Geodesic flow, left-handedness and templates

PIERRE DEHORNOY

We establish that for every hyperbolic orbifold of type $(2, q, \infty)$ and for every orbifold of type $(2, 3, 4g + 2)$, the geodesic flow on the unit tangent bundle is left handed. This implies that the link formed by every collection of periodic orbits (i) bounds a Birkhoff section for the geodesic flow, and (ii) is a fibered link. We also prove similar results for the torus with any flat metric. We also observe that the natural extension of the conjecture to arbitrary hyperbolic surfaces (with non-trivial homology) is false.

1 Introduction

In this paper, we investigate the dynamical properties of certain particular 3–dimensional flows, namely the geodesic flows attached to surfaces and 2–dimensional orbifolds. If Σ is a Riemannian surface or, more generally, a Riemannian 2–dimensional orbifold, that is, a space locally modeled on quotients of surfaces under the action of discrete rotation groups, the unit tangent bundle $T^1\Sigma$ is a 3–manifold, and the geodesics of Σ induce a natural complete flow in $T^1\Sigma$. This flow is called the *geodesic flow* of $T^1\Sigma$, hereafter denoted by Φ_Σ. What we do here is to specifically study the way the periodic orbits of Φ_Σ may wrap one around the other.

In every 3–dimensional manifold M, the linking number of two disjoint links can be defined in a non-ambiguous way whenever the links are null-homologous, that is, have a trivial image in $H_1(M; \mathbb{Q})$ (Kaiser [11]). When the latter group is trivial, that is, when M is a rational homology sphere, the linking number is always defined, and it yields a topological invariants of links.

If Σ is a 2–dimensional orbifold, every geodesic on Σ can be lifted to $T^1\Sigma$ in two ways, yielding a pair of orbits of Φ_Σ. It follows from Birkhoff's results [3] that the linking number of any two such pairs of orbits is the opposite of the number of intersections of the geodesics, hence is nonpositive. This implies that, in a geodesic flow, there are always many pairs of orbits with a negative linking number. By contrast, there is no simple construction necessarily leading to collections of orbits with a positive linking number, and it makes sense to raise the following question:

Question 1-1 Assume that Σ is a Riemannian 2–dimensional orbifold. Let γ, γ' be two null-homologous collections of periodic orbits of Φ_Σ. Does $\mathrm{Lk}(\gamma, \gamma') < 0$ necessarily hold?

There are two cases when the answer to Question 1-1 is known to be positive, namely when Σ is a sphere $\mathbb{S}^2$ with a round metric and when Σ is the modular surface $\mathbb{H}^2/\mathrm{PSL}_2(\mathbb{Z})$ (Ghys [10]). In the latter article, Étienne Ghys actually proves stronger results involving the natural extension of the linking number to arbitrary measures. Namely, he defines a complete flow Φ in a homology 3–sphere M to be *left handed* if the linking number of every pair of Φ–invariant measures whose supports are disjoint is always negative, and proves that the above two flows are left handed. It is then natural to raise this question:

Question 1-2 (Ghys) Assume that Σ is a Riemannian 2–dimensional orbifold satisfying $H_1(T^1\Sigma, \mathbb{Q}) = 0$. Is the geodesic flow Φ_Σ on $T^1\Sigma$ necessarily left handed?

By definition, a positive answer to Question 1-2 implies a positive answer to Question 1-1. As we shall explain, the converse implication, that is, the fact that the negativity of the linking number for pairs of periodic orbits implies the negativity of the linking number for arbitrary invariant measures, is true whenever the flow has sufficiently many periodic orbits, in particular when the flow is of Anosov type.

The aim of this paper is to provide positive answers to Questions 1-1 and 1-2 in new cases, namely when Σ is a hyperbolic orbifold of type $(2, q, \infty)$ with $q \geq 3$ and when Σ is a hyperbolic orbifold of type $(2, 3, 4g + 2)$ with $g \geq 2$.

Theorem A Assume that Σ is (a) *either an orbifold of type $(2, q, \infty)$ with $q \geq 3$, equipped with a negatively curved metric, or* (b) *an orbifold of type $(2, 3, 4g + 2)$ with $g \geq 2$, equipped with a negatively curved metric. Then:*

(i)   *Any two null-homologous collections of periodic orbits of Φ_Σ have a negative linking number.*

(ii) *The geodesic flow of $T^1\Sigma$ is left handed.*

In the case of a good orbifold with zero curvature, that is, a quotient of a torus with a flat metric, the unit tangent bundle always has non-trivial homology. Nevertheless it makes sense to address Question 1-1. In this case as well, the answer is (almost) always positive.

Theorem B Assume that Σ is a quotient of the torus $\mathbb{T}^2$ equipped with a flat metric. Then any two collections γ, γ' of orbits of Φ_Σ whose projections on Σ intersect have a negative linking number.

On the other hand, we give two examples showing that, when Σ is not a homology sphere or its curvature has a non-constant sign, Question 1-1 has a negative answer.

Proposition 1-3 (i) *Let Σ be a hyperbolic surface. Then there exist two null-homologous collections γ, γ' of periodic orbits of Φ_Σ with $\text{Lk}(\gamma, \gamma') > 0$.*

(ii) *Let Σ be a sphere with two non-intersecting simple geodesics. Then there exist two null-homologous collections γ, γ' of periodic orbits of Φ_Σ with $\text{Lk}(\gamma, \gamma') > 0$. The geodesic flow Φ_Σ is not left handed.*

When Questions 1-1 and 1-2 have positive answers, an important consequence is the existence of many Birkhoff sections. A Birkhoff section for a non-singular flow on a 3–manifold is a compact surface whose boundary is the union of finitely many periodic orbits of the flow, whose interior is transverse to the flow and intersects every orbit infinitely many times. The existence of a Birkhoff section for a flow is very useful as, in this case, studying the dynamics of the flow essentially reduces to studying the first return map on the section. Therefore, it is natural to wonder whether a flow admits Birkhoff sections. Now, as explained by Ghys [10], the left-handedness of a flow implies the existence, for every finite collection of periodic orbits, of a Birkhoff section bounded by this collection. Thus our current results imply:

Corollary 1-4 *If Σ is one of the orbifolds mentioned in Theorem A, every finite null-homologous collection of periodic orbits of Φ_Σ bounds a Birkhoff section.*

Next, it is known (Dehornoy [5]) that every link that is the boundary of a Birkhoff section for a flow is fibered. Therefore, a direct consequence of Corollary 1-4 is:

Corollary 1-5 *If Σ is one of the orbifolds mentioned in Theorem A, every link in $T^1\Sigma$ formed by a null-homologous collection of periodic orbits of the flow Φ_Σ is fibered.*

Similar statements hold in the case of the flat torus (see Theorem 3-12), with, in addition, an explicit simple formula for the genus of the involved Birkhoff sections.

Let us give a few hints about proofs. The case of the torus $\mathbb{T}^2$ is the most simple one. It can be solved by elementary means, and it appears as a sort of warm-up. The key point is to encode every null-homologous collection γ of periodic orbits of $\Phi_{\mathbb{T}^2}$ into some convex polygon Pol_γ in the affine plane $\mathbb{R}^2$ with integral vertices. Using Pol_γ and Van Horn-Morris' helix boxes [17], we classify Birkhoff sections up to isotopy and derive their existence and the explicit formulas for the genus and the linking number of two null-homologous collections of periodic orbits (Theorem 3-12).

Once these formulas are available, the negativity of the linking numbers easily follows (Corollary 3-13).

For Theorem A, the proofs rely on a common principle but require specific ingredients depending on the orbifold. Our strategy decomposes in two steps. We first develop a general method for investigating the geodesic flow on a hyperbolic orbifold. A *multitemplate* is a geometric 2–dimensional branched surface carrying a flow. This notion generalizes Birman and Williams' notion of template [4], that has been introduced for studying hyperbolic flows. Here we prove that, given an orbifold Σ, for every tessellation $\mathcal{T}$ of the hyperbolic plane adapted to Σ, there exists a multitemplate $\mathcal{B}_{\mathcal{T}}$ embedded in $T^1\Sigma$ such that the set of periodic orbits of Φ_Σ is isotopic to a subset of the periodic orbits of $\mathcal{B}_{\mathcal{T}}$ (Theorem 4-9). Moreover, if the orbifold Σ has at least one cusp, we can choose the tessellation $\mathcal{T}$ so that the set of periodic orbits of the geodesic flow is isotopic to the whole set of periodic orbits of the template $\mathcal{B}_{\mathcal{T}}$. This result provides a combinatorial description of the isotopy classes of the periodic orbits of Φ_Σ in terms of some finite data specifying the orbifold. Note that the construction of the multi-template follows the strategy proposed by Birman and Williams for hyperbolic flows [4]. In one sentence: we choose a Markov partition for the flow and contract the stable direction.

To complete the proof in the case when Σ is an orbifold of type $(2, q, \infty)$ with $q \geq 3$, we start from the fact that $T^1\Sigma$ is diffeomorphic to the complement of a certain knot K_∞ in some lens space, and we choose a particular compactification. Then, choosing an adapted tessellation of the hyperbolic plane and using the template provided by Theorem 4-9, we estimate the linking number of an arbitrary pair of collections of periodic orbits and see that it is always negative. Along the way, we also compute the linking number of a geodesic with the knot K_∞ (Proposition 5-7), a function of interest in number theory.

To complete the proof in the case of the orbifolds $\Sigma_{2,3,4g+2}$, the most delicate case, we use a covering of $\Sigma_{2,3,4g+2}$ by some explicit genus-g surface Σ_g. Then we use the template of Theorem 4-9 to bound the linking number of two collections of periodic orbits of Φ_{Σ_g} in terms of some associated combinatorial data. More precisely, we start from a tessellation of $\mathbb{H}^2$ by $(4g + 2)$–gons. For every periodic geodesic γ in Σ_g and for every pair of edges (e_i, e_j) in a tile of the tessellation, we denote by $b_{i,j}(\gamma)$ the number of times the projection of γ goes from e_i to e_j. Then, for every pair of geodesics γ, γ', we show that the linking number $\mathrm{Lk}(\gamma, \gamma')$ is bounded above by a certain bilinear form S_{4g+2} involving the coefficients $b_{i,j}(\gamma)$ and $b_{i,j}(\gamma')$. The form S_{4g+2} is not negative on the whole cone of vectors with positive coordinates (a manifestation of Proposition 1-3). What we do here is to show that the form S_{4g+2} is negative on the subcone of vectors that come from liftings of geodesics of $\Sigma_{2,3,4g+2}$,

which is enough to deduce the main result (Proposition 6-14). The reason why the proof works in this case, unlike for general families of geodesics on Σ_g, is that a family of geodesics on $\Sigma_{2,3,4g+2}$ lifts to a family on Σ_g that admits many symmetries, and that these symmetries force the associated coefficients $b_{i,j}$ to live in a small subcone where the bilinear form Q_{4g+2} is negative.

It should be noted that, in the case of orbifolds of type $(2, q, \infty)$, a result similar to Theorem 4-9 has been established by Tali Pinsky [14] in a previous work. Precisely, when Σ is the orbifold $\mathbb{H}/\mathrm{PSL}_2(\mathbb{Z})$, Ghys [9] proved that the periodic orbits of the geodesic flow can be distorted on a template that coincides with the geometric Lorenz template, so that periodic orbits are Lorenz knots [4]. His construction corresponds to ours when Σ is the orbifold $\mathbb{H}/\mathrm{PSL}_2(\mathbb{Z})$ (which is of type $(2, 3, \infty)$) and $\mathcal{T}$ the tessellation of $\mathbb{H}^2$ by ideal triangles. Later, Pinsky [14] generalized Ghys' construction to orbifolds of type $(2, q, \infty)$. Her construction can be recovered in our setting using a tiling of $\mathbb{H}^2$ by ideal regular q–gons. The presentations of Ghys and Pinsky differ from ours in the sense that they construct a template by opening the cusp in the associated orbifold, thus distorting the underlying manifold $T^1\Sigma$, and then contracting the stable direction of the geodesic flow. The notion of discretization of geodesics (Definition 4-3) allows us to construct multitemplates even when the considered orbifold has no cusp.

The plan of the article is as follows. First, we recall some basic definitions — linking number, orbifold, unit tangent bundle, geodesic flow — and prove two general lemmas on left handed flows in Section 2. We then treat the case of the torus in Section 3. Next, we turn to hyperbolic orbifolds and construct a template for the geodesic flow on every orbifold in Section 4, where we prove Theorem 4-9. We then complete the case of orbifolds of type $(2, q, \infty)$ in Section 5. We investigate the geodesic flows on surfaces of genus g and complete the case of the orbifolds of type $(2, 3, 4g + 2)$ in Section 6. Finally, we construct the counter-examples of Proposition 1-3 and discuss further questions in Section 7.

2 Definition and motivation

Here we set the general context. We recall the needed definitions, and establish some preliminary results.

2.1 Orbifolds and their unit tangent bundles

A *Riemannian, orientable,* 2*–dimensional orbifold* Σ is a topological surface locally modeled on a Riemannian surface modulo actions by finite subgroups of rotations [16]. More precisely Σ consists of a topological surface X_Σ with an atlas of covering charts $\phi_i \colon V_i \to U_i$, where $\{U_i\}$ is a collection of open sets of X_Σ closed under finite intersections, $\{V_i\}$ is a collection of open sets of a Riemannian surface, such that to each V_i is associated a finite group Γ_i of rotations of V_i identifying U_i with V_i / Γ_i, and such that every change of charts $\phi_i^{-1} \circ \phi_j$, when defined, consist of isometries.

In the sequel we will restrict ourselves to orbifolds that are also *good*, meaning that the whole underlying space X_Σ admits a finite degree covering by a surface (which needs not to be compact), say Σ_0. In this case, the orbifold Σ can be identified with the quotient Σ_0 / Γ_0 for some discrete subgroup Γ_0 of $\mathrm{Isom}^+(\Sigma_0)$. The universal cover of Σ_0 is defined as the universal cover of Σ, hereafter denoted by $\widetilde{\Sigma}$. One can then identify Σ with the quotient $\widetilde{\Sigma} / \Gamma$ for some discrete subgroup Γ of $\mathrm{Isom}^+(\widetilde{\Sigma})$. The latter subgroup is called the *fundamental group* of Σ. If Σ has a constant curvature, then $\widetilde{\Sigma}$ is either the sphere $\mathbb{S}^2$, the Euclidean plane $\mathbb{R}^2$ or the hyperbolic plane $\mathbb{H}^2$. Accordingly, the orbifold Σ is said to be *spherical*, *Euclidean* or *hyperbolic*.

By definition, the orbifold structure transports the metrics, so that each point x of a good 2–orbifold admits a neighborhood of the form V_x / Γ_x, where V_x is an open disc in $\widetilde{\Sigma}$ and Γ_x a finite group of rotations. The order of Γ_x is called the *index* of x. A point with index 1 is called *regular*, otherwise it is called *singular*. It is important to note that singular points are isolated.

We now turn to the unit tangent bundle of an orbifold. Let Σ be a good 2–orbifold with fundamental group Γ. Then the action of Γ on $\widetilde{\Sigma}$ by isometries is properly discontinuous. The *unit tangent bundle* $T^1\Sigma$ of Σ is defined to be the quotient of the total space $T^1\widetilde{\Sigma}$ of the unit tangent bundle of $\widetilde{\Sigma}$ by the action of Γ on the tangent space of $\widetilde{\Sigma}$, ie, $T^1\Sigma = (T^1\widetilde{\Sigma}) / \Gamma$.

Let us illustrate this definition with two examples which are important for the sequel. Assume that $\mathbb{D}^2$ is an open disc. Its unit tangent bundle $T^1\mathbb{D}^2$ then consists of the set of unit vectors tangent to $\mathbb{D}^2$. The unit tangent vectors based at a given point form a circle, so that the manifold $T^1\mathbb{D}^2$ is a solid torus.

Consider the action of $\mathbb{Z}/p\mathbb{Z}$ on $\mathbb{D}^2$ by rotations of angles that are multiple of $2\pi/p$. The action is not free because the center of $\mathbb{D}^2$ is fixed. It is the only point with non-trivial stabilizer. The quotient $\mathbb{D}^2/(\mathbb{Z}/p\mathbb{Z})$ is then an orbifold. Denote it by $\mathbb{D}^2_p$. Since the action of $\mathbb{Z}/p\mathbb{Z}$ is by isometries, it can be extended to the unit tangent bundle $T^1\mathbb{D}^2$. Given a point with polar coordinates (x, θ) on $\mathbb{D}^2$, and a unit tangent

vector making an angle ϕ with the horizontal direction, an element $\bar{k}$ of $\mathbb{Z}/p\mathbb{Z}$ then acts by $\bar{k} \cdot (r, \theta, \phi) = (r, \theta + 2k\pi/p, \phi + 2k\pi/p)$. The action on $T^1\mathbb{D}^2$ is therefore free, and the quotient $T^1\mathbb{D}^2/(\mathbb{Z}/p\mathbb{Z})$ is a manifold. It is the unit tangent bundle $T^1\mathbb{D}_p^2$ to $\mathbb{D}_p^2$. It is also a solid torus (see Figure 1).
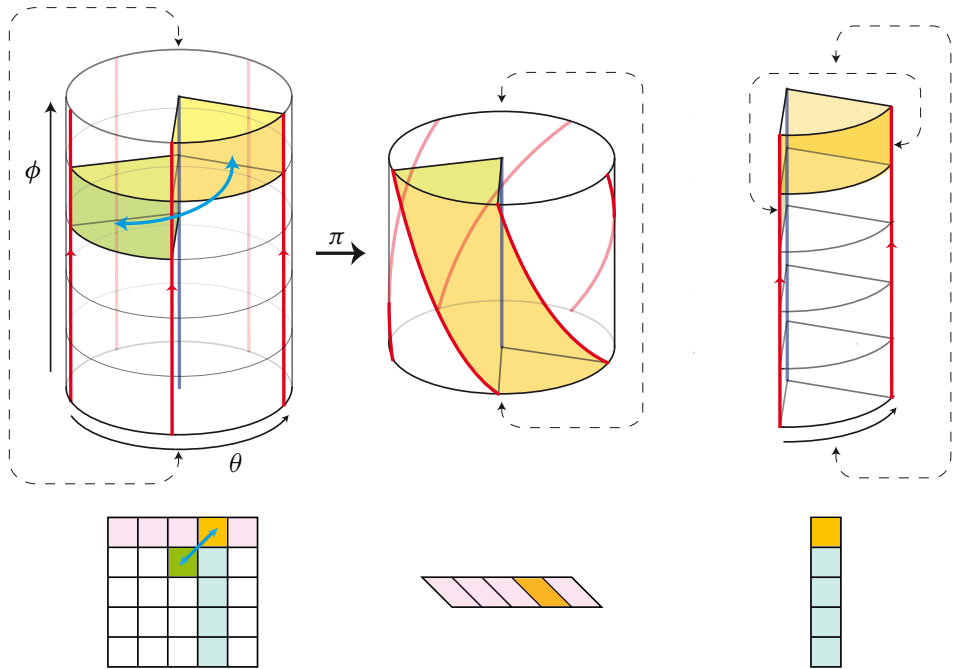

Figure 1: On the top left, the unit tangent bundle $T^1\mathbb{D}^2$ to a disc $\mathbb{D}^2$. It is a solid torus. The action of $\mathbb{Z}/p\mathbb{Z}$ (here with $p = 5$) is indicated with a blue arrow. It is a screw-motion. Thus $T^1\mathbb{D}^2$ can be seen as a tower formed of p^2 pieces of cheese, where the generator of $\mathbb{Z}/p\mathbb{Z}$ acts by a vertical translation plus a $2\pi/p$–rotation. On the bottom left, the boundary of $T^1\mathbb{D}^2$ with the rind of the p^2 pieces. A horizontal story of p pieces of cheese is then a fundamental domain for the action (in the center). The quotient is obtained by identifying the floor and the ceiling of the story with a $-2\pi/p$–rotation. Every meridian disc intersects each fiber p times, except the central fiber, which it intersects only once. This model (called the *story model*) shows that the unit tangent bundle is a Seifert fibered bundle [12]. The p pieces of cheese located between two vertical walls form another fundamental domain (on the right). The quotient is obtained by identifying two vertical walls with a vertical translation of length $2\pi/p$ (assuming the thickness of the cake to be 2π). We call this model the *slice-of-cake model*. Figures 13, 19, 20 and 23 are drawn using this model.

As every point in an orbifold admits a neighborhood of the form $\mathbb{D}^2$ or $\mathbb{D}_p^2$ for some p, the unit tangent bundle of every orbifold is obtained by gluing solid tori of type $T^1\mathbb{D}^2$ or $T^1\mathbb{D}_p^2$.

2.2 The geodesic flow on the unit tangent bundle

Assume that Σ is a good 2–dimensional orbifold. The orientation of Σ defines an orientation on the tangent planes, whence an orientation on $T^1\Sigma$.

Assume now that γ is an oriented curve drawn on Σ. For p lying on γ, let $T_p(\gamma)$ be the unit tangent vector to γ at p. Then the family of all pairs $(p, T_p(\gamma))$ is an oriented curve in $T^1\Sigma$, the *lifting* of γ in $T^1\Sigma$. In particular, the oriented geodesics of Σ are canonically lifted to $T^1\Sigma$. More precisely, for every point p in Σ and every direction v in $\mathbb{S}^1$, there exists a unique geodesic $\gamma_{p,v}$ of Σ going through p with the direction v. Now, for t in $\mathbb{R}$ and (p, v) in $T^1\Sigma$, let us define $\Phi_\Sigma(t, (p, v))$ to be (p', v') where x' is the unique point of $\gamma_{p,v}$ at distance t from p and v' is the unit tangent vector to $\gamma_{p,v}$ at p'. Then Φ_Σ is a continuous map of $\mathbb{R} \times T^1\Sigma$ to $T^1\Sigma$ and, by construction, it is additive in the first coordinate. Hence Φ_Σ is what is called a complete flow on $T^1\Sigma$, and it is naturally called the *geodesic flow* on $T^1\Sigma$. By construction, the liftings of the geodesics of Σ in $T^1\Sigma$ are the orbits of the geodesic flow (but they are not geodesic in $T^1\Sigma$, since no metric has been defined there).

2.3 Linking number and left handed flows

Assume that M is a 3–manifold, and that K, K' are two null-homologous links in M. Then there exists an oriented surface S (or even a simplicial 2–chain) with boundary K that is transverse to K'. The intersection points between S and K' then have an orientation, and their sum defines the algebraic intersection number $\mathrm{Int}(S, K')$. Adding a closed 2–chain to S does not change the intersection number since K' is null-homologous, so that $\mathrm{Int}(S, K')$ depends on K and K' only. It is the *linking number* of the pair K, K', denoted by $\mathrm{Lk}(K, K')$.

In the last fifty years, several works [1; 2; 7; 15] have emphasized the interest of considering a vector field as a long knot, or, more precisely, of considering invariant measures under the flow as (infinite) invariant knots. Following this idea, given a flow Φ on a rational homology sphere M, one can generalize the standard definition of the linking number for pairs of periodic orbits to pairs of invariant measures (see Arnol'd's work on asymptotic linking number [1]). Ghys then suggested to look at those flows for which this linking number is always negative, and called them *left handed flows*. We refer to the original article [10] for a discussion about the motivations

and the properties of these flows. Below we only mention the result explaining that, for a flow with many periodic orbits, left handedness can be deduced from the negativity of the linking numbers of pairs of periodic orbits only. A flow Φ is said *knot-shadowable* if, for every Φ–invariant measure μ, there exists a sequence (γ_n) of periodic orbits of Φ such that the sequence of the Dirac measures on γ_n weakly converges to μ.

Lemma 2-1 *Assume that Φ is a knot-shadowable flow. If the linking number of every pair of periodic orbits of Φ is negative, then Φ is left handed.*

Proof Assume that μ, μ' are two invariant measures. Let $(\gamma_n), (\gamma'_n)$ be two distinct sequences of knots that converge to μ, μ'. Write t_n, t'_n for the lengths of γ_n, γ'_n respectively. Then it is known [10] that the sequence $\frac{1}{t_n t'_n} \mathrm{Lk}(\gamma_n, \gamma'_n)$ converges to $\mathrm{Lk}(\mu, \mu')$, which is therefore negative. □

Lemma 2-1 is useful only for flows that are knot-shadowable. This is the case for flows of Anosov type, and in particular for the geodesic flows on hyperbolic 2–orbifolds. Thus a positive answer to Question 1-2 follows from a positive answer to Question 1-1. In short, if the curvature is negative, we only have to compute linking numbers of pairs of knots for proving left-handedness.

2.4 Coverings

We complete this introductory section with an observation about the behavior of linking numbers under quotient. The result is easy, but useful, as it gives new left handed flows from old ones. It will be crucial for establishing the left-handedness of $\Phi_{\Sigma_{2,3,7}}$ (Proposition 6-14).

Lemma 2-2 *Assume that $M, \widehat{M}$ are two 3–manifolds with a covering map $\pi\colon \widehat{M} \to M$ of index n. Let K, K' be two null-homologous links in M. Write $\widehat{K}, \widehat{K}'$ for the π–equivariant lifts of K, K' in $\widehat{M}$. Then the links $\widehat{K}, \widehat{K}'$ are null-homologous, and we have $\mathrm{Lk}(K, K') = \frac{1}{n}\mathrm{Lk}(\widehat{K}, \widehat{K}')$.*

Proof Let S be an oriented surface with boundary K. Write $\widehat{S}$ for its π–equivariant lift in $\widehat{M}$. Then we have $\pi(\partial \widehat{S}) = K$, hence $\partial \widehat{S} = \widehat{K}$. Therefore $\widehat{K}$ is also null-homologous. Since $\widehat{S}$ and $\widehat{K}'$ are π–equivariant, every intersection point of S with K' lifts to n intersection points of $\widehat{S}$ with $\widehat{K}'$, so that $\mathrm{Lk}(K, K') = \frac{1}{n}\mathrm{Lk}(\widehat{K}, \widehat{K}')$ holds. □

If we have a covering map between two orbifolds $\widehat{\Sigma} \to \Sigma$, then the map extends to the unit tangent bundles and it commutes with the geodesic flow. Lemma 2-2 then implies

that the sign of the linking numbers in $T^1\Sigma$ are the same as those in $T^1\widehat{\Sigma}$, so that, if the geodesic flow $\Phi_{\widehat{\Sigma}}$ is left handed, so does Φ_Σ. For instance, as the geodesic flow on $T^1\mathbb{S}^2$ is left handed [10], we deduce that the same holds for any quotient of $\mathbb{S}^2$, such as the Poincaré sphere $\Sigma_{2,3,5}$.

Corollary 2-3 *Let Σ be a spherical 2–orbifold. Then the geodesic flow Φ_Σ is left handed.*

3 Birkhoff sections for the geodesic flow on a flat torus

This section is devoted to the geodesic flow $\Phi_{\mathbb{T}^2}$ on a torus with a flat metric. Our aim is to establish Theorem B. On the way, we shall completely classify Birkhoff sections up to isotopy and show that (almost) every collection of periodic orbits bounds a Birkhoff section (Theorem 3-12 and Corollary 3-13).

We first parametrize the geodesic flow on a flat torus and define the polygon Pol_γ associated with a finite collection γ of periodic orbits (Section 3.1). Next, we describe how Birkhoff sections may look like, first in the neighborhood of so-called regular levels (Section 3.2), then in the neighborhood of critical levels with the help of helix boxes (Section 3.3). Finally, pieces are glued together in Section 3.4.

3.1 The polygon associated with a collection of periodic orbits

We show how to encode finite collections of periodic orbits of the geodesic flow $\Phi_{\mathbb{T}^2}$ in $T^1\mathbb{T}^2$ using polygons whose vertices have integral coordinates.

