

MODULI SPACES OF FLAT TORI WITH PRESCRIBED HOLONOMY

SELIM GHAZOUANI · LUC PIRIO



Abstract. We generalise to the genus one case several results of Thurston concerning moduli spaces of flat Euclidean structures with conical singularities on the two dimensional sphere. More precisely, we study moduli spaces of flat tori with n cone points and a prescribed holonomy ρ . In his paper ‘*Flat Surfaces*’ Veech has established that under some assumptions on the cone angles, such a moduli space $\mathcal{F}_{[\rho]} \subset \mathcal{M}_{1,n}$ carries a natural geometric structure modeled on the complex hyperbolic space $\mathbb{C}\mathbb{H}^{n-1}$ which is not metrically complete. Using surgeries for flat surfaces, we prove that the metric completion $\overline{\mathcal{F}_{[\rho]}}$ is obtained by adjoining to $\mathcal{F}_{[\rho]}$ certain strata that are themselves moduli spaces of flat surfaces of genus 0 or 1, obtained as degenerations of the flat tori whose moduli space is $\mathcal{F}_{[\rho]}$. We show that the $\mathbb{C}\mathbb{H}^{n-1}$ -structure of $\mathcal{F}_{[\rho]}$ extends to a complex hyperbolic cone-manifold structure of finite volume on $\overline{\mathcal{F}_{[\rho]}}$ and we compute the cone angles associated to the different strata of codimension 1. Finally, we address the question of whether or not the holonomy of Veech’s $\mathbb{C}\mathbb{H}^{n-1}$ -structure on $\mathcal{F}_{[\rho]}$ has a discrete image in $\text{Aut}(\mathbb{C}\mathbb{H}^{n-1}) = \text{PU}(1, n - 1)$. We outline a general strategy to find moduli spaces $\mathcal{F}_{[\rho]}$ whose $\mathbb{C}\mathbb{H}^{n-1}$ -holonomy gives rise to lattices in $\text{PU}(1, n - 1)$ and eventually we give a finite list of $\mathcal{F}_{[\rho]}$ ’s whose holonomy is a complex hyperbolic arithmetic lattice.

1 Introduction

For any non-negative integers g and n such that $2g - 2 + n > 0$, we denote by $\mathcal{M}_{g,n}$ the moduli space of genus g Riemann surfaces with n marked points, viewed as a complex orbifold (see [ACG11, Chap. XII] for instance).

★

In their paper [DM86] on the monodromy of Appell–Lauricella hypergeometric functions, Deligne and Mostow bring to light complex hyperbolic structures on $\mathcal{M}_{0,n}$ for $n \geq 4$, parametrised by a n -tuple $\mu = (\mu_1, \dots, \mu_n) \in]0, 1[^n$ such that $\sum_{i=1}^n \mu_i = 2$. They prove that if μ verifies the arithmetic criterion

$$\text{(INT)} \quad \forall i, j \text{ with } i \neq j : \mu_i + \mu_j < 1 \implies (1 - \mu_i - \mu_j)^{-1} \in \mathbb{Z},$$

then the holonomy of the associated complex hyperbolic structure is a lattice in the automorphism group

$$\text{PU}(1, n - 3) = \text{Aut}(\mathbb{C}\mathbb{H}^{n-3})$$

of the $(n - 3)$ -dimensional complex hyperbolic space $\mathbb{C}\mathbb{H}^{n-3}$. The above criterion can be refined to the following (see [Mos86]):

$$(\Sigma\text{INT}) \quad \forall i, j \text{ with } i \neq j : \mu_i + \mu_j < 1 \implies (1 - \mu_i - \mu_j)^{-1} \in \begin{cases} \mathbb{Z} & \text{if } \mu_i \neq \mu_j, \\ \frac{1}{2}\mathbb{Z} & \text{if } \mu_i = \mu_j. \end{cases}$$

In [Thu98], Thurston gives a geometric interpretation of these complex hyperbolic structures in terms of flat metrics with cone type singularities on the sphere S^2 . Define $\theta = (\theta_1, \dots, \theta_n) \in]0, 2\pi[^n$ by $\theta_i = 2\pi(1 - \mu_i)$ for $i = 1, \dots, n$. One can think of $\mathcal{M}_{0,n}$ as the set of flat metrics with n cone points of respective angles $\theta_1, \dots, \theta_n$ on the sphere with area 1 (up to isometry), which we denote by $\mathcal{M}_{0,\theta}$. Parametrising such flat structures naturally endows $\mathcal{M}_{0,\theta}$ with a complex hyperbolic structure (see [Thu98], [Sch] or [Par06]) which coincides with the one considered in [DM86]. Thurston describes the metric completion of $\mathcal{M}_{0,\theta}$ in terms of degenerations of flat spheres and recovers that the criterion (ΣINT) is essentially equivalent¹ to the metric completion of $\mathcal{M}_{0,\theta}$ being an orbifold and therefore a lattice quotient of the complex hyperbolic space $\mathbb{C}\mathbb{H}^{n-3}$.

In [Vee93], Veech extends to compact (oriented) surfaces of arbitrary genus several basic results of Thurston’s approach. The starting point is a theorem of Troyanov [Tro86] asserting that, given $g \geq 0$ and $n > 0$ such that $2g - 2 + n > 0$, if $\theta = (\theta_i)_{i=1}^n \in]0, \infty[^n$ satisfies the following *discrete Gauß-Bonnet formula*

$$\sum_{i=1}^n (2\pi - \theta_i) = 2\pi(2 - 2g), \tag{1}$$

then, given a genus g closed oriented surface N_g , a conformal structure on it and n distinct points p_1, \dots, p_n on N_g , there exists a unique flat structure of area 1 on N_g which is compatible with the given conformal structure, singular exactly at p_1, \dots, p_n and such that it is locally isometric at p_i to a Euclidean cone of angle θ_i , this for every $i = 1, \dots, n$.

Troyanov’s theorem gives a natural isomorphism between $\mathcal{M}_{g,n}$ and the set, denoted by $\mathcal{M}_{g,\theta}$, of isomorphism classes of flat structures on N_g with n cone points of angle data θ . The naive hope that a complex hyperbolic structure would arise when parametrising such a moduli space is doomed to failure. Such a fact actually happens in the genus 0 case because in that case prescribing the cone angles is equivalent to prescribing the parallel transport along any closed curve on the punctured surface. But this does not hold for a n -punctured surface $N_{g,n} = N_g \setminus \{p_i\}$ of higher genus.

In [Vee93], Veech shows that the level sets of the (locally well defined) linear holonomy map

$$\mathcal{M}_{g,\theta} \longrightarrow H^1(N_{g,n}, \mathbb{U})$$

¹ ‘Essentially equivalent’ means ‘up to some particular cases’ which all have been classified (in [Mos88]). In particular, when $n \geq 5$, there is only a finite number of such particular cases.

form a real analytic foliation of $\mathcal{M}_{g,\theta}$ whose leaves are holomorphically embedded complex manifolds of dimension $2g - 3 + n$. Actually, this map is not well defined at the orbifold points of $\mathcal{M}_{g,\theta}$. To bypass this difficulty, one has to work, as Veech did with great care, not on this moduli space but on its orbifold universal covering. When $g = 0$, this foliation is trivial and has only one leaf, which is the whole moduli space $\mathcal{M}_{0,\theta}$, on which one can put a natural geometric structure modeled on a homogeneous space. A geometric way to do this is as follows: given a flat sphere with n prescribed conical singularities, one can develop it into the Euclidean plane and get a $(2n - 2)$ -gon from which the original flat sphere can be reconstructed. The conical angles θ_i being prescribed, the polygons obtained this way depend only on $n - 2$ complex parameters and the area form is a non-degenerate Hermitian form in these parameters if none of the θ_i 's is an integer multiple of 2π . This method, which was first introduced by Thurston in [Thu98] in the genus 0 case, extends very naturally to the leaves of Veech's foliation, whatever the genus is: one can parametrise locally such a leaf by means of Euclidean $(4g + 2n - 2)$ -gons which depend only on $2g + n - 2$ complex parameters. We call *linear parametrisation* such a parametrisation.

Given two integers $p, q \in \mathbb{N}$, let $U(p, q)$ be the group of linear automorphisms of \mathbb{C}^{p+q} which leave invariant the standard Hermitian form $h_{p,q}$ of signature (p, q) . The projectivisation of the set of $z \in \mathbb{C}^{p+q}$ such that $h_{p,q}(z) > 0$ is known as the (*indefinite when $p > 1$*) *complex hyperbolic space of type (p, q)* and will be denoted by $\mathbb{C}\mathbb{H}_p^{p+q-1}$. It is homogeneous under $PU(p, q)$, cf. [Wol11, §12.2].

Let $\theta \in (\mathbb{R}_+^* \setminus 2\pi\mathbb{Z})^n$ be such that the Gauß-Bonnet relation (1) holds true.

Theorem ([Vee93]). *There exists $(p_\theta, q_\theta) \in \mathbb{N}^2$ with $p_\theta + q_\theta = 2g + n - 2$ such that the natural linear parametrisations of the leaves of Veech's foliation together with their area form endow them with a $(\mathbb{C}\mathbb{H}_{p_\theta}^{2g+n-3}, PU(p_\theta, q_\theta))$ -structure.*

Moreover, in [Vee93, §14], Veech performs a lengthy explicit calculation leading to the conclusion that the geometric structure on the leaves of the preceding theorem is complex hyperbolic (*i.e.* $p_\theta = 1$) in exactly two cases:

- (i) $g = 0$ and all the conical angles θ_i are in $]0, 2\pi[$; or
- (ii) $g = 1$ and all the angles θ_i are in $]0, 2\pi[$ except one which lies in $]2\pi, 4\pi[$.

As said above, the former case was treated in [Thu98] (as well as in [DM86] but with the approach involving hypergeometric functions). In this paper we investigate the latter case.

Let $\theta = (\theta_i)_{i=1}^n$ satisfying (1) for $g = 1$ and $n > 1$ and such that condition (ii) above holds true. For any linear holonomy ρ , we denote by $\mathcal{F}_{[\rho]}$ the leaf of Veech's foliation on $\mathcal{M}_{1,\theta}$ that corresponds to (the orbit $[\rho]$ through the action of the pure mapping class group of) ρ . The analytic and very explicit description of Veech's foliation carried out in the twin paper [GP] shows that when ρ is rational (meaning that the subgroup $\text{Im}(\rho) \subset \mathbb{U}$ is finite), $\mathcal{F}_{[\rho]}$ is an algebraic suborbifold of $\mathcal{M}_{1,n}$.

In this paper we give an extrinsic geometric description of the metric completion of such a leaf $\mathcal{F}_{[\rho]}$ for the complex hyperbolic structure given by Veech's Theorem

above. The main theorem of the paper is a generalisation of a result of Thurston in [Thu98].

Theorem 1. *Let $\rho \in H^1(N_{1,n}, \mathbb{U})$ be a rational linear holonomy data.*

- (1) *The metric completion of $\mathcal{F}_{[\rho]}$ has a stratified analytic structure whose strata are finite unramified covers of lower dimensional rational leaves of Veech foliations on $\mathcal{M}_{g',n'}$ with $g' = 0$ and $n' \leq n + 1$ or with $g' = 1$ and $n' \leq n - 1$; there is a finite number of such strata.*
- (2) *This metric completion, denoted by $\overline{\mathcal{F}_{[\rho]}}$ hereafter, is a complex hyperbolic cone manifold of dimension $n - 1$, whose volume is finite.*
- (3) *The cone angles around strata of complex codimension 1 of $\overline{\mathcal{F}_{[\rho]}}$ can be computed using appropriate surgeries.*

The genus 0 case invites us to wonder if some of these leaves $\overline{\mathcal{F}_{[\rho]}}$ are lattice quotients of $\mathbb{C}\mathbb{H}^{n-1}$. Unfortunately, the computation of the cone angles around codimension 1 strata (point (3) of the previous theorem) shows that as soon as $n > 2$, the cone angle around a certain stratum of codimension 1 (formed by collisions involving the only cone point whose angle is larger than 2π) is bigger than 2π . This prevents any leaf $\overline{\mathcal{F}_{[\rho]}}$ from being a lattice quotient provided that $n \geq 3$.

Nevertheless, it does not exclude the possibility that the holonomy of the complex hyperbolic structure is a lattice in $\mathrm{PU}(1, n - 1)$, as both Mostow and Sauter showed that it can happen when $g = 0$ (see [Mos88, Sau90]). As a nice corollary of Theorem 1, we obtain that the holonomy of a finite number of moduli spaces $\mathcal{F}_{[\rho]}$ is an arithmetic lattice. More precisely:

Corollary. *Let $\rho \in H^1(N_{1,n}, \mathbb{U})$ be such that $\mathbb{Z}[\mathrm{Im}(\rho)] \subset \mathbb{C}$ is discrete. Then the image of the complex hyperbolic holonomy of (each connected component of) $\mathcal{F}_{[\rho]}$ is an arithmetic lattice in $\mathrm{PU}(1, n - 1)$.*

This leads us to ask the question of determining all ρ such that the complex hyperbolic holonomy of $\mathcal{F}_{[\rho]}$ is a lattice. The moduli spaces $\mathcal{F}_{[\rho]}$ are not always connected; consequently it is more relevant to ask the aforementioned question for the connected components of $\mathcal{F}_{[\rho]}$. Of course, one must first determine these components, which already seems interesting and not completely trivial.

We give in Section 11 some necessary conditions for a component of $\mathcal{F}_{[\rho]}$'s holonomy to be a lattice. These conditions should reduce the problem to the study of a finite number of candidates.

Finally, we would like to draw attention on a possible interpretation of our work. If $\mathrm{Im}(\rho) = \langle \exp(2i\pi/n) \rangle \subset \mathbb{U}$, the leaf $\mathcal{F}_{[\rho]}$ can be seen as a stratum of the space of meromorphic differential forms of order n on elliptic curves.

1.1 Organisation of the Paper. Section 1 is the present Introduction.

Section 2 and Section 3 are dedicated to introducing the central objects of the article: flat surfaces ρ and Veech isoholonomic foliations on $\mathcal{M}_{g,n}$ respectively.

Building on [Sch, Thu98, Vee93], we introduce in **Section 4** natural parametrizations of the leaves of Veech's foliations that will be used in the sequel.

Section 5 is devoted to proving technical lemmas on the geometry of flat surfaces that are crucial for describing the metric completion of $\mathcal{F}_{[\rho]}$. According to us, some of them, such as Lemma 5.7, are missing in [Thu98] and could help to complete some proofs in the genus 0 case.

We describe in **Section 6** several surgeries on flat surfaces which are the major tools of the paper. They allow us to reinterpret some results of [Thu98] and to formally understand the possible ways flat tori can geometrically degenerate. This leads to a definition of '*geometric convergence*' for flat surfaces distinguishing limits by taking into account not only the isometry class of the limit metric space but also the way to degenerate to it in $\mathcal{F}_{[\rho]}$. This definition coincides with the one of convergence for the complex hyperbolic metric but is susceptible to be generalised to cases when Veech's $\mathbb{C}\mathbb{H}_p^{p+q-1}$ -structure of $\mathcal{F}_{[\rho]}$ is not Riemannian. Finally, using these surgeries and a simple inductive process, we compute by a geometric argument the signature of Veech's area form, recovering Veech's result.

Sections from 7 to 9 are devoted to analysing the geometric structure of $\mathcal{F}_{[\rho]}$. In **Section 7**, we describe the metric completion of $\mathcal{F}_{[\rho]}$ in terms of the surgeries introduced in Section 6, while in **Section 8**, we prove that the complex hyperbolic volume of $\mathcal{F}_{[\rho]}$ is finite by performing an explicit calculation using special coordinates. We finally prove in **Section 9** that the metric completion of $\mathcal{F}_{[\rho]}$ has the structure of a complex hyperbolic cone manifold, which is a refinement of the stratified structure brought to light in Section 7.

After describing a general algorithm to determine the strata appearing in the metric completion of $\mathcal{F}_{[\rho]}$, we analyse in **Section 10** the specific case of tori with two cone points. In particular, we show how the rather abstract material developed in this article is used to analyse the (one dimensional in that case) complex hyperbolic structure. We prove that the $\mathcal{F}_{[\rho]}$ are hyperbolic surfaces with a finite number of cone points and we compute their angles. This section strongly echoes the article [GP]. In particular, this analysis shows that some of them are lattice quotients of $\mathbb{C}\mathbb{H}^1$.

In **Section 11**, we draw a strategy to answer the following question: in the case of tori with $n \geq 3$ conical points, when is the holonomy of the $(n - 1)$ -dimensional complex hyperbolic structure on $\mathcal{F}_{[\rho]}$ a lattice in $\text{PU}(1, n - 1)$? We give necessary conditions for the answer to be positive and use them to outline a strategy to reduce the question to a finite number of candidates. We also exhibit cases, for $n = 3, 4, 5, 6$, where the holonomy group of $\mathcal{F}_{[\rho]}$ is an arithmetic lattice in $\text{PU}(1, n - 1)$.

The paper ends with two short appendices. In **Appendix A**, after recalling some basic points of complex hyperbolic geometry, we introduce some special coordinates which appear useful in our study. Finally, **Appendix B** is devoted to the notion of cone manifold. We focus in particular on the case of complex hyperbolic cone manifolds.

1.2 Notes and References. We think it could be helpful to the reader to mention the main other mathematical works to which the present paper is linked.

As is more than obvious from the previous lines, this text must be seen as an attempt to generalize some results of Thurston's seminal paper [Thu98] concerning moduli spaces of flat spheres to the case of tori. Even if the term does not appear formally in Thurston's paper, we believe it is fair to say that the crucial geometric tools used by Thurston are '*surgeries*' for objects of this type. This is a standard but powerful technique to study flat surfaces which has been widely used in the more specific realm of (half-)translation surfaces, see for instance [MS91, Section 6], [MZ08] or [EMZ03] among many other papers of this field. It is then not so surprising that surgeries play a central role in the present paper as well.

Thurston's article [Thu98] has been very influential. Among the papers deeply relying on it about the theory of conical flat structures on the Riemann sphere, one can mention [Web93], [Par06], [GLL11], [BP15] and [Pas16] where some particular cases are considered in detail. The recent paper [McM17] deserves to be mentioned as well: in it, the author gives a more detailed treatment of the notion of cone-manifold than in [Thu98] and obtains a nice version of the Gauß-Bonnet theorem for complex hyperbolic cone-manifolds that he eventually uses to compute the volumes of the Picard/Deligne–Mostow/Thurston's moduli spaces.

As is well known, some of the main results of [Thu98] coincide with some results obtained previously by Deligne and Mostow in the celebrated papers [DM86] and [Mos88] (see also their book [DM93]). In the true masterpiece [DM86], they pursue and obtain definitive results that conclude researches on the monodromy groups of Appell–Lauricella hypergeometric functions going back at least to Picard. Their approach is not geometric as in [Thu98] but relies essentially on arguments of analytic and/or cohomological nature.

In addition to [Thu98] and [DM86], the other starting point of our research is the remarkable paper [Vee93] by Veech that concerns moduli spaces of flat surfaces with conical singularities and seems to have been deeply influenced by the two former articles. In it, Veech establishes basic and important results concerning flat surfaces of arbitrary genus. One can say that the methods used by Veech are a mix of the geometric ones of Thurston, and of the analytic ones of Deligne and Mostow.

It seems to us that this paper by Veech has not received the attention it deserved despite the importance of the results obtained therein and the interesting problems it suggests. Some of the reasons for this could be that [Vee93] is quite long and technical. If most of its arguments are basically elementary, the analytic treatment used by Veech as well as some long computations at some points hide at first sight the geometrical beauty of its main results. Note also that the topic discussed in [Vee93] is quite general since the linear holonomies of the flat surfaces considered in it, if unitary, are not just ± 1 . It seems that the researchers interested in this subject, Veech included, have focused on the case of (half-)translations surfaces that

is nowadays very popular and for which a lot of deep results have been obtained during the last twenty years.

We believe that an interesting fact highlighted by our work is that both Thurston's geometric approach and Deligne–Mostow's hypergeometric one can be generalised to the genus 1 case. As said above, this is what is done for Thurston's approach in the present paper. The hypergeometric approach à la Deligne–Mostow is developed in the dizygotic twin paper [GP]. In it, we first prove that, as in the genus 0 case for which this is well-known, Veech's constructions of [Vee93] can be made completely explicit when working with elliptic curves. We then specialise to the case of elliptic curves with two marked points and are able to describe exactly the moduli spaces of such marked tori that are algebraic subvarieties of the moduli space $\mathcal{M}_{1,2}$: these are the modular curves $Y_1(N)$ for $N \geq 2$ and we can describe very precisely the complex hyperbolic structure constructed by Veech which each of them carries.

Readers are encouraged to take a look at [GP] and to compare the methods and the results of the latter to the ones of the present text.

2 Flat Surfaces

We collect in this section some well-known notions and basic results on flat surfaces. For some general references, see [Tro86, Vee93, Tro07].

2.1 Generalities. The Gauß-Bonnet formula ensures that the only compact orientable surface carrying a flat metric is the torus. Nevertheless, relaxing the requirement that the metric is flat everywhere and allowing singular points make it possible to build flat surfaces in every genus.

2.1.1 We first define the kind of singularities that will be allowed for flat surfaces in this paper. For any $\theta > 0$ distinct from 2π , the **Euclidean cone of angle θ** , denoted by C_θ throughout the paper, is the quotient of $\mathbb{R}^+ \times (\mathbb{R}/\theta\mathbb{Z})$ obtained by contracting $\{0\} \times (\mathbb{R}/\theta\mathbb{Z})$ onto a point (called the **apex** of the cone) endowed with the flat metric $dr^2 + r^2 dt^2$ in the standard coordinates $(r, t) \in \mathbb{R}_{>0} \times (\mathbb{R}/\theta\mathbb{Z})$ (see Figure 1).




Figure 1: The Euclidean cone C_θ of angle $\theta \in]0, 2\pi[$ embeds in \mathbb{R}^3 .


For any positive ϵ , we denote by $C_\theta(\epsilon)$ the image of $[0, \epsilon] \times (\mathbb{R}/\theta\mathbb{Z})$ into C_θ and the superscript symbol $*$ will mean that the apex has been removed.

DEFINITION 2.1 ([Tro86]). A **flat surface with conical singularities** is an orientable compact surface endowed with a flat Riemannian metric singular at n points p_1, \dots, p_n , such that any p_i has a neighbourhood isometric to a Euclidean cone.

For the sake of simplicity, we will use **flat surface** throughout the paper instead of *flat surface with conical singularities*. A singular point p is called a **cone point** or a **conical point** and the angle θ_p of the associated Euclidean cone its **cone angle**. The quantity $2\pi - \theta_p$ is called the **curvature at p** .

2.2 Examples. We describe below some classical examples of flat surfaces.

2.2.1 A very intuitive example of a flat structure is given by the surface of a cube embedded in \mathbb{R}^3 . The pull-back of the ambient metric defines a flat metric on the 2-dimensional sphere away from the edges and the vertices. On the edges, away from the vertices, the pulled-back metric can be extended in such a way that it is still flat on each edge (this corresponds to the intuitive operation of bending the faces around an edge). We have defined a flat metric on the 8-punctured sphere. A neighbourhood of each vertex is isometric to a neighbourhood of a Euclidean cone of angle $3\pi/2$.



2.2.2 The case of the cube considered above generalises in a straightforward manner to the boundary of any polyhedron P in the 3-dimensional Euclidean space: the natural flat structures of the polygonal 2-faces of ∂P glue together along the straight edges of ∂P and induce a global flat structure which is regular outside the vertices of P and with conical singularities at these points.

2.2.3 Another way to build flat surfaces consists in gluing isometrically the sides of only one Euclidean polygon. We will see later on in Section 4 that, in some sense to be made precise, every flat surface can be built this way. This approach is quite useful and will be extensively used throughout this paper.

In order to give a concrete example, we consider the case of a hexagon. One can glue its sides in three essentially different ways (see [JV01, p.89]) to build (topologically) a torus as in Figure 2 below which respectively give after gluing the three tori with two marked points of Figure 3.

Now assume that H is a Euclidean hexagon and choose one of the three gluing patterns of Figure 2, say Pattern 2. We choose H such that the sides which are




Figure 2: Gluing patterns for flat tori with two singular points.

glued together have the same length (see Figure 4). The Euclidean metric on the hexagon can be extended to the whole torus, except at the points corresponding to the vertices of the hexagon. These points have a punctured neighbourhood isometric to a Euclidean cone with angle $\alpha + \gamma + \mu$ for the first point and $\beta + \delta + \nu$ for the second one. Since H can be triangulated using four Euclidean triangles, one obtains that $\alpha + \beta + \gamma + \delta + \mu + \nu = 4\pi$.

Rewriting this equality

$$(2\pi - (\alpha + \gamma + \mu)) + (2\pi - (\beta + \delta + \nu)) = 0,$$

one obtains that the sum of the curvatures at the singular points vanishes, which is exactly Gauß-Bonnet formula in this case (see §2.3.2 below).

2.2.4 A popular and very much studied example of flat surfaces is given by the so-called **(half-)translation surfaces**, namely pairs (X, ω) (resp. (X, η)) where X is a compact Riemann surface and ω (resp. η) an abelian (resp. a quadratic) differential on it. The flat metric associated to such a pair is just $|\omega|^2$ (resp. $|\eta|^2$). Note that these objects can be characterised as the flat surfaces whose linear holonomy (cf. Section 2.4 below) is trivial (resp. has values in $\{\pm 1\}$) hence they form a very particular class of flat surfaces.

2.2.5 A **flat cylinder** C is the metric space one gets by gluing two opposite sides of a Euclidean rectangle. It is a flat surface with two totally geodesic boundary components. Its **length** is the length of the sides glued together and its **width** is the length of one of its boundary component. More intrinsically, the length of C is the distance between its two boundary components and its width is its systole (its systole being the length of the shortest essential closed curve).

2.2.6 A **saddle connection** is a totally geodesic path joining two singular points and meeting no singular points in its interior.




Figure 3: The flat tori with 2 cone points corresponding to the gluing patterns of Figure 2.




Figure 4: A flat torus with two cone points built from gluing the sides of the Euclidean hexagon at the top.

2.3 On the geometry of flat surfaces. In the subsections below, we collect some classical material about the geometry of flat surfaces and fix some definitions and notations that will be used in the sequel.

2.3.1 Flat surfaces as length spaces. For a general exposition of the theory of length spaces, we refer to [BH99] or to [Gro99] for the proofs of the results stated in this subsection.

Let N be an arbitrary flat surface with cone type singularities. If $\gamma : [0, 1] \rightarrow N$ is a piecewise C^1 -path, one defines its **length** as

$$L(\gamma) = \int_{[0,1]} |\gamma'(t)| dt .$$

As usual, the **distance between two points** $x, y \in N$ is defined as

$$d(x, y) = \inf_{\gamma} L(\gamma)$$

where γ is taken amongst all the C^1 -paths such that $\gamma(0) = x$ and $\gamma(1) = y$.

The following basic results will be extensively used throughout the article:

- The map $d : N \times N \longrightarrow \mathbb{R}_+$ is a distance on N . Whenever we refer to a distance on the flat surface N , it will be this one.
- For any two points $x, y \in N$, there exists a piecewise geodesic path γ from x to y such that $d(x, y) = L(\gamma)$.
- For any non-trivial free homotopy class c of closed curves on N , there exists a closed, piecewise geodesic path φ whose free homotopy class is c such that $L(\varphi) = \inf_{\{\gamma \mid [\gamma]=c\}} L(\gamma)$.
- There exists a closed, piecewise geodesic path σ whose free homotopy class is non-trivial such that $L(\sigma) = \inf_{\{\gamma \mid [\gamma] \neq 0\}} L(\gamma)$.

From now on, we will use the notation N for a surface, which depending on the context will be understood as endowed with either a topological, a flat or a conformal structure. We will also use the notation $N_{g,n}$ when we will need to specify the genus and/or the number of cone points (resp. marked points) of the flat (resp. conformal) structure. Finally whenever we will refer to the (co)homology groups $H_1(N_{g,n}, G)$ or $H^1(N_{g,n}, G)$ for a given group G , $N_{g,n}$ will stand for the underlying topological surface of genus g with n punctures.

2.3.2 The Gauß-Bonnet formula. A regular Riemannian metric h on a surface Σ enjoys the fact that its curvature function κ_h satisfies

$$\int_{\Sigma} \kappa_h d\mu_h = 2\pi\chi(\Sigma)$$

where μ_h is the measure on Σ induced by h and $\chi(\Sigma)$ is the Euler characteristic of Σ . This relation is called the Gauß-Bonnet formula and can be generalised to the case of flat surfaces with conical singularities. One must think of the associated singular Riemannian metric as a metric whose curvature is concentrated at its singular locus, and therefore think of its curvature function as a linear combination of Dirac masses at the singular points.

If N is a compact orientable flat surface with n singular points of respective cone angles $\theta_1, \dots, \theta_n$, the following **Gauß-Bonnet formula** holds true:

$$\sum_{i=1}^n (2\pi - \theta_i) = 2\pi\chi(N).$$

We refer to [Tro86, §3] or [Vee93, §3] for proofs and more details on this matter.

2.3.3 Exponential maps. Let p be a regular point of N and denote by r_p the distance from p to the set of singularities. For $r > 0$, one denotes by $D(r)$ the Euclidean disk of radius r centered at the origin. We will say that ‘the’ **exponential map i_p at p** is the map (well defined and unique up to rotations)

$$i_p : D(r_p) \longrightarrow N$$

such that $i_p(0) = p$ and which is a local isometry. This map can be extended to the whole Euclidean plane except for a countable union of semi-lines which correspond

to the geodesics starting at p which cannot be extended because they meet a singular point. The proof is elementary and left to the reader.

This definition generalises at a singular point p of N . If θ stands for the cone angle at p , let r_p be the biggest $r > 0$ such that the portion of cone $C_\theta(r)$ (cf. §2.1.1) can be isometrically embedded in N at p . Then one defines the **exponential map at the cone point p** as the corresponding embedding

$$i_p : C_\theta(r_p) \longrightarrow N$$

which is unique, up to the isometries (*i.e.* rotations) of the cone C_θ . This map enjoys the same properties as the exponential map at a regular point.

