

Similarity structures on the torus and the Klein bottle via triangulations

Stefano Francaviglia

(Communicated by G. Gentili)

Abstract. In this paper we study the possibility of defining a similarity structure on the torus and the Klein bottle using the combinatorial data of a triangulation. Given a choice of moduli for the triangles of a triangulation of a surface, the problem is to decide whether such moduli are compatible with a global similarity structure on the surface.

We study this problem under two different viewpoints. From one side we look at the combinatorial data of triangulations, and we develop an algorithmic method, which allows us to reduce the general problem to a simpler one, which is easily solved. From the other side we study the problem more algebraically, looking at the properties of the holonomy, and we give a complete characterization of the choices of moduli defining global similarity structures on the torus (or on the Klein bottle).

1 Introduction

A similarity structure on a surface is a smooth atlas in which the changes of charts are similarities. A similarity structure on an oriented triangle can be described by a complex number, called modulus, that lies in $\pi_+ = \{z \in \mathbb{C} : \text{Im}(z) > 0\}$. Therefore, once an orientation for each triangle has been fixed, a similarity structure on a triangulated surface determines a modulus for each triangle. Conversely, given a choice of moduli for the triangles of a triangulation of a surface, it is natural to ask when such moduli are compatible with a global similarity structure on the surface.

The problem of the study of similarity structures on the torus T and the Klein bottle K via triangulations arises in a natural way as a “boundary” problem in the study of hyperbolic 3-manifolds. First of all because if M is a complete hyperbolic manifold of finite volume, then the ends of M are homeomorphic to $T \times (0, \infty)$ or $K \times (0, \infty)$ and horospherical sections give Euclidean structures on T or K . Moreover, the triangulations of tori and Klein bottles are strictly related to the ideal triangulations of 3-manifolds.

More precisely, let M be a 3-manifold homeomorphic to the interior of a compact manifold \overline{M} whose boundary consists of tori and Klein bottles. A topological ideal triangulation of M is a description of M as a gluing of ideal tetrahedra, i.e. tetrahe-

dra with vertices removed. Any ideal triangulation of M induces a triangulation of $\partial\bar{M}$. Given an ideal triangulation of M one can try to define a finite-volume hyperbolic structure on M by defining it on each tetrahedron and then by requiring that such structures are compatible with a global one on M . As in the case of similarity structures on triangles, a hyperbolic structure on an ideal tetrahedron can be described by a complex modulus, so the problem can be described in terms of choices of moduli for a given triangulation. Moreover a choice of moduli for the triangulation of M induces a choice of moduli for the triangulation of $\partial\bar{M}$.

It is well-known that if the moduli are all in π_+ and define a hyperbolic structure on M , then they satisfy a system \mathcal{C} of algebraic equations called compatibility equations, and the structure is complete if and only if the moduli satisfy also the so-called completeness equations \mathcal{M} . See [10], [8], [7], [9], [1], [5], [4], [3] for details about ideal triangulations of hyperbolic manifolds.

Restricted to a component $C \subset \partial\bar{M}$, the equations \mathcal{C} and \mathcal{M} have the following meaning: if the moduli induced on the triangulation of C define a similarity structure on C then they satisfy \mathcal{C} , and such a structure is Euclidean (up to scaling) if and only if they satisfy also \mathcal{M} . In particular, to require the structures on the components of $\partial\bar{M}$ to be Euclidean corresponds to require the completeness of the hyperbolic structure of M .

One of the main applications of the technique of ideal triangulations is the proof of the Thurston hyperbolic Dehn filling Theorem ([10], [8], [7]). The idea of the proof is to start with a geodesic ideal triangulation of a complete finite-volume hyperbolic 3-manifold M , in which the tetrahedra have structures whose moduli satisfy \mathcal{C} and \mathcal{M} and then to perturb the moduli along the space of solutions of \mathcal{C} . Finally, one looks at the completions of the structures obtained in this way near the complete structure. This line leads to the study of solutions of \mathcal{C} near a solution of \mathcal{C} and \mathcal{M} . This problem has been studied in [8] using particular ideal triangulations in which the moduli of the initial solution of \mathcal{C} and \mathcal{M} lie in $\pi_+ \cup (\mathbb{R} \setminus \{0, 1\})$. The notion of modulus of a triangle (or an ideal tetrahedron) can be further generalized, allowing the moduli to lie in $\pi_- = \{z \in \mathbb{C} : \text{Im}(z) < 0\}$. The geometric meaning of a triangle with negative modulus is that its orientation is reversed, and this produces overlapping phenomena like the one pictured in Figure 2. The problem of whether a choice of moduli in $\mathbb{C} \setminus \{0, 1\}$ defines a (even incomplete) hyperbolic structure on M is surprisingly difficult if compared with the classical one in which only positive moduli appear, and in general the answer is still unknown (see [7], [8], [9], [5], [4], [3] for details).

At the level of the boundary, this translates to understanding when a solution of \mathcal{C} defines a similarity structure on the torus or the Klein bottle and to the study of similarity structures near Euclidean ones.

The main results of this paper are the following. We show that a choice of moduli defines a Euclidean structure on the torus if and only if it is a solution of \mathcal{C} and \mathcal{M} with non-zero area (Theorems 10.8 and 10.9 and Proposition 10.11). Moreover we show that the set of solutions of \mathcal{C} defining structures on T is a open subset of the set of solutions of \mathcal{C} (see Proposition 10.10). These results are particularly interesting if compared with the corresponding statements in the three-dimensional case (where, for example, there exists solutions of \mathcal{C} and \mathcal{M} with non-zero volume that do not define hyperbolic structures; see [5], [4], [3]).

The paper is structured as follows. In Sections 2–4 we give some definitions and preliminary facts. In Sections 5–9 we describe the algorithm that we use to manipulate triangulations on the torus, based on a geometric version of the topological diagonal swap. Under a supplementary hypothesis, using our algorithm we are able to reduce a given triangulation to one having only two triangles and which is equivalent to the initial one (see below for details). In Section 10 we describe a necessary and sufficient algebraic condition on the moduli in order for a solution of \mathcal{C} or \mathcal{C} and \mathcal{M} to define on T a similarity structure or a Euclidean one up to scaling. We notice that this section is self-contained and can be read independently from Sections 5–9. In Section 11 we reduce the case of the Klein bottle to the case of the torus.

Acknowledgement. I would like to thank Professor C. Petronio for bringing the problem of this paper to my attention and for the interesting and helpful discussions.

2 Preliminaries

First of all, we explain what we mean by triangulation. A classical triangulation of a surface is a subdivision of the surface in embedded triangles. In our triangulations we allow the triangles to have multiple and self-adjacencies, and so, for example, cases like those pictured in Figure 1 possibly do appear.

Combinatorially, a triangulation of a surface is a presentation of the surface as the quotient $\bigsqcup \Delta_i / \sim$, where $\bigsqcup \Delta_i$ is the disjoint union of a finite number of copies of the standard 2-simplex and \sim is the equivalence relation generated by a set of face-pairing maps between the edges of the triangles.

In the sequel, if there are no ambiguities, we will omit the distinction between the abstract and the immersed copy of a triangle of a triangulation.

From now on we fix the following notation: T will be an oriented torus, τ will be a triangulation on T with triangles Δ_i , $i \in I$, where I has a finite cardinality denoted by $|I|$.

We study when it is possible to define a similarity structure on T by defining it separately on each triangle, and then by gluing the geometric version of the triangles using the rules of τ . In other words, we study when it is possible to realize τ as a triangulation in which each triangle has a fixed similarity structure and the face pairing maps consist of similarities between the edges.

A similarity structure on an oriented triangle Δ is a Euclidean structure up to scaling. Once a vertex v of Δ has been fixed, similarity structures on Δ correspond




Figure 1. Auto-adjacencies and multiple adjacencies

to complex numbers in the upper half-plane π_+ : given $z \in \pi_+$ one takes in \mathbb{C} the triangle t_z with vertices in $0, 1, z$ and the only orientation-preserving simplicial homeomorphism $\varphi_z : \Delta \rightarrow t_z$ such that $\varphi(v) = 0$, and defines the similarity structure on Δ induced by z as the pull-back of the structure of t_z via φ_z . Given a modulus $z \in \pi_+$, and provided that a vertex v of Δ is fixed, we call a map $\varphi : \Delta \rightarrow \mathbb{C}$ a *similarity map* if there exists an orientation-preserving similarity Φ of \mathbb{C} such that $\Phi \circ \varphi = \varphi_z$.