Throughout this section, $\mathbb{T}^2$ denotes the torus equipped with a flat metric. By definition, $\mathbb{T}^2$ is a quotient $\mathbb{R}^2/\mathbb{Z}^2$ of the Euclidean plane. For all p, p' in $\mathbb{T}^2$, the translation by $p' - p$ carries the tangent plane at p to the tangent plane at p'. Therefore, the unit tangent bundle $T^1\mathbb{T}^2$ is $\mathbb{T}^2 \times \mathbb{S}^1$. Next, the geodesics of $\mathbb{T}^2$ are induced by those of $\mathbb{R}^2$. Their liftings in $T^1\mathbb{T}^2$ are horizontal, that is, lie in some level $\mathbb{T}^2 \times \{\theta\}$ for some θ in $\mathbb{S}^1$. Hence we have $\Phi_{\mathbb{T}^2}(t, (x, y, \theta)) = (x + t\cos\theta, y + t\sin\theta, \theta)$. If $\tan\theta$ is a rational number, then, for every initial value of (x, y), the associated orbit goes back to (x, y) in finite time, and, conversely, every finite orbit of $\Phi_{\mathbb{T}^2}$ is of this type. In such a case, we define θ to be the *slope* of the orbit, and the unique pair (p, q) of coprime numbers verifying $\tan\theta = p/q$ and p is of the same sign as $\cos\theta$ to be the *code* of the orbit.

Assume that γ is a finite collection of periodic orbits of $\Phi_{\mathbb{T}^2}$. We define the *combinatorial type* of γ to be the sequence $((n_1, \theta_1, p_1, q_1), \dots, (n_k, \theta_k, p_k, q_k))$, such

that γ consists of n_1 orbits of slope θ_1, plus n_2 orbits of slope $\theta_2, \ldots,$ plus n_k orbits of slope θ_k, we have $\tan\theta_1 = p_1/q_1, \ldots, \tan\theta_k = p_k/q_k$, and $\theta_1, \ldots, \theta_k$ make an increasing sequence in $[0, 2\pi)$.

Lemma 3-1 *Assume that γ is a finite collection of periodic orbits in $\Phi_{\mathbb{T}^2}$. Let $((n_1,, \theta_1, p_1, q_1), \ldots, (n_k, \theta_k, p_k, q_k))$ be the combinatorial type of γ. Then the image of γ in $H_1(T^1\mathbb{T}^2; \mathbb{Z})$ is zero if and only if $\sum n_i(p_i, q_i) = (0, 0)$ holds.*

Proof The image of an orbit with slope (p, q) in $H_1(T^1\mathbb{T}^2; \mathbb{Z})$ admits the coordinates $(p, q, 0)$ in the standard basis. Indeed, the class of a straight line with code (p, q) on $\mathbb{T}^2$ is (p, q) in this basis. As the lifts of the geodesics of $\mathbb{T}^2$ in $T^1\mathbb{T}^2$ are horizontal, the third coordinate of the lift of a geodesic in $T^1\mathbb{T}^2$ is constant. Therefore the third coordinate of its image in $H_1(T^1\mathbb{T}^2; \mathbb{Z})$ is zero. The result then follows from the additivity of homology. □

Here comes the main definition of this section.

Definition 3-2 (See Figure 2.) Assume that γ is a null-homologous collection of periodic orbits in $\Phi_{\mathbb{T}^2}$ with combinatorial type $((n_1, \theta_1, p_1, q_1), \ldots, (n_k, \theta_k, p_k, q_k))$. The *polygon* Pol_γ of γ is the k–vertex polygon of $\mathbb{R}^2$ whose j^{th} vertex is $\sum_{i=1}^{j} n_i(p_i, q_i)$ for $j = 1, \ldots, k$.
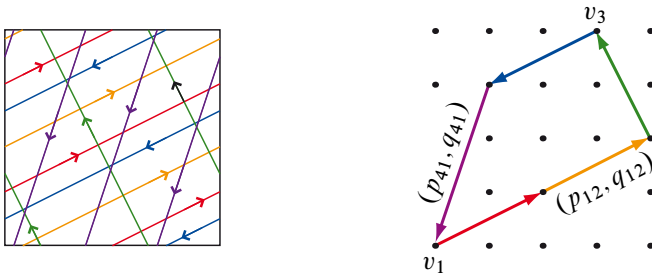

Figure 2: A null-homologous family γ of periodic orbits of the geodesic flow, and the associated polygon Pol_γ

Owing to the order condition on the slopes in the combinatorial type, Pol_γ is a convex polygon and, as p_i and q_i are coprime for every i, the only points on the boundary of Pol_γ that have integral coordinates are the vertices plus the intermediate points of the form $\sum_{i=1}^{j-1} n_i(p_i, q_i) + m(p_j, q_j)$ with $m < n_j$.

3.2   Transverse surfaces and regular levels

We now turn to surfaces in $T^1\mathbb{T}^2$ transverse to $\Phi_{\mathbb{T}^2}$, with the aim of connecting the existence of such a surface with boundary γ with the properties of the associated polygon Pol_γ.

Hereafter, for every θ in $\mathbb{R}/2\pi\mathbb{Z}$, the subset of $T^1\mathbb{T}^2$ made of the points whose last coordinate is θ will be called the θ^{th} *level* of $T^1\mathbb{T}^2$, denoted by L_θ. As $T^1\mathbb{T}^2$ is trivial, every level is a copy of $\mathbb{T}^2$. If γ is a null-homologous collection of periodic orbits in $\Phi_{\mathbb{T}^2}$ with combinatorial type $((n_1, \theta_1, p_1, q_1), \ldots, (n_k, \theta_k, p_k, q_k))$, the k angles θ_i, as well as the associated levels of $T^1\mathbb{T}^2$, will be called γ–*critical*, whereas the other angles will be called γ–*regular*.

Lemma 3-3   *Assume that γ is a null-homologous collection of periodic orbits of $\Phi_{\mathbb{T}^2}$ and S is a surface with boundary γ whose interior is transverse to $\Phi_{\mathbb{T}^2}$. For θ in $\mathbb{R}/2\pi\mathbb{Z}$, let S_θ be the intersection of S with the level L_θ. Then, if θ is γ–regular, S_θ is a union of disjoint circles. If θ, θ' are γ–regular and the interval (θ, θ') contains no γ–critical angle, S_θ and $S_{\theta'}$ are homologous.*

Proof   By construction, the geodesic flow $\Phi_{T^1\mathbb{T}^2}$ is tangent to L_θ whereas, by assumption, S is transverse to $\Phi_{T^1\mathbb{T}^2}$. Hence S and L_θ are transverse. Therefore their intersection is a closed 1–dimensional submanifold of L_θ, hence a union of parallel disjoint circles. The family $(S_t)_{t \in [\theta, \theta']}$ provides an isotopy between S_θ and S'_θ. These (multi)-curves are therefore homologous. $\square$

In the above context, the multicurve S_θ is called a *stratum* of S. For every γ–regular value θ, the stratum S_θ is cooriented by the geodesic flow. By convention, we orient it so that the concatenation of the chosen orientation and the orientation of the flow gives the standard orientation on the torus L_θ. With this choice, the class $[S_\theta]$ is a well-defined element of the group $H_1(L_\theta; \mathbb{Z})$, the latter being canonically identified with $H_1(\mathbb{T}^2; \mathbb{Z})$. Then, Lemma 3-3 implies that $[S_\theta]$ is constant when θ describes an interval of γ–regular values. Our goal now is to understand how $[S_\theta]$ evolves when θ passes a γ–critical value.

3.3   Packing into helix boxes

Van Horn-Morris [17] constructed open book decompositions of the torus bundles over the circle by using special boxes and controlling how they match with each other. We use now the same elementary boxes for decomposing and describing the surfaces whose boundary is transverse to $\Phi_{\mathbb{T}^2}$ around critical levels.

Definition 3-4 A *positive (resp. negative) helix box* is a cube containing an oriented surface isotopic to the surface depicted on Figure 3, called the *helix*. The oriented boundary of the helix is made of seven oriented segments lying in the faces of the cube, plus one segment, called the *binding*, lying inside the cube and connecting two opposite faces of the cube.
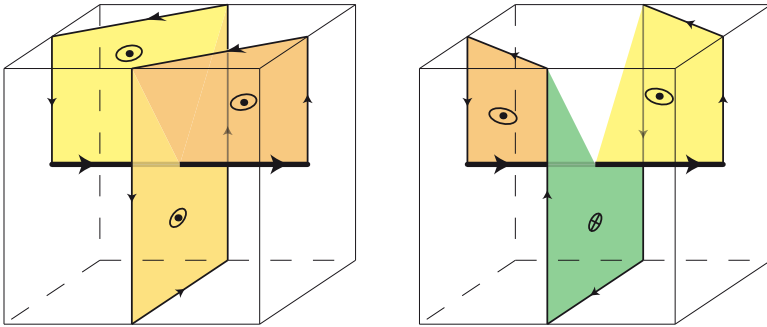

Figure 3: A positive helix box on the left, a negative helix box on the right. The bindings are in bold. The orientations of the helices are represented by dotted and crossed circles. The interiors of the helices are transverse to the direction of the binding, positively or negatively oriented according to the sign of the box.

The next result asserts that almost every surface transverse to $\Phi_{\mathbb{T}^2}$ is locally made of helices.

Lemma 3-5 *(See Figures 3 and 4.) Assume that γ is a null-homologous collection of periodic orbits of $\Phi_{\mathbb{T}^2}$ and S is a surface with boundary γ whose interior is transverse to $\Phi_{\mathbb{T}^2}$. Let γ_i be an element of γ. Denote by L_{θ_i} the γ–critical level containing γ_i. Then there exists a small tubular neighborhood N_{γ_i} of γ_i of the form $]\gamma_i - \epsilon, \gamma_i + \epsilon[\times]\theta_i - \eta, \theta_i + \eta[$ in $\mathbb{T}^2 \times \mathbb{S}^1$ such that:*

(i) *If the interior of S does not intersect the level L_{θ_i}, then the surface S is negatively transverse to $\Phi_{\mathbb{T}^2}$ and is locally isotopic to $\gamma_i \times [\theta, \theta + \epsilon[$ or to $\gamma_i \times]\theta - \epsilon, \theta]$.*

(ii) *Otherwise N_{γ_i} can be decomposed as the union of a positive number t_{γ_i} of helix boxes, which are all positive (resp. negative) if S is positively (resp. negatively) transverse to $\Phi_{\mathbb{T}^2}$, and such that γ_i is identified with the union of the bindings, S is the union of the helices, and the horizontal and vertical faces of N_{γ_i} are identified with the horizontal and vertical faces of the helix boxes.*
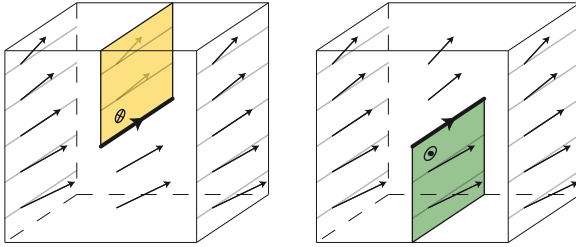
Figure 4: Case (i) of Lemma 3-5 when the surface S is negatively transverse to the flow and the vector $\vec{n}_p$ always points in the same half-space. The boundary ∂S is in bold.

Proof We write $N_{\epsilon,\eta}$ for the tubular neighborhood $]\gamma_i - \epsilon, \gamma_i + \epsilon[\times]\theta_i - \eta, \theta_i + \eta[$ of γ_i in $\mathbb{T}^2 \times \mathbb{S}^1$. For every point p on γ_i, we denote by $\vec{n}_p$ the unique unit vector orthogonal to γ_i, tangent to S, and pointing inside S. If ϵ and η are small enough, then the intersection of S with $N_{\epsilon,\eta}$ is isotopic to the surface generated by $p + t\vec{n}_p$ when p describes γ_i and t is non-negative. We choose for N_{γ_i} such a neighborhood. The surface L_{θ_i} induces a trivialization of the unit normal bundle $\nu_p(\gamma_i)$ of γ_i, so that we can define $\psi(p)$ to be the angle between $\vec{n}_p$ and L_{θ_i}. We then set d_{γ_i} to be the degree of the map $\psi \colon \gamma_i \simeq \mathbb{S}^1 \to \nu_p(\gamma_i) \simeq \mathbb{S}^1$.

If S is positively transverse to $\Phi_{\mathbb{T}^2}$, then $\psi(p)$ increases as p describes the curve γ_i. Therefore the degree d_{γ_i} of ψ is positive. We then obtain the helix boxes by cutting N_{γ_i} at each point where $\vec{n}_p$ points upward. This happens d_{γ_i} times, thus yielding d_{γ_i} positive helix boxes. The result when S is positively transverse follows with $t_{\gamma_i} = d_{\gamma_i}$.

If S is negatively transverse to the flow, then ψ is a non-increasing function. Indeed, since the geodesic flow is not parallel to γ_i, but rotates when level changes, the vector $\vec{n}_p$ can be constant and the application ψ can be of degree 0; see Figure 4. If so, the surface S lies on one side of L_{θ_i} only. It is therefore isotopic to $\gamma_i \times [\theta_i, \theta_i + \epsilon[$ or to $\gamma_i \times]\theta_i - \epsilon, \theta_i]$, and we are in case (i) of the statement. Otherwise, the degree d_{γ_i} of ψ is negative, and the situation is similar to that in the positive case. The only difference is that the negativity of the intersection of S with $\Phi_{\mathbb{T}^2}$ forces S to wind in the other direction, so that we obtain $-d_{\gamma_i}$ negative boxes. The result then follows with $t_{\gamma_i} = -d_{\gamma_i}$.   □

In the above context, the tubular neighborhood N_{γ_i} of γ_i is called a *product-neighborhood* of γ_i. If the interior of S does not intersect the level L_{θ_i} (Case (ii)), then N_{γ_i} is assumed to be decomposed as a union of t_{γ_i} helix boxes.

Lemma 3-5 gives the structure of a surface transverse to the flow around its boundary. The next result decomposes such a surface around an entire critical level.

Lemma 3-6 *Assume that γ is a null-homologous collection of periodic orbits of $\Phi_{\mathbb{T}^2}$ with combinatorial type $((n_1, \theta_1, p_1, q_1), \ldots, (n_k, \theta_k, p_k, q_k))$, and S is a surface with boundary γ whose interior is transverse to $\Phi_{\mathbb{T}^2}$. Let i be an element of $\{1, \ldots, k\}$. Call $\gamma_{i,1}, \ldots, \gamma_{i,n_i}$ the n_k elements of γ lying in the γ–critical level L_{θ_i}, and suppose that $N_{\gamma_{i,1}}, \ldots, N_{\gamma_{i,n_i}}$ are the associated product-neighborhoods. Then all the curves $\gamma_{i,1}, \ldots, \gamma_{i,n}$ are parallel, and all the numbers $t_{\gamma_{i,1}}, \ldots, t_{\gamma_{i,n_i}}$ are equal to some number, say t_{θ_i}. Moreover, if t_{θ_i} is not zero, there exists a neighborhood of L_{θ_i} of the form $]L_{\theta_i-\epsilon}, L_{\theta_i+\epsilon}[$, which is tiled by $n_i \times t_{\theta_i}$ helix boxes such that, in each helix box, the surface S coincides with the helix.*

Proof By definition of $\Phi_{\mathbb{T}^2}$, every orbit in the level L_{θ_i} has direction θ_i. At the expense of possibly restricting some of them, we can suppose that all rectangular tubular neighborhoods $N_{\gamma_{i,j}}$ have the same height 2η. Then the complement of their union $N_{\gamma_{i,1}} \cup \cdots \cup N_{\gamma_{i,n_i}}$ in the horizontal thick torus $]L_{\theta_i-\eta}, L_{\theta_i+\eta}[$ is also the union of n_i solid tori admitting a rectangular section. We denote these tori by $M_{i,1}, \ldots, M_{i,n_i}$. At the expense of possibly permuting the names, we can suppose that, for every j, the torus $M_{i,j}$ lies between the tori $N_{\gamma_{i,j}}$ and $N_{\gamma_{i,j+1}}$. Since S is transverse to the flow, its intersection with $M_{i,j}$ is transverse to the direction θ_i. Therefore it is the union of a certain number, say $s_{i,j}$, of discs whose boundaries are meridian circles in the solid torus $M_{i,j}$.

If, for some j, the number $t_{\gamma_{i,j}}$ is zero, then the two vertical boundaries of $N_{\gamma_{i,j}}$ do not intersect S. Therefore, the intersection of S with $M_{i,j}$ is empty, which implies $s_{i,j} = 0$. Considering the other boundary of $M_{i,j}$, we get $t_{\gamma_{i,j+1}} = 0$. By induction, we get $t_{\gamma_{i,j}} = 0$ for every j.

If, for some j, the number $t_{\gamma_{i,j}}$ is not zero, then $N_{\gamma_{i,j}}$ is tiled into $t_{\gamma_{i,j}}$ helix boxes. Therefore the boundary between $M_{i,j}$ and $N_{\gamma_{i,j}}$ is an annulus that intersects S along $t_{\gamma_{i,j}}$ vertical segments, and we deduce $s_{i,j} = t_{\gamma_{i,j}}$. Considering the other vertical boundary of $M_{i,j}$, we get $s_{i,j} = t_{\gamma_{i,j+1}}$, and therefore $t_{\gamma_{i,j+1}} = t_{\gamma_{i,j}}$. By induction, all numbers $t_{\gamma_{i,j}}$ are equal to some fixed number, say t_{θ_i}. Finally, since the intersection of S with $M_{i,j}$ consists of discs only, we can extend the solid tori $N_{\gamma_{i,j}}$ so that their union covers the whole neighborhood $]L_{\theta_i-\eta}, L_{\theta_i+\eta}[$. Since every $N_{\gamma_{i,j}}$ is tiled by t_{θ_i} helix boxes, the thick torus $]L_{\theta_i-\eta}, L_{\theta_i+\eta}[$ is tiled by $n_i \times t_{\theta_i}$ helix boxes. $\qquad\square$

Considering for a moment the angular parameter θ as a (periodic) time, a surface transverse to $\Phi_{\mathbb{T}^2}$ can be seen as the movie of its strata. By Lemma 3-3, the strata vary continuously as long as θ is regular. Using Lemma 3-6, we can now describe how the strata evolve when θ crosses a critical value.

Lemma 3-7 In the context of Lemma 3-6, if the surface S is negatively transverse $\Phi_{\mathbb{T}^2}$, then for every γ–critical angle θ_i, the homology classes of the strata $S_{\theta_i-\eta}$ and $S_{\theta_i+\eta}$ are related by the relation

$$(3\text{-}8) \qquad\qquad [S_{\theta_i+\eta}] = [S_{\theta_i-\eta}] + n_i(p_i, q_i).$$

If S is positively transverse to $\Phi_{\mathbb{T}^2}$, then we have

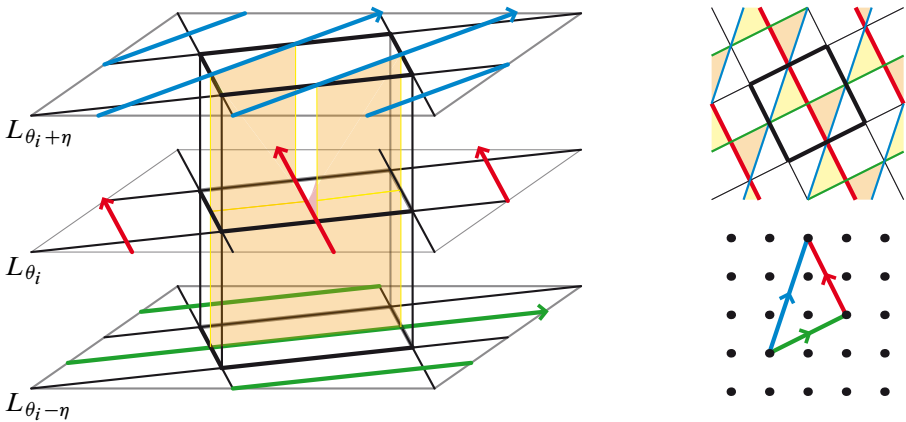$$(3\text{-}9) \qquad\qquad [S_{\theta_i+\eta}] = [S_{\theta_i-\eta}] - n_i(p_i, q_i).$$

Figure 5: On the left, a surface S with boundary γ transverse to $\Phi_{\mathbb{T}^2}$ in a neighborhood of a γ–critical level L_{θ_i} of the form $]L_{\theta_i} - \eta, L_{\theta_i} + \eta[$. It is tiled by five negative helix boxes. Here, there is only one component, say $\gamma_{i,1}$, of γ in L_{θ_i} (in red), that is, we have $n_i = 1$. The intersection of S with one of the five helix boxes is depicted. Its boundary consists of one fifth of the curve $\gamma_{i,1}$, one fifth of the stratum $S_{\theta_i+\eta}$ (on the top, in blue), one fifth of $S_{\theta_i-\eta}$ (on the bottom, in green), and of vertical segments which are glued to the four other boxes. On the top right is the projection on a horizontal torus. On the bottom right is the homological relation between $n_i(p_i, q_i)$, $[S_{\theta_i-\eta}]$ and $[S_{\theta_i+\eta}]$ stated in Lemma 3-7, here with $n_i = 1$, $(p_i, q_i) = (-1, 2)$, $[S_{\theta_i-\eta}] = (2, 1)$ and $[S_{\theta_i+\eta}] = (1, 3)$. According to Lemma 3-11(ii), the area of this homological triangle $\frac{5}{2}$ is half the number of helix boxes involved in the tiling of the neighborhood of the γ–critical level L_{θ_i}.

Proof We continue with the notation of Lemma 3-6. In particular, we assume that the neighborhood $]L_{\theta_i-\eta}, L_{\theta_i+\eta}[$ of L_θ is tiled with $n_i \times t_{\theta_i}$ helix boxes. The boundary of the intersection of the surface S with $]L_{\theta_i-\eta}, L_{\theta_i+\eta}[$ consists of pieces of three types: the curves $\gamma_{i,1}, \ldots, \gamma_{i,n_i}$, the stratum $S_{\theta_i-\eta}$ and the stratum $S_{\theta_i+\eta}$. Therefore, the

sum of these curves, with the orientations induced by the surface S, is null-homologous in $T^1\mathbb{T}^2$. After projection on $\mathbb{T}^2$, this sum is still zero.

When S is negatively transverse to $\Phi_{\mathbb{T}^2}$, then the two orientations on $S_{\theta_i-\eta}$ given by S and by $\Phi_{\mathbb{T}^2}$ agree, whereas the two orientations on $S_{\theta_i+\eta}$ are opposite. We thus get $n_i(p_i, q_i) + [S_{\theta_i-\eta}] - [S_{\theta_i+\eta}] = 0$. Similarly, when S is positively transverse to $\Phi_{\mathbb{T}^2}$, the two orientations on $S_{\theta_i-\eta}$ are opposite, whereas the two orientations on $S_{\theta_i+\eta}$ agree, yielding Equation (3-9). □

3.4 Correspondence between pointed polygons and transverse surfaces

We can now associate with every surface transverse to $\Phi_{\mathbb{T}^2}$ a polygon in the lattice $H_1(\mathbb{T}^2; \mathbb{Z})$ that encodes the homology classes of all strata simultaneously.

Definition 3-10 Assume that γ is a null-homologous collection of periodic orbits of $\Phi_{\mathbb{T}^2}$ and S is a surface with boundary γ whose interior is transverse to $\Phi_{\mathbb{T}^2}$. The *pointed polygon* Pol°_S of S is the polygon of $\mathbb{R}^2$ whose vertices are the points $[S_\theta]$ for θ a γ–regular angle.

Lemma 3-11 (i) *In the above context, let* $((n_1, \theta_1, p_1, q_1), \ldots, (n_k, \theta_k, p_k, q_k))$ *be the combinatorial type of* γ *and* Pol_γ *be the polygon associated with* γ. *If* S *is negatively transverse to* $\Phi_{\mathbb{T}^2}$, *then the polygon* Pol°_S *is a pointed copy of* Pol_γ. *If* S *is positively transverse to* $\Phi_{\mathbb{T}^2}$, *then* Pol°_S *is obtained from* Pol_γ *by a reflection.*

(ii) *For every* γ–*critical angle* θ_i, *the number of helix boxes used for tessellating a neighbourhood* $]L_{\theta_i-\eta}, L_{\theta_i+\eta}[$ *is equal to the area of the parallelogram spanned by the vectors* $[S_{\theta_i-\eta}]$ *and* $[S_{\theta_i+\eta}]$.

Proof By Lemma 3-3, the polygon Pol°_S has at most k vertices. Now if S is negatively transverse to $\Phi_{\mathbb{T}^2}$, then (3-8) implies that, for every i, the two vertices $[S_{\theta_i-\eta}]$ and $[S_{\theta_i+\eta}]$ of Pol°_S differ by the vector $n_i(p_i, q_i)$. On the other hand, if S is positively transverse to $\Phi_{\mathbb{T}^2}$, then $[S_{\theta_i-\eta}]$ and $[S_{\theta_i+\eta}]$ differ by $-n_i(p_i, q_i)$. This proves (i).

For (ii), we see on Figure 5 that, for every i, every helix box used in the tiling of the neighbourhood $]L_{\theta_i-\eta}, L_{\theta_i+\eta}[$ of L_{θ_i} is above an intersection point between the projection of the curve $S_{\theta_i+\eta}$ and the projection of one component of γ in L_{θ_i}. Therefore the number of helix boxes in the tiling is the absolute value of the intersection number of the classes $[S_{\theta_i+\eta}]$ and $n_i(p_i, q_i)$ in $H_1(\mathbb{T}^2; \mathbb{Z})$. As depicted in Figure 5 (right), this number coincides with the absolute value of the intersection number of $[S_{\theta_i-\eta}]$ and $[S_{\theta_i+\eta}]$, which is the area of the parallelogram spanned by these two vectors. □

Assume that S is a surface transverse to $\Phi_{\mathbb{T}^2}$ and L_θ is a regular level of $T^1\mathbb{T}^2$ for S, so that the stratum S_θ is a smooth multicurve. Then we obtain another surface S' transverse to $\Phi_{\mathbb{T}^2}$ by cutting S along S_θ, gluing a copy of L_θ and smoothing. We say that S' is obtained from S by *horizontal surgery*. It is easy to check that the polygons Pol_S° and $\mathrm{Pol}_{S'}^\circ$ coincide, although the surfaces S and S' are not isotopic. Therefore pointed polygons do not encode all information about the isotopy type of transverse surfaces. Nevertheless, we will see that horizontal surgeries are the only freedom left by polygons.

For γ a null-homologous collection of periodic orbits of $\Phi_{\mathbb{T}^2}$ with associated polygon Pol_γ, we write $A(\gamma)$ for the area of Pol_γ (which is an integer by Pick's formula) and $I(\gamma)$ for the number of integer points in the strict interior of Pol_γ. We can now state and establish the main result.

