? Is every generalized diagonal in P the image of a saddle connection?
2. For which triangles P does $M(P)$ not have saddle connections?
3. Show that if P is a triangle whose interior angles are all irrational multiples of π , then $V_{hol}(M(P))$ is an infinite non-discrete subset of the plane that is contained in the complement of a disc $\{(x, y) \in \mathbb{R}^2 \mid |z| < r\}$ for some $r > 0$ that depends only on P .

Question 2.3.3

Let P be an irrational polygon. Is $V_{hol}(M(P)) \subset \mathbb{R}^2$ bounded above? Is the interior of the closure of $V_{hol}(M(P))$ empty?

In the preceding chapter we introduced conical, infinite angle and wild singularities (see definition 1.1.22) and the notion of a tame translation surface M : these are surfaces for which $\text{Sing}(M)$ does not contain wild singularities. If M is a tame translation surface we can describe a small neighbourhood of any point in its metric completion: it is either isometric to a neighbourhood of $0 \in \mathbb{C}$ or to a cyclic (maybe infinite) covering of a neighbourhood of $0 \in \mathbb{C}$ ramified over 0 . The class of tame translation surfaces includes all compact translation surfaces, those arising from billiards on irrational polygons, wind-tree models, and, as the following exercise shows, tameness is preserved when taking coverings.

Exercise 2.3.4

Prove that if $p : \widetilde{M} \rightarrow M$ is a G -covering whose base is a tame translation surface (not necessarily of finite topological type), then \widetilde{M} is also a tame translation surface.

Example 2.3.5. The following example, inspired in the infinite-step billiard studied in [DEDML98], shows that contrary to intuition *infinite-type tame translation surfaces of finite area exist*. We refer the reader to figure 2.4. Consider the step function:

$$f(x) = \begin{cases} \frac{1}{n^2} & \text{for } x \in [n-1, n) \text{ and } n \in \mathbb{N}, \\ f(-x) & \text{for } x < 0. \end{cases}$$

³This nomenclature is justified by remark 1.1.29.

⁴A marked point or a puncture is a regular point in \widetilde{M} that does not appear in M .

and let P be the infinite polygon in \mathbb{R}^2 obtained by taking the closure of the open set bounded by the graph of f and the real line, which we consider as the polygonal line defined by the union of the segments $[n, n+1]$, $n \in \mathbb{Z}$. Let M be the translation surface obtained from P by identifying opposite sides using translations. By construction the area of M is equal to $2 \sum_{i=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{3}$. It is not difficult to check that the diameter of M is infinite, and hence the only points defining singularities are those coming from the vertices of P . More precisely, $\text{Sing}(M)$ is a countable set of conic singularities, each of total angle 6π .




Figure 2.4: An infinite-type tame translation surface of finite area.

2.3.1 The space of linear approaches

Every translation surface M can be endowed with a distance coming from the intrinsic flat metric which allows us to consider, as we defined in the preceding chapter, its metric completion \widehat{M} and $\text{Sing}(M) \subset \widehat{M}$ the set of singularities. Given that $M^0 := \widehat{M} \setminus \text{Sing}(M)$ has no singularities, its tangent bundle TM^0 is isomorphic to the product $M^0 \times \mathbb{R}^2$. In the following paragraphs we introduce a continuous extension of this bundle, deprived of its zero section, to \widehat{M} . At the level of the *unit* tangent (sub)bundle, this extension produces a topological space called the *space of linear approaches* of M , which we use to describe the geometry of a translation surface near any singularity. The germ of the idea behind the space of linear approaches can be found in R. Chamanara's work [Cha04]. As seen in example 1.2.2, the metric completion of baker's surface B_α with respect to its natural is achieved by adding one wild singularity z_∞ . R. Chamanara points out that “[g]eometrically, the surface spirals infinitely many times around this point”. The space of linear approaches formalizes this intuition. Moreover, one can extract enough information from this space to calculate the Veech group of baker's surface, as done in Section 3.5. The approach and notations we use to present the space of linear approaches slightly differ from those used in [BV13] and [Ran16], which are the principal references on the subject.

