Winter Braids Lecture Notes

Juan González-Meneses

Geometric approaches to braid groups and mapping class groups

Vol. 2 (2015), Course nº III, p. 1-25.

<http://wbln.cedram.org/item?id=WBLN_2015__2__A3_0>

Winter Braids Lecture Notes
Vol. **2** (2015) 1-25
Course n° III

Geometric approaches to braid groups and mapping class groups

JUAN GONZÁLEZ-MENESES

Abstract

These are Lecture Notes of a course given by the author at the School Winter Braids, held at the Université de Pau et des Pays de L'Adour (France), on February 2015. It is explained how mapping class groups, and in particular braid groups, act on some interesting geometric spaces like the hyperbolic plane and the complex of curves, and how this allows to obtain some algebraic properties of the groups. A proof of the hyperbolicity of the graph of curves, following Hensel-Przytycki-Webb, is given.

1. Introduction

Given $n > 1$, the braid group on n strands B_n, introduced by Artin [2, 3], is given by the following group presentation:

$$B_n = \left\langle \sigma_1, \dots, \sigma_{n-1} \;\middle|\; \begin{array}{ll} \sigma_i\sigma_j = \sigma_j\sigma_i & |i-j| > 1 \\ \sigma_i\sigma_j\sigma_i = \sigma_j\sigma_i\sigma_j & |i-j| = 1 \end{array} \right\rangle$$

These groups have very natural generalizations, both from the algebraic and from the topological points of view.

On the algebraic side, Artin-Tits groups [5, 27] are groups which can be defined by a presentation that looks like the one above: A finite number of generators, and at most one relation for each pair of generators, in which they alternate forming words of equal length which are identified. Artin-Tits groups are closely related to Coxeter groups [5, 10]. Namely, imposing that the square of each generator is trival produces a Coxeter group from any given Artin-Tits group. When the associated Coxeter group is finite, the Artin-Tits group is said to be of spherical type, and in that case it satisfies very interesting divisibility properties, admitting what is called a Garside structure [12, 13]. Artin-Tits groups of spherical type are then the main examples of Garside groups and, among them, the braid group B_n is the most representative one, as the Coxeter group associated to the braid group B_n is precisely the symmetric group Σ_n.

On the topological side, braids can be seen as mapping classes on a punctured disc. More precisely, let $\mathbb{D}_n$ be a disc with n punctures, that is, with n interior points removed. Denote $\mathrm{Homeo}^+(\mathbb{D}_n, \partial\mathbb{D}_n)$ the set of orientation preserving homeomorphisms of $\mathbb{D}_n$, fixing the boundary pointwise. One can endow this set with a natural compact-open topology, so one has the notion of what means 'deforming' one homeomorphism into another, keeping the boundary

$\partial\mathbb{D}_n$ fixed through the deformation. Considering these homeomorphisms up to deformation just means considering the set $\pi_0\big(\mathrm{Homeo}^+(\mathbb{D}_n, \partial\mathbb{D}_n)\big)$, whose elements are called *mapping classes* (fixing the boundary). The deformation of homeomorphisms is compatible with their composition, and this turns the set of mapping classes into a group. It is well known [4] that the braid group on n strands is isomorphic to this group:

$$B_n \simeq \pi_0\big(\mathrm{Homeo}^+(\mathbb{D}_n, \partial\mathbb{D}_n)\big)$$

The goal of these notes is to give some examples on how Geometry can help to understand a group, focusing on the case of braid groups. In general, one can use geometric properties to study a group G by finding a suitable geometric space, so that G acts on it in a *nice* way. One important example is obtained when the space is the Cayley graph of the group G, with respect to a given set of generators, and the action is given by left multiplication. This is one of the main tools of Geometric Group Theory. But in this notes we will deal with other kinds of examples.

We will explain two main examples of such an action. In Section 2, we will see how the braid groups, and mapping class groups in general, act on the hyperbolic plane, and how this produces an action on the space of geodesic laminations from which we can extract many interesting properties. We will follow the exposition in [9] of the work by Thurston [29], and we will see how elements of the mapping class groups are classified into periodic, reducible and pseudo-Anosov. We will require from the reader just some basic notions of covering maps (see for instance [21]).

In Section 3 we will present a very interesting geometric space introduced by Harvey [19], called the *curve complex*, on which braid groups and mapping class groups act by isometries. We will explain a recent result by several authors, showing that the complexes of curves are uniformly hyperbolic. We will follow the arguments in [22], accompanied by a good amount of pictures.

Acknowledgements: I am very grateful to Paolo Bellingeri, Vincent Florens, JB Meilhan and Emmanuel Wagner for organizing Winter Braids, for inviting me to give a course, for their extraordinary patience while waiting for this manuscript, and for many valuable comments. Thanks also to Juan Souto for very interesting conversations at very interesting places. Finally, I thank the anonymous referee for his/her many wise observations, suggestions and remarks, which have improved considerably the quality of these notes. The author is partially supported by the Spanish research project MTM2013-44233-P and FEDER.

2. Mapping class groups acting on the hyperbolic plane

Throughout these notes, braids will be seen as mapping classes, as explained in the Introduction. That is, a braid in B_n will be an orientation preserving homeomorphism of the n-times punctured disc $\mathbb{D}_n$, fixing the boundary pointwise, up to deformations which also fix the boundary pointwise. This means that there is an action of the braid group on the topological space $\mathbb{D}_n$. But we want B_n to act on a geometric space, so we will put some geometry on $\mathbb{D}_n$.

Before that, we will start by studying a simpler kind of space. Let $S = S_g$ be a closed orientable surface of genus $g > 0$. We will study the mapping class group $MCG(S) = \pi_0(Homeo^+(S))$. This time, as there is no boundary, we do not need to fix any subset of S. It will be easier to explain how to put a geometry on S, as we can identify the universal cover of S with the hyperbolic plane $\mathbb{H}^2$.

The hyperbolic plane $\mathbb{H}^2$ is a space that can be identified with the interior of the unit disc in $\mathbb{R}^2$. The geometry of $\mathbb{H}^2$ can be described by stating that geodesics in $\mathbb{H}^2$ correspond either to arcs of circles which are orthogonal to the boundary of the unit disc or, in the extreme case, to diameters of the unit disc (see Figure 2.1).

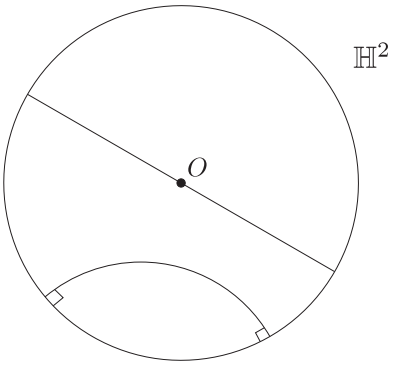

Figure 2.1: Geodesics in the hyperbolic plane

This assumption on geodesics allows to deduce the relation between the Euclidean metric in the unit disc, say ds, and the hyperbolic metric dh in $\mathbb{H}^2$. Up to a scaling factor, it is $dh = \frac{2\,ds}{1-r^2}$, where r is the Euclidean distance to the origin (that is, to the center of the unit disc). This means that, close to the origin, both metrics are quite similar, but as one approaches the border of the unit disc, points which are close with respect to the Euclidean metric become very distant with respect to the hyperbolic metric. This representation of $\mathbb{H}^2$ is called Poincaré's disc.

Although Euclidean and hyperbolic metrics in the unit disc are quite different, angles in $\mathbb{H}^2$ are measured exactly in the same way as in $\mathbb{R}^2$. One of the most particular properties of a hyperbolic space is that the sum of the angles of a (geodesic) triangle is *smaller* than π, while in the Euclidean plane this sum is always π (see Figure 2.2). Even more surprising is the fact that the area of a (geodesic) triangle in $\mathbb{H}^2$ depends only on its angles!

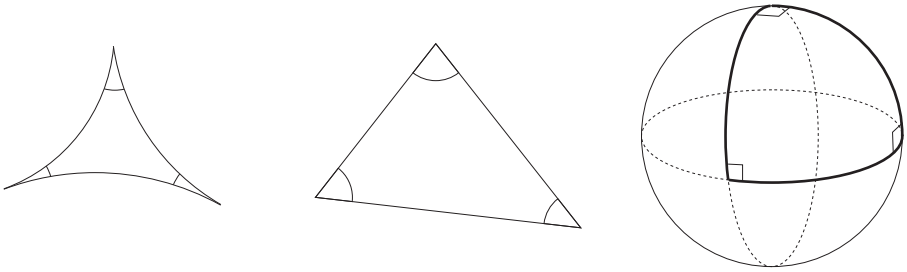

Figure 2.2: Triangles in a hyperbolic, Euclidean and elliptic space, respectively. The letters α, β and γ represent the angles of a geodesic triangle

Theorem 2.1 (Gauss-Bonnet)**.** *The area of a (geodesic) triangle in $\mathbb{H}^2$ whose angles are α, β and γ, is equal to $\pi-(\alpha+\beta+\gamma)$.*

There are other spaces in which the sum of the angles of a geodesic triangle is bigger than π. For instance, in a 2-dimensional sphere $\mathbb{S}^2$ with the metric inherited from $\mathbb{R}^3$, geodesics are long circles (i.e. circles belonging to planes which contain the center of the sphere). In Figure 2.2 one can see a geodesic triangle in $\mathbb{S}^2$ with three right angles. We are not going to be interested in these *elliptic* (or positively curved) spaces, but we will concentrate on hyperbolic (or negatively curved) spaces, namely on $\mathbb{H}^2$.

2.1. Hyperbolic structure on a closed surface

A closed orientable surface S of genus $g > 0$ can be obtained from a regular polygon of $4g$ sides, identifying the sides in a suitable way, like in Figure 2.3.

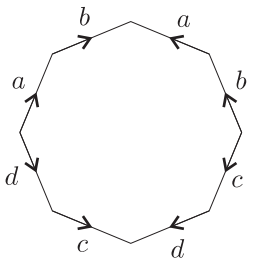

Figure 2.3: How to identify the sides of an octagon to get a closed surface of genus 2. In general, to get a surface of genus g, the sides of the polygon could be $a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1}$ reading clockwise.

Each interior angle in this polygon measures $\frac{2g-1}{4g} 2\pi$. If $g = 1$, S is the torus $\mathbb{T}^2$, the polygon is a square, and one can tile the Euclidean plane with these squares identifying the corresponding sides, as in Figure 2.4. This is how one can endow the torus $\mathbb{T}^2$ with an Euclidean geometry: Geodesics in the torus correspond to geodesics in the Euclidean plane, that is, to straight segments. This tiling of the plane determines a map π from $\mathbb{R}^2$ to $\mathbb{T}^2$ (each point in $\mathbb{R}^2$ corresponds to a point in $\mathbb{T}^2$, and this map is well defined at the squares' overlaps, as the squares' sides are identified in the right way). The map π is a covering map, and $\mathbb{R}^2$ is simply connected, hence $\mathbb{R}^2$ is the universal cover of the torus $\mathbb{T}^2$.