Theorem 3-12   (i)   *The map $S \mapsto \mathrm{Pol}_S^\circ$ induces a one-to-one correspondence between surfaces negatively transverse to the flow $\Phi_{\mathbb{T}^2}$ with boundary made of periodic orbits, up to isotopy and horizontal surgeries, and convex polygons with integer vertices in $\mathbb{R}^2$ containing the origin in their interior or on their boundary.*

   (ii)   *The map $S \mapsto \mathrm{Pol}_S^\circ$ induces a one-to-one correspondence between negative Birkhoff sections for the flow $\Phi_{\mathbb{T}^2}$ and convex polygons with integer vertices in $\mathbb{R}^2$ containing the origin their (strict) interior.*

   (iii)   *There is no surface positively transverse to $\Phi_{\mathbb{T}^2}$ with boundary made of periodic orbits.*

   (iv)   *Assume that γ is a null-homologous collection of periodic orbits of $\Phi_{\mathbb{T}^2}$ with associated polygon Pol_γ. Then for every surface S transverse to $\Phi_{\mathbb{T}^2}$ with boundary γ, the Euler characteristic of S is $-2A(\gamma)$ and the genus of S is $I(\gamma)$.*

   (v)   *Assume that γ, γ' are two null-homologous collections of periodic orbits of $\Phi_{\mathbb{T}^2}$. Then their linking number $\mathrm{Lk}(\gamma, \gamma')$ is equal to $A(\gamma) + A(\gamma') - A(\gamma \cup \gamma')$.*

Proof   (i)   (See Figure 6.) Assume that S is a surface negatively transverse to $\Phi_{\mathbb{T}^2}$. Let γ be its boundary and $((n_1, \theta_1, p_1, q_1), \ldots, (n_k, \theta_k, p_k, q_k))$ be the combinatorial type of γ. For every γ–regular angle θ, the stratum S_θ is transverse to the direction θ. Therefore, if S_θ is non-empty and with the orientation of S_θ defined in Section 3.2, the basis formed by a vector tangent to S_θ and a vector with direction θ is direct. Hence the basis formed by the direction of $[S_\theta]$ and the direction θ is also direct. Let D_θ be the line oriented by θ passing through the vertex S_θ of Pol_S°. The previous observation implies that the point $(0,0)$ is on the left of D_θ. Let θ_i be the smallest γ–critical angle larger than θ. Then, when θ tends to θ_i, the line D_θ tends to the line supporting the edge of Pol_S° with direction θ_i. Therefore, the point $(0,0)$ is also

on the left of the edge of Pol_S° with direction θ_i (the boundary of Pol_S° being oriented trigonometrically). If the stratum S_θ is empty, we have $[S_\theta] = 0$, and $(0, 0)$ is also on the left the line D_θ. Taking all critical values of θ into account, we deduce that the point $(0, 0)$ is on the left of all oriented edges of Pol_S°, and therefore lies in the interior or on the boundary of Pol_S°. Thus the map Pol° associates with every surface transverse to $\Phi_{\mathbb{T}^2}$ a polygon in $H_1(\mathbb{T}^2; \mathbb{Z})$ containing $(0, 0)$ in its interior.
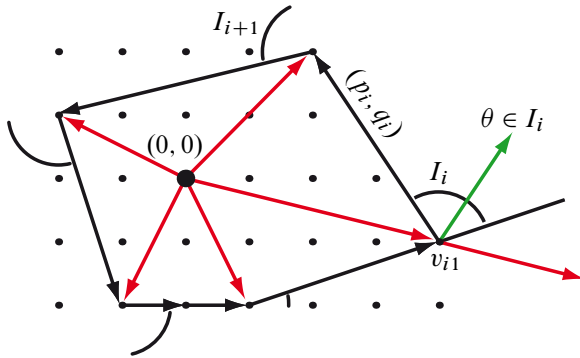

Figure 6: The Polygon Pol_S° for a surface S negatively transverse to the geodesic flow. For a γ–regular angle θ, the directions $[S_\theta]$ (indicated by a red arrow) and θ (indicated by a green arrow) form a direct basis. The point $(0, 0)$ is on the left of all edges of Pol_S°, and therefore in the interior or on the boundary of Pol_S°.

For the surjectivity of the map Pol°, suppose that a convex polygon P containing $(0, 0)$ is given. Let θ_0 be a γ–regular angle. Let V be the unique vertex of P such that the line of slope θ_0 passing through V lies on the right of P. Then construct a surface S a follows. Start with a stratum S_{θ_0} that is transverse to the θ_0–direction and whose homology class is V. Let (p_1, q_1) denote the edge of P starting at V. Then erect helix boxes whose bindings have direction (p_1, q_1) so that their bottom faces match with S_{θ_0}. By Lemma 3-7, the boundary of the helices in the top faces form a curve whose homology class is $V + (p_1, q_1)$, so that the stratum corresponds to the second vertex of P. By continuing this procedure of gluing helix boxes whose direction is prescribed by the edges of P and whose number is dictated by the strata that have been constructed previously, we erect a surface which is negatively transverse to $\Phi_{\mathbb{T}^2}$ and whose associated polygon is P.

For the injectivity, note that the surface S can be recovered from Pol_S° by the above procedure. The only choice arises when θ has described the whole circle $\mathbb{S}^1$ and comes back to θ_0: we have to glue the last floor of helix boxes to the stratum S_θ. This gluing is not unique, but two such gluings precisely differ by a horizontal surgery.

(ii)   Assume that S is a negative Birkhoff section for $\Phi_{\mathbb{T}^2}$. As S is transverse to the flow, we can apply the result of (i) and deduce that the polygon Pol_S° contains $(0,0)$ in its interior or on its boundary. Since S is a Birkhoff section, it intersects all orbits of $\Phi_{\mathbb{T}^2}$. In particular, this implies that for every γ–regular value of θ, the stratum S_θ is non-empty. This excludes the case where $(0,0)$ lies on the boundary of Pol_S°.

(iii)   Assume that S is a surface with boundary positively transverse to $\Phi_{\mathbb{T}^2}$. Then we can apply the same argument as in the negative case (i). The only difference is that, for every γ–regular angle θ, the basis formed by $[S_\theta]$ and θ is indirect. Therefore, the point $(0,0)$ lies on the right of the line with direction θ passing through the vertex $[S_\theta]$. Thus $(0,0)$ is on the right of all edges of Pol_S°, whereas the boundary is oriented trigonometrically, a contradiction.

(iv)   In every helix box, the helix surface consists of a topological disc, of eight edges, seven of them being on the boundary of the box, and of eight vertices, two of them being in the center of a face of the box and the six others in the middle of an edge of the box. Therefore, the contribution of a helix box to the Euler characteristic is $1 - (1 + \frac{7}{2}) + (\frac{2}{2} + \frac{6}{4}) = -1$. Assume that S is a surface transverse to $\Phi_{\mathbb{T}^2}$ with boundary γ, and let Pol_S° be the associated polygon. Let θ, θ' be two γ–regular angles such that there is exactly one γ–critical value in $]\theta, \theta'[$. Then, according to Lemma 3-11, the number of helix boxes used for tiling the thick torus lying between the two levels L_θ and $L_{\theta'}$ is twice the area of the triangle whose vertices are $(0,0)$, $[S_\theta]$ and $[S_{\theta'}]$. By summing over all γ–critical levels, we obtain that the total number of helix boxes in twice the area of Pol_S°, hence twice the area of Pol_γ. As the genus of S is given by the formula $\chi(S) = 2 - 2g(S) - \sum n_i$, Pick's formula for the area of a polygon with integral vertices gives the formula for the genus.

(v)   By definition, the linking number $\mathrm{Lk}(\gamma, \gamma')$ is the intersection number of a surface with boundary γ and the collection γ'. It is well-defined when γ is null-homologous, since, in this case, the intersection number does not depend on the choice of the surface. Here, let us pick a Birkhoff section for $\Phi_{\mathbb{T}^2}$ with boundary γ, and call it S_γ. Let $((n_1', \theta_1', p_1', q_1'), \ldots, (n_k', \theta_k', p_k', q_k'))$ be the combinatorial type of γ'. Then the intersection number of S_γ with a periodic orbit of $\Phi_{\mathbb{T}^2}$ of slope (p_i', q_i') is the opposite of the area of the parallelogram spanned by the vectors $[S_{\theta_i'}]$ and (p_i', q_i'). Since the area of Pol_S° equals the area of Pol_γ, the jigsaw puzzle depicted on Figure 7 shows that the sum of the areas of these parallelograms is equal to $A(\gamma \cup \gamma') - A(\gamma) - A(\gamma')$. □

Corollary 3-13 (Theorem B)   *Assume that Σ is a quotient of $\mathbb{T}^2$ on which any two geodesics intersect. Then for every pair γ, γ' of periodic orbits of Φ_Σ, the inequality $\mathrm{Lk}(\gamma, \gamma') < 0$ holds.*
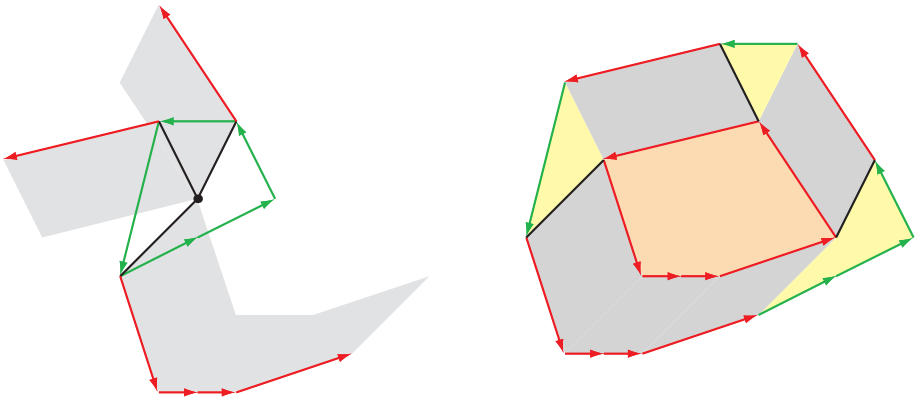
Figure 7: On the left, the polygon Pol_S° associated to a Birkhoff section S of $\Phi_{\mathbb{T}^2}$, and the homology classes of the elements of a family γ', each of them based at the vertex of Pol_S° corresponding to the class of the intersected stratum of S. In gray, the parallelograms whose areas add up to the intersection number of S with the collection γ'. On the right, the polygon $\mathrm{Pol}_{\gamma \cup \gamma'}$ is decomposed into three parts whose areas respectively are the area of the gray zone, the area of Pol_S° and the area of $\mathrm{Pol}_{\gamma'}$.

Proof By Lemma 2-2, it is enough to shows that the lifts $\hat{\gamma}, \hat{\gamma}'$ of γ, γ' in $T^1 \mathbb{T}^2$ have a negative linking number. As the projections of γ, γ' on Σ intersect, the projections of $\hat{\gamma}, \hat{\gamma}'$ on $\mathbb{T}^2$ also intersect. Theorem 3-12(v) shows that the linking number of two collections is zero if and only if the latter consists of parallel lifts of one geodesics on $\mathbb{T}^2$. The hypothesis on the intersection then discards this situation. $\square$

Theorem 3-12(ii) implies that almost every null-homologous collection of periodic orbits of $\Phi_{\mathbb{T}^2}$ bounds a Birkhoff section. The exceptions are the collections whose associated polygon contains no point with integral coordinates.

For example, let γ be an unoriented periodic geodesics on $\mathbb{T}^2$. Let (p, q) be its code. Denote by γ_+, γ_- its two lifts in $T^1 \mathbb{T}^2$ (one for each of the two possible orientations of γ). Then γ_+ and γ_- are periodic orbits of $\Phi_{\mathbb{T}^2}$, and their sum is null-homologous. The associated polygon is made of one segment with coordinates (p, q) only. As predicted by Theorem B(i), the union of γ_+ and γ_- bounds two non-isotopic surfaces that are transverse to $\Phi_{\mathbb{T}^2}$, namely the two vertical ribbons in $T^1 \mathbb{T}^2$ consisting of the unit tangent vectors which are based on γ and which point into one of the two sides of γ. None of these two ribbons is a Birkhoff section for $\Phi_{\mathbb{T}^2}$ since each of them only intersect half of the orbits.

For another example, consider the three orbits with respective slopes $(1, 0)$, $(0, 1)$ and $(-1, -1)$. They bound three non-isotopic surfaces transverse to $\Phi_{\mathbb{T}^2}$, but they do not bound any Birkhoff section, since the associated polygon is a triangle whose interior contains no point with integral coordinates.

A last example, which was a surprise for us, is given by the four orbits with slopes $(\pm 1, 0)$ and $(0, \pm 1)$, in which case the associated polygon is the unit square, again containing no integral point inside.

As explained in the introduction, Birkhoff sections give rise to open book decompositions for the underlying 3–manifold, here for unit tangent bundle $T^1 \mathbb{T}^2$, a 3–torus. Planar open book decompositions, that is, decompositions where the pages are of genus 0, have been often investigated. Theorem B(iv) implies that none of them comes from Birkhoff sections of the geodesic flow on the torus.

Corollary 3-14 *The geodesic flow on $T^1 \mathbb{T}^2$ contains no Birkhoff section of genus 0.*

Since helix boxes contribute -1 to the Euler characteristics, and since every helix box involves one boundary component, Birkhoff sections with genus 1 are very peculiar.

Corollary 3-15 *A Birkhoff section of genus 1 for $\Phi_{\mathbb{T}^2}$ is made of exactly one helix box per boundary component.*

In the article where he introduced what are now called Birkhoff sections [3], Birkhoff gave examples by constructing sections for the geodesic flow on the unit tangent bundle of every surface Σ. More precisely, a collection of periodic orbits of Φ_{Σ} is said to be *symmetric* if, for every element of the collection, the orbit corresponding to the opposite orientation of the underlying geodesics also belongs to the collection. Birkhoff showed that every large enough symmetric collection γ of periodic orbits of $\Phi_{\mathbb{T}^2}$ bounds a section. In the case when Σ is a torus with a flat metric, the symmetry hypothesis implies that the polygon Pol_γ is symmetric. The section constructed by Birkhoff corresponds to the surface S whose associated polygon Pol_S° is pointed in the center, that is, contains $(0, 0)$ as symmetry center.

4 Templates for the geodesic flow of a hyperbolic orbifold

We turn to hyperbolic orbifolds. The aim of this section is to show how the geodesic flow associated with an arbitrary hyperbolic 2–orbifold Σ can be distorted onto a certain multitemplate (Definition 4-6) lying inside $T^1 \Sigma$. The important property of this distortion is that its restriction to periodic orbits is an isotopy (Theorem 4-9), so

that the topological properties of the periodic orbits of the geodesic flow can be studied using this multitemplate. What makes the construction possible is that distinct periodic geodesics on a hyperbolic orbifold never point in the same direction at infinity. Our strategy is similar to that of Birman and Williams [4], who contract the stable direction of a hyperbolic flow. The characteristic here is that the explicit nature of the geodesic flow makes it possible to perform the construction in full detail.

Let Σ be a good hyperbolic 2–orbifold, and let Γ denote its fundamental group. Our strategy for constructing the template adapted to the geodesic flow Φ_Σ is as follows. We first choose an adapted tessellation of the universal cover $\mathbb{H}^2$ of Σ, namely, a Γ–invariant tesselation such that every tile contains at most one point whose stabilizer has order larger than 2. We also choose in every tile a smooth immersed graph pairwise connecting the sides in such a way that the graphs associated with adjacent tiles match on their common side. We then distort all geodesics in the hyperbolic plane into quasi-geodesics consisting of edges of the graphs so constructed (Section 4.1). Next, we lift this deformation in the unit tangent bundle $T^1\Sigma$ by forcing every tangent vector to always point toward its initial direction at infinity (Section 4.2). Then the image of the deformation at time 1 provides the expected (multi)template. It naturally carries a flow, namely the image of the geodesic flow by the deformation (Section 4.3).

4.1 Discretization of geodesics

The construction starts with a tessellation of the hyperbolic plane that behaves nicely with respect to the orbifold.

Definition 4-1 Assume that Γ is a Fuchsian group. Let Σ denote the orbifold $\mathbb{H}^2/\Gamma$. A tessellation $\mathcal{T}$ of $\mathbb{H}$ is *adapted* to Σ if:

(i) $\mathcal{T}$ is Γ–invariant.

(ii) Every tile of $\mathcal{T}$ is a convex polygon (with possibly some vertices on $\partial\mathbb{H}^2$).

(iii) Every tile of $\mathcal{T}$ contains at most one singular point in its interior, and points of index at most 2 on its boundary.

(iv) Every tile of $\mathcal{T}$ has a finite stabilizer in Γ.

(v) If T, T' are adjacent tiles of $\mathcal{T}$ separated by a side e_0, then, for all other sides e of T and e' of T' not both adjacent to e_0, the two geodesics respectively containing e and e' do not intersect.

For example, assume that Σ is a hyperbolic compact surface. Consider a convex polygonal fundamental domain D for the action of $\pi_1(\Sigma)$ on $\mathbb{H}^2$. Then the tessellation

formed by the images of D under the action of $\pi_1(\Sigma)$ is adapted to Σ. Note that condition (iii) prevents fundamental domains from providing tessellations adapted to arbitrary orbifolds. However, it is easy to see that, when a Fuchsian group Γ and a Γ–invariant tessellation $\mathcal{T}$ are given, one can always subdivide $\mathcal{T}$ and adapt it to $\mathbb{H}^2/\Gamma$. Condition (v) in Definition 4-1 may look strange. It is nevertheless important in order to guarantee that the ribbons of the template we will subsequently construct do not intersect (Lemma 4-8).

It would be natural to add a sixth constraint, namely that no periodic geodesic goes through a vertex (see Definition 4-3 below). However this is not always possible, in particular for the triangle groups we will be interested in Sections 5 and 6.

We now define, for every tessellation that is adapted to some orbifold, a graph that is *dual* to the tessellation, and on which we will then distort the geodesics of $\mathbb{H}^2$. We have to choose some additional data, namely to pick points on the sides of the tessellation and to choose edges connecting them, but the construction will not depend on these choices, ie, the templates we will eventually associate to two such choices will be isotopic. In the sequel, we use the word "side" when referring to the tiles of a tessellation, and "edge" when referring to a graph.

Assume that P is a polygon in $\mathbb{H}^2$ with finitely many sides $e_1, \ldots, e_n$, and suppose that $v_1, \ldots, v_n$ are points on $e_1, \ldots, e_n$ respectively. Let $\mathcal{G}_P^0$ be a complete unoriented graph with vertices $v_1, \ldots, v_n$, which is immersed in P so that its edges are orthogonal to the sides of P, and such that two edges intersect at most once (see Figure 8). Call *internal graph* of P associated to $v_1, \ldots, v_n$ the orientation cover $\mathcal{G}_P$ of $\mathcal{G}_P^0$, that is, the oriented graph with twice as many edges as $\mathcal{G}_P^0$, each of them corresponding to an edge of $\mathcal{G}_P$ oriented in one of the two possible ways. If e_i, e_j are two distinct sides of P_{4g+2}, we denote by $c_{e_i}^{e_j}$ the oriented edge of $\mathcal{G}_P$ connecting e_i to e_j.

Assume now that Γ is a Fuchsian group, and that $\mathcal{T}$ is a tessellation adapted to $\mathbb{H}^2/\Gamma$. A set $\mathcal{V}$ of points in $\mathbb{H}^2$ is called a $\mathcal{T}$–*marking* if every point in $\mathcal{V}$ lies on the common boundary between two tiles of $\mathcal{T}$, every side between two tiles of $\mathcal{T}$ contains exactly one element of $\mathcal{V}$, and $\mathcal{V}$ is Γ–invariant.

Definition 4-2 Assume that Γ is a Fuchsian group, that $\mathcal{T}$ is a tessellation of $\mathbb{H}^2$ adapted to $\mathbb{H}^2/\Gamma$ and that $\mathcal{V}$ is a $\mathcal{T}$–marking. Suppose that in every tile T of $\mathcal{T}$, there is an internal graph $\mathcal{G}_T$ associated to $\mathcal{V}$, and that the set of internal graphs is Γ–invariant. Then the union $\mathcal{G}_\mathcal{T}$ of all internal graphs $\mathcal{G}_T$ is said to be a graph *dual to* $\mathcal{T}$ and *associated with* $\mathcal{V}$.

It is easy to see that dual graphs exist for every tessellation. In the sequel, we will omit to mention the set $\mathcal{V}$ of marked points, since its choice does not influence the
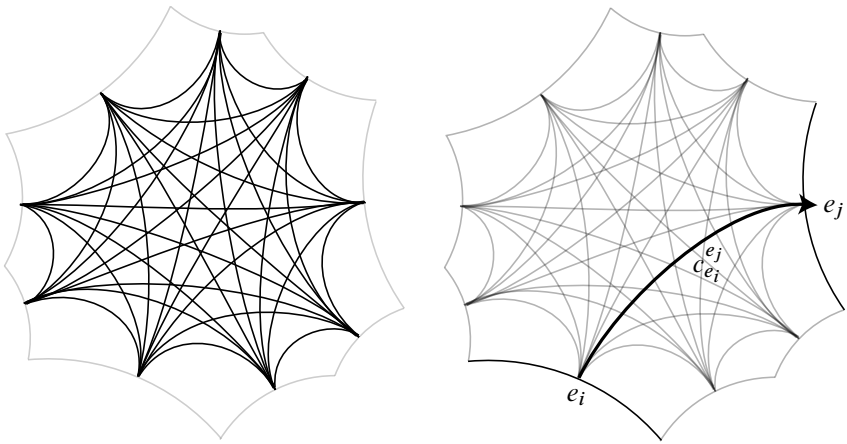
Figure 8: An internal graph. On the left, the unoriented graph $\mathcal{G}_P^0$. On the right, an oriented edge of $\mathcal{G}_P$.

construction. A graph dual is a sort of discretization of the hyperbolic plane adapted to a given Fuchsian group. If the Fuchsian group is of the first kind, that is, when its limit set is the whole boundary at infinity $\partial_\infty \mathbb{H}^2$, the limit set of any graph dual to any adapted tessellation is also the whole circle $\partial_\infty \mathbb{H}^2$. We now introduce a procedure that distorts the geodesics of $\mathbb{H}^2$ to curves included in the dual graph $\mathcal{G}_\mathcal{T}$.

Definition 4-3 Assume that Γ is a Fuchsian group of the first kind, and that $\mathcal{T}$ is a tessellation adapted to $\mathbb{H}/\Gamma$. Let $\mathcal{G}_\mathcal{T}$ be a graph dual to $\mathcal{T}$. Then a *discretization of geodesics on* $\mathcal{G}_\mathcal{T}$ is a family consisting, for every geodesics $\underline{\gamma}$ in $\mathbb{H}^2$, of:

(i)   A curve $\underline{\gamma}_\mathcal{T}$ embedded in $\mathcal{G}_\mathcal{T}$ that crosses the same tiles of $\mathcal{T}$ as $\underline{\gamma}$ (or a small perturbation of $\underline{\gamma}$ in case $\underline{\gamma}$ goes through a vertex of $\mathcal{G}_\mathcal{T}$).

(ii) An isotopy $f_{\mathcal{T},\gamma} \colon [0,1] \times \mathbb{R} \to \mathbb{H}^2$ between $\underline{\gamma}$ and $\underline{\gamma}_\mathcal{T}$, ie a smooth map such that $f^0_{\mathcal{T},\gamma}(t)$ describes $\underline{\gamma}$ when t describes $\mathbb{R}$, $f^1_{\mathcal{T},\gamma}(t)$ describes $\underline{\gamma}_\mathcal{T}$ when t describes $\mathbb{R}$ and, for every s in $[0,1]$, the curve $f^0_{\mathcal{T},\gamma}(\mathbb{R})$ is a smooth embedded curve in $\mathbb{H}$.

In addition, the family is supposed to be Γ–invariant in the sense that, if $g(\underline{\gamma}) = \underline{\gamma}'$ holds for some g in Γ, then $g(\underline{\gamma}_\mathcal{T}) = \underline{\gamma}'_\mathcal{T}$ and $g(f^s_{\mathcal{T},\gamma}(t)) = f^s_{\mathcal{T},\gamma'}(t)$ hold for every (s,t) in $[0,1] \times \mathbb{R}$.

The invariance condition implies in particular that, if $\underline{\gamma}$ is the lift of periodic geodesics on $\mathbb{H}/\Gamma$, then $\underline{\gamma}_\mathcal{T}$ projects on a periodic curve on $\overline{\mathbb{H}}/\Gamma$. More generally, it implies that all choices commute with the covering map $\mathbb{H} \to \mathbb{H}/\Gamma$.

Also, assume that a geodesic $\underline{\gamma}$ enters a tile T by a side e_i and leaves it by e_j, then its discretization $\gamma_{\mathcal{T}}$ visits the same tiles as $\underline{\gamma}$ before and after T. Therefore $\gamma_{\mathcal{T}}$ contains the edge $c_{e_i}^{e_j}$ of $\mathcal{G}_{\mathcal{T}}$.

Given a hyperbolic 2–orbifold and an adapted tessellation, the existence of discretization of geodesics easily follows from the definition.

A discretization of geodesics contracts many geodesics together. Indeed, if two oriented geodesics γ, γ' have one end in common, their discretization will necessarily coincide on some neighborhood of their positive end. Discretization will nevertheless be useful for studying Γ–periodic geodesics.

4.2 Lifting the discretization to the unit tangent bundle

Given an orbifold $\mathbb{H}^2/\Gamma$ and some additional data, using the discretization procedure of Definition 4-3, we distorted the geodesics of $\mathbb{H}^2$ onto some discrete graph. We now lift this procedure to the unit tangent bundle, in view of subsequently constructing the expected template for $\Phi_{\mathbb{H}/\Gamma}$.

Definition 4-4 Assume that Γ is Fuchsian group of the first kind, that $\mathcal{T}$ is an adapted tessellation of $\mathbb{H}^2$, that $\mathcal{G}_{\mathcal{T}}$ is a graph dual to $\mathcal{T}$ and that a discretization of geodesics on $\mathcal{G}_{\mathcal{T}}$ has been chosen. Then the associated *tearing map* of the unit tangent bundle is the map $F_{\mathcal{T}}$ from $[0, 1] \times T^1\mathbb{H}^2$ to $T^1\mathbb{H}^2$ defined as follows. For (p, v) in $T^1\mathbb{H}^2$, let $\underline{\gamma}$ denote the geodesics containing p and oriented by v, let $f_{\gamma,\mathcal{T}}$ denote the associated isotopy, let t_p be the real parameter such that $p = f_{\gamma,\mathcal{T}}^0(t_p)$, and let γ_+ be the positive extremity of γ in $\partial\mathbb{H}^2$. Then $F_{\mathcal{T}}^s(p, v)$ is defined to be the unique unit tangent vector based at $f_{\gamma,\mathcal{T}}^s(t_p)$ and pointing in the direction of $\underline{\gamma}_+$.

Note that a tearing map is not continuous. Indeed, since the graph $\mathcal{G}_{\mathcal{T}}$ is discrete, there are pairs of arbitrarily close tangent vectors that are mapped to different edges of $\mathcal{G}_{\mathcal{T}}$. Also, a tearing map can be injective when the time s is close to 0, but its time-1 map may, for instance, collapse some horocyle. For these two reasons, a tearing map is not an isotopy. Nevertheless, if we restrict to Γ–periodic geodesics, that is, to geodesics which are g–invariant for some g in Γ, we have:

Lemma 4-5 *In the above context, the restriction of $F_{\mathcal{T}}$ to vectors tangent to Γ– periodic geodesics is an isotopy.*

Proof Suppose that $F_{\mathcal{T}}^s(p_1, v_1) = F_{\mathcal{T}}^s(p_2, v_2)$ holds for some s in $[0, 1]$. Let $\underline{\gamma}_1, \underline{\gamma}_2$ denote the two geodesics tangent to v_1 and v_2 at p_1 and p_2, respectively. As the

vectors $F_{\mathcal{T}}^s(p_1, v_1)$ and $F_{\mathcal{T}}^s(p_2, v_2)$ point in the directions $(\underline{\gamma}_1)_+$ and $(\underline{\gamma}_2)_+$, their equality implies $(\underline{\gamma}_1)_+ = (\underline{\gamma}_2)_+$. Therefore $\underline{\gamma}_1$ and $\underline{\gamma}_2$ get closer with an exponential rate. By hypothesis, both are Γ–periodic, hence compact in $\mathbb{H}/\Gamma$. Therefore they coincide. By definition of $F_{\mathcal{T}}$, the equality $F_{\mathcal{T}}^s(p_1, v_1) = F_{\mathcal{T}}^s(p_2, v_2)$ implies

$$f_{\mathcal{T}, \underline{\gamma}_1}^s(p_1) = f_{\mathcal{T}, \underline{\gamma}_1}^s(p_2).$$

Since $f_{\mathcal{T}, \underline{\gamma}_1}^s$ is an isotopy, we deduce $p_1 = p_2$. Finally, since the vectors v_1, v_2 point in the same direction, they also coincide. $\qquad\square$

4.3 Multitemplates for the geodesic flow

We have now constructed a deformation of the unit tangent bundle that preserves the topology of periodic geodesics. Our task is now to determine the image of the deformation. In particular, we want to show that it lies inside some specific object that we call a *multitemplate*.

Definition 4-6 (See Figure 9.) Assume that M is a 3–manifold. A *ribbon* in M is an embedded surface in M diffeomorphic to $[0, 1]^2$ equipped with the horizontal flow generated by $\frac{\partial}{\partial x}$. If Rib is a ribbon, we denote by X_{Rib} the vector field on it.

A *multitemplate* S in M is a branched surface equipped with a vector field X_S that is locally a union of finitely many ribbons, and is such that:

(i)   Two distinct ribbons $\text{Rib}_1, \text{Rib}_2$ of S can only intersect along their vertical edges, which are then called *branching segments*.

(ii) At every point on a branching segment, there are finitely many ribbons, and the associated vector fields all coincide.

(iii) For every ribbon Rib of S, the vector field X_{Rib} coincide with X_S on Rib.

An *orbit* of a multitemplate S is a complete immersion of the real line $\mathbb{R}$ is S that is everywhere tangent to X_S.