Definition 2.3.6. Given $\varepsilon > 0$, let

$$\mathcal{L}^\varepsilon(M) := \{\text{geodesic trajectories } \gamma : (0, \varepsilon) \rightarrow M^0\},$$

where geodesics⁵ are taken with respect to the flat metric on M . Two geodesics $\gamma_1 \in \mathcal{L}^{\varepsilon_1}(M)$ and $\gamma_2 \in \mathcal{L}^{\varepsilon_2}(M)$ are said to be equivalent if and only if $\gamma_1(t) = \gamma_2(t)$ for all $t \in (0, \min\{\varepsilon_1, \varepsilon_2\})$. We denote this equivalence relation by \sim and define:

$$T\widehat{M} := \bigsqcup_{\varepsilon > 0} \mathcal{L}^\varepsilon(M) / \sim \quad (2.3)$$

The equivalence class of γ will be denoted by $[\gamma]$.

To understand better what elements of $T\widehat{M}$ are about, endow $TM^0 = M^0 \times \mathbb{R}^2$ with its natural product topology and suppose that there exists $\gamma : (0, \varepsilon) \rightarrow M^0$ in $\mathcal{L}^\varepsilon(M)$ for which $\lim_{t \rightarrow 0} \gamma(t) = z_0$ exists in M^0 . Then, given that geodesics are parametrized with constant non-zero speed, the limit $\lim_{t \rightarrow 0} (\gamma(t), \gamma'(t)) = (z_0, v_0)$ exists in TM^0 (deprived of the zero section) and is independent of the representative chosen within $[\gamma]$. Hence we can associate to $[\gamma]$ the tangent vector γ' based at z_0 and think that this class encodes the fact that one can *approach* the point z_0 using the a translation flow with speed v_0 .

In general $\lim_{t \rightarrow 0} \gamma(t)$ only exists in the metric completion \widehat{M} and is independent from the representative chosen in $[\gamma]$. Hence we can think of $(\lim_{t \rightarrow 0} \gamma(t), \gamma')$ as the tangent vector γ' based at the point $z_0 = \lim_{t \rightarrow 0} \gamma(t)$, keeping in mind that strictly speaking there is no tangent space at z_0 for in general one cannot extend the differentiable structure of M^0 to all points in the metric completion.

⁵Instead of speaking in terms of geodesics one can phrase the definitions in this section in terms of trajectories of the translation flows F_θ^t on M . We have decided not to do this to be more coherent with the exposition found in [BV13] and to keep the notation simple.

Let us now define a topology on $T\widehat{M}$. The uniform convergence of functions induces the uniform topology on $\mathcal{L}^\varepsilon(M)$. For each $\varepsilon \geq \varepsilon'$ the restriction of element in $\mathcal{L}^\varepsilon(M)$ to the interval $(0, \varepsilon')$ defines a continuous injection:

$$\rho_{\varepsilon'}^{\varepsilon} : \mathcal{L}^\varepsilon(M) \rightarrow \mathcal{L}^{\varepsilon'}(M)$$

and hence $\langle \mathcal{L}^\varepsilon(M), \rho_{\varepsilon'}^{\varepsilon} \rangle$ is a direct system of topological spaces over (\mathbb{R}^+, \geq) . Since for every $\varepsilon \geq \varepsilon' > 0$ the projection map $\gamma \mapsto [\gamma]$ from $\mathcal{L}^\varepsilon(M)$ to $T\widehat{M}$ is injective and commutes with $\rho_{\varepsilon'}^{\varepsilon}$, we have the equality of sets

$$T\widehat{M} = \varinjlim \mathcal{L}^\varepsilon(M) \quad (2.4)$$

It is then natural to endow $T\widehat{M}$ with the limit topology. Let us describe more precisely this topological space. For every $v \in TM^0$, $r, t > 0$ define:

$$B(v, t, r) := \{[\gamma] \mid d(\gamma(t), v) < r\}. \quad (2.5)$$

Remark that if $[\gamma] \in B(v, t, r)$, then there exist a representative $\gamma : (0, \varepsilon) \rightarrow M^0$ for which $\varepsilon > t$. We claim that the family of sets $\mathcal{B} := \{B(v, t, r)\}_{v \in TM^0; r, t > 0}$ is a basis for the limit topology. Indeed, it is sufficient to remark that for every fixed $\varepsilon > 0$, the family of sets

$$\mathcal{B}_\varepsilon := \{B_\varepsilon(v, t, r) := \{\gamma \in \mathcal{L}^\varepsilon(M) \mid d(\gamma(t), v) < r\}\}_{v \in TM^0; r, t > 0}$$

is a basis for the topology on $\mathcal{L}^\varepsilon(M)$.