We now extend our notions of similarity structure and similarity map by allowing flat triangles (but with distinct vertices) and inverted triangles. In terms of moduli this corresponds to taking z in $\mathbb{R} \setminus \{0, 1\}$ or in $\pi_- = -\pi_+$. Even in these cases z determines a map $\varphi_z : \Delta \rightarrow t_z$ (provided v is fixed as above). For $z \in \mathbb{R} \setminus \{0, 1\}$ we require φ_z to be simplicial and orientation-preserving on all edges. Note that φ_z cannot be a homeomorphism, nevertheless φ_z is regular, i.e. for each $x \in t_z$ the fiber $\varphi_z^{-1}(x)$ is connected and the set of the fibers is a codimension-1 foliation of Δ . Finally, since the triangle t_z changes its orientation when z passes from π_+ to π_- , and since we want that the maps φ_z to depend continuously on z , for $z \in \pi_-$ we require φ_z to be orientation-reversing.

Therefore, a similarity structure on Δ is defined by three moduli z_1, z_2, z_3 (one for each vertex). It is easy to see that if $z_1 \in \pi_+$ (respectively π_- , $\mathbb{R} \setminus \{0, 1\}$) then also z_2 and z_3 are in π_+ (respectively π_- , $\mathbb{R} \setminus \{0, 1\}$). Moreover, if z_1, z_2, z_3 are chosen in a counterclockwise way (according to the orientation of $\partial\Delta$), then they are in the following relationship

$$z_2 = \frac{1}{1 - z_1} \quad z_3 = 1 - \frac{1}{z_1}. \quad (1)$$

In the sequel we will tacitly assume a vertex to be fixed on each triangle we consider, so we parameterize the (possibly flat or inverted) similarity structures on the triangles by complex numbers different from 0 and 1.

Definition 2.1. We call *positive*, *negative* or *flat* a modulus lying respectively in π_+ , π_- or $\mathbb{R} \setminus \{0, 1\}$. We call a triangle *positive*, *negative* or *flat* if it has respectively positive, negative or flat modulus. We say that z and w are inverse moduli if $zw = 1$.

To find a structure on T , we start by assigning a modulus z_i to each triangle Δ_i of τ . We write $\mathbf{z} = (z_1, \dots, z_{|I|})$ to mean a choice of moduli. Both τ and \mathbf{z} lift to a triangulation with moduli of the universal covering $\mathbb{R}^2 \rightarrow T$. If there are no ambiguities we will use the same notation for τ, \mathbf{z} , a triangle Δ_i , and their lifts.

The question is when does \mathbf{z} define a structure on T . If the moduli are all positive or all negative, then by gluing the geometric versions of the triangles according to the identification rules of τ , we obtain a torus, and the problem is when is the choice of moduli coherent. When the moduli are of different nature, then it is not clear a priori how to define the identification space itself, when this space is a torus, and in which sense \mathbf{z} defines a structure on T . Namely, when a positive triangle and a negative one are glued along an edge, then they geometrically overlap and such an overlapping might involve other triangles which might overlap with other triangles and so on (see Figure 2).



Figure 2. Overlapping of two triangles

We omit here a discussion about (X, G) -structures of topological spaces. We recall that an (X, G) -structure on a topological space M is an atlas of charts modeled on a base-model space X in which the change of charts lie in a fixed group G of isomorphisms of X . A similarity structure on a surface is a $(\mathbb{C}, \text{Aff}(\mathbb{C}))$ -structure.

For such structures the notions of developing map and holonomy are well-defined. Roughly speaking, one can think of the developing map as a map D from \tilde{M} , the universal covering of M , to X constructed by lifting the atlas of M to \tilde{M} and then by gluing together the charts of such an atlas. The fundamental group of M acts by deck transformations on \tilde{M} . The holonomy is a representation $h : \pi_1(M) \rightarrow G$ such that $D(\alpha(x)) = h(\alpha)(D(x))$ for all $x \in \tilde{M}$ and $\alpha \in \pi_1(M)$.

We introduce now a notion of developing map for a triangulation with moduli of T , in which the triangles with moduli play the role of the charts of an atlas. The concept of holonomy will be discussed in the next section.

Definition 2.2. Let $\mathbb{R}^2 \rightarrow T$ be the universal covering of T . A *developing map* (for (τ, \mathbf{z})) is a map $D : \mathbb{R}^2 \rightarrow \mathbb{C}$ such that for each Δ of τ the restriction $D|_{\Delta}$ is a similarity map.

Definition 2.3. Let T' be a torus equipped with a similarity structure. Let \tilde{T}' be the universal covering of T' and let $D : \tilde{T}' \rightarrow \mathbb{C}$ be a developing map for the similarity structure of T' .

A map $\varphi : (T, \tau, \mathbf{z}) \rightarrow T'$ is called a *similarity map* if it lifts to a map $\tilde{\varphi} : \mathbb{R}^2 \rightarrow \tilde{T}'$ such that $D \circ \tilde{\varphi}$ is a developing map for (τ, \mathbf{z}) .

Definition 2.4. We say that \mathbf{z} defines a *similarity structure* on T if there exists a similarity map φ of degree ± 1 from (T, τ, \mathbf{z}) to a torus T' equipped with a similarity structure. We say that \mathbf{z} defines a *Euclidean structure* (up to scaling) on T if T' is equipped with a Euclidean structure.

If \mathbf{z} defines a similarity structure on T then we can map the triangles Δ_i on T' one after each other, by following the rules of τ and according with the moduli z_i .

Definition 2.5. The *dual graph* of τ is the trivalent graph Γ whose vertices and edges are respectively the triangles and the edges of τ . Two vertices are joined by an edge e if and only if they are glued together along the edge e in τ .

If Δ is a vertex of Γ we define an *angle* at Δ as a pair of oriented edges (e, f) such that e ends and f starts from Δ . We consider equivalent the angles (e, f) and $(-f, -e)$ (where the minus sign inverts the orientation). A pair $(e, -e)$ is not an angle (while, if e starts and ends at the same vertex, then (e, e) is an angle).

From the definition it follows that angles of Γ correspond to pairs (v, Δ) in τ where v is a vertex of Δ . An orientation of Δ corresponds to a cyclic ordering of the (oriented) edges ending at Δ in Γ . An orientation on T induces, for each vertex v of τ , a cyclic order of the set of the edges of τ ending at v . A modulus for a vertex of Δ in τ corresponds to a modulus for an angle at Δ in Γ .

Remark 2.6. It is easy to see that Γ can be naturally embedded into the 1-skeleton of the first barycentric subdivision of τ . Such an embedding induces a surjective homomorphism of the first homology group.

3 Developing map, holonomy and equations

In the rest of the paper, let \mathbf{z} denote a choice of moduli for τ . In this section we introduce the notion of holonomy for a triangulation with moduli and we discuss the existence of developing maps. To do that we need some natural conditions on the moduli. Such conditions are necessary conditions in order for \mathbf{z} to define a similarity structure on T . We discuss also other conditions on the moduli, which are necessary in order for \mathbf{z} to define a Euclidean structure on T . All these conditions can be expressed by algebraic equations in $(\mathbb{C} \setminus \{0, 1\})^{|\Gamma|}$.

First of all we have the so-called *compatibility equations* \mathcal{C} which say that the product of moduli around each vertex of τ is 1. Historically there are also equations, named \mathcal{C}^* , which require \mathcal{C} to hold and the sum of arguments of moduli around each vertex of τ to be 2π . We do not use these conditions about the arguments because on one hand they are not natural using the techniques we are going to describe, on the other hand they are not necessary (see Section 4 for details).