The difference with the usual notion of a template [4; 8] is that there is no uniquely defined semi-flow, but a multiflow. Indeed, at a point of a branching segment, there may be several escaping ribbons, and therefore several possible futures. If there were at most one escaping ribbon at every branching point, we would speak of a template. This will only happen in our construction when the starting tessellation consists of ideal polygons. Note also that there may be points that are visited by no orbit of the multitemplate, as for instance the points on the right of the branching segment on Figure 9.
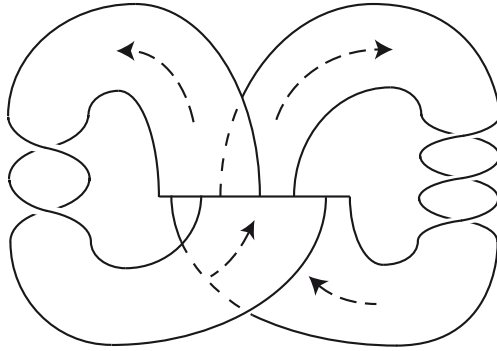
Figure 9: A multitemplate in $\mathbb{R}^3$. Along the branching segment, there are, from left to right, 1, then 2, then 1 and then 0 escaping ribbons.
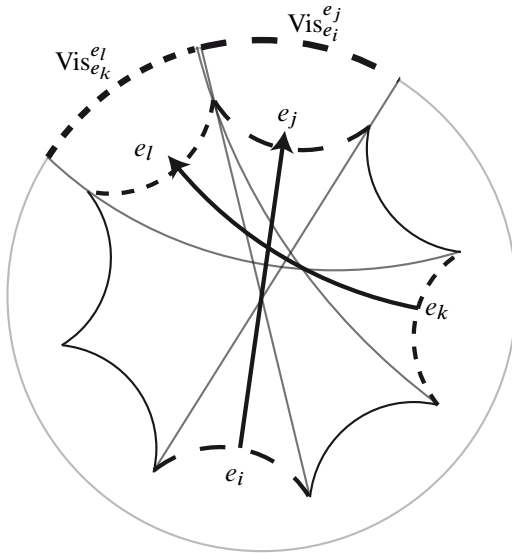

Figure 10: The visual intervals $\mathrm{Vis}^{e_j}_{e_i}$, $\mathrm{Vis}^{e_l}_{e_k}$ associated to two intersecting edges of $\mathcal{G}_{\mathcal{T}}$. Since $\mathrm{Vis}^{e_j}_{e_i}$ and $\mathrm{Vis}^{e_l}_{e_k}$ are disjoint, the associated ribbons $\mathrm{Rib}^{e_j}_{e_i}$, $\mathrm{Rib}^{e_l}_{e_k}$ do not intersect.

Let us go back to the construction. In order to specify the ribbons making the expected multitemplate, we describe the set of directions at infinity that are pointed in by elements in the image of a deformation $F^1_{\mathcal{T}}$.

Definition 4-7 (See Figure 10.) Assume that Γ is Fuchsian group of the first kind, that $\mathcal{T}$ is an adapted tessellation of $\mathbb{H}^2$, and that $\mathcal{G}_{\mathcal{T}}$ is a graph dual to $\mathcal{T}$. Let T be a

tile of $\mathcal{T}$, and e_i, e_j be two sides of T. Then the *visual interval* associated to (e_i, e_j) is the interval consisting of the positive extremities of geodesics connecting a point of e_i to a point e_j in $\partial_\infty \mathbb{H}^2$. We denote it by $\mathrm{Vis}_{e_i}^{e_j}$. The associated *product-ribbon* is the product of the oriented edge $c_{e_i}^{e_j}$ connecting e_i to e_j in $\mathcal{G}_\mathcal{T}$ by the interval $\mathrm{Vis}_{e_i}^{e_j}$ in $T^1 \mathbb{H}^2$, seen as the product $\mathbb{H}^2 \times \partial_\infty \mathbb{H}^2$. We denote it by $\mathrm{Rib}_{e_i}^{e_j}$. It is equipped with the horizontal vector field whose flow goes along the curves $c_{e_i}^{e_j} \times \{*\}$ at speed 1.

In the above context, we denote by $\widetilde{\mathcal{B}}_\mathcal{T}$ the union in $T^1 \mathbb{H}^2$ of the product-ribbons associated with all oriented edges of $\mathcal{G}_\mathcal{T}$. Its quotient under the action of Γ is denoted by $\mathcal{B}_{\Gamma,\mathcal{T}}$.

Lemma 4-8 *In the context of Definition 4-7, $\widetilde{\mathcal{B}}_\mathcal{T}$ is a multitemplate in $T^1 \mathbb{H}^2$.*

Proof By definition, the set $\widetilde{\mathcal{B}}_\mathcal{T}$ is the union of several ribbons, which are in one-to-one correspondence with the oriented edges of the graph $\mathcal{G}_\mathcal{T}$. Let T_1, T_2 be two adjacent tiles of $\mathcal{T}$. Call e the common side of T_1 and T_2, and let p be a vertex of $\mathcal{G}_\mathcal{T}$ lying on e. Since the tiles of $\mathcal{T}$ are supposed to have finitely many sides, there are finitely many ribbons that intersect the fiber $T^1\{p\}$ of p. Since all the edges of $\mathcal{G}_\mathcal{T}$ with p as an extremity are orthogonal to e, the associated product-ribbons all are tangent in $T^1\{p\}$. Now the product-ribbons that have an extremity in $T^1\{p\}$ decompose into four classes depending on whether they lie above T_1 or above T_2, and on whether they correspond to edges of $\mathcal{G}_\mathcal{T}$ starting at p or ending at p. Let (p, v) be a tangent vector based at p. Suppose that v points into T_1. Then the only ribbons that may contain (p, v) are those coming from geodesics with a positive extremity on the same side of e as T_2. In this case, the vector field on any such ribbon at (p, v) is the unit vector orthogonal to e, and pointing into T_1. Therefore the vector fields on all such ribbons coincide. Similarly, if v points into T_2, the vector fields of all ribbons that contain (p, v) are equal at (p, v) with the unit vector orthogonal to e and pointing into T_2.

It remains to show that product-ribbons are disjoint outside the fiber of the vertices of $\mathcal{G}_\mathcal{T}$. Since product-ribbons are in the fibers of edges of $\mathcal{G}_\mathcal{T}$, this is equivalent to showing that, if two edges $c_{e_i}^{e_j}, c_{e_k}^{e_l}$ of $\mathcal{G}_\mathcal{T}$ intersect inside a tile, say T, of $\mathcal{T}$, then the associated visual intervals $\mathrm{Vis}_{e_i}^{e_j}$ and $\mathrm{Vis}_{e_k}^{e_l}$ are disjoint. Indeed, in this situation, at the expense of possibly exchanging the indices and performing a symmetry, we can suppose that the edges e_i, e_k, e_j, e_l are cyclically ordered. Let $\underline{\gamma}_{i,j}^l$ be the geodesics joining the right extremity of e_i to the left extremity of e_j, and $\underline{\gamma}_{i,j}^r$ be the geodesics connecting the left extremity of e_i to the right extremity of e_j. Define $\gamma_{k,l}^l$ and $\gamma_{k,l}^r$ similarly. Then $\mathrm{Vis}_{e_i}^{e_j}$ is the interval $[(\gamma_{i,j}^r)_+, (\gamma_{i,j}^l)_+]$, and $\mathrm{Vis}_{e_k}^{e_l}$ is the interval $[(\gamma_{k,l}^r)_+, (\gamma_{k,l}^l)_+]$ (see Figure 10). The geodesics $\gamma_{i,j}^l$ and $\gamma_{k,l}^r$ intersect inside T, so that $(\gamma_{k,l}^r)_+$ lies on the left of $(\underline{\gamma}_{i,j}^l)_+$ on $\partial_\infty \mathbb{H}^2$. Therefore $\mathrm{Vis}_{e_i}^{e_j}$ and $\mathrm{Vis}_{e_k}^{e_l}$ are disjoint. $\square$
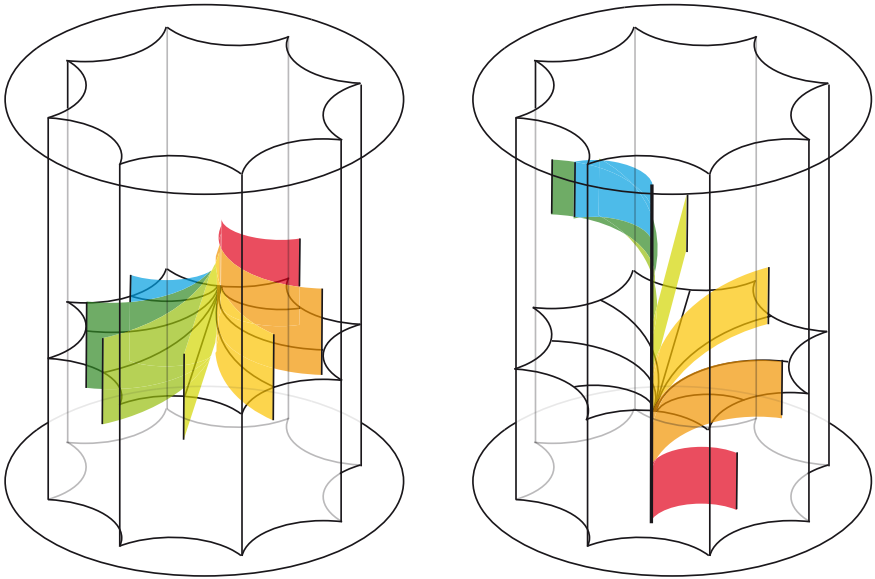
Figure 11: Some ribbons of a template above a tile that is not an ideal polygon. Both incoming ribbons (on the left) and outgoing ribbons (on the right) overlap, since the associated visual intervals overlap.

Assume now that the tiles of $\mathcal{T}$ all are ideal polygons. Let p be a vertex of $\mathcal{G}_\mathcal{T}$. Then all visual intervals associated with the edges of $\mathcal{G}_\mathcal{T}$ ending at p are disjoint. Hence, for every tangent vector v in the fiber $T^1\{p\}$, there is at most one escaping ribbon. Therefore $\widetilde{\mathcal{B}}_\mathcal{T}$ is a template, and so is $\mathcal{B}_{\Gamma,\mathcal{T}}$.

In the above context, since all steps in the construction of $\widetilde{\mathcal{B}}_\mathcal{T}$ are Γ–invariant, the quotient $\mathcal{B}_{\Gamma,\mathcal{T}}$ is also a multitemplate. We can now state the main result of this section.

Theorem 4-9 *Assume that Γ is Fuchsian group of the first kind, $\mathcal{T}$ is an adapted tessellation of $\mathbb{H}^2$, $\mathcal{G}_\mathcal{T}$ is a graph dual to $\mathcal{T}$, and a discretization of geodesics on $\mathcal{G}_\mathcal{T}$ has been chosen. Let $F_\mathcal{T}$ denote the associated tearing map of $T^1\mathbb{H}^2$, and $\mathcal{B}_{\Gamma,\mathcal{T}}$ denote the associated multitemplate in $T^1\mathbb{H}^2/\Gamma$. Then the action of $F_\mathcal{T}$ on $T^1\mathbb{H}/\Gamma$ induces an isotopy of every collection of periodic orbits of the geodesic flow $\Phi_{\mathbb{H}^2/\Gamma}$ onto a collection of periodic orbits of $\mathcal{B}_\mathcal{T}$. Moreover, if all tiles of $\mathcal{T}$ are ideal polygons, then $\mathcal{B}_{\Gamma,\mathcal{T}}$ is a template, and $F_\mathcal{T}$ is a one-to-one correspondence between the periodic orbits of $\Phi_{\mathbb{H}^2/\Gamma}$ and the periodic orbits of $\mathcal{B}_{\Gamma,\mathcal{T}}$ that do not lie in the boundary of $\mathcal{B}_{\Gamma,\mathcal{T}}$.*

Proof Let $\underline{\gamma}$ be a geodesic of $\mathbb{H}^2$. Then its discretization $\underline{\gamma}_\mathcal{T}$ is included in the graph $\mathcal{G}_\mathcal{T}$. Let p be point on $\underline{\gamma}$ and v be the tangent vector to $\underline{\gamma}$ at p. Then the
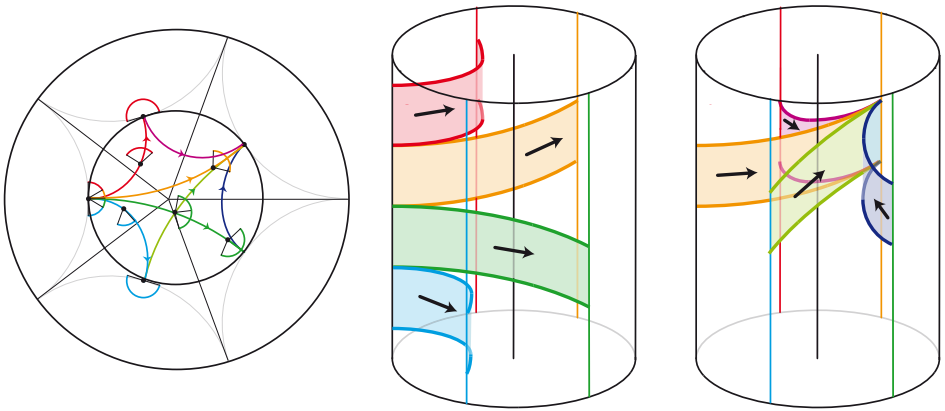
Figure 12: Some ribbons of a template above a tile that is an ideal polygon. The ribbons emerging from the same side do not overlap, since the associated visual intervals are disjoint.

vector $F_{\mathcal{T}}^1((p, v))$ lies in a fiber over $\mathcal{G}_{\mathcal{T}}$. By construction, for every edge $c_{e_i}^{e_j}$ of $\mathcal{G}_{\mathcal{T}}$ contained in $\gamma_{\mathcal{T}}$, the direction γ_+ belongs to the visual interval $\mathrm{Vis}_{e_i}^{e_j}$, so that the part of the curve $F_{\mathcal{T}}^1(\gamma)$ above $\overline{c}_{e_i}^{e_j}$ lies in the ribbon $\mathrm{Rib}_{e_i}^{e_j}$ and points towards γ_+. Therefore $F_{\mathcal{T}}^1(\gamma)$ sits in the multitemplate $\widetilde{\mathcal{B}}_{\mathcal{T}}$, and is everywhere tangent to the vector field $X_{\widetilde{\mathcal{B}}_{\mathcal{T}}}$. By Lemma 4-5, the restriction of $F_{\mathcal{T}}$ to Γ–invariant geodesics is an isotopy. Since everything commutes with the action of Γ, we can mod out by Γ, so that the projection of $F_{\mathcal{T}}$ realizes an isotopy between the periodic orbits of the geodesic flow $\Phi_{\mathbb{H}^2/\Gamma}$ and their images.

Suppose now that all tiles of $\mathcal{T}$ are ideal polygons. Let $\gamma(t)$ be an orbit of $X_{\widetilde{\mathcal{B}}_{\mathcal{T}}}$ not lying in the boundary of $\widetilde{\mathcal{B}}_{\mathcal{T}}$ and g–invariant for some g in Γ. Let γ_0 be its projection on $\mathbb{H}^2$. It is a g–invariant curve in $\mathcal{G}_{\mathcal{T}}$. Since all tiles of $\mathcal{T}$ are ideal polygons, γ_0 is a simple curve. The assumption that $\gamma(t)$ does not lie in the boundary of $\widetilde{\mathcal{B}}_{\mathcal{T}}$ implies that the two extremities of γ_0 are distinct. Let γ_1 be the unique geodesics in $\mathbb{H}$ connecting $(\gamma_0)_-$ to $(\gamma_0)_+$. Then γ_1 is also g–invariant. It turns out that γ_0 is then the discretization of γ_1. Therefore, $F_{\mathcal{T}}$ maps the vectors that are tangent to γ_1 to vectors that are tangent to $\gamma(t)$. $\qquad\square$

To conclude this section, we introduce some terminology that will be useful when Theorem 4-9 is applied in the sequel. Assume that $\mathcal{T}$ is a tessellation of the hyperbolic plane, and that $\mathcal{G}_{\mathcal{T}}$ is an associated graph. Let T_0 be a tile of $\mathcal{T}$ with n sides. Then the part of the template $\widetilde{\mathcal{B}}_{\mathcal{T}}$ that lies above T_0, that is, the intersection of $\widetilde{\mathcal{B}}_{\mathcal{T}}$ with $T^1 T_0$, consists of $n(n-1)$ ribbons, as depicted in Figures 11 and 12. In particular, there are n branching segments on which the template flow enters the solid torus $T^1 T_0$,
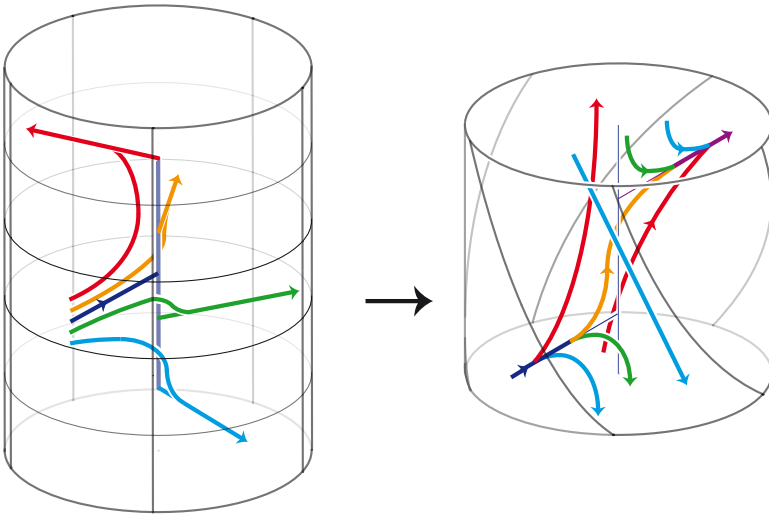
Figure 13: A simplified version of some ribbons of a template above some tile with an order 5 symmetry. On the left we displayed only 4 of the 20 ribbons, and only one orbit per ribbon; the other ribbons are obtained by iterating a screw-motion (remember Figure 1). On the right is the quotient of the unit tangent bundle by the order 5 symmetry, in the story model. This is what we call an *elevator*. For example the red ribbon on the left goes one floor up (that is, it goes in the next fundamental domain for the story model), so we see it crossing once the horizontal disc in the story model.

which we call *incoming segments*, and n branching segments where the template flow escapes $T^1 T_0$, which we call *outgoing segments*. We call such a part of a template a *switch tower*. If T_0 has a trivial stabilizer in Γ, then the part of the template $\mathcal{B}_{\Gamma,\mathcal{T}}$ above the quotient of T_0 by Γ is also a complete star.

Suppose now that T_0 has a non-trivial stabilizer, say Γ_{T_0}, in Γ. Then the part of $\mathcal{B}_{\Gamma,\mathcal{T}}$ above T_0/Γ is the quotient of $\widetilde{\mathcal{B}}_{\mathcal{T}}$ by Γ_{T_0}. If Γ_{T_0} has order d, then the part of the template has $n(n-1)/d$ ribbons. In particular, if T_0 is a regular n–gon and if its stabilizer Γ_{T_0} is of order n, then there are only $n-1$ ribbons in the quotient, all of them joining a unique incoming segment to a unique outgoing segment, see Figure 13. We call such a part of a template an *elevator*.

5 Geodesic flow for the orbifolds of type (p, q, ∞)

We now turn to the linking properties of orbits associated with hyperbolic orbifolds of type $(2, q, \infty)$ with $q \geq 3$. The goal of this section is to prove the first case of

Theorem A, that is, to prove that the linking number of every two orbits of $\Phi_{\Sigma_{2,q,\infty}}$ is negative.

The idea is to apply the construction of Section 4, thus obtaining a template that describes the topology of the periodic orbits of $\Phi_{\Sigma_{2,q,\infty}}$ (Section 5.1), and then to compute the linking number of a pair of periodic orbits. Actually, we do more and first compactify the unit tangent bundle into a lens space (Section 5.2). As a lens space is a rational homology sphere, the linking number is defined for every pair of links. We then show that the linking number of every pair of periodic orbits of the template is negative (Proposition 5-9 and case (a) of Theorem A). By the way, we consider a slightly more general context and construct a template for every orbifold of type (p, q, ∞) with $p \geq 2$. The advantage of this approach is to also provide a precise formula for the linking number of a periodic orbit of $\Phi_{\Sigma_{p,q,\infty}}$ with the fiber of the cusp in the unit tangent bundle, that is, with the link that has been added for the compactification (Proposition 5-7).

5.1 A template for $\Phi_{\Sigma_{p,q,\infty}}$

Here we introduce orbifolds of type (p, q, ∞), choose adapted tessellations of the hyperbolic plane, and describe the associated templates. As we will recall, the space $T^1\Sigma_{p,q,\infty}$ is obtained by gluing two solid tori along their boundary, and what we will do is to describe a template that lies in a neighborhood of the gluing torus. In the case $p = 2, q = 3$, we recover Ghys' template for the geodesic flow on the modular surface [9], and, in the more general case $p = 2, q \geq 3$, we recover Pinsky's template [14].

Until the end of Section 5, we assume that p, q are fixed integers satisfying $p \geq 2$ and $q \geq 3$. Since $1/p + 1/q < 1$ holds, there exists a hyperbolic triangle PQZ in $\mathbb{H}^2$, with the two vertices P, Q inside $\mathbb{H}^2$ with respective angles $2\pi/p$ and $2\pi/q$, and the vertex Z lying on $\partial\mathbb{H}^2$. For convenience, we also suppose P, Q, Z are trigonometrically ordered. Let $\Gamma_{p,q}^*$ be the group generated by the symmetries around the sides of PQZ, and let $\Gamma_{p,q}$ be its index 2 subgroup consisting of orientation preserving isometries, often called the *Hecke triangular group*. The group $\Gamma_{p,q}$ acts properly and discontinuously on $\mathbb{H}^2$. The action is not free since, for example, P and Q have stabilizers of order p and q respectively. The quotient $\mathbb{H}^2/\Gamma_{p,q}$ is then an orbifold, with two singular points of order p and q, and one cusp. We call it $\Sigma_{p,q,\infty}$.

For $k = 1, \ldots, q-1$, let Z_q^k be the image of Z by a rotation of center Q and of angle $2k\pi/q$ (see Figure 15). Then $Z, Z_q^1, \ldots, Z_q^{q-1}$ are the vertices of an ideal q–gon, say Δ_Q. Let Γ_Q be the stabilizer of Q in $\Gamma_{p,q}$. Then Δ_Q is invariant under the action of Γ_Q.
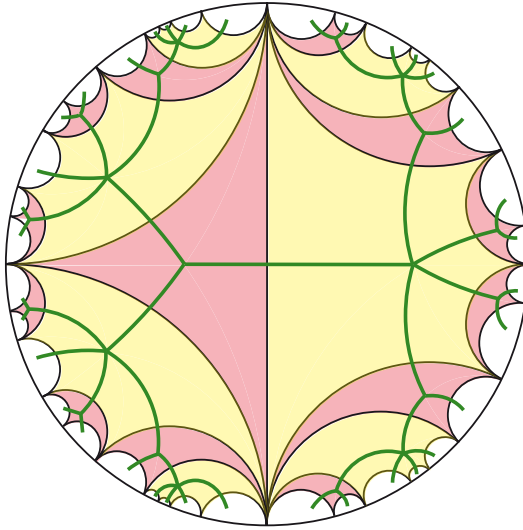
Figure 14: The tessellation of $\mathbb{H}^2$ by copies of the ideal polygons Δ_P and Δ_Q, here with $p = 3$ and $q = 5$

Assume now $p > 2$. Define similarly the points $Z_p^1, \ldots, Z_p^{p-1}$ on $\partial_\infty \mathbb{H}^2$ and the polygon Δ_P. Note that the points Z_q^1 and Z_p^{p-1} coincide. Call e the geodesics ZZ_q^1. Then the polygons Δ_P, Δ_Q lie on different sides of e, hence they are distinct. One easily sees that the images of Δ_P and Δ_Q under $\Gamma_{p,q}$ cover the whole hyperbolic plane, and therefore form a tessellation (Figure 14). We denote it by $\mathcal{T}_{P,Q}$. The sides of the tiles of $\mathcal{T}_{P,Q}$ exactly are the images of e under $\Gamma_{p,q}$. Since all tiles are ideal polygons, no two sides in the tessellation intersect inside $\mathbb{H}^2$. Also, every tile is a copy of either Δ_P or Δ_Q, and therefore contains exactly one singular point in its interior.

The unit tangent bundles to Δ_Q/Γ_Q and Δ_P/Γ_P are both non-compact solid tori (remember Figure 1). The unit tangent bundle $T^1\Sigma_{p,q,\infty}$ is then obtained by identifying the tangent vectors that constitute the boundaries of the unit tangent bundles to Δ_Q/Γ_Q and Δ_P/Γ_P. These are exactly the images in the quotient of the tangent vectors based on e, that is, the image of T^1e in the orbifold $\Sigma_{p,q,\infty}$.

We still assume $p > 2$. Let M be the intersection of the segment PQ with e, let $\mathrm{BS}_{P \to Q}$ be the set of all tangent vectors at M pointing into Δ_Q, and let $\mathrm{BS}_{Q \to P}$ be the set of all tangent vectors at M pointing into Δ_P. Then the template $\mathcal{B}_{\Gamma_{p,q}, \mathcal{T}_{P,Q}}$ given by Theorem 4-9 consists of two parts: one elevator (Figure 13) sitting inside the solid torus $T^1\Delta_Q/\Gamma_Q$ with $q-1$ ribbons, say $\mathrm{Rib}_q^1, \ldots, \mathrm{Rib}_q^{q-1}$, all connecting $\mathrm{BS}_{P \to Q}$ to $\mathrm{BS}_{Q \to P}$, and one elevator sitting inside $T^1\Delta_P/\Gamma_P$ with $p-1$ ribbons, say $\mathrm{Rib}_p^1, \ldots, \mathrm{Rib}_p^{p-1}$, all connecting $\mathrm{BS}_{Q \to P}$ to $\mathrm{BS}_{P \to Q}$.
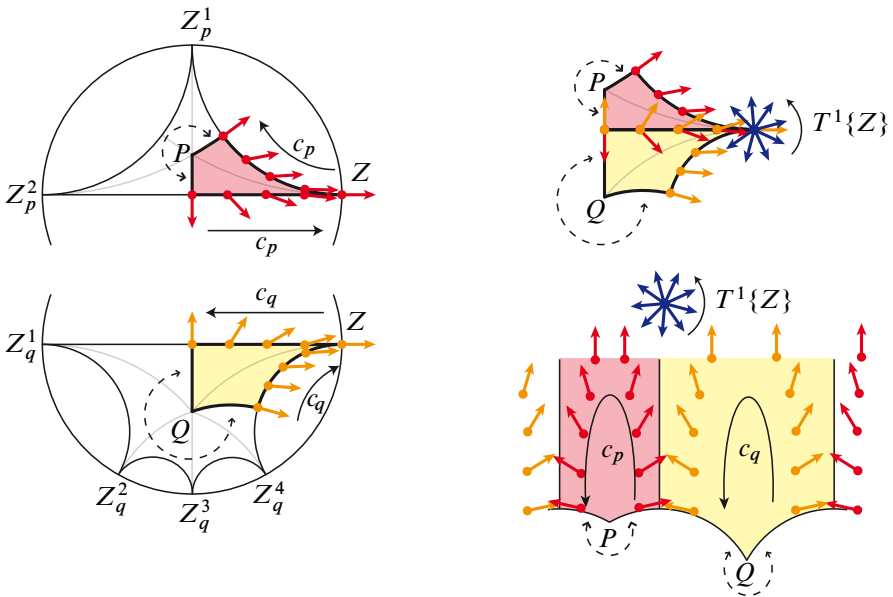
Figure 15: On the left are the tiles Δ_P and Δ_Q, with $p = 3$ and $q = 5$. On the right is a fundamental domain for the action of $\Gamma_{p,q}$ on $\mathbb{H}^2$. The curves c_p, c_q and $T^1\{Z\}$ are also depicted. They lie on the common boundary $\overline{T^1 e}$ of the two solid tori $\overline{T^1\Delta_P}/\Gamma_P$ and $\overline{T^1\Delta_Q}/\Gamma_Q$.