Exercise 2.3.7

Let M be a translation surface. To every $[\gamma] \in T\widehat{M}$ we can associate its base point $\text{bp}[\gamma] := \lim_{t \rightarrow 0} \gamma(t) \in \widehat{M}$ and its direction^a $\text{dir}[\gamma] = \gamma' \in \mathbb{R}^2$. This association defines two maps $\text{bp} : T\widehat{M} \rightarrow \widehat{M}$ and $\text{dir} : T\widehat{M} \rightarrow \mathbb{R}^2$ called the *basepoint* and *direction* map respectively. Prove that:

1. The space $T\widehat{M}$ is Hausdorff and second countable.
2. Let $(TM^0)^*$ denote the tangent bundle TM^0 deprived of its zero section. Show that the map

$$i : (TM^0)^* \rightarrow T\widehat{M}$$

that associates to each (z, v) the class $[\gamma]$ defined by $(\text{bp}[\gamma], \text{dir}[\gamma]) = (z, v)$ is a topological embedding.

Hint: use the fact that M is a Riemann surface and hence second countable; prove that the maps bp and dir are continuous.

^aHere the convention is to send the horizontal direction to zero.

Functoriality. Let $f : M_1 \rightarrow M_2$ be an affine map whose derivative lies in $\text{GL}(2, \mathbb{R})$ and $\widehat{f} : \widehat{M}_1 \rightarrow \widehat{M}_2$ its (unique) continuous extension to the metric completion. We define the map $\widehat{T}f : T\widehat{M}_1 \rightarrow T\widehat{M}_2$ by setting $\widehat{T}f[\gamma] := [f \circ \gamma]$. It is clear from the definition that $\widehat{T}Id = Id_{T\widehat{M}}$ and that $\widehat{T}(f \circ g) = \widehat{T}f \circ \widehat{T}g$. Moreover, if $\lim_{t \rightarrow 0} \gamma(t) = z_0$ for any given $\gamma \in \mathcal{L}^\varepsilon(M)$, then $\text{bp}(\widehat{T}f[\gamma]) = \widehat{f}(z_0)$ and $\text{dir}(\widehat{T}f[\gamma]) = Df \cdot \gamma'$. It is in this last sense that $\widehat{T}f$ is an extension of the tangent map $Tf : TM_1^0 \rightarrow TM_2^0$.

Exercise 2.3.8

Prove that the tangent map $\widehat{T}f$ defined above is continuous. *Disclaimer:* as shown in example 2.3.17 below the space $T\widehat{M}$ is not in general a regular (T_3) space and hence this exercise does not follow immediately from classical extension theorems.

We now introduce the space of linear approaches of a translation surface. The key point here is that in order to understand the geometry near a singularity it is sufficient to consider approximations at unit speed.

Definition 2.3.9. Let M be a translation surface. The space of linear approaches of M is the subspace of $T\widehat{M}$ defined by:

$$T^1\widehat{M} := \{[\gamma] \in T\widehat{M} \mid \gamma' \text{ is a unit vector}\} \quad (2.6)$$

Every $[\gamma] \in T^1\widehat{M}$ is called a *linear approach* to the point $\lim_{t \rightarrow 0} \gamma(t) \in \widehat{M}$. For every $z \in \widehat{M}$, the subspace $T_z^1\widehat{M} := \text{bp}^{-1}(z) \cap T^1\widehat{M}$ is called the space of linear approaches to the point $z \in \widehat{M}$.