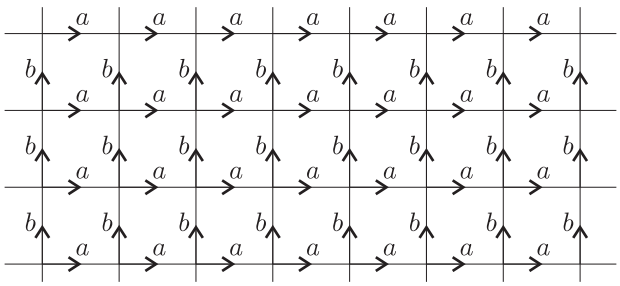

Figure 2.4: Tiling $\mathbb{R}^2$ with squares representing a torus.

If $g > 1$, then $2g-1$ does not divide $4g$, so the size of each angle of the polygon in Figure 2.3 (or the analogous polygon for bigger genus) is not an integral divisor of 2π. Hence, one

cannot tile the Euclidean plane using this kind of polygons. But one can tile the hyperbolic plane $\mathbb{H}^2$, as we are going to see.

In order to identify the sides of the *tiling pieces* in the right way, $4g$ copies of the tiling piece should be incident around each vertex. This means that we need to draw a regular $4g$-gon in $\mathbb{H}^2$, whose angles equal $\frac{2\pi}{4g}$.

In order to achieve this, consider a regular $4g$-gon in $\mathbb{H}^2$ centered at the origin. It can be obtained by considering $4g$ evenly distributed radii of the unit disc, and by taking a point on each radius, all $4g$ points being at the same distance from the origin. The geodesics joining consecutive points form the desired regular $4g$-gon (See Figure 2.5). Notice that if the chosen vertices ar very close to the origin, we have a tiny polygon whose sides are similar to straight segments, hence its interior angles are close to $\frac{2g-1}{4g}2\pi$. But if the chosen vertices are very far from the origin, the obtained polygon will be similar to the one on the right hand side of Figure 2.5, hence its interior angles would be close to 0. If we move the vertices continuously from the former situation to the latter, the size of its angles also varies continuously, so there will be a position in which the angles are exactly $\frac{2\pi}{4g}$, as desired. Let us denote P this regular polygon in $\mathbb{H}^2$, the sum of whose angles is 2π.

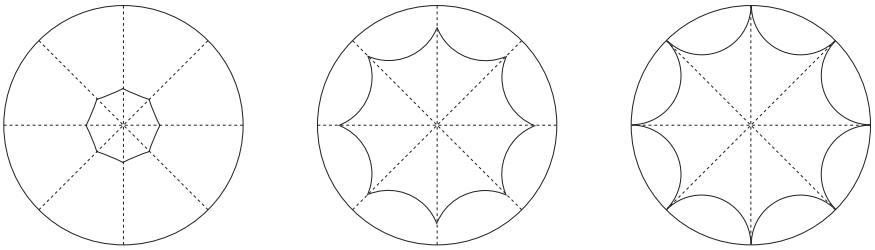

Figure 2.5: Regular octagons centered at the origin of $\mathbb{H}^2$. Small polygons have angles close to $3\pi/4$. Huge polygons have angles close to 0. In between some polygon has angles of size $\pi/4$.

Now we can tile the whole hyperbolic plane $\mathbb{H}^2$ with copies of P just by translating P (using translations with respect to the hyperbolic metric). The angles of P have the exact size so that this tiling can be done. See Figure 2.6 to get an idea of the tiling.

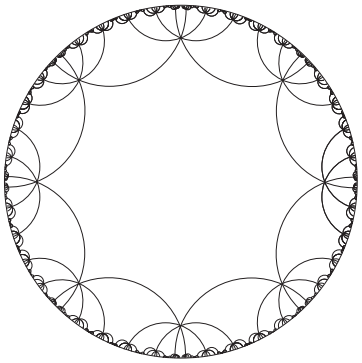

Figure 2.6: Tiling the hyperbolic plane $\mathbb{H}^2$ with isometric copies of the regular octagon.

In the same way as it was done for $\mathbb{R}^2$ and the torus $\mathbb{T}^2$, one can define a map π from $\mathbb{H}^2$ to the surface S, using the just constructed tiling. This is a covering map and, as $\mathbb{H}^2$ is simply connected, it follows that $\mathbb{H}^2$ is the universal cover of S.

The existence of this covering map implies that the fundamental group $\pi_1(S)$ acts freely on $\mathbb{H}^2$ by isometries (deck transformations). Here we have an example of a group acting on a geometric space, so that we can extract some properties of the group. On one side, the deck transformation associated to an element of $\pi_1(S)$ is a *hyperbolic isometry*. This means that there is a geodesic γ in $\mathbb{H}^2$ (called the *axis*) in which the isometry acts like a translation, and the image of any other point of $\mathbb{H}^2$ is determined by the fact that orthogonal geodesics must be sent to orthogonal geodesics, and distances in $\mathbb{H}^2$ must be preserved. See Figure 2.7.

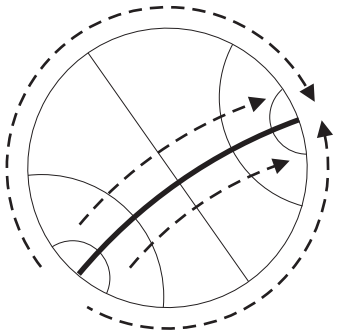

Figure 2.7: A hyperbolic isometry. The dark geodesic is the axis. Every geodesic perpendicular to the axis is sent to a geodesic perpendicular to the axis.

Remark 2.2. *As suggested by the referee, it could have been more intuitive to use the model of the genus 2 surface as a decagon with opposite sides identified. But we just wanted to show that the usual way to model the surface as the octagon in Figure 2.3 can also be used. Also, the model we use is a natural generalization of the case of the torus.*

Using this very particular structure of deck transformations, one can deduce some algebraic properties of the fundamental group $\pi_1(S)$. First, $\pi_1(S)$ has no torsion, as deck transformations have infinite order. Second, $\mathbb{Z}^2$ is not a subgroup of $\pi_1(S)$, as two deck transformations do not commute if they have distinct axes, and they generate a cyclic subgroup if they have the same axis.

The covering map $\pi : \mathbb{H}^2 \to S$ endows S with a hyperbolic structure, that is, a geometry coming from the geometry of $\mathbb{H}^2$.

2.2. Geodesics in a closed surface

Let S be a closed surface of genus $g > 1$, and let π be the covering map defined in the previous section. Given a curve γ in S, it can be lifted to a curve in $\mathbb{H}^2$ via the map π. There are infinitely many ways to lift γ, but all of them are related by deck transformations, and the lifting of γ depends only on the lifting of a single point. A geodesic in S is defined as a curve whose lifting by π is a geodesics in $\mathbb{H}^2$.

If γ is a closed curve, the lifting of γ in $\mathbb{H}^2$ can be extended indefinitely, so if γ is a geodesic in S then it can be lifted to a complete geodesic in $\mathbb{H}^2$. We shall be interested in *essential*

curves, that is, closed curves which are not null-homotopic. One of the most interesting facts about essential curves in S is that they can always be deformed to become a geodesic.

Theorem 2.3. *Every essential closed curve in S is freely homotopic to a unique closed geodesic.*

See [9] for a detailed proof. We will just explain how one can obtain the closed geodesic associated to an essential closed curve c. Namely, as c is essential, it determines a non-trivial element of $\pi_1(S)$ (up to conjugacy). This element corresponds to a deck transformation, which has an axis $\widetilde{\gamma}$. The projection $\gamma = \pi(\widetilde{\gamma})$ is the geodesic in S freely homotopic to c. Notice that, if we take a conjugate element in $\pi_1(S)$, the axis could be different, but in any case it would be the image of $\widetilde{\gamma}$ by a deck transformation, so its projection by π would be the same geodesic γ.

Another way to see the geodesic γ homotopic to c is to consider the 'walls' of the polygon P which are crossed by c. This allows to lift c to $\mathbb{H}^2$, and we see that the lifting approaches (with the Euclidean metric) two points in the boundary of $\mathbb{H}^2$. These two points determine a geodesic (a circle orthogonal to the boundary), whose image by π is precisely γ.

2.3. Laminations and foliations. Nielsen-Thurston classification

Recall that the Mapping Class Group of a closed surface S is $MCG(S) = \pi_0(\mathrm{Homeo}^+(S))$. Hence $MCG(S)$ acts on S by orientation preserving homeomorphisms, up to deformation. But this is not a group action, as the image of a point of S under an element of $MCG(S)$ is not even well defined. However, we can concentrate on how this transforms some suitable subsets of the surface. For instance, consider a simple closed curve c in S. Any homeomorphism of S will send c to a simple closed curve. If c is essential (i.e. not null-homotopic), then both c and its image are freely homotopic to unique closed geodesics. Therefore, a mapping class sends simple closed geodesics to simple closed geodesics.

We can then concentrate on the set of closed geodesics of S, and see how $MCG(S)$ acts on it. But it will be more convenient to enlarge this set, to include disjoint unions of simple closed geodesics, and even non-closed simple geodesics (which will appear as limits of iterated images of a simple closed geodesic by a mapping class). If we have such a non-closed simple geodesic, we will consider also its closure, which is actually an union of simple closed geodesics [9]. In such a way, we will always consider closed subsets of S:

Definition 2.4. *A geodesic lamination (or just a lamination) on S is a closed subset of S which is a nonempty disjoint union of simple geodesics.*

It is known [9] that the decomposition of a lamination into an union of simple geodesics is unique. So one can say without ambiguity that a geodesic belongs to a given lamination.

The set of laminations on S is denoted $\Lambda(S)$ and, clearly, $MCG(S)$ acts on $\Lambda(S)$. The reason why one includes non-closed geodesics is to obtain a compact set. Actually, in $\Lambda(S)$ we can consider the Hausdorff distance: The distance between two laminations is at most ϵ if each one is contained in an ϵ-neighborhood of the other. Then one has:

Theorem 2.5. [9] *$\Lambda(S)$ is compact (with respect to the Hausdorff distance), and $MCG(S)$ acts on $\Lambda(S)$ by homeomorphisms.*

Here we have a very interesting action of $MCG(S)$ on a geometric space. Elements in $MCG(S)$ are classified according to the way they act on $\Lambda(S)$ [29]. This is called the Nielsen-Thurston classification, and it goes as follows.

An element in $MCG(S)$ is called **periodic** if it has finite order. That is, if it is represented by an automorphism h of S such that h^n is homotopic to the identity, for some $n > 0$.

An element in $MCG(S)$ is called **reducible** if it preserves a simple closed 1-submanifold of S. That is, it is represented by a homeomorphism h which sends a simple closed 1-submanifold c of S to another one which is homotopic to c. This 1-manifold c correspond to an element of $\Lambda(S)$ which is an union of simple closed geodesics (corresponding to simple closed curves in S), and is a fixed point under the action of the mapping class of h.

Finally, if an element $[h]$ in $MCG(S)$ represented by an automorphism h is neither periodic nor reducible, its action on $\Lambda(S)$ is very special. By the work of Thurston [29], we know that $[h]$ preserves two laminations, the *stable lamination* L^s, and the *unstable lamination* L^u. Moreover, the action of $[h]$ pushes every lamination (except the stable and unstable one) away from the unstable lamination and closer and closer to the stable one.