2.4 Affine and Linear Holonomy of a Flat Surface. If N is a flat surface, the punctured surface $N^* = N \setminus S$ (where S is the set of singular points of N) is endowed with a (non complete) flat metric which is everywhere regular. Another way to phrase this is to say that $N \setminus S$ carries a $(\mathbb{C}, \text{Iso}^+(\mathbb{C}))$ -structure ($\text{Iso}^+(\mathbb{C})$ denotes the group of orientation preserving Euclidean isometries of $\mathbb{C} \simeq \mathbb{R}^2$).

With such a structure comes a **holonomy representation**

$$\text{Hol} : \pi_1(N^*) \longrightarrow \text{Iso}^+(\mathbb{C}), \tag{2}$$

whose class for the action by conjugation of $\text{Iso}^+(\mathbb{C})$, is a geometric invariant of the flat structure. The group $\text{Iso}^+(\mathbb{C})$ being the set of affine transformations of the form $z \mapsto az + b$ with $a \in \mathbb{U}$ and $b \in \mathbb{C}$, it is isomorphic to the semi-direct product $\mathbb{U} \ltimes \mathbb{C}$. The projection onto the first factor \mathbb{U} is a group homomorphism and post-composing hol by this projection produces a new representation

$$\rho : \pi_1(N^*) \longrightarrow \mathbb{U}$$

called the **linear holonomy** of the considered flat surface.

The group \mathbb{U} being commutative, ρ factors through the abelianisation of $\pi_1(N^*)$ namely the first homology group $H_1(N^*, \mathbb{Z})$ of the punctured surface N^* . Let n be the cardinality of S and denote by p_1, \dots, p_n the n cone points of N . If $\theta = (\theta_1, \dots, \theta_n)$ is the associated angle datum (*i.e.* the cone angle at p_k is θ_k for any k) and if δ_k is a simple closed curve turning anticlockwise around p_k , then necessarily $\rho(\delta_k) = e^{i\theta_k}$ for $k = 1, \dots, n$. We denote by $H^1(N^*, \mathbb{U}, \theta)$ the set of \mathbb{Z} -linear forms on $H_1(N^*, \mathbb{Z})$ which maps δ_k onto $e^{i\theta_k}$ for every k :

$$H^1(N^*, \mathbb{U}, \theta) = \left\{ \rho \in \text{Hom}\left(H_1(N^*, \mathbb{Z}), \mathbb{U}\right) \mid \rho(\delta_k) = e^{i\theta_k} \text{ for } k = 1, \dots, n \right\}.$$

In what follows, we will consider ρ as an element of this space. Remark that basically, ρ is nothing else but the parallel transport of the flat Riemannian metric on N^* being considered.

2.5 Isometries. We end this section with a few words about the group $\mathbf{Iso}^+(N)$ of direct isometries of a given compact flat surface N .

First, remark that this group is finite since it embeds into the group of biholomorphisms of the underlying Riemann surface which is known to be finite. Second, the subgroup of $\mathbf{Iso}^+(N)$ made of elements fixing a given cone point must be cyclic, since its elements are completely determined by their differential at the fixed point which is a rotation. It follows easily that the subgroup $\mathbf{PIso}^+(N)$ formed by pure direct isometries of N (here ‘pure’ means that the considered isometries fix pointwise the set of cone points) is necessarily cyclic.

3 Veech’S Isoholonomic Foliations on $\mathcal{M}_{g,n}$

3.1 Moduli Spaces of Flat Surfaces and Troyanov’s Theorem. Let N be a compact oriented surface of genus g with n marked points p_1, \dots, p_n and $\theta = (\theta_i)_{i=1}^n$ a set of angle data satisfying the Gauß-Bonnet relation (1). Since we are only interested in this case, and because making such an assumption will simplify the exposition, we will always assume that

$$\text{none of the angles } \theta_i \text{ is an integer multiple of } 2\pi. \quad (3)$$

We define $\mathcal{E}_{g,\theta}$ as the set of flat structures on N such that the metric is singular at p_i with a cone angle θ_i at this point, up to the action of $\text{Diff}_0^+(N, S)$ where $S = \{p_1, \dots, p_n\}$. One can think of $\mathcal{E}_{g,\theta}$ as the set of flat surfaces of genus g with cone angles θ with a marking of its fundamental group. For more details on this construction and the ones to come, we refer to [Vee93, Theorem 1.13].

Notice that a flat structure defines canonically a conformal structure on N . Away from the singularities, this conformal structure is given by the regular flat structure. At a singular point p of cone angle θ_p , there is an essentially unique local coordinate z centered at p such that the flat metric is $|z^{\alpha_p} dz|^2$ with $\alpha_p = (\theta_p/2\pi) - 1$. By means of z , one extends the conformal structure of the punctured surface through p . Since this can be done for every conical singularity of N , one obtains a well-defined map

$$\mathcal{E}_{g,\theta} \longrightarrow \mathcal{Teich}_{g,n}, \quad (4)$$

where $\mathcal{Teich}_{g,n}$ denotes the usual Teichmüller space of conformal structures on a surface of genus g with n marked points. The remarkable fact is that this map is one-to-one. This is a consequence of Troyanov’s theorem stated below.

Theorem ([Tro86]). *Every conformal structure on N is induced by a flat metric with conical singularities of angle θ_i at p_i for $i = 1, \dots, n$. Moreover, this flat metric is unique up to normalization.*

The proof (given in [Tro86]) essentially consists in solving the PDE that the metric tensor associated to a given conformal structure must satisfy. For more details, we refer to the original article [Tro86] which is very pleasant to read.

Consequently, one has a one-to-one correspondance $\mathcal{E}_{g,\theta} \simeq \mathcal{T}eich_{g,n}$ allowing these two moduli spaces to be identified. In particular, this endows $\mathcal{E}_{g,\theta}$ with the structure of a complex manifold of dimension $3g - 3 + n$.

3.2 Veech's Foliations. Since we are considering marked flat structures, the **linear holonomy map**

$$\text{hol} = \text{hol}_\theta : \mathcal{T}eich_{g,n} \simeq \mathcal{E}_{g,\theta} \longrightarrow \mathbb{H}^1(N_{g,n}, \mathbb{U}) \quad (5)$$

which associates its linear holonomy morphism to a flat structure, is well defined. Clearly, hol maps $\mathcal{T}eich_{g,n}$ into $\mathbb{H}^1(N_{g,n}, \mathbb{U}, \theta)$.

From hypothesis (3), it follows that the trivial character (the one sending any holonomy class onto $1 \in \mathbb{U}$) does not belong to $\mathbb{H}^1(N_{g,n}, \mathbb{U}, \theta)$. This case being excluded, the following theorem holds true:

Theorem ([Vee93]). *The linear holonomy map (5) is an open real-analytic submersion. Moreover, for any $\rho \in \text{Im}(\text{hol})$, the level set $\text{hol}^{-1}(\rho)$ is a complex submanifold of $\mathcal{T}eich_{g,n}$ of complex dimension $2g - 3 + n$.*

This result implies in particular that the level sets $\mathcal{F}_\rho = \text{hol}^{-1}(\rho)$ for $\rho \in \text{Im}(\text{hol})$ form a real-analytic foliation by complex submanifolds of $\mathcal{T}eich_{g,n}$. This foliation will be denoted by $\mathcal{F}(\theta)$ (or just \mathcal{F} for short, when θ has been fixed) and will be called the **Veech foliation of $\mathcal{T}eich_{g,n}$ associated to θ** .

3.3 Invariance by the Pure Mapping Class Group. We now explain how this foliation descends to $\mathcal{M}_{g,n}$. The pure mapping class group $\text{PMCG}_{g,n}$ acts on $\mathcal{T}eich_{g,n}$ preserving Veech's foliation: namely any element $f \in \text{PMCG}_{g,n}$ sends \mathcal{F}_ρ onto $\mathcal{F}_{f^*\rho}$. Hence the foliation $\mathcal{F}(\theta)$ factors through the projection

$$\mathcal{T}eich_{g,n} \longrightarrow \mathcal{M}_{g,n} = \mathcal{T}eich_{g,n}/\text{PMCG}_{g,n} \quad (6)$$

to define a singular foliation on the moduli space $\mathcal{M}_{g,n}$. The latter is denoted by $\mathcal{F}(\theta)$ (or just by \mathcal{F} when θ is fixed) and will also be called **Veech's foliation**. Strictly speaking, since $\text{PMCG}_{g,n}$ acts with fixed points on $\mathcal{T}eich_{g,n}$, one should more rigorously speak of $\mathcal{F}(\theta)$ as an 'orbifold' on $\mathcal{M}_{g,n}$. However, because it will not be the source of real problems, we will ignore this subtlety in the whole paper.

- We will now refer to a specific leaf $\mathcal{F}_{[\rho]}$ where $[\rho]$ is the orbit of an element of $\mathbb{H}^1(N_{g,n}, \mathbb{U}, \theta)$ under the action of $\text{PMCG}_{g,n}$. Note that it is the image of $\mathcal{F}_\rho \subset \mathcal{T}eich_{g,n}$ by the quotient map (6).
- We say that $\mathcal{F}_{[\rho]}$ is a leaf of Veech's foliation. That is not rigorously correct, because usually, in foliation theory, one demands that leaves be connected. It is actually proven in [GP, §4.2.5], through some explicit analytic computations, that $\mathcal{F}_{[\rho]}$ can have several distinct connected components. Nevertheless, we will refer below to the $\mathcal{F}_{[\rho]}$'s as leaves for convenience.
- Since $\text{PMCG}_{g,n}$ acts on $\mathbb{H}^1(N_{g,n}, \mathbb{U}, \theta)$ preserving its symplectic form, the foliation has a transverse symplectic structure of dimension $2g$.

3.4 Geometric Structures on the Leaves. Let p and q be non-negative integers and let

$$h_{p,q} : (z, w) \mapsto \sum_{i=1}^p z_i \bar{w}_i - \sum_{j=p+1}^{p+q} z_j \bar{w}_j$$

be the standard Hermitian form of signature (p, q) on $V = \mathbb{C}^{p+q}$. Let V_+ be the set of elements $z \in \mathbb{C}^{p+q}$ such that $h_{p,q}(z, z) > 0$ and let $\mathbb{C}\mathbb{H}_p^{p+q-1}$ be the image of V_+ in $\mathbb{C}\mathbb{P}^{p+q-1}$. The group of automorphisms of $h_{p,q}$, namely $\text{PU}(p, q)$, acts transitively by biholomorphisms on $\mathbb{C}\mathbb{H}_p^{p+q-1}$, see [Wol11, §12.2]. Note that for $p = 1$, $\mathbb{C}\mathbb{H}_1^q$ is nothing else but the usual complex hyperbolic space $\mathbb{C}\mathbb{H}^q$.

Recall that we are assuming that hypothesis (3) holds true: the angle datum $\theta = (\theta_i)_{i=1}^n \in]0, +\infty[^n$ is supposed to be such that $\theta_i \notin 2\pi\mathbb{Z}$ for any $i = 1, \dots, n$.

Theorem ([Vee93]). *There exists a pair of integers $(p, q) = (p_\theta, q_\theta) \in \mathbb{N}^2$ with $p + q = 2g - 2 + n$, such that the leaves of Veech’s foliation $\mathcal{F}(\theta)$ on $\mathcal{T}eich_{g,n}$ are endowed with natural $(\mathbb{C}\mathbb{H}_p^{p+q-1}, \text{PU}(p, q))$ -structures. These geometric structures are invariant under the action of the pure mapping class group hence can be pushed-forward on the leaves of Veech’s foliation $\mathcal{F}(\theta)$ on $\mathcal{M}_{g,n}$.*

In [Vee93, §14], Veech gives an explicit closed formula for the signature (p_θ, q_θ) as a function of θ . We explain briefly where this geometric structure comes from. Consider $\rho \in H^1(N_{g,n}, \mathbb{U}, \theta)$ in the image of hol_θ and consider the associated leaf $\mathcal{F}_\rho \subset \mathcal{T}eich_{g,n}$. Given a flat surface in it, one can consider its full Euclidean holonomy (2). Since its linear part is fixed (and equal to ρ), the meaningful geometric information is contained in the translation part of this full holonomy, which can be viewed as an element of the projectivisation of a certain twisted cohomology group denoted here by $H_\rho^1(N_{g,n}, \mathbb{C})$. One can then construct a **relative period map**

$$\mathcal{F}_\rho \longrightarrow \mathbf{P}\left(H_\rho^1(N_{g,n}, \mathbb{C})\right). \tag{7}$$

Veech (and previously Thurston in [Thu98] for flat surfaces of genus 0) proves that the preceding map is a local biholomorphism (see [Vee93, Theorem 0.6]). Moreover, one can define a non-degenerate Hermitian form h_ρ on $H_\rho^1(N_{g,n}, \mathbb{C})$ (which is actually the area of the corresponding flat surface) such that the relative period map lands in $\mathbf{P}_+(H_\rho^1(N_{g,n}, \mathbb{C}))$, where the latter stands for the set of complex lines in $H_\rho^1(N_{g,n}, \mathbb{C})$ on which h_ρ is positive. Thus the target space of (7) is nothing else but a model of $\mathbb{C}\mathbb{H}_p^{p+q-1}$ and any element f of the $\text{PMCG}_{g,n}$ induces an isomorphism of $(\mathbb{C}\mathbb{H}_p^{p+q-1}, \text{PU}(p, q))$ -structures

$$f : (\mathcal{F}_\rho, h_\rho) \longrightarrow (\mathcal{F}_{f^*\rho}, h_{f^*\rho})$$

(see [Vee93, Theorem 0.7]; this amounts to saying that changing the marking of a flat surface does not change its area).

Except for very few cases, those $\mathbb{C}\mathbb{H}_p^{p+q-1}$ -structures are known not to be complete. The term ‘complete’ has to be understood here in the sense of geometric structures, cf. [Thu97, §3.5]. For geometric structures modeled on a (possibly indefinite) complex hyperbolic space $\mathbb{C}\mathbb{H}_p^{p+q-1}$, this coincides with the fact of being geodesically complete for the associated Levi–Civita connection (see Proposition 1.2 of [Tho15] for instance). A more geometric description of some of those structures in the following sections will make this fact obvious.

4 Linear Charts on the Leaves of Veech’s Foliation

In this section we present material about local parametrisations of moduli spaces of flat surfaces. Although this material is well known, there is no standard point of view or unified theory of these parametrisations. Depending on the context, parameters obtained from gluing Euclidean polygons, analytic calculations, twisted cohomology or a combination of several of these are better suited to formalize an idea or to simply perform a computation. Nevertheless all these points of view (to be detailed) are essentially the same. References developing various material are [Thu98, Section 3], [GP], [Sch] and [Vee93, Sections 9,10 and 11].

For the remainder of the section g, n and $\rho : H_1(N_{g,n}, \mathbb{Z}) \rightarrow \mathbb{U}$ are fixed. We also suppose that $2g + n - 3 > 0$ in order for \mathcal{F}_ρ to have positive dimension.

4.1 Polygonal Models for Flat Surfaces. We describe here a geometrically intuitive local parametrisation of \mathcal{F}_ρ . Take N a flat surface in \mathcal{F}_ρ such that N can be recovered from gluing isometrically suitable sides of a Euclidean polygon P with $2k$ sides whose vertices project to singular points. Necessarily, $k = 2g - 1 + n$. We identify to \mathbb{C} the Euclidean plane in which P lies in and associate to each side the corresponding complex number $z_i, 1 \leq i \leq 2k$ (with the convention that P is positively oriented relatively to its interior) defined by

$$z_i = \text{end point of the side} - \text{initial point of the side}.$$

Assuming that (the side associated to the complex number) z_i is paired with (the side associated to) z_{k+i} , the $2k$ -tuple (z_1, \dots, z_{2k}) must satisfy the following relations:

$$\sum_{i=1}^{2k} z_i = 0 \quad \text{and} \quad |z_i| = |z_{k+i}| \quad \text{for } i = 1, \dots, k.$$

Note that since we require that z_i be paired with z_{i+k} , the z_j ’s for $j = 1, \dots, 2k$ do not necessarily appear in cyclic order (see Figure 5 for instance). Each complex number $\rho_i = z_i/z_{k+i} \in \mathbb{U}$ is the holonomy of a curve (which is closed in the corresponding flat surface N) joining the middle of z_i to the middle of z_{k+i} and therefore belongs to $\text{Im}(\rho)$. One rewrites the previous equations as

$$\sum_{i=1}^{2k} z_i = 0 \quad \text{and} \quad z_i = \rho_i z_{k+i} \quad \text{for } i = 1, \dots, k. \tag{8}$$




Figure 5: Polygonal model for a genus 2 surface with two cone points.

After eliminating $z_k, z_{k+1}, \dots, z_{2k}$, one sees that $z = (z_1, \dots, z_{k-1}) \in \mathbb{C}^{k-1}$ completely characterises the polygon P and therefore the associated flat surface N .

Any $(k-1)$ -tuple $u = (u_1, \dots, u_{k-1})$ close to z in \mathbb{C}^{k-1} defines a polygon P_u whose sides satisfy the equations above. Performing the associated gluing (meaning that one glues the side associated to u_i to the one associated to $u_{i+k} = \rho_i^{-1}u_i$ for $i = 1, \dots, k-1$) builds another element of \mathcal{F}_ρ .

Let $U \subset \mathbb{C}^{k-1}$ be a small open subset containing z such that all the $2k$ -gons corresponding to elements of U are non-degenerate. One defines a map

$$\varphi : U \longrightarrow \mathcal{F}_\rho$$

by associating to each $u \in U$ the **renormalized flat surface associated to P_u** , that is the one of area one. Notice that φ is not locally injective since $\varphi(\lambda u) = \varphi(u)$ for all $(\lambda, u) \in \mathbb{C}^* \times U$ such that $\lambda u \in U$. That being said, φ induces a map $\psi : V \longrightarrow \mathcal{F}_\rho$ where $V = \mathbf{P}U$ is the image of U in $\mathbf{P}(\mathbb{C}^{k-1})$. It is possible to prove that ψ is a local biholomorphism for the structure inherited as a leaf of a foliation of $\mathcal{T}eich_{g,n}$ as has been done by Veech, see [Vee93, Lemma 10.23]. Nonetheless, we want to adopt an intrinsic point of view on the geometry of \mathcal{F}_ρ and will therefore ignore Veech's results.

We remark that $\psi : V \longrightarrow \mathcal{F}_\rho$ is a local homeomorphism:

- The fact that ψ is one-to-one onto its image follows straightforwardly from the following remark: since we are looking at marked flat structures, any isometry preserving the marking between close surfaces in the parametrisation V must come from an isometry of the polygons themselves being the identity on the boundary of the polygon; and therefore be the identity.
- The fact that ψ is onto is a consequence of the fact that the polygonal model survives small deformations.

We will actually ignore the second point and define a structure of (complex) manifold on \mathcal{F}_ρ using ψ . Two details remain to be settled:

- (1) we have been able to build ψ only if N is built out from gluing sides of a polygon. We now need to extend this construction to the general case;

- (2) we need to prove that if two charts have overlapping images, then the transition maps are biholomorphisms.

The first difficulty can be settled by introducing the notion of *pseudo-polygon*. We follow here [Sch]. A **pseudo-polygon** is a flat metric on a (closed) disk whose boundary is locally isometric to a piecewise geodesic path in \mathbb{C} . By developing a pseudo-polygon, we can also define it as an immersion of the closed disk into the plane whose boundary is piecewise geodesic.

PROPOSITION 4.1. *Every flat surface N can be built out from gluing sides of a pseudo-polygon.*

The proof of this proposition is carefully done in the case $g = 0$ in [Sch], and in the general case in [Vee93]. The crucial point is the existence of a totally geodesic triangulation (see Lemma 6.23 in [Vee93] or the construction of the Delaunay decomposition that we will detail in Section 5) for a given flat surface N). Starting from there, one easily checks that for any graph Γ in the 1-skeleton of such a triangulation such that $N \setminus \Gamma$ is simply connected, then (the metric completion for the length metric of) the latter is a pseudo-polygon.

The main remark at this point is that the parametrisation built when N comes from a polygonal model straightforwardly generalises to the case when N is built out from a pseudo-polygonal model, simply by immersing (using the developing map of the flat structure) such a pseudo-polygon in \mathbb{C} . According to Proposition 4.1, every surface has a pseudo-polygonal model and therefore the maps ψ built this way form an atlas of charts for \mathcal{F}_ρ .

From now on, a local parametrisation (z_1, \dots, z_{k-1}) arising in this way will be referred to as a **polygonal parametrisation**.

4.2 Area form and Linear Parametrisation. Another very important remark at this point is that a polygonal parametrisation comes with a natural Hermitian form which is the signed area of the corresponding flat surface. If U is an open subset of \mathbb{C}^{k-1} on which is defined a polygonal parametrisation φ of \mathcal{F}_ρ , we denote by $A_{\varphi,U}$ the corresponding Hermitian form.

The proof that $A_{\varphi,U}$ is actually a Hermitian form in $z = (z_1, \dots, z_{k-1})$ goes the following way: every immersed pseudo-polygon can be triangulated in such a way that each side is a geodesic path joining two edges. Let T_1, \dots, T_L be the triangles of the triangulation. For any l , the area of T_l is

$$A(T_l) = \frac{1}{2} \operatorname{Im}(z_{T_l} \overline{w_{T_l}})$$

where z_{T_l} and w_{T_l} are the complex numbers associated to two consecutive sides of T_l , oriented in such a way that they form a direct basis of \mathbb{C} (seen as a 2-dimensional real vector space). Both z_{T_l} and w_{T_l} are linear combinations of z_1, \dots, z_{2k} and therefore of z_1, \dots, z_{k-1} thanks to (8). For any l , the area $A(T_l)$ is a Hermitian form in z .

Since the area of the whole surface is given by

$$A_{\varphi,U}(N) = \sum_{l=1}^L A(T_l),$$

it follows that $A_{\varphi,U}$ is indeed a Hermitian form in (z_1, \dots, z_{k-1}) .

The next proposition describes the regularity of the transition maps and settles point (2) of Section 4.1.

PROPOSITION 4.2. *Let (φ_1, U_1) and (φ_2, U_2) be two polygonal parametrisations of \mathcal{F}_ρ such that $W = \varphi_1(U_1) \cap \varphi_2(U_2) \subset \mathcal{F}_\rho$ is non-empty, connected and sufficiently small for the projectivisations $\psi_i : \mathbf{P}U_i \rightarrow \mathcal{F}_\rho$ of the φ_i 's (see above) to induce isomorphisms between $\psi_i^{-1}(W)$ and W , for $i = 1, 2$.*

Then $\psi_2^{-1} \circ \psi_1 : \psi_1^{-1}(W) \rightarrow \psi_2^{-1}(W)$ is the restriction of the projectivisation of a linear map $g \in \mathrm{GL}_{k-1}(\mathbb{C})$ such that $g^ A_{\varphi_1, U_1} = A_{\varphi_2, U_2}$.*

Proof. Let P and Q be two polygonal models for a flat surface N , and immerse P in \mathbb{C} . Let z_1, \dots, z_{2k} be the complex numbers associated to the sides of P . Consider now a side w_i of Q , and develop it in \mathbb{C} starting from an initial copy of P (say P_0) and gluing a copy P_{i+1} of P to a side of P_i every time it is necessary to keep track of w_i . Thus one can express w_i as a linear combination of the complex number associated to the sides of P and find an expression of w_i of the form

$$w_i = \sum_{j=1}^{k-1} \alpha_{i,j} z_j$$

where the $\alpha_{i,j}$ are constants depending only on ρ and the combinatorics of the side w_i relatively to P . Therefore the coordinates (w_1, \dots, w_{k-1}) depend linearly on (z_1, \dots, z_{k-1}) . Swapping the roles of the two charts, one gets that the transition map actually lies in $\mathrm{GL}_{k-1}(\mathbb{C})$. The area only depends on the underlying surface and therefore does not depend on the parametrisation. \square

The proposition above tells us that the polygonal charts endow \mathcal{F}_ρ with a complex projective structure (and with an additional structure coming from the preserved area form, which will be investigated later). The previous analysis invites us to define a more general class of parametrisations:

DEFINITION 4.1. *A local holomorphic parametrisation (z_1, \dots, z_{k-1}) of \mathcal{F}_ρ is called a **linear parametrisation** if it depends linearly on a polygonal parametrisation.*

This class is much more convenient than the class of polygonal parametrisation because it is the larger class of holomorphic charts enjoying the property that the area form is Hermitian in the associated coordinates. We will also see in Section 4.4 that it is possible to build other such linear parametrisations in a natural way which will be extensively used throughout the article.

4.3 Projection onto $\mathcal{F}_{[\rho]}$. We have built in this section projective charts on $\mathcal{F}_\rho \subset \mathcal{Teich}_{g,n}$. If $N \in \mathcal{F}_{[\rho]}$ is a regular point of $\mathcal{F}_{[\rho]}$ *i.e.* if the projection

$$\pi : \mathcal{F}_\rho \longrightarrow \mathcal{F}_{[\rho]}$$

is a local homeomorphism at N , any chart at $\tilde{N} \in \pi^{-1}(N)$ can be pushed forward and gives a chart at N . The fact that N is not regular is equivalent to the fact that $\text{PIso}^+(\tilde{N})$, the group of pure direct isometries of the flat surface \tilde{N} (see §2.5), is non-trivial. In that case any chart at \tilde{N} gives a non-injective local parametrisation of a neighbourhood of N in $\mathcal{F}_{[\rho]}$ whose transformation group is the stabilizer of \tilde{N} in $\text{PMCG}_{1,n}$ which is isomorphic to $\text{PIso}^+(\tilde{N})$.

4.4 Parametrisations Coming from Topological Gluing. Here we describe parametrisations which are generalisations of polygonal parametrisations: we are just going to relax the condition that the sides of the polygon we are gluing be geodesic.

Consider a (topological) triangulation \mathcal{T} of $N_{g,n}$ such that the set of vertices is exactly the set of cone points of N . As explained in [Thu98] in genus 0 (see [Sch] for details) and [Vee93, §10] in arbitrary genus, one can find a graph in the 1-skeleton of \mathcal{T} , such that its complement Q in $N_{g,n}$ is simply connected. Q is a topological disk endowed with a flat metric whose boundary corresponds to consecutive edges of triangulation. Let $F : Q \longrightarrow \mathbb{C}$ be a developing map of the flat metric on Q and let $q_1, \dots, q_{2k}, q_{2k+1} = q_1$ be the vertices of the boundary ∂Q of the metric completion \overline{Q} of Q for the length metric induced by the flat structure of Q . The map F extends continuously to \overline{Q} and one sets $\xi_i = F(q_{i+1}) - F(q_i)$ for $i = 1, \dots, 2k$. The following proposition holds true:

PROPOSITION 4.3. *For an appropriate choice of pairwise distinct indices i_1, \dots, i_{k-1} in $\{1, \dots, 2k\}$, the parameters $(\xi_{i_1}, \dots, \xi_{i_{k-1}})$ form a linear parametrisation of \mathcal{F}_ρ .*

Notice that if the triangulation \mathcal{T} used to construct them was totally geodesic then these coordinates would form a polygonal parametrisation. The proof uses arguments similar to those of the proof of Proposition 4.2.

5 Geometric Properties of Flat Surfaces and Characteristic Functions

In this section we develop material and prove several technical lemmas about the intrinsic geometry of flat surfaces which will be used in Sections 7 and 8 in order to understand the geometry of the moduli spaces $\mathcal{F}_{[\rho]}$. Most of the work done in this paper is about reinterpreting questions regarding the geometry of these moduli spaces, in terms of how flat surfaces can degenerate. The material developed below goes some way to answering these questions.

We denote by h the flat metric on a given flat surface N and by d_h (or just by d for short) the induced distance (see Section 2.3). We also denote by $S \subset N$ the set of conical points of N (for the flat structure induced by h).

5.1 Characteristic Functions. We define four quantities associated to N :

– its **systole**²:

$$\sigma(N) = \sigma(N, h) = \inf \left\{ L_h(\gamma) \mid \gamma \text{ simple essential closed curve} \right\};$$

– its **relative systole**:

$$\delta(N) = \delta(N, h) = \inf \left\{ L_h(\gamma) \mid \gamma \text{ joining two distinct singular points} \right\};$$

– its **diameter**:

$$D(N) = D(N, h) = \sup_{x, y \in N} d_h(x, y);$$

– its **relative diameter**:

$$s(N) = s(N, h) = \sup_{x \in N} d_h(x, S).$$

(The terminology *relative* is inspired by the terminology used for translation surfaces, where a *relative period* of an abelian form on a Riemann surface is the value of the integral of this 1-form on a path linking two of its zeroes).

Note that these four quantities all depend linearly on a rescaling of h . Most of the time, we will consider them under the supplementary assumption that the area of N is 1. In this case, one gets geometric invariants attached to N .

A classical fact from Riemannian geometry (see Section 2.3) is that σ , δ , D and s all are realised by piecewise geodesic paths, singular only at points where they cross singular points of N .

PROPOSITION 5.1. *The following four inequalities hold true:*

- | | |
|-----------------------------|-------------------------------|
| (1) $D(N) \geq \delta(N)$; | (2) $D(N) \geq \sigma(N)/2$; |
| (3) $D(N) \geq s(N)$; | (4) $s(N) \geq D(N)/(2n)$. |

Proof. The first and third inequalities are obvious. We now prove the second one. Consider c a curve realising $\sigma(N)$. Let p and q be two points on c diametrically opposed (by this we mean that they cut c into two parts of equal length). We claim that $d(p, q) = \sigma(N)/2$. Otherwise there would be a path of length strictly smaller than $\sigma(N)/2$ going from p to q . This path completed with one of the parts of c going from p to q would form an essential closed curve of length smaller than $\sigma(N)$. Since $d(p, q) = \sigma(N)/2$, we have $D(N) \geq \sigma(N)/2$.

Finally we prove (4). Let now p and q be two points realising $D(N)$. The point p can be joined to a point $s_p \in S$ by a path of length at most $s(N)$, and q to $s_q \in S$ by a path of length at most $s(N)$. Note that given $s' \in S$, there exists $s'' \in S$ distinct from s' which can be joined to the latter by a path of length at most $2s(N)$.

² There is no systole when $g = 0$.