Recall that $\mathbb{S}^1 \subset \mathbb{C}$ denotes the unit circle given by the image of the segment $[0, 2\pi]$ under the exponential map. If z is a regular point of M the restriction $\text{dir}_| : T_z^1\widehat{M} \rightarrow \mathbb{S}^1$ defines a homeomorphism and hence the space of linear approaches to a regular point is naturally parametrized by $\mathbb{R}/2\pi\mathbb{Z}$. If z is a conical singularity of angle $2k\pi$ the map $\text{dir}_|$ can be lifted to a homeomorphism between $T_z^1\widehat{M}$ and the natural $k : 1$ covering of \mathbb{S}^1 , hence in this case $T_z^1\widehat{M}$ is parametrized by $\mathbb{R}/2k\pi\mathbb{Z}$. In the same line of thought it can be proven that if z is an infinite angle singularity, $T_z^1\widehat{M}$ is naturally parametrized by \mathbb{R} . A simple way to understand these facts is to pick a linear approach $[\gamma]$ in $T_z^1\widehat{M}$ and think of it as a geodesic segment anchored at the point $z = \text{bp}[\gamma]$. The aforementioned parametrizations are achieved by “turning the geodesic segment around x ” clockwise and counter-clockwise until we come back to the initial position or we browse the whole space $T_z^1\widehat{M}$.

Exercise 2.3.10

Is there a translation surface M for which there exists a point z in the metric completion \widehat{M} such that $T_z^1\widehat{M} = \emptyset$?

We summarize the preceding discussion in the following:

Theorem 2.3.11

Let M be a translation surface, \widehat{M} its metric completion w.r.t. the flat metric and $M^0 = M \setminus \text{Sing}(M)$. There exists continuous extensions of $(TM^0)^*$ (the tangent bundle TM^0 deprived of the zero section) and the unit tangent T^1M^0 to \widehat{M} , that we denote by $T\widehat{M}$ and $T^1\widehat{M}$ respectively. These spaces are Hausdorff, second countable but in general not regular. Both $(TM^0)^*$ and T^1M^0 are dense in $T\widehat{M}$ and $T^1\widehat{M}$. Moreover, these extensions are functorial in the category of translation surfaces: the derivative of any every affine map $f : M_1 \rightarrow M_2$ whose derivative lies in $\text{GL}(2, \mathbb{R})$ can be extended to a continuous map $\widehat{T}f : T\widehat{M}_1 \rightarrow T\widehat{M}_2$. This extension is such that $\widehat{T}Id = Id_{T\widehat{M}}$ and that $\widehat{T}(f \circ g) = \widehat{T}f \circ \widehat{T}g$.

Remark 2.3.12. The notions of linear approach and the space of linear approaches can be also defined using the formalism of germs and stalks coming from algebraic geometry.

Rotational components. The idea of “turning a geodesic segment around a point” is formalised with the notion of *rotational component*, which we define in the following paragraphs. First we need the notion of an *angular sector*. We follow the approach of R. Schwartz in [Sch11].

Consider two infinite rays r_1 and r_2 emanating from the origin in \mathbb{R}^2 . A *sector* in \mathbb{R}^2 is one of the connected components of $\mathbb{R}^2 \setminus \{r_1, r_2\}$ and its angle is defined as the angle between the corresponding rays measured from the interior of the sector. An Euclidean angular sector is a space obtained by gluing together *along rays in the boundary* at most countably many angular sectors in a consecutive pattern. This consecutive pattern might be cyclic, case in which we say that the Euclidean angular sector is *closed*. For example, for every $n \in \mathbb{N} \cup \{\infty\}$, the n to 1 translation covering of \mathbb{R}^2 ramified over the origin is considered to be an Euclidean angular sector. The cone point of an Euclidean angular sector is the equivalence class of the origin(s) under gluing and it is the only point in the Euclidean angular sector that does not have a neighbourhood that is locally isometric to the plane. The *angle* of an Euclidean angular sector is the sum of the angles of the sectors forming it (which might be divergent).

Exercise 2.3.13

Prove that a *closed* Euclidean angular sector defines a translation surface if and only if its total angle is an integer multiple of 2π . Is this translation surface of finite type?

Definition 2.3.14 (Angular sector). An angular sector A in a translation surface M is a subset that is isometric to a *punctured* neighbourhood of the cone point of an Euclidean angular sector.