These laminations can be very easily understood thanks to the following result. Roughly speaking, if one takes any essential simple closed curve c in S, and one considers the sequence of simple closed curves $h^n(c)$ for $n \geq 0$, this sequence tends to the stable lamination L^s. Actually, this is not exactly true as stated, as a finite number of geodesics not belonging to L^s can also appear in the limit, and they can be permuted at every iteration, but this is solved by stating the result in the following way:

Theorem 2.6. [9] *If $[h] \in MCG(S)$ is neither periodic nor reducible, then there exists $m > 0$ so that, denoting $g = h^m$, for every essential simple closed curve c in S, the sequence $[g^n](C)$ tends to K as n tends to infinity, where K is one of the finitely many laminations containing L^s.*

This means that one can *see* the stable lamination by applying h many times to any essential simple closed curve. The unstable lamination L^u is obtained in a symmetric way, applying h^{-1} to any essential simple closed curve.

But mapping classes not only act nicely on laminations. There are other objects which can be constructed from laminations, and provide yet another very interesting geometric action. First, it can be shown that, given a lamination L on the surface S which is preserved by an automorphism, its complement $S \backslash L$ is a subsurface of S all of whose connected components are polygons with a finite number of geodesic sides. If $[h]$ is neither periodic nor reducible, the complement of the stable and unstable laminations can be collapsed in such a way that the laminations become singular foliations, where the singular points (called *prong singularities*) look like in Figure 2.8.

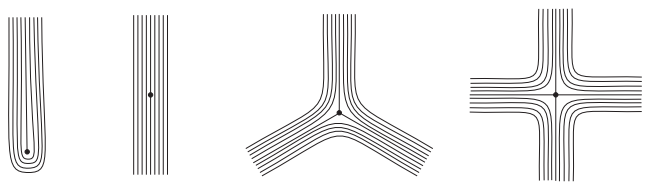

Figure 2.8: Foliations at singular points. Here we see just k-prong singularities for $k = 1, \ldots, 4$, but there are similar pictures for bigger k.

These two foliations, $\mathcal{F}^s$ and $\mathcal{F}^u$, obtained from the stable and the unstable laminations, respectively, can be endowed with transverse measures, μ^s and μ^u, giving a sense of distance between leaves. The amazing fact is that the mapping class $[h]$ act on these transverse

measured foliations by preserving the foliations as a set, and scaling the measures by a real factor $\lambda > 1$ in the case of the unstable foliation, and by λ^{-1} is the case of the stable one. Roughly speaking, under the action of $[h]$, leaves in $\mathcal{F}^s$ get closer to each other, and leaves in $\mathcal{F}^u$ get farther to each other, in both cases by a factor $\lambda > 1$.

It is then said that a homeomorphism of S is **pseudo-Anosov** if it preserves two transverse measured foliations, scaling the measure of one of them by λ, and the other by λ^{-1}. A mapping class is pseudo-Anosov if it can be represesnted by a pseudo-Anosov homeomorphism. After the previous discussion, we have the famous Nielsen-Thurston classification of mapping classes.

Theorem 2.7. [29] *Every element in MCG(S) is either periodic, or reducible, or pseudo-Anosov.*

We notice that an element can be, at the same time, periodic and reducible, and that is the only possible intersection between the above kinds of mapping classes. In order to get an actual trichotomy, the second kind could be replaced by that of reducible and non-periodic mapping classes.

2.4. Hyperbolic structure on braids

We saw in the previous section how mapping classes of a closed surface act on geodesic laminations. But, what about braids? Braids are mapping classes of the punctured disc, which is not a closed surface, but we can use a similar approach.

First of all, in order to be able to apply the same theory of mapping classes, we will collapse the boundary of the disc to be a new puncture. In terms of mapping classes, this means that rotating the boundary of the disc by a full twist (which corresponds to a non-trivial braid called Δ^2) is now a trivial mapping class. In other words, this corresponds to replace the group B_n by $B_n/\langle\Delta^2\rangle$. As the subgroup $\langle\Delta^2\rangle$ is precisely the center of B_n, collapsing the boundary of $\mathbb{D}_n$ to a puncture corresponds to quotient B_n by its center.

The space obtained when collapsing the boundary of $\mathbb{D}_n$ to a puncture is the $(n+1)$-times punctured sphere $\mathbb{S}^2_{n+1}$. Now we can tile the hyperbolic plane with cut-open copies of $\mathbb{S}^2_{n+1}$, as we previously did with cut-open copies of a closed surface. First, consider segments $a_1, \dots, a_n$ in $\mathbb{D}_n$ joining all punctures and the boundary along the diameter of the disc (these segments are denoted a, b, c in Figure 2.9). They can also be seen in $\mathbb{S}^2_{n+1}$, and we can cut this latter surface along these paths. The resulting surface is a polygon P with $2n$-sides, with the vertices removed. In order to obtain $\mathbb{S}^2_{n+1}$ one must identify the sides of the polygon P as in Figure 2.9.

Now we can place the polygon P in the hyperbolic plane $\mathbb{H}^2$ as in Figure 2.10: The sides correspond to geodesics, and the missing vertices correspond to points in the boundary of the disc representing $\mathbb{H}^2$. This is called an ideal polygon. One can tile the whole hyperbolic plane $\mathbb{H}^2$ just by reflecting this polygon along its geodesic sides, and keep reflecting the newly obtained copies. Labeling the edges in the appropriate way, we see that $\mathbb{H}^2$ is the universal cover of $\mathbb{S}^2_{n+1}$, and this endows $\mathbb{S}^2_{n+1}$ with a hyperbolic structure.

The main difference with respect to the case of closed surfaces is that, as the punctures of $\mathbb{S}^2_{n+1}$ correspond to points in the boundary of the disc representing $\mathbb{H}^2$, points which are close to some puncture of P with respect to the Euclidean distance are very distant with respect to the hyperbolic distance. Hence, in a neighborhood of each puncture the metric 'explodes', and the geometry is quite different to what one could expect.

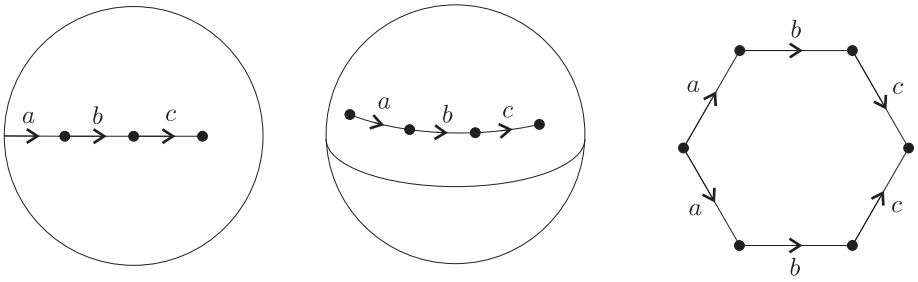

Figure 2.9: A punctured disc, a punctured sphere obtained collapsing the boundary to a puncture, and a polygon representing the punctured sphere.

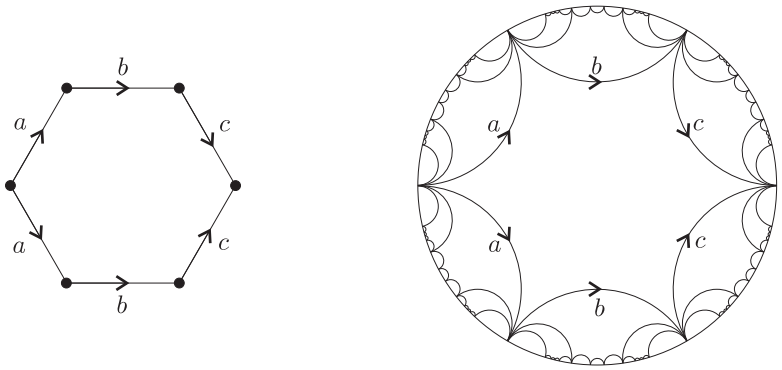

Figure 2.10: How to tile the hyperbolic plane with ideal polygons corresponding to 4-punctured spheres.

As an example, one can show that every puncture has a neighborhood which is disjoint from all possible simple closed geodesics in $\mathbb{S}^2_{n-1}$. In other words, simple closed geodesics are far apart from the punctures. To see this, it is better to represent $\mathbb{H}^2$ as the upper half plane of $\mathbb{R}^2$, where geodesics correspond to vertical lines and to arcs of circles which are orthogonal to the real line. This representation can be obtained from the one we are using, by transforming the complex plane using the Möbius transformation $z \mapsto \frac{-iz-i}{z-1}$.

We can assume that the point $1 \in \mathbb{C}$ is one of the ideal vertices of our tiling, as in Figure 2.10, corresponding to one of the punctures p_i of $\mathbb{S}^2_{n+1}$. Then the tiling of the upper half plane representing $\mathbb{H}^2$ will look like in Figure 2.11. We can consider neighborhoods of p_i consisting of points in $\mathbb{R}^2$ whose vertical coordinate is greater than a certain number.

Now suppose that a closed geodesic γ in $\mathbb{S}^2_{n-1}$ has a point p which is close to the puncture p_i. This means that if we consider a lift $\tilde{p}$ of p in the half-plane $\mathbb{H}^2$, in a region with two vertical straight line boundaries, $\tilde{p}$ will have a big vertical coordinate, and γ will be lifted to a big half-circle, orthogonal to the real line. But we could also lift p to another copy of the polygon P, which is obtained from the previous one by a horizontal translation, and γ would be lifted to a big half-circle which is obtained from the previous one by a horizontal translation. Now, if the vertical coordinate of $\tilde{p}$ is big enough, those half-circles will intersect, meaning that the geodesic γ intersects itself, hence it is not simple. Therefore, there is a neighborhood of p_i (determined by a given height of in the upper half-plane), which is disjoint from every simple

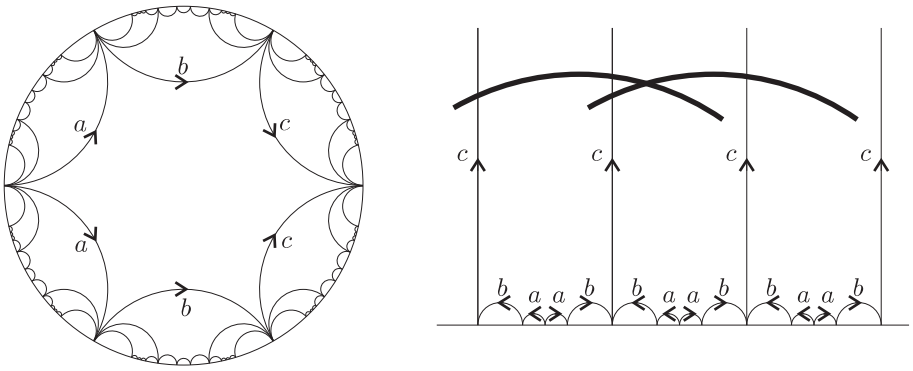

Figure 2.11: Tiling the upper half plane with ideal polygons shows that simple geodesics cannot be close to a puncture: If they were, distinct lifts would intersect.

closed geodesic of $\mathbb{S}^2_{n+1}$, as we wanted to show. This argument was shown to me by Juan Souto.