We now prove that one can join s_p and s_q by a path going from singular point to singular point with leaps of length less than $2s(N)$. Let Γ be the graph whose set of vertices is S and for which there is an edge between two singular points if they are at distance less than $2s(N)$. We claim that Γ is connected. If not, let (Γ_1, Γ_2) be a pair of two distinct connected components of Γ with $d(\Gamma_1, \Gamma_2)$ minimal and pick two singular points $s_i \in \Gamma_i$ for $i = 1, 2$ such that $d(s_1, s_2) = d(\Gamma_1, \Gamma_2)$. Since the distance between the two considered connected components has been chosen minimal, the (piecewise) geodesic path realising the distance between s_1 and s_2 cannot meet any other singular point on the way. If its length was more than $2s(N)$, its middle point would be at a distance larger than $s(N)$ from the set of singular points which is impossible. This means that we can build the announced path. Remark that we can find such a path which visits each singular point only once. Such a path has length at most $(n - 1)2s(N)$, hence $D(N) \leq 2s(N) + (n - 1)2s(N) = 2ns(N)$. \square

5.2 Voronoi Decomposition and Delaunay Triangulation. We explain briefly a well-known but important construction in the realm of flat surfaces. We omit the proofs below and refer to [MS91] for a careful and detailed treatment.

The **Voronoi decomposition** of N is defined as follows:

- the 2-cells are the connected components of the set of points $p \in N$ such that $d(p, S)$ is realised by a unique geodesic path;
- the 1-cells are the connected components of the set of points $p \in N$ such that $d(p, S)$ is realised by exactly two distinct geodesic paths;
- the 0-cells are the connected components of the set of points $p \in N$ such that $d(p, S)$ is realised by at least three distinct geodesic paths.

It is checked in [MS91] (see Proposition 4.1) that 0-cells are points and 1-cells are totally geodesic paths.

The **Delaunay decomposition** is defined as the polygonal decomposition which is dual to the Voronoi decomposition in the following way. One checks that D_p , the Euclidean disk of radius $d(p, S)$, injects at p for any p being a 0-cell of the Voronoi decomposition. A Delaunay 2-cell is defined as the convex hull of the elements of S belonging to ∂D_p . A 1-cell is a connected component of the boundary in $N \setminus S$ of such a convex hull and a 0-cell is an element of S .

In [MS91, Lemma 4.3 and Theorem 4.4], it is checked that:

- the set of 0-cells is exactly S ;
- 1-cells are saddle connections;
- for each 1-cell C_1 , there are two distinct 2-cells C_2 and C'_2 such that $C_1 \sqcup C_2 \sqcup C'_2$ is a neighbourhood of C_1 in N ;
- a Delaunay 2-cell is isometric to a convex Euclidean polygon inscribed in a circle of radius less than $s(N)$;
- Delaunay 1-cells have length smaller than or equal to $2s(N)$.

From the Delaunay decomposition (which is unique and only depends on the geometry of N) one can get a **Delaunay triangulation** by subdividing the 2-cells

into triangles. Notice that a Delaunay triangulation is not necessarily a simplicial triangulation since a triangle might not be determined by its vertices. We have now as an immediate corollary of this construction and of Proposition 5.1:

PROPOSITION 5.2. *The length of any 1-cell of any Delaunay triangulation of N is always smaller than $2D(N)$.*

We also prove the following lemma:

LEMMA 5.3. *The interior of any path in N realising $\delta(N)$ is a 1-cell of the Delaunay decomposition of N (hence is a 1-cell of any Delaunay triangulation of N).*

Proof. Remark that if a saddle connection is such that the only paths realising the distance of its middle point to S are the two paths connecting the middle point to the end points, then it is a 1-cell of the Delaunay decomposition. This is a direct consequence of the construction of the latter. We now check that such a saddle connection γ realising $\delta(N)$ must verify the above property.

Assume that there is a second path u going from $p \in S$ to the middle point of γ whose length is less than $\delta(N)/2$. The point p must be different from one of the two endpoints of γ , and if concatenating the half of γ starting from this point and u , one gets a path v of length less than $\delta(N)$ going from two distinct elements of S . Being singular at the middle of γ , v can be shortened in order to get a path whose length is strictly less than $\delta(N)$ which is impossible. Therefore γ must be a 1-cell of the Delaunay decomposition. \square

5.3 Surfaces with Large Diameter. The aim of this subsection is to prove that flat surfaces with large diameter and finite linear holonomy must necessarily contain long flat cylinders. If one dismisses the hypothesis that the linear holonomy is finite, one can build counterexamples by gluing cones of very small angle. This was already known for spheres (see [Thu98]) or when the linear monodromy ranges in $\{-1, 1\}$ (see [MS91, Corollary 5.5]). The proof of Proposition 5.6 below is highly inspired by the techniques developed in [MS91].

Elementary facts about cones. We remind the reader that C_θ stands for the (Euclidean) cone of angle $\theta \in]0, +\infty[$, namely the metric space obtained by gluing the sides of a plane sector of angle θ . Its vertex is denoted by 0 and one sets $C_\theta^* = C_\theta \setminus \{0\}$. This cone with the apex removed, does not contain closed regular geodesic but, when $\theta < \pi$, it contains piecewise geodesic paths with only one angular point. More precisely:

PROPOSITION 5.4. *If $\theta < \pi$ then for any point $p \in C_\theta^*$:*

- *there exists a unique closed simple piecewise geodesic path in C_θ^* singular only at p ;*
- *the interior angle of the latter at the angular point is $\pi - \theta$;*
- *the length of this piecewise geodesic path is $2 \sin(\theta/2) \cdot d(0, p)$.*

Proof. The proof of the proposition is straightforward after noticing that such a cone is obtained after doing the gluing pictured in Figure 6. \square

LEMMA 5.5. *Let N be a flat surface and γ be a piecewise geodesic path of length $L(\gamma)$ on N with one angular point which avoids conical points. Assume that in a small neighbourhood of its angular point, γ cuts N into two angular sectors of angles $\pi + \theta$ and $\pi - \theta$ respectively, with $0 < \theta < \pi$. Then*

- (1) *the linear holonomy along γ is $e^{i\theta}$ or $e^{-i\theta}$;*
- (2) *there is a cone point q of N such that $d(q, \gamma) \leq L(\gamma)/(2 \tan(\theta/2))$.*

Proof. The point is that such a geodesic γ has a neighbourhood that is isometric to a neighbourhood of the unique (up to isometry) closed geodesic of length $L(\gamma)$ of the cone of angle θ . The only obstruction for this isometry to extend to the whole cone is that the boundary of its definition domain meets a singular point of N (one can use the exponential map along γ). Otherwise γ is on the cone of a cone point of N whose associated conical angle is θ . In any case, there is a singular point of N , whose distance to γ is less than the distance from the geodesic of length $L(\gamma)$ in C_θ to the apex of C_θ . This distance is exactly $L(\gamma)/(2 \tan(\theta/2))$. \square

PROPOSITION 5.6. *Let $\rho \in H^1(N, \mathbb{U}, \theta)$ be such that $\text{Im}(\rho)$ is finite. There exist two positive constants $K_1(\rho)$ and $K_2(\rho)$ such that for every flat surface $N \in \mathcal{F}_\rho$ normalised such that its area is 1, if $D(N) > K_1(\rho)$ then N contains an embedded flat cylinder of length at least $K_2(\rho)D(N)$.*

Proof. Let N be an element of \mathcal{F}_ρ . Let $p \in N$ be a point maximizing the distance s to the set of singularities, i.e. such that $s = s(N) = d(p, S)$ where $S \subset N$ stands for the set of cone points of N . Throughout the proof, we will mainly work with $s = s(N)$ which, as a function, is of same order as $D(N)$ according to the last two points of Proposition 5.1.

Let r_p be the injectivity radius at p . Then $r_p < 1/\sqrt{\pi}$ since the area of N is one. If $s > r_p$ then $\overline{D(r_p)}$, the closed Euclidean disk of radius r_p , can be immersed in N at p (since s is realised at p). There are two distinct points a and b on the boundary




Figure 6: The simple closed piecewise geodesic path with one angular point at p on C_θ (in green).

of $\overline{D(r_p)}$ which project onto the same point in N and the immersion $i : \overline{D(r_p)} \rightarrow N$ is injective on $D(r_p)$, by definition of r_p . Therefore the chord joining a and b maps to a piecewise closed geodesic γ path in N , with one angular point at $i(a) = i(b)$.

We claim that if s is large enough, then the linear holonomy along γ must be trivial. This is a corollary of Lemma 5.5. More precisely, if $\text{Im}(\rho) = \langle e^{2i\pi/m} \rangle$ and $s > r_p(1 + \tan(\pi/m)^{-1})$, γ cuts N at $i(a)$ into two angular sectors both of angles π . Hence γ is a closed regular geodesic which belongs to a flat cylinder C and the holonomy along γ is 1. Moreover, a and b must be diametrically opposed and γ must have length $2r_p$. Otherwise one side of the cylinder C would be covered by $D(r_p)$. But then r_p would not be the injectivity radius at p . The closed geodesic γ contains p and the cylinder C containing γ has length at least $2\sqrt{s^2 - r_p^2}$, because any cylinder on a flat surface can be extended until its boundary meets a singular point.

If one assumes that $s = s(N) \geq 2/\sqrt{\pi} \geq 2r_p$, then the cylinder we have found has length at least $s(N)\sqrt{3}$ hence at least $D(N)\sqrt{3}/(2n)$ according to Proposition 5.1. Recall that to ensure that the linear holonomy along γ is trivial and therefore that γ is a closed geodesic belonging to a cylinder, we had to assume $s(N) > r_p(1 + \tan(\pi/m)^{-1})$. Since $r_p \leq \frac{1}{\sqrt{\pi}}$, the statement of the lemma follows if one takes $K_1(\rho) = \frac{2n}{\sqrt{\pi}} \max \{2, (1 + \tan(\pi/m)^{-1})\}$ and $K_2(\rho) = \sqrt{3}/(2n)$. \square

5.4 Collisions. An important issue to describe the metric completion of the moduli spaces $\mathcal{F}_{[\rho]}$ is to characterize geometrically what happens when two singular points collide, *i.e.* when $\delta(N)$ goes to zero. We prove below two results describing situations when such a collision cannot occur, at least without the diameter going to infinity.

LEMMA 5.7. *Let θ_1 and θ_2 be two positive angles such that $\theta_1 + \theta_2 < 2\pi$. There exists a constant $K(\theta_1, \theta_2) > 0$ such that if Σ is any flat sphere with n conical singularities satisfying the three following conditions:*

- *all the cone points p_1, \dots, p_n of Σ have positive curvature;*
- *the cone angles of Σ at p_1 and p_2 are θ_1 and θ_2 respectively;*
- *the area of Σ is 1;*

then the following holds true: $d(p_1, p_2) \geq K(\theta_1, \theta_2)$.

This lemma tells us that two too positively curved singular points cannot collide. We would like to draw attention to the fact that, in the authors' opinion, this lemma (and in particular, its proof!) is missing in [Thu98].

Proof. The idea of the proof is to compare this situation to the case of the sphere Σ^0 with three cone points, of respective angles θ_1 , θ_2 and $2\pi - \theta_1 - \theta_2$. Such a sphere is unique up to dilatation of the flat metric and is built by gluing two isometric triangles of angles $\theta_1/2$, $\theta_2/2$ and $(2\pi - \theta_1 - \theta_2)/2$.

Let p_1^0 , p_2^0 and p_3^0 be the cone points on Σ^0 of respective angles θ_1 , θ_2 and $2\pi - \theta_1 - \theta_2$. Normalise Σ^0 so that the length of the unique geodesic l^0 from p_1^0 to p_2^0 has

the same length as the one from p_1 to p_2 on Σ , denoted by l . Remark that Σ^0 is the disjoint union of geodesic paths going from p_3^0 to points of l^0 .

A neighbourhood of l^0 in Σ^0 is isometric to a neighbourhood of l in Σ . We extend such an isometric identification using the remark above, developing the geodesics of the decomposition. The only obstruction to do so appears if such a geodesic meets a singular point, which can only happen for a finite number of such geodesics. We denote by A the finite union of those parts of geodesics on which the isometry cannot be extended.

We have thus defined a local isometry

$$i : \Sigma^0 \setminus A \longrightarrow \Sigma .$$

Since all the singular points of Σ have positive curvature, the closure of $i(\Sigma^0)$ must also be open and since i is a local isometry, one gets

$$\text{area}(\Sigma^0) \geq \text{area}(\Sigma) = 1 .$$

The uniform bound on the area of Σ^0 gives a uniform bound on $d(p_1^0, p_2^0)$ which equals $d(p_1, p_2)$ by construction. □

LEMMA 5.8. *Let M be a flat torus with two cone points p_1 and p_2 . There exists a pseudo-hexagon P such that M is isometric to P / \sim where \sim is one of the three gluing patterns of Figure 2.*

Proof. Let Γ be a connected graph in the 1-skeleton of the Delaunay decomposition of M such that $M \setminus \Gamma$ is connected and simply connected. Γ has exactly for vertices the two cone points of M . By a Euler characteristic argument, its number of edges e must satisfy $2 - e + 1 = \chi(M) = 0$ and therefore $e = 3$.

One easily checks that the only connected graphs with two vertices and three edges that one can draw on a torus are the three graphs represented on Figure 3. Then cutting along Γ gives the expected pseudo-polygonal model for M . □

PROPOSITION 5.9. *Let $(M_\ell)_{\ell \in \mathbb{N}}$ be a sequence of flat tori with two cone points belonging to a leaf $\mathcal{F}_{[\rho]}$ with $\text{Im}(\rho)$ finite. Assume that for all $\ell \in \mathbb{N}$, M_ℓ has area 1. If $\lim_{\ell \rightarrow +\infty} \delta(M_\ell) = 0$ then $D(M_\ell) \rightarrow +\infty$ as ℓ goes to infinity.*

Proof. Suppose that $(M_\ell)_{\ell \in \mathbb{N}}$ and ρ are as in the statement and assume that $D(M_\ell)$ does not go to infinity although $\delta(M_\ell)$ tends to zero when $\ell \rightarrow +\infty$. Then, up to extracting an appropriate subsequence, we can assume that the $D(M_\ell)$'s are bounded. For any ℓ , consider the Delaunay decomposition of M_ℓ and take in its 1-skeleton a graph Γ_ℓ such that

- Γ_ℓ contains a curve realising $\delta(M_\ell)$, as guaranteed by Lemma 5.3;
- the set of vertices of Γ_ℓ is equal to the set of singular points of M_ℓ ;
- $Q_\ell = M_\ell \setminus \Gamma_\ell$ is simply connected.

According to Lemma 5.8, for any $\ell \in \mathbb{N}$, the metric completion $\overline{Q_\ell}$ of Q_ℓ is a pseudo-hexagon (*i.e.* a pseudo-polygon with six sides) whose lengths of the sides are uniformly bounded (according to Proposition 5.2) and the gluing pattern to recover M_ℓ is one of the three patterns of Figure 2.

Again up to extracting a subsequence, we can assume that for any $\ell \in \mathbb{N}$:

- (1) M_ℓ can be obtained from $\overline{Q_\ell}$ by using the same gluing pattern;
- (2) the sides glued together always form the same angle;
- (3) the length of each side converges.

(To assume (2), one has to use that ρ has finite image. The fact that one can assume that (3) holds true as well follows from Proposition 5.2.)

Since the lengths of two sides go to zero (the ones which are identified by the gluing with the curve realising $\delta(M_\ell)$ in M_ℓ), the sequence of pseudo-hexagons $(Q_\ell)_{\ell \in \mathbb{N}}$ converges to a quadrilateral whose opposite sides have the same length and therefore are parallel. Since $\text{Im}(\rho)$ is finite, this implies that the corresponding sides in $\partial Q_\ell = \overline{Q_\ell} \setminus Q_\ell$ were parallel for all ℓ sufficiently large. This forces the gluing pattern to be Pattern 1 or 2 of Figure 2. But a hexagon glued according to one of these patterns and having two parallel pairs of sides glued together must be a regular torus with no singular point (this is an easy exercise left to the reader). This would force $\mathcal{F}_{[\rho]}$ to contain regular tori, which is impossible since we have supposed that its elements have exactly two singular points. Therefore the sequence of diameters $(D(M_\ell))_{n \in \mathbb{N}}$ must go to infinity as ℓ does. \square

5.5 Closed Curves Realising the Systole. As well as collisions, the ways in which simple closed curves can collapse are also very important to characterise.

LEMMA 5.10. *Let N be a flat torus with $n \geq 2$ cone points and suppose that p_1 is the only cone point which has negative curvature. The systole $\sigma(N)$ is realised by a simple closed piecewise geodesic which meets the set of cone-points only once at p_1 . Moreover, the only point at which it might not be smooth is p_1 .*

Proof. Consider the set of non homotopically trivial closed curves. Classical Riemannian geometry (see Section 2.3) ensures there exists a minimiser of the length functional on this set and that it is piecewise geodesic.

We claim that such a minimiser is simple. Otherwise it could be decomposed into two closed curves of strictly shorter length with at least one of these two being essential.

A minimiser cannot pass through a point of positive curvature because otherwise one can deform it in order that it avoids the cone point and that its length is shorter. Therefore the only cone point it might pass through is p_1 , and one can always make sure that there is a minimising path passing through p_1 : otherwise the path is actually totally geodesic and a neighbourhood of this path is a flat cylinder which can be extended until meeting a cone point which must be p_1 . Any boundary component of this extended flat cylinder would be a required path. \square

A path realising $\sigma(N)$ cuts the surface at p_1 in two angle sectors, whose angle must be bigger than π (otherwise one can shorten the path by passing on the side where the angle is smaller than π). Two possibilities can occur:

- (1) one of the angle equals π ; in this case such a path bounds a flat cylinder;
- (2) both angles are strictly bigger than π .

For our purpose, it is important to distinguish these two situations.

In the case we are mostly interested in (when $g = 1$ and $\theta = (\theta_i)_{i=1}^n$ is such that only the point of cone angle θ_1 carries negative curvature), they actually correspond to two geometric aspects of $\mathcal{F}_{[\rho]}$: flat tori verifying (1) are in a cusp while those verifying (2) are close to a stratum corresponding to the Devil's surgery \mathcal{S}_3 , see Section 6.3. The proposition below proves that, in the very specific case when $g = 1$ and ρ is rational, if the diameter remains bounded and the systole goes to zero, we are in situation (2).

PROPOSITION 5.11. *Assume that $g = 1$, $\theta = (\theta_i)_{i=1}^n$ is such that only θ_1 is bigger than 2π and $\rho \in \mathbb{H}^1(N, \mathbb{U}, \theta)$ has finite image. For all $K > 0$, there exists a constant $\epsilon(K) > 0$ such that for $N \in \mathcal{F}_{[\rho]}$ normalised such that its area equals 1 the following holds true. If $D(N) \leq K$ and $\sigma(N) \leq \epsilon(K)$, then any curve c realising the systole and passing through p_1 the point of negative curvature of N cuts p_1 into two angular sectors whose angles both are strictly larger than π .*

Proof. We argue by contradiction. Assume that there exist a constant K and a sequence $(N_m)_{m \in \mathbb{N}}$ of flat tori such that for all $m \in \mathbb{N}$, one has:

- $N_m \in \mathcal{F}_{[\rho]}$ and its area is equal to 1;
- $\sigma(N_m) \leq \frac{1}{m}$;
- N_m contains a cylinder C_m of width $\sigma(N_m)$ (i.e. we are in situation (1) described above);
- $D(N_m) \leq K$.

For all $m \in \mathbb{N}$, $N_m \setminus C_m$ is a sphere whose boundary is the union of two piecewise geodesic closed curves of the same length $\sigma(N_m)$ touching at the only point where they both are singular, namely $p_1(m)$ the cone point of negative curvature of N_m . One can cut at the point where the two boundary circles touch, and glue together the two geodesic parts of the new boundary circles (which have the same length) to get a flat sphere S_m . Since the cone angle θ_1 at $p_1(m)$ is supposed to belong to $]2\pi, 4\pi[$, the resulting sphere has only positively curved cone points. The angles $\theta'_1(m)$ and $\theta'_2(m)$ at the two 'new' cone points of S_m created by the previous cutting and pasting operation, must satisfy $\theta'_1(m) + \theta'_2(m) + \pi + \pi = \theta_1$. Since θ_1 is strictly smaller than 4π , we get:

$$\theta'_1(m) + \theta'_2(m) < 2\pi.$$

For $k = 1, 2$, the cone angle $\theta'_k(m)$ must be such that $e^{i\theta'_k(m)} \in \text{Im}(\rho)$ because $e^{i\theta'_k(m)}$ is the linear holonomy of a curve in the free homotopy class of the curve

realising the systole. Therefore these two angles can take only a finite number of values. So, up to extracting a subsequence, one can assume that these two cone angles are independent of m . The fact that the sequence of diameters $(D(M_m))_{m \in \mathbb{N}}$ is bounded by K implies that the length of C_m is bounded by $2K$. Therefore the area of S_m , which is larger than $1 - 2\sigma(N_m)K$, is bigger than $1/2$ provided that m is large enough. Hence $\theta'_1(m) + \theta'_2(m) < 2\pi$ and the distance between the associated cone points, which equals $\sigma(N_m)$ by construction, goes to zero. This contradicts Lemma 5.7 and proves the proposition. \square

6 Surgeries

A **surgery** is a procedure through which a new flat surface with conical singularities is produced from another one by means of geometrical gluing and pasting relying on elementary Euclidean geometry. This notion naturally appears when studying moduli spaces of flat surfaces (implicitly in [Thu98] but also more explicitly in [KZ03]). In the previous section, we have studied different ways for flat surfaces to degenerate, namely sequences of surfaces containing very large embedded flat cylinders, essential curves which collapse or cone points colliding. In the present section, we introduce several surgeries which are to be seen as the inverse processes of the aforementioned degenerations.

We will distinguish five distinct types of surgeries:

- the first one, denoted by \mathcal{S}_1 , was known and implicitly considered by Thurston in [Thu98]. It consists in blowing up a singular point of positive curvature into two singular points of positive curvature;
- the second surgery, denoted by \mathcal{S}_2 , is a straightforward generalisation of the first one, which allows to blow up points of negative curvature. We will therefore refer to both \mathcal{S}_1 and \mathcal{S}_2 as **Thurston's surgeries**;
- it seems to us that the third surgery \mathcal{S}_3 is new. We call it **Devil's surgery**. It consists in creating a handle by removing the neighbourhoods of two singular points and gluing their boundaries together;
- the fourth surgery \mathcal{S}_4 consists in blowing up a regular point into three singular points. We call it the **Kite surgery**;
- the last surgery \mathcal{S}_5 consists in creating a handle by adding a long flat cylinder to any flat surface having two isometric totally geodesic boundary components.

Surgeries \mathcal{S}_1 , \mathcal{S}_2 and \mathcal{S}_4 could have been seen as the same in a more general presentation but we find it more convenient to differentiate them for our purpose. At the end of this Section, we compute the signature of the area form in the case we are interested in by using surgeries and we give a definition of the notion of *geometric convergence* which will be central in the description of the metric completion of $\mathcal{F}_{[\rho]}$ carried on in Section 7.

6.1 Thurston’s Surgery \mathcal{S}_1 for a Cone Angle Smaller Than 2π . Let N be a flat surface of genus g with cone angles $\theta = (\theta_1, \dots, \theta_n)$ at $p_1, \dots, p_n \in N$. Let $\mathcal{F}_{[\rho]}$ be the leaf of Veech’s foliation to which N belongs. Suppose that $\theta_1 < 2\pi$ and let θ'_1 and θ''_1 be two angles smaller than 2π such that

$$2\pi - \theta_1 = (2\pi - \theta'_1) + (2\pi - \theta''_1). \tag{9}$$

In this subsection, we describe a surgery building out flat surfaces of genus g with $n + 1$ singular points of cone angle $\theta' = (\theta'_1, \theta''_1, \theta_2, \dots, \theta_n)$ from N (note that because we have assumed (9), the new angle datum θ' still satisfies Gauß-Bonnet formula (1)). The surgery is local on N , in the sense that it is performed on a small neighbourhood of p_1 without modifying the rest of the surface.

Choose a point p in a small neighbourhood C of p_1 isomorphic to the portion of cone $C_{\theta_1}(\epsilon)$ for a certain $\epsilon > 0$ (see §2.1). As θ'_1 is bigger than θ_1 , there are exactly two distinct segments of the same length issuing from p which meet at their endpoints and form an interior angle equal to $2\pi - \theta'_1$ at p (see Figure 7).

The surgery works the following way: delete the bigon on C which corresponds to the quadrilateral B in grey on Figure 7. Its sides are two geodesics which have the same length and the same endpoints p'_1 and p''_1 . Removing the bigon and gluing these two segments together, one gets a new flat surface N' having two cone points of angle θ'_1 and θ''_1 at p'_1 and p''_1 respectively.

Recall that $\mathcal{F}_{[\rho]}$ is the leaf of Veech’s foliation to which N belongs. There exists a neighbourhood U of N in $\mathcal{F}_{[\rho]}$ and $\epsilon > 0$ sufficiently small so that for all flat surfaces in U , the previous surgery can be performed for all p in a disc of radius ϵ centered at p_1 , the cone point of angle θ_1 . Remark that the class $[\rho']$ such that $N' \in \mathcal{F}_{[\rho']}$ does not depend on the choice of N in U .

This allows us to define a map

$$\begin{aligned} \mathcal{S}_1 : C_{\theta_1}^*(\epsilon) \times U &\longrightarrow \mathcal{F}_{[\rho']} & (10) \\ (p, N) &\longmapsto N', \end{aligned}$$

where $C_{\theta_1}^*(\epsilon)$ is $C_{\theta_1}(\epsilon)$ minus its apex (see Section 2). This definition requires an identification of $C_{\theta_1}(\epsilon)$ with a neighborhood of the cone point of angle θ_1 in each flat surface element of U . We do this by choosing a geodesic path c joining p_1 and p_2 . This path survives in a neighbourhood of N in U . We decide that $C_{\theta_1}(\epsilon)$ is embedded in an element of U in such a way that it always meets the previous geodesic path in the same locus - this latter requirement defining unambiguously such an embedding. Let $z_0 \in \mathbb{C}$ be a (germ of) linear parametrisation of $C_{\theta_1}^*(\epsilon)$ such that $z_0 \in \mathbb{R}_+$ if and only if the corresponding point p belongs to the aforementioned geodesic path joining p_1 to p_2 .

There is a little ambiguity for the choice of the path c whenever some elements of U have non-trivial isometries; in that case the identification of $C_{\theta_1}^*(\epsilon)$ on elements of U cannot be made continuous. We choose to ignore this difficulty for a moment and then we will address it in Remark 6.2 below.




Figure 7: Thurston’s surgery \mathcal{S}_1 consists in (1) cutting the surface along the dashed blue segment; then (2) removing the grey piece of the surface and gluing the two red segments together.

PROPOSITION 6.1. *We use the notations introduced just above.*

- (1) *If (z_1, \dots, z_m) is a linear parametrisation of U then (z_0, z_1, \dots, z_m) is a linear parametrisation of the image $U' \subset \mathcal{F}_{[\rho']}$ of the map (10).*
- (2) *The map \mathcal{S}_1 is a local biholomorphism.*
- (3) *If all elements of U have no non-trivial isometry then \mathcal{S}_1 is one-to-one.*

Proof. Consider a topological polygonal model of $N \in U$ which is such that c is an edge of this polygon. This model can be extended to a model of N' by adding a point on the side representing the class of c , see Figure 8.

Let (z_1, \dots, z_m) be a linear parametrisation of U such that the geodesic path C from p_1 to p_2 develops on z_1 , and let z_0 be the complex number onto which the geodesic path from p_1 to p (which is going to become p'_1 after the surgery) develops. Let N' be a flat surface obtained after applying a \mathcal{S}_1 surgery to N . Let (w_0, \dots, w_m) be a linear parametrisation of a neighbourhood of N' in $\mathcal{F}_{[\rho']}$ associated to the extended polygonal model of N such that w_0 represents the shortest geodesic path from p'_1 to p''_1 on N' and w_1 represents a path joining p'_1 to p_2 , w_2 a path joining p''_1 to p_3 (note that p_3 and w_2 do not appear explicitly on Figures 8 and 9).

According to Figure 9, which is a superposition of the developing maps of N and N' near the point p_1 , where the parametrisations w and z correspond respectively




Figure 8: The surface N before surgery \mathcal{S}_1 on the left and the surface N' obtained after surgery on the right.




Figure 9: A superposition of the parametrizations w and z before and after surgery, near p_1 .

to before (the surface N) and after (the surface N') proceeding to the surgery, the following relations hold true

$$w_0 = \rho_0 z_0, \quad w_1 = z_1 + z_0 \quad \text{and} \quad w_2 = z_2 + \rho_1 z_0,$$

where ρ_0 and ρ_1 are constants (which can be made explicit by means of elementary geometry of Euclidean triangles) depending only on θ_1, θ_1' and θ_1'' . All the other w_i 's can be expressed in a similar fashion. Furthermore, if w_i represents a path involving end points different from p_1' and p_1'' , it is equal to one of the z_j 's.

Therefore (z_0, z_1, \dots, z_m) is a linear parametrisation of a neighbourhood of N' in $\mathcal{F}_{[\rho']}$. This implies directly the two first points of the proposition, in particular the fact that \mathcal{S}_1 is a local biholomorphism. It remains to prove the injectivity of \mathcal{S}_1 under the additional hypothesis that all the elements of U have a trivial isometry group. The length of the shortest path from p_1' to p_1'' (which equals $|z_0|$ provided that the latter is small enough in the area 1 normalisation) is a geometric invariant. The surface N from which N' is obtained from also is a geometric invariant. Assume that there exist two points p and p' on $C_{\theta_1}^*(\epsilon)$ such that the resulting surfaces from the surgery at p and p' are the same. This would imply that the initial surface has an isometry fixing p_1 and sending p to p' . The (pure) isometry group of a surface being finite, \mathcal{S}_1 is a local biholomorphism which is one-to-one if all the elements of U are all isometry free. □

REMARK 6.2. (1) It is worth giving a more abstract and intrinsic description of the surgery introduced above. Let $U \subset \mathcal{F}_{[\rho]}$ as above and assume that none of its elements admits a nontrivial isometry. Then there exists a ‘universal flat curve $\nu_U : \mathcal{T}_U \rightarrow U$ over U ’: it is a map such that the fiber over a flat surface N viewed as a point of U is N itself, but this time viewed as a 2-dimensional flat surface with conical singularities. This family of surfaces comes with n sections $p_i : U \rightarrow \mathcal{T}_U$ which are such that $p_i(N)$ is the i -th cone point of N for every $i = 1, \dots, n$. One denotes by P_i the image of p_i for every i and by $\mathcal{T}_U^* = \mathcal{T}_U \setminus \cup_{i=1}^n P_i$ the ‘ n -punctured universal flat curve over U .’