We describe now the *holonomy representation*. Let Δ_0 be a base-point of Γ and let γ be a path in Γ . Such a path can be viewed as a sequence $(\Delta_0, \dots, \Delta_k)$ of simplices of τ equipped with a fixed set of face-pairing map of τ between Δ_i and Δ_{i+1} . For each i let φ_i be a similarity map $\varphi_i: \Delta_i \rightarrow \mathbb{C}$. We require $\varphi_i = \varphi_j$ if $\Delta_i = \Delta_j$. Let Φ_i be the only orientation-preserving similarity of \mathbb{C} which realizes the face-pairing between $\varphi_i(\Delta_i)$ and $\varphi_{i+1}(\Delta_{i+1})$, i.e. such that $\Phi_i(\varphi_{i+1}(\Delta_{i+1}))$ is glued to $\varphi_i(\Delta_i)$. We define $h(\gamma) = \Phi_1 \circ \dots \circ \Phi_k$. The map h is defined from the set of paths in Γ to $\text{Aff}(\mathbb{C})$ and depends on the choices of the maps φ_i , but it is easy to see that if we restrict to the set of loops in Γ , i.e. with $\Delta_0 = \Delta_k$, then h is well-defined up to conjugation by a similarity of \mathbb{C} .

The fact that equations \mathcal{C} hold implies that h projects to a map $h: \pi_1(T) \rightarrow \text{Aff}(\mathbb{C})$ which is a representation, well-defined up to conjugation. Since for the torus we have $\pi_1(T) \cong H_1(T)$, we can consider h as a map $H_1(T) \rightarrow \text{Aff}(\mathbb{C})$. We call h a *holonomy representation*, *holonomy map* or simply *holonomy*.

Lemma 3.1. *If equations \mathcal{C} hold, then there exists a developing map for (τ, \mathbf{z}) .*

Proof. Let $\mathbb{R}^2 \rightarrow T$ be the universal covering. We recall that we call (τ, \mathbf{z}) both the triangulations on T and on \mathbb{R}^2 . If equations \mathcal{C} hold for T then they hold also for \mathbb{R}^2 . Moreover, since \mathbb{R}^2 is simply connected, its holonomy is trivial. We consider

now τ under the combinatorial point of view, i.e. we consider \mathbb{R}^2 as the quotient of a disjoint union of triangles $\bigsqcup \Delta_i$ modulo the relations given by the face-pairing maps of τ . First we construct a developing map on $\bigsqcup \Delta_i$ as follows: we take a triangle Δ_0 and we map it to \mathbb{C} via a similarity map φ_0 . We consider now a triangle Δ_1 adjacent to Δ_0 along an edge e and we map Δ_1 to \mathbb{C} via the only similarity map φ_1 such that $\varphi_0(\Delta_0)$ and $\varphi_1(\Delta_1)$ coherently glue along e . We go on in this way and we define a similarity map φ_i for each Δ_i .

The fact that the holonomy is trivial implies that the choice of the maps φ_i depends only on the choice of the first map φ_0 and not on the order in which we define the maps. So the map $\bigsqcup \varphi_i : \bigsqcup \Delta_i \rightarrow \mathbb{C}$ projects to a map $\varphi : \mathbb{R}^2 \rightarrow \mathbb{C}$ which is obviously a similarity map. □

Remark 3.2. It is easy to see that two developing maps differ by a similarity of \mathbb{C} . Moreover, if we fix one triangle Δ in T and one lift $\tilde{\Delta}$ of Δ , then a choice of a developing map determines a representative of the holonomy by using $\tilde{\Delta}$ as the base-point.

Once a representative of the holonomy has been fixed, its image consists of maps which commute with each other. So either they are translations or they have a common fixed point. The existence of such a point does not depend on the choice of the representative.

Definition 3.3. If a common fixed point of the holonomy exists, we call it an *axis* of the holonomy.

Remark 3.4. Since the holonomy is defined up to conjugation, the dilation component of the holonomy is a well-defined map $\bar{h} : H_1(T) \rightarrow \mathbb{C}^*$.

When equations \mathcal{C} hold (and so the holonomy is defined), we can write down the so-called *completeness* equations \mathcal{M} which say that $\bar{h} = 1$. This means that the image of the holonomy consists of translations. Equations \mathcal{M} are necessary in order to have a Euclidean structure on T .

Remark 3.5. Both \mathcal{C} and \mathcal{M} equations can be expressed in terms of products of moduli along loops in the dual graph Γ . For equations \mathcal{C} , consider all triangles around a vertex v in τ as a loop γ in Γ , then the product of moduli around v is the product of moduli of angles of γ .

Regarding equations \mathcal{M} , consider an element $[\gamma] \in H_1(T)$ represented by a loop $\gamma = (\Delta_0, \Delta_1, \dots, \Delta_0)$ in Γ . Let e_i be the edge between Δ_{i-1} and Δ_i in γ . Let us define coefficients n_i as follows: set $n_i = 1$ if e_i follows e_{i+1} (without passing through the third edge) in the cyclic ordering around Δ_i , and $n_i = -1$ otherwise. Let z_i be the modulus associated to the angle (e_i, e_{i+1}) ; set $z_i = -1$ if $e_{i+1} = -e_i$. Then $\bar{h}([\gamma]) = \prod_i z_i^{n_i}$ (see Figure 3). Since equations \mathcal{C} hold, \bar{h} does not depend on the representative γ .

Are the equation \mathcal{C} or \mathcal{C} and \mathcal{M} sufficient to have a structure on T ? In general the answer is no. For example, for a triangulation of T with only two triangles, equations




Figure 3. Moduli along paths in Γ

\mathcal{C} are always satisfied and if both moduli are flat then the identification space cannot be a torus, even if equations \mathcal{M} hold. For more interesting examples, see [9].

In the following we give some sufficient conditions in order for a solution of \mathcal{C} (or \mathcal{C} and \mathcal{M}) to define a similarity (or Euclidean) structure on T . We notice that these conditions are necessary only “near” the Euclidean case.

4 Equations \mathcal{C}^*

As announced above, we do not use conditions on the arguments of moduli, but we discuss them here.

First of all we note that if there exist only non-negative triangles, then we can define in a natural way the arguments of the moduli, allowing the flat moduli to have arguments 0 or π depending on whether they are positive or negative as real numbers. Whereas, if one has also negative triangles, then it is not clear how to define the arguments. One could simply define arguments to be negative, or require the sum of inner angles of each triangle to be π , and other ways are possible. Moreover, in general it is not possible to use only arguments in $[0, \pi]$ or $[0, 2\pi]$.

The following examples show that the conditions on the arguments are not necessary for \mathbf{z} to define a similarity structure on T .

Example 4.1. Suppose we have only positive triangles around a vertex v . Then a geometric picture near v looks like Figure 4a).




Figure 4. Geometric and topological situations




Figure 5. Equations \mathcal{C}^* are not necessary

We now add and remove a triangle Δ_1 , i.e. we add two copies of Δ_1 with inverse moduli. The topological picture is like in Figure 4f), but the geometric result is the same as before (Figure 4b)). We add and remove a second triangle Δ_2 (Figures 4c) and 4g)).

Now we observe that $(-\Delta_1) \cup (-\Delta_2)$ is geometrically equivalent to a big negative triangle $-\Delta$ plus a positive triangle Δ_3 (Figure 4d)–4e)). By replacing $(-\Delta_1) \cup (-\Delta_2)$ with $(-\Delta) \cup (+\Delta_3)$, the topological picture looks like in Figure 4h), the geometric situation is not changed, around v there are only positive triangles, but now the sum of arguments around v is 4π (if the convention about the arguments of positive moduli is the natural one).

Example 4.2. Suppose that \mathbf{z} defines a similarity structure on T . Suppose that there are no negative moduli. Then the arguments of moduli are defined. Consider two triangles glued together and change the triangulation by adding Δ_1 and Δ_2 as in Figure 5.

Assign now moduli r and r^{-1} to the \bullet -vertex of Δ_1 and Δ_2 , with $r \in \mathbb{R}^+$. It is easy to check that such a choice of moduli defines a similarity structure on T and that equations \mathcal{C}^* are not satisfied.

5 The moves

For Sections 5–9 we fix the hypothesis that $\mathbf{z} \in (\mathbb{C} \setminus \{0, 1\})^{|I|}$ is a solution of \mathcal{C} , that is we require the product of moduli around each vertex to be 1.

In the following we define a strategy which allows us to find a minimal triangulation of the torus equipped with a choice of moduli which will be “equivalent” to (τ, \mathbf{z}) in the sense that it defines a structure on T if and only if \mathbf{z} does, and if so the structures coincide.

Before proceeding we make two supplementary hypotheses which will be crucial in the sequel:

- H5.1 We suppose that the image of the holonomy has rank 2, i.e. it is not cyclic.
- H5.2 We suppose that if the holonomy has an axis, then it lies outside the image of a developing map.