We can now define laminations in $\mathbb{S}^2_{n+1}$ in the same way as we did for closed surfaces, and consider the action of the braid group (modulo Δ^2) on the set of laminations $\Lambda(\mathbb{S}^2_{n+1})$. Using this action, we can classify a braid in B_n by sending it to $B_n/\langle\Delta^2\rangle$ and seeing how it acts on $\Lambda(\mathbb{S}^2_{n+1})$. Hence, braids are classified into periodic, reducible and pseudo-Anosov.

Periodic braids correspond to elements which are sent to an element of finite order in $B_n/\langle\Delta^2\rangle$. That is, a **periodic** braid is a root of Δ^m, for some $m > 0$. A braid is **reducible** if it preserves a family of essential simple closed curves in $\mathbb{D}_n$. And a braid is **pseudo-Anosov** if it preserves two families of transverse measured foliations, scaling the measure of one of them by a real number $\lambda > 1$, and the measure of the other foliation by λ^{-1}. These foliations can be obtained, as in the case of closed surfaces, by iterating the action of the braid (or its inverse) on a given curve (homotopic to a geodesic), taking the limit of the obtained sequence of geodesics in the Hausdorff metric, and collapsing the complement of this limit lamination, obtaining singular points and leaves as in Figure 2.8.

In Figure 2.12 we can see the image of a curve under iterated application of the pseudo-Anosov braid $\sigma_1\sigma_2^{-1} \in B_3$, and its inverse, respectively. We can see that the resulting foliations have four 1-prong singularities: one for each puncture, and one corresponding to the boundary of the disc.

This geometric interpretation of braids can be used to prove algebraic facts in the group B_n. For instance, in [25] (see also [23]) this is used to show that the centralizer of a pseudo-Anosov element in a mapping class group is virtually cyclic. Translated to braids (see [15]), this implies that the centralizer of a pseudo-Anosov braid β is isomorphic to $\mathbb{Z}^2$, with a basis consisting of a pseudo-Anosov element (usually β itself), and a periodic element (usually Δ^2).

Another algebraic result which can be easily proved using these geometric techniques is the following:

Theorem 2.8. [16] *Pseudo-Anosov elements in B_n have unique roots. That is, if η is a pseudo-Anosov braid and $\alpha^n = \beta^n = \eta$ for some braids α, β and some $n > 0$, then $\alpha = \beta$.*

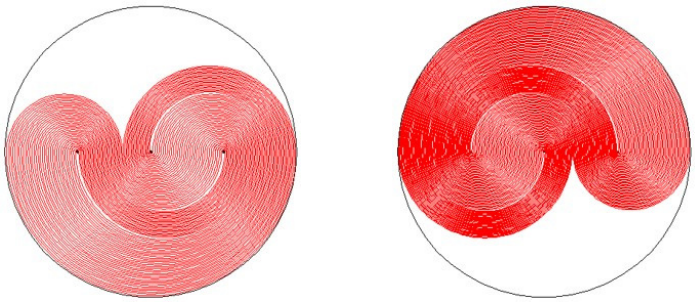

Figure 2.12: The limit laminations of the braid $\sigma_1\sigma_2^{-1} \in B_3$ look like the curves in this figure. On the left (resp. right) hand side, the image of a curve under iterated application of $\sigma_1\sigma_2^{-1}$ (resp. $(\sigma_1\sigma_2^{-1})^{-1}$).

Proof. First, α and β must be pseudo-Anosov as this property is preserved by taking powers. They also preserve the same foliations as η, and the stretch factor λ must be the same for both braids, as it is the n-th root of the stretch factor of η.

Then $\alpha\beta\alpha^{-1}\beta^{-1}$ also preserves the same foliations, with stretch factor equal to 1. This implies that $\alpha\beta\alpha^{-1}\beta^{-1}$ is a periodic element. But in B_n the periodic elements are well known, and the only one whose sum of exponents is 0 is the trivial element. Hence $\alpha\beta\alpha^{-1}\beta^{-1} = 1$, so α and β commute.

Therefore $1 = \eta\,\eta^{-1} = \alpha^n\beta^{-n} = (\alpha\beta^{-1})^n$. As the braid group B_n has no torsion, it follows that $\alpha = \beta$. □

We have then seen how these geometric aspects of elements of the braid group, coming from the action of B_n on a geometric space, can help to extract algebraic information about the group itself.

3. The curve complex

When studying actions of a group on a geometric space, it is crucial that the maps determined by the elements of the group satisfy some geometric properties: For instance, to be isometries of the geometric space. In that case we say that the group act on the space by isometries.

The actions of B_n that we saw in previous sections are not by isometries. But one just needs to look at the appropriate space. We will find a geometric space with an interesting metric, on which B_n acts by isometries.

Let S be an orientable surface. We saw that if the universal cover of S is $\mathbb{H}^2$, the action of $MCG(S)$ on S sends closed geodesics to closed geodesics. Now instead of looking at the space of laminations (closed unions of geodesics), we will just look at the set of simple, closed, essential curves, and we will construct a geometric space out of it. This is the *complex of curves*, or the *curve complex*, defined by Harvey [19].

The **curve complex** of the surface S is a simplicial complex, denoted $\mathcal{C}(S)$, which is a family of subsets of the set of simple, closed, essential curves in S, up to isotopy. That is, a 0-simplex corresponds to the isotopy class of a simple, closed, essential curve in S.

For $k > 0$, a k simplex is a collection of $k+1$ simple, closed, pairwise non-isotopic essential curves, up to isotopy, which can be realized to be **disjoint**. For instance, an edge of $\mathcal{C}(S)$, or a 1-simplex, is given by a pair of disjoint essential curves in S, which are not isotopic to each other. The 1-skeleton of this complex, that is, the set of 0-simplices and 1-simplices, is called the *curve graph* of S, denoted $\mathcal{C}_1(S)$.

In order to give some examples, let $S_{g,n}$ be the orientable surface of genus g with n boundary components (or, alternatively, n punctures, which yields the same curve complex).

Example 3.1. *If $S = S_{0,n}$, $n \leq 3$, is a sphere with at most 3 boundary components, then $\mathcal{C}(S)$ is empty. Indeed, if S is a closed sphere, every simple curve is homotopic to a point, hence it is not essential. If S has some boundary component (or some puncture), one can imagine S as a disc with at most two holes. If a simple closed curve encloses no hole, it is homotopic to a point. If it encloses just one hole, it is homotopic to the boundary component (or puncture) corresponding to that hole. If it encloses two holes, it is homotopic to the boundary component (or puncture) corresponding to the boundary of the disc. In any case, one can see that there is no essential, simple, closed curve in S. So $\mathcal{C}(S)$ is empty.*

Example 3.2. *If $S = S_{0,4}$ then $\mathcal{C}(S)$ is an infinite collection of isolated 0-simplices. In the same way as above, one can imagine $S_{0,4}$ as a disc with 3 holes. An essential curve is a curve enclosing exactly two punctures. There is an infinite number of such curves. But once that such a curve is chosen, the complement consists of two copies of $S_{0,3}$ (two pairs of pants), hence there is no essential curve disjoint to the chosen one. That is, there is no edge in $\mathcal{C}(S)$.*

Example 3.3. *If $S = S_{1,0}$ or $S = S_{1,1}$ then $\mathcal{C}(S)$ is an infinite collection of isolated 0-simplices. A simple closed curve in the torus $S_{1,0}$ is isotopic to a geodesic, which can be seen in the universal cover $\mathbb{R}^2$ of the torus as a straight line of rational slope. We can then have an infinite number of non-isotopic simple closed curves. But if two such curves are not isotopic, they have distinct slopes, hence they intersect. Therefore there are no edges in $\mathcal{C}(S)$. The case of $S_{1,1}$ is analogous, as every simple closed curve is isotopic to a geodesic avoiding the boundary component. With more than one boundary component this argument does not work anymore, and there exist pairs of disjoint simple, closed, essential curves.*

The above examples, $S_{0,n}$ for $n \leq 4$ and $S_{1,n}$ for $n = 0, 1$ are called the *sporadic surfaces*. The remaining cases share a common property:

Theorem 3.4. [19] *If $S = S_{g,n}$ is not a sporadic surface, then $\mathcal{C}(S)$ is connected.*

Remark 3.5. *In the sporadic cases $S_{0,4}$, $S_{1,0}$ and $S_{1,1}$, in which the usual definition of $\mathcal{C}(S)$ yields an infinite set of points, the definition of the complex of curves is sometimes modified, placing edges between vertices corresponding to curves of smallest possible intersection number. This means intersection number 2 for $S_{0,4}$ and intersection number 1 for $S_{1,0}$ and $S_{0,2}$. With this modified definition, the complex of curves becomes the well-known Farey graph, which is connected (and hyperbolic)* [24, 26].

We can now transform $\mathcal{C}(S)$ into a metric space. Actually, in order to show the geometric results we are interested in, and to simplify the discussion, it suffices to deal with the 1-skeleton of $\mathcal{C}(S)$ (the *curve graph*), denoted $\mathcal{C}_1(S)$. It consists of vertices (isotopy classes of essential simple closed curves), and edges (pairs of isotopy classes of such curves which can be realized disjointly). Imposing the each edge has length 1, providing it with the standard Euclidean metric, it follows that $\mathcal{C}_1(S)$ is a geodesic metric space. This is a natural metric: The distance between two vertices in $\mathcal{C}_1(S)$ is precisely the minimal number of edges in a path from one vertex to the other. As $\mathcal{C}_1(S)$ is (path) connected, if S is not sporadic, every two vertices are at finite distance, and can be joined by at least one geodesic. We denote by $d(\alpha, \beta)$ the distance between two vertices of $\mathcal{C}_1(S)$.

Example 3.6. *Let $S = S_{2,0}$ and consider the curves α, β and γ in Figure 3.1. The isotopy classes of α and γ cannot be represented by disjoint curves, hence the pair (α, γ) is not an edge of $\mathcal{C}(S)$. But α is disjoint from β, and β is disjoint from γ, hence there is an edge in $\mathcal{C}(S)$ from α to β, and an edge from β to γ. Therefore $d(\alpha, \gamma) = 2$.*

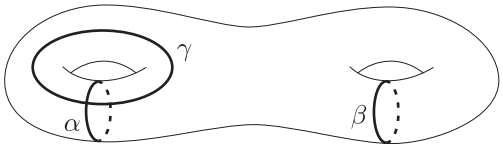

Figure 3.1: Curves α, β and γ are vertices of the curve complex of $S_{2,0}$.