If $p = 2$, then, with the above definition, Δ_P is a bigon with an empty interior. In this case, the tessellation $\mathcal{T}_{P,Q}$ consists of copies of Δ_Q only. In the quotient of $\mathbb{H}^2$ by $\Gamma_{p,q}$, the edges of Δ_Q are quotiented by order 2 rotations, so that the unit tangent bundle of $\Sigma_{2,q,\infty}$ is obtained by considering the solid torus $T^1\Delta_Q/\Gamma_Q$, and identifying pairs of points on the boundary with the order 2 rotation around P.

For convenience (especially in view of the pictures in Section 5.4), we slightly modify the tessellation in this case. We consider a tile Δ'_P, that is the ϵ–neighborhood of Δ_P, and we change Δ_Q accordingly. If ϵ is positive, the sides of the tiles are no longer geodesic, so that the construction of Section 4 does not apply. We rather see ϵ as infinitely small. The unit tangent bundle of $\Sigma_{2,q,\infty}$ is then the union of $T^1\Delta'_Q/\Gamma_Q$, which is infinitesimally smaller than $T^1\Delta_Q/\Gamma_Q$, with the infinitesimally small solid torus $T^1\Delta'_P/\Gamma_P$. The role of the latter solid torus is to identify pairs of points on the boundary of $T^1\Delta'_Q/\Gamma_Q$.

Mimicking the case $p > 2$, we denote by M the point on the segment $[PQ]$ that is at distance ϵ from P, by $\mathrm{BS}_{P \to Q}$ the set of all tangent vectors at M pointing into Δ'_Q and by $\mathrm{BS}_{Q \to P}$ the set of all tangent vectors at M pointing into Δ'_P. Then $\mathcal{B}_{\Gamma_{2,q}, \mathcal{T}_{P,Q}}$ consists of one elevator in $T^1\Delta'_Q/\Gamma_Q$ with $q - 1$ ribbons, say $\mathrm{Rib}^1_q, \ldots, \mathrm{Rib}^{q-1}_q$, all

connecting $\mathrm{BS}_{P \to Q}$ to $\mathrm{BS}_{Q \to P}$, and one ribbon in $T^1 \Delta'_P / \Gamma_P$ connecting $\mathrm{BS}_{Q \to P}$ to $\mathrm{BS}_{P \to Q}$.

In the sequel, it will be imported to visualize how the ribbons $\mathrm{Rib}_p^1, \ldots, \mathrm{Rib}_p^{p-1}$ can be distorted on the torus $\partial T^1 \Delta_P / \Gamma_P = \partial T^1 \Delta_Q / \Gamma_Q$ (and similarly for $\mathrm{Rib}_q^1, \ldots, \mathrm{Rib}_q^{q-1}$). Figure 16 shows two ways of deforming every such ribbon by pushing it to the left or to the right.

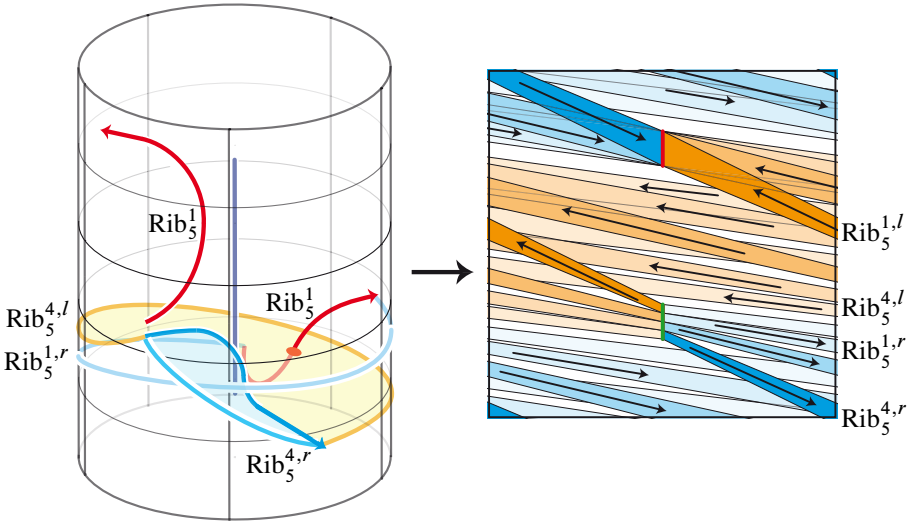Figure 16: The two possible deformations $\mathrm{Rib}_p^{i,l}$ and $\mathrm{Rib}_p^{i,r}$ of a ribbon of type Rib_p^i on $\partial T^1 \Delta_P / \Gamma_P$. On the left, with $p = 5$, the ribbon Rib_p^{p-1} (in blue) can be pushed in $T^1 \Delta_P$ either to the right (in light blue) or to the left (in orange). Another ribbon (here Rib_p^1 in red) and one of its images under the $\mathbb{Z}/p\mathbb{Z}$–action are shown. Since the rightmost ribbon (here the blue one) goes the lowest, when distorting the ribbons on $\partial T^1 \Delta_P$, all crossings that appear are positive. On the right, the two projections of all ribbons $\mathrm{Rib}_p^1, \ldots, \mathrm{Rib}_p^{p-1}$ on $\partial T^1 \Delta_P / \Gamma_P$, seen in the slice-of-cake model (the one whose fundamental domain is the space located between two vertical walls). The blue ribbons are obtained when pushing to the right, and the orange ones when pushing to the left.

5.2 Compactification and coordinates

The unit tangent bundle $T^1 \Sigma_{p,q,\infty}$ is a non-compact 3–manifold with first homology group $\mathbb{Z}$. This can be seen in the previous discussion by considering a loop of tangent vectors based along a horocycle centered at Z, and checking that this loop is not

null-homologous. For addressing Question 1-1, we want to compute linking numbers in $T^1\Sigma_{p,q,\infty}$.

As stated in the introduction, we will make a more general computation by first compactifying $T^1\Sigma_{p,q,\infty}$ into a rational homology sphere, and then compute linking in the resulting manifold. Since $\Sigma_{p,q,\infty}$ has one cusp, a natural compactification that does not change the homology type consists in adding a boundary-circle. For $T^1\Sigma_{p,q,\infty}$, this corresponds to the addition of a boundary-torus. As we want a compactification with trivial first rational homology group, we need to fill this torus. A natural choice is to fill the boundary-circle with a disc and to lift this filling. But this choice is not appropriate for the hyperbolic structure, and a more adapted choice is to force all vectors tangent to a given horocycle to bound a disc in the compactification. These two compactifications are defined according to whether we see the boundary circle as a hole or a cusp. Actually, there is one filling of the torus, and therefore one compactification by a circle for every choice of a Euler number, thus leading to a fiber bundle with the chosen Euler class (as explained by Pinsky [14]). The hole-like filling corresponds to Euler number 0, while the cusp-like filling has Euler number -1. This leads to the following:

Definition 5-1 The *hyperbolic compactification* $\overline{T^1\Sigma_{p,q,\infty}}$ of $T^1\Sigma_{p,q,\infty}$ is obtained by adding a fiber associated to the cusp Z, that is, by considering the topology induced by the compactification of $\Sigma_{p,q,\infty}$ in the hyperbolic disc.

The compactification $\overline{T^1\Sigma_{p,q,\infty}}$ is obtained by gluing the two solid tori $\overline{T^1\Delta_Q/\Gamma_Q}$ and $\overline{T^1\Delta_P/\Gamma_P}$ (with Δ' instead of Δ in the case $p=2$) along their boundaries. It is then a lens space. In order to describe it, let us introduce some notation (see Figure 15). We write $\overline{T^1e}$ for the 2–torus that is the boundary between $\overline{T^1\Delta_Q/\Gamma_Q}$ and $\overline{T^1\Delta_P/\Gamma_P}$. We define a_Z to be the loop in $\overline{T^1e}$ describing the fiber $T^1\{Z\}$ with the trigonometric orientation, and c_P to be the curve consisting of tangent vector based on $\bar{e}$ and oriented by the geodesics going through P. We define c_Q in the same way. We also consider the set $\hat{D}_P$ of all vectors based on points of Δ_P and pointing in the direction of Z, and its quotient D_P under the projection $T^1\Delta_P \to T^1\Delta_P/\Gamma_P$ with the induced orientation (see Figure 17). We write ∂D_P for the oriented boundary of D_P. We define D_Q and ∂D_Q in the same way.

Lemma 5-2 (i) The set D_P is a meridian disc of the solid torus $\overline{T^1\Delta_P/\Gamma_P}$.

(ii) The homology classes $[c_P]$ and $[c_Q]$ form a basis of $H_1(\overline{T^1e}; \mathbb{Z})$. In this basis, we have the decompositions $[a_Z] = (1,1)$, $[\partial D_P] = (p-1,-1)$ and $[\partial D_Q] = (-1, q-1)$.

Proof (i) The disc $\widehat{D}_P$ is contractible in $\overline{T^1 \Delta_P}$ and its boundary belongs to the boundary $\overline{T^1(\partial \Delta_P)}$. Therefore its quotient D_P is also contractible in $\overline{T^1 \Delta_P / \Gamma_P}$, and its boundary ∂D_P belongs to the boundary $\overline{T^1(\partial \Delta_P / \Gamma_P)}$, which is, by definition, the 2–torus $\overline{T^1 e}$. The loop ∂D_P is not contractible in $\overline{T^1 e}$ because its projection on the basis is not. Therefore D_P is a meridian disc in $\overline{T^1 e}$.
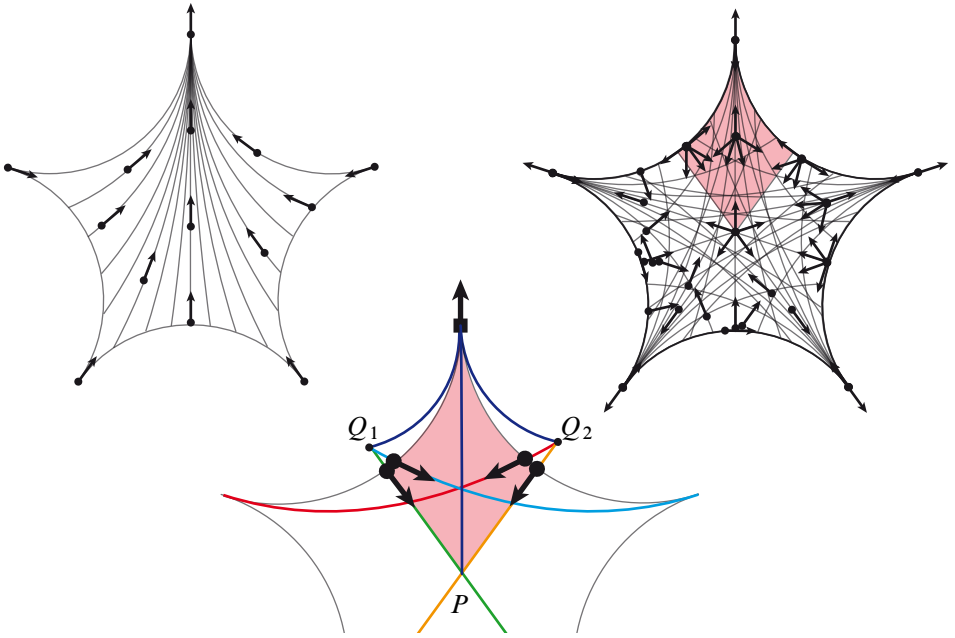

Figure 17: On the left is the meridian disc $\widehat{D}_P$, with $p = 5$. It is the set of all vectors pointing at Z. On the right is the union of its iterated images under the rotation of angle $2\pi/p$. It is the set of all vectors pointing at one of the p vertices of Δ_P. The meridian disc D_P (Lemma 5-2) is obtained by restricting to a fundamental domain, for example the tinted part. On the bottom, the $p-1$ intersection points between c_Q and ∂D_P (the leftmost and the rightmost vectors are identified in the quotient). The unique intersection point between c_P and ∂D_P is the vector based at Z with a squared origin.

(ii) We write · for the intersection form on the torus $H_1(\overline{T^1 e}; \mathbb{Z})$. By definition, and as indicated on Figure 15, the three curves c_P, c_Q and a_Z have one point in common, namely the unit tangent vector based at Z and oriented by outgoing geodesics. Therefore, we have

$$|[c_P] \cdot [c_Q]| = |[a_Z] \cdot [c_P]| = |[a_Z] \cdot [c_Q]| = 1,$$

so that the classes $[c_P]$ and $[c_Q]$ form a basis of $H_1(\overline{T^1 e}; \mathbb{Z})$. We orient the 2–torus $\overline{T^1 e}$ in such a way that the basis $([c_P], [c_Q])$ is positive. The signs of the intersections $[a_Z] \cdot [c_P]$ and $[a_Z] \cdot [c_P]$ can be determined by checking that the concatenation of the loops c_P and c_Q is homotopic to a_Z, so that, in the basis $([c_P], [c_Q])$, we have $[a_Z] = (1, 1)$.

In order to determine the coordinates of $[\partial D_P]$, we compute the intersection numbers with the basis vectors. For the intersection between $[\partial D_P]$ and $[c_P]$, we see on Figure 17 that there is only one vector in $[\partial D_P] \cap [c_P]$, namely the vector based at Z and oriented by outgoing geodesics. For the intersection between $[\partial D_P]$ and $[c_Q]$, we have to count the vectors emerging from Q and pointing into one of the vertices Z_p^k. There are $p - 1$ such vectors, depicted on Figure 17. Once again, the signs can be determined by checking that the loop ∂D_P is isotopic to the concatenation of p times c_P and one time a_Z, taken backwards, whence the relation $[\partial D_P] = (p - 1, -1)$.

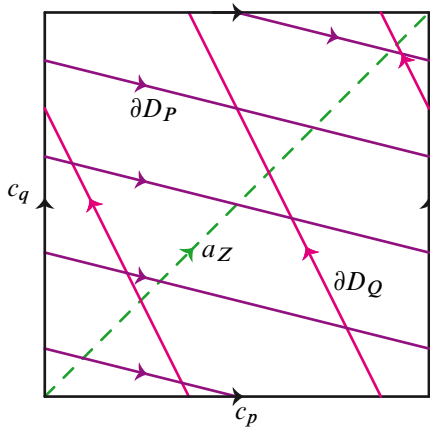The coordinates of $[\partial D_Q]$ are determined in the same way. $\qquad \square$

Figure 18: The median torus $\overline{T^1 e}$, in the basis (c_p, c_q), is shown here with $p = 5$ and $q = 3$. The hyperbolic compactification $\overline{T^1 \Sigma_{p,q,\infty}}$ is obtained by gluing two solid tori $\overline{T^1 \Delta_P / \Gamma_P}$ and $\overline{T^1 \Delta_Q / \Gamma_Q}$, with respective meridian ∂D_P and ∂D_Q, along $\overline{T^1 e}$. The fibers of the points of e are the curves on $\overline{T^1 e}$ that are parallel to a_Z.

We can now deduce the topology of $T^1 \Sigma_{p,q,\infty}$.

Lemma 5-3 *The hyperbolic compactification $\overline{T^1 \Sigma_{p,q,\infty}}$ of the unit tangent bundle to the orbifold $\Sigma_{p,q,\infty}$ is diffeomorphic to the lens space $L_{pq-p-q,p-1}$, the circle added when compactifying being a (p, q)–torus knot drawn on a median torus of $L_{pq-p-q,p-1}$.*

Proof (See Figure 18.) We continue with the same notation. Since $\overline{T^1\Sigma_{p,q,\infty}}$ is obtained by gluing the two solid tori $\overline{T^1\Delta_P}/\Gamma_P$ and $\overline{T^1\Delta_Q}/\Gamma_Q$, it is a lens space. By Lemma 5-2, the two curves $\partial D_P, \partial D_Q$ are respective meridians in the two solid tori. Using their coordinates, we deduce that their intersection number is

$$\begin{vmatrix} p-1 & -1 \\ -1 & q-1 \end{vmatrix} = pq - p - q.$$

As the curve c_Q intersects ∂D_Q once, it is a parallel for the solid torus $\overline{T^1\Delta_Q}$. As c_Q intersects $p-1$ times ∂D_P, the 3–manifold $\overline{T^1\Sigma_{p,q,\infty}}$ is the lens space $L_{pq-p-q,p-1}$.

The circle that has been added when compactifying is the fiber a_Z of the point Z. By Lemma 5-2, it intersects p times the circle ∂D_P, and q times ∂D_Q. Therefore it is a (p,q)–torus knot. $\square$

Remark 5-4 Since $(p-1)(q-1) \equiv 1 \mod pq - p - q$, Brody's theorem asserts that the lens spaces $L_{pq-p-q,p-1}$ and $L_{pq-p-q,q-1}$ are diffeomorphic. This can be seen in the above proof by exchanging p and q.

Remark 5-5 One can check that the alternative compactifications of $T^1\Sigma_{p,q,\infty}$ associated with other Euler numbers can be obtained by cutting along T^1e, making a transvection along the curve a_Z, and gluing back. This changes the manifold into $L_{kpq-p-q,kp-1}$ for some k in $\mathbb{Z}$ (see Pinsky [14] for more detail).

We now have a full description of the template $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ and of how it embeds into $\overline{T^1\Sigma_{p,q,\infty}}$ (see also Figure 19). It is worth noting that in the case $p=2, q=3$, the compactification $\overline{T^1\Sigma_{p,q,\infty}}$ is the 3–sphere, the fiber a_Z of the cusp is a trefoil knot, and the template $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ is Lorenz' template, as stated by Ghys [9].

5.3 Linking with the fiber of the cusp

For $p \geq 2, q \geq 3$, we use now the template $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ for computing the linking number in $\overline{T^1\Sigma_{p,q,\infty}}$ between a periodic orbit of $\Phi_{\Sigma_{p,q,\infty}}$ and the (p,q)–torus knot a_Z that has been added when compactifying $T^1\Sigma_{p,q,\infty}$ (Proposition 5-7). This computation has been done in the case $p=2$, $q=3$ by Ghys [9]. In this case, the linking number equals the Rademacher function of the underlying geodesics; a function of interest in number theory [13]. As before, we assume that we are given a triangle PQZ in $\mathbb{H}^2$, that $\Gamma_{p,q}$ is the associated Hecke triangular group, that $\mathcal{T}_{P,Q}$ is the associated adapted tessellation of $\mathbb{H}^2$, and that $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ is the associated template.

Let γ be a geodesic of $\mathbb{H}$ whose extremities are not lifts of the cusp of $\Sigma_{p,q,\infty}$, that is, γ_+ and γ_- are not in the orbit $\Gamma_{p,q}(Z)$. Then picking an arbitrary starting point on it,
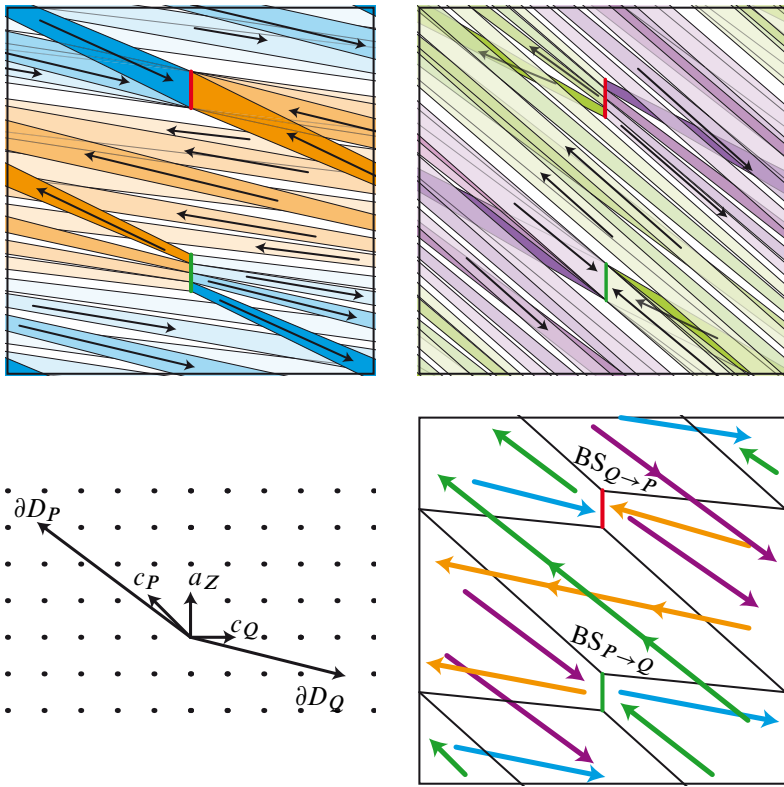
Figure 19: The projection of the template $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ on the 2–torus $\overline{T^1 e}$ is shown here with $p = q = 4$. The sources of the projection are the fibers $T^1\{P\}$ for the part that lies inside the solid torus $\overline{T^1 \Delta_P / \Gamma_P}$, and $T^1\{Q\}$ for what lies inside $\overline{T^1 \Delta_Q / \Gamma_Q}$. On the top left: the part of $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ lying inside $\overline{T^1 \Delta_Q / \Gamma_Q}$. On the top right: the part of $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ lying inside $\overline{T^1 \Delta_P / \Gamma_P}$. The two pictures differ by a transvection. This is due to the choice of the compactification. Changing the compactification of $T^1 \Sigma_{p,q,\infty}$ leads to another transvection for the identification. Since the two solid tori are glued outgoing normal vs. incoming normal, the two pictures have opposite orientations, namely the front/back order of the ribbons is reversed. On the bottom left are the vectors $[c_P], [c_Q], [a_Z], [\partial D_P]$ and $[\partial D_Q]$ in $H_1(\overline{T^1 e}; \mathbb{Z})$. The slope -1 of $[c_P]$ explains the transvection on the top right picture. On the bottom right: The directions of the two possible deformations of the ribbons that constitute $\mathcal{B}_{\Gamma_{p,q},\mathcal{T}_{P,Q}}$ on $\overline{T^1 e}$. Each of the four colors corresponds to one to the four types $\mathrm{Rib}_p^{i,g}$, $\mathrm{Rib}_p^{i,d}$, $\mathrm{Rib}_q^{i,g}$ and $\mathrm{Rib}_q^{i,d}$. The key point for proving the negativity of linking numbers (Proposition 5-9) is that, in each of the two vertical intervals between $\mathrm{BS}_{P \to Q}$ and $\mathrm{BS}_{Q \to P}$, all ribbons go in the same direction.

γ is determined by a starting tile T_0 and a bi-infinite code $\cdots u^{i-1} v^{j-1} u^{i_0} v^{j_0} u^{i_1} v^{j_1} \cdots$ describing how γ behaves in each tile of the tessellation $\mathcal{T}_{P,Q}$. Precisely, if γ enters a copy of Δ_P by a side, and goes out by another side that is obtained from the entering one by a rotation of angle $2i\pi/p$, then the corresponding letter is u^i. Similarly, when γ enters a copy of Δ_Q, the corresponding letter v^j describes how to pass from the entering side to the outgoing side. As Δ_P has p sides, every index i_k is between 1 and $p-1$. Similarly, every index j_k lies between 1 and $q-1$. Considering another starting tile induces a shift of the code. If two geodesics are obtained one from the other by the action of an element g of $\Gamma_{p,q}$, then their starting tiles are also obtained from one another by g, and their codes coincide. Therefore, there is a one-to-one correspondence between codes up to shift and geodesics on the orbifold $\Sigma_{p,q,\infty}$ not pointing into the cusp. Moreover, if a geodesic on $\Sigma_{p,q,\infty}$ is periodic, then its code is periodic, that is, of the form $(u^{i_1} \cdots v^{j_m})^{\mathbb{Z}}$. In this case, we call the word $u^{i_1} \cdots v^{j_m}$, which is assumed to be of minimal possible length, a *reduced code* of the periodic geodesic. Different reduced codes for a given periodic geodesic differ by a cyclic permutation of the letters.

We now define an invariant of periodic geodesics that will be useful for expressing the linking number of their liftings in $T^1 \Sigma_{p,q,\infty}$ with the fiber of the cusp. Assume that γ is a geodesic in $\mathbb{H}^2$ with code $\cdots u^{i-1} v^{j-1} u^{i_0} v^{j_0} u^{i_1} v^{j_1} \cdots$. For a more symmetric expression, we set $i'_k = i_k - p/2$ and $j'_k = j_k - q/2$. Then the discretization $\gamma_{\mathcal{T}_{P,Q}}$ of γ lies in the tree depicted in Figure 14. By definition, for every index k, the discretization $\gamma_{\mathcal{T}_{P,Q}}$ turns by an angle $2\pi i'_k/p$ in the corresponding copy of Δ_P and by an angle $2\pi j'_k/q$ in the corresponding copy of Δ_Q.

Definition 5-6 Assume that $\underline{\gamma}$ is a periodic geodesic on $\Sigma_{p,q,\infty}$. Let $u^{i_1} v^{j_1} u^{i_2} \cdots v^{j_m}$ be a reduced code of $\underline{\gamma}$. Then the *wheel turn* $\Theta_{\text{wheel}}(\underline{\gamma})$ of $\underline{\gamma}$ is the rational number $\sum_{i=1}^{m} i'_k/p + j'_k/q$.

Here is the expected evaluation of the linking number between a geodesic of $\Phi_{\Sigma_{p,q,\infty}}$ and the fiber of the cusp in terms of an analog of the Rademacher function.

Proposition 5-7 *Assume $p \geq 2$, $q \geq 3$. Then, for every periodic orbit γ of the geodesic flow $\Phi_{\Sigma_{p,q,\infty}}$, we have*

$$\text{Lk}(\gamma, a_Z) = \frac{pq}{pq - p - q} \Theta_{\text{wheel}}(\underline{\gamma}),$$

where $\underline{\gamma}$ is the projection of γ on $\overline{\Sigma_{p,q,\infty}}$.

The principle of the proof is as follows. Write r for the number $pq - p - q$. Since the first homology group of $L_{pq-p-q,p-1}$ is $\mathbb{Z}/r\mathbb{Z}$, we know that for every element $[c]$

of $H_1(L_{pq-p-q,p-1}; \mathbb{Z})$, the cycle $r[c]$ is a boundary of an integral 2–chain. The idea will be to construct a 2–chain with boundary $r[\gamma]$ that is transverse to a_Z, and then to count the intersection number with a_Z. Since γ is isotopic in the complement of a_Z in an orbit of the template $\mathcal{B}_{\Gamma_{p,q}, T_{P,Q}}$, we can then make use of the available information about the position of the latter in $T^1 \Sigma_{p,q,\infty}$.

In order to implement the argument, let us write h for the orbit of $\mathcal{B}_{\Gamma_{p,q}, T_{P,Q}}$ whose code is $(u^1 v^1)^{\mathbb{Z}}$. Note that h is one of the two periodic orbits of $\mathcal{B}_{\Gamma_{p,q}, T_{P,Q}}$ that is not isotopic to a periodic orbit of the geodesic flow, but to a periodic orbit of the horocyclic flow. Write a_P for the curve that describes the fiber $T^1 P$. It is the core of the solid torus $\overline{T^1 \Delta_P / \Gamma_P}$. Similarly, write a_Q for curve describing the fiber $T^1 Q$. We begin with a preliminary computation. Remember that $u^{i_1} v^{j_1} u^{i_2} \cdots v^{j_m}$ denotes a reduced code of γ.

Lemma 5-8 *In the above context, the cycle $[\gamma]$ is homologous in $T^1 \Sigma_{p,q,\infty}$ to the 1–cycle*

$$\sum_{k=1}^m \big([h] + (i_k - 1)[a_P] + (j_k - 1)[a_Q]\big).$$

Proof Let γ^1 be the image of γ under the deformation $F^1_{T_{P,Q}}$. Then γ^1 is an orbit of the template $\mathcal{B}_{\Gamma_{p,q}, T_{P,Q}}$. Suppose that γ first travels along the ribbon Rib^1_p, and then along Rib^1_q. Then it is homologous to h in $T^1 \Sigma_{p,q,\infty}$ during the corresponding interval of time, and its code starts with $u^1 v^1$. Otherwise, the homology class of γ in the complement of a_Z during one period is obtained by adding to h the cycles consisting in traveling along Rib^i_p backwards and then along Rib^{i+1}_p frontwards, for every i between 1 and i_k, and by also adding the cycles consisting in traveling along Rib^j_q backwards and then along Rib^{j+1}_q forwards, for every j between 1 and j_k. Every cycle in the first category is actually equal to $[a_P]$. Indeed, the ribbons are not the same, but the annuli of the form $\mathrm{Rib}^{j+1}_p - \mathrm{Rib}^j_p$ are homologous in the quotient $\overline{T^1 \Sigma_{p,q,\infty}}$: they correspond to curves turning once around the point P on $\Sigma_{p,q,\infty}$. So $[a_P]$ is added $i_k - 1$ times. Similarly, every cycle in the second category is equal to $[a_Q]$, so $[a_Q]$ is added $j_k - 1$ times. □

We can then complete the argument.