Lemma 5.3. *Conditions H5.1 and H5.2 do not depend on the choice of the representatives of the developing and holonomy maps.*



Figure 6. Topological and geometric diagonal swap




Figure 7. The moduli in the GDS

Proof. Condition H5.1 does not depend on the choice of the representative of the holonomy because conjugation does not change the cyclicity of a subgroup.

Let D be a developing map. By Remark 3.2, once one triangle Δ and one of its lifts $\tilde{\Delta}$ has been fixed, then the map D determines a representative of the holonomy. Call x the axis of the holonomy. If we change D to $\Phi \circ D$, where $\Phi \in \text{Aff}(\mathbb{C})$, we have that the holonomy changes by the conjugation by Φ . Then the axis for the new holonomy is $\Phi(x)$ and $x \notin \text{Im}(D)$ if and only if $\Phi(x) \notin \text{Im}(\Phi \circ D)$. If we change the choices of Δ and $\tilde{\Delta}$, then the holonomy changes by the conjugation with an element of its image, so the axis does not change and this completes the proof of the second assertion. \square

We use the classical move of *topological diagonal swap (TDS)* to manipulate triangulations (Figure 6). This move extends in an obvious way in a geometric setting to give a *geometric diagonal swap (GDS)* like in Figures 6 and 7. Any GDS can be viewed as a function from the space of triangulations with moduli on the torus to itself.

Remark 5.4. With notation as in Figure 6, since we are considering only moduli in $\mathbb{C} \setminus \{0, 1\}$, we can apply the GDS only if $z' \neq z^{-1}$. And in this case the GDS is continuous as a function from the space of moduli on τ to the space of moduli on the resulting triangulation.

Let γ be a path (loop) in the dual graph of τ . Then by applying a GDS to two consecutive vertices of γ we obtain a path (loop) γ' as Figure 8 shows.

Simple calculations lead to the following facts:




Figure 8. Effects of a GDS on a path in Γ




Figure 9. The developed image of a GDS

Lemma 5.5. *The GDS does not change the products of moduli along paths in the dual graph from which we obtain equations \mathcal{C} and \mathcal{M} . So \mathbf{z} satisfies equations \mathcal{C} (or \mathcal{C} and \mathcal{M}) if and only if the resulting moduli after the move do.*

Remark 5.6. Any GDS lifts to a set of moves on the universal covering $\mathbb{R}^2 \rightarrow T$.

Lemma 5.7. *Let D be a developing map. The application of a GDS does not change the restriction of D to the 0-skeleton. It follows that the holonomy does not change under the GDS. Then a GDS preserves hypothesis H5.1.*

Remark 5.8. A GDS-move in general does not preserve hypothesis H5.2. Figure 9 shows a particular case of GDS, in which the axis is being incorporated into the image of the developing map.

Remark 5.9. The moves do not change the number of simplices of triangulations.

6 Loops and e-loops

As mentioned in Section 2, in our setting, a simplex of a triangulation might not be embedded in T . For a 1-simplex, this implies that it is a loop, that is, it starts and ends at the same point. This section is devoted to study the edges of τ that are loops. In this section we enforce the hypothesis H5.1.

Definition 6.1. We call *e-loop* an edge of the triangulation which starts and ends at the same vertex. We say that a vertex v has an *e-loop* if there exists an e-loop starting at v .

Remark 6.2. Since an e-loop is an edge of the triangulation, it is an embedded loop.

Definition 6.3. For any vertex v , we set $S(v)$ to be the number of triangles (with multiplicity) having v as a vertex.

Lemma 6.4. For each vertex v we have $S(v) > 1$.

Proof. Since equations \mathcal{C} hold, if $S(v) = 1$ the modulus of the only triangle around v must be 1. But this contradicts the fact that we are considering only moduli in $\mathbb{C} \setminus \{0, 1\}$. \square

Proposition 6.5. Let l be an e-loop. Then $[l] \neq 0$ as elements of $H_1(T)$.

Proof. Suppose the contrary. Then l bounds a sub-complex B homeomorphic to a disc. Thus l lifts to a loop in \mathbb{R}^2 . But the developed image of l is a straight segment which cannot be a loop. \square

Remark 6.6. The notion of parallelism between loops is well-defined for disjoint loops. Since we deal with loops that can share one point, we say that two loops on a surface are *parallel* if they jointly bound either an embedded annulus or an embedded pinched annulus.

Lemma 6.7. In a torus, the relation of parallelism between non-contractible loops sharing at most one point is transitive.

Proof. Let α, β, γ be non-contractible loops so that α is parallel to β which is parallel to γ . Cutting the torus along β we obtain a cylinder in which α and γ are parallel to the boundary. Since α and γ share at most one point, it follows that $\alpha \cup \gamma$ bounds either an annulus or a pinched annulus, and so they are parallel. \square

Remark 6.8. We claim that e-loops at different vertices are disjoint. By Proposition 6.5 and an argument as in Lemma 6.7, two disjoint e-loops are topologically parallel. So e-loops at different vertices are parallel.

Proposition 6.9. If there exist more than one vertex with e-loops, then e-loops at the same vertex are parallel.

Proof. Let l_1 and l_2 be two different e-loops at a vertex v . Let $v' \neq v$ be a vertex which has an e-loop l . By Remark 6.8, the e-loop l is parallel to both l_1 and l_2 . By Lemma 6.7 it follows that l_1 is parallel to l_2 . \square

Proposition 6.10. Suppose that each vertex has an e-loop. Then two different e-loops l_1 and l_2 at the same vertex v are not parallel.

Proof. Suppose the contrary. Then l_1 and l_2 jointly bound a region R whose fundamental group is isomorphic to \mathbb{Z} and generated by $[l_1] = [l_2]$. So R cannot contain




Figure 10. The triangles Δ_1 and Δ_2

any vertex in its interior because each vertex has an e-loop and e-loops are not contractible. It follows that R is a bigon, but this cannot happen in a triangulation. \square

Corollary 6.11. *Suppose that each vertex has an e-loop. If a vertex v has two different e-loops then there exists only one vertex. Equivalently, if there exists more than one vertex, then each vertex has exactly one e-loop.*

Lemma 6.12. *In the current hypotheses, suppose moreover that H5.2 holds. Let Δ_1 and Δ_2 be two triangles, glued along an edge and let v be a vertex of such edge. Suppose that Δ_1 and Δ_2 have inverse moduli at v , and that v has no e-loop. Suppose moreover, with notation as in Figure 10, that $v_1 = v_2$. Then $[\gamma_1\gamma_2^{-1}] = 0$ as an element of $H_1(T)$.*

Proof. If $\gamma_1 = \gamma_2$ the assertion is obvious. Otherwise, suppose $[\gamma_1\gamma_2^{-1}] \neq 0$. Since v has no e-loops, $\gamma_1\gamma_2^{-1}$ is a loop embedded in T . Then we choose $\gamma_1\gamma_2^{-1}$ as an element of a \mathbb{Z} -basis of $H_1(T)$. Now the fact that $v_1 = v_2$ implies that the developed image of $v_1 = v_2$ is a fixed point of $h(\gamma_1\gamma_2^{-1})$. So either the image of the holonomy is cyclic or the axis lies in the image of a developing map, but both cases are impossible because of hypotheses H5.1 and H5.2. \square

Remark 6.13. In the proof of Lemma 6.12, we used the hypotheses H5.1 and H5.2. This is not only a technical trick but it has relevant topological aspects, see Section 8 for more details.

7 The strategy

We say that a triangulation of T is *minimal* if it has only one vertex (and then exactly two triangles).

In this section we define a recursive algorithm based on six steps in order to find a minimal triangulation equivalent to the initial one. The algorithm will either get the requested triangulation or stop saying that hypothesis H5.2 has been violated. We will see that the algorithm must stop in a finite time. We mean by a *strategy* a complete application of the algorithm. In this section we enforce the hypothesis H5.1.

Remark 7.1. Suppose that a vertex v has no e-loops. If we can apply the GDS to two triangles having v as a vertex and that are consecutive around v , then after the move we have that:

- (1) $S(v)$ is decreased by 1.
- (2) v remains without e-loops.

Remark 7.2. If $S(v) = 2$ then equations \mathcal{C} imply that the two triangles have inverse moduli at v .