Remark 2.3.15. It is important to stress that our definition requires the angular sector to be a subset of a translation surface. If $A \subset M$ is an angular sector of total angle $\theta \in (0, \infty]$, the point $z \in \widehat{M}$ in the metric

completion corresponding to the cone point of the Euclidean angular sector in question is well-defined. For this reason we denote sometimes $A = A(\theta, z)$ and call z the cone point of the angular sector. It is important to remark that the cone point of an angular sector *is not necessarily a cone angle singularity*, for it might be a wild singularity.

Definition 2.3.16 (Rotational component). Let $[\gamma_1]$ and $[\gamma_2]$ be two linear approaches in $T_z^1\widehat{M}$. We say that $[\gamma_1]$ and $[\gamma_2]$ belong to the same rotational component if and only if there exist representatives $\gamma_i : (0, \varepsilon_i) \rightarrow M^0$, $i = 1, 2$ and an angular sector $A(\theta, z)$ containing the traces of γ_1 and γ_2 to M^0 . We denote by $\text{Rot}[\gamma]$ the set of all linear approaches belonging to the same rotational component as $[\gamma]$ and we call it *the rotational component of $T_z^1\widehat{M}$ containing $[\gamma]$* . The *total angle* of $\text{Rot}[\gamma]$ is the supremum of all θ for which there exist $[\gamma_1], [\gamma_2] \in \text{Rot}[\gamma]$ with representatives γ_1, γ_2 whose traces are contained in an angular sector $A(\theta, z)$. By convention, the total angle of a singleton is zero.

Loosely speaking, if two linear approaches $[\gamma_1]$ and $[\gamma_2]$ in $T_z^1\widehat{M}$ are in the same rotational component and correspond to rays r_1 and r_2 in some Euclidean angular sector isometric to $A(\theta, z)$ then one can “rotate $[\gamma_1]$ into $[\gamma_2]$ ”, by considering an angular displacement that takes r_1 into r_2 . This intuition can be formalized using the function dir defined in exercise 2.3.7. Indeed, if we fix $[\gamma_0]$ in some rotational component $\text{Rot}[\gamma]$ and $\text{dir}[\gamma_0] = \theta_0 \in \mathbb{R}/2\pi\mathbb{Z}$, then the choice of a point $\tilde{\theta}_0 \in \mathbb{R}$ in the fiber over θ_0 of the universal covering $\mathbb{R} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ defines an injective lift

$$\widetilde{\text{dir}} : \text{Rot}[\gamma] \rightarrow \mathbb{R} \quad (2.7)$$

In this context if $\tilde{\theta}_1$ and $\tilde{\theta}_2$ are the images of $[\gamma_1]$ and $[\gamma_2]$ by the map $\widetilde{\text{dir}}$, then the angular displacement from one into the other is achieved by a translation in \mathbb{R} sending $\tilde{\theta}_1$ to $\tilde{\theta}_2$.

Given that the map $\widetilde{\text{dir}}$ defined above is injective, one can endow any rotational component with the standard topology of the real line. Moreover, using this map it is not difficult to prove that every rotational component admits a (G, X) -structure, where $X = \mathbb{R}$ and $G = 2\pi\mathbb{Z}$ acts by translations, possibly with non-empty boundary, see [BV13]. However this (G, X) -structure is somehow imposed artificially as every rotational component carries a natural subspace topology coming from the limit topology of $T\widehat{M}$. As the following example shows, these topologies in general do not agree. As a matter of fact even in simple cases rotational components with their subspace topology are not metrizable spaces.

Example 2.3.17. The following description relies on Figure 2.5 and was first described by Bowman and Randecker in [Bow12] and [Ran16] respectively. Let $W = \{w_n\}_{n \geq 0}$ be a strictly monotonic decreasing sequence of positive real numbers that converges to 0. For each $n \in \mathbb{N}$ consider the rectangle R_n of height 1 and width w_n . For each $n \geq 1$ glue the lower edge of R_{n+1} to the left part of the upper side of R_n and the vertical sides of R_n together. The result looks like the pile of rectangles shown in the figure. For each n there is still a segment D_{n+1} to be identified: in order to do this we subdivide the lower edge of R_0 into segments of size $w_{n+1} - w_n$ whose extremities are labeled by A_{n+1}, A_n , $n \geq 1$. We assign the label A_0 to the left corner of R_0 . It is not difficult to see that the result of these gluings is a wild translation surface M having only one singularity z . Moreover, $T_z^1\widehat{M}$ is formed by just one rotational component that can be (set theoretically) identified using the map (2.7) with $(0, \infty)$. The convention in this identification is that image of the linear approach in the vertical direction at A_0 corresponds to π . Using this correspondence we can write $T_z^1\widehat{M} = \{[\gamma_s] \mid s \in (0, \infty)\}$.