Actually, there is an upper bound for the distance, in $\mathcal{C}_1(S)$, between two given vertices. Denote $i(\alpha, \beta)$ the intersection number of α and β, that is, the minimal number of intersection points between two realizations of α and β. Then one has:

Theorem 3.7. [24] *If S is not sporadic and α, β are two essential, simple closed curves in S, then:*

$$d(\alpha, \beta) \leq \log_2(i(\alpha, \beta)) + 2$$

One can define a natural metric in $\mathcal{C}(S)$ in such a way that $\mathcal{C}_1(S)$ is quasi-isometric to $\mathcal{C}(S)$ (distances between two points in $\mathcal{C}_1(S)$ are roughly the same when considered in $\mathcal{C}_1(S)$ than when considered in $\mathcal{C}(S)$). Many geometric properties are preserved by quasi-isometry: This is why we will be interested only in $\mathcal{C}_1(S)$.

Now there is an easy observation: An automorphism of S sends disjoint curves to disjoint curves. Hence, it preserves the distance between 0-simplices in $\mathcal{C}_1(S)$. In other words, $MCG(S)$ acts on $\mathcal{C}_1(S)$ by isometries.

But it does not suffice to have a metric space on which a group act by isometries. The space should have some interesting geometry. This is what we explain in the next section.

3.1. The curve complex is hyperbolic

We saw that the hyperbolic plane $\mathbb{H}^2$ is a very interesting geometric space. It would be nice to have geometric spaces sharing some of the good properties of $\mathbb{H}^2$. These spaces are called Gromov-hyperbolic, or δ-hyperbolic. The main property, from which many other properties can be deduced, is one that we observed in $\mathbb{H}^2$: geodesic triangles are *thin*, in some sense. This is one of the several equivalent definitions:

Definition 3.8. *A geodesic metric space X is Gromov-hyperbolic with constant $\delta \geq 0$ (or it is δ-hyperbolic) if for every geodesic triangle, there is a point $c \in X$ which is at distance at most δ from each of its three sides. See Figure 3.2.*

In the Euclidean plane $\mathbb{R}^2$, triangles are flat, not thin. For any given δ, one can take a triangle big enough so that no ball of radius δ touches the three sides. But in the hyperbolic plane $\mathbb{H}^2$, geodesic triangles are thin [7].

Another example of Gromov-hyperbolic space is a tree, with the natural distance in which every edge has length one. Such a tree is actually a 0-hyperbolic space, as in every geodesic triangle, there is always a point contained in the three sides (see Figure 3.3). As a particular case, the real line $\mathbb{R}$ is 0-hyperbolic.

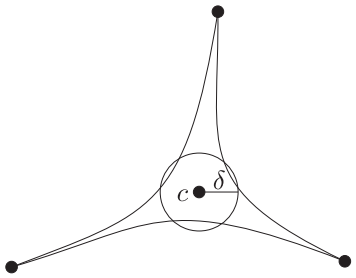

Figure 3.2: A geodesic triangle is δ-thin if its sides are touched by a circle of radius δ. A geodesic metric space is δ-hyperbolic if all geodesic triangles are δ-thin.

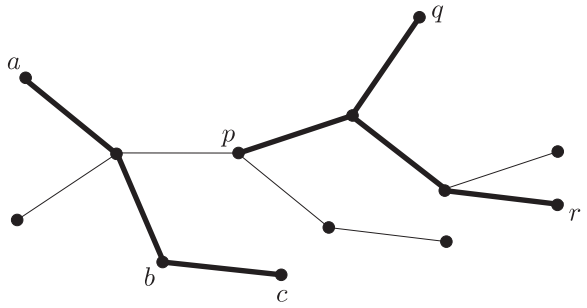

Figure 3.3: Geodesic riangles in a tree are always 0-hyperbolic. In this picture, the geodesic triangles *abc* and *pqr*.

Gromov-hyperbolic spaces are very common in Geometric Group Theory: A group is said to be **hyperbolic** if its Cayley graph with respect to some generating set is Gromov-hyperbolic. This property is actually independent of the generating set, so one can consider any Cayley graph of the group. By the above discussion about trees it follows that free groups are hyperbolic. Actually, they are the main example of hyperbolic groups.

Hyperbolic groups have many interesting algebraic properties. For instance, they have solvable word and conjugacy problems, and moreover they are biautomatic.

One should note that, if a group is hyperbolic, it does not contain subgroups isomorphic to $\mathbb{Z}^2$. This is one of the tools used to prove that some group is not hyperbolic. This shows that, in general, mapping class groups are not hyperbolic, as there are usually many pairs of commuting mapping classes not generating a cyclic subgroup.

As we are interested mainly in the braid group B_n, or more generally in a mapping class group $MCG(S)$, which is not hyperbolic, we will not consider its Cayley graph as a metric space. But we already know a space, $\mathcal{C}(S)$, on which $MCG(S)$ act by isometries. Fortunately this space is interesting enough, as shown by Masur and Minsky:

Theorem 3.9. [24] *The curve complex $\mathcal{C}(S)$ is Gromov-hyperbolic.*

Hyperbolicity is preserved by quasi-isometry, so showing that $\mathcal{C}(S)$ is Gromov-hyperbolic is equivalent to show that so is $\mathcal{C}_1(S)$. Later we will sketch a proof of the hyperbolicity of $\mathcal{C}_1(S)$. Now let us mention that Masur and Minsky not only showed that the curve complex is Gromov-hyperbolic, but also that the action of a pseudo-Anosov element on $\mathcal{C}(S)$ is *hyperbolic*, having *North-South* dynamics (similar to the action we studied on the hyperbolic plane), and that the orbits are quasi-geodesics. This is a very interesting kind of action, which allows to show several important algebraic consequences, like the following:

Theorem 3.10. [24] *Given an orientable surface S, there exists a constant K such that, if $h_1, h_2 \in MCG(S)$ are conjugate pseudo-Anosov elements, there is a conjugating element ω with $|\omega| \leq K(|h_1| + |h_2|)$.*

In other words, the length of a minimal conjugator from one pseudo-Anosov to another, is bounded linearly by the lengths of the conjugated elements. To show this result, one needs to use not only the preceding results, but also some more, very deep and difficult theory.

Unfortunately, the above is an existence result, giving no method to compute the constant K. If the constant K could be explicitly computed, we would have a *fast* algorithm to decide whether a braid is periodic, reducible or pseudo-Anosov, as shown by Calvez in [8]. To be more precise, the time taken by the algorithm would be a quadratic function with respect to the length of the input braid and its number of strands.

In any case, the fact that the curve complex is hyperbolic is an extraordinary fact. Recently, several authors independently [1, 6, 11, 22] proved something even more surprising: The constant δ which bounds the thinness of triangles in $\mathcal{C}_1(S)$ is independent of the surface S! That is, thinness of triangles is uniformly bounded, regardless of the genus or the number of boundary components of S. Moreover, one can take a small number like $\delta = 17$ [22].

All these results can be applied to the braid group, considering the complex of curves on the punctured disc, as the action of an element of B_n on the simple closed curves on the punctured disc remains the same if we consider the element in $B_n/\langle\Delta^2\rangle$, via the natural projection.

In the following section, we will sketch the proof in [22] showing that $\mathcal{C}_1(S)$ is 17-hyperbolic.

3.2. Hyperbolicity of the curve graph

In this final section we will describe the proof by Hensel, Przytycki and Webb [22], showing that the curve graph $\mathcal{C}_1(S)$ is 17-hyperbolic. Instead of giving all the details of the proof, we will just give the arguments that may lead the reader to try to prove that his/her favorite geometric space is hyperbolic.

In order to do this, another graph closely related to $\mathcal{C}_1(S)$ is studied in [22]. At first, only surfaces with boundary are considered, but the curve graph of a closed surface S is a 1-Lipschitz retract of the curve graph of the punctured surface $S\backslash\{p\}$ [18, 28], and punctures are equivalent to boundary components from the point of view of curve complexes, so this is actually not a restriction. Hence we can consider that S has some boundary component.

The *arc graph*, $\mathcal{A}_1(S)$, is a graph whose vertices are simple arcs in S with endpoints lying in the boundary of S. Two such vertices are joined by an edge if the curves can be realized disjointly: the same definition as for the curve graph. When one considers the graph whose vertices are either arcs or closed curves, with the same adjacency condition, one obtains the *arc and curve graph* $\mathcal{AC}_1(S)$.

To avoid cumbersome notation, as we will just study graphs in this section, we will denote $\mathcal{C}(S)$, $\mathcal{A}(S)$ and $\mathcal{AC}(S)$ the curve graph, the arc graph, and the arc and curve graph, respectively.

In [22] it is shown that $\mathcal{A}(S)$ is 7-hyperbolic, and then this is used to deduce that $\mathcal{C}(S)$ is 17-hyperbolic. We are going to focus on the proof that $\mathcal{A}(S)$ is 7-hyperbolic, trying to generalize the strategy as much as possible, so that it can be adapted to other spaces.

The main strategy, if it is difficult to determine geodesics in a geometric space, is to try to find a *nice* family of paths in the space, forming thin triangles, which is closed under taking subpaths. In this way one can show that these paths are not far from geodesics, and this will imply thinness of geodesic triangles.

In the graph $\mathcal{A}(S)$, these paths will be called *unicorn paths*, and were introduced by Hatcher [20]. Given simple arcs a and b in S (two vertices of $\mathcal{A}(S)$), we will define a unicorn path in $\mathcal{A}(S)$ going from a to b, denoted $\mathcal{P}(a, b)$.

Let us orient the arcs a and b, and denote $\alpha \in \partial(S)$ the starting point of a, and $\beta \in \partial(S)$ the starting point of b. If a and b can be realized disjointly, the unicorn path $\mathcal{P}(a, b)$ is just the path (a, b) (we will denote a path in $\mathcal{A}(S)$ by a sequence of vertices).

Now suppose that the arcs a and b cannot be realized disjointly (i.e. they are not adjacent in $\mathcal{A}(S)$). We will realize them by curves having minimal intersection (we say that the curves are in minimal position). Let $a_1, \dots, a_k$ be the intersection points, ordered according the chosen orientation of a. The same points can be denoted $b_1, \dots, b_k$ following the orientation of b, and the order is not necessarily the same as before. See Figure 3.4.

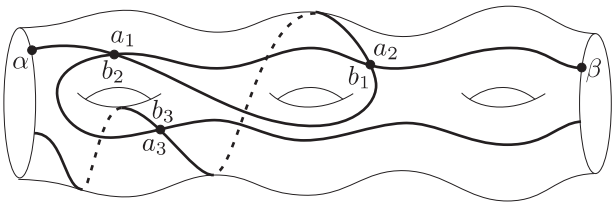

Figure 3.4: Two arcs in minimal position.

Let us represent the situation regarding the intersection points of a and b using the diagram in Figure 3.5. The vertical lines correspond to the curves a and b, and we draw segments connecting points of a and b which intersect. Notice that we orient a upwards and b downwards, and that the transversal segments may cross each other, depending on the ordering of $a_1, \dots, a_k$ with respect to the ordering of $b_1, \dots, b_k$.