Within this formalism, one can verify that Thurston's surgery \mathcal{S}_1 admits an intrinsic definition as the (germ of) map $(\mathcal{T}_U^*, P_1) \rightarrow \mathcal{F}_{[\rho']}$ which, for any $p \in \mathcal{T}_U^*$ sufficiently close to P_1 associates the flat surface N' obtained by performing the surgery described by Figure 7 above on the surface $N = \nu_U(p)$ with respect to p and the cone point $p_1(N)$. Clearly, obtaining the more explicit definition (10) just amounts to trivializing $\mathcal{T}_U \rightarrow U$ along P_1 .

(2) The interest of the preceding, more conceptual, approach is that it points out the main issue when some of the elements of U admit nontrivial isometries and how to deal with it. Indeed, in this case, there is no universal curve over U but one will exist over a non-trivial orbifold cover \tilde{U} of U and working with the latter, one can define Thurston's surgery the same way than above.

For instance and more concretely, if $N_0 \in U$ is such that $\text{PIso}^+(N_0)$ is non-trivial, then it is necessarily cyclic of finite order, say m , according to §2.5. In this case there exists $\tilde{U} \rightarrow U$ an orbifold cover of order m of U , whose deck transformation group is isomorphic to the isometry group of N_0 , on which the identification of $C_{\theta_1}^*(\epsilon)$ with some neighborhoods of the corresponding cone points in flat surfaces belonging to U can be made in a continuous way.

Therefore the surgery still defines a map

$$\begin{aligned} \mathcal{S}_1 : C_{\theta_1}^*(\epsilon) \times \tilde{U} &\longrightarrow \mathcal{F}_{[\rho']} \\ (p, N) &\longmapsto N' \end{aligned}$$

which is equivariant under the action of the isometry group of N_0 .

6.2 Thurston's Surgery \mathcal{S}_2 for a Cone Angle Greater Than 2π . Assume now that $\theta_1 > 2\pi$. Let $\theta'_1 > 2\pi$ and $\theta''_1 < 2\pi$ be such that

$$2\pi - \theta_1 = (2\pi - \theta'_1) + (2\pi - \theta''_1).$$

Let V_1 be a neighbourhood of p_1 isometric to a portion of cone $C_{\theta_1}(\epsilon)$ for a certain $\epsilon > 0$. Define $\eta = \theta_1 - 2\pi$ and let p be a point of V_1 . If p is close enough to the singular point p_1 there is a unique 4-gon P in V_1 having the following properties (see Figure 10):

- p and p_1 are opposite vertices of P and the external angles of the latter at these two points are θ''_1 and η respectively;
- the external angles at the two other vertices of P both are $\theta'_1/2$;
- the sides of P meeting at p (resp. at p_1) have the same length.

We build a new flat surface N' in the following way: we remove the interior of P and glue together the sides meeting at p_1 and p . We obtain a flat surface with a singularity of angle θ''_1 and another singularity of angle θ'_1 . As when $\theta_1 < 2\pi$, the class $[\rho']$ such that $N' \in \mathcal{F}_{[\rho']}$ does not depend on the choice of N . For a neighbourhood U in $\mathcal{F}_{[\rho]}$, this allows us to define a map

$$\begin{aligned} \mathcal{S}_2 : C_{\theta_1}^*(\epsilon) \times U &\longrightarrow \mathcal{F}_{[\rho']} \\ (p, N) &\longmapsto N'. \end{aligned}$$

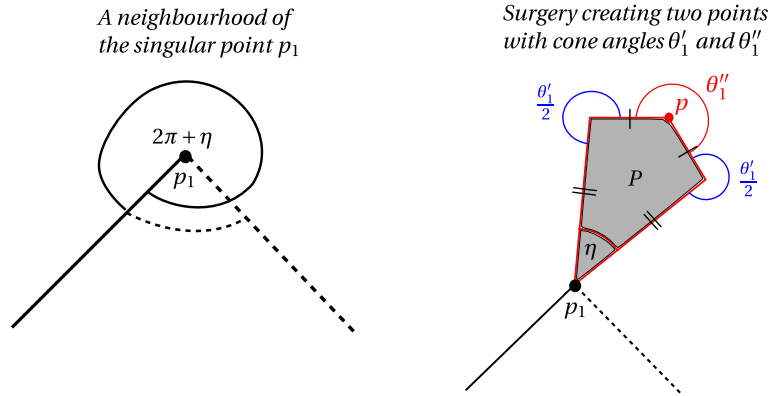


Figure 10: Thurston's surgery \mathcal{S}_2 .

As in the case when $\theta_1 < 2\pi$, the map \mathcal{S}_2 is a local biholomorphism on its image, and for any linear parametrisation (z_1, \dots, z_m) of U , (z_0, z_1, \dots, z_m) is a linear parametrisation of the image of \mathcal{S}_2 , where z_0 is the complex number on which the segment $[p_1, p]$ develops. The proof is exactly the same as in the $\theta_1 < 2\pi$ case and is left to the reader. A similar remark to Remark 6.2 also holds true for \mathcal{S}_2 as well.

The following remark will play a crucial role in the proof of Proposition 7.7 which is one of the main results of the paper.

REMARK 6.3. If l is the length of the segment between the points of angle θ'_1 and θ''_1 and L the length between p_1 and the point of angle θ'_1 (after surgery), then

$$L = \frac{\sin\left(\frac{\theta''_1}{2}\right)}{\sin\left(\frac{\theta'_1 + \theta''_1}{2}\right)} l. \tag{11}$$

Conversely, given a flat surface N with two conical points p'_1 and p''_1 of angles $\theta'_1 > 2\pi$ and $\theta''_1 < 2\pi$ which are linked by a saddle connection c of length l , one can wonder when it is possible to reverse Thurston's surgery \mathcal{S}_2 . The only obstruction to doing so is that c can be extended on the side of p'_1 (the point of negative curvature) on a distance equal to the right-hand side of (11), while cutting θ'_1 in half. The proof of this claim is elementary and left to the reader.

6.3 Devil's Surgery \mathcal{S}_3 . Let $\theta'_1, \theta''_1, \theta_2, \dots, \theta_n$ be the respective cone angles at the cone points $p'_1, p''_1, p_2, \dots, p_n$ of a flat surface N . We make the simplifying assumption that θ'_1 and θ''_1 are (strictly) less than π , but we will see later on that the surgery we are going to describe can still be performed for θ'_1 and θ''_1 less than 2π .

Let C'_1 and C''_1 be two neighbourhoods of p'_1 and p''_1 , isometric to $C_{\theta'_1}(\epsilon)$ and $C_{\theta''_1}(\epsilon)$ respectively, with $\epsilon > 0$ sufficiently small. Consider two points q' and q'' in C'_1 and C''_1 respectively such that the unique closed geodesic paths c' and c'' in C'_1

and C_1'' respectively, singular only at q' and q'' have the same length. Removing the ‘upper parts’ of the two cones C_1' and C_1'' by cutting along c' and c'' respectively, one can glue c' and c'' isometrically in such a way that q' and q'' are glued together. One gets a flat surface \widehat{N} of genus $g + 1$ with n cone points of angle $\theta_1, \theta_2, \dots, \theta_n$ with $\theta_1 = 2\pi + \theta'_1 + \theta''_1$ (see Figure 11).

Let β be a simple curve in N , avoiding its conical points, joining two regular points of c' and c'' respectively which are glued together by the identification considered above. Denote by $\widehat{\beta}$ the simple loop in the regular part of \widehat{N} obtained from

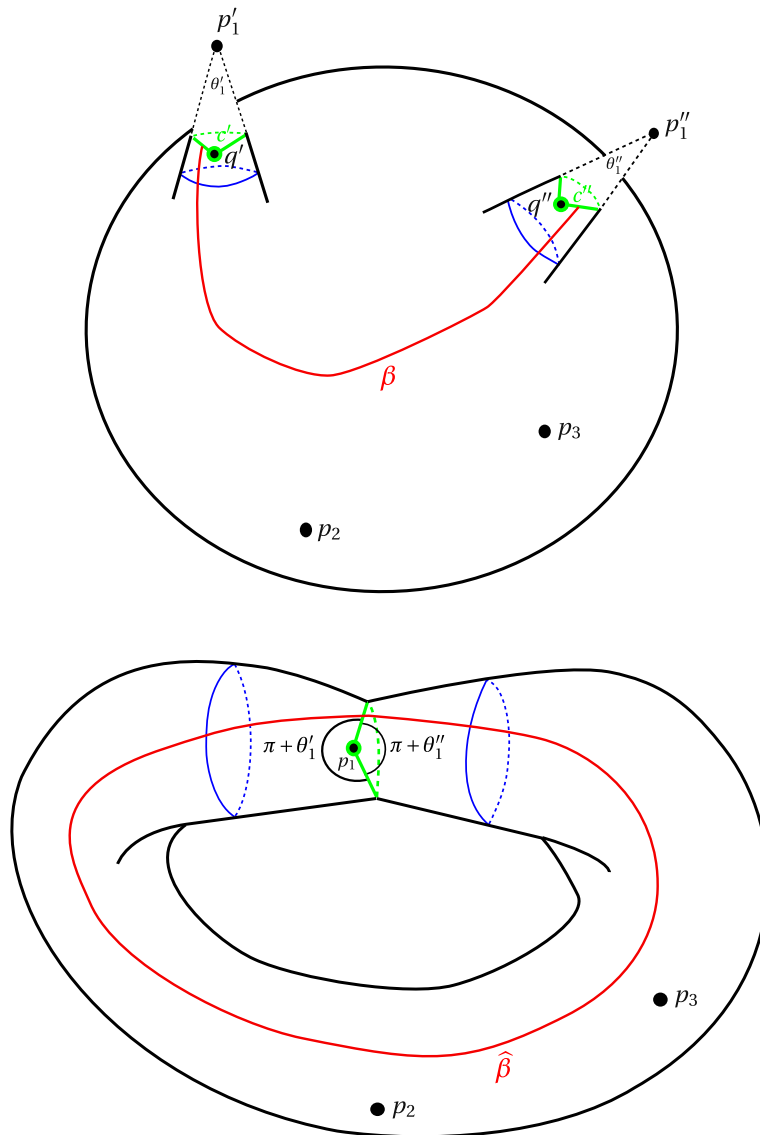


Figure 11: Devil's surgery from a flat sphere to a flat torus.

β by this gluing. We want to perform the surgery in such a way that the linear holonomy of the resulting flat surface does not change when q' and q'' move on the portions of cones C'_1 and C''_1 respectively and equals a certain (class under the action of the pure mapping class group of) $\widehat{\rho}$. It is equivalent to the fact that the holonomy along $\widehat{\beta}$ does not vary, because the holonomy of the resulting flat surface \widehat{N} is totally determined by the holonomy of the original flat surface N and the holonomy along $\widehat{\beta}$.

When q' moves on the circle it belongs to (namely, the set of points on C'_1 whose distance to the apex p'_1 is precisely $d(p'_1, q')$), (a lift to $\mathbb{R} = \widetilde{U}$ of) the linear holonomy along $\widehat{\beta}$ increases exactly by the angle that q' makes relatively to its initial position. Hence if we want to keep the linear holonomy along $\widehat{\beta}$ constant, we have to move q'' by the same angle as q' .

This allows us to build a map:

$$\begin{aligned} \widetilde{\mathcal{F}}_3 : \widetilde{C_{\theta'_1}^*(\epsilon)} \times U &\longrightarrow \mathcal{F}_{[\widehat{\rho}]} \\ (p, N) &\longmapsto \widehat{N} \end{aligned}$$

where U is an open subset of $\mathcal{F}_{[\widehat{\rho}]}$, the moduli space to which the original surface belongs.

PROPOSITION 6.4. *Suppose that every element of U has no non-trivial isometry. Then $\widetilde{\mathcal{F}}_3$ is a covering map onto its image. If θ'_1 and θ''_1 are not commensurable (i.e. if $\theta'_1/\theta''_1 \notin \mathbb{Q}$) then it is a biholomorphism onto its image.*

However, if $\theta'_1/\theta''_1 = k/l$ for some coprime positive integers k and l , then the deck group is the group generated by the rotation of angle $l\theta'_1 = k\theta''_1$. In particular $\widetilde{\mathcal{F}}_3$ factors through

$$\begin{aligned} \mathcal{S}_3 : C_{l\theta'_1}^*(\epsilon) \times U &\longrightarrow \mathcal{F}_{[\widehat{\rho}]} \\ (p, N) &\longmapsto \widehat{N} \end{aligned}$$

which is a biholomorphism onto its image.

Proof. The fact that $\widetilde{\mathcal{F}}_3$ is a local biholomorphism just relies on the fact that one can get a linear parametrisation of $\mathcal{F}_{[\widehat{\rho}]}$ by adding the parameter z_0 to any linear parametrisation of U , with z_0 a linear parametrisation of $C_{l\theta'_1}^*(\epsilon)$.

The key fact is that two surfaces resulting from the surgery under consideration are isometric if and only if the two points q' and q'' are the same because we have supposed that the elements of U do not have non-trivial isometries (see the proof of Proposition 6.1). When q' varies in the universal covering of $C_{\theta'_1}^*(\epsilon)$, the point q'' eventually comes back to its initial position if and only if the two cone angles θ'_1 and θ''_1 are rationally related: the lack of injectivity appears when q' and q'' come back for the first time to their initial position as q' turns around p'_1 and q'' follows, which happens if and only if an equation of the form $l\theta'_1 = k\theta''_1$ (for some non-trivial pair of coprime integers (k, l)) is satisfied. In this case, l is the exact number of times q'' turns around p''_1 while q' turns k times around p'_1 . □




Figure 12: The modified cone of angle $\theta > \pi$.

When θ'_1 (resp. θ''_1) is greater than π , we cannot cut the cone of angle θ'_1 (resp. θ''_1) in the way it has been done previously. We let the reader verify that one only has to replace the truncated cone of angle θ'_1 at p'_1 (resp. of angle θ''_1 at p''_1) by the metric space obtained by gluing the sides a and b on Figure 12. This metric space is the cone of angle θ'_1 (resp. θ''_1) to which one has added the triangle T appearing in grey on Figure 12 with the aforementioned identifications.

Regarding the lack of injectivity of the map $\widetilde{\mathcal{S}}_3$ (or \mathcal{S}_3) when some elements of U have non-trivial isometries, one can make a statement similar to Remark 6.2 to address the question.

REMARK 6.5. In contrast with Thurston's surgery, there might be **essentially different** ways to perform Devil's surgery when $\text{Im}(\rho')$ is finite. What we mean by that is the following (where we continue to use the notations introduced above): given a base surface N , a target moduli space $\mathcal{F}_{[\widehat{\rho}]}$ in which one lands when performing Devil's surgery on N relatively to two cone points p'_1 and p''_1 on it, different initial choices of pairs (q', q'') might yield surgery maps \mathcal{S}_3 whose image are pairwise distinct open subsets of $\mathcal{F}_{[\widehat{\rho}]}$, a fact which might be a bit surprising at first glance.

This curiosity comes from the fact that the surfaces obtained via Devil's surgery which belong to the same $\mathcal{F}_{[\widehat{\rho}]}$ are exactly the ones corresponding to admissible pairs $(q', q'') \in C_{\theta'_1}(\epsilon) \times C_{\theta''_1}(\epsilon)$ such that $\widehat{\rho}(\widehat{\beta}) \in \text{Im}(\rho)$. The solutions to this equation is a subsurface of $C_{\theta'_1}(\epsilon) \times C_{\theta''_1}(\epsilon)$ which might have more than one connected component. These connected components **are** the different ways to perform Devil's surgery and yield different open subsets in $\mathcal{F}_{[\widehat{\rho}]}$. We carry an explicit analysis of this phenomenon in the paragraph 'Cone points' of §10.5 which, even if it concerns only a particular case, should allow the reader to understand in full generality the 'curiosity' we are talking about here.

6.4 The Kite Surgery \mathcal{S}_4 . Fix $\theta_1 \in]2\pi, 4\pi[$ and $\theta_2, \theta_3 \in]0, 2\pi[$ such that

$$(2\pi - \theta_1) + (2\pi - \theta_2) + (2\pi - \theta_3) = 0.$$

We describe in this section a local surgery building from a regular flat torus a new one with three conical points of respective angles θ_1, θ_2 and θ_3 .

Let T be a regular flat torus (*i.e.* without singular points for the flat metric), and p_2 a point on T . If p_3 is a point close enough to p_2 , there exists a unique kite contained in T with opposite vertices p_2 and p_3 and such that the external angle at these points are θ_2 and θ_3 respectively. The external angles at the two other vertices of the kite are necessarily equal to the half of θ_1 , see Figure 13.

The **kite surgery** consists in removing the kite (in grey on the above picture) and gluing the adjacent sides in order to get three singular points of respective angles θ_1, θ_2 and θ_3 . Since the flat torus T is determined by a lattice in \mathbb{C} , *i.e.* by two \mathbb{R} -linearly independent complex numbers z_1, z_2 such that $T = \mathbb{C}/(\mathbb{Z}z_1 + \mathbb{Z}z_2)$, one can perform the kite surgery on T by placing p_2 at 0 and p_3 at z_0 for any given z_0 sufficiently small.

Up to renormalisation, we can assume that $z_1 = 1$ and $z_2 = \tau \in \mathbb{H}$. Let T' be the resulting torus (more precisely the class of tori up to renormalisation by an element of \mathbb{C}). Let $\mathcal{F}_{[\rho']}$ be the leaf to which this surgery sends T' . Because z_0 and $-z_0$ are equivalent under the action of the hyperelliptic involution, the kite surgery performed at these two parameters gives isometric surfaces. As in the previous cases, one can build a map

$$\mathcal{S}_4 : C_\pi^*(\epsilon) \times \mathbb{H} \longrightarrow \mathcal{F}_{[\rho]}$$

$$(z_0, \tau) \longmapsto T'.$$

which is a local biholomorphism onto its image, except at points having exceptional symmetries where it is an orbifold covering onto its image. The proof is similar to the proof of Proposition 6.1, one just has to take a suitable topological model for T that makes (z_0, z_1, z_2) a linear parametrisation of $\mathcal{F}_{[\rho]}$.




Figure 13: Performing the kite surgery.

6.5 The Surgery \mathcal{S}_5 : Building Flat Surfaces with a Euclidean Cylinder.

In this section we do not make any assumptions on the cone angles $\theta_1, \dots, \theta_n$. We explain a simple surgery (to which we shall refer as \mathcal{S}_5) building flat surfaces of genus $g + 1$ and with $n - 1$ cone points having an arbitrarily long Euclidean cylinder out of an initial flat surface N of genus g and with n cone points, the new cone point being of angle larger than 2π .

Let γ be a geodesic path joining p_1 and p_2 , the two conical points of N of respective angles θ_1 and θ_2 . Cut along γ to get a flat surface with one boundary component and then glue p_1 and p_2 together. The resulting surface has a boundary consisting of two simple closed geodesics touching at one point where they are singular. Then we glue a flat cylinder along these two boundary components to get a new flat surface N' of genus $g + 1$ with an embedded cylinder (see Figure 14). Note that the cone angle of N' at its new singular point (namely the one obtained after having identified p_1 and p_2) is easily seen to be $\theta_1 + \theta_2 + 2\pi$.

There are two real parameters for the aforementioned gluing of the flat cylinder: the length of the latter and a twisting parameter (starting from one given way to glue a given cylinder, one can compose the gluing function at one of its extremities by a rotation, the angle of this rotation being the twisting parameter). Both can be encoded by a single complex number z_0 whose imaginary part is positive, such that the cylinder we glue identifies with the one of base 1 and height z_0 . This makes sense because there always exists a normalization of N such that the geodesic joining p_1 and p_2 in the initial surface develops onto the segment $[0, 1]$.

If U is a neighbourhood of the initial surface in $\mathcal{F}_{[\rho]}$, we can build a natural map

$$\begin{aligned} \mathcal{S}_5 : \mathbb{A} \times U &\longrightarrow \mathcal{F}_{[\rho]} \\ (z_0, N) &\longmapsto N', \end{aligned}$$

where \mathbb{A} stands for the infinite cylinder $\mathbb{H}/(z \sim z + 1)$.

The surgery \mathcal{S}_5 associates to N and z_0 the surface obtained after gluing the flat cylinder of height the parameter z_0 to (the good normalisation of) N .

PROPOSITION 6.6. *The surgery \mathcal{S}_5 is a local biholomorphism onto its image.*

Proof. Let (z_1, \dots, z_m) be a linear parametrisation of U such that z_1 parametrises the geodesic segment along which the surgery is performed. Then one verifies that $(z_0 z_1, z_1, \dots, z_m)$ is a linear parametrisation of the image of \mathcal{S}_5 . Indeed, one needs to rescale the cylinder of base and height $(1, z_0)$ to $(z_1, z_1 z_0)$ so it fits the segment parametrised by z_1 onto which it is glued. The rest of the proof works in the same way as in Proposition 6.1. \square

6.6 Calculation of the Signature of the Area form in Particular Cases.

A corollary of the description of these surgeries is an easy inductive computation of the signature of Veech's form in the specific cases we are interested in. In order to perform this computation, we will need the following definition:




Figure 14: The surgery \mathcal{S}_5 performed on a flat sphere with four cone points: (1) we cut along the geodesic segment between two of them; (2) we identify the two corresponding points on the boundary; (3) then we glue a flat cylinder in order to obtain a flat tori with three cone points.

DEFINITION 6.1. A unitary character $\tilde{\rho} \in H^1(M, \mathbb{U})$ is a **reduction** of $\rho \in H^1(N, \mathbb{U})$ if there exists an injective diffeomorphism $i : M \rightarrow N$ such that $\tilde{\rho} = i^* \rho$.

In particular, the linear holonomy of the base surface of one of the surgeries we have described is a reduction of the linear holonomy of the surface obtained by one of these surgeries. We have the following:

PROPOSITION 6.7. (1) Suppose that $g = 0$, $n \geq 3$ and that $0 < \theta_i < 2\pi$ for all $i = 1, \dots, n$. Then Veech's area form has signature $(1, n - 3)$.
 (2) Suppose that $g = 1$, $n \geq 2$, $2\pi < \theta_1 < 4\pi$ and that $0 < \theta_i < 2\pi$ for all i such that $1 < i \leq n$. Then Veech's area form has signature $(1, n - 1)$.

Proof. The proof goes by induction in both cases. We explain only (2) since the proof of (1) is basically the same but simpler and roughly sketched in [Thu98].

We suppose that $g = 1$, $n \geq 3$ and $\rho \in H^1(N, \mathbb{U}, \theta)$. Let ρ' be the reduction (see Section 7) associated to a collision between two points of angles θ_i and θ_j . Such a collision can actually happen if and only if $(2\pi - \theta_i) + (2\pi - \theta_j) < 2\pi$, and two such points always exist since there is initially only one conical point of negative curvature which is also smaller than 4π , provided that $n \geq 3$. Therefore the leaf associated to ρ' is not empty and one can perform Thurston's surgery on elements of $\mathcal{F}_{[\rho']}$ to get elements of the n -dimensional leaf $\mathcal{F}_{[\rho]}$. Let $(z_0, z_1, \dots, z_{n-1})$ be a linear parametrisation of $\mathcal{F}_{[\rho]}$ such that (z_1, \dots, z_{n-1}) is a linear parametrisation of $\mathcal{F}_{[\rho']}$ and z_0 is such that $(z_0, z_1, \dots, z_{n-1})$ represents the element that one gets after performing Thurston's surgery with parameter z_0 on the surface represented by (z_1, \dots, z_{n-1}) . If A is the area form for the parametrisation $(z_0, z_1, \dots, z_{n-1})$ and A' the one for (z_1, \dots, z_{n-1}) , then one has

$$A(z_0, \dots, z_{n-1}) = A'(z_1, \dots, z_{n-1}) - \mu|z_0|^2$$

for a certain positive constant μ which depends only on the value of the angle of the cone point on which the surgery is performed. Indeed, $\mu|z_0|^2$ is the area of the portion of cone removed when proceeding to the surgery which is an isosceles triangle with fixed centre angle and whose base has length linearly depending on $|z_0|$. The induction hypothesis ensures that A' has signature $(1, n - 2)$ and therefore A has signature $(1, n - 1)$.

The case $n = 2$ remains to be handled. With a similar argument, but using Devil's surgery instead of Thurston's one, we find that in that case the signature is $(1, 1)$. One starts with a flat sphere with three cone points. The set of such flat spheres can be parametrised by a complex number z_1 such that the area of the associated sphere is $|z_1|^2$ (and recovering that up to projectivising, there is only one such flat sphere). Devil's surgery consists in removing two portions of cones in a sphere. If z_0 is the parameter of the surgery, the area of the resulting torus will be of the form

$$|z_1|^2 - \mu|z_0|^2$$

for a certain $\mu > 0$. This completes the proof of the proposition. \square

6.7 Cone Angle Around a Codimension 1 Stratum. As we will see in detail later, the surgery maps \mathcal{S}_i (for $i = 1, \dots, 4$) describe the cone-manifold structure of the metric completion of $\mathcal{F}_{[\rho]}$ close to a codimension 1 stratum, when Veech's area form endows $\mathcal{F}_{[\rho]}$ with a complex hyperbolic structure. In particular, they allow the computation of the associated cone-manifold angles.

- (1) In case of both Thurston's surgeries \mathcal{S}_1 of \mathcal{S}_2 , the cone-manifold angle around the codimension 1 stratum is the angle of the Euclidean cone angle on which the surgery is performed.
- (2) In the case of Devil's surgery, when both angles are rational multiples of 2π , say $2\pi m'/M$ and $2\pi m''/M$, the cone-manifold angle is $2\pi \text{lcm}(m', m'')/M$.

- (3) In the case of the Kite surgery \mathcal{S}_4 , the cone angle always equals π since the parameter space for the surgery is the neighbourhood of a regular point of angle 2π on which the hyperelliptic involution acts.

6.8 Geometric Convergence. We end this section dedicated to surgeries by a paragraph on a notion of geometric convergence for flat surfaces. Whether two Riemannian manifolds (in a moduli space) are close or not depends on an *a priori* definition.

DEFINITION 6.2. Let N' be a flat surface of area 1 obtained from a surgery \mathcal{S} on a surface N . The **width** of this surgery at N is:

- the distance between the two new cone points in N' if $\mathcal{S} \in \{\mathcal{S}_1, \mathcal{S}_2\}$;
- the length of the short essential curve created on N' if $\mathcal{S} = \mathcal{S}_3$;
- the distance between p_2 and p_3 in N' (see §6.4 for the notations) if $\mathcal{S} = \mathcal{S}_4$.

The width of a surgery is a positive parameter whose square depends linearly on the Euclidean area of the removed part of the initial surface on which one performs the surgery. When g and θ are such that $\mathcal{F}_{[\rho]}$ has a complex hyperbolic structure, this width has a geometric interpretation in terms of the distance to the strata of the metric completion which will be made explicit in the next section (see Lemma 7.9).

DEFINITION 6.3 (**Geometric convergence**). A sequence of flat surfaces $(M_\ell)_{\ell \in \mathbb{N}} \in (\mathcal{F}_{[\rho]})^{\mathbb{N}}$ is said to be **geometrically converging** either if it converges in $\mathcal{F}_{[\rho]}$ or if there exist

- a flat surface M_∞ belonging to a leaf $\mathcal{F}_{[\rho']}$ for a reduction ρ' of ρ ;
- a small neighbourhood U of M_∞ in $\mathcal{F}_{[\rho']}$ on which a surgery map $\mathcal{S} = \mathcal{S}_i$ (for some $i = 1, 2, 3, 4$) is well defined;
- $(\epsilon_\ell)_{\ell \in \mathbb{N}}$ a sequence of positive numbers going to 0;
- a truncated sequence $(M'_\ell)_{\ell \gg 1}$ of flat surfaces elements of U which converge to M_∞ in $U \subset \mathcal{F}_{[\rho']}$,

such that M_n is obtained after a surgery \mathcal{S} of width ϵ_ℓ on M'_ℓ for $\ell \gg 1$.

We say that the sequence $(M_\ell)_{\ell \in \mathbb{N}}$ **geometrically converges** to the pair (M_∞, \mathcal{S}) , or just to M_∞ if $(M_\ell)_{\ell \in \mathbb{N}}$ converges in $\mathcal{F}_{[\rho]}$.

The true interest of this definition is that it will allow us to make the difference between two sequences of surfaces in $\mathcal{F}_{[\rho]}$ whose limits are isometric metric spaces but lying at different places in (the metric completion of) $\mathcal{F}_{[\rho]}$.

7 The Metric Completion

In this section, N stands for a surface of genus $g = 0$ with $n + 3$ marked points or of genus 1 with $n + 1$ marked points (we write the number of marked points this way in order that $\mathcal{F}_{[\rho]}$ be of complex dimension n in these two cases).

We also make the assumption that

- if $g = 0$, then all the cone angles $\theta_1, \dots, \theta_{n+3}$ belong to $]0, 2\pi[$;
- if $g = 1$ then $\theta_1 \in]2\pi, 4\pi[$ and all the other cone angles θ_i are in $]0, 2\pi[$.

Therefore, for any $\rho \in H^1(N, \mathbb{U}, \theta)$ in the image of the linear holonomy map (5), the leaf $\mathcal{F}_{[\rho]}$ of Veech's foliation in the corresponding moduli space of marked curves is endowed with a complex hyperbolic structure of dimension n (see [Thu98] for the case $g = 0$ and Proposition 6.7 or [Vee93] for the case $g = 1$).

In the present section, we are interested in the structure of the metric completion of $\mathcal{F}_{[\rho]}$ endowed with this complex hyperbolic structure. Therefore every mention of a geometric property of $\mathcal{F}_{[\rho]}$ will now be relative to this structure.