The claim is that $T_z^1\widehat{M}$ with the topology induced from $T^1\widehat{M}$ is not a regular space. That is, there exist a closed subset F and a point $[\gamma]$ in $T_z^1\widehat{M}$ which cannot be separated by neighbourhoods. We define first the closed subset. For every fixed $\varepsilon > 0$ let $\rho_\varepsilon : \mathcal{L}^\varepsilon(M) \rightarrow T^1\widehat{M}$ be the natural projection. By definition $\rho_\varepsilon(\mathcal{L}^\varepsilon)$ is an open subset. For every fixed $\varepsilon' > 0$ let

$$F_{\varepsilon'} := T_z^1\widehat{M} \setminus \bigcup_{\varepsilon > \varepsilon'} \rho_\varepsilon(\mathcal{L}^\varepsilon)$$

The set $F_{\varepsilon'}$ is closed and is formed by all linear approaches $[\gamma]$ to z for which the length of any representative is bounded above by ε' .

Exercise 2.3.18

Show that if $\varepsilon' > 0$ is small enough, every neighbourhood of $[\gamma_\pi]$ and $F_{\varepsilon'}$ intersect. Conclude that $T_z^1\widehat{M}$ is not a regular space and hence not metrizable. *Hint:* show first that for small $\varepsilon' > 0$ the set $F_{\varepsilon'}$ is formed by horizontal saddle connections contained in the lower side of the rectangle R_0 and

whose length tends to 0.




Figure 2.5: The stack of boxes surface.

In what follows we illustrate with a series of examples how simple or complicated the spaces $T_z^1 \widehat{M}$ (and their rotational components) can be when z is in the singular locus $\text{Sing}(M)$. Except for the double parabola and baker’s surface, all examples are due to A. Randecker. See [Ran16] for details.

Example 2.3.19 (Double parabola). As we see in the following lines, rotational components can be singletons. Let $\pm I_n$ be a family of segments in the xy -plane whose endpoints are $(\pm 2^n, 2^{2n})$ and $(\pm 2^{n+1}, 2^{2(n+1)})$, $n \in \mathbb{Z}$. Let $\pm J_n$ be the family of segments whose endpoints are $(\pm 2^n, -2^{2n})$ and $(\pm 2^{n+1}, -2^{2(n+1)})$, $n \in \mathbb{Z}$. Let P_- be closure of the connected component of

$$\mathbb{R}^2 \setminus \{ \{ \pm I_n \}_{n \in \mathbb{Z}} \cup (0, 0) \cup \{ \pm J_n \}_{n \in \mathbb{Z}} \} \tag{2.8}$$

containing the negative x -axis. Analogously, let P_+ be the closure of the connected component of (2.8) containing the positive x -axis. By construction $\partial P_- = \{ -I_n \}_{n \in \mathbb{Z}} \cup (0, 0) \cup \{ -J_n \}_{n \in \mathbb{Z}}$ and $\partial P_+ = \{ I_n \}_{n \in \mathbb{Z}} \cup (0, 0) \cup \{ J_n \}_{n \in \mathbb{Z}}$. Remove all vertices (and the origin) from P_- and P_+ and identify this two disjoint domains along parallel sides of the same length using translations. This produces a translation surface M that we call the *double parabola*.

Exercise 2.3.20

Let M be the double parabola defined above.

1. Show that $\text{Sing}(\widehat{M})$ is just one wild singularity z_∞ . *Hint:* $\lim_{n \rightarrow -\infty} |\pm I_n| = \lim_{n \rightarrow -\infty} |\pm J_n| = 0$ and these segments accumulate to the origin.
2. Determine all rotational components in $T_{z_\infty}^1 \widehat{M}$. In particular, show that each rotational component defined by $\pm \gamma(t) = (\pm t, 0)$, $t \in (0, 1)$ consists of only one point.