Step 1. If hypothesis H5.2 does not hold, then we stop here. If H5.2 is satisfied, and each vertex has an e-loop, then we go to Step 6. Otherwise we go to Step 2.

Step 2. Let v be a vertex without e-loops, and let Δ_1 and Δ_2 be two consecutive triangles around v . If Δ_1 and Δ_2 have inverse moduli, then we go to Step 3. Otherwise we apply a GDS to Δ_1 and Δ_2 . Then, if hypothesis H5.2 does not hold we stop here, otherwise we repeat this step keeping the vertex v fixed.

Remark 7.3. Note that by Remarks 7.1 and 7.2 it is sufficient to repeat Step 2 only a finite number of times.

Remark 7.4. When we come to Step 3, we have that the hypothesis H5.2 holds and that the two adjacent triangles Δ_1 and Δ_2 around the vertex v have inverse moduli.

Remark 7.5. With notation as in Figure 10, if $v_1 = v_2$ then the hypotheses of the Lemma 6.12 are satisfied. Then $\gamma_1\gamma_2^{-1}$ bounds an embedded disc B inside which there are no e-loops because e-loops are not contractible. Moreover if Δ_1 and Δ_2 are glued along only one edge, then B contains a vertex $w \neq v, v_1, v_3$.

Step 3. If either $v_1 \neq v_2$ or Δ_1 and Δ_2 are glued along two edges, then we go to Step 4. Otherwise we look only at the vertices $w \in B$ with $w \neq v, v_1, v_3$ as described in Remark 7.5 and we return to Step 1.

Remark 7.6. Note that applying a GDS around vertices different from v, v_1, v_3 does not change Δ_1 and Δ_2 . Then, since τ is finite, it is sufficient to repeat Steps 1–3 only a finite number of times.

Recall that to have inverse moduli, geometrically it means that Δ_1 and Δ_2 completely overlap with inverse orientations.

Step 4. We delete Δ_1 and Δ_2 from τ and we change the gluing rules as described in Figure 11. Then we go to Step 5.

To be more precise, the cancellation of Step 4 corresponds to Figure 11a) if Δ_1 and Δ_2 are glued along one edge, and to Figure 11b) if they are glued along two edges.

Remark 7.7. The choice made in Step 3 is necessary in order to avoid changes of topology of the torus. Namely, suppose that we go directly to Step 4 without going




Figure 11. Cancellations in the triangulation and in the dual graph

through Steps 1–3. Then the cancellation might disconnect the dual graph (see again Figure 11a)). We will see in the next section that all the cancellations we do in the strategy actually do not change the topology of the torus.

Step 5. We return to Step 1 and apply this step to the new triangulation obtained in Step 4. We repeat these steps until each vertex has an e-loop or the hypothesis H5.2 breaks. If H5.2 holds and each vertex has an e-loop, then we go to Step 6.


Remark 7.8. Note that applying a GDS around vertices without e-loops, the number of e-loops of τ does not decrease. Then, by induction on the number of triangles and on the number of vertices without e-loops, it is sufficient to repeat Steps 1–5 only a finite number of times.

Remark 7.9. In Step 4, it is easy to see that a choice of moduli satisfies equations \mathcal{C} or \mathcal{C} and \mathcal{M} if and only if the resulting moduli after the cancellation do. Moreover, the cancellations of Step 4 preserve the conjugacy class of the holonomy map and the hypothesis H5.2.

The last part of the strategy consists in reducing the triangulation to a minimal one.

Step 6. We recall that when we come to this step, hypothesis H5.2 holds and each vertex has an e-loop.

If the triangulation is minimal, we stop here. Otherwise we choose an e-loop l and an embedded closed simplicial path α so that α is a generator of $H_1(T)$ that meets l once. It follows that α meets once any e-loop parallel to l and so, by Remark 6.8 and Corollary 6.11, it meets any e-loop. Since each e-loop contains only one vertex, it follows that α contains all vertices. Cutting the torus along l and α we get a disc, whose triangulation looks like the one given in Figure 12.

Figure 12. Cutting along α and l

Since the triangulation is not minimal, we can choose an e-loop $l' \neq l$ at a vertex v . Now we try to apply a GDS to the two triangles having l' as an edge. If the move is not possible, i.e. if the triangles have inverse moduli at v , then we cancel the triangles as in Step 4 and we restart from Step 1. If the GDS is possible then we apply it. Note that after the move the vertex v has no e-loops. Then we restart from Step 1.

Remark 7.10. Note that Steps 4 and 6 involve a cancellation of pairs of two triangles. Then by induction on the number of triangles it follows that it is sufficient to apply only a finite number of steps of our strategy.

Remark 7.11. A strategy following the rules of Steps 1–6 as described stops only when either we lose H5.2 or the triangulation is reduced to a minimal one. Moreover if we consider a strategy \mathfrak{S} as a map between triangulations with moduli, then we have that (τ, \mathbf{z}) satisfies \mathcal{C} or \mathcal{C} and \mathcal{M} if and only if $\mathfrak{S}((\tau, \mathbf{z}))$ does.

Definition 7.12. We say that a strategy works if it leads to a minimal triangulation.

8 The effects of cancellations

In general, a cancellation might produce a degeneration of the topology of the torus (see Proposition 8.2). In this section we see that the cancellations that occur in our strategy do not change the topology of the torus. To prove this we simply check all possible cases. We recall that the hypothesis H5.1 is enforced.

Let Δ_1 and Δ_2 be the triangles we are going to cancel. For this section we fix the notation of Figure 13. If the triangles have two common edges, then they are either both embedded or both not embedded in the torus.

If they are embedded in the torus, then we are collapsing an embedded disc to its diameter and this does not change the topology. See Figure 13a).

If Δ_1 and Δ_2 have two common edges and are not embedded in the torus, then the unique possibility is that one vertex is in the interior of $\Delta_1 \cup \Delta_2$ and the other two coincide. Since the e-loops are not contractible, in this case we are collapsing a pinched annulus to a loop and this does not change the topology of the torus.

Now suppose that Δ_1 and Δ_2 have only one common edge. If $\Delta_1 \cup \Delta_2$ is embedded in the torus, then we are collapsing an embedded disc and there are no problems. If




Figure 13. The triangles Δ_1 and Δ_2




Figure 14. The cancellation in the case $v_1 = v_3$

$\Delta_1 \cup \Delta_2$ is not embedded, then we have two cases. Either the cancellation occurs in a Step 4 or in a Step 6.

If the cancellation is performed in a Step 4, then v has no e -loops and one can easily see that the only possible case is that $v_1 = v_3 \neq v_2$ (or $v_2 = v_3 \neq v_1$). Then γ_1 is an e -loop and so $[\gamma_1] \neq 0 \in H_1(T)$. To see that this cancellation does not change the topology of the torus see Figure 14.

If the cancellation is performed in a Step 6 then, using hypotheses H5.1 and H5.2 as in Lemma 6.12, one can see that $v_1 \neq v_2$. In this case we are collapsing a pinched annulus to a loop and this does not change the topology of the torus.

Remark 8.1. We used hypotheses H5.1 and H5.2 to say that we have not a cancellation with $v_1 = v_2$ and $[\gamma_1\gamma_2^{-1}] \neq 0 \in H_1(T)$. Actually this is a very bad case.

Proposition 8.2. *Suppose that two triangles Δ_1 and Δ_2 have inverse moduli at v and are glued along only one edge. Suppose moreover that $v_1 = v_2$ and $[\gamma_1\gamma_2^{-1}] \neq 0 \in H_1(T)$. Then a cancellation of Δ_1 and Δ_2 as in Step 4 produces a degeneration of the topology of the torus.*

Proof. Let P be the space obtained from T by removing $\Delta_1 \cup \Delta_2$ and changing the gluing rules as in Figure 11a) (i.e. the resulting space after the cancellation). Since $[\gamma_1\gamma_2^{-1}] \neq 0$, by cutting T along $\gamma_1\gamma_2^{-1}$ we obtain a cylinder. By removing $\Delta_1 \cup \Delta_2$ and changing the gluing rules, we obtain a sphere. Now in order to reconstruct P we have only to glue v_1 to v_2 , and then P is not a torus. \square

9 The existence of similarity maps

In the preceding section we have seen that any strategy leading to a minimal triangulation preserves the topology of the torus; we see now that it preserves also similarity structures. In this section we enforce the hypothesis H5.1.