Our goal here is to prove the following theorem:

Theorem 7.1. *Let X be the metric completion of $\mathcal{F}_{[\rho]}$. If $\text{Im}(\rho)$ is finite then:*

- (1) X has a stratified structure $X = X_0 \sqcup X_1 \sqcup \dots \sqcup X_n$ with $X_0 = \mathcal{F}_{[\rho]}$;
- (2) the topological closure of X_i in X is $X_i \sqcup X_{i+1} \sqcup \dots \sqcup X_n$;
- (3) for $i = 0, \dots, n$, X_i is a smooth complex hyperbolic manifold of complex dimension $n - i$ which carries a natural $\mathbb{C}H^{n-i}$ -structure;
- (4) each X_i is a finite union of finite covers of $\mathcal{F}_{[\tilde{\rho}]}$ for some reductions $\tilde{\rho}$ of ρ .

From now on, we assume that $\text{Im}(\rho)$ is finite.

7.1 Strata. In Section 6, we have introduced various surgeries, describing different ways flat surfaces can degenerate and how to parametrise these surgeries. The degenerate flat surfaces we see appearing in these ways belong to the metric completion of $X_0 = \mathcal{F}_{[\rho]}$. More precisely, they appear in copies of (finite coverings of) $\mathcal{F}_{[\tilde{\rho}]}$ of complex dimension $n - 1$ where $\tilde{\rho}$ is a reduction of ρ in the sense of Definition 6.1.

We define inductively the strata X_1, \dots, X_n appearing in the description of X . The first stratum X_1 is the set of pairs surface/surgery (N, \mathcal{S}) such that there exists a sequence of elements of $X_0 = \mathcal{F}_{[\rho]}$ geometrically converging (in the sense of Definition 6.3) to a pair (N, \mathcal{S}) not already in X_0 .

LEMMA 7.2. *The stratum X_1 is a union of (finite covers of) $\mathcal{F}_{[\tilde{\rho}]}$ s for some reductions $\tilde{\rho}$ of ρ and such that the complex dimension of $\mathcal{F}_{[\tilde{\rho}]}$ is $n - 1$.*

Proof. The holonomy $\tilde{\rho}$ of the limit of a sequence geometrically converging is a reduction of ρ . Since $\text{Im}(\rho)$ is finite, there are only finitely many such reductions and therefore only finitely $\mathcal{F}_{[\tilde{\rho}]}$ to which such a limit can belong.

Now let $\mathcal{F}_{[\tilde{\rho}]}$ be such that one of its elements (N_0, \mathcal{S}_0) appears as the limit of a geometrically converging sequence. We claim that Z , the connected component of X_1 to which (N_0, \mathcal{S}_0) belongs, is a finite cover of a connected component of $\mathcal{F}_{[\tilde{\rho}]}$. One can define a local homeomorphism from a neighbourhood of (N_0, \mathcal{S}_0) in Z to the component of $\mathcal{F}_{[\tilde{\rho}]}$ containing N_0 associating to any pair $(N, \mathcal{S}_0) \in Z$ sufficiently close to (N_0, \mathcal{S}_0) the associated flat surface N . This map is a covering map which is finite since the fiber over a point is included in the set of different ways to perform the corresponding surgery (see Remark 6.5), which is finite. This fiber is actually trivial for surgeries different from Devil's surgery.

More precisely let $\tilde{\rho}$, be such that one element N_0 of $\mathcal{F}_{[\tilde{\rho}]}$ appears as the limit of a geometrically converging sequence along a surgery \mathcal{S} . The set of elements of X_1 whose associated flat surface belongs to $\mathcal{F}_{[\tilde{\rho}]}$ is exactly the covering map over $\mathcal{F}_{[\tilde{\rho}]}$ whose fiber over a point is the number of ways to perform the surgery \mathcal{S} at this point. \square

The next step is to prove that the disjoint union $X_0 \sqcup X_1$ embeds into the metric completion of $X_0 = \mathcal{F}_{[\rho]}$. This is a direct consequence of the following proposition:

PROPOSITION 7.3. *If $(M_\ell)_{\ell \in \mathbb{N}} \in (\mathcal{F}_{[\rho]})^{\mathbb{N}}$ converges geometrically, then it is a Cauchy sequence for the metric induced by the complex hyperbolic structure on $\mathcal{F}_{[\rho]}$.*

Proof. The statement is clear if $(M_\ell)_{\ell \in \mathbb{N}}$ converges in $\mathcal{F}_{[\rho]}$ so we assume that it is not the case. We consider a linear parametrisation (z_0, z_1, \dots, z_n) such that

- z_0 is the surgery parameter (*i.e.* the small segment linking the two new cone points after Thurston’s surgery, or the small closed broken geodesic segment appearing after a Devil’s surgery which is such that the width of the surgery is $|z_0|$ in a normalisation of area 1, etc.); and
- (z_1, \dots, z_n) is a linear parametrisation of the leaf to which the surface on which the surgery is done belongs.

For any $\ell \in \mathbb{N}$, let $z(\ell) = (z_i(\ell))_{i=0}^n$ stand for the coordinates of M_ℓ in the considered linear parametrisation, normalised so that the corresponding area of M_ℓ is 1. From the very definition of metric convergence, we have that $(z_0(\ell))_{\ell \in \mathbb{N}}$ goes to 0 and $((z_i(\ell))_{i=1}^n)_{\ell \in \mathbb{N}}$ converges in \mathbb{C}^n since the associated sequence of flat surfaces in the corresponding stratum converges. It follows that $z(\ell)$ converges in \mathbb{C}^{n+1} as ℓ tends to infinity.

On the other hand, the normalisation of the areas of the M_ℓ ’s ensures that the $z(\ell)$ ’s stay away from the boundary of the model of the complex hyperbolic space associated with the considered linear parametrisation (z_0, \dots, z_n) . The proposition follows. \square

Thanks to the preceding result, one has a map:

$$i_1 : X_0 \sqcup X_1 \longrightarrow \overline{\mathcal{F}_{[\rho]}}.$$

PROPOSITION 7.4. *The map i_1 defined just above is injective.*

Proof. Clearly, the restriction of i_1 to $X_0 = \mathcal{F}_{[\rho]}$ is the identity. Since $i_1(X_1) \cap \mathcal{F}_{[\rho]} = \emptyset$, it suffices to show that $i_1|_{X_1}$ is injective to get the proposition.

Let (A_1, \mathcal{S}_1) and (A_2, \mathcal{S}_2) be two distinct points in X_1 . The surgery maps

$$\begin{aligned} \mathcal{S}_i : C_{\theta_1}^*(\epsilon_i) \times U_i &\longrightarrow \mathcal{F}_{[\rho]} \\ (z_0, N) &\longmapsto N', \end{aligned}$$

defined in Section 6, where U_i is a neighbourhood of A_i for $i = 1, 2$, extend continuously to $\overline{\mathcal{S}_i} : C_{\theta_1}(\epsilon_i) \times U_i \longrightarrow \overline{\mathcal{F}_{[\rho]}}$ which are homeomorphisms onto their images

and whose respective images are neighbourhoods of $i_1(A_1)$ and $i_1(A_2)$ in X . If U_i and ϵ_i are chosen small enough, the images of \mathcal{S}_1 and \mathcal{S}_2 do not overlap which implies that $i(A_1)$ and $i(A_2)$ are separated and therefore different. \square

The distance induced on X_1 by this embedding is nothing else but the one induced by its natural complex hyperbolic structure: a neighbourhood of X_1 in $X = \overline{\mathcal{F}_{[\rho]}}$ can be described by a finite number of linear parametrisations of the form (z_0, \dots, z_n) in which X_1 corresponds to the locus $\{z_0 = 0\}$.

REMARK 7.5. We would like to stress that in the preceding assertion, one has to be aware that (z_0, \dots, z_n) does not induce a local system of coordinates on a neighbourhood in X of a small open subset of X_1 . Actually, what must be understood is that the equation $z_0 = 0$ cuts out something (a piece of X_1 as it happens) in the boundary of the definition domain of the chart induced by the linear parametrisation (z_0, \dots, z_n) . We will not dwell again on this subtlety in what follows but will only make reference to the present remark.

The definition of X_2 is slightly more subtle because it is possible that two essentially different geometrically converging sequences in X_1 converge to the same point in $\overline{\mathcal{F}_{[\rho]}}$: consider for instance a case when $g = 1$ and $n = 3$. We can distinguish two types of components in X_1 : the ones which are moduli spaces of tori with two cone points, and those which are Thurston–Deligne–Mostow’s moduli spaces of flat spheres with four cone points. Both can degenerate on flat spheres with three cone points. It can happen that these a priori different limits are identified in $\overline{\mathcal{F}_{[\rho]}}$. One must think of such points as parts of the intersection locus of the closures of two connected components of X_1 on which two different surgeries can be performed, each leading to a different component of X_1 .

In order to define correctly X_2 , we proceed in two steps: we first define in an analogous way Y_2 , that one shall think to be roughly the set of pairs flat surface/surgery (M, \mathcal{S}) such that there exists a sequence in X_1 geometrically converging to (M, \mathcal{S}) . However, the fact that we are dealing with finite covers of leaves prevents us from giving such a straightforward definition. Bypassing this difficulty is rather easy: X_1 is a finite union of finite covers of some leaves $\mathcal{F}_{[\tilde{\rho}]}$, for some reductions $\tilde{\rho}$ of ρ . Such a finite cover Z_0 can be partially metrically completed by adjoining a codimension 1 stratum Z_1 (possibly with several connected components) in order that the covering map $\pi : Z_0 \rightarrow \mathcal{F}_{[\tilde{\rho}]}$ extends to a map

$$\tilde{\pi} : Z_0 \sqcup Z_1 \rightarrow \mathcal{F}_{[\tilde{\rho}]} \sqcup X'_1$$

which is a covering map, possibly ramified along X'_1 , where X'_1 is analogous to X_1 associated to $\mathcal{F}_{[\tilde{\rho}]}$ and such that $\tilde{\pi}^{-1}(X'_1) = Z_1$: X'_1 is a finite union of some unramified finite covers of some leaves $\mathcal{F}_{[\tilde{\rho}']}$ for some reduction $\tilde{\rho}'$ of $\tilde{\rho}$. Then one defines Y_2 as the (finite) union of all such Z_1 ’s associated to all the finite covers appearing in X_1 and is itself a finite union of finite covers of some leaves $\mathcal{F}_{[\tilde{\rho}]}$ for some reduction $\tilde{\rho}$ of ρ (a reduction of a reduction of ρ is again a reduction of ρ as it follows immediately from Definition 6.1).

Defined this way, $X_1 \sqcup Y_2$ maps into the metric completion of X_1 as a complex hyperbolic manifold with each of its connected components endowed with the induced complex hyperbolic distance. But note that this construction does not take into account how close these connected components can be in $\overline{\mathcal{F}_{[\rho]}}$.

As we did for X_1 , we define a map:

$$i_2 : X_1 \sqcup Y_2 \longrightarrow \overline{\mathcal{F}_{[\rho]}} \tag{12}$$

whose restriction to X_1 coincides with the one of i_1 . The main difference between the maps i_1 and i_2 is that the latter is not injective: some components of Y_2 are identified. We define X_2 as the image of Y_2 under that map or equivalently, Y_2 with the aforementioned components identified. The crucial point is that X_0 has a distance only defined on each of its connected components by the complex hyperbolic metric, while the distance on X_1 takes into account the way in which X_1 is embedded in X .

The following property allows to identify the irreducible components of X_2 with some unramified coverings of some reductions of $\mathcal{F}_{[\rho]}$.

PROPOSITION 7.6. *If A and B are distinct points of Y_2 such that $i_2(A) = i_2(B)$ then*

- (1) *the flat surfaces associated to A and B are isometric;*
- (2) *there exists $\widehat{\rho}$, a reduction of ρ , and a connected component Z of $\mathcal{F}_{[\widehat{\rho}]}$ such that the images by i_2 of the components of Y_2 containing A and B are both equal to the same finite cover of Z .*

Proof. Let N_A and N_B be the two flat surfaces associated to A and B respectively.

Assume that N_A and N_B are not isometric. Consider two sequences in X_1 geometrically converging to A and B respectively. They can be approximated by two Cauchy sequences in $\mathcal{F}_{[\rho]}$ converging to the same point in $\overline{\mathcal{F}_{[\rho]}}$. But their associated flat surfaces converge towards two different metric spaces, which is impossible. This proves (1).

Let Z be the component of $\mathcal{F}_{[\widehat{\rho}]}$ to which $N = N_A = N_B$ belongs. We first remark that there are neighbourhoods of A and B in Y_2 which are identified under i_2 : there are two surgeries \mathcal{S}_A and \mathcal{S}_B on N which produce images under i_2 of neighbourhoods of A and B in Y_2 . This identification can be extended to a cover of Z . This cover must be finite since it is covered by a component of Y_2 which is finite according to Lemma 7.2. □

From the map (12) and by the very definition of X_2 , one deduces an injective map $j_2 : X_1 \sqcup X_2 \rightarrow \overline{\mathcal{F}_{[\rho]}}$. Since the restrictions of i_1 and j_2 to X_1 coincide, one can consider their fiber product over X_1 in order to get an injective map:

$$X_0 \sqcup X_1 \sqcup X_2 \longrightarrow \overline{\mathcal{F}_{[\rho]}}$$

Inductively, one defines Y_{j+1} from X_j in exactly the same way we defined Y_2 from X_1 . Then one defines X_{j+1} by identifying some components of Y_{j+1} using the natural map $i_{j+1} : X_j \sqcup Y_{j+1} \longrightarrow \overline{\mathcal{F}_{[\rho]}}$. Note at this point that the analog of Proposition 7.6

for Y_{j+1} holds true, the proof being completely similar. Since a reduction $\widehat{\rho}$ of a reduction $\widetilde{\rho}$ of ρ is still a reduction of ρ , we get that X_i is a complex hyperbolic manifold of dimension $n - i$ whose connected components are some coverings of some leaves $\mathcal{F}_{[\widetilde{\rho}]}$ for some reduction $\widetilde{\rho}$ of ρ .

Putting all pieces together we get that

- $X_0 \sqcup X_1 \sqcup \dots \sqcup X_n$ embeds into the metric completion X of $X_0 = \mathcal{F}_{[\rho]}$;
- $\forall i$, X_i is a finite union of finite covers of $\mathcal{F}_{\widetilde{\rho}}$ for some reductions $\widetilde{\rho}$ of ρ ;
- $\forall i$, the distance of X induces on X_i its natural structure of complex hyperbolic manifold of dimension $\dim(X_0) - i$.

7.2 Proof of the Surjectivity.

PROPOSITION 7.7. *Assume that $\text{Im}(\rho)$ is finite. Then the embedding*

$$X_0 \sqcup X_1 \sqcup \dots \sqcup X_n \longrightarrow X = \overline{\mathcal{F}_{[\rho]}}$$

is onto.

This proposition says in substance that the metric space obtained by adding to $\mathcal{F}_{[\rho]}$ the degenerate surfaces that one sees when reversing the surgeries studied in Section 6 is complete. Before giving the proof, we have to state two technical lemmas relating the complex hyperbolic geometry of $\mathcal{F}_{[\rho]}$ to the geometry of the underlying flat surfaces parametrized by this leaf.

LEMMA 7.8. *If $\text{Im}(\rho)$ is finite then the two following assertions hold true for any sequence $(M_\ell)_{\ell \in \mathbb{N}}$ of flat surfaces in $\mathcal{F}_{[\rho]}$ normalised so that their area is 1:*

- (1) *if $(M_\ell)_{\ell \in \mathbb{N}}$ is a Cauchy sequence then $(D(M_\ell))_{\ell \in \mathbb{N}}$ is bounded;*
- (2) *if $(D(M_\ell))_{\ell \in \mathbb{N}}$ is bounded then $(M_\ell)_{\ell \in \mathbb{N}}$ is a Cauchy sequence (up to passing to a subsequence).*

Proof. We postpone the proof of (1) to Section 8 in which we provide a description of the parts of $\mathcal{F}_{[\rho]}$ on which the diameter function D is large.

The proof of (2) consists in remarking that using the Delaunay decomposition of M_ℓ , we can assume that, up to passing to a subsequence:

- all the M_ℓ 's are recovered by gluing the sides of a pseudo-polygon through the same gluing pattern;
- the side glued together always form the same angle (that this can be assumed follows from the fact that $\text{Im}(\rho)$ is finite by assumption);
- the lengths of each side converge (since the lengths of the edges of the Delaunay triangulation are smaller than $2D(M_\ell)$ by Proposition 5.2).

In the chart defined by the gluing pattern, the coordinates of the M_ℓ 's form a Cauchy sequence. Then using the fact that their areas all have been assumed to be 1, one can argue in the same way as at the end of the proof of Proposition 7.3 and get that $(M_\ell)_{\ell \in \mathbb{N}}$ is a Cauchy sequence for the metric on $\mathcal{F}_{[\rho]}$ induced by the complex hyperbolic structure it carries. \square

LEMMA 7.9. *There exists a positive constant $K = K_{[\rho]}$ such that if $M \in X_0 = \mathcal{F}_{[\rho]}$ (which is supposed to be normalised such that its area equals 1) is obtained from a surgery \mathcal{S} of width ϵ from an element of X_1 then*

$$d(M, X_1) \leq K \epsilon,$$

where d denotes the extension of the complex hyperbolic distance on X_0 to X .

Proof. Let (z_0, z_1, \dots, z_n) be a linear parametrisation compatible with the surgery \mathcal{S} (see Section 6) which is such that

- the parameter z_0 is the surgery parameter, in particular $|z_0| = \epsilon$ is the width of the surgery;
- (z_1, \dots, z_n) is a linear parametrisation of $U \subset X_1$;
- in the coordinates z_0, \dots, z_n , the area form A writes down

$$A(z_0, z_1, \dots, z_n) = A'(z_1, \dots, z_n) - \mu|z_0|^2$$

where

- $\mu = \mu_{\mathcal{S}}$ is a positive real constant depending on the surgery \mathcal{S} (it is the constant such that $\mu|z_0|^2$ is the area of the part of the surface removed while processing the surgery);
- A' is the area form on U expressed in the coordinates z_1, \dots, z_n .

(Note that since the image of ρ is assumed to be finite, the set of such $\mu_{\mathcal{S}}$'s is finite and thus $\mu_{\mathcal{S}}$ is uniformly bounded from above).

One can compute the complex hyperbolic distance between two points in the complex hyperbolic space using formulas involving A (see [Gol99, p.77] for instance). If $a : \mathbb{C}^{n+1} \times \mathbb{C}^{n+1} \rightarrow \mathbb{C}$ stands for the polarisation of A , namely the Hermitian form such that $A(X) = a(X, X)$ for every $X \in \mathbb{C}^{n+1}$, the complex hyperbolic distance $d(X, Y)$ between two points $[X], [Y] \in \mathbb{C}\mathbb{H}^n \subset \mathbb{C}\mathbb{P}^n$ satisfies

$$\cosh^2 \left(\frac{d(X, Y)}{2} \right) = \frac{a(X, Y)a(Y, X)}{a(X, X)a(Y, Y)}.$$

If $X = (x_0, X') \in \mathbb{C}^{n+1}$ and $Y = (y_0, Y') \in \mathbb{C}^{n+1}$ with $X' = (x_1, \dots, x_n)$ and $Y' = (y_1, \dots, y_n)$ in \mathbb{C}^n , the formula for $a(X, Y)$ is

$$a(X, Y) = a'(X', Y') - \mu x_0 \overline{y_0}$$

where a' stands for the polarisation of A' .

We want to estimate $\alpha = d((z_0, z_1, \dots, z_n), (0, z_1, \dots, z_n))$. Since M is supposed to have area 1, it follows from the discussion above that we have

$$\cosh^2(\alpha/2) = \frac{(1 + \mu \epsilon^2)^2}{1 \cdot (1 + \mu \epsilon^2)} = 1 + \mu \epsilon^2.$$

Since for all $a > 0$, one has $1 + a^2/2 \leq \cosh(a)$, it comes

$$1 + \frac{\alpha^2}{8} \leq \sqrt{1 + \mu \epsilon^2} \leq 1 + \frac{\mu}{2} \epsilon^2$$

from which we deduce that $\alpha \leq 2\sqrt{\mu}\epsilon$. Since μ is bounded from above by a constant only depending on the image of ρ , the proposition is proved. \square

We end this section with the proof of Proposition 7.7. We still suppose that $\text{Im}(\rho)$ is finite, which is the crucial hypothesis on which everything done in this paper relies on. We just say a word on the general strategy. In order to show that any Cauchy sequence accumulates to one point in a stratum, we first prove that, since the diameter along a Cauchy sequence is bounded, if such a Cauchy sequence does not converge in $\mathcal{F}_{[\rho]}$, it implies that it degenerates in the sense that either its systole or its relative systole goes to zero. If the latter occurs, we show that such a surface having a sufficiently short systole or relative systole is obtained from one of the four surgeries described in Section 6 and therefore is very close to X_1 . Then we conclude with an inductive argument.

Proof of Proposition 7.7. Let $M_\bullet = (M_\ell)_{\ell \in \mathbb{N}}$ be a Cauchy sequence in $X_0 = \mathcal{F}_{[\rho]}$ for the complex hyperbolic metric. For the remainder of the proof, we set

$$D_\ell = D(M_\ell), \quad \sigma_\ell = \sigma(M_\ell) \quad \text{and} \quad \delta_\ell = \delta(M_\ell)$$

for any $\ell \in \mathbb{N}$. We aim at proving that M_\bullet converges in $\overline{\mathcal{F}_{[\rho]}}$ to a point belonging to the image of the embedding $X_0 \sqcup \dots \sqcup X_n \longrightarrow \overline{\mathcal{F}_{[\rho]}}$.

The proof goes by induction on $\dim(X_0)$. Recall that the following inequalities hold true for any ℓ (see Proposition 5.1):

$$\delta_\ell \leq D_\ell \quad \text{and} \quad \sigma_\ell \leq 2D_\ell.$$

We distinguish three cases:

(1). **Both the two sequences $(\delta_\ell)_{\ell \in \mathbb{N}}$ and $(\sigma_\ell)_{\ell \in \mathbb{N}}$ do not converge to zero.** According to Lemma 7.8, the sequence of diameters $(D_\ell)_{\ell \in \mathbb{N}}$ is bounded. The Delaunay decomposition provides polygonal models of M_ℓ such that the length of each side is bounded (see Proposition 5.2). One can extract a subsequence such that all polygonal models have the same gluing pattern, and therefore, since $\text{Im}(\rho)$ is finite, extract a subsequence whose polygonal model converges towards a non degenerate pseudo-polygon whose associated surface in $X_0 = \mathcal{F}_{[\rho]}$ is the limit of the Cauchy sequence M_\bullet .

(2). **The sequence $(\delta_\ell)_{\ell \in \mathbb{N}}$ converges to zero while $(\sigma_\ell)_{\ell \in \mathbb{N}}$ does not.** In that case, one proves that $(d(M_\ell, X_1))_{\ell \in \mathbb{N}}$ converges to zero. First remark that necessarily $\dim(X_0) \geq 2$ in that case. Indeed, according to Proposition 5.9, if $\dim(X_0) = 1$ we have that δ_ℓ converging to zero implies that D_ℓ goes to infinity which would contradict Lemma 7.8. For every ℓ , consider two singular points p_ℓ, q_ℓ of M_ℓ of respective cone angles θ_ℓ and θ'_ℓ , such that $d(p_\ell, q_\ell) = \delta_\ell$. Three subcases (2.i), (2.ii) and (2.iii) are to be distinguished in this situation:

- (2.i). both curvatures $(2\pi - \theta_\ell)$ and $(2\pi - \theta'_\ell)$ are positive (*i.e.* p_ℓ and q_ℓ carry positive curvature). In that case one can always reverse Thurston's surgery \mathcal{S}_1 , with width of order δ_ℓ (see Section 6.1).

In the next two subcases (2.ii) and (2.iii), only one of the two curvatures $(2\pi - \theta_\ell)$ or $(2\pi - \theta'_\ell)$, say the former, is assumed to be negative: one has $\theta_\ell > 2\pi$ while $\theta'_\ell < 2\pi$. Note that in this case $(2\pi - \theta_\ell) + (2\pi - \theta'_\ell)$ is negative since there is at least one other cone point which must have positive curvature.

- (2.ii). In this subcase, we assume that there is 'enough room to reverse Thurston's surgery \mathcal{S}_2 ', *i.e.* one can extend the geodesic line from q_ℓ to p_ℓ after the point of negative curvature on a distance of

$$\frac{\sin\left(\frac{\theta'_\ell}{2}\right)}{\sin\left(\frac{\theta_\ell + \theta'_\ell}{2}\right)} \cdot \delta_\ell \quad (13)$$

so that the extended line cuts the cone angle at p_ℓ into two equal angles (see Remark 6.3). If this is possible, one can cut along the extended line and fill with an appropriate Euclidean kite, and therefore reverse Thurston's surgery \mathcal{S}_2 with small width of order δ_ℓ .

- (2.iii). We now prove that if the two subcases (2.i) and (2.ii) do not occur, we are in a situation where the kite surgery \mathcal{S}_4 can be reversed. If we cannot extend the geodesic line from q_ℓ to p_ℓ , it must be either because it meets another cone point or that the line self-intersects. The latter case cannot happen if ℓ is large enough otherwise the systole would be smaller than (13). The fact that the line self-intersecting gives rise to a non-essential curve is not totally obvious. Actually this very curve could turn around a singular point r_ℓ of cone angle smaller than π . But since there can be only one such point³ whose cone angle is smaller than π , we can play the same game with r_ℓ and p_ℓ being sure that this situation will not occur. Since θ_ℓ and θ'_ℓ range in a finite set, that would imply that σ_ℓ goes to zero. So the extended line meets a singular point r_ℓ . One can try to reverse Thurston's surgery with p_ℓ and r_ℓ . If $[p_\ell, q_\ell]$ is long enough we are brought back to the previous case. In the case when it is not long enough, we are going to prove that

$$(2\pi - \theta_\ell) + (2\pi - \theta'_\ell) + (2\pi - \theta''_\ell) = 0.$$

In that case $\dim(X_0) = 2$ and one can reverse the kite surgery with very small width.

More precisely, let $l_\ell = \delta_\ell$ be the length of the geodesic segment from q_ℓ to p_ℓ and l'_ℓ be the length of the one from r_ℓ to p_ℓ (see Figure 15).

³ Indeed, since we have supposed that the cone point of negative curvature has cone angle $< 4\pi$, there cannot be two cone points of angles $< \pi$, for otherwise the total curvature would exceed 0 and violate the Gauß-Bonnet equality.




Figure 15: Geometric picture in subcase (2.iii).

As a consequence of Lemma 7.10, we have the following trichotomy:

(1) if

$$l'_\ell > \frac{\sin\left(\frac{\theta'_\ell}{2}\right)}{\sin\left(\frac{\theta_\ell + \theta'_\ell}{2}\right)} l_\ell,$$

then perform the surgery \mathcal{S}_2 on the pair of points p_ℓ, q_ℓ ;

(2) if

$$l_\ell > \frac{\sin\left(\frac{\theta'_\ell}{2}\right)}{\sin\left(\frac{\theta_\ell + \theta'_\ell}{2}\right)} l'_\ell,$$

then perform the surgery \mathcal{S}_2 on the pair of points p_ℓ, r_ℓ ;

(3) if

$$l'_\ell \leq \frac{\sin\left(\frac{\theta'_\ell}{2}\right)}{\sin\left(\frac{\theta_\ell + \theta'_\ell}{2}\right)} l_\ell \quad \text{and} \quad l_\ell \leq \frac{\sin\left(\frac{\theta'_\ell}{2}\right)}{\sin\left(\frac{\theta_\ell + \theta'_\ell}{2}\right)} l'_\ell,$$

then perform the kite surgery on the points p_ℓ, q_ℓ and r_ℓ .

(3). **The sequence $(\sigma_\ell)_{\ell \in \mathbb{N}}$ converges to zero.**

This case is the easiest. Since $(D_\ell)_{\ell \in \mathbb{N}}$ is bounded and $\sigma_\ell \rightarrow 0$ as ℓ goes to infinity, Proposition 5.11 applies for ℓ large enough. This implies that a Devil's surgery \mathcal{S}_3 of small width can be reversed.

We have proven so far that either $(M_\ell)_{\ell \in \mathbb{N}}$ converges to a point in X_0 or that for ℓ large enough M_ℓ can be recovered from a point of X_1 by a surgery of width going to zero as ℓ goes to infinity. In that latter case, Proposition 7.9 ensures that $d(M_\ell, X_1)$ converges to zero. Applying the induction hypothesis to a sequence $(M'_\ell)_{\ell \in \mathbb{N}}$ of flat surfaces $M'_\ell \in X_1$ which are such that $d(M'_\ell, M_\ell) \leq d(M_\ell, X_1) \leq 1/\ell$ for any $\ell \gg 1$, one gets that the limit of the sequence $(M_\ell)_{\ell \in \mathbb{N}}$ in $\overline{\mathcal{F}_{[\rho]}}$ belongs to $X_0 \sqcup X_1 \sqcup \dots \sqcup X_n \subset \overline{\mathcal{F}_{[\rho]}}$.

The proof of Proposition 7.7 is over. \square

LEMMA 7.10. *Let T be a flat torus and p_1, p_2 and p_3 three distinct singular points on it, of respective cone angles θ_1, θ_2 and θ_3 . Assume that p_1 is the only point of negative curvature among all the cone points of T and that p_2, p_1 and p_3 sit, in this order, on a geodesic line broken at p_1 , cutting the cone angle θ_1 into two equal*

angles. Denote by l the length of the part of the line from p_1 to p_2 , and l' the length of the part of the line from p_1 to p_3 . Assume also that

$$l' \leq \frac{\sin\left(\frac{\theta_2}{2}\right)}{\sin\left(\frac{\theta_1+\theta_2}{2}\right)} l \quad \text{and} \quad l \leq \frac{\sin\left(\frac{\theta_3}{2}\right)}{\sin\left(\frac{\theta_1+\theta_3}{2}\right)} l'. \tag{14}$$

Then the four following assertions hold true:

- (1) $(2\pi - \theta_1) + (2\pi - \theta_2) + (2\pi - \theta_3) = 0$;
- (2) T has no other cone point than p_1, p_2 and p_3 ;
- (3) both inequalities in (14) actually are equalities;
- (4) T can be recovered by a kite surgery from a regular flat torus.