Proof of Proposition 5-7 In $\overline{T^1 \Sigma_{p,q,\infty}}$, the cycle h bounds a disc whose intersection number with a_Z equals -1. Indeed, since h is homologous to a horocyle, the latter bounds a horodisc, say d_h, which is foliated by horocycles parallel to h. By definition of the compactification, the family of all vectors tangent to these horocycles extends to the

fiber of the cups, and therefore form a disc in $\overline{T^1\Sigma_{p,q,\infty}}$ that intersects a_Z in exactly one point, namely the limit of the tangent vectors. Hence we have $\mathrm{Lk}(h, a_Z) = -1/r$.

Let us turn to $\mathrm{Lk}(h, a_Z)$. Write

$$c_P^+, \quad \partial D_p^+ \quad \text{and} \quad \partial D_q^+$$

for the curves c_P, ∂D_p and ∂D_q slightly pushed away from $\overline{T^1 e}$ in $\overline{T^1\Delta_P / \Gamma_P}$, so that they do not intersect a_Z. As c_Q is a parallel for the solid torus $\overline{T^1\Delta_P / \Gamma_P}$, the cycle $r[a_P]$ is homologous in $T^1\Sigma_{p,q,\infty}$ to $r[c_P^+]$. The latter has coordinates

$$(pq - p - q, 0)$$

in the basis $([c_P], [c_Q])$ of $H_1(\overline{T^1 e}; \mathbb{Z})$. By Lemma 5-2, the cycles $[\partial D_p^+]$ and $[\partial D_q^+]$ have coordinates $(p-1, -1)$ and $(-1, q-1)$, so that $r[c_P^+]$ equals

$$(1-q)[\partial D_p^+] - [\partial D_q^+].$$

On the one hand, $[\partial D_p^+]$ bounds a meridian disc for $\overline{T^1\Delta_P / \Gamma_P}$ that does not intersect a_Z. On the other hand, since ∂D_q and a_Z intersect q times on $\overline{T^1 e}$, the curve ∂D_q^+ bounds a meridian disc for

$$\overline{T^1\Delta_Q / \Gamma_Q}$$

that intersects $-q$ times a_Z. Therefore we have $\mathrm{Lk}(a_P, a_Z) = q/r$. We obtain in the same way $\mathrm{Lk}(a_Q, a_Z) = p/r$. The result then follows from Lemma 5-8. $\quad\square$

5.4 Linking number between collections of geodesics

We now restrict to the case $p = 2$, and study the linking number between two collections of periodic geodesics of $\Phi_{\Sigma_{2,q,\infty}}$. Our goal is to show:

Proposition 5-9 (Case (a) of Theorem A) *Assume $q \geq 3$. Then, for all collections of periodic orbits γ, γ' of the geodesic flow $\Phi_{\Sigma_{2,q,\infty}}$ in $T^1\Sigma_{2,q,\infty}$, the linking number between γ and γ' is negative.*

The proof of this statement will occupy the rest of Section 5. The strategy is as follows. Owing to Theorem 4-9, it is enough to show that the linking number of every pair γ, γ' of collections of periodic orbits of the template $\mathcal{B}_{\Gamma_{2,q}, \mathcal{T}_{P,Q}}$ is negative. By Lemma 5-3, the first homology group of $\overline{T^1\Sigma_{2,q,\infty}}$ is $\mathbb{Z}/(q-2)\mathbb{Z}$, so that the 1–cycle $(q-2)[\gamma]$ is the boundary of some 2–chain. What we shall do is to explicitly construct a 2–chain S whose boundary is $(q-2)[\gamma]$, and show that the intersection number of S with γ' is negative. As the family γ' lies in the template $\mathcal{B}_{\Gamma_{2,q}, \mathcal{T}_{P,Q}}$, working with the 1–skeleton of the template as in the proof of Proposition 5-7 is impossible. Instead, we shall choose

a particular projection of $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$ on $\overline{T^1 e}$ and reduce the problem to computing intersection numbers on $\overline{T^1 e}$. Practically, we shall construct the 2–chain S as the union of three parts, namely a 2–chain S_Q^γ lying inside the solid torus $T^1 \Delta'_Q / \Gamma_Q$, a 2–chain S_2^γ lying inside the solid torus $T^1 \Delta'_P / \Gamma_Q$, and a 2–chain S_e^γ in the torus $\overline{T^1 e}$. Then we shall show that the intersection number between S_Q^γ and γ' is slightly positive, that the intersection number between S_2^γ and γ' is zero, and that the intersection number between S_e^γ and γ' is very negative, so that the sum of these three numbers is negative, as expected.

Let us turn to the construction of the 2–chains S_2^γ, S_Q^γ and S_e^γ. They will be defined by gluing discs whose boundaries will consist of *elementary arcs*, some particular segments drawn inside the ribbons of the multitemplate $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$.

As depicted in Figure 19, every ribbon Rib_q^j of $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$ can be distorted in two ways on $\overline{T^1 e}$, according to whether the ribbon is pushed on its right or on its left (see Figure 13). We denote by $\mathrm{Rib}_q^{i,l}$ and $\mathrm{Rib}_q^{i,r}$ the two ribbons in $\overline{T^1 e}$ produced that way. Similarly, the ribbon Rib_p can be pushed on the right or on the left, and can thus be distorted on two ribbons on $\overline{T^1 e}$. We denote them by Rib_p^l and Rib_p^r. What we shall do is to decompose the orbits of $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$ into pieces lying inside a ribbon, and choose for every such piece a combination of the two possible projections, so that the sum of these projections is null-homologous in $\overline{T^1 e}$ (see Figure 20). Here is the precise notion.

Definition 5-10 We say that α is an *elementary arc* (of $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$) if α is

- (type 1) either a segment of an orbit in $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$ that goes from a point A_0 of $\mathrm{BS}_{P \to Q}$ to a point A_1 of $\mathrm{BS}_{Q \to P}$ and travels along the ribbon Rib_q^i for some i between 1 and $q-1$; then we write α^l and α^r for the segments of $\overline{T^1 e}$ that connect A_0 to A_1 and are orbits in the ribbons $\mathrm{Rib}_q^{i,l}, \mathrm{Rib}_q^{i,r}$ respectively,

- (type 2) or a segment of an orbit in $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$ that goes from $\mathrm{BS}_{Q \to P}$ to $\mathrm{BS}_{P \to Q}$ by traveling along Rib_p; then we write α^l and α^r for the deformations of α that are orbits of the ribbons Rib_p^l and Rib_p^r.

We now choose a canonical projection of every elementary arc to a convenient multicurve. So assume that α, α' are elementary of type 1 and 2 respectively, and that the end of α coincides with the origin of α'. Note that the condition about the ends implies that α' is uniquely determined by α. Then we denote by α_π the multicurve consisting of i times α^l and $q-2-i$ times α^r, followed by i times α'^l and $q-2-i$ times α'^r. The reason for this particular choice is the following:

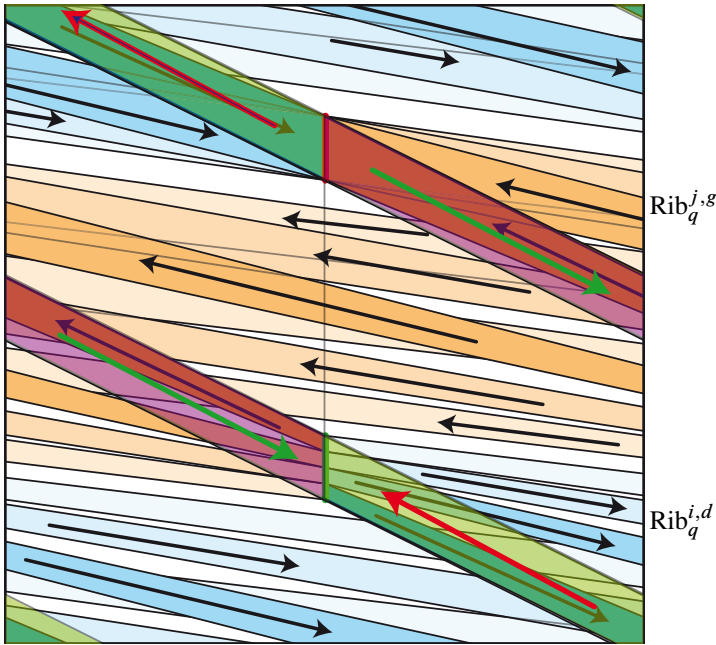Figure 20: The two projections of $\mathcal{B}_{\Gamma_{2,q},\mathcal{T}_{P,Q}}$ on $\overline{T^1 e}$, with $q = 5$, in the slice-of-cake model. The curve a_Z is the vertical boundary of the depicted square, while the curve c_Q is the horizontal boundary. The solid torus $T^1 \Delta'_P / \Gamma_P$ is in front of the picture, so that we see on the front the two projections of the ribbon Rib_2^1. On the back are the two projections of each of the four ribbons $\mathrm{Rib}_q^1, \ldots, \mathrm{Rib}_q^4$.

Lemma 5-11 *Let* $\alpha_1, \alpha'_1 \ldots, \alpha_n, \alpha'_n$ *be the decomposition of* γ *into a concatenation of elementary arcs of type 1 and 2 alternately. Then the union* γ_π *of the multicurves* $(\alpha_1)_\pi, (\alpha'_1)_\pi, \ldots, (\alpha_n)_\pi, (\alpha'_n)_\pi$ *is a multicurve on* $\overline{T^1 e}$ *that is trivial in homology.*

Proof We see on Figure 20 that, for every i, the ribbon $\mathrm{Rib}_q^{i,r}$ (blue on the picture) cuts the curve a_Z (the vertical boundary on the picture) i times and the curve c_Q (the horizontal boundary) -1 times. Similarly, $\mathrm{Rib}_q^{i,l}$ (orange) cuts a_Z (vertical) $q - i$ times and c_Q zero times. In the same way, Rib_p^r cuts a_Z negative one times and c_Q one time, whereas Rib_p^l cuts a_Z one time and c_Q zero times.

Suppose that α, α' are two consecutive elementary arcs of γ, with α lying on the ribbon Rib_q^i for some i and α' lying on Rib_p. Then the above remark implies that the (non-closed) multi-curve $\alpha \cup \alpha'$ has zero-intersection with both a_Z and c_Q. By adding the contributions of all elementary arcs of γ, we deduce that γ is null-homologous in $\overline{T^1 e}$. $\qquad\square$

We are now going to define the 2–chains S_2^γ and S_Q^γ. If α is an elementary arc of type 1, we denote by S_q^α the 2–cycle consisting of i times a disc in $\overline{T^1\Delta_Q'/\Gamma_Q}$ with boundary $\alpha \cup -\alpha^l$ plus $q-2-i$ times a disc with boundary $\alpha \cup -\alpha^r$. Symmetrically, if α is of type 2, we denote by S_p^α the 2–cycle consisting of i times a disc in $\overline{T^1\Delta_P'/\Gamma_P}$ with boundary $\alpha \cup -\alpha^l$ plus $q-2-i$ times a disc with boundary $\alpha \cup -\alpha^r$.

Definition 5-12 With the above notation, we define S_Q^γ to be the union of the 2–cycles $S_q^{\alpha_1}, \ldots, S_q^{\alpha_n}$, and S_2^γ to be the union of the 2–cycles $S_p^{\alpha_1'}, \ldots, S_p^{\alpha_n'}$.

The next step is to complete $S_2^\gamma \cup S_Q^\gamma$ into a 2–chain with boundary γ. Owing to Lemma 5-11, this can be done inside $\overline{T^1e}$. Indeed, the multi-curve γ_π divides $\overline{T^1e}$ into a finite number of regions, say $R_1, \ldots, R_n$, that can be seen as 2–chains. Since $[\gamma_\pi]$ is zero in $H_1(\overline{T^1e}; \mathbb{Z})$, there exists an integral linear combination $\sum \lambda_k [R_K]$ with boundary γ_π. In fact, the coefficients λ_k are defined up to a constant only. With our particular choice of the projection γ_π, at every point of $\mathrm{BS}_{P\to Q}$ or $\mathrm{BS}_{Q\to P}$, the number of segments of γ_π that come from the left (resp. right) equals the number of segments that leave to the left (resp. right).

Definition 5-13 Let us choose numbers λ_k so that, for every region R_i intersecting $\mathrm{BS}_{Q\to P}$, the associated coefficients λ_i is zero. Then we define S_e^γ to be the 2–chain $\sum \lambda_k R_k$.

Note that, by construction, the boundary of the 2–chain S_e^γ is the multicurve γ_π.

At this point, we have associated with the first collection of periodic orbits γ a certain 2–chain $S_2^\gamma \cup S_Q^\gamma \cup S_e^\gamma$ that, by construction, has boundary γ. Let us now consider the second collection of periodic orbits γ', which is assumed to be disjoint from γ. We shall estimate the intersection number between γ' and each of the 2–chains S_2^γ, S_Q^γ and S_e^γ, and prove that their sum is negative. For this, we introduce specific combinatorial data encoding the position of the collections γ and γ' inside the template $\mathcal{B}_{\Gamma_{2,q}, \mathcal{T}_{P,Q}}$.

Lemma 5-14 *The collection γ' does not intersect the 2–chain S_2^γ, and the intersection number between γ' and S_Q^γ is at most*

$$(5\text{-}15) \qquad \sum_{1\le i<j\le q-1} \left(\left\lfloor \frac{j-i}{2}\right\rfloor + 1\right)(i-1)\,b_i b_j' + \sum_{1\le j<i\le q-1} \left(\left\lfloor \frac{i-j}{2}\right\rfloor + 1\right)(q-1-i)\,b_i b_j',$$

where, for every i between 1 and $q-1$, b_i (resp. b_i') is the number of elementary arcs of γ (resp. γ') lying in the ribbon Rib_q^i.
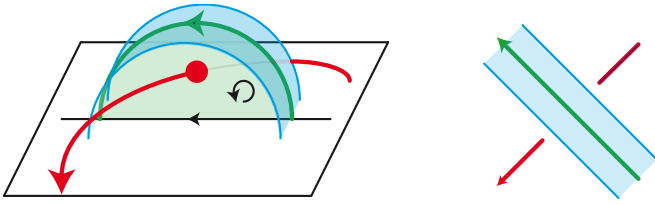
Figure 21: An elementary arc of γ' (red) may intersect the 2–chain S_Q^γ (green) only if it lies on a ribbon under the ribbon containing γ (blue). Since the projection of this intersection point corresponds to a positive crossing (see Figure 20), the intersection number is $+1$.

Proof Every intersection point between γ' and the 2–chains S_2^γ, S_Q^γ is the intersection between one of the discs bounded by an elementary arc, say α, of γ and one of its two projections α^l or α^r, and an elementary arc, say α', of γ'. This implies (see Figure 21) that $\overline{\alpha}, \overline{\alpha'}$ project on $\overline{T^1 e}$ on a double point, and that the ribbon containing α' is between $\overline{T^1 e}$ and the ribbon containing α. In particular, these two ribbons have to be different.

On Figure 20, one sees that all intersections between projected ribbons on $\overline{T^1 e}$ correspond to two ribbons of type $\mathrm{Rib}_q^{i,l}$ and $\mathrm{Rib}_q^{j,l}$, or to two ribbons of type $\mathrm{Rib}_q^{i,r}$ and $\mathrm{Rib}_q^{j,r}$. Therefore no intersection point comes from Rib_p^l or Rib_p^r, so that γ' does not intersect S_2^γ. We also see that, for every i, j, the two projected ribbons $\mathrm{Rib}_q^{i,l}$ and $\mathrm{Rib}_q^{j,l}$ intersect $\lfloor |i - j|/2 \rfloor$ times transversely, and overlap just before the gluing segment $\mathrm{BS}_{Q \to P}$.

The collection of the numbers b_i does not determine the position of the orbit γ on $\mathcal{B}_{\Gamma_{2,q}, \mathcal{T}_{P,Q}}$ completely. In particular, it does not say whether two orbits on $\mathrm{Rib}_q^{i,l}$ and $\mathrm{Rib}_q^{j,l}$ respectively will overlap before $\mathrm{BS}_{Q \to P}$. Nevertheless, since all projected crossings are positive, we obtain an upper bound for the intersection number when assuming that two such elementary arcs always overlap before $\mathrm{BS}_{Q \to P}$.

By construction, there are $i\, b_i$ elementary arcs of γ_π on $\mathrm{Rib}_q^{i,l}$, and $(q - 2 - i)\, b_i$ on $\mathrm{Rib}_q^{i,r}$. Each elementary arc of type 1 yields at most $(\lfloor (j - i)/2 \rfloor + 1) b_j'$ intersection points with elementary arcs of γ' lying on Rib_q^j if $j > i$, and no intersection point for $j \leq i$. Similarly, for $j < i$, each elementary arc of type 2 yields at most $(\lfloor (j - i)/2 \rfloor + 1) b_j'$ intersection points with elementary arcs of γ' lying on Rib_q^j. All intersection points are positive, and (5-15) follows. $\qquad \square$

We now compute the contribution of S_e^γ to the linking number of γ and γ'. For convenience, we set $\Delta = \sum_{1 \leq i \leq q-1} (i - 1)(q - 1 - i)\, b_i$.

Lemma 5-16 *The intersection number between γ' and S_e^γ is at most*

$$(5\text{-}17) \qquad \sum_{j \le q/2} \Big(-\Delta + \sum_{k \le j} (k-1)\, b_k\Big) b_j' + \sum_{j > q/2} \Big(-\Delta + \sum_{k > j} (q-1-k)\, b_k\Big) b_j'.$$

Proof Since γ' intersects the torus $\overline{T^1 e}$ on $\mathrm{BS}_{P \to Q}$ and $\mathrm{BS}_{Q \to P}$ only, we have to estimate the coefficients λ_k of the associated regions in the 2–chain S_e^γ. By definition of S_e^γ, the coefficient of every region intersecting $\mathrm{BS}_{Q \to P}$ is zero. Since every elementary arc of γ' that intersects $\mathrm{BS}_{P \to Q}$ goes from the solid torus $T^1 \Delta'_Q / \Gamma_Q$ into $T^1 \Delta'_P / \Gamma_P$, the intersection number between γ' and S_e^γ is exactly the sum of the levels of the intersection points of γ' with $\mathrm{BS}_{P \to Q}$. Let us cut the segment $\mathrm{BS}_{P \to Q}$ into $q-1$ segments, say $[M_1 M_2], \dots, [M_{q-1} M_q]$, corresponding to the origins of the ribbons $\mathrm{Rib}_q^1, \dots, \mathrm{Rib}_q^{q-1}$.

We claim that the level of the points M_1 and M_q is $-\Delta$. Indeed, starting from M_q (the top point in the segment $\mathrm{BS}_{P \to Q}$ on Figure 20), and following the fiber until we reach $\mathrm{BS}_{Q \to P}$, we intersect the projections of all several ribbons of type $\mathrm{Rib}_q^{i,l}$. For every i, there are $q-1-i$ such intersections, all positive. Since the ribbon $\mathrm{Rib}_q^{i,l}$ contains $i\, b_i$ elementary arcs of γ, we cross γ exactly Δ times along the path. The same argument works for M_1.

Now we claim that, for $i \le q/2$, the level at every point of $[M_i M_{i+1}]$ is at most $\Delta + \sum_{k \le i} (k-1)\, b_k$. Indeed, when starting from M_1 and following $\mathrm{BS}_{P \to Q}$, the level changes when we cross an intersection point of γ with $\mathrm{BS}_{P \to Q}$. Let B be such a point. Then there are $q-2$ elementary arcs of γ arriving at B from the ribbon Rib_p. Depending on the ribbon Rib_q^j followed before Rib_p, the $q-2$ elementary arcs of the projection γ_π arriving at B decompose into $l-1$ of them arriving from the left along $\mathrm{Rib}_q^{j,r}$, and $q-1-l$ arriving from the right along $\mathrm{Rib}_q^{j,l}$. Similarly, since γ leaves B along Rib_q^i, there are $q-2$ elementary arcs of γ_π that leave B, $i-1$ of them on the left along $\mathrm{Rib}_q^{i,l}$, and $q-1-i$ of them on the right along $\mathrm{Rib}_q^{i,r}$. Therefore the difference of level under and above B is $i-l$. In particular, it is at most i. Using an induction on i, we deduce that the level is at most $\Delta + \sum_{k \le i} (k-1)\, b_k$ at M_{i+1}, and *a fortiori* at every point on $[M_i M_{i+1}]$. We get a similar result for i larger that $q/2$. Equation (5-17) easily follows. $\square$

We are now able to complete the argument.

Proof of Proposition 5-9 (Case (a) of Theorem A) We continue with the same notations. Equation (5-17) bounding the intersection number between γ' and S_e^γ

expands into

$$
- \sum_{1 \le j \le q/2} \left(\sum_{i \le j} (i-1)(q-2-i)\, b_i + \sum_{i > j} (i-1)(q-1-i)\, b_i \right) b_j'
$$
$$
- \sum_{q/2 < j \le q-1} \left(\sum_{i < j} (i-1)(q-1-i)\, b_i + \sum_{i \ge j} (i-2)(q-1-i)\, b_i \right) b_j'.
$$

By adding Equation (5-15), we obtain the expression

$$
\sum_{1 \le j \le q/2} \left(\sum_{i < j} -(i-1)(q-2-i-\lfloor (j-i)/2 \rfloor - 1)\, b_i \right.
$$
$$
\left. + \sum_{i > j} -(q-1-i)(i-\lfloor (i-j)/2 \rfloor - 1)\, b_i \right) b_j'
$$
$$
+ \sum_{q/2 < j \le q-2} \left(\sum_{i < j} -(i-1)\big(q-1-i-\lfloor (j-i)/2 \rfloor - 1\big)\, b_i \right.
$$
$$
\left. + \sum_{i > j} -(q-1-i)\big(i-2-\lfloor (i-j)/2 \rfloor - 1\big)\, b_i \right) b_j'
$$

plus some terms in $b_i b_j'$ whose coefficients all are negative. Therefore, the intersection number between γ' and S is bounded from above by a quadratic form in the families $(b_{i,j})$, $(b_{i,j}')$, all of whose coefficients are negative. Therefore, the linking number $\mathrm{Lk}(\gamma, \gamma')$ is negative. $\qquad\square$

6 Surfaces and orbifolds of type $(2, 3, 4g+2)$

We now turn to the hyperbolic 2–orbifolds $\Sigma_{2,3,4g+2}$ and to case (b) in Theorem A, namely the result that every two collections of periodic orbits of the geodesic flow on $\Sigma_{2,3,4g+2}$ are negatively linked. We recall from the introduction that, as the unit tangent bundle $T^1\Sigma_{2,3,4g+2}$ is a quotient of the unit tangent bundle of a specific hyperbolic surface Σ_g of genus g, our strategy will be to lift the question to $T^1\Sigma_g$, estimate the linking number between lifts of orbits of $\Phi_{\Sigma_{2,3,4g+2}}$, and eventually use Lemma 2-2.

In the whole section, g denotes a fixed integer larger than or equal to 2. The successive steps are as follows. We start in Section 6.1 with a $4g+2$–gon in the hyperbolic plane and consider the multitemplate $\mathcal{B}_{4g+2}$ provided by Theorem 4-9. Mimicking the method of the previous section, we bound in Section 6.2 the linking number of a pair of collections of periodic orbits of $\mathcal{B}_{4g+2}$ by a quadratic form Q_{4g+2} in terms of the number of arcs that travel along every ribbon of $\mathcal{B}_{4g+2}$. The form Q_{4g+2} is not negative on the cone of admissible coordinates for geodesics on Σ_g, but, using symmetries to reduce the set of possible coordinates, we introduce a refined form S_{4g+2} in Section 6.3, and show that the linking form is negative on the reduced cone.

6.1 A template for Φ_{Σ_g}

From now on, we fix a regular $4g + 2$–gon P_{4g+2} in the hyperbolic plane whose angles all are equal to $2\pi/(2g + 1)$. We write $e_1, \ldots, e_{4g+2}$ for the sides of P_{4g+2}. For every side e_i, we write e_i^l for its left extremity (when looking at e_i from inside P_{4g+2}), and e_i^r for its right extremity. We also write $e_{\bar{i}}$ for the side opposite to e_i (that is, we set $\bar{i} = i + 2g + 1 \mod 4g + 2$). We call Σ_g the genus-g surface obtained by identifying opposite sides of P_{4g+2}. The vertices of type e_{2k}^r then project to one point of Σ_g, say V_0. Similarly, the vertices of type e_{2k}^l project to one point, say V_1. The unit tangent bundle to P_{4g+2} is the product $P_{4g+2} \times \partial_\infty \mathbb{H}^2$, where a tangent vector is identified with its direction on $\partial_\infty \mathbb{H}^2$. Then $T^1 P_{4g+2}$ is a solid torus whose boundary is made of the $4g + 2$ annuli $T^1 e_1, \ldots, T^1 e_{4g+2}$. The unit tangent bundle $T^1 \Sigma_g$ is obtained from $T^1 P_{4g+2}$ by identifying opposite annuli *via* homographies of $\mathbb{H}^2$. Precisely, if $g_{i,\bar{i}}$ denotes the isometry that maps e_i to $e_{\bar{i}}$, then $g_{i,\bar{i}}$ extends to $\partial_\infty \mathbb{H}^2$, and the fibers of two paired points of e_i and $e_{\bar{i}}$ are identified using the extension of $g_{i,\bar{i}}$ to $\partial_\infty \mathbb{H}^2$. We also introduce two small discs D_0, D_1 on Σ_3 centered at V_0, V_1 respectively. We write $\widehat{P}_{4g+2}$ for the complement of $D_0 \cup D_1$ in P_{4g+2}. This is a domain whose boundary is made of $4g + 2$ geodesic segments and $4g + 2$ arcs of circle are small radius.
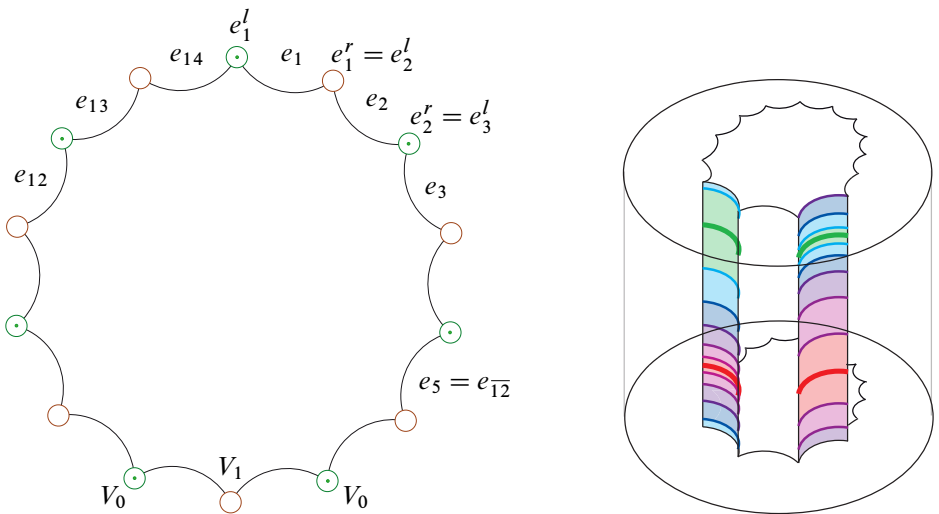

Figure 22: On the left, the regular 14–gon P_{14}. The surface Σ_3 is obtained by identifying opposite sides. On the right, the unit tangent bundle is obtained in the standard coordinates $P_{4g+2} \times \partial_\infty \mathbb{H}^2$ by gluing opposite walls using homographies.