Example 2.3.21. There exist finite-area translation surfaces with only one wild singularity z_∞ and $T_{z_\infty}^1 \widehat{M}$ is formed by two rotational parametrized by \mathbb{R} . Indeed, consider the infinite polygon P depicted in figure 2.6: its vertices are $\{(n, \pm 2^{-|n|}) \mid n \in \mathbb{Z}\}$ and its interior contains the real axis. The sides of P can be grouped into pairs of parallel sides; these are labeled in the figure with the same labels. Let P^0 be P deprived of its vertices and M the translation surface obtained by gluing parallel sides of P using translations. The surface M is called the *exponential surface*. By construction the metric completion \widehat{M} is obtained by adding to M all vertices of P , and by the way we glued the sides of P to obtain M , these vertices merge *a priori* to two points $z_1, z_2 \in \widehat{S}$. However, since every saddle connection in the vertical direction joins z_1 to z_2 and the length of these accumulates to zero we conclude that $z_1 = z_2$ is the only singularity in the metric completion \widehat{M} . Let us denote this singularity by z_∞ . We claim that z_∞ is a wild singularity. Indeed, given that the total area of S is finite, z_∞ cannot be an infinite angle singularity and it cannot be a finite angle singularity since the number of saddle connections emerging from it is infinite. The space of linear approaches $T_{z_\infty}^1 \widehat{M}$ is formed by two rotational components which are isometric to \mathbb{R} , as depicted in figure 2.7.




Figure 2.6: The exponential surface.




Figure 2.7: The space of linear approaches to the wild singularity x_∞ of the exponential surface.

Example 2.3.22. As we saw in 1.2.2, baker’s surface B_α has only one wild singularity $z_\infty \in \widehat{B}_\alpha$. Champanara’s remark “[g]eometrically, the surface spirals infinitely many times around this point” can now be rephrased in the terminology we just introduced as *the space of linear approaches $T_{z_\infty}^1 \widehat{M}$ contains an unbounded rotational component*. As a matter of fact we now have the tools to be completely precise. For the sake of clarity we set $\alpha = \frac{1}{2}$ and we refer henceforth to figure 2.2 in the preceding section. Suppose that the intersection of the diagonals in the unit square defining baker’s surface is the origin and the corners **a** and **c** have coordinates $(-\frac{1}{2}, \frac{1}{2})$ and $(\frac{1}{2}, -\frac{1}{2})$ respectively. In this coordinates the geodesic segments $\gamma_1(t) := (1-t)(-\frac{1}{2}, \frac{1}{2})$ and $\gamma_2(t) := -\gamma_1(t)$, $t \in (0, \varepsilon)$, are linear approaches to z_∞ that define two bi-infinite rotational components $[\gamma_1]$ and $[\gamma_2]$, each isometric to \mathbb{R} . These rotational components are drawn in red and blue in figure 1.11 in section 1.2.2. On the other hand, $\eta_1(t) = (1-t)(\frac{1}{2}, \frac{1}{2})$ and $\eta_2(t) = -\eta_1(t)$, define two bounded rotational components whose total angle is $\pi/4$. The boundary of these rotational components is formed by the horizontal and vertical saddle connections labeled A_i and B_i , $i \in \mathbb{N}$ in the same figure. In the next section we see how the action of the Veech group on $T_{z_\infty}^1 \widehat{M}$ can be used to describe all rotational components, see lemma 3.5.3 in Section 3.5.

Example 2.3.23. One of the main issues when studying the translation flows F_θ^t on M is to determine whether they are defined for all directions on sets of full measure. The only obstruction for these flows to be defined for all times are points in $\text{Sing}(M)$. Intuition tells us that if $\text{Sing}(M)$ is a singleton, then the translation flow should be defined for all directions on sets of full measure. In what follows we present an example that shows that this intuition is wrong: there exist a finite area translation surface M of infinite type with only two singularities such that the *only* direction on which the translation flow F_θ^t is defined on a set of full measure is the horizontal direction. In other words, for infinite-type translation surfaces of finite area, having a finite set of singularities does not imply that the translation flow is defined in a set of full measure for almost every direction.