Proof. The two inequalities of (14) together yield

$$\sin\left(\frac{\theta_1 + \theta_2}{2}\right) \sin\left(\frac{\theta_1 + \theta_3}{2}\right) \leq \sin\left(\frac{\theta_2}{2}\right) \sin\left(\frac{\theta_3}{2}\right)$$

or equivalently $\cos\left(\theta_1 + \frac{\theta_2+\theta_3}{2}\right) - \cos\left(\frac{\theta_2+\theta_3}{2}\right) \geq 0$ which in its turn is equivalent to

$$\sin\left(\frac{\theta_1 + \theta_2 + \theta_3}{2}\right) \sin\left(\frac{\theta_1}{2}\right) \leq 0. \tag{15}$$

On the one hand, we have $-2\pi < (2\pi - \theta_1) + (2\pi - \theta_2) + (2\pi - \theta_3) \leq 0$ because of the Gauß-Bonnet formula. This implies that $4\pi > (\theta_1 + \theta_2 + \theta_3)/2 \geq 3\pi$ and therefore that $\sin((\theta_1 + \theta_2 + \theta_3)/2) \leq 0$, with equality if and only if θ_1, θ_2 and θ_3 sum up to 6π . But on the other hand, $2\pi < \theta_1 < 4\pi$ according to our hypothesis hence $\pi < \theta_1/2 < 2\pi$ and $\sin(\theta_1/2) < 0$. Inequality (15) forces $\sin((\theta_1 + \theta_2 + \theta_3)/2)$ to vanish. Therefore $\theta_1 + \theta_2 + \theta_3 = 6\pi$ or equivalently

$$(2\pi - \theta_1) + (2\pi - \theta_2) + (2\pi - \theta_3) = 0.$$

This implies in particular that

$$\frac{\sin\left(\frac{\theta_2}{2}\right)}{\sin\left(\frac{\theta_1+\theta_2}{2}\right)} = \left(\frac{\sin\left(\frac{\theta_3}{2}\right)}{\sin\left(\frac{\theta_1+\theta_3}{2}\right)}\right)^{-1}$$

and therefore one obtains that the inequalities in (14) actually are equalities.

Note that the lengths of two consecutive sides of a kite of external angles $\theta_1/2, \theta_2, \theta_3$ satisfy the above equalities and therefore one can cut along the aforementioned geodesic line and fill with the appropriate kite to reverse the kite surgery. □

8 Finiteness of the Volume of $\mathcal{F}_{[\rho]}$

In this section, we continue to use the notations of the preceding one: $\mathcal{F}_{[\rho]}$ stands for a n -dimensional moduli space of flat surfaces of genus 0 or 1.

Below, we prove that the volume of $\mathcal{F}_{[\rho]}$ is finite under the hypothesis that ρ has finite image. Without the latter assumption (that we shall assume to hold true for the remainder of the section), it is possible to prove that the volume of $\mathcal{F}_{[\rho]}$ must be infinite. We will only be interested in the genus 1 case, the genus 0 case having already been dealt with by Thurston in [Thu98].

Proposition 5.6 essentially tells us that the lack of compactness of the metric completion of $\mathcal{F}_{[\rho]}$ is characterised by the property of having large embedded flat cylinders. Surfaces satisfying this property can be recovered by performing a surgery on a flat sphere along a distinguished geodesic segment between two conical points, see Section 6.5.

8.1 Cylindrical Coordinates. Let $T_0 \in \mathcal{F}_{[\rho]}$ be a torus containing a flat embedded systolic cylinder. It is built up from a flat sphere S_0 on which the surgery \mathcal{S}_5 described in Section 6 has been performed along a geodesic segment between two conical points of S_0 . Let $\tilde{\rho}$ be such that $S_0 \in \mathcal{F}_{[\tilde{\rho}]}$ and let (z_0, \dots, z_{n-1}) be a local linear parametrisation of $\mathcal{F}_{[\tilde{\rho}]}$ at S_0 such that z_0 represents the geodesic path along which the surgery is performed, and let z_n be the complex number such that the inserted cylinder has sides z_0, z_n . Then (z_0, \dots, z_n) is a linear parametrisation of $\mathcal{F}_{[\rho]}$. We call any such parametrisation a **cylindrical parametrisation** whose existence is guaranteed by Proposition 4.3.

Let A be the area form of the flat tori in $\mathcal{F}_{[\rho]}$ close to T_0 expressed in the coordinates z_0, \dots, z_n and denote by B the area form of the associated flat spheres in $\mathcal{F}_{[\tilde{\rho}]}$, expressed in the coordinates z_0, \dots, z_{n-1} .

The (signed) area of the aforementioned embedded flat cylinder is $\text{Im}(z_n \bar{z}_0)$ therefore the two area forms A and B are linked by the following relation:

$$A(z_0, \dots, z_n) = B(z_0, \dots, z_{n-1}) + \text{Im}(z_n \bar{z}_0).$$

Normalising with $z_0 = 1$, we get a genuine parametrisation of $\mathcal{F}_{[\rho]}$ (resp. $\mathcal{F}_{[\tilde{\rho}]}$) (z_1, \dots, z_n) (resp. (z_1, \dots, z_{n-1})) and the preceding relation becomes

$$A(1, z_1, \dots, z_n) = B(1, z_1, \dots, z_{n-1}) + \text{Im}(z_n).$$

8.2 Finiteness of the Volume. The strategy to estimate the complex hyperbolic volume of $\mathcal{F}_{[\rho]}$ is to restrict ourselves to parts of $\mathcal{F}_{[\rho]}$ where the diameter is large (*i.e.* where corresponding flat tori have large embedded flat cylinders, see Proposition 5.6) and use the cylindrical coordinates defined above to perform some quasi-explicit estimations.

In what follows, all the flat cylinders which we will consider will be assumed to be ‘*maximal*’ in the sense that none of them is a proper subcylinder of an embedded flat cylinder of the same width but of strictly higher length.

For every positive ϵ , one sets

$$A_\epsilon = \{ T \in \mathcal{F}_{[\rho]} \mid \sigma(T) = \epsilon \text{ and } T \text{ contains a flat cylinder of width } \epsilon \}$$

and $B_\epsilon = \{ T \in \mathcal{F}_{[\rho]} \mid \sigma(T) \leq \epsilon \text{ and } T \text{ contains a flat cylinder of width } \sigma(T) \}$.

As is often implicitly assumed in a large part of the paper, the points of $\mathcal{F}_{[\rho]}$ are flat surfaces which are supposed to be normalised in order that their area is 1. In particular we assume this hypothesis in the definitions above.

Both A_ϵ and B_ϵ are closed subsets of $\mathcal{F}_{[\rho]}$. Moreover, from Section 6.5, it comes that when non-empty, A_ϵ is a smooth real-analytic hypersurface in $\mathcal{F}_{[\rho]}$.

For a given $\epsilon > 0$, the elements of A_ϵ can be modified by thickening of a length t the embedded flat cylinder of width ϵ (by thickening, we mean replacing the cylinder of length l by a cylinder of length $l + t$). When renormalising in order for the area to be 1, the width of the cylinder becomes smaller than ϵ . This defines a map

$$A_\epsilon \times \mathbb{R}_{\geq 0} \longrightarrow B_\epsilon, \tag{16}$$

which is a local diffeomorphism (see Section 6.5). The fact that this map is well-defined relies on the uniqueness of the maximal cylinder of width $\epsilon = \sigma(T)$ (the ‘systolic cylinder’) in any $T \in A_\epsilon$ for ϵ small enough. This fact follows easily from Lemma 8.4 which is proved in Subsection 8.3. Note that from this Lemma, it also can be deduced that A_ϵ is precisely the boundary of B_ϵ .

PROPOSITION 8.1. *For any ϵ sufficiently small, the map $A_\epsilon \times \mathbb{R}_{\geq 0} \longrightarrow B_\epsilon$ is onto.*

Proof. First remark that it is sufficient to prove the proposition for a fixed $\epsilon_0 > 0$, because the statement will then hold true for every smaller ϵ .

Second, since B_ϵ is the disjoint union of the A_η ’s for $\eta \in]0, \epsilon]$, the proposition follows from the fact that, for every $t > 0$, the image of $A_\epsilon \times \{t\}$ by (16) is the whole hypersurface $A_{\epsilon/\sqrt{1+\epsilon t}}$, as soon as ϵ is taken sufficiently small. This technical assertion is proved in Subsection 8.4 below. □

For the remainder of this section, we fix ϵ such that (16) is surjective. Now remark that the closure of $\mathcal{F}_{[\rho]} \setminus B_\epsilon$ in $\overline{\mathcal{F}_{[\rho]}}$ is compact. Indeed, a sequence in $\mathcal{F}_{[\rho]} \setminus B_\epsilon$ must have bounded diameter according to Proposition 5.6. But then, up to passing to a subsequence, it is a Cauchy sequence by Lemma 7.8 and therefore converges in $\overline{\mathcal{F}_{[\rho]}}$. Since $A_\epsilon = \partial B_\epsilon$, its closure in $\overline{\mathcal{F}_{[\rho]}}$ must be compact as well.

We are now able to prove the

PROPOSITION 8.2. *The volume of B_ϵ is finite.*

Proof. Since the closure of A_ϵ in $\overline{\mathcal{F}_{[\rho]}}$ is compact, A_ϵ can be recovered by a finite union of simply-connected open sets $(U_i)_{i \in I}$ such that for each $i \in I$:

- (1) the diameter of U_i for the complex hyperbolic metric is finite;
- (2) there are cylindrical coordinates defined on U_i .

More precisely, each element in U_i can be recovered from surgery \mathcal{S}_5 on a sphere of a certain leaf $\mathcal{F}_{[\bar{\rho}]}$ along a geodesic joining two singular points and we have a linear parametrisation (z_0, \dots, z_n) of U_i such that

- z_0 parametrises the geodesic along which the surgery is performed;
- (z_0, z_n) parametrises the added cylinder;
- (z_0, \dots, z_{n-1}) is a linear parametrisation of $\mathcal{F}_{[\bar{\rho}]}$.

The area form therefore writes down the following way

$$A(z_0, \dots, z_n) = B(z_0, \dots, z_{n-1}) + \text{Im}(z_n \bar{z}_0). \quad (17)$$

Normalising with $z_0 = 1$, we get a parametrisation of $\mathcal{F}_{[\rho]}$ (resp. of $\mathcal{F}_{[\bar{\rho}]}$) by (z_1, \dots, z_n) (resp. by (z_1, \dots, z_{n-1})) and the preceding relation becomes

$$A(1, z_1, \dots, z_n) = B(1, z_1, \dots, z_{n-1}) + \text{Im}(z_n).$$

In this chart the local diffeomorphism $A_\epsilon \times \mathbb{R}_{\geq 0} \longrightarrow B_\epsilon$ is given by

$$((z_1, \dots, z_n, \theta), t) \longmapsto (z_1, \dots, z_{n-1}, z_n + it + \theta).$$

where θ is the twist parameter of the cylinder of width ϵ in A_ϵ (see Figure 16 below). At this point, we would like to stress that $(z_1, \dots, z_n, \theta)$ is not a system of coordinates on A_ϵ , but the latter written in these coordinates is a real-analytic submanifold of codimension 2.

In view of (17), $(z_1, \dots, z_{n-1}, z_n + it + \theta)$ is a system of pseudo-horospherical coordinates on B_ϵ (see Appendix A where this notion is introduced and discussed). Since the diameter of each U_i is finite, the image of $B_i = (A_\epsilon \cap U_i) \times \mathbb{R}_{\geq 0}$ by the preceding map is included into a domain U_{K_i, λ_i} introduced in Appendix A, for some $K_i, \lambda_i > 0$. It follows (from Lemma A.3) that the complex hyperbolic volume $\text{vol}(B_i)$ of B_i is finite for any $i \in I$. There are only a finite number of U_i 's covering A_ϵ and since ϵ has been taken such that the map (16) is onto, one gets

$$\text{vol}(B_\epsilon) \leq \sum_{i \in I} \text{vol}(B_i),$$

which implies that the complex hyperbolic volume of B_ϵ is finite. \square

As mentioned above, the finiteness of the volume of $\mathcal{F}_{[\rho]}$ follows from the preceding proposition, hence we have proved the following theorem:

Theorem 8.3. *Assume $g = 1$ and θ is such that $2\pi < \theta_1 < 4\pi$ and $\theta_i < 2\pi$ for $i \geq 2$. If ρ has finite image, then the volume of $\mathcal{F}_{[\rho]}$ for its complex hyperbolic structure is finite.*




Figure 16: Parametrisation of the extension of the cylinder.

8.3 A Uniqueness Result for Cylinders of Small Width. We now prove the following lemma which implies the result announced in Subsection 8.2 (namely, for ϵ sufficiently small, the unicity of the systolic cylinder in any element of A_ϵ):

LEMMA 8.4. *There exists $\epsilon_\rho > 0$ (depending only on ρ) such that for any positive $\epsilon < \epsilon_\rho$ and any flat torus $T \in \mathcal{F}_{[\rho]}$, the following holds true: if T belongs to B_ϵ , i.e. if T contains an embedded flat cylinder of width $\sigma(T) \leq \epsilon$, then the latter is unique among the flat cylinders embedded in T of width strictly less than ϵ_ρ .*

Proof. Since $\text{Im}(\rho)$ is finite, the number of genus 0 moduli spaces $\mathcal{M}_{0,\theta'}$ that can be obtained from elements of $\mathcal{F}_{[\rho]}$ by reversing the surgery \mathcal{S}_5 is finite. Consequently, the minimum κ of the set of positive constants $K(\theta'_1, \theta'_2)$ given by Lemma 5.7 for the corresponding angle data θ' (of course each time with respect to the two cone points involved in the surgery), is positive as well.

Let $T \in \mathcal{F}_{[\rho]}$ be an element of B_ϵ , for a fixed $\epsilon > 0$ supposed to be strictly less than κ . By definition of B_ϵ , T contains a flat cylinder C_1 of width $\epsilon_1 = \sigma(T)$ with $\epsilon_1 \leq \epsilon$. To prove the lemma (with $\epsilon_\rho = \kappa$), we argue by contradiction by assuming that T contains another flat cylinder C_2 , of width $\epsilon_2 < \kappa$.

First assume that the interiors of these two cylinders intersect. Then there exists a boundary component ∂_2 of C_2 which intersects C_1 . Since ∂_2 is totally geodesic, it must enter C_1 through one of its boundary components and exit by the other. But ∂_2 has length ϵ_2 so C_1 has length at most ϵ_2 . On the other hand, since the width of C_1 coincides with the systole of T , it follows from Lemma 5.10 that C_1 is a ‘systolic cylinder’. In particular, it is a flat embedded cylinder as in Figure 14 with respect to which the surgery \mathcal{S}_5 can be reversed. But removing C_1 from T and inverting the surgery \mathcal{S}_5 would give a flat sphere with a short geodesic between two of its cone points. After renormalization of the area, this sphere and this geodesic together would contradict Lemma 5.7 since $\epsilon_2 < \kappa$ (the computational details are left to the reader). This shows that C_1 and C_2 have disjoint interiors.

Let c_1 (resp. c_2) be a closed geodesic of length ϵ_1 (resp. ϵ_2) contained in the interior of C_1 (resp. of C_2). Cutting T along c_1 and c_2 gives us two connected components Σ and Σ' , each of them being a flat sphere with a boundary formed by two disjoint totally geodesic circles (in other terms: both Σ and Σ' are cylinders too). Note that since C_1 and C_2 are of maximal length, any component of their boundary contains a conical point from which it follows that both Σ and Σ' contain conical points. Now assume that the unique cone point of negative curvature of T belongs to Σ' . Then Σ contains only cone points of positive curvature. But this is impossible since, according to Gauß-Bonnet formula (applied to the flat surface with geodesic boundary Σ), the total curvature of Σ is equal to $2\pi \cdot \chi(\Sigma) = 0$. The lemma follows from this contradiction. \square

8.4 A Technical Lemma. We fix $\epsilon > 0$. For any positive $\eta < \epsilon$, it is easily seen that the preimage of $A_\eta \subset B_\epsilon$ by (16) is $A_\epsilon \times \{t_\eta\}$ with $t_\eta = (\epsilon^2 - \eta^2)/(\epsilon\eta^2) > 0$.

LEMMA 8.5. *For ϵ sufficiently small, any map $A_\epsilon \times \{t_\eta\} \rightarrow A_\eta$ is surjective.*

Proof. We claim that the statement of the lemma holds true for any $\epsilon < \kappa$, where κ stands for the positive constant introduced in the first paragraph of the proof of Lemma 8.4. Indeed, if it were not the case, there would exist $T \in A_\eta$ for some $\eta < \epsilon$, which was not in the image of $A_\epsilon \times \{t_\eta\} \rightarrow A_\eta$. For such a T , one verifies that the length of the systolic flat cylinder of T (of width η) is necessarily less than or equal to $\eta t_\eta / \epsilon$. Then removing this cylinder from T and inverting the surgery \mathcal{S}_5 would give a flat sphere with a short geodesic between two of its cone points. After renormalization of the area, this sphere and this geodesic together would contradict Lemma 5.7 since $\epsilon < \kappa$ (the computational details are left to the reader). \square

8.5 Proof of Lemma 7.8. We finally explain how the above description of the parts of $\mathcal{F}_{[\rho]}$ consisting of tori with long embedded cylinders gives a proof of the first point of Lemma 7.8, namely that the diameters of the elements of any Cauchy sequence in $\mathcal{F}_{[\rho]}$ are uniformly bounded.

Consider a path $\gamma : [0, 1] \rightarrow B_\epsilon$. We have the following estimate

$$L(\gamma) \geq \left| \log(c(\gamma(1))) - \log(c(\gamma(0))) \right|.$$

where $c(T)$ is the length of the cylinder of width at most ϵ (with $T \in B_\epsilon$). This is a direct consequence of Lemma A.3 of Appendix A.

Assume that we have a Cauchy sequence whose diameter goes to infinity. We can assume that all its elements belong to a subset B_ϵ of $\mathcal{F}_{[\rho]}$ considered above (this follows from Proposition 5.6). Applying the above estimate to paths linking elements of the sequence leads to a contradiction.

9 The Metric Completion is a Cone-Manifold

In this section, we prove a theorem describing the structure of the metric completion of $\mathcal{F}_{[\rho]}$. We refer to Appendix B for further precisions and references on the notion of cone-manifold.

We turn back to the notations used before the two preceding sections : we are dealing with flat surfaces of genus $g = 0, 1$ with n cone points. More precisely, we assume that either

- $g = 0$ and $\theta_i \in]0, 2\pi[$ for all $i = 1, \dots, n$; or
- $g = 1$, $\theta_1 \in]2\pi, 4\pi[$ and $\theta_i \in]0, 2\pi[$ for $i = 2, \dots, n$,

so that Veech’s geometric structure on $\mathcal{F}_{[\rho]}$ is complex hyperbolic.

Theorem 9.1. *Let $\rho \in H^1(N, \mathbb{U}, \theta)$ be such that $\text{Im}(\rho)$ is finite. The metric completion of $\mathcal{F}_{[\rho]}$ is a complex hyperbolic cone-manifold.*

The proof goes by induction on $m = \dim(\mathcal{F}_{[\rho]})$. Assume that it has been proven that the theorem holds true for all $\mathcal{F}_{[\rho]}$ carrying a complex hyperbolic structure such that $\dim(\mathcal{F}_{[\rho]}) \leq m - 1$. The case $g = 0$ has been dealt with by Thurston in [Thu98]. The base case of the induction is when $m = 0$ that is $\mathcal{F}_{[\rho]}$ is a point in which case the theorem holds. Note that $m = 0$ can only happen if $g = 0$ and $n = 3$.

Consider $\rho \in H^1(N, \mathbb{U}, \theta)$ such that $\dim(\mathcal{F}_{[\rho]}) = m$. We have proven in Section 7 (see Theorem 7.1) that $X = \overline{\mathcal{F}_{[\rho]}}$ is a disjoint union of X_0, \dots, X_m such that

- $X_0 = \mathcal{F}_{[\rho]}$;
- X_i is a complex hyperbolic manifold of dimension $m - i$;
- the metric completion of X_i in X is $X_i \sqcup X_{i+1} \sqcup \dots \sqcup X_m$.

We prove by induction on i that any point in X_i has a neighbourhood in X isometric to a complex hyperbolic cone-manifold. Let p be a point of X_i . It has a neighbourhood U_p in X_i which is isometric to an open subset of $\mathbb{C}H^{m-i}$. Following Thurston in [Thu98], we define an ‘orthogonal projection’ $\pi : V_p \longrightarrow U_p$ from a neighbourhood V_p of p in X onto U_p the following way: we have seen in Section 7 that there exists a neighbourhood V_p of p in X such that for any $q \in V_p$ there exists a unique $r \in U_p$ such that q can be recovered from r by performing a finite number of the four surgeries $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$ and \mathcal{S}_4 described in Section 6, and one has $\pi(q) = r$.

Thurston calls this map ‘orthogonal projection’ because, in a sense which is made precise in Appendix B, the fibers of π are orthogonal to its image U_p .

LEMMA 9.2. *Let p, U_p, V_p and π be defined as above.*

- (1) *For all $r \in U_p$, $V(r) = \pi^{-1}(r) \setminus \{r\}$ is foliated by geodesics ending at r .*
- (2) *For all $r \in U_p$, the intersection $V(r) \cap (X_0 \sqcup \dots \sqcup X_{i-1})$ is a totally geodesic sub-cone-manifold of $X_0 \sqcup \dots \sqcup X_{i-1}$.*
- (3) *For all $r \in U_p$, $V(r)$ is orthogonal to U_p .*

Proof. We fix $r \in U_p$ and consider an element q of $V(r) = \pi^{-1}(r)$. By using an appropriate topological polygonation (cf. Proposition 4.3), one can find a linear parametrisation (z_1, \dots, z_m) at q such that

- if $(\xi_0, \dots, \xi_{m-i}, \dots, \xi_m)$ are the coordinates of q in this parametrisation, then $\pi(q)$ has coordinates $(\xi_0, \dots, \xi_{m-i}, 0, \dots, 0)$;
- the area form A in the z_i 's can be written out

$$A(z_0, \dots, z_m) = A_1(z_0, \dots, z_{m-i}) - A_2(z_{m-i+1}, \dots, z_m) \quad (18)$$

where A_1 has signature $(1, m-i)$ and A_2 is positive-definite (A_2 is the total area removed by the successive surgeries).

- in the local coordinates z_1, \dots, z_m , the stratum X_j for $j > i$ is cut out by the equations $z_{m-(j+1)} = z_{m-j+2} = \dots = z_m = 0$ (there is a subtlety here about the precise location of the locus cut out by these equations with respect to the domain of definition of the linear parametrisation we consider. We let the reader state a remark analogous to Remark 7.5 for the case under scrutiny).

The image of $[0, 1] \ni t \longrightarrow (\xi_0, \dots, \xi_{m-i}, t\xi_{m-i-1}, \dots, t\xi_m) \in \mathbb{C}^{m+1}$ projects onto a geodesic path in V_p joining q to $r = \pi(q)$ (see Lemma A.1.(2)). Remark that this geodesic does not depend on the choice of the z_i 's: it is the one pointing in the direction of r hence it is intrinsic (because there is a unique geodesic segment linking to distinct points in a complex hyperbolic space). The collection of those geodesics for all $q \in V(r)$ gives the announced foliation of $V(r)$. Taking V_p small enough, one can ensure that the foliation is globally well defined by using for instance a finite number of linear parametrisations whose pairwise intersections are 1-connected. The first point of the lemma is proved.

A neighbourhood of q in $V(r) \cap X_0$ consists of the submanifold parametrised by z'_1, \dots, z'_m such that $z'_1 = \xi_1, \dots, z'_{m-i} = \xi_{m-i}$ and therefore projects onto a totally geodesic subspace $\mathbb{C}\mathbb{H}^i$ (we use the fact that a complex affine submanifold of the complex hyperbolic space is totally geodesic).

Finally, the splitting (18) of A gives us that $V(r)$ and U_p are orthogonal. \square

We continue to use the notations of the previous lemma.

PROPOSITION 9.3. *With the same notations as before, $V(r)$ is a $\mathbb{C}\mathbb{H}^i$ -cone-manifold with r as its unique cone point.*

Proof. Define

$$B(\epsilon) = \{q \in V(r) \mid \exists \text{ a geodesic of length } \leq \epsilon \text{ joining } q \text{ to } r\}$$

and $S(\epsilon) = \partial B(\epsilon) = \{q \in V(r) \mid \exists \text{ a geodesic of length } \epsilon \text{ joining } q \text{ to } r\}$

For ϵ small enough, $S(\epsilon)$ does not meet X_j for $j \geq i$ then it lives in

$$X \setminus (\cup_{j \geq i} X_j) = X_0 \sqcup X_1 \sqcup \dots \sqcup X_{i-1}$$

which is a complex hyperbolic cone-manifold according to the induction hypothesis.

In particular $S(\epsilon)$ is locally a totally geodesic sub-cone-manifold intersected with a piece of a complex hyperbolic sphere whose center belongs to $V(r)$. It is therefore, according to Lemma B.3, a $(S^{2i-1}, U(i))$ -cone-manifold. According to the previous lemma, $B(\epsilon)$ is a cone over this cone-manifold and the proposition is proved. \square

PROPOSITION 9.4. *For any $r \in U_p$, there exists a neighbourhood of r in X which is a complex hyperbolic cone-manifold.*

Proof. There exists an $\epsilon > 0$ such that for all $r \in U$, the ball of radius ϵ at r in V_r is an embedded cone. There is a neighbourhood of U in X which has the product structure $U \times B(\epsilon)$ satisfying the hypothesis of Proposition B.5. Hence $U \times B(\epsilon)$ is a complex hyperbolic cone-manifold and this proves the proposition. \square

The induction process can be carried on which proves that $X = \overline{\mathcal{F}_{[\rho]}}$ is a complete complex hyperbolic cone-manifold.

10 Listing the $\mathcal{F}_{[\rho]}$'s and Their Codimension 1 Strata

We have given so far a rather abstract analysis of the geometric structure of a leaf $\mathcal{F}_{[\rho]}$ when $\text{Im}(\rho)$ is finite. We now give a list of all such $\mathcal{F}_{[\rho]}$ associated to a rational angle datum θ when $g = 1$. Let $G_\theta \subset \mathbb{U}$ be the subgroup generated by $e^{i\theta_1}, \dots, e^{i\theta_n}$ and ω_ρ a root of unity such that $\text{Im}(\rho) = \langle \omega_\rho \rangle$.

Let M be the smallest positive integer such that $G_\theta = \langle \omega_\rho^M \rangle$.

10.1 Listing the $\mathcal{F}_{[\rho]}$'s Associated to θ . The starting point of our description is the following lemma:

LEMMA 10.1. *Consider ρ and ρ' two elements of $H^1(N, \mathbb{U}, \theta)$ such that:*

- (1) $\text{Im}(\rho) = \text{Im}(\rho')$;
- (2) $\text{Im}(\rho)$ (and therefore $\text{Im}(\rho')$) is finite.

Then ρ and ρ' are equivalent under the action of the pure mapping class group.

Proof. Let $\rho \in H^1(N, \mathbb{U}, \theta)$ be such that $\text{Im}(\rho)$ is finite, equals to $\langle \exp(2i\pi/q) \rangle$ for some positive integer q .

We consider two simple closed curves a, b avoiding the marked points of N which form a symplectic basis of the homology of the unmarked torus N , and c_1, \dots, c_n curves that circle the marked points, chosen in such a way that $\rho_k = \rho(c_k) = \exp(i\theta_k)$ for $k = 1, \dots, n$. Up to the action of an element of the pure mapping class group, we can replace a and b by $a^\alpha b^\beta$ and $a^\gamma b^\delta$ with $\begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} \in \text{SL}_2(\mathbb{Z})$. We can this way arrange that $\rho_a = \rho(a) = 1$ with $\rho_b = \rho(b)$ such that $\text{Im}(\rho) = \langle G_\theta, \rho_b \rangle$. From this

equality and thanks to our assumption, there exist integers ν_1, \dots, ν_n and $\nu \neq 0$ such that $(\rho_b)^\nu \cdot \rho_1^{\nu_1} \cdots \rho_n^{\nu_n} = e^{2i\pi/q}$ hence generates $\text{Im}(\rho)$. Then replacing a and b by $ab^{\nu-1}$ and $ab^\nu c_1^{\nu_1} \cdots c_n^{\nu_n}$ respectively, one can assume that $\text{Im}(\rho) = \langle \rho_b \rangle$ with $\rho_b = e^{2i\pi/q}$. Then $\rho_a = (\rho_b)^\mu$ for a certain integer p . Then considering ab^{q-p} instead of a , one eventually obtains that the element of $H^1(N, \mathbb{U}, \theta)$ uniquely characterized by assignating the values 1 and $e^{2i\pi/q}$ to a and b respectively, is a representative of the orbit of the initial character ρ under the action of the pure mapping class group. As an immediate consequence of the preceding fact, we get that the class of an element $\rho \in H^1(N, \mathbb{U}, \theta)$ under the action of the pure mapping class group only depends on its image, if the latter is finite. This gives the lemma. \square

A leaf $\mathcal{F}_{[\rho]}$ is therefore only determined by its associated angle data θ and the smallest integer $M \geq 1$ which is such that G_θ is generated by ω_ρ^M . From now on, we refer to such a leaf/moduli space as $\mathcal{F}_\theta(M)$. We are now going to give a description of its codimension 1 strata.

We distinguish three types of such strata:

- the **P-strata** which are obtained from Devil's surgery. Here 'P' stands for *pinching*;
- the **C-strata** which are obtained from a Thurston's type surgery. Here 'C' stands for *colliding*;
- the **K-strata** which are obtained from a kite surgery. Here 'K' stands for *kite*.

(Note that since a K -stratum appears when three cone points collide together, it can be considered as a particular kind of C -stratum).