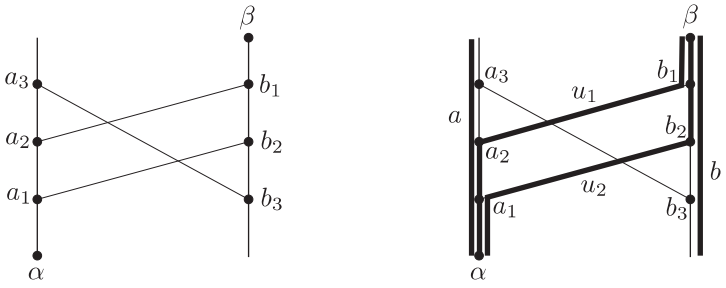

Figure 3.5: A diagram representing the intersection points of the curves in Figure 3.4. And the unicorn path $\mathcal{P}(a, b) = (a, u_1, u_2, b)$.

Definition 3.11. *With the above notations, a* **unicorn arc** *obtained from a and b is an arc that starts at α, goes along a until an intersection point $a_i = b_j$, and then goes along b up to β, provided the defined arc is simple.*

Notice that a unicorn arc is determined by an intersection point (once the orientations of a and b are chosen). Looking at the diagram in the right hand side of Figure 3.5, it is easy to see which intersection points determine a unicorn arc: A transversal segment correspond

to a unicorn arc if and only if it does not cross a segment of higher slope. For instance, in Figure 3.5, one sees that a_1 and a_2 determine unicorn arcs, but a_3 does not, as the segments a_1 and a_2 cross a_3 and have higher slope (so they correspond to intersection points of the possible unicorn arc).

It follows from this interpretation that the collection of all unicorn arcs obtained from a and b correspond to segments that do not cross. One can determine all intersection points a_i corresponding to unicorn arcs as follows: First take a_1, which is always unicorn. Then, at every step, take the lowest a_i whose segment does not intersect the previously chosen segments. At the end you will take the segment corresponding to b_1, and you are done. We can denote $u_r, u_{r-1}, \cdots, u_1$ the obtained unicorn arcs, which correspond to vertices of $\mathcal{A}(S)$.

Definition 3.12. *In the above situation, the* **unicorn path** *from a to b is the sequence of vertices $(a = u_0, u_1, \ldots, u_r, u_{r+1} = b)$ of $\mathcal{A}(S)$.*

It is important to notice that two consecutive vertices of a unicorn path are arcs that can be realized disjointly in S. Indeed, both curves coincide, except on a subarc of a going from some a_i to some a_j, and in a subarc of b going from some b_k to some b_l, where $a_i = b_k$ and $a_j = b_l$ (see Figure 3.5). These two segments cannot intersect, otherwise there would be a a segment corresponding to a unicorn arc, between the two chosen ones. Now one can perturb slightly the coinciding subarcs, so that they become parallel: One must slide each one to the left or to the right, depending on the incidence of the subarcs of a and b at the intersection point (see Figure 3.6). In this way, both arcs are realized disjointly, and this means that a unicorn path $\mathcal{P}(a, b)$ is actually a path in $\mathcal{A}(S)$.

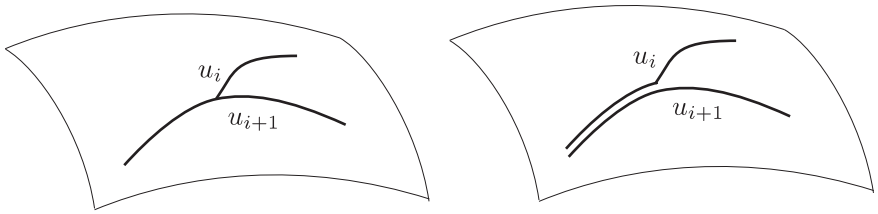

Figure 3.6: Two consecutive unicorn arcs in a unicorn path can always be realized disjointly: The common subarcs can can be isotoped so that they become parallel.

Recall the two curves a and b in Figure 3.4. The unicorn path $\mathcal{P}(a, b)$ is given in Figure 3.7.

To see that unicorn paths are good candidates to show the hyperbolicity of $\mathcal{A}(S)$, let us prove that they form thin triangles.

Lemma 3.13. [22] *Let a, b, c be three vertices of $\mathcal{A}(S)$. Realize them in minimal position, choose an orientation for each one, and let $\mathcal{P}(a, b)$, $\mathcal{P}(b, c)$ and $\mathcal{P}(c, a)$ be the corresponding unicorn paths. Then, for every $u \in \mathcal{P}(a, b)$, there exists some $v \in \mathcal{P}(b, c) \cup \mathcal{P}(c, a)$ which is adjacent to u.*

Proof. We can draw the situation as in Figure 3.8, where we have just represented segments corresponding to unicorn arcs. The missing segments (if they exist) would have smaller slope than the ones joining the same pair of arcs.

Consider a unicorn arc $u \in \mathcal{P}(a, b)$. If u can be realized disjoint from c, then $d(c, u) \leq 1$ and c is at distance at most 1 from $\mathcal{P}(a, b)$ (the distance from c to $\mathcal{P}(a, b)$ is by definition the minimal length among all paths in $\mathcal{A}(S)$ going from c to a vertex in $\mathcal{P}(a, b)$). Otherwise, let us move along c starting at its initial point γ, and let c_i be the first intersection point with u

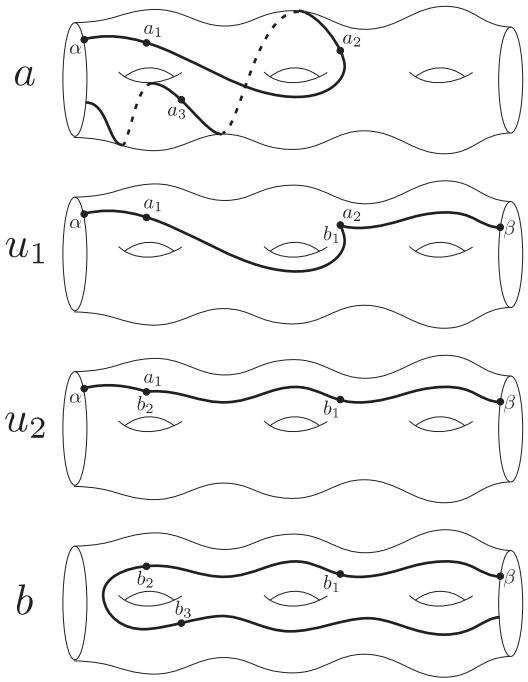

Figure 3.7: The unicorn path $\mathcal{P}(a, b)$ for the curves a and b in Figure 3.4. Notice that consecutive arcs can be realized disjointly on the surface.

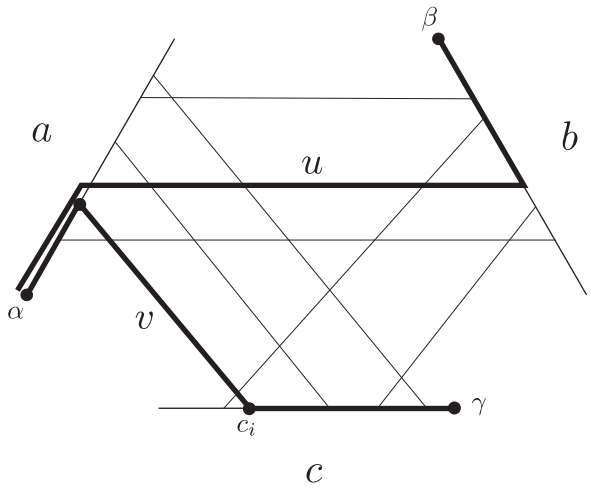

Figure 3.8: A diagram representing a triangle of unicorn paths with vertices a, b and c. A unicorn arc $u \in P(a, b)$ and a unicorn arc $v \in \mathcal{P}(b, c) \cup \mathcal{P}(a, c)$ which is adjacent to u as a vertex in $\mathcal{A}(S)$.

encountered along the way. Up to interchanging a and b, we can assume that c_i belongs to a, as in Figure 3.8.

Notice that c_i corresponds to a unicorn arc $v \in \mathcal{P}(a, c)$, otherwise there would be a segment connecting a and c and crossing the segment corresponding to c_i with higher slope, but then c_i would not be the first encountered intersection point of c and u.

Let us see that $d(u, v) = 1$. By construction, the initial part of v (which is the initial part of c, up to c_i) does not touch u. The remaining part of v coincides with the subarc of a going from c_i to α, hence it can be isotoped to be parallel to a, and this realizes u and v disjointly. □

The above result shows that triangles made of unicorn paths (where the vertices' orientations coincide) are 2-thin. Indeed, let $\mathcal{P}(a,b) = (a = u_0, u_1, \ldots, u_{r+1} = b)$. Every vertex is at distance ≤ 1 to either $\mathcal{P}(a,c)$ or $\mathcal{P}(b,c)$, or both. As $a \in \mathcal{P}(a,c)$ and $b \in \mathcal{P}(b,c)$, then either some vertex u_i is at distance ≤ 1 from the other two sides, or there are two consecutive vertices u_i, u_{i+1} at distances ≤ 1 from $\mathcal{P}(a,c)$ and $\mathcal{P}(b,c)$, respectively. It follows that the 2-ball centered at u_i touches the three sides.

The situation is even better: Unicorn triangles are 1-thin. This is not necessary to show that the arc graph is hyperbolic, but it allows to reduce the hyperbolicity constant from 8 to 7, as we will see in Theorem 3.17. We will give the following result, which implies immediately that unicorn triangles are 1-thin:

Lemma 3.14. [22] *Let a, b, c be three vertices of $\mathcal{A}(S)$. Realize them in minimal position, choose an orientation for each one, and let $\mathcal{P}(a,b)$, $\mathcal{P}(b,c)$ and $\mathcal{P}(c,a)$ be the corresponding unicorn paths. Then there are vertices $u \in \mathcal{P}(a,b)$, $v \in \mathcal{P}(b,c)$ and $w \in \mathcal{P}(c,a)$ which are at distance ≤ 1 from each other.*

Proof. If a and b are adjacent, one can take $u = a = w$ and $v = b$. So we can assume that $\mathcal{P}(a,b)$ has length bigger than 1. Also, if some vertex $z \in P(a,b)$ is adjacent to c, then we can take $u = z$ and $v = w = c$. Hence we also assume that no vertex of $\mathcal{P}(a,b)$ is adjacent to c, so each vertex in $\mathcal{P}(a,b)$ represents a curve which intersects the curve c.

Let $\mathcal{P}(a,b) = (a = u_0, u_1, \ldots, u_k = b)$. For each u_i we will look at its intersection points with c. Going along c from its starting point γ, the first intersection point with u_i that one encounters will belong to either a or b (and it will correspond to a unicorn arc). If $i = 0$, the first intersection point belongs to a, and if $i = k$ it belongs to b. Hence, there is some j such that the first intersection point of c with u_j (say p_j) belongs to a, and the first intersection point of c with u_{j+1} (say p_{j+1}) belongs to b. Up to interchanging the roles of a and b, we can assume that p_{j+1} comes before p_j (i.e. p_{i+1} is closer to δ than p_i along c). We will see that u_j is the vertex u promised in the statement.