Theorem 9.1. *Let (T, τ, \mathbf{z}) be a solution of \mathcal{C} satisfying the two hypotheses H5.1 and H5.2. Suppose that we can follow a strategy successfully to obtain a minimal triangulation. Also assume that this minimal triangulation defines a similarity structure on T . Then the original parameters \mathbf{z} also defines a similarity structure on T . Moreover, such a structure does not depend on the strategy used.*

Proof. Let us fix a strategy that works. By hypothesis, for the minimal triangulation there exists a torus T' endowed with a similarity structure and a similarity map $\varphi : T \rightarrow T'$ of degree ± 1 .

We will follow backward the steps of the strategy, showing that at each step there exists a similarity map from T to T' , having the same degree of φ . More precisely, we show that if there exists a similarity map φ after a GDS or a cancellation, then before the GDS or the cancellation there exists a similarity map, which we still call φ , which agrees with the previous one where the triangulation is not changed.

Let \tilde{T}' be the universal covering of T' and let $D' : \tilde{T}' \rightarrow \mathbb{C}$ be a developing map for T' .

Suppose that we have done a cancellation. Let Δ_1 and Δ_2 be the canceled triangles. Then we have to define φ on $\Delta_1 \cup \Delta_2$.

The map φ lifts to a map $\tilde{\varphi} : \tilde{T} \rightarrow \tilde{T}'$ such that $D' \circ \tilde{\varphi}$ is a developing map D . Let D_1 be a developing map which coincides with D outside the lifts of Δ_1 and Δ_2 before the cancellation. By checking all the possible cancellations, one can easily see that D_1

always exists. Moreover, since H5.2 holds, it splits along a map $\tilde{\varphi}_1 : \tilde{T} \rightarrow \tilde{T}'$ such that $D' \circ \tilde{\varphi}_1 = D_1$. Such a $\tilde{\varphi}_1$ projects to the requested similarity map φ . We notice that if H5.2 does not hold, then the image of D_1 contains the axis of the holonomy, which lies outside of the image of D' . If this is the case, the map D_1 cannot split along $\tilde{\varphi}_1$.

Now suppose that we have done a GDS which changes two triangles, say Δ_1 and Δ_2 , into two new triangles Δ'_1 and Δ'_2 . Then we have to change φ on $\Delta_1 \cup \Delta_2 = \Delta'_1 \cup \Delta'_2$.

We proceed exactly as above considering a lift of the map $\varphi|_{T \setminus (\Delta_1 \cup \Delta_2)}$. As above, since H5.2 holds, such a lift extends to the lifts of $\Delta_1 \cup \Delta_2 = \Delta'_1 \cup \Delta'_2$ and projects to the requested similarity map φ .

Regarding the degree of such maps, if we have done a cancellation, then passing from φ to φ_1 we are adding two triangles, having inverse moduli, whose developed images completely overlap. So $\deg(\varphi) = \deg(\varphi_1)$. Similarly, a GDS does not change the contribution of $\Delta_1 \cup \Delta_2$ to the degree of φ .

We now prove the second assertion. Let \mathbb{X} be the universal covering of the image of D' , equipped with the pull-back similarity structure. The map D' lifts to a map $\tilde{D}' : \tilde{T}' \rightarrow \mathbb{X}$. It turns out that \tilde{D}' is a global homeomorphism. Moreover, since D' is a developing map for T' , the holonomy of T' lifts to a representation $h' : H_1(T) \rightarrow \text{Aff}(\mathbb{X})$ such that \tilde{D}' is h' -equivariant. As a $(\mathbb{C}, \text{Aff}(\mathbb{C}))$ -space, the torus T' is isomorphic to the quotient of \mathbb{X} under the action of the image of h' .

Similarly, the holonomy of T lifts to a map $\tilde{h} : H_1(T) \rightarrow \text{Aff}(\mathbb{X})$ such that $\tilde{h} = h' \circ \varphi_*$. Since $\deg(\varphi) = \pm 1$, the map φ_* is an isomorphism. It follows that the image of h' is the same of the one of \tilde{h} .

Finally observe that a cancellation does not change the image of the holonomy. By Lemma 5.7 also a GDS does not change the image of the holonomy. So the structure of T' is independent of the strategy used. □

Remark 9.2. One can easily obtain a partial converse of this theorem. That is, if there exists a similarity map φ from (T, τ, \mathbf{z}) to a torus T' with a similarity structure, then \mathcal{C} and H5.2 hold, but in general the hypothesis H5.1 may be not satisfied (see Example 9.6).

Remark 9.3. If (T, τ, \mathbf{z}) is a solution of the equations \mathcal{C} and \mathcal{M} , then we can coherently choose the sizes of the triangles. It follows that the algebraic sum A of the areas of the triangles of τ with moduli \mathbf{z} is well-defined up to multiplication by a positive factor.

Corollary 9.4. *Suppose (T, τ, \mathbf{z}) is a solution of the equations \mathcal{C} and \mathcal{M} and let A be as in Remark 9.3. If $A \neq 0$, then \mathbf{z} defines a Euclidean structure on the torus.*

Proof. If \mathbf{z} is a solution of \mathcal{M} then the holonomy has no axis, so H5.2 is always satisfied. We have to check H5.1.

Suppose the contrary. Since $\text{Im}(h)$ consists of translations, it is isomorphic either to \mathbb{Z} or to the trivial group. It follows that $\text{Im}(h)$ acts freely and properly discontinuously on \mathbb{C} .




Figure 15. In a minimal triangulation, equations \mathcal{M} imply $z_1 = z_2$




Figure 16. A similarity structure without hypothesis H5.1

Let C be the quotient of \mathbb{C} under the action of $\text{Im}(h)$. The map D descends to a map $\mathcal{D} : T \rightarrow C$ because D is h -equivariant. Let ω be the area-form on C induced by \mathbb{C} . Then we have

$$A = \int_T \mathcal{D}^* \omega = \int_{\mathcal{D}(T)} \text{deg}(\mathcal{D}) \omega.$$

As T is compact and C is not compact, $\text{deg}(\mathcal{D}) = 0$. Thus $A = 0$, that is a contradiction.

Then H5.1 is satisfied and, since H5.2 is satisfied, it follows that any strategy works. Moreover in the minimal triangulation we obtain, equations \mathcal{M} imply that the two triangles have the same modulus (see Figure 15). Thus, the two triangles are either both positive or both negative, and it is easily checked that gluing together such triangles, one obtains a Euclidean torus. \square

Remark 9.5. To prove these results we used H5.1 and H5.2 in a crucial way. As we have seen in the proof of Theorem 9.1, the hypothesis H5.2 is necessary in order to have similarity maps while the hypothesis H5.1 is necessary only in order to apply a strategy. Namely, if H5.1 is not satisfied, then it may be that \mathbf{z} defines a similarity structure on T , but we cannot use a strategy to define similarity maps.

Example 9.6. Take $\mathbb{C} \setminus \{0\}$ and make the quotient by the multiplication by 2. Then $\mathbb{C} \setminus \{0\} / \sim$ is a torus equipped with a similarity structure, and we can triangulate a fundamental domain as in Figure 16. We can then assign moduli to the triangles in the obvious way and it is clear that they are a solution of \mathcal{C} but do not satisfy H5.1. If we try to apply a strategy, we lose H5.2 at the initial steps. Moreover in this case it is easy to check that no strategy works.

10 Algebraic conditions on the moduli

In this section we give a complete characterization of the choices of moduli defining similarity structures on the torus. This section is almost independent from Sections 5–9 (we refer only to Lemma 5.3, Remark 9.3 and Corollary 9.4). We refer to Sections 2–4 for the general definitions.

Let τ be a triangulation of the torus T and let \mathbf{z} be a choice of moduli which is a solution of equations \mathcal{C} . Consider the universal covering \mathbb{R}^2 of T . Then τ lifts to a triangulation with moduli on \mathbb{R}^2 . Let $D : \mathbb{R}^2 \rightarrow \mathbb{C}$ be a developing map which maps each triangle Δ_i in a straight one with modulus z_i . Let $h : H_1(T) \rightarrow \text{Aff}(\mathbb{C})$ be the representative of the holonomy map compatible with D .