The fundamental group $\pi_1(\Sigma_g)$ is generated by the isometries $g_{1,\bar{1}}, \ldots, g_{4g+1,\overline{4g+1}}$. We write $\mathcal{T}_g$ for the tessellation of $\mathbb{H}^2$ induced by the images of P_{4g+2} under $\pi_1(\Sigma_g)$. One easily checks that $\mathcal{T}_g$ is adapted to Σ_g (only point (v) in Definition 4-1 requires some attention). Finally, we choose a graph $\mathcal{G}_{\mathcal{T}_g}$ dual to $\mathcal{T}_g$ and an associated discretization of geodesics. We then write $\mathcal{B}_{4g+2}$ for the corresponding template in $T^1\Sigma_g$. By definition, it consists of $(4g+2)(4g+1)$ ribbons connecting every pair of distinct boundary annuli. For every i, j, we denote by $\mathrm{Rib}_{i,j}$ the ribbon that connects $T^1 e_i$ to $T^1 e_j$. Above every side of P_{4g+2}, there are two branching segments, corresponding to geodesics crossing the side in both directions. The length of each branching segment is half the length of the fiber. Since we are interested in the topology of $\mathcal{B}_{4g+2}$ only, we can distort it using an isotopy, so that each branching segment has a small length, say ϵ, and consists of vectors that are almost orthogonal to e_i. We then obtain a template similar to the one depicted on Figure 23. For every edge e_i of P_{4g+2}, we denote by $\mathrm{BS}_{i,\bar{i}}$ the branching segment that contains the orbits arriving on the side e_i and leaving from the side $e_{\bar{i}}$, and by $\mathrm{BS}_{\bar{i},i}$ the other branching segment that contains the orbits arriving on the side $e_{\bar{i}}$ and leaving from the side e_i.
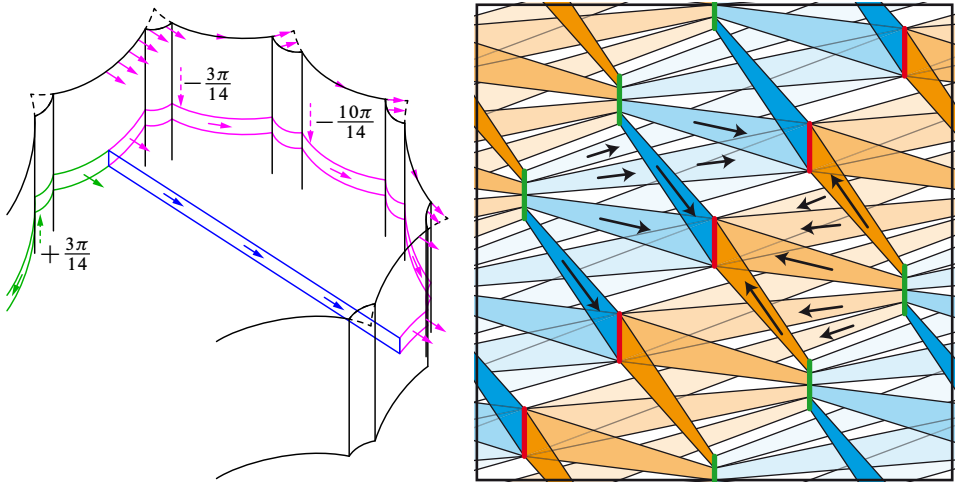

Figure 23: On the left, a ribbon $\mathrm{Rib}_{i,j}^l$ and its projection $\mathrm{Rib}_{i,j}^l$ on $T^1\partial P_{4g+2}$. By definition, it is horizontal in every wall of type $T^1 e_i$. As stated in Lemma 6-9, it goes down around every vertex. This long descent (which is reminiscent of a picture by Escher) stems from the observation that a vector traveling along the left boundary of P_{4g+2} and staying tangent to ∂P_{4g+2} has to turn right at every vertex. On the right, the two projections of the whole template $\mathcal{B}_{4g+2}$ on $T^1\partial P_{4g+2}$ (with 5 instead of $4g+2$): every ribbon has a blue and an orange projection.

In the sequel, we need two particular deformations of $\mathcal{B}_{4g+2}$ on $T^1 \partial \widehat{P}_{4g+2}$ that we describe now. Assume that $\mathrm{Rib}_{i,j}$ is a ribbon of $\mathcal{B}_{4g+2}$. We isotope $\mathrm{Rib}_{i,j}$ to the boundary of $T^1 \widehat{P}_{4g+2}$ without changing the extremities in two ways. For the first deformation, we push $\mathrm{Rib}_{i,j}$ to the left until reaching $T^1 \widehat{P}_{4g+2}$ (see Figure 23 left). The image is denoted by $\mathrm{Rib}_{i,j}^l$. Moreover, we choose the isotopy so that

- the part of $\mathrm{Rib}_{i,j}^l$ lying in $T^1 e_i$ has width ϵ and consists of vectors almost orthogonal to e_i,

- for every c satisfying $j < c < i$ in the cyclic order, the part of $\mathrm{Rib}_{i,j}^l$ lying in $T^1 e_c$ has width ϵ and consists of vectors almost parallel to e_c,

- the part of $\mathrm{Rib}_{i,j}^l$ lying in $T^1 e_j$ has width ϵ and consists of vectors almost orthogonal to e_i.

We construct $\mathrm{Rib}_{i,j}^r$ similarly by pushing $\mathrm{Rib}_{i,j}$ to the right in $T^1 \widehat{P}_{4g+2}$. We write $\mathcal{B}_{4g+2}^l$ for the union of all left projections of ribbons of $\mathcal{B}_{4g+2}$, and $\mathcal{B}_{4g+2}^r$ for the union of all right projections (see Figure 23 right).

6.2 Bounds for the linking number of orbits in Φ_{Σ_g}

Our goal is now to estimate and to bound the linking number between two null-homologous collections of periodic orbits of $\mathcal{B}_{4g+2}$. We will do that by considering the number of times the given collections travel along every ribbon of $\mathcal{B}_{4g+2}$. The formula may look convoluted, but hopefully the meaning of every term should be clear from the proof. The key point is that the bound we establish is bilinear in the number of times each collection travels along every ribbon, so that it can be easily estimated. We use Knuth's convention and write $\{\cdot\}$ for the characteristic function of a property. Also the inequality signs refer to the cyclic order in $\mathbb{Z}/(4g+2)\mathbb{Z}$. The functions v_0, v_1, h_0, h_1 will be defined in Definitions 6-10 and 6-11 below.

Definition 6-1 For every i, j, k, l in $\{1, \ldots, 4g+2\}$ with $i \neq j$ and $k \neq l$, we define the real number $q_{i,j,k,l}$ by

$$\frac{1}{2}\big(\{i < k < l \leq j\} + \{k < i < j \leq l\}\big) - \frac{1}{8}\big(\{k \neq i, j\} + \{k \neq \bar{i}, \bar{j}\}\big)$$

$$+ v_0(i,j)h_0(k,l) + v_1(i,j)h_1(k,l) + \frac{1}{2g-2}(v_0(i,j) + v_1(i,j))(v_0(k,l) + v_1(k,l));$$

we write Q_{4g+2} for the bilinear form on $\mathbb{R}^{(4g+2)(4g+1)}$ whose coefficients are the $q_{i,j,k,l}$.

Definition 6-2 Assume that γ is a null-homologous collection of periodic orbits of the template $\mathcal{B}_{4g+2}$. For every i, j in $\{1, \ldots, 4g+2\}$, let $b_{i,j}$ denote the number of arcs of γ that travel along the ribbon $\mathrm{Rib}_{i,j}$, respectively. The family $(b_{i,j})_{1 \le i \ne j \le 4g+2}$ consists of $(4g+2)(4g+1)$ non-negative integers; it is called the *linear code* of γ.

Proposition 6-3 *Assume that γ, γ' are two null-homologous collections of periodic orbits of the template $\mathcal{B}_{4g+2}$. Denote by $(b_{i,j})$ and $(b'_{i,j})$ their linear codes. Then the linking number $\mathrm{Lk}(\gamma, \gamma')$ is at most $\sum_{1 \le i,j,k,l \le 4g+2} q_{i,j,k,l} b_{i,j} b'_{k,l}$.*

Note that, in the expression for $q_{i,j,k,l}$ given in Definition 6-1, the roles of γ and γ' are not symmetric. This is connected with our subsequent choice of a particular 2–chain, and with the fact that the coefficients $b_{i,j}$ satisfy some linear constraints, so that the above formula is one among many other possible expressions.

The idea of the proof of Proposition 6-3 is to construct a rational 2–chain S^γ with boundary γ, and to bound its intersection number with γ'. The 2–chain S^γ will consist of four parts, denoted by $S^\gamma_\pi, S^\gamma_\partial, S^\gamma_{V_0}$ and $S^\gamma_{V_1}$, each being a combination of several rational 2–cells.

We now establish several intermediate results consisting in evaluating various intersection numbers. First, we consider the above defined projections $\mathcal{B}^l_{4g+2}$ and $\mathcal{B}^r_{4g+2}$ of $\mathcal{B}_{4g+2}$. We write γ^l_π for the image of γ that lies in $\mathcal{B}^l_{4g+2}$, and γ^r_π for the image that lies in $\mathcal{B}^r_{4g+2}$.

Definition 6-4 Let γ_π be the combination $\frac{1}{2}\gamma^l_\pi + \frac{1}{2}\gamma^r_\pi$. Then we define S^γ_π to be the sum, for each elementary arc α of γ, of a (rational) disc d^l_α with boundary $\frac{1}{2}(\alpha \cup -\alpha^l_\pi)$ and of a (rational) disc d^r_α with boundary $\frac{1}{2}(\alpha \cup -\alpha^r_\pi)$.

It follows from the definition that S^γ_π connects γ to γ_π.

Lemma 6-5 *The intersection number between the collection γ' and the rational 2–chain S^γ_π is at most $\sum_{i,j,k,l} \frac{1}{2}(\{i < k < l \le j\} + \{k < i < j \le l\}) b_{i,j} b'_{k,l}$.*

Proof We have to estimate, for every pair of elementary arcs (α, α') of γ and γ' respectively, whether α' intersects the discs d^l_α and d^r_α defined above, how many times it possibly does, and what is the sign of the intersection points. Let $\mathrm{Rib}_{i,j}$ denote the ribbon containing α, and let $\mathrm{Rib}_{k,l}$ the ribbon containing α'.

First, suppose $i \ne k$ and $j \ne l$. Figure 23 right then shows that α' intersects d^r_α if and only if $i < k < l < j$ in the cyclic order. In this case, there is only one intersection point, and its sign is positive (Figure 21 is also relevant here). Since the disc d^r_α has a

coefficient $\frac{1}{2}$, the contribution of this intersection point to the total intersection number is $+\frac{1}{2}$. Similarly, α' intersects d_α^l if and only if $i < j < l < k$ in cyclic order, and the contribution is then $+\frac{1}{2}$.

Second, suppose $i \neq k$ and $j = l$. Then α' may intersect d_α^l or d_α^r or not, depending on which arc is above the other on $\mathrm{BS}_{j,\bar{j}}$, and which arc comes from the right or the left before reaching $T^1 e_j$. Since we look for an upper bound on the linking number, and since the sign of the intersection, if any, is positive, we can assume that there is always an intersection, so that the contribution is $+\frac{1}{2}$. This happens if $i < k < l = j$ or $i < j = l < k$ in the cyclic order. (Note that this is the only approximation that makes our computation of the linking number not exact. It will be refined for symmetric collections of orbits in the next section.)

Third, suppose $i = k$ and $j \neq l$. Then, as in the previous case, the arc α' may intersect d_α^l or d_α^r or not. But, unlike the previous case, we can ignore this potential intersection point. Indeed, let A, B denote the respective starting points of α and α', which are located in the branching segment $\mathrm{BS}_{\bar{i},i}$. Then there is an intersection point if A is under B and at the same time we have $j > l$, or if A is above B and we have $j < l$. At the expense of possibly performing a symmetry, we may restrict to the first case. A under B means that α points on the right of α' on $\partial_\infty \mathbb{H}^2$, whereas $j > l$ means that α escapes on the left of α'. This is possible for $j = l + 1$, but this implies that the geodesics of $\mathbb{H}^2$ that have been distorted onto γ and γ' intersect after crossing e_i. As they are geodesics, they cannot intersect twice, so that they did not intersect before crossing e_i. Therefore there was a pair of arcs that lie before α and α' on γ and γ' that was counted in the previous paragraph (since $i = k$) and should not have. So we can compensate this factor $+\frac{1}{2}$ by ignoring the current intersection.

Fourth, suppose $i = k$ and $j = l$. Then α, α' lie on the same ribbon, and α' does not intersect the discs d_α^r and d_α^l.

Summing up, we obtain the announced upper bound. $\qquad\square$

The second part of S^γ will lie in the $2g + 1$ annuli $T^1 e_c$ with $1 \leq c \leq 2g + 1$ (we recall that e_c is identified with e_{c+2g+1}). Its boundary will be made of γ_π plus some curves lying in $T^1 \partial D_0$ and $T^1 \partial D_1$. Before describing it, we must describe γ_π in more detail.

Lemma 6-6 (See Figure 23 left.) Let e_c be a side of P_{4g+2}. Then the part of γ_π that lies in $T^1 e_c$ consists of

(i) $\frac{1}{2} \sum_{i \neq c} b_{i,c}$ arcs joining the fiber $T^1 e_c^l$ to the branching segment $\mathrm{BS}_{c,\bar{c}}$, plus another $\frac{1}{2} \sum_{i \neq c} b_{i,c}$ arcs joining $T^1 e_c^r$ to $\mathrm{BS}_{c,\bar{c}}$, plus $\frac{1}{2} \sum_{j \neq \bar{c}} b_{c,j}$ arcs joining $\mathrm{BS}_{c,\bar{c}}$ to $T^1 e_c^l$, plus $\frac{1}{2} \sum_{j \neq \bar{c}} b_{c,j}$ arcs joining $\mathrm{BS}_{c,\bar{c}}$ to $T^1 e_c^r$, all these

arcs lying at a height that corresponds to vectors escaping from P_{4g+2} almost orthogonally,

(ii) $\frac{1}{2} \sum_{i \neq \bar{c}} b_{i,\bar{c}}$ arcs joining $T^1 e_c^l$ to $\mathrm{BS}_{\bar{c},c}$, plus $\frac{1}{2} \sum_{i \neq \bar{c}} b_{i,\bar{c}}$ arcs joining $T^1 e_c^r$ to $\mathrm{BS}_{\bar{c},c}$, plus $\frac{1}{2} \sum_{j \neq c} b_{\bar{c},j}$ joining $\mathrm{BS}_{\bar{c},c}$ to $T^1 e_c^l$, plus $\frac{1}{2} \sum_{j \neq c} b_{\bar{c},j}$ arcs joining $\mathrm{BS}_{\bar{c},c}$ to $T^1 e_c^r$, all these arcs lying at a height that corresponds to vectors entering P_{4g+2} almost orthogonally,

(iii) $\frac{1}{2} (\sum_{i < c < j < i} b_{i,j} + \sum_{i < j < \bar{c} < i} b_{i,j})$ arcs joining the fiber $T^1 e_c^r$ to $T^1 e_c^l$, all these arcs lying at a height that corresponds to vectors almost tangent to e_c and pointing toward e_c^l,

(iv) $\frac{1}{2} (\sum_{i < j < c < i} b_{i,j} + \sum_{i < \bar{c} < j < i} b_{i,j})$ arcs joining the fiber $T^1 e_c^l$ to $T^1 e_c^r$, all these arcs lying at a height that corresponds to vectors almost tangent to e_c and pointing toward e_c^r.
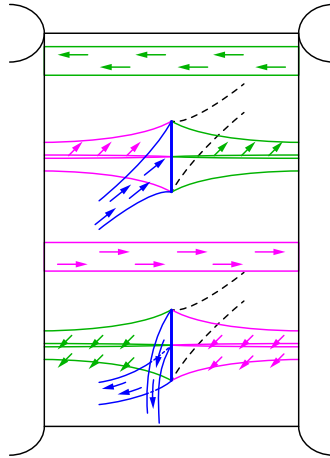

Figure 24: The templates $\mathcal{B}_{4g+2}^l$ and $\mathcal{B}_{4g+2}^r$ inside a wall of type $T^1 e_i$. Between the two branching segments, all ribbons have the same orientation. If γ is a null-homologous collection of orbits of $\mathcal{B}_{4g+2}$, then there are as many arcs of γ_π traveling from left to right (along the pink ribbons) as arcs of γ_π traveling from right to left (along the green ribbons).

Proof Let $\mathrm{Rib}_{i,j}$ be a ribbon of $\mathcal{B}_{4g+2}$. Then every arc of $\gamma \cap \mathrm{Rib}_{i,j}$ projects on an arc of $\gamma^l \cap \mathrm{Rib}_{i,j}^l$ and on an arc of $\gamma^r \cap \mathrm{Rib}_{i,j}^r$. First suppose that the index c differs from both i and j. If e_c lies on the right of $\mathrm{Rib}_{i,j}$, then the arcs of $\gamma^r \cap \mathrm{Rib}_{i,j}^r$ travel along $T^1 e_c$, from $T^1 e_c^r$ to $T^1 e_c^l$. By construction of $\mathrm{Rib}_{i,j}^r$, they are at the height of vectors almost tangent to e_c. Therefore they contribute to (iii). Similarly, if e_c is on the left of $\mathrm{Rib}_{i,j}$, then the arcs of $\gamma^l \cap \mathrm{Rib}_{i,j}^l$ travel along $T^1 e_c$, and contribute to

(iv). In the same vein, we obtain the two others terms of (iii) and (iv) by recalling that $T^1 e_c$ is identified with $T^1 e_{\bar{c}}$, so that, if $\bar{c}$ differs from both i and j, the arcs of $\gamma \cap \text{Rib}_{i,j}$ also project on $T^1 e_{\bar{c}}$ in the same way.

Suppose now $c = j$. Then the arcs of $\gamma^l \cap \text{Rib}_{i,j}^l$ finish their travel by connecting $T^1 e_c^l$ to $\text{BS}_{c,\bar{c}}$, and the arcs of $\gamma^r \cap \text{Rib}_{i,j}^r$ connect $T^1 e_c^r$ to $\text{BS}_{c,\bar{c}}$. Thus they all contribute to (i). Similarly if $\bar{c} = i$, then the arcs of $\gamma^l \cap \text{Rib}_{i,j}^l$ begin their travel by connecting $\text{BS}_{c,\bar{c}}$ to $T^1 e_{\bar{c}}^r = T^1 e_c^l$ to, and the arcs of $\gamma^r \cap \text{Rib}_{i,j}^r$ connect $\text{BS}_{c,\bar{c}}$ to $T^1 e_c^r$, thus all contributing (i). Similarly, we get the expression for (ii) by considering the cases $c = i$ and $\bar{c} = j$. □

As the collection γ is null-homologous, the collection γ_π is also null-homologous, so that for every side e_c of P_{4g+2}, the number of arcs of γ_π that travel along e_c in one direction is equal to the number of arcs in the other direction. This implies that the numbers of arcs given by Lemma 6-6(iii) and (iv) are equal. We then define Δ_c to be their common value, which then admits the more symmetric expression

$$\frac{1}{4}\left(\sum_{i<c<j<i} b_{i,j} + \sum_{i<j<\bar{c}<i} b_{i,j} + \sum_{i<j<c<i} b_{i,j} + \sum_{i<\bar{c}<j<i} b_{i,j} \right),$$

or simply $\frac{1}{4}(\sum_{i,j\neq c} b_{i,j} + \sum_{i,j\neq \bar{c}} b_{i,j})$. Also, since every arc of γ that arrives on $\text{BS}_{c,\bar{c}}$ is followed by an arc that leaves $\text{BS}_{c,\bar{c}}$, the numbers $\sum_{i\neq \bar{c}} b_{i,\bar{c}}$ and $\sum_{j\neq c} b_{\bar{c},j}$ are equal. Hence it is possible to choose a 2-chain in $T^1 e_c$ whose boundary is $\gamma_\pi \cap T^1 e_c$, plus some arcs in the fibers $T^1 e_c^l$ and $T^1 e_c^r$. This 2-chain is unique up to adding multiples of $T^1 e_c$, so that we can make a specific choice that will be convenient for estimating the contributions of the last two components $S_{V_0}^\gamma$ and $S_{V_1}^\gamma$ of S^γ.

Definition 6-7 (See Figure 24.) With the above notation, we define S_∂^γ to be the 2-chain consisting, for every side c of P_{4g+2}, of Δ_c cells in $T^1 e_c$ whose oriented boundary consists of the Δ_c arcs of γ_π that join $T^1 e_c^l$ to $T^1 e_c^r$, plus the Δ_c arcs of γ_π that join $T^1 e_c^r$ to $T^1 e_c^l$, plus $\Delta_c/2$ arcs that go up and $\Delta_c/2$ arcs that go down in the fiber $T^1 e_c^l$, plus $\Delta_c/2$ arcs that go up and $\Delta_c/2$ arcs that go down in the fiber $T^1 e_c^r$.

Lemma 6-8 *The intersection number between the collection γ' and the rational 2–chain S_∂^γ is equal to*

$$-\sum_{i,j,k,l} \tfrac{1}{8}(\{i \neq k \text{ and } j \neq k\} + \{i \neq \bar{k} \text{ and } j \neq \bar{k}\}) b_{i,j} b_{k,l}'.$$

Proof The collection γ' intersects S_∂^γ only on branching segments. Figure 24 then shows that all intersection points have negative sign. For every side e_k of P_{4g+2}, there are $\sum_{l \neq k} b'_{k,l}$ arcs of γ' that cross $T^1 e_k$ in each direction. Every such arc then intersects $\Delta_k/2$ cells of S_∂^γ negatively, so that the total contribution of $T^1 e_k$ to the intersection number is $\Delta_k \sum_{l \neq k} b'_{k,l}$. Therefore the total intersection number is the sum over all sides e_k of P_{4g+2} of the terms $\Delta_k \sum_{l \neq k} b'_{k,l}$. As the sides e_k and $e_{\overline{k}}$ coincide, the latter sum admits the more symmetric expression $\frac{1}{2} \sum_k \Delta_k \sum_{l \neq k} b'_{k,l}$. We then find the expected value by expanding Δ_k. □

The boundary of the 2–chain $S_\pi^\gamma \cup S_\partial^\gamma$ that was constructed above is γ, plus some multiples of the fibers $T^1 V_0$ and $T^1 V_1$ that we now determine.

Lemma 6-9 *Assume that* $\mathrm{Rib}_{i,j}$ *is a ribbon of* $\mathcal{B}_{4g+2}$. *Then:*

(i) *The part of* $\mathrm{Rib}_{i,j}^l$ *that lies in the neighborhood of* $T^1 e_i^l$ *goes down by a height* $(2g-3)\pi/(4g+2)$, *the part of* $\mathrm{Rib}_{i,j}^l$ *that lies in the neighborhood of* $T^1 e_j^r$ *goes down by a height* $(2g-3)\pi/(4g+2)$, *except if* $j = i+1$, *in which case the part of* $\mathrm{Rib}_{i,j}^l$ *that lies in the neighborhood of* $T^1 e_i^l = T^1 e_j^r$ *goes up by a height* $4\pi/(4g+2)$.

(ii) *The part of* $\mathrm{Rib}_{i,j}^l$ *that lies in the neighborhood of* $T^1 e_c^l$, *for* $j+1 < c < i$, *goes down by a height* $(4g-2)\pi/(4g+2)$.

(iii) *The part of* $\mathrm{Rib}_{i,j}^r$ *that lies in the neighborhood of* $T^1 e_i^r$ *goes up by a height* $(2g-3)\pi/(4g+2)$, *the part of* $\mathrm{Rib}_{i,j}^r$ *that lies in the neighborhood of* $T^1 e_j^l$ *goes up by a height* $(2g-3)\pi/(4g+2)$, *except if* $j = i-1$, *in which case the part of* $\mathrm{Rib}_{i,j}^r$ *that lies in the neighborhood of* $T^1 e_i^r = T^1 e_j^l$ *goes up by a height* $4\pi/(4g+2)$.

(iv) *The part of* $\mathrm{Rib}_{i,j}^r$ *that lies in the neighborhood of* $T^1 e_c^r$, *for* $i < c < j-1$, *goes up by a height* $4\pi/(4g+2)$.

Proof The proof is illustrated in Figures 23 and 25. It relies on the assumption that the angle between adjacent sides of P_{4g+2} is $2\pi/(2g+1)$, and on the height we chose for the parts of the ribbons $\mathrm{Rib}_{i,j}^l$ and $\mathrm{Rib}_{i,j}^r$ above each edge of $\widehat{P}_{4g+2}$. The values follow from the equalities $\pi/2 - 2\pi/(2g+1) = (2g-3)\pi/(4g+2)$ and $\pi - 2\pi/(2g+1) = (4g-2)\pi/(4g+2)$. □

Definition 6-10 For i, j in the range $\{1, \ldots, 4g+2\}$, we define $v_0(i, j)$ to be the sum over all even vertices of P_{4g+2} of the increases of $\mathrm{Rib}_{i,j}^l$ and of $\mathrm{Rib}_{i,j}^r$ around this vertex.

For example, if $i = 1$ and $j = 2$, then $\mathrm{Rib}^l_{1,2}$ contributes $4\pi/(4g+2)$ to $v_0(0,1)$ and 0 to $v_1(0,1)$. On the other hand, $\mathrm{Rib}^r_{1,2}$ contributes $2g$ times $+(4g-2)\pi/(4g+2)$ to $v_0(0,1)$ and 2 times $(2g-3)\pi/(4g+2)$ plus $2g-1$ times $(4g-2)\pi/(4g+2)$ to $v_1(0,1)$. Therefore we have $v_0(0,1) = (8g^2 - 4g + 4)\pi/(4g+2)$ and $v_1(0,1) = (8g^2 - 4g - 4)\pi/(4g+2)$. With the above notation, the boundary of the 2–chain $S^\gamma_\pi \cup S^\gamma_\partial$ consists of the union of γ, of $-\sum_{i,j} v_0(i,j)b_{i,j}$ times the fiber $T^1 V_0$ and of $-\sum_{i,j} v_1(i,j)b_{i,j}$ times the fiber $T^1 V_1$.
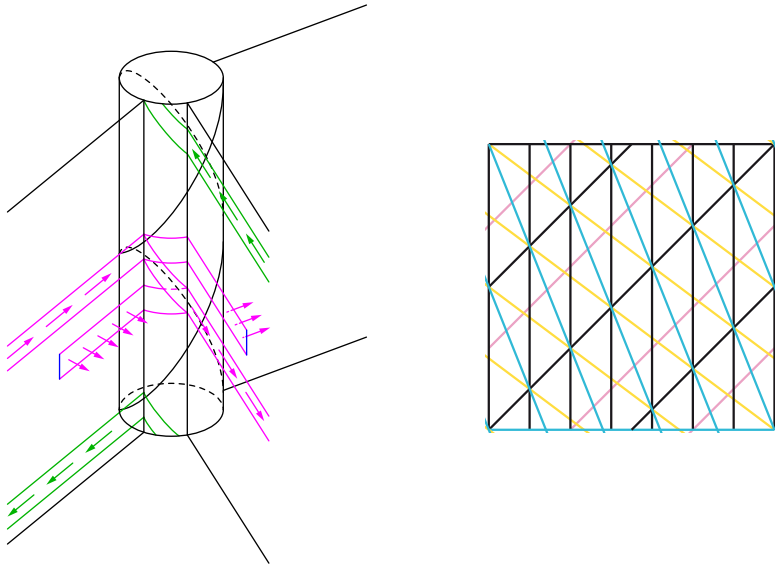

Figure 25: The templates $\mathcal{B}^l_{4g+2}$ and $\mathcal{B}^r_{4g+2}$ on the neighbourhood of a vertex V_i. For $\mathcal{B}^l_{4g+2}$, most of the ribbons go down by $(4g-2)\pi/(4g+2)$ (in blue), or they are close to a branching segment and they go down by $(2g-3)\pi/(4g+2)$ (in yellow), unless they are close to both branching segments and they go up by $4\pi/(4g+2)$ (in pink). For $\mathcal{B}^r_{4g+2}$, the signs are reversed.