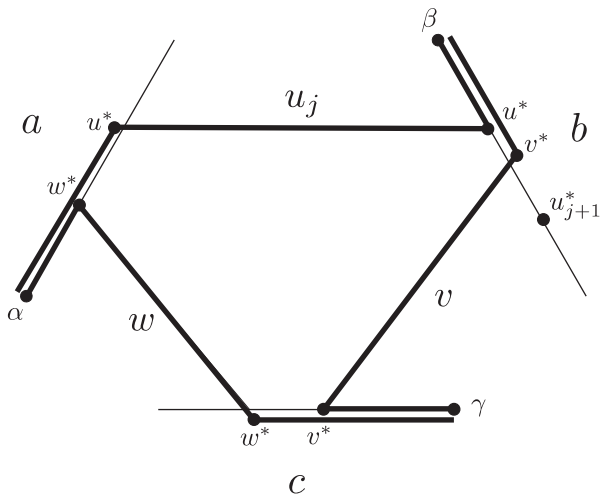

Figure 3.9: A triangle whose sides are unicorn paths is 1-thin.

If we move along c starting at γ, the first intersection with u_{j+1} that we encounter is p_{j+1}, but we can encounter several other intersections points with $u_{j+1} \cap b$ before arriving to p_j. Let v^* be the one which is closest to β along b. By construction, it corresponds to a unicorn arc v. As v^* is closer to γ in c than p_j, it does not belong to u_j, so the segment corresponding to v in the diagram we draw to represent the situation (see Figure 3.9) does not cross the

segment corresponding to u_j. Hence the situation is exactly as shown in Figure 3.9, where we have denoted by $u^* = u_j^*$ the intersection point of a and b corresponding to u_j, by u_{j+1}^* the intersection point of a and b corresponding to u_{j+1}, and by w^* the point p_j, corresponding to a unicorn arc w. We will show that $u = u_j$, v and w are pairwise adjacent.

See Figure 3.9. We can slightly deform the coincident parts of u, v and w so that they become parallel. The three obtained curves are disjoint by the following reason: Any intersection point belonging to a and b would yield a unicorn arc between u_j and u_{j+1}, which is impossible. Any intersection point belonging to a and c would be an intersection point of u_j and c closer to γ in c than w^*, which is also impossible. And any intersection point belonging to b and c would contradict the choice of v^*, which is as close to β as possible in b. Therefore, the three curves u, v and w are pairwise disjoint, as we wanted to show. □

Lemma 3.13 allows to show the following very important technical result. If we have a sequence of m consecutive unicorn paths in $\mathcal{A}(S)$, then a single unicorn path joining the initial and final vertices is at bounded distance (at most $\lceil \log_2 m \rceil$) from the original sequence of unicorn paths. Here is the explicit statement:

Corollary 3.15. *Let $x_0, \dots, x_m$ be a sequence of (not necessarily adjacent) vertices in $\mathcal{A}(S)$, with $m \leq 2^k$. For every $u \in \mathcal{P}(x_0, x_m)$ there is some i and some $v \in \mathcal{P}(x_i, x_{i+1})$ with $d(u, v) \leq k$.*

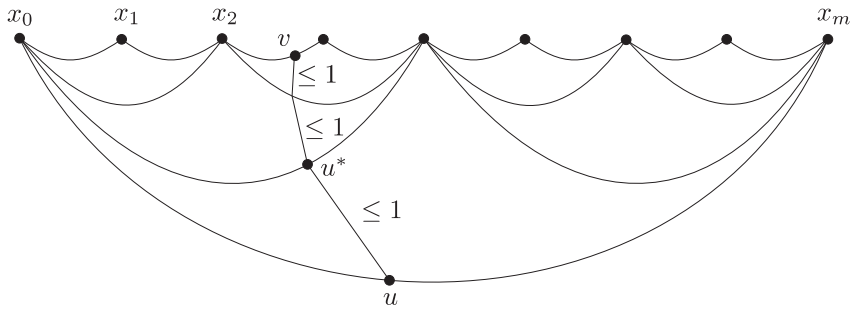

Figure 3.10: Every vertex in a unicorn path is at distance at most k from some vertex in a sequence of m unicorn paths joining the same endpoints, provided $m \leq 2^k$.

Proof. See Figure 3.10, where the represented arcs correspond to unicorn paths in $\mathcal{A}(S)$. Let u be some vertex in $\mathcal{P}(x_0, x_m)$. We can consider the triangle formed by $\mathcal{P}(x_0, x_m)$, $\mathcal{P}(x_0, x_{\lfloor \frac{m}{2} \rfloor})$, $\mathcal{P}(x_{\lfloor \frac{m}{2} \rfloor}, x_m)$. By Lemma 3.13 there is some vertex u^* in either $\mathcal{P}(x_0, x_{\lfloor \frac{m}{2} \rfloor})$ or $\mathcal{P}(x_{\lfloor \frac{m}{2} \rfloor}, x_m)$ which is at distance ≤ 1 from u. The number of unicorn paths of the form $\mathcal{P}(u_i, u_{i+1})$ connecting the endpoints of either $P(x_0, x_{\lfloor \frac{m}{2} \rfloor})$ or $\mathcal{P}(x_{\lfloor \frac{m}{2} \rfloor}, x_m)$ is at most 2^{k-1}. Hence, by induction, u^* is at distance at most $k-1$ from some $\mathcal{P}(u_i, u_{i+1})$. This finishes the proof. □

In [22] it is shown that a subpath of a unicorn path is again a unicorn path, except when the subpath has length two and its two endpoints are adjacent. This can sound odd: One could think, looking at our diagrams, that a subpath of a unicorn path is always a unicorn path. But the key point is to realize that, if u_i and u_j belong to the same unicorn path, they are not necessarily in minimal position. An example is shown in Figure 3.11: The unicorn path $\mathcal{P}(a, b)$ is $(a = u_0, u_1, u_2, u_3 = b)$, but the curves u_0 and u_2 can be realized disjointly (by sliding the endpoint α).

Now one can show, using an argument by Hamenstädt [17], that unicorn paths are not far from geodesics. Notice that in the proof of the following result we are only using that unicorn

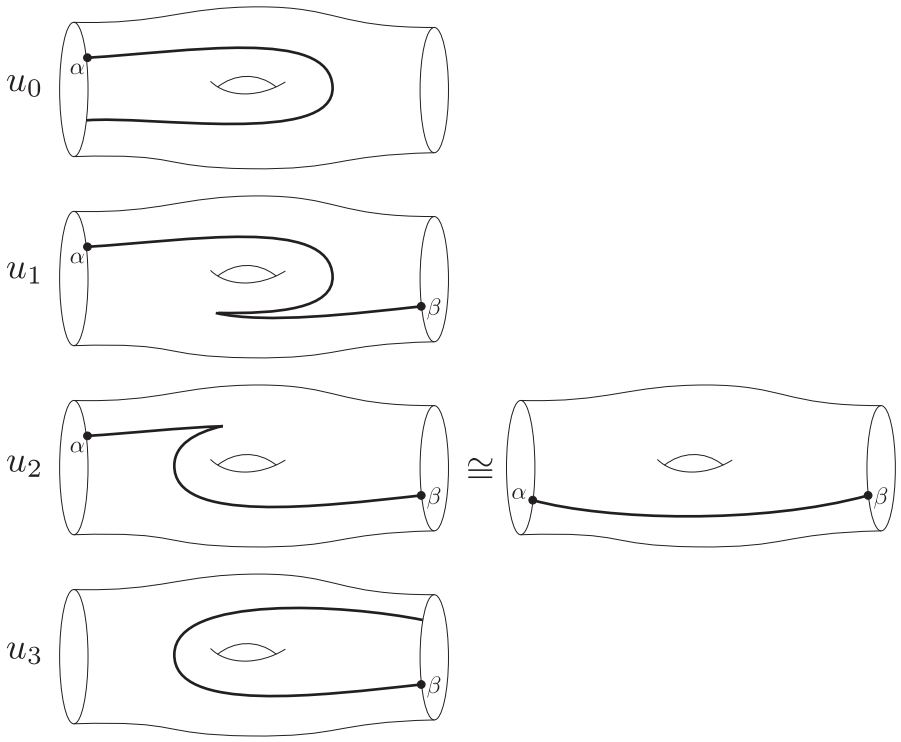

Figure 3.11: A subpath of a unicorn path may not be a unicorn path.

paths form thin triangles and that subpaths of unicorn paths are close to unicorn paths. Hence, the arguments can be applied in a much more general situation, in which a family of paths in a metric space satisfies the above two conditions, to show that such paths are not far from geodesics. We will discuss this later.

Proposition 3.16. [22] *Let a, b be vertices of $\mathcal{A}(S)$. Then every vertex in $\mathcal{P}(a,b)$ is at distance ≤ 6 from any geodesic $\mathcal{G}$ from a to b.*

Proof. See Figure 3.12. Let u be a vertex of $\mathcal{P}(a,b)$ which is at maximal distance k from the geodesic $\mathcal{G} = \mathcal{G}(a,b)$. We need to show that $k \leq 6$. Let us consider the ball centered at u with radius $2k$. The maximal subpath of $\mathcal{P}(a,b)$ contained in this ball and containing u is a unicorn path $\mathcal{P}(a',b')$ (as we can assume $k > 1$, otherwise the result would be trivially true).

Now consider the path $(x_0,\dots,x_m)$, where $\mathcal{G}_1 = (a' = x_0,\dots,x_p = a'')$ is a geodesic realizing the minimal distance from a' to $\mathcal{G}$, $\mathcal{G}_3 = (b'' = x_q,\dots,x_m = b')$ is a geodesic realizing the minimal distance from $\mathcal{G}$ to b', and $\mathcal{G}_2 = (a'' = x_p,\dots,x_q = b'')$ is the subpath of $\mathcal{G}$ joining a'' and b'' (hence a geodesic from a'' to b'').

We claim that $d(u,x_i) \geq k$ for $i = 0,\dots,m$. Indeed, as k is the distance from u to $\mathcal{G}$, $d(u,x_i) \geq k$ for all $x_i \in \mathcal{G}_2$. Also, by maximality of k, the geodesics $\mathcal{G}_1$ and $\mathcal{G}_3$ have length at most k. If there is some $x_i \neq a$ in $\mathcal{G}_1$ (that is, if $a' \neq a$) then $d(u,a') = 2k$, hence $d(u,x_i) \geq d(u,a') - d(a',x_i) \geq 2k - k = k$. And the same happens for every $x_i \in \mathcal{G}_3$, so the claim holds.

Finally, notice that $\mathcal{G}_2$ is a geodesic of length at most $6k$, as $d(a'',b'') \leq d(a'',a') + d(a',u) + d(u,b') + d(b',b'') \leq k + 2k + 2k + k$. Together with the fact that the geodesics $\mathcal{G}_1$ and $\mathcal{G}_3$ have length at most k, this implies that the path $(x_0,\dots,x_m)$ has length at most $8k$. That is, $m \leq 8k$.