Let $\langle \gamma_1, \gamma_2 \rangle$ be a fixed generator of $H_1(T)$. Since $h(\gamma_1)$ and $h(\gamma_2)$ commute, we have two cases:

- 1) Both $h(\gamma_1)$ and $h(\gamma_2)$ are translations.
- 2) $h(\gamma_1)$ and $h(\gamma_2)$ have a common fixed point, which we call axis. In this case we can and will suppose that $0 \in \mathbb{C}$ is the axis of the holonomy. Thus h acts on \mathbb{C} by multiplications, that is, for each $\gamma \in H_1(T)$ and $x \in \mathbb{C}$, we have $h(\gamma)(x) = x \cdot h(\gamma)(1)$, where the dot denotes the multiplication of \mathbb{C} .

Definition 10.1. Let \mathbb{X} be defined as follows. In Case 1) let $\mathbb{X} = \mathbb{C}$, considered as the universal covering of itself. In Case 2) let \mathbb{X} be the universal covering of $\mathbb{C}^* = \mathbb{C} \setminus \{0\}$. We put on \mathbb{X} the pull-back similarity structure.

Remark 10.2. In both Cases 1) and 2), we have $\mathbb{X} = \mathbb{C}$. In Case 1) the covering map is the identity and the similarity structure on \mathbb{X} is the usual one. In Case 2), the covering map $\mathbb{X} \rightarrow \mathbb{C}^*$ is the usual exponential map $\exp : \mathbb{C} \rightarrow \mathbb{C}^*$, the similarity structure on \mathbb{C}^* is the usual one, while the similarity structure we consider on \mathbb{X} is not the usual one.

Proposition 10.3. *In both Cases 1) and 2) the translations of \mathbb{X} are similarity maps.*

Proof. There is nothing to say in Case 1). In Case 2) the thesis follows because a translation $\tau(x) = x + \xi$ descends to the map $\theta(y) = ye^\xi$ which is a similarity map on \mathbb{C}^* . □

Proposition 10.4. *Suppose that in Case 2) the image of D does not contain the element 0. Then the map D lifts to a map $\tilde{D} : \mathbb{R}^2 \rightarrow \mathbb{X}$ and h to a representation $\tilde{h} : H_1(T) \rightarrow \text{Aff}(\mathbb{X})$ such that \tilde{D} is \tilde{h} -equivariant. Moreover, the image of \tilde{h} consists of translations.*

Proof. This is tautological in Case 1). Let us deal with Case 2). The map \tilde{D} exists because $0 \notin \text{Im}(D)$. We define \tilde{h} as follows. For each $\gamma \in H_1(T)$ and $x \in \mathbb{X}$ we have

$$\exp(\tilde{D}(\gamma x)) = D(\gamma x) = h(\gamma)(D(x)) = D(x) \cdot h(\gamma)(1) = \exp(\tilde{D}(x)) \cdot h(\gamma)(1).$$

Thus $\tilde{D}(\gamma x) = \tilde{D}(x) + \xi(\gamma, x)$ with $\exp(\xi(\gamma, x)) = h(\gamma)(1)$. The function $\xi(\gamma, x)$, as a function of x , is a continuous map from a connected space to a discrete space, and then it is constant. It follows that $\xi(\gamma, x) = \xi(\gamma)$. The function $\gamma \rightarrow \tilde{h}(\gamma)$ defined by $\tilde{h}(\gamma)(x) = x + \xi(\gamma)$ is the requested representation. The second claim immediately follows. \square

Definition 10.5. We set $\tilde{h}(\gamma_i)(x) = x + \xi_i$, $i = 1, 2$. We say that \tilde{h} has rank 2 over \mathbb{R} if ξ_1 and ξ_2 are linearly independent over \mathbb{R} .

Remark 10.6. The previous definition is equivalent to saying that $\xi_1/\xi_2 \notin \mathbb{R} \cup \{\infty\}$ or that $x_1 y_2 \neq x_2 y_1$, if we set $\xi_k = x_k + iy_k$, for $k = 1, 2$.

Remark 10.7. As in Lemma 5.3, one can see that the conditions that $0 \notin \text{Im}(D)$ and that \tilde{h} has rank 2 over \mathbb{R} do not depend on the choice of D .

Theorem 10.8. Let \mathbf{z} be a choice of moduli for a triangulation of T such that \mathcal{C} is satisfied and, in Case 2), suppose that $0 \notin \text{Im}(D)$. Suppose that \tilde{h} has rank 2 over \mathbb{R} . Then \mathbf{z} defines a similarity structure on T .

Proof. Since ξ_1 and ξ_2 are linearly independent, the action of $H_1(T)$ on \mathbb{X} via \tilde{h} is free and properly discontinuous. So $X = \mathbb{X}/\tilde{h}$ is well-defined and is a torus with a similarity structure. Moreover a map $f : T \rightarrow X$ obtained by pushing-down \tilde{D} is well-defined. Obviously f induces on each Δ_i the same structure defined by z_i .

We prove now that $\deg f = \pm 1$. The homotopy class of f is completely determined by f_* . Since f_* is an isomorphism between $\pi_1(T)$ and $\pi_1(X)$, it is easy to construct a map $g : T \rightarrow X$ of degree ± 1 such that $g_* = f_*$. Then f is homotopic to g and thus has degree ± 1 . \square

The converse of Theorem 10.8 is also true, so its hypotheses are necessary and sufficient conditions in order for \mathbf{z} to define a similarity structure on T .

Theorem 10.9. Suppose that a choice of moduli defines a similarity structure on T , then the hypotheses of Theorem 10.8 are satisfied.

Proof. Consider a similarity map $f : T \rightarrow T'$ of degree ± 1 , where T' is a torus equipped with a similarity structure. The existence of f implies that equations \mathcal{C} hold. Moreover f lifts to a map $\tilde{\varphi} : \mathbb{R}^2 \rightarrow \tilde{T}'$ (\tilde{T}' and \mathbb{R}^2 are the universal covering of T' and T) such that the obvious diagram is commutative.

Since $\deg(f) = \pm 1$, it is easy to check that f_* is an isomorphism. Moreover, since the diagram is commutative, the holonomy of T is the composition of f_* with the holonomy of T' .

If T' is a Euclidean torus, then \tilde{T}' is exactly \mathbb{C} with its similarity structure, so $\tilde{\varphi}$ is a developing map, the holonomy consists of translations and so it has no axis.

Otherwise consider a developing map $D' : \tilde{T}' \rightarrow \mathbb{C}$ for T' . It is clear that $\text{Im}(D')$ does not contain the axis of the holonomy. The map $D' \circ \tilde{\varphi}$ is a developing map for T and then its image does not contain the axis of the holonomy.

We have now to check that the holonomy has rank 2 over \mathbb{R} . It is easy to check that T' is the quotient of \mathbb{X} under the image of \tilde{h} . Suppose that ξ_1 and ξ_2 are linearly dependent on \mathbb{R} . Then the image of \tilde{h} is isomorphic to a subgroup of \mathbb{R} and then it is either cyclic or dense in \mathbb{R} . In both cases T' cannot be a torus. \square

Theorems 10.8 and 10.9 completely characterize the choices of moduli \mathbf{z} defining similarity structures on the torus. Moreover, as the following proposition shows, the set of moduli defining structures on T is open in the space of solutions of equations \mathcal{C} .

Proposition 10.10. *If \mathbf{z} defines a similarity structure on T , then all choices of moduli near \mathbf{z} which satisfy \mathcal{C} define similarity structures on T .*

Proof. Let \mathbf{z}' be a solution of \mathcal{C} sufficiently close to \mathbf{z} and let D' and h' be the corresponding developing and holonomy maps. First, suppose that the holonomy corresponding to \mathbf{z} has an axis. Since to have an axis that lies outside the image of a developing map is an open condition, it follows that also h' has an axis outside the image of a developing map. Moreover, also the map \tilde{D}' used to define \tilde{h}' is close to the map \tilde{D} , so the ξ'_i 's are close to the ξ_i 's. Since the condition on the ξ_i 's is an open one, also \tilde{h}' has rank 2 over \mathbb{R} , and the assertion follows.