In order to complete the chain S^γ, it suffices that we add a 2–chain whose boundary is $(\sum_{i,j} v_0(i,j)b_{i,j})T^1 V_0$ and a 2–chain whose boundary is $(\sum_{i,j} v_1(i,j)b_{i,j})T^1 V_1$. Since Σ_g has Euler characteristic $2-2g$, there exists a vector field on Σ_g with only one singularity at V_0, the index of the latter being $2-2g$. By lifting this vector field in $T^1\Sigma_g$, we obtain a surface with boundary $(2g-2)T^1 V_0$. We then define $S^\gamma_{V_0}$ to be $\frac{1}{2g-2}(\sum_{i,j} v_0(i,j)b_{i,j})$ times this surface. Similarly, we can construct a surface with boundary $(2g-2)T^1 V_1$, and we then define $S^\gamma_{V_1}$ to be $\frac{1}{2g-2}(\sum_{i,j} v_1(i,j)b_{i,j})$ times the latter surface. We have now only to determine the intersection number of γ' with both $S^\gamma_{V_0}$ and $S^\gamma_{V_1}$. For this it is enough to determine the linking number of γ'

with the fibers $T^1 V_0$ and $T^1 V_1$, and then to multiply by $\frac{1}{2g-2}(\sum_{i,j} v_0(i,j)b_{i,j})$ and $\frac{1}{2g-2}(\sum_{i,j} v_1(i,j)b_{i,j})$ respectively.

Definition 6-11 For i, j in the range $\{1, \ldots, 4g + 2\}$, we define $h_0(i, j)$ as the number of even vertices of P_{4g+2} on the left of $\mathrm{Rib}_{i,j}$, minus the number of even vertices on the right of $\mathrm{Rib}_{i,j}$, divided by $2g + 1$. Similarly, we define $h_1(i, j)$ as the number of odd vertices of P_{4g+2} on the left of $\mathrm{Rib}_{i,j}$, minus the number of odd vertices on the right, divided by $2g + 1$.

The precise expressions for $h_0(i, j)$ and $h_1(i, j)$ are $((j - i)[4g + 2] - (2g + 1) + j[2] - i[2])/(2g+1)$ and $((j-i)[4g+2] - (2g+1) - j[2] + i[2])/(2g+1)$, respectively. Moreover, we fix two arbitrary points V_0' and V_1' on the boundaries of D_0 and D_1 respectively. We also choose two meridians m_0 and m_1 of the solid tori $T^1 D_0$ and $T^1 D_1$.

Lemma 6-12 *The collection γ' is homologous, in the complement of $T^1 V_0 \cup T^1 V_1$, to*

$$
\left(\sum_{k,l} v_0(k,l)b'_{k,l} \right) T^1 V_0' + \left(\sum_{k,l} v_1(k,l)b'_{k,l} \right) T^1 V_1'
$$
$$
+ \left(\sum_{k,l} h_0(k,l)b'_{k,l} \right) m_0 + \left(\sum_{k,l} h_1(k,l)b'_{k,l} \right) m_1.
$$

Proof A construction similar to the construction of the 2–chain $S_\pi^\gamma \cup S_\partial^\gamma$, applied to γ' instead of γ, realizes a cobordism between γ and the announced collection of curves. □

Lemma 6-13 *The intersection number between $S_{V_0}^\gamma \cup S_{V_1}^\gamma$ and γ' is equal to*

$$
\sum_{i,j,k,l} \Big[v_0(i,j)h_0(k,l) + v_1(i,j)h_1(k,l)
$$
$$
+ \frac{1}{2g+1}(v_0(i,j) + v_1(i,j))(v_0(k,l) + v_1(k,l)) \Big] b_{i,j} b'_{k,l}.
$$

Proof The curve m_0 bounds a meridian disc for $T^1 D_0$, so that its linking numbers with $T^1 V_0$ and $T^1 V_1$ are 0 and 1 respectively. Similarly one has $\mathrm{Lk}(m_1, T^1 V_0) = 0$ and $\mathrm{Lk}(m_1, T^1 V_1) = 1$. The lift of the vector field on Σ_3 with only one singularity p defines a surface in $T^1 \Sigma_g$ whose boundary is $(2g - 2)T^1 p$ and which intersects every other fiber once. Therefore, $\mathrm{Lk}(T^1 p, T^1 p') = \frac{1}{2g-2}$ holds for every point p' distinct from p. □

Proposition 6-3 now follows from Lemmas 6-5, 6-8 and 6-13, which together give the expected bound for $\mathrm{Lk}(\gamma, \gamma')$ directly.

The set of linear codes $(b_{i,j})$ that correspond to geodesics on Σ_g is a subset of $\mathbb{R}^{(4g+2)(4g+1)}$. Actually, it is a cone included in $\mathbb{R}_+^{(4g+2)(4g+1)}$ that we denote by C_g. It is not hard to see that C_g is a proper subset of $\mathbb{R}_+^{(4g+2)(4g+1)}$, ie, that there are more constraints on the possible values of $(b_{i,j})$ than the positivity of the coefficients. For example, there are linear equality constraints coming from the fact that every arc of the associated collection that crosses a side of P_{4g+2} continues on the other side, as well as linear equalities coming from the fact the collection is null-homologous. There are also inequality constraints coming from the fact that the collection consists of geodesics, so that it cannot always wind around a vertex. Precisely, some coefficients of the form $b_{i,j}$ with $|i - j| \geq 2$ cannot be too small when compared with the coefficients of the form $b_{i,i+1}$.

Implementing the above constraints in a computer program leads to numerical bounds for the linking numbers of orbits of Φ_{Σ_g}. However, as we shall see in Section 7, some collections of orbits have a positive linking number, so there is no hope of proving a uniform negativity result.

6.3 Linking of geodesics on the orbifolds $\Sigma_{2,3,4g+2}$

We now consider the case of the orbifold $\Sigma_{2,3,4g+2}$. Our goal is to establish upper bounds for the linking numbers of pairs of orbits in the associated geodesic flow. We shall prove:

Proposition 6-14 (Case (b) of Theorem A) *Let* γ, γ' *be two orbits of* $\Phi_{\Sigma_{2,3,4g+2}}$ *in* $T^1\Sigma_{2,3,4g+2}$. *Then we have* $\mathrm{Lk}(\gamma, \gamma') < 0$.

The proof relies on a more precise study of the template $\mathcal{B}_{4g+2}$ and refinement of Proposition 6-3. The starting point is that $\Sigma_{2,3,4g+2}$ admits a covering of index $3(4g + 2)$ by a genus-g surface Σ_g obtained by identifying sides of a regular $4g + 2$–gon (see Figure 26). So, by the behavior under quotient of the linking number (Lemma 2-2), in order to establish Proposition 6-14, it is enough to prove that $\Gamma_{2,3,4g+2}$–invariant geodesics of Σ_g have a negative linking number. These $\Gamma_{2,3,4g+2}$–invariant geodesics have three advantages that are needed in the proof. First, their symmetry properties allow to use reduced linear codes with $4g + 1$ coordinates instead of $(4g + 2)(4g + 1)$, thus also simplifying the matrix Q_{4g+2}, bounding the linking number to a more simple $(4g + 1) \times (4g + 1)$ matrix (Lemma 6-15). Second, it is possible to refine the bounds on the linking number by refining the intersection number between the 2–chain S_π^γ and
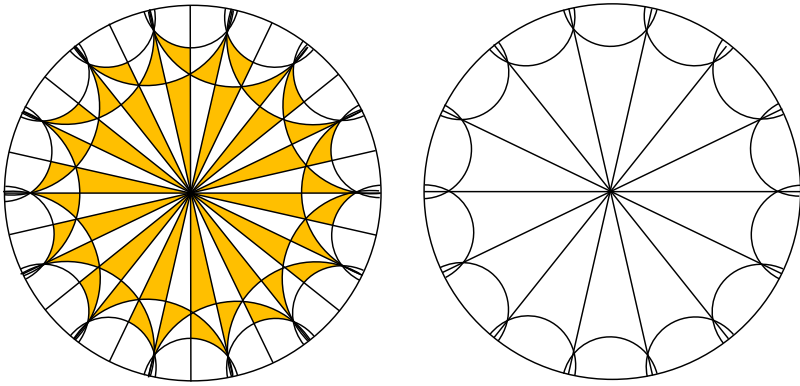
Figure 26: On the left, the covering of $\Sigma_{2,3,4g+2}$ by a genus-g surface, for $g = 3$. A fundamental domain for $\Sigma_{2,3,4g+2}$ is obtained by taking the union of any white triangle with a neighboring orange triangle. On the right the intermediate tiling $\mathcal{T}_{4g+2}$ of the hyperbolic plane by equilateral triangles with angles $\pi/(2g+1)$. Each triangle is a 3–fold cover of $\Sigma_{2,3,4g+2}$. The polygon P_{4g+2} (and therefore the surface Σ_g) is obtained by gluing $4g+2$ triangles that are adjacent to a vertex.

the curve γ', thus refining the first term in Definition 6-1. The price to pay is to add $2g$ coordinates to the reduced linear code that describes how many consecutive times the family takes the rightmost and leftmost ribbons of the template. These two first steps then associate to every collection of $\Gamma_{2,3,4g+2}$–invariant geodesics a reduced linear code with $6g+1$ coordinates, so that the linking number between two collections in bounded by a bilinear form S_{4g+2} in the reduced linear code. Third, we determine a cone $C_{2,3,4g+2}$ in $\mathbb{R}^{6g+1}$ that (strictly) contains all reduced linear codes, and whose extremal rays are easy to determine. The proof of Proposition 6-14 then consists of proving that the form S_{4g+2} is negative on all pairs of extremal rays of $C_{2,3,4g+2}$.

So, let γ, γ' be two orbits of $\Phi_{\Sigma_{2,3,4g+2}}$. Let $\hat{\gamma}, \hat{\gamma}'$ be the images in the template $\mathcal{B}_{4g+2}$ of the $\Gamma_{2,3,4g+2}$–invariant lifts of γ and γ' in $T^1\Sigma_g$. Denote by $b_{i,j}$ and $b'_{i,j}$ their linear codes, as defined in Definition 6-2. Since the collection $\hat{\gamma}$ is invariant under an order $4g+2$ rotation around the center of P_{4g+2}, we have $b_{i,j} = b_{i+1,j+1}$ for every i, j. Therefore, one can consider a simpler code $\hat{b}_{i,j}$ defined for $j = 1, \ldots, 4g+1$ by $\hat{b}_j = \sum_{i=0,\ldots,4g+1} b_{i,i+j}$. Similarly, we introduce a reduced form $\hat{Q}_{2,3,4g+2}$ on $\mathbb{R}^{4g+1}$ whose coefficients $\hat{q}_{j,l}$ are defined by $\hat{q}_{j,l} = \sum_{i,k=0,\ldots,4g+1} q_{i,i+j,k,k+l}$.

Lemma 6-15 *With the above definitions, for $j, l = 1, \ldots, 4g+1$, we have*

$$(6\text{-}16) \qquad \hat{q}_{j,l} = (2g+1)|j-l| - 2g(2g+1) + \frac{1}{2g-2}(j-2g-1)(l-2g-1).$$

Proof We start from the formula for $q_{i,j,k,l}$ given by Definition 6-1 with replacing j by $i + j$ and l by $k + l$. The first term

$$-\tfrac{1}{2}(\{i < k < k + l \leq i + j\} + \{k < i < i + j \leq k + l\})$$

equals $\tfrac{1}{2}$ if the two chords connecting the edges e_i to e_{i+j}, and e_k to e_{k+l} do not intersect and are parallel, or if they have a common head (see Figure 27). When the differences j and l are fixed, they are $4g + 2$ possible choices for the first chord, and then there are $|j - l|$ positions for the second chord that give an admissible position. This gives the first term of Equation (6-16).

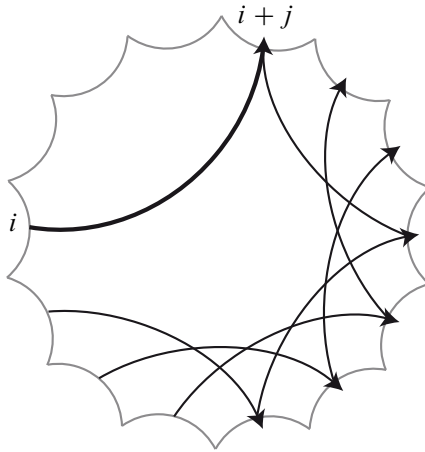

Figure 27: Once the chord connecting e_i to e_{i+j} is chosen (with $j = 4$ on the picture), there are $|l - j|$ choices for k, so that the chord connecting e_k to e_{k+l} contributes to $\{i < k < k + l \leq i + j\}$ or to $\{k < i < i + j \leq k + l\}$ (with $l = 11$ on the picture).

For the second term, we note that when $i, i + j$ are fixed, there are $4g$ values of k that add $\tfrac{1}{8}$ to the sum, and $4g$ values of $\bar{k}$ that also add $\tfrac{1}{8}$. So this yields a contribution of g when i is fixed. By summing over all i, we obtain the second term.

The third term in Definition 6-1 depends of the parity of $i, i + j, k, k + l$, because we are considering the rotation amount of the chord with respect to the two different vertices of P_{4g+2}. When summed over all i, j, these two rotation amounts are equal, so that we only consider the mean rotation of the chords. These are equal to $j - 2g - 1$ and $l - 2g - 1$ respectively. Then the contribution to $\hat{q}_{j,l}$ is a multiple of $(j - 2g - 1)(l - 2g - 1)$. The constant is given by Lemma 6-9. □

The symmetry of the families $\hat{\gamma}, \hat{\gamma}'$ now allows us to refine Lemma 6-5, at the expense of expanding the code. The idea is that if several consecutive arcs of $\hat{\gamma}$ all travel along

the rightmost ribbon, then they cannot cross as many arcs of $\widehat{\gamma}'$ as the bound (and the proof) of Lemma 6-5 suggests. For the sequel, it is important to remember that the families $\widehat{\gamma}$ and $\widehat{\gamma}'$ are invariant by a rotation of order $4g+2$ of P_{4g+2}.

Definition 6-17 For $m = 1, \ldots, 2g$, let c_m (resp. d_m) denote the number of arcs of $\widehat{\gamma}$ that travel exactly m consecutive times along $\mathrm{Rib}_{0,1}$ (resp. $\mathrm{Rib}_{0,4g+1}$). Define c'_m and d'_m similarly. The set $((b_j)_{j=1,\ldots,4g+1}, (c_m)_{m=1,\ldots,2g}, (d_m)_{m=1,\ldots,2g})$ is called the *linear reduced code* of $\widehat{\gamma}$.

For $m, n = 1, \ldots, 2g$, define $r_{m,n}$ to be $-2g-1$ if $i = j \neq 1$ and 0 otherwise. Let R_{4g+2} denote the bilinear form on $\mathbb{R}^{2g}$ with coefficients $r_{m,n}$, and let S_{4g+2} denote the bilinear form on $\mathbb{R}^{8g+1}$ that is the direct sum $\widehat{Q}_{4g+2} \oplus R_{4g+2} \oplus R_{4g+2}$.

Note that if an arc travels m consecutive times along a ribbon, then it travels $m-1$ times at it next move. Thus we have $c_{m-1} \geq c_m$ and $d_{m-1} \geq d_m$ for every m. Note also that some orbits of the template could travel more than g times along the leftmost ribbon, thus making more than one half-turn around the corresponding vertex of P_{4g+2}. These orbits do not interest us, since they cannot come from geodesics.

Lemma 6-18 *With the above notation, the intersection number between $\widehat{\gamma}'$ and the 2–chain S_π^γ is at most*

$$\sum_{j,l=1}^{4g+1} (2g+1)|j-l|b_j b'_l - (2g+1) \sum_{m=2}^{g} (c_m c'_m + d_m d'_m).$$

The linking number $\mathrm{Lk}(\widehat{\gamma}, \widehat{\gamma}')$ *is smaller than*

$$S_{4g+2}\big(((b_j), (c_m), (d_m)), ((b_j), (c_m), (d_m))\big).$$

Proof (See Figure 28.) We use the notation introduced in the proof of Lemma 6-5. In the second case of this proof ($i \neq k$, $j = l$), we assumed that there was always an intersection between the considered arc α' on $\mathrm{Rib}_{k,l}$ and any elementary piece d_α^l or d_α^r of the 2–chain S_π^γ. Actually, if α' is an arc that lies in the rightmost ribbon $\mathrm{Rib}_{j+1,j}$ there is an intersection point with d_α^l or d_α^r if and only if α does not lie in $\mathrm{Rib}_{j+1,j}$ and the head of α' in the vertical branching segment $\mathrm{BS}_{j,\bar{j}}$ is under the head of α. In particular, we know that there are c'_2 elementary arcs of $\widehat{\gamma}'$ in $\mathrm{Rib}_{j+1,j}$ whose heads are above all tails of arcs of $\widehat{\gamma}$ that will travel more than one time along the rightmost ribbon. Indeed, if an arc travels two or more times along the rightmost ribbon, then its direction at infinity is on the right of the direction of an arc travels only once on the rightmost ribbon (see Figure 28). Since there are at least d_2 such arcs
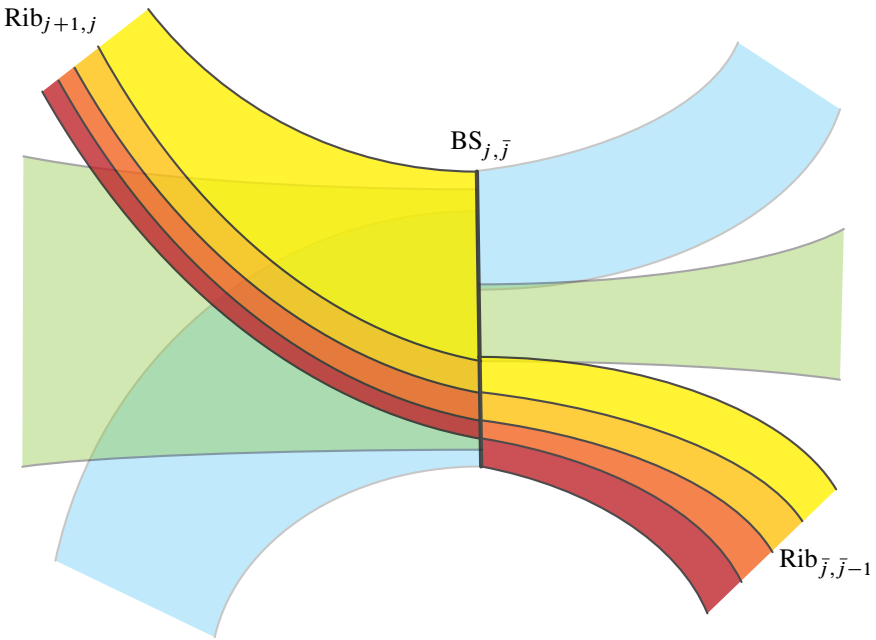
Figure 28: A neighborhood of a branching segment $\mathrm{BS}_{j,\bar{j}}$ in $T^1\Sigma_g$. Both ribbons $\mathrm{Rib}_{j+1,j}$ and $\mathrm{Rib}_{\bar{j},\bar{j}-1}$ are subdivided into subribbons containing arcs that travel $1, 2, 3, \ldots$ consecutive times respectively along the rightmost ribbon. For a family of orbits of the template that is invariant by rotation of P_{4g+2}, the arcs on $\mathrm{Rib}_{j+1,j}$ and on $\mathrm{Rib}_{\bar{j},\bar{j}-1}$ are in one-to-one correspondence.

in $\hat{\gamma}$ at each branching segment, we can add a term $-\frac{1}{2}(4g+2)d_2d_2'$ to the previous bound on the intersection number between S_π^γ and $\hat{\gamma}'$. Similarly, we can consider the d_3' arcs of $\hat{\gamma}'$ that reach $\mathrm{BS}_{j,\bar{j}}$ along the rightmost arc and that will travel along it two more times. Their heads cannot be above the tails of the d_3 arcs of $\hat{\gamma}$ that arrive at $\mathrm{BS}_{j,\bar{j}}$ from a different ribbon and that travel two or more times along the rightmost ribbon. At the end, we can then add a term $-(2g+1)(d_2d_2' + d_3d_3' + \cdots + d_gd_g')$. Considering also the leftmost ribbons gives the announced extra-term.

The formula for total linking number then follows by replacing the first term in Equation (6-16) by the above one. $\qquad\square$

The goal is now to bound the value of the quadratic form S_{4g+2} on the set of linear reduced codes that come from geodesics of $\Sigma_{2,3,4g+2}$. In order to do this, we first determine a cone in $\mathbb{R}^{6g+1}$ that contains the set of linear reduced codes.

Definition 6-19 For x, y in $\{1, \ldots, 2g\}$, let $V_{x,y}$ be the vector in $\mathbb{R}^{4g+1} \oplus \mathbb{R}^g \oplus \mathbb{R}^g$ with coordinates

$$((x - 1, 0, \ldots, 0, 1, 0, \ldots, 0, 1, 0, \ldots, 0, y - 1), (2, \ldots, 2, (1), 0, \ldots, 0),$$

$$(2, \ldots, 2, (1), 0, \ldots)),$$

where the two 1 in the first block are in position $y + 1$ and $4g - x + 1$, where there are $\lfloor (x - 1)/2 \rfloor$ coefficients 2 in the second block, one 1 if x is even, and there are $\lfloor (y - 1)/2 \rfloor$ coefficients 2 in the last block, and one 1 if y is even.

Let C_{4g+2} be the conway hull in $\mathbb{R}^{6g+1}$ of the rays generated by the $4g^2$ vectors $V_{x,y}$.

Lemma 6-20 *With the above definition, the reduced linear code of every collection of $\Gamma_{2,3,4g+2}$–periodic geodesics belongs to $C_{4g+2} \setminus \{0\}$.*
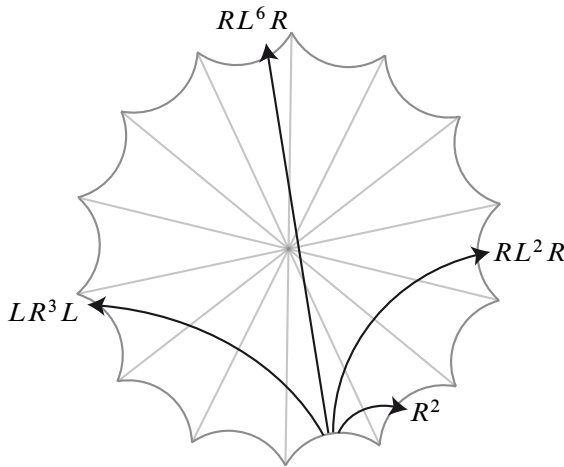

Figure 29: Some dynamical codes. The length of the code equals the number of crossed triangles.

Proof (See Figure 29.) Let $\mathcal{T}_{4g+2}$ denote the $\Gamma_{2,3,4g+2}$–invariant tessellation of $\mathbb{H}^2$ by equilateral triangles with angles $2\pi/(4g + 2)$. Note that a fundamental domain for the action of $\Gamma_{2,3,4g+2}$ on $\mathbb{H}^2$ is given by a third a tile of $\mathcal{T}_{4g+2}$. Note also that by considering the $4g + 2$ triangular tiles that are adjacent to a given vertex, we obtain a fundamental domain P_{4g+2} for the surface Σ_g.

As before, let $\hat{\gamma}$ be a $\Gamma_{2,3,4g+2}$–periodic geodesics, considered in $\mathbb{H}^2$. We associate to it a *dynamical code* in the following way. Starting from an arbitrary intersection point between $\hat{\gamma}$ and an edge of $\mathcal{T}_{4g+2}$, we follow the geodesics $\hat{\gamma}$. Every time we cross a triangle of $\mathcal{T}_{4g+2}$, we add a letter L to the dynamical code if $\hat{\gamma}$ goes to the

left in this triangle, a letter R if it goes to the right. Of course we stop after one period. At the expense of a cyclic permutation, the dynamical code can then be assumed to be of the form $L^{x_1} R^{y_1} L^{x_2} \cdots R^{y_n}$.

The key point is that $1 \le x_k \le 2g$ and $1 \le y_k \le 2g$ hold for every k. Indeed, a curve that goes more than $2g$ consecutive times on the left crosses one of the geodesics containing edges of the tiling more than once, and therefore it cannot be a geodesic.

The second point is that the linear reduced code depends linearly of the exponents x_k, y_k in an explicit way. Indeed, every arc α of $\hat{\gamma}$ in P_{4g+2} is associated to a unique position in the dynamical code that describes the dynamical code when starting at the tail of α. Conversely, to every position in the dynamical code are associated $4g + 2$ arcs of $\hat{\gamma}$ that are obtained one from another by a rotation about the center of P_{4g+2}.

Now, if an arc α goes from an edge e_i to the edge e_{i+j} in P_{4g+2} with $2 \le j \le 2g+1$, then the corresponding dynamical code is $LR^{j-1}L$, while the linear reduced code contains only a 1 in $j-1^{\text{st}}$ position. Similarly if an arc goes from e_i to e_{i+j} with $2g+1 \le j \le 4g$, then the dynamical code is $RL^{4g+1-j}R$ and the linear reduced code contains only a 1 in $j-1^{\text{st}}$ position. (There is an ambiguity in the case $j = 2g+1$ for the dynamical code, depending on which side of the center of P_{4g+2} the geodesics go. But both give rise to the same the linear reduced code, so that this ambiguity is of no consequence.)

In the remaining case, if an arc goes from e_i to e_{i+1}, then the dynamical code begins with L^2, and the linear reduced code begins with a 1 in 1^{st} position. However, the second block of coordinates (that corresponds to the vector (c_m)) can be non-zero, depending on how many times the geodesics $\hat{\gamma}$ will go on the left after escaping P_{4g+2}. The point here is that the dynamical code actually begins with $L^x R$, and the number of times that $\hat{\gamma}$ will travel along the leftmost ribbon is $\lfloor x/2 \rfloor$. Therefore the second block there contains a 1 in $\lfloor x/2 \rfloor^{\text{th}}$ position. The case $j = 4g+1$ is treated similarly.

Finally, we truncate the dynamical code of $\hat{\gamma}$ into the n blocks

$$RL^{x_1} R^{y_1-1}, \quad RL^{x_2} R^{y_2-1}, \quad \ldots, \quad RL^{x_n} R^{y_n-1}.$$

The linear reduced code that corresponds to a block $RL^{x_k} R^{y_k-1}$ is the sum of the linear codes corresponding to each of the $x_k + y_k$ letters, which turns out to be V_{x_k, y_k} by the above discussion. Therefore the linear code associated to $\hat{\gamma}$ is the sum of n such vectors. Thus it belongs to $C_{4g+2} \setminus \{0\}$. □

Lemma 6-21 *The form S_{4g+2} is negative on $C_{4g+2} \setminus \{0\}$.*

Proof The form S_{4g+2} is bilinear, so that it is enough to show that it is negative when evaluated on every pair $(V_{x,y}, V_{x',y'})$ of extremal vectors. Now we note that $S_{4g+2}(V_{x,y}, V_{x',y'})$ is a function of x, x', y, y' which is almost linear in each of the coordinates. Indeed, we can expend Equation (6-16) corrected with the term of Lemma 6-18 and obtain for $S_{4g+2}(V_{x,y}, V_{x',y'})$ the value

$$
\begin{aligned}
(2g+1)&\big((x-1)(y'+4g-x') + (y-1)(4g-y'+x') + (x'-1)(y+4g-x) \\
&+ (y'-1)(4g-y+x) + 4g(x-1)(y'-1) \\
&+ 4g(x'-1)(y-1) + |y-y'| + |y-4g+x'| + |x-x'| + |y'-4g+x|\big) \\
&- (2g+1)\Big(4\min\Big(\Big\lfloor\frac{x-1}{2}\Big\rfloor, \Big\lfloor\frac{x'-1}{2}\Big\rfloor\Big) + 4\min\Big(\Big\lfloor\frac{y-1}{2}\Big\rfloor, \Big\lfloor\frac{y'-1}{2}\Big\rfloor\Big) + \cdots\Big) \\
&- 2g(2g+1)(x+y)(x'+y') \\
&+ \frac{2g-1}{2g-2}\big(-2g(x-1) + (y-2g) + (2g-x) + 2g(y-1)\big) \\
&\qquad\qquad \times \big(-2g(x'-1) + (y'-2g) + (2g-x') + 2g(y'-1)\big).
\end{aligned}
$$

The second term — which corresponds to the