At this point we must mention an aspect of the description that we have so far ignored: we have been using throughout the article the terminology 'leaf' in a non-standard way. While in foliation theory, a leaf is automatically supposed to be connected, the definition we use allows $\mathcal{F}_\theta(M)$ not to be. Such a non connectedness phenomena can indeed happen when N has genus 1 as the explicit description of the case $g = 1$, $n = 2$ carried on in [GP] reveals (see [GP, §4.2.5] for an explicit example). The determination of the connected components in the general case is an open problem that seems interesting to the authors for the reasons explained in Section 11.

In what follows, we fix a leaf $\mathcal{F}_{[\rho]}$ with $\text{Im}(\rho)$ finite and we explain below several algorithms to determine the strata of complex codimension 1 appearing in the completion $\overline{\mathcal{F}_{[\rho]}}$.

10.2 Finding the P-Strata. Let m be the positive integer such that G_θ is the subgroup of \mathbb{U} generated by $e^{\frac{2i\pi}{m}}$. With these notations, one has

$$\text{Im}(\rho) = \left\langle e^{\frac{2i\pi}{mM}} \right\rangle.$$

Since $\theta_1 > 2\pi$ there exists p such that

$$\theta_1 = 2\pi \left(1 + \frac{p}{m}\right).$$

A P -stratum of codimension 1 is a (finite cover of a) moduli space of flat spheres whose angles datum is $(\theta'_1, \theta''_1, \theta_2, \dots, \theta_n)$ with θ'_1 and θ''_1 such that

- (1) $\theta'_1 + \theta''_1 = \theta_1 - 2\pi$;
- (2) both $e^{i\theta'_1}$ and $e^{i\theta''_1}$ belong to $\text{Im}(\rho) = \langle e^{\frac{2i\pi}{mM}} \rangle$.

This condition is sufficient for such a (finite cover of a) moduli space of flat spheres to appear as a stratum of the metric completion of $\mathcal{F}_{[\rho]}$. There is therefore a P -stratum for each way of decomposing the integer pM as a sum of two positive integers, this number being $\lfloor \frac{pM}{2} \rfloor$.

The stratum associated to a decomposition

$$pM = r' + r'' \tag{D}$$

with $r', r'' > 0$ is a finite cover of the moduli space of flat spheres whose angles datum is

$$\left(e^{\frac{2i\pi r'}{mM}}, e^{\frac{2i\pi r''}{mM}}, \theta_2, \dots, \theta_n \right).$$

According to Section 6.7, the cone angle around the stratum associated to the decomposition (D) is $2\pi \cdot \text{lcm}(r', r'') / (mM)$.

10.3 Finding the C -Strata. A C -stratum of codimension 1 is a moduli space of flat tori with $n - 1$ cone points corresponding to the collision of two cone points p_k and p_l of respective angles θ_k and θ_l . The new angle datum $\theta' = (\theta'_i)_{i=1}^{n-1}$ is such that θ_k and θ_l have been replaced by $\theta_k + \theta_l - 2\pi$, the other θ_i 's staying unchanged.

A holonomy character $\rho' \in \mathbb{H}^1(N, \mathbb{U}, \theta')$ is such that a finite cover of $\mathcal{F}_{[\rho']}$ is a stratum of $\mathcal{F}_{[\rho]}$ if and only if $\text{Im}(\rho)$ is generated by $\text{Im}(\rho')$, $e^{i\theta_k}$ and $e^{i\theta_l}$ (see Subsection 6.1). We describe now the positive integers M' which are such that $\mathcal{F}_{\theta'}(M')$ appears as a C -stratum of the metric completion of $\mathcal{F}_{[\rho]} = \mathcal{F}_\theta(M)$.

Let m' be the positive integer such that $G_{\theta'}$ is generated by $e^{\frac{2i\pi}{m'}}$. Remark that m' divides m . We are trying to find the integers M' such that $e^{\frac{2i\pi}{m'M'}}$ and $e^{\frac{2i\pi}{m}}$ generate $\text{Im}(\rho) = \langle e^{\frac{2i\pi}{mM}} \rangle$. This is equivalent to find the integers M' such that $\text{lcm}(m, m'M') = mM$. Since $m'|m$, this is equivalent to determine the positive integers M' verifying

$$\text{lcm} \left(M', \frac{m}{m'} \right) = M \frac{m}{m'}.$$

The list of solutions to the preceding relation viewed as an equation in M' , provides the list of leaves $\mathcal{F}_{\theta'}(M')$, associated to the collision between p_k and p_l which appear as C -strata of $\mathcal{F}_{[\rho]}$. For any such C -strata, it comes from Section 6.7 that the conifold angle around it is $\theta' = \theta_k - \theta_l - 2\pi$.

Remark that, as in the genus 0 case considered by Thurston, the complex hyperbolic conifold angle θ' coincides with the new cone angle of the flat surfaces whose isomorphism classes belong to the considered stratum.

10.4 Finding the K -Strata. A K -stratum appears in codimension 1 in the metric completion of a leaf $\mathcal{F}_{[\rho]}$ if and only if $n = 3$ and $G_\theta = \text{Im}(\rho)$. In this case, the conifold angle around such a stratum is π according to the third point of Section 6.7.

10.5 The 1-Dimensional Case. We now consider the case when $n = 2$. We assume $\theta = (\theta_1, \theta_2) \in 2\pi\mathbb{Q}^2$ with $\theta_1 + \theta_2 = 4\pi$. A leaf $\mathcal{F}_\theta(M)$ is a 1-dimensional complex hyperbolic manifold or equivalently, a real hyperbolic surface.

According to Theorem 9.1, the metric completion of $\mathcal{F}_\theta(M)$ is a hyperbolic surface of finite volume, with a finite number of cone points and a finite number of cusps. We give in this section a refinement of the description of the P -strata appearing in $\overline{\mathcal{F}_\theta(M)}$ (there is actually no C -stratum in the metric completion of $\mathcal{F}_\theta(M)$ when $n = 2$ according to Proposition 5.9) and give a list of the cusps by geometric means. We finally give explicit details in the case when $\theta = (3\pi, \pi)$.

A leaf $\mathcal{F}_\theta(M)$ when $n = 2$ shall be thought of as a generalisation of the modular surface $\mathbb{H}/\text{PSL}(2, \mathbb{Z})$ which is the moduli space of regular flat tori. Veech's hyperbolic structure matches its standard one. It is not very surprising that the analytical analysis carried on in [GP] shows that the connected components of $\mathcal{F}_\theta(M)$ are conformally equivalent to modular curves of the form $Y_1(N) = \mathbb{H}/\Gamma_1(N)$ for certain integers $N \geq 2$, see [GP, §4.2.4] for more details.

Cone points. Strata of $\mathcal{F}_\theta(M)$ correspond to flat spheres S with 3 cone points whose associated angle datum $(\theta'_1, \theta''_1, \theta_2)$ is such that $\theta'_1 + \theta''_1 = \theta_1 - 2\pi$. Listing such flat spheres has been done in the preceding section. We also saw that P -strata are finite covers of moduli spaces of flat spheres, in this specific case such a finite cover is a union of points. There are as many copies of S appearing in the metric completion $\overline{\mathcal{F}_\theta(M)}$ as ways of performing Devil's surgery on S .

Let $T \in \mathcal{F}_\theta(M)$ be a torus built by Devil's surgery on S . Let $\widehat{\beta}$ be a simple curve on T avoiding the singular points that intersects the systole only once (see Section 6.3 for more details and pictures). We remark that $\rho(\widehat{\beta})$ has to be such that $\langle e^{i\theta'_1}, e^{i\theta''_1}, e^{i\theta_2}, \rho(\widehat{\beta}) \rangle = \text{Im}(\rho)$.⁴

With the notation of the previous section, if

$$\theta'_1 = \frac{2\pi r'}{mM} \quad \text{and} \quad \theta''_1 = \frac{2\pi r''}{mM},$$

then there are exactly $\text{gcd}(r', r'')$ different ways to perform Devil's surgery in such a way that the holonomy along $\widehat{\beta}$ belongs to $\text{Im}(\rho)$. Amongst these ways, only $\varphi(\text{gcd}(r', r'', mM))$ (where φ is Euler's totient function) are such that $\langle e^{i\theta'_1}, e^{i\theta''_1}, e^{i\theta_2} \rangle = \text{Im}(\rho)$ and this number is the exact number of times that S appears in the metric completion of $\mathcal{F}_\theta(M) = \mathcal{F}_{[\rho]}$.

We now explain how to perform this counting. As explained in Subsection 6.3 (to which we refer for the notations used below), the parameters for Devil's surgery

⁴ Note that here and in what follows, the unitary holonomy character ρ corresponds to the one which was denoted by $\widehat{\rho}$ in §6.3.

are pairs of points (q', q'') on two circles (respectively identified with) $\mathbb{R}/\theta'_1\mathbb{Z}$ and $\mathbb{R}/\theta''_1\mathbb{Z}$, around the two singular points of the flat sphere involved in the surgery and a positive parameter r equal to the radius of the aforementioned circle (up to constants depending on θ'_1 and θ''_1). These points q' and q'' are the points which are going to be identified together to create a new point of negative curvature. Remark that all such parameters do not give rise to surfaces necessarily belonging to $\mathcal{F}_{[\rho]}$: for this to happen, q' and q'' have to move along $\mathbb{R}/\theta'_1\mathbb{Z}$ and $\mathbb{R}/\theta''_1\mathbb{Z}$ by the same amount. This way the holonomy along $\widehat{\beta}$ (of the flat structure of the surfaces obtained after surgery) remains constant. The number of ways of performing Devil's surgery is the number of connected components of the space of such parameters (q', q'') giving rise to flat surfaces belonging to the same leaf $\mathcal{F}_{[\rho]}$. Remark that any point q' in $\mathbb{R}/\theta'_1\mathbb{Z}$ belongs to a pair (q', q'') for which the associated surface belongs to such a chosen connected component of parameters (if this set is non-empty in the first place). In order to count the number of ways to invert Devil's surgery, it suffices then to count, given a point $q' \in \mathbb{R}/\theta'_1\mathbb{Z}$, the number of points $q'' \in \mathbb{R}/\theta''_1\mathbb{Z}$ such that the couple (q', q'') gives rise to an element of $\mathcal{F}_{[\rho]}$ and distinguish those belonging to different components of parameters.

Chose any point q''_0 such that (q', q''_0) gives rise to an element of $\mathcal{F}_\theta(M)$. Denote by $\widehat{\beta}(q', q'')$ the curve $\widehat{\beta}$ ⁵ on the torus associated to the pair (q', q'') . Up to multiplication by a power of $e^{i\theta''_1}$, we have

$$\rho(\widehat{\beta}(q', q''_0 + \vartheta)) = \rho(\widehat{\beta}(q', q''_0)) \cdot e^{i\vartheta}$$

for any $\vartheta \in \mathbb{R}/\theta''_1\mathbb{Z}$ (the addition referring to the standard group law of $\mathbb{R}/\theta''_1\mathbb{Z}$).

We find this way that a pair (q', q'') gives rise to a surface such that $\rho(\widehat{\beta}) \in \langle e^{\frac{2i\pi}{mM}} \rangle$ (such a pair of parameters is said to be **admissible**) if and only if $q'' - q''_0 \in \frac{2i\pi}{mM}\mathbb{Z}$ which gives exactly r'' possibilities for a point q'' to pair with q' and giving rise to an element of $\mathcal{F}_{[\rho]}$ (recall that $\theta''_1 = \frac{2\pi r''}{mM}$). Amongst these possibilities, are in the same components of the set of admissible parameters those differing by a multiple of r' modulo r'' . This gives at most $\gcd(r', r'')$ different components of admissible parameters (see Remark 6.5).

The last thing that we have to do to make sure that the parameters we are considering truly correspond to elements of $\mathcal{F}_\theta(M) = \mathcal{F}_{[\rho]}$ is to ensure that

$$\langle e^{i\theta'_1}, e^{i\theta''_1}, e^{i\theta_2}, \rho(\widehat{\beta}) \rangle = \langle e^{\frac{2i\pi}{mM}} \rangle.$$

If β is such that $\rho(\widehat{\beta}) = e^{\frac{2ik\pi}{mM}}$ for a certain positive integer k , then

$$\langle e^{i\theta'_1}, e^{i\theta''_1}, e^{i\theta_2}, \rho(\widehat{\beta}) \rangle = \left\langle e^{\frac{2i\pi \gcd(r', r, k, mM)}{mM}} \right\rangle$$

⁵ 'The' curve $\widehat{\beta}$ is a priori not well-defined. One can make an arbitrary choice for (q', q''_0) and then define $\widehat{\beta}$ for every nearby parameter using the Gauss–Manin connection. The initial choice does not really matter because the computation of its holonomy would give results well determined up to multiplication by an element of $\langle e^{i\theta'_1}, e^{i\theta''_1} \rangle$.

which leads to $\varphi(\gcd(r', r, k, mM))$ essentially different ways to perform the surgery. Finally, the $\mathbb{C}\mathbb{H}^1$ -cone angle around such a point is $2\pi\text{lcm}(r', r'')/(mM)$ (it is a direct application of the description of Section 6.3).

Cusps. According to Subsection 6.5, the cusps of $\mathcal{F}_\theta(M)$ are in one-to-one correspondence with pairs (S, γ) where

- (1) S is a flat sphere with three cone points of angles θ'_1 , θ''_1 and θ_2 , such that $\theta'_1 + \theta''_1 = \theta_1 - 2\pi$ and $\text{Im}(\rho) = \langle e^{i\theta'_1}, e^{i\theta''_1}, e^{i\theta_2} \rangle$;
- (2) γ is a regular geodesic in S between the cone point of angle θ'_1 and the one of angle θ''_1 .

Such a geodesic always exists and is unique, we therefore simply have to count the number of flat spheres with three cone points such that $\text{Im}(\rho) = \langle e^{i\theta'_1}, e^{i\theta''_1}, e^{i\theta_2} \rangle$. This reduces to counting the number of pairs of positive integers (r', r'') such that $r' + r'' = pM$ and $\gcd(r', r'') = 1$.

10.6 An Example: $\theta = (3\pi, \pi)$. We are now going to compute the number of conical points and cusps of $\mathcal{F}_\theta(M)$ in the special case when $\theta = (3\pi, \pi)$. In this case $p = 1$ and $m = 2$.

- Each $\mathbb{C}\mathbb{H}^1$ -cone point of $\mathcal{F}_\theta(M) = \mathcal{F}_{(3\pi, \pi)}(M)$ corresponds to a partition $r + s = M$ with $r, s > 0$. To such a partition are associated $\varphi(\gcd(r, s) = \gcd(r, M))$ cone points, all of the same cone angle $2\pi\text{lcm}(r, M - r)/M$. A particular case is when $M = 2M'$ is even. In this case the partition $M = M' + M'$ only gives rise to half of the cone points predicted by the above paragraph, namely $\frac{1}{2}\varphi(M')$. Indeed the underlying sphere on which the surgery is performed has angles $(\pi, \pi/2, \pi/2)$ and has a symmetry of order two. This symmetry permutes the different inversions of the surgery, except for the case $M = 4$ where there is only $\varphi(2) = 1$ way to perform the surgery and in this particular case it only halves the cone angle at the underlying $\mathbb{C}\mathbb{H}^1$ -cone point.
- In particular the number of cone points is

$$\frac{1}{2} \sum_{r=1}^{2M-1} \varphi(\gcd(r, M)) + \frac{1}{2}\varphi(M)$$

for $M > 4$ and is 2 when $M = 3, 4$.

- There are as many cusps as proper partitions $M = r + s$ such that $r > 0$ and M are coprime. Thus the number of cusps is exactly

$$\frac{1}{2}\varphi(M)$$

for $M \geq 3$ and is equal to 1 for $M = 2$.

The number of cusps and cone points put together gives us the number of punctures of $\mathcal{F}_{(3\pi,\pi)}(M)$ which is equal to

$$\frac{1}{2} \sum_{r=1}^M \varphi(\gcd(r, M)) \tag{19}$$

for $M > 4$, to 2 for $M = 2$ and to 3 for $M = 3, 4$. When $M > 4$, a reordering of the sum (19) gives us that the total number of punctures of $\mathcal{F}_{(3\pi,\pi)}(M)$ is actually

$$\frac{1}{2} \sum_{d|M} \varphi(d)\varphi(M/d).$$

This number is equal to the number of cusps of the modular curve $Y_1(M) = \mathbb{H}/\Gamma_1(M)$ (see [DS05]). This is not a coincidence: it is proved in [GP, §4.2.5.3] that the conformal type of $\mathcal{F}_{(3\pi,\pi)}(M)$ is actually the same as the one of $Y_1(M)$.

11 Holonomy of the $\mathbb{C}\mathbb{H}^{n-1}$ -Structure: Discreteness

11.1 Previous Results in the Genus 0 Case. Thurston proves in [Thu98] (recovering by geometric methods results of Deligne and Mostow from [DM86]) that when $g = 0$ and $n \geq 4$, if the angle datum $\theta = (\theta_1, \dots, \theta_n) \in]0, 2\pi[^n$ verifies

$$\forall i, j = 1, \dots, i \neq j, \quad 2\pi < \theta_i + \theta_j \implies \theta_i + \theta_j - 2\pi \text{ divides } 2\pi, \tag{INT}$$

then the metric completion of $\mathcal{F}_\theta \simeq \mathcal{M}_{0,n}$ (the unique leaf of Veech’s foliation in this case) is a connected complex hyperbolic orbifold of finite volume and therefore a quotient $\mathbb{C}\mathbb{H}^{n-3}/\Gamma_\theta$ where Γ_θ is a lattice in $\text{PU}(1, n - 3)$. This lattice Γ_θ is exactly the image of holonomy morphism

$$\text{hol} : \pi_1(\mathcal{M}_{0,n}) \longrightarrow \text{PU}(1, n - 3)$$

of the $(\mathbb{C}\mathbb{H}^{n-3}, \text{PU}(1, n - 3))$ -structure of \mathcal{F}_θ . Sometimes it happens that the image $\Gamma_\theta = \text{hol}(\pi_1(\mathcal{F}_\theta))$ of the holonomy is a lattice in $\text{PU}(1, n - 3)$ even if the metric completion of \mathcal{F}_θ is not an orbifold. The combined works of Picard, LeVavasseur, Terada, Deligne–Mostow, Mostow, Thurston and Sauter (see [LV93, Ter73, Thu98, DM86, Mos88, Sau90]) lead to the following results:

- (1) there exist 94 angles data θ for which \mathcal{F}_θ is an orbifold, and therefore $\Gamma_\theta = \text{hol}(\pi_1(\mathcal{F}_\theta))$ is a lattice;
- (2) this builds lattices in $\text{PU}(1, N)$ for all $N = n - 3 = 1, \dots, 9$, some of them being non-arithmetic for $N = 1, 2$ and 3 (these have been for a long time the only known examples of non-arithmetic complex hyperbolic lattices until the recent work of Deraux, Parker and Paupert [DPP15]);
- (3) if $N \geq 3$, Γ_θ is a lattice if and only if θ verifies (ΣINT) , apart from one exception (we recall that (ΣINT) is the refinement of (INT) stated in the Introduction above; see also [Mos86, §1] or [Thu98, Theorem 0.2]);
- (4) when $N = 2$, there exist 9 angles data failing (ΣINT) for which $\Gamma_\theta = \text{hol}(\pi_1(\mathcal{F}_\theta))$ is a lattice.

11.2 Genus 1 and $n = 3$. We now address the following question, which must seem natural at this point:

Question. *Let T be a torus with three marked points, and θ an admissible rational angle datum. Does there exist $\rho \in \mathbf{H}^1(T, \mathbb{U}, \theta)$ such that $\text{hol}(\pi_1(\mathcal{F}_{[\rho]}))$ is a lattice in $\text{PU}(1, 2)$?*

As we have seen in Section 10, such a leaf $\mathcal{F}_{[\rho]}$ is only determined by θ and an integer M . We denote such a leaf by $\mathcal{F}_\theta(M)$. The first difficulty to address is the question of the connectedness of the leaves $\mathcal{F}_{[\rho]}$: $\mathcal{F}_\theta(M)$ may have several connected components (see Subsection 3.3 for a short discussion of this matter) and it is possible that the holonomy of one of these is a lattice and that it is not the case for the others.

The following lemma, whose proof in the genus 0 case can be found in [Mos88], outlines a strategy to search for connected components of $\mathcal{F}_\theta(M)$ whose holonomy is a lattice:

LEMMA 11.1. *Let \mathcal{F} be a connected component of $\mathcal{F}_\theta(M)$ whose complex hyperbolic holonomy is a lattice in $\text{PU}(1, 2)$. Then the complex hyperbolic holonomy of every codimension 1 stratum is a lattice in $\text{PU}(1, 1)$.*

This lemma provides necessary conditions on such a connected component \mathcal{F} to have discrete holonomy in $\text{PU}(1, 2)$, conditions which hold true in several cases. To find candidates whose holonomy is a lattice in $\text{PU}(1, 2)$, an optimistic strategy goes as follows:

- (1) identify the different connected components of $\mathcal{F}_\theta(M)$;
- (2) verify if the criterion given by Lemma 11.1 is verified, using the list of Deligne–Mostow and Thurston for genus 0 type codimension 1 strata or the strategy suggested in the next paragraph for genus 1 codimension 1 strata;
- (3) amongst the isolated candidates, compute the complex hyperbolic holonomy and verify that it is discrete.

(We call ‘**genus g strata**’ a strata whose elements are flat surfaces of genus g).

The last step seems to be the most difficult to achieve so far, since the methods used in the genus 0 case (see for example [Par06]) tends to become algorithmically too complicated in our case and strongly rely on the knowledge of simple generators of the fundamental group of $\mathcal{M}_{0,n}$ (a finite family of distinguished Dehn twists).

11.3 Some Cases when the Holonomy is an Arithmetic Lattice. We remark that for a certain number of connected components \mathcal{F} of leaves of Veech’s foliation, the holonomy is an arithmetic lattice in $\text{PU}(1, 2)$.

This follows from the following lemma:

LEMMA 11.2. *Let \mathcal{F} be a connected component of a leaf $\mathcal{F}_{[\rho]}$. Then, up to a suitable conjugation, the coefficients of the matrices of $\text{hol}(\pi_1(\mathcal{F}))$ lie in $\mathbb{Z}[\text{Im}(\rho)]$.*

This lemma is an easy consequence of the fact that the matrices of the transitions maps of an atlas of linear parametrisations coming from topological polygonations

Table 1 Data θ and M such that $\text{Im}(\rho) = \langle \exp(2i\pi/m) \rangle$ with $m \in \{3, 4, 6\}$ (n stands for the number of cone points and coincides with the dimension of the leaf $\mathcal{F}_\theta(M)$ plus one)

n	$m \theta / 2\pi = (m \theta_i / 2\pi)_{i=1}^n$	m	M	label
3	(5, 2, 2)	3	1	a
	(6, 3, 3)	4	1	b
	(8, 5, 5)	6	1	c
3	(7, 3, 2)	4	1	d
	(9, 5, 4)	6	1	e
	(10, 5, 3)	6	1	f
	(11, 5, 2)	6	1	g
	(11, 4, 3)	6	1	h
	(5, 2, 2)	3	2	i
4	(7, 3, 3, 3)	4	1	j
	(9, 5, 5, 5)	6	1	k
	(10, 5, 5, 4)	6	1	l
	(11, 5, 4, 4)	6	1	m
5	(10, 5, 5, 5, 5)	6	1	n
	(11, 5, 5, 5, 4)	6	1	o
6	(11, 5, 5, 5, 5, 5)	6	1	p

must have coefficients in $\mathbb{Z}[\text{Im}(\rho)]$ (see Section 4 and the proof of Proposition 4.2). In particular if $\mathbb{Z}[\text{Im}(\rho)]$ is discrete in \mathbb{C} , then for any connected component \mathcal{F} of $\mathcal{F}_{[\rho]}$, the image of its holonomy is discrete in $\text{PU}(1, 2)$.

The developing map $\tilde{\mathcal{F}} \rightarrow \mathbb{C}\mathbb{H}^2$ of such a \mathcal{F} factors through a local isometry $\mathcal{F} \rightarrow \mathbb{C}\mathbb{H}^2/\text{hol}(\pi_1(\mathcal{F}))$. Since \mathcal{F} has finite volume (cf. Section 8), $\text{hol}(\pi_1(\mathcal{F}))$ is necessarily a lattice which must be arithmetic since it belongs to $\text{PU}(1, 2) \cap \text{SL}_3(\mathbb{Z}[\text{Im}(\rho)])$. This situation actually happens: if $\text{Im}(\rho) = \langle \exp(2i\pi/m) \rangle$ for $m = 3, 4$ or 6 then $\mathbb{Z}[\text{Im}(\rho)]$ is discrete. Note that the argument does apply to higher dimensions as well.

By straightforward computations, these cases can be completely determined:

- PROPOSITION 11.3. (1) *The data θ and M such that $\text{Im}(\rho) = \langle \exp(2i\pi/m) \rangle$ with $m \in \{3, 4, 6\}$ are exactly those given in TABLE 1.*
 (2) *For any such θ and M , the (image of the) complex hyperbolic holonomy of any connected component of $\mathcal{F}_\theta(M)$ is an arithmetic lattice.*

We think it is relevant to make the following remarks regarding TABLE 1:

Table 2 Coincidences with some Picard/Deligne–Mostow lattices

ℓ	$\theta(\ell)$	$\mu(\ell)$	label in [DM86, p. 86]
a	$(\frac{10\pi}{3}, \frac{4\pi}{3}, \frac{4\pi}{3})$	$(\frac{1}{6}, \frac{1}{3}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$	6
b	$(3\pi, \frac{3\pi}{2}, \frac{3\pi}{2})$	$(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$	2
c	$(\frac{8\pi}{3}, \frac{5\pi}{3}, \frac{5\pi}{3})$	$(\frac{1}{3}, \frac{1}{6}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$	6

REMARK 11.4. (1) Using the same approach as the one considered in [GP, §4.3.1], it is not difficult to establish that $\mathcal{F}_\theta(M)$ is connected for any one of the elements (θ, M) of TABLE 1.

(2) From the preceding remark, it follows that to any label ℓ among the sixteen of TABLE 1 corresponds precisely one arithmetic complex hyperbolic lattice in $\mathrm{PU}(1, n-1)$, which will be denoted by Γ_ℓ .

(3) The lattices Γ_ℓ associated to the first three labels a, b and c were previously known. Indeed, for such a label, the associated angle datum $\theta(\ell) = (\theta_i(\ell))_{i=1}^3$ is such that $\theta_2(\ell) = \theta_3(\ell)$. From this and because $M = 1$, one deduces that the flat structure of a flat tori T whose class belongs to $\mathcal{F}_{\theta(\ell)}(1)$ is invariant by the elliptic involution $i \circlearrowleft T$. Consequently, the flat structure of T comes from a flat structure on $T/i \simeq \mathbb{P}^1$ with five cone points. It follows that there exists a (possibly orbifold) covering $\mathcal{F}_{\theta(\ell)}(1) \rightarrow \mathcal{M}_{0, \tilde{\theta}(\ell)}$ with $\tilde{\theta}(\ell) = (\theta_1(\ell)/2, \theta_2(\ell), \pi, \pi, \pi)$ (see [GP, §4.2.5] for more details).

This eventually gives us that Γ_ℓ coincides with a Picard/Deligne–Mostow lattice $\Gamma_{\mu(\ell)}$ associated to a 5-tuple $\mu(\ell)$ (we use here the notations of [DM86] and [Mos88]) which depends only on ℓ , see TABLE 2. Note that it shows that Γ_a and Γ_c coincide (up to conjugacy).

(4) Excluding the cases associated to the labels a, b and c , one gets thirteen *a priori* new arithmetic complex hyperbolic lattices. It would be interesting to know more about them. To achieve this goal, the following approach, albeit computational, could be fruitful: for any label $\ell \in \{a, b, \dots, p\}$, using the results of [GP, §4.2.4], it is possible to give an explicit basis of the fundamental group of the corresponding leaf $\mathcal{F}_\ell = \mathcal{F}_{\theta(\ell)}(M(\ell))$. Then, combining the results of [GP, §4.4] with the monodromy formulae of [Man08, §6], it should be possible to construct an explicit version $h_\ell : \pi_1(\mathcal{F}_\ell) \rightarrow \mathrm{PU}(1, n-1)$ of the holonomy map which would allow to study $\Gamma_\ell = \mathrm{Im}(h_\ell)$ quite concretely.

11.4 Holonomy in the 1-Dimensional Case. Subsection 10.5 is the first step towards a geometric description of the moduli spaces $\mathcal{F}_{[\rho]}$ when $g = 1$ and $n = 2$. A comprehensive description of this case is carried on in the paper [GP].

The following Proposition, which is just a suitable reformulation of Poincaré’s theorem on fundamental domains of Fuchsian groups, gives an easily verifiable

sufficient criterion for a (connected component of a) leaf \mathcal{F} to have discrete holonomy in $\mathrm{PU}(1, 1)$.

PROPOSITION 11.5. *The metric completion $\overline{\mathcal{F}}$ of \mathcal{F} is a lattice quotient of $\mathbb{C}\mathbb{H}^1$ if and only if all the cone angles at points of $\overline{\mathcal{F}} \setminus \mathcal{F}$ are integer parts of 2π .*

This Proposition combined with the analysis carried on in [GP] allows us to find several such \mathcal{F} which are lattice quotients (cf. [GP, §6.1] for more details).

Appendix A. Complex Hyperbolic Geometry

A.1 Complex Hyperbolic Space. On the complex vector space \mathbb{C}^{n+1} of dimension $n + 1$, we consider the Hermitian form $\langle \cdot, \cdot \rangle$ of signature $(1, n)$ defined by

$$\langle z, w \rangle = z_0 \overline{w_0} - \sum_{i=1}^n z_i \overline{w_i}$$

for $z = (z_0, \dots, z_n)$ and $w = (w_0, \dots, w_n)$ in \mathbb{C}^{n+1} .