The example we present is known as the *icicled surface*. The construction of this example starts with the closed rectangle $R = [0, 1] \times [0, 2]$. Remove from R the top and bottom sides $[0, 1] \times \{2\}$, $[0, 1] \times \{0\}$ and identify vertical sides using a translation. For every $n \geq 1$ and odd $i \in \{1, \dots, 2^n - 1\}$ consider the vertical segments $I_{i,n}^{top}$, $I_{i,n}^{bot}$ inside R of length 2^{-n} with one extremity at $(\frac{i}{2^n}, 2)$ and $(\frac{i}{2^n}, 0)$ respectively.

These are the *icicles* and are depicted in figure 2.8. We now define gluings on icicles with one extremity in $[0, \frac{1}{2}] \times \{2\}$:

1. First subdivide each side of the icicle $I_{1,1}^{top}$ at $(\frac{1}{2}, 2)$ following the geometric progression $\frac{1}{2^k}$, $k > 1$ from bottom to top. This creates an infinite subdivision of the icicle $I_{1,1}^{top}$ into segments of length $\frac{1}{2^k}$.
2. Glue the left side of the lower half of the icicle $I_{1,1}^{top}$ to the right side of the icicle $I_{1,2}^{top}$ (which has one extremity at $(\frac{1}{4}, 1)$).
3. For every $n > 2$ and ever odd $i \in \{3, \dots, 2^{n-1} - 1\}$ such that $I_{i,n}^{top}$ has one extremity in $[0, \frac{1}{2}] \times \{1\}$, glue the left side of $I_{i,n}^{top}$ to the right side of $I_{i-2,n}^{top}$.
4. For every $n > 2$, the left side of the icicle at $(\frac{1}{2^{n-1}}, 2)$ is cut into two segments of the same length. The lower segment obtained by this subdivision is glued to the right side of the icicle at $(\frac{2^{n-1}-1}{2^n}, 2)$ and the upper segment to the right side of the only segment at $(\frac{1}{2}, 2)$ that has the same length.

The gluings described above are sketched in figure 2.9. The gluings on the icicles with one extremity on $[\frac{1}{2}, 1] \times \{2\}$ are defined in an analogous way, the only difference is that we need to change *left* for *right* in all instructions above. The gluings for icicles with one extremity on $[0, 1] \times \{0\}$ are just a mirror image (w.r.t. the real axis) of the gluings defined above. The result of these gluings is the *icicled surface*.




Figure 2.8: The icicles in the icicled surface




Figure 2.9: Gluings of the icicles to produce the icicled surface.

Figure 2.10: The icicled surface.

Exercise 2.3.24

Prove that the icicled surface M has only two wild singularities $\{\sigma_{top}, \sigma_{bot}\}$. *Hint:* by the way the gluings were defined, all the tips of the icicles $I_{i,n}^{top}$ are identified into a point that we denote by σ_{top} . Show that every non-dyadic point in $[0, 1] \times \{2\}$, *i.e.* which is not an extremity of an icicle, is the limit of a sequence of tips of icicles. In other words, every non-dyadic point in $[0, 1] \times \{2\}$ is identified with σ_{top} in the metric completion \widehat{M} . The same argument works for dyadic points in $[0, 1] \times \{0\}$

Exercise 2.3.25

Prove that the space of linear approaches $T_{\sigma_{top}}^1 \widehat{M}$ contains *uncountably* many rotational components

which are singletons, and at least two rotational components of total angle $\frac{\pi}{4}$. Does it contain unbounded rotational components (*i.e.* of infinite total angle)?

Remark that for every flat point z in the iced surface the geodesic flow $F_\theta^t(x)$ is not defined in the future and in the past whenever $\theta \neq 0$. Indeed, by the way the gluings defining the iced surface were performed, every such trajectory reaches the top or the bottom of the rectangle R in finite time. On the other hand, for $\theta = 0$ the geodesic flow F_θ^t decomposes the iced surface into an infinite family of cylinders and saddle connections.