Now suppose that h consists of translations. If also h' consists of translations, then as above \tilde{h}' has rank 2 over \mathbb{R} . Suppose that h' has an axis. We have to check that the axis lies outside the image of D' and that \tilde{h}' has rank 2 over \mathbb{R} . Since h consists of translations, if \mathbf{z}' is sufficiently close to \mathbf{z} , then

$$h'(\gamma_1)x = (1 + \beta_1)x + \eta_1, \quad h'(\gamma_2)x = (1 + \beta_2)x + \eta_2$$

with $\beta_i \sim 0$ and $\eta_i \sim \xi_i$. Therefore, the fixed point of $h'(\gamma_i)$ is $-\eta_i/\beta_i$. It follows that the axis of h' goes to ∞ as $\mathbf{z}' \rightarrow \mathbf{z}$. If $F \subset \mathbb{R}^2$ is a fundamental domain for T , then the developed image of F depends continuously on \mathbf{z} , that is, if \mathbf{z}' is close to \mathbf{z} , then $D'(F)$ is close to $D(F)$. Since the axis of h' goes to ∞ as $\mathbf{z}' \rightarrow \mathbf{z}$, for \mathbf{z}' sufficiently close to \mathbf{z} it lies outside $D'(F)$ and so, by h' -equivariance of D' , we have that the axis of h' lies outside the image of D' .

We check now that \tilde{h}' has rank 2 over \mathbb{R} . Since the image of h' is an abelian group, it follows that $\beta_1/\beta_2 = \eta_1/\eta_2$. With this notation, $e^{\xi'_i} = 1 + \beta_i$. Using the determination of the logarithm such that $\log(x)$ has imaginary part in $(-\pi, \pi)$ it turns out that $\xi_i = \log(1 + \beta_i)$. Thus for \mathbf{z}' sufficiently close to \mathbf{z} we have

$$\frac{\xi'_1}{\xi'_2} = \frac{\log(1 + \beta_1)}{\log(1 + \beta_2)} \sim \frac{\beta_1}{\beta_2} = \frac{\eta_1}{\eta_2} \sim \frac{\xi_1}{\xi_2} \notin \mathbb{R} \cup \{\infty\}.$$

So \tilde{h}' has rank 2 over \mathbb{R} . \square

In Case 1) the hypothesis that \tilde{h} has rank 2 over \mathbb{R} can be easily checked, as the following proposition shows. We recall that if equations \mathcal{M} hold, then we can define the algebraic area A as in Remark 9.3.

Proposition 10.11. *Suppose that \mathbf{z} is a solution of \mathcal{C} and \mathcal{M} . Then \tilde{h} has rank 2 over \mathbb{R} if and only if $A \neq 0$.*

Proof. The *only if* part follows from Theorem 10.8. The *if* part follows because by Corollary 9.4, \mathbf{z} defines a Euclidean structure on T , then Theorem 10.9 applies. \square

11 The Klein bottle

In this section we show how the problem of similarity structures on the Klein bottle can be reduced to a problem on the torus.

Let τ be a triangulation of the Klein bottle K and let us fix an orientation for each triangle of τ (clearly such orientations cannot be coherent because K is not orientable). If we consider the orienting double covering $\pi : T \rightarrow K$, then we can pull-back τ to a triangulation $\sigma = \pi^*(\tau)$ on T . Clearly σ can be coherently oriented (as we do). Each triangle Δ of τ is covered by exactly two triangles of σ , mapped to Δ one with the opposite orientation to the other.

Let \mathbf{z} be a choice of moduli on τ . Then \mathbf{z} induces a choice of moduli \mathbf{z}' on σ in a natural way. Namely, if $\pi^{-1}(\Delta) = \Delta_1 \cup \Delta_2$, if π preserves the orientation of Δ_1 and z is the modulus of Δ , then the modulus of Δ_1 is z and that of Δ_2 is \bar{z} .

Now suppose that there exists a similarity map φ from (K, τ, \mathbf{z}) to a Klein bottle J equipped with a similarity structure. Let $p : Y \rightarrow J$ be the orienting double covering. The structure of J lifts to a structure of the torus Y . Moreover, as Lemma 11.2 shows, we can lift the map φ to a similarity map between the orienting tori. If such a map has degree ± 1 , then we say that \mathbf{z} defines a similarity structure on K .

Clearly if \mathbf{z}' defines a similarity structure on the torus, then we can push-down the structure on Y and then \mathbf{z} defines a similarity structure on K . So we have the following theorem:

Theorem 11.1. *The choice of moduli \mathbf{z} defines a similarity structure on K if and only if \mathbf{z}' does so on the torus.*

Lemma 11.2. *Let $f : K \rightarrow J$ be a map between two Klein bottles and let $\pi : T \rightarrow K$ and $p : Y \rightarrow J$ be their double orienting coverings. Then f lifts to a map φ such that the following diagram is commutative.*

$$\begin{array}{ccc} T & \xrightarrow{\varphi} & Y \\ \pi \downarrow & & \downarrow p \\ K & \xrightarrow{f} & J \end{array}$$

Proof. If there are no ambiguities, let us denote both $\pi_1(K)$ and $\pi_1(J)$ by $\langle a, b; abab^{-1} \rangle$. So we have the commutation rule $ab = ba^{-1}$. Each element in $\pi_1(K)$ can be written in a unique way in the form $b^\beta a^\alpha$. Now $\pi_*(\pi_1(T)) \subset \pi_1(K)$ is the set of elements of the form $b^{2k} a^x$ and the same for $p_*(\pi_1(Y)) \subset \pi_1(J)$. The map φ exists if

and only if $f_*\pi_*(\pi_1(T)) \subset p_*(\pi_1(Y))$. Then we only have to check that $f_*\pi_*(\gamma)$ is of the form $b^{2k}a^x$ for all $\gamma \in \pi_1(T)$.

Let $f_*(a) = b^m a^n$ and $f_*(b) = b^s a^t$. Since f_* is a homomorphism, we have $1 = f_*(abab^{-1}) = b^m a^n b^s a^t b^m a^n b^{-s} a^{(-1)^{s+1}t}$ and it easily follows that $m = 0$. Finally, since $(b^p a^q)^r = b^{rp} a^{rq}$ with $y \in \mathbb{Z}$, then $f_*(b^{2k} a^x) = b^{2ks} a^z$ with $z \in \mathbb{Z}$ and so it is of the requested form for all $k, x \in \mathbb{Z}$. \square

References

- [1] R. Benedetti, C. Petronio, *Lectures on hyperbolic geometry*. Springer 1992.
MR1219310 (94e:57015) Zbl 0768.51018
- [2] D. B. A. Epstein, R. C. Penner, Euclidean decompositions of noncompact hyperbolic manifolds. *J. Differential Geom.* **27** (1988), 67–80. MR918457 (89a:57020)
Zbl 0611.53036
- [3] S. Francaviglia, Hyperbolicity equations for cusped 3-manifolds and volume-rigidity of representations. PhD Thesis, Scuola Normale Superiore, Pisa, 2003.
- [4] S. Francaviglia, Algebraic and geometric solutions of hyperbolicity equations. *Topology Appl.* **145** (2004), 91–118. MR2100867 (2005k:57033)
- [5] S. Francaviglia, Hyperbolic volume of representations of fundamental groups of cusped 3-manifolds. *Int. Math. Res. Not.* no. **9** (2004), 425–459. MR2040346 (2004m:57032)
Zbl 02207650
- [6] M. Kapovich, *Hyperbolic manifolds and discrete groups*. Birkhäuser 2001.
MR1792613 (2002m:57018) Zbl 0958.57001
- [7] W. D. Neumann, D. Zagier, Volumes of hyperbolic three-manifolds. *Topology* **24** (1985), 307–332. MR815482 (87j:57008) Zbl 0589.57015
- [8] C. Petronio, J. Porti, Negatively oriented ideal triangulations and a proof of Thurston’s hyperbolic Dehn filling theorem. *Expo. Math.* **18** (2000), 1–35.
MR1751141 (2001c:57017) Zbl 0977.57011
- [9] C. Petronio, J. R. Weeks, Partially flat ideal triangulations of cusped hyperbolic 3-manifolds. *Osaka J. Math.* **37** (2000), 453–466. MR1772844 (2001g:57032)
Zbl 0952.57003
- [10] W. P. Thurston, The geometry and topology of 3-manifolds. Mimeographed notes, Princeton University Mathematics Department, 1979.

Received 14 April, 2004

S. Francaviglia, Departament de Matemàtiques, Edifici C, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain
Email: s.francaviglia@sns.it