All the definitions to come do not depend on the choice of the Hermitian metric of signature $(1, n)$ since two such forms are linearly conjugate. Recall that $\mathbb{C}\mathbb{P}^n$ is the set of complex lines in \mathbb{C}^{n+1} . We define $\mathbb{C}\mathbb{H}^n$, the **complex hyperbolic space of dimension n** , as the subset of $\mathbb{C}\mathbb{P}^n$ formed by the lines in \mathbb{C}^{n+1} on which $\langle \cdot, \cdot \rangle$ is positive:

$$\mathbb{C}\mathbb{H}^n = \left\{ [z] \in \mathbb{C}\mathbb{P}^n \mid z \in \mathbb{C}^{n+1}, \langle z, z \rangle > 0 \right\}.$$

We denote by $\mathrm{PU}(1, n)$ the set of linear automorphisms of \mathbb{C}^{n+1} which preserve $\langle \cdot, \cdot \rangle$. It acts projectively on $\mathbb{C}\mathbb{H}^n$ and satisfies the following properties:

- its action on $\mathbb{C}\mathbb{H}^n$ is transitive;
- $\mathrm{PU}(1, n)$ is exactly the group $\mathrm{Aut}(\mathbb{C}\mathbb{H}^n)$ of biholomorphisms of $\mathbb{C}\mathbb{H}^n$;
- there exists a Riemannian metric on $\mathbb{C}\mathbb{H}^n$ (unique up to rescaling), for which $\mathrm{PU}(1, n)$ is exactly the set of holomorphic isometries. This metric is called the **complex hyperbolic metric**;
- this metric has sectional curvature comprised between $-\frac{1}{4}$ and -1 . Its holomorphic sectional curvature is constant.

The stabiliser of a point in $\mathbb{C}\mathbb{H}^n$ (which is exactly the stabiliser of a positive line in $\mathrm{PU}(1, n)$) is conjugate to $\mathrm{U}(n) \subset \mathrm{PU}(1, n)$, which is the maximal compact subgroup of $\mathrm{PU}(1, n)$. The complex hyperbolic space $\mathbb{C}\mathbb{H}^n$ is therefore isometric to the rank one (Hermitian) symmetric space $\mathrm{PU}(1, n)/\mathrm{U}(n)$. It is the non-compact dual of $\mathbb{C}\mathbb{P}^n$.

The distance for the complex hyperbolic metric can be explicitly computed by means of the initial Hermitian form:

LEMMA A.1. Let $[z]$ and $[w]$ be two points in $\mathbb{C}\mathbb{H}^n \subset \mathbb{C}\mathbb{P}^n$ with $z, w \in \mathbb{C}^{n+1}$.

(1) The complex hyperbolic distance α between $[z]$ and $[w]$ satisfies

$$\cosh^2\left(\frac{\alpha}{2}\right) = \frac{\langle z, w \rangle \cdot \langle w, z \rangle}{\langle z, z \rangle \cdot \langle w, w \rangle}.$$

(2) The geodesic curve linking $[z]$ to $[w]$ in $\mathbb{C}\mathbb{H}^n$ is the projectivisation of the linear segment $[z, w] = \{z + tw \mid t \in [0, 1]\}$ linking z to w in \mathbb{C}^{n+1} .

(For some proofs, see [Gol99, §3.3.5])

A.2 Coordinates.

A.2.1 The Ball Model. In order to have coordinates on $\mathbb{C}\mathbb{H}^n$, one can take affine coordinates of $\mathbb{C}\mathbb{P}^n$. Since $z_0 \neq 0$ if $[z] = [z_0 : \dots : z_n]$ belongs to $\mathbb{C}\mathbb{H}^n$, the latter is contained in the affine chart $\{z_0 \neq 0\}$ of $\mathbb{C}\mathbb{P}^n$.

In the $z_0 = 1$ normalisation, it comes that z_1, \dots, z_n provide a global system of holomorphic coordinates which identify $\mathbb{C}\mathbb{H}^n$ with the complex n -ball:

$$\left\{ (z_i)_{i=1}^n \in \mathbb{C}^n \mid \sum_{i=1}^n |z_i|^2 < 1 \right\}.$$

In this model of the complex hyperbolic space, the hyperbolic metric identifies with the Bergman metric of the complex n -ball.

Although we do not use it in the present text, the complex ball is a very classical model for $\mathbb{C}\mathbb{H}^n$ which is worth being mentioned. We will not say anything more about it but one can find a comprehensive presentation in [Gol99].

A.2.2 Pseudo-Horospherical Coordinates. More important for our purpose is a special kind of affine coordinates on $\mathbb{C}\mathbb{H}^n$ which are very close, in spirit, to the *horospherical coordinates* introduced by Goldman and Parker in [GP92].

Let $\xi = (\xi_0, \dots, \xi_n)$ be a system of linear coordinates on \mathbb{C}^{n+1} such that the expression of the Hermitian form $\langle \cdot, \cdot \rangle$ in these can be written out

$$\langle \xi, \xi \rangle = \frac{i}{2} (\xi_n \bar{\xi}_0 - \xi_0 \bar{\xi}_n) + a(\widehat{\xi}, \widehat{\xi})$$

for a Hermitian form a of signature $(1, n-1)$ and where $\widehat{\xi}$ stands for $(\xi_0, \dots, \xi_{n-1})$.

LEMMA A.2. If $\xi = (\xi_i)_{i=0}^n$ is such that $\langle \xi, \xi \rangle > 0$ then $\xi_0 \neq 0$.

Proof. One verifies that, up to a linear change of coordinates letting ξ_0 invariant, one can assume that $\langle \xi, \xi \rangle = \frac{i}{2} (\xi_n \bar{\xi}_0 - \xi_0 \bar{\xi}_n) + \sum_{j=0}^{n-1} \epsilon_j \xi_j \bar{\xi}_j$ for some ϵ_j belonging to $\{-1, 0, 1\}$. By assumption, $a(\widehat{\xi}, \widehat{\xi}) = \sum_{j=0}^{n-1} \epsilon_j \xi_j \bar{\xi}_j$ has signature $(1, n-1)$ hence exactly one of the ϵ_j 's is equal to 1, all the others being equal to -1.

If $\epsilon_0 = -1$, then $\frac{i}{2} (\xi_n \bar{\xi}_0 - \xi_0 \bar{\xi}_n) - \xi_0 \bar{\xi}_0$ has signature $(1, 1)$. Since $\sum_{j=1}^{n-1} \epsilon_j \xi_j \bar{\xi}_j$ has signature $(1, n-2)$ (because $\epsilon_j = 1$ for some $j \geq 1$), this would imply that $\langle \cdot, \cdot \rangle$ has signature $(2, n-1)$, a contradiction. \square

From the preceding lemma, it follows that the complex hyperbolic space admits a model contained in the affine chart $\{\xi_0 \neq 0\}$ of $\mathbb{C}\mathbb{P}^n$. Then, under the normalization $\xi_0 = 1$, the ξ_k 's for $k = 1, \dots, n$ provide global affine coordinates on this model which will be called **pseudo-horospherical coordinates**.

In such coordinates, the associated quadratic form is given by $\langle \xi, \xi \rangle = \text{Im}(\xi_n) + a(\widehat{\xi}, \widehat{\xi})$ with $\widehat{\xi} = (1, \xi_1, \dots, \xi_{n-1})$ and consequently, this model of the complex hyperbolic space $\mathbb{C}\mathbb{H}^n$ consists in the set of $\xi = (\widehat{\xi}, \xi_n) \in \mathbb{C}^n$ such that

$$\text{Im}(\xi_n) > -a(\widehat{\xi}, \widehat{\xi}).$$

In the standard (homogeneous) coordinates $z = (z_0, z_1, \dots, z_n)$ on \mathbb{C}^{n+1} , the formula for the complex hyperbolic metric is the following

$$g = -\frac{4}{\langle z, z \rangle^2} \begin{vmatrix} \langle z, z \rangle & \langle dz, z \rangle \\ \langle z, dz \rangle & \langle dz, dz \rangle \end{vmatrix}.$$

A straightforward calculation gives the following formula for the expression of this metric in pseudo-horospherical coordinates:

$$g = -\frac{4}{\langle \xi, \xi \rangle^2} \left(\langle \xi, \xi \rangle \cdot a(d\widehat{\xi}, d\widehat{\xi}) - a(\widehat{\xi}, d\widehat{\xi}) \cdot a(d\widehat{\xi}, \widehat{\xi}) - \text{Im}(d\xi_n \cdot a(\widehat{\xi}, d\widehat{\xi})) - |d\xi_n|^2 \right).$$

Introducing $u = \langle \xi, \xi \rangle$ and $s = \text{Re}(\xi_n)$, we therefore have $\xi_n = s + i(u - a(\widehat{\xi}, \widehat{\xi}))$. In the coordinates system $(s, u, \xi_1, \dots, \xi_{n-1})$ on the pseudo-horospherical model of $\mathbb{C}\mathbb{H}^n$ we are considering, the metric tensor g writes down

$$g = \frac{4}{u^2} \left(\frac{du^2}{4} + \left(\frac{ds}{2} + \text{Im}(\omega) \right)^2 + \text{Re}(\omega)^2 - u \cdot \Omega \right) \tag{20}$$

where $\omega = a(\widehat{\xi}, d\widehat{\xi})$ and $\Omega = a(d\widehat{\xi}, d\widehat{\xi})$.

We now introduce the family of open sets in $\mathbb{C}\mathbb{H}^n$:

$$U_{K,\lambda} = \left\{ [1, \xi_1, \dots, \xi_n] \in \mathbb{C}\mathbb{H}^n \mid |\xi_1|, \dots, |\xi_{n-1}|, |\text{Re}(\xi_n)| < K \text{ and } \text{Im}(\xi_n) > \lambda \right\}$$

with $K, \lambda > 0$.

LEMMA A.3. *Let K and λ be arbitrary positive constants.*

- (1) *The complex hyperbolic volume of $U_{K,\lambda}$ is finite.*
- (2) *If $\gamma : [0, 1] \rightarrow U_{K,\lambda}$ is a path such that $\gamma(t) = (\xi_1(t), \dots, \xi_n(t))$ for any $t \in [0, 1]$, then its length $L(\gamma)$ for the complex hyperbolic metric satisfies*

$$L(\gamma) \geq \left| \log(\xi_n(1)) - \log(\xi_n(0)) \right|.$$

Proof. In the coordinates system $(s, u, \xi_1, \dots, \xi_{n-1})$ on $U_{K,\lambda}$, the complex hyperbolic volume element writes down

$$\sqrt{\det(g)} ds du d\xi_1 d\bar{\xi}_1 \cdots d\xi_{n-1} d\bar{\xi}_{n-1}.$$

Since both ω and Ω depend continuously on ξ_1, \dots, ξ_{n-1} , one gets that

$$\sqrt{\det(g)} = \frac{f(\xi_1, \dots, \xi_{n-1})}{u^{2n+2}}$$

for some positive and continuous function f which thereby is bounded on $U_{K,\lambda}$. The finiteness of the volume of $U_{K,\lambda}$ follows directly from evaluating the associated integral.

The second point of the lemma follows directly from the fact that $g \geq u^{-2} du^2$ on $U_{K,\lambda}$. To see this, one has to prove that Ω is negative. But if Ω was not, since $du^2/4 + (ds/2 + \operatorname{Im}(\omega))^2 + \operatorname{Re}(\omega)^2$ does not depend on u , one would deduce from (20) that g would not be positive for large values of u , a contradiction. \square

Appendix B. Cone-Manifolds

B.1 Generalities. This section strongly builds on [McM17], in particular the use of joints for describing spherical cone-manifolds.

Let X be a complete homogeneous Riemannian manifold and let G be its isometry group (or more generally a subgroup of its isometry group). We develop material on cone-manifolds in this specific case. For any point $p \in X$, one denotes by X_p the set of geodesic rays emanating from it and $G_p = \operatorname{Stab}_G(p)$ stands for its stabiliser.

A **(X, G) -cone-manifold** is a geometric object built inductively as follows:

- if X is 1-dimensional, a (X, G) -cone-manifold is just a (X, G) -manifold;
- otherwise, a (X, G) -cone-manifold is a topological space such that any point in it has a neighbourhood isomorphic to a cone over a (X_p, G_p) -cone-manifold.

One remarks that X_p is just the unit sphere at p in X and therefore G_p can naturally be seen as a subgroup of $O(n)$ where n is the dimension of X .

A simple example of a non trivial cone-manifold is a Euclidean cone. If $X = \mathbb{R}^2$ and $G = \operatorname{Iso}(\mathbb{R}^2)$, $X_p = S^1$ and $G_p = O(2)$. A (X_p, G_p) -manifold is nothing else but a circle of length θ and a cone over it is a cone of angle θ . Finally, remark that any (X, G) -manifold is also a (X, G) -cone-manifold in a natural way.

B.2 Cones are Cone-Manifolds. Let X be a connected Riemannian manifold such that G is the component of the identity of its isometry group. Let X' be a totally geodesic submanifold of codimension 2 in X such that $\operatorname{Stab}_G(X')$ is $S^1 = \mathbb{R}/\mathbb{Z}$, *i.e.* it acts by rotation of angle θ around X' for any $\theta \in S^1$.

We explain the general construction of the *cone of angle θ over X'* . The metric completion Y of the universal covering of $X \setminus X'$ is an infinite cyclic cover of X branched along X' . There is a group \mathbb{R} of isometries lifting the action of S^1 by

rotation to Y and if $\theta \in]0, +\infty[$, one defines X_θ the **cone of angle θ over X'** as the quotient of Y by the action of the ‘rotation’ of angle θ on Y . The image in X_θ of the preimage of X' in Y is called the **singular locus of the cone**.

PROPOSITION B.1. X_θ is a (X, G) -cone-manifold.

Proof. The proof goes by induction on the dimension of X . Away from its singular locus, X_θ is a (X, G) -manifold hence the proposition is clear here.

Let p be a point of the singular locus. The set W of points of X_θ that can be joined to p by a geodesic path of length 1 happens to be a cone of angle θ for a sphere S of radius 1 at a point $q \in X'$ with isometry group $\text{Stab}_G(q)$. A neighbourhood of p in X_θ is then the cone over W . We want to show that W is actually a $(S, \text{Stab}_G(q))$ -cone-manifold. This will be done by showing that W is actually a cone of angle θ and applying the induction hypothesis.

The intersection $S' = X' \cap S$ is a totally geodesic submanifold of S for the metric induced by X and $S^1 \subset \text{Stab}_G(X') \subset \text{Stab}_G(q)$. The universal cover of $S \setminus S'$ embeds in the one of $X \setminus X'$ and therefore the metric completion of the universal cover of $S \setminus S'$ embeds in the metric completion Y of the universal cover of $X \setminus X'$. W is then the quotient of the metric completion of the universal cover of $S \setminus S'$ by the rotation of angle θ . Hence W is a $(S, \text{Stab}_G(q))$ -cone and since $\dim(S) = \dim(X) - 1$, is a $(S, \text{Stab}_G(q))$ -cone-manifold. \square

B.3 Joints. We now restrict to the case when $X = \mathbb{C}\mathbb{H}^n$ and $G = \text{PU}(1, n)$. The unit sphere at a point p in X is $S^{2n-1} = \partial(B^n)$ where B^n is the unit ball at p and its isometry group is $\text{U}(n) \subset G$. For every k in $\{1, \dots, n\}$, we can carry on the construction detailed below.

The joint $A * B$ of two topological spaces A and B is the space you get by adjoining to every pair of points $(a, b) \in A \times B$ a segment $[a, b]$. This operation can be made geometrical if A and B are spherical manifolds. One remarks that $S^{2(n+k)-1}$ is the joint of $S^{2n-1} * S^{2k-1}$ where S^{2n-1} and S^{2k-1} are embedded in $S^{2(n+k)-1}$ in a essentially unique way such that every points $x \in S^{2n-1}$ and $y \in S^{2k-1}$ are joined by a unique geodesic path of length $\frac{\pi}{2}$. This makes it very clear how one can endow the joint of X a $(S^{2n-1}, \text{U}(n))$ -manifold and Y a $(S^{2k-1}, \text{U}(k))$ -manifold with the structure of a $(S^{2(n+k)-1}, \text{U}(n+k))$ -manifold. A good reference that deals with this construction is [BH99, Chapter I.5, p.63].

This property of naturality extends in some way to cone-manifolds.

LEMMA B.2. Let M be a $(S^{2k-1}, \text{U}(k))$ -cone-manifold. Then the joint $S^{2(n-k)-1} * M$ has a natural structure of $(S^{2n-1}, \text{U}(n))$ -cone-manifold.

Proof. The proof goes by double induction on n and $i = (n - k)$. To be more precise, we assume that the lemma is true for all (n', k') such that either $n' < n$ or $n' = n$ and $k < k'$. Take $p \in S^{2(n-k)-1} * M$. We distinguish two cases:

- (1) **p does not belong to $S^{2(n-k)-1}$.** In that case p belongs to an arc $[x, y]$ with $x \in S^{2(n-k)-1}$ and $y \in M$. Denote by M_i the union of the strata of codimension i of M . If $i = 0$, i.e. if y is a regular point in M , then p is a regular point of

$S^{2(n-k)-1} * M$. For $i \geq 1$, $S^{2(n-k)-1} * M_i$ is a $(S^{2n-3}, U(n-1))$ -cone-manifold by the induction hypothesis. In that case p has a neighbourhood which is a cone over the joint $S^{2(k+i)-1} * V(y)$ where $V(y)$ is a $(S^{2(n-k-i)-1}, U(n-k-i))$ -cone-manifold over which a neighbourhood of y in M_i is a cone.

- (2) **p belongs to $S^{2(n-k)-1}$.** In that case a neighbourhood of p in $S^{2(n-k)-1} * M$ is a cone over the joint $S^{2(n-k-1)-1} * M$ and the induction hypothesis allows to conclude. \square

B.4 Strata. A $\mathbb{C}\mathbb{H}^n$ -cone-manifold X has a stratified structure $X_0 \sqcup X_1 \sqcup \dots \sqcup X_n$ where X_k is a $\mathbb{C}\mathbb{H}^{n-k}$ -manifold whose metric completion is $X_k \sqcup \dots \sqcup X_n$. Every point $p \in X$ has a neighbourhood which is the cone over the joint $S^{2(n-i)-1} * X(p)$ where $X(p)$ is a $(S^{2i-1}, U(i))$ -cone-manifold. X_k is defined as the set of points for which the biggest integer i for which a neighbourhood of p has the latter structure is equal to k .

B.5 Totally Geodesic Subcone-Manifolds. We assume here that X is a Riemannian manifold which is either $\mathbb{C}\mathbb{H}^n$ or S^k and G is either $\text{PU}(1, n)$ or a subgroup of $O(k)$. X_p is the unit sphere at a point $p \in X$ and $G_p = \text{Stab}_G(\{p\})$. If X is a (X, G) -cone-manifold, a totally geodesic sub-cone-manifold Y of X is a subset of X such that the intersection of Y with each stratum of X is a totally geodesic submanifold of the stratum.

LEMMA B.3. *Let p be a point of X and Y be a totally geodesic submanifold of X such that $p \in Y$. Then $X_p \cap Y$ is a totally geodesic submanifold of X_p .*

Proof. This is a consequence of the fact that in all the cases we are considering there exists a subgroup G'_p of G_p such that $\text{Stab}(G'_p) = Y$. \square

PROPOSITION B.4. *A totally geodesic subcone-manifold Y of a Riemannian cone-manifold M endowed with the natural metric structure coming from its embedding is also a cone-manifold.*

Proof. The proof goes by induction on $\dim(Y)$. Take q in Y . This point has a neighbourhood in M which is a cone over a (X_p, G_p) -manifold X' , where X_p is the unit sphere at a point $p \in X$ and $G_p = \text{Stab}_G(\{p\})$. According to Lemma B.3, $X' \cap Y$ is also a totally geodesic cone manifold of dimension $\dim(Y) - 1$. The induction hypothesis ensures that $X' \cap Y$ is also a cone-manifold and therefore p has a neighbourhood which is a cone over a cone-manifold. \square

B.6 Higher Dimensional Complex Hyperbolic Cones. We now give local models for some specific complex hyperbolic cone-manifolds. In particular we generalise the notion of cone previously defined in the particular case of complex hyperbolic geometry. Let X be a complete complex hyperbolic cone-manifold of dimension k and let p a point being on a stratum of codimension k . We denote by X_0 the set of regular points, which is open in X . Consider the trivial product $\mathbb{C}\mathbb{H}^n \times X_0$. There is a unique complex hyperbolic structure on $\mathbb{C}\mathbb{H}^n \times X_0$ such that

- each fiber $\{*\} \times X_0$ is locally totally geodesic;
- any fiber $\{*\} \times X_0$ intersects $\mathbb{C}\mathbb{H}^n \times \{p\}$ orthogonally.

The metric completion of $\mathbb{C}\mathbb{H}^n \times X_0$ is then $\mathbb{C}\mathbb{H}^n \times X$. Now is the good moment to explain the notion of orthogonality in a $(\mathbb{C}\mathbb{H}^n, \text{PU}(1, n))$ -cone-manifold. Let Y and Z be two totally geodesic sub-cone-manifolds of X , a $(\mathbb{C}\mathbb{H}^n, \text{PU}(1, n))$ -cone-manifold. Assume that Y and Z intersect only at a point p . We say that they intersect orthogonally if every pair of regular points $p \in Y$ and $q \in Z$ is contained in an open set U of X such that

- U is isometric to an open subset of $\mathbb{C}\mathbb{H}^n$;
- $Y \cap U$ and $Z \cap U$ are respectively identified with open subsets of copies of $\mathbb{C}\mathbb{H}^i$ and $\mathbb{C}\mathbb{H}^j$ in $\mathbb{C}\mathbb{H}^n$ which intersect orthogonally.

PROPOSITION B.5. *$\mathbb{C}\mathbb{H}^n \times X$ seen as the metric completion of $\mathbb{C}\mathbb{H}^n \times X_0$ is a complex hyperbolic cone-manifold.*

Proof. Let q be a point in X which has maximal codimension. There is a neighbourhood of q in X which is a cone over a $(S^{2k-1}, U(k))$ -cone-manifold X' . According to Lemma B.2, the spherical joint $X' * S^{2n-1}$ has a natural structure of $(S^{2(n+k)-1}, U(n+k))$ -cone-manifold, of which a neighbourhood of q in $\mathbb{C}\mathbb{H}^n \times X$ is a cone over. \square

Acknowledgments

We are thankful to Adrien Boulanger for some interesting discussions about the notion of holonomy. We are also grateful to Richard Schwartz for kindly answering several technical questions about the notion of cone manifold, as well as to John Parker and Curtis McMullen for useful correspondences. We are very thankful to Bertrand Deroin for the constant support and deep interest he has shown in our work since its very beginning. We are indebted to Irene Pasquinelli and John Parker for very many useful comments and suggestions on the substance and structure of this paper, which have rendered it much clearer. Finally, the second author thanks Brubru for her many corrections.

While investigating moduli spaces of flat metrics, we have taken an interest in historical questions related to the genesis of the papers [DM86] and [Thu98]. We are thankful to Nicolas Bergeron, Pierre Deligne, Étienne Ghys, Curtis McMullen, Jean-Pierre Otal, John Parker, Marc Troyanov and William Veech for taking the time to answer some of our historically flavoured questions. Finally, we are grateful to the two anonymous referees for their careful readings of the paper and valuable comments and suggestions

References

- [ACG11] E. ARBARELLO, M. CORNALBA, and P.A. GRIFFITHS. *Geometry of algebraic curves. Volume II*, volume 268 of *Grundlehren der Mathematischen Wis-*

- senschaften [Fundamental Principles of Mathematical Sciences]*. Springer, Heidelberg, (2011). With a contribution by Joseph Daniel Harris.
- [BH99] M. R. BRIDSON and A. HAEFLIGER. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, (1999).
- [BP15] R. BOADI and J. PARKER. Mostow's lattices and cone metrics on the sphere. *Adv. Geom.*, 15(1) (2015), 27–53
- [DM86] P. DELIGNE and G. D. MOSTOW. Monodromy of hypergeometric functions and nonlattice integral monodromy. *Inst. Hautes Études Sci. Publ. Math.*, (63) (1986), 5–89
- [DM93] P. DELIGNE and G. D. MOSTOW. *Commensurabilities among lattices in $PU(1, n)$* , volume 132 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, (1993).
- [DPP15] M. DERAUX, J. PARKER, and J. PAUPERT. New non-arithmetic complex hyperbolic lattices. *Invent. Math.*, 203(3) (2016), 681–771
- [DS05] F. DIAMOND and J. SHURMAN. *A first course in modular forms*, volume 228 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, (2005).
- [EMZ03] A. ESKIN, H. MASUR, and A. ZORICH. Moduli spaces of abelian differentials: the principal boundary, counting problems, and the Siegel-Veech constants. *Publ. Math. Inst. Hautes Études Sci.*, (97) (2003), 61–179
- [GLL11] A. GONZÁLEZ and J. LÓPEZ-LÓPEZ. Shapes of tetrahedra with prescribed cone angles. *Conform. Geom. Dyn.*, 15 (2011), 50–63
- [Gol99] W. M. GOLDMAN. *Complex hyperbolic geometry*. Oxford Mathematical Monographs. The Clarendon Press, Oxford University Press, New York, (1999). Oxford Science Publications.
- [GP] S. GHAZOUANI and L. PIRIO. Moduli spaces of flat tori and elliptic hypergeometric functions. Preprint [arXiv:1605.02356](https://arxiv.org/abs/1605.02356).
- [GP92] W. M. GOLDMAN and J. R. PARKER. Dirichlet polyhedra for dihedral groups acting on complex hyperbolic space. *J. Geom. Anal.*, 2(6) (1992), 517–554
- [Gro99] M. GROMOV. *Metric structures for Riemannian and non-Riemannian spaces*, volume 152 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 1999. Based on the 1981 French original [MR0682063 (85e:53051)], With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.
- [JV01] D. JACKSON and T. VISENTIN. *An atlas of the smaller maps in orientable and nonorientable surfaces*. CRC Press Series on Discrete Mathematics and its Applications. Chapman & Hall/CRC, Boca Raton, FL, (2001).
- [KZ03] M. KONTSEVICH and A. ZORICH. Connected components of the moduli spaces of Abelian differentials with prescribed singularities. *Invent. Math.*, 153(3) (2003), 631–678
- [LV93] R. LE VAVASSEUR. Sur le système d'équations aux dérivées partielles simultanées auxquelles satisfait la série hypergéométrique à deux variables $F_1(\alpha, \beta, \beta', \gamma; x, y)$. *Ann. Fac. Sci. Toulouse Sci. Math. Sci. Phys.*, 7(3) (1893), 1–120
- [Man08] TOSHIYUKI MANO. The Riemann-Wirtinger integral and monodromy-preserving deformation on elliptic curves. *Int. Math. Res. Not. IMRN*, pages Art. ID rnn110, 19, 2008.
- [McM17] C. T. MCMULLEN. The gauss-bonnet theorem for cone manifolds and volumes of moduli spaces. *Amer. J. Math.*, 139(1) (2017), 261–291

- [Mos86] G. D. MOSTOW. Generalized Picard lattices arising from half-integral conditions. *Inst. Hautes Études Sci. Publ. Math.*, (63) (1986), 91–106
- [Mos88] G. D. MOSTOW. On discontinuous action of monodromy groups on the complex n -ball. *J. Amer. Math. Soc.*, 1(3) (1988), 555–586,
- [MS91] H. MASUR and J. SMILLIE. Hausdorff dimension of sets of nonergodic measured foliations. *Ann. of Math. (2)*, 134(3) (1991), 455–543
- [MZ08] H. MASUR and A. ZORICH. Multiple saddle connections on flat surfaces and the principal boundary of the moduli spaces of quadratic differentials. *Geom. Funct. Anal.*, 18(3) (2008), 919–987
- [Par06] J. R. PARKER. Cone metrics on the sphere and Livné’s lattices. *Acta Math.*, 196(1) (2006), 1–64
- [Pas16] I. PASQUINELLI. Deligne-Mostow lattices with three fold symmetry and cone metrics on the sphere. *Conform. Geom. Dyn.*, 20 (2016), 235–281
- [Sau90] J. SAUTER. Isomorphisms among monodromy groups and applications to lattices in $PU(1, 2)$. *Pacific J. Math.*, 146(2) (1990), 331–384
- [Sch] R. E. SCHWARTZ. Notes on shape of polyhedra. *Preprint arXiv:1506.07252*
- [Ter73] T. TERADA. Problème de Riemann et fonctions automorphes provenant des fonctions hypergéométriques de plusieurs variables. *J. Math. Kyoto Univ.*, 13 (1973), 557–578
- [Tho15] N. THOLOZAN. Sur la complétude de certaines variétés pseudo-riemanniennes localement homogènes. *Ann. Inst. Fourier.*, 65(5) (2015), 1921–1952
- [Thu97] W. P. THURSTON. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, (1997). Edited by Silvio Levy.
- [Thu98] W. P. THURSTON. Shapes of polyhedra and triangulations of the sphere. In: *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 511–549. Geom. Topol. Publ., Coventry, (1998).
- [Tro86] M. TROYANOV. Les surfaces euclidiennes à singularités coniques. *Enseign. Math. (2)*, 32(1-2) (1986), 79–94
- [Tro07] M. TROYANOV. On the moduli space of singular Euclidean surfaces. In: *Handbook of Teichmüller theory. Vol. I*, volume 11 of *IRMA Lect. Math. Theor. Phys.*, pages 507–540. Eur. Math. Soc., Zürich, (2007).
- [Vee93] W. VEECH. Flat surfaces. *Amer. J. Math.*, 115(3) (1993), 589–689
- [Web93] M. WEBER. *Fundamentaltbereiche komplex hyperbolischer Flächen*. Bonner Mathematische Schriften [Bonn Mathematical Publications], 254. Universität Bonn, Mathematisches Institut, Bonn, (1993).
- [Wol11] J. WOLF. *Spaces of constant curvature*. AMS Chelsea Publishing, Providence, RI, sixth edition, (2011).

SELIM GHAZOUANI, Mathematics Institute University of Warwick Coventry CV4 7AL,UK
s.ghazouani@warwick.ac.uk

LUC PIRIO, LMV (Laboratoire de Mathématiques de Versailles), CNRS, UMR 8100 & Université de Versailles Saint-Quentin, 45 avenue des États-Unis, 78035, Versailles cedex, France
luc.pirio@uvsq.fr

Received: March 21, 2017

Revised: July 25, 2017

Accepted: October 13, 2017