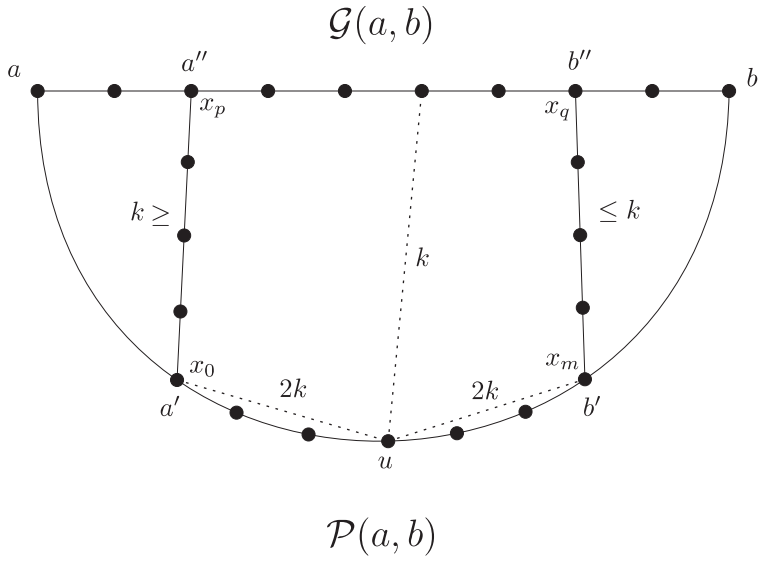

Figure 3.12: Unicorn paths are not far from geodesics.

We then have a unicorn path $\mathcal{P}(a', b')$, and another path $(x_0, \dots, x_m)$ going from a' to b' which has length at most $8k$. As x_i is adjacent to x_{i+1} for each i, we can consider $(x_0, \dots, x_m)$ as the concatenation of m unicorn paths $\mathcal{P}(x_i, x_{i+1})$. Then, by Corollary 3.15, the vertex $u \in \mathcal{P}(a', b')$ is at distance at most $\lceil \log_2 m \rceil$ from some x_i. That is, $d(u, x_i) \leq \lceil \log_2 8k \rceil$. As we have shown that $k \leq d(u, x_i)$, we finally obtain $k \leq \lceil \log_2 8k \rceil$, which implies $k \leq 6$. □

And finally, we can prove what we had promised:

Theorem 3.17. [22] *The arc complex $\mathcal{A}(S)$ is 7-hyperbolic.*

Proof. Consider a geodesic triangle in $\mathcal{A}(S)$ with vertices a, b and c. Put the three arcs in minimal position, orient them in some way, and consider the corresponding unicorn paths $\mathcal{P}(a, b)$, $\mathcal{P}(b, c)$ and $\mathcal{P}(a, c)$. See Figure 3.13. By Lemma 3.14, the unicorn triangle is 1-thin, so there is a ball of radius 1 touching the three unicorn paths. By Proposition 3.16, each intersection point of the ball with one unicorn side, is at distance ≤ 6 from the corresponding geodesic side. Therefore, there is a ball of radius 7 touching the three geodesic sides, as we wanted to show. □

We recall that, using Lemma 3.13 instead of Lemma 3.14, we would have obtained that $\mathcal{A}(S)$ is 8-hyperbolic instead of 7-hyperbolic, but we would have shown uniform hyperbolicity anyway.

Finally, we will just give a rough idea of how to pass from the 7-hyperbolicity of the arc graph $\mathcal{A}(S)$ to the 17-hyperbolicity of the curve graph $\mathcal{C}(S)$ (see [22] for details). Given a geodesic triangle T in $\mathcal{C}(S)$, the idea is to look at T in $\mathcal{AC}(S)$. One can retract $\mathcal{AC}(S)$ to $\mathcal{C}(S)$, sending each arc to a boundary component of a regular neighborhood of its union with $\partial(S)$ (if the two endpoints of the arc lie on the same component of $\partial(S)$, one of the two obtained curves can be used, otherwise the arc would be non-essential or S would be sporadic; if the endpoints lie in distinct components of $\partial(S)$, the obtained curve is essential, otherwise S would be sporadic). This retract is 2-Lipschitz. Hence, the sides of T are 2-quasigeodesics in $\mathcal{AC}(S)$. Now, one can take three arcs (vertices of $\mathcal{A}(S)$) which are adjacent to the vertices of T, and consider the triangle made by *unicorn paths*, joining those three arcs. One can show that the

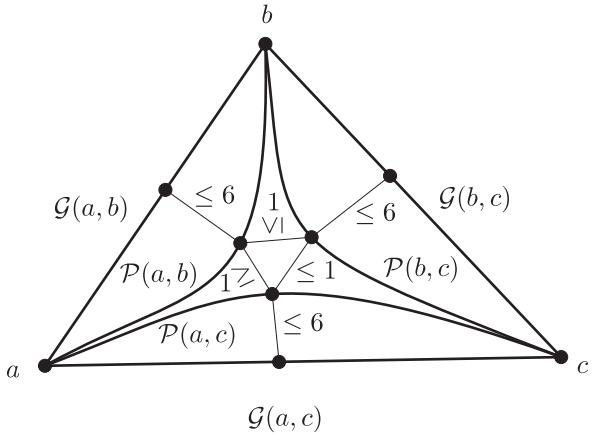

Figure 3.13: A triangle whose sides are unicorn paths is 1-thin.

distance from any vertex in such unicorn path to the corresponding side of T is at most 8 (by an argument similar to the one in Proposition 3.16). This implies that there is a ball of radius 9 which touches the three sides of T in $\mathcal{AC}(S)$, and by the above retraction it follows that $\mathcal{C}(S)$ is 17-hyperbolic. Hence we have:

Theorem 3.18. [22] *The curve graph $\mathcal{C}(S)$ is 17-hyperbolic.*

The main goal of this section was to give some tools that could allow the reader to prove that some metric space is hyperbolic. We point out that the hyperbolicity of the curve graph was obtained *just* by finding some family of paths between each pair of vertices (in this case the unicorn paths) satisfying the following conditions:

1. Paths in the family form thin triangles (Lemma 3.13).
2. The family is closed under the operation of taking subpaths.

Actually, unicorn paths are not closed by subpaths, but almost: subpaths of length greater than 2 are unicorn. But we also used in our proofs that subpaths of length 1 are unicorn. Hence, some weaker conditions one could impose to a family of paths, in order to be able to apply the above arguments, are the following:

1. Paths in the family form thin triangles.
2. The family is closed under the operation of taking subpaths of length 1.
3. The family is closed under the operation of taking subpaths of length greater than k, for some $k > 0$.

These properties can be used to show that the paths in the family are not far from geodesics (Proposition 3.16), and then to show the hyperbolicity of the space (Theorem 3.17).

References

[1] T. Aougab, *Uniform hyperbolicity of the graphs of curves*, Geom. Topol. 17 (2013), no. 5, 2855–2875.
[2] E. Artin, *Theorie der Zöpfe*, Abh. Math. Sem. Univ. Hamburg 4 (1925), no. 1, 47-72.

[3] E. Artin, *Theory of braids*, Ann. of Math. 48 (1947), no. 2, 101-126.
[4] J. S. Birman, *Braids, links, and mapping class groups*, Annals of Mathematics Studies, No. 82. Princeton University Press, Princeton, N.J. 1974.
[5] N. Bourbaki, *Éléments de mathématique: groupes et algèbres de Lie. Chapitres 4, 5 et 6*, Actualités Scientifiques et Industrielles, No. 1337, Hermann, Paris 1968.
[6] B. H. Bowditch, *Uniform hyperbolicity of the curve graphs*, Pacific J. Math. 269 (2014), no. 2, 269–280.
[7] M. R. Bridson, A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 319. Springer-Verlag, Berlin, 1999.
[8] M. Calvez, *Fast nielsen-thurston classification of braids*, Algebr. Geom. Topol. 14 (2014) 1745-1758.
[9] A. Casson, S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts, 9. Cambridge University Press, Cambridge, 1988.
[10] H. S. M. Coxeter, *Discrete groups generated by reflections*, Ann. of Math. (2) 35 (1934), no. 3, 588–621.
[11] A. Clay, K. Rafi, S. Schleimer, *Uniform hyperbolicity of the curve graph via surgery sequences*, Algebr. Geom. Topol. 14 (2014), no. 6, 3325–3344.
[12] P. Dehornoy, F. Digne, E. Godelle, D. Krammer, J. Michel, *Foundations of Garside Theory*, EMS Tracts in Mathematics, volume 22, European Mathematical Society, 2015.
[13] P. Dehornoy, L. Paris, *Gaussian groups and Garside groups, two generalisations of Artin groups*, Proc. London Math. Soc. (3) 79 (1999), no. 3, 569–604.
[14] A. Fathi, F. Laudenbach, V. Poénaru, *Travaux de Thurston sur les surfaces*, Astérisque No. 66-67 (1991). SMF, Paris, 1991.
[15] J. González-Meneses, B. Wiest, *On the structure of the centralizer of a braid*, Ann. Sci. Norm. Sup. (4) 37 (2004), no. 5, 729–757.
[16] J. González-Meneses, *The nth root of a braid is unique up to conjugacy*, Algebr. Geom. Topol. 3 (2003), 1103–1118.
[17] Ursula Hamenstädt, *Geometry of the complex of curves and of Teichmüller space*, Handbook of Teichmüller theory. Vol. I, IRMA Lect. Math. Theor. Phys., vol. 11, Eur. Math. Soc., Zürich, 2007, 447–467.
[18] J. L. Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math 84 (1986), no. 1, 157-176.
[19] W. J. Harvey, *Boundary structure of the modular group*, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), pp. 245–251, Ann. of Math. Stud., 97, Princeton Univ. Press, Princeton, N.J., 1981.
[20] A. Hatcher, *On triangulation of surfaces*, Topology Appl. 40 (1991), no. 2, 189-194.
[21] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002.
[22] S. Hensel, P. Przytycki, R. C. Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, J. Eur. Math. Soc. (JEMS) 17 (2015), no. 4, 755–762.
[23] N. V. Ivanov, *Subgroups of Teichmüller modular groups*, Translations of Mathematical Monographs, 115. American Mathematical Society, Providence, RI, 1992.
[24] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. 138 (1999), no. 1, 103–149.
[25] J. D. McCarthy, *Normalizers and centralizers of pseudo-Anosov mapping classes*, Unpublished. Available at: `http://users.math.msu.edu/users/mccarthy/publications/selected.papers.html`
[26] Y. Minsky, *A geometric approach to the complex of curves*, Proceedings of the 37th Taniguchi Symposium on Topology and Teichmüller Spaces (S. Kojima et. al., ed.), World Scientific, 1996, pp. 149–158.
[27] L. Paris, *Braid groups and Artin groups*, Handbook of Teichmüller theory. Vol. II, 389–451, IRMA Lect. Math. Theor. Phys., 13, Eur. Math. Soc., Zürich, 2009.
[28] K. Rafi, S. Schleimer, *Curve complexes are rigid*, Duke Math. J. 158 (2011), no. 2, 225-246.
[29] W. P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) 19 (1988), no. 2, 417-431.

Departamento de Álgebra, Facultad de Matemáticas, Instituto de Matemáticas (IMUS), Universidad de Sevilla, Av. Reina Mercedes, s/n, 41012 Sevilla, Spain • meneses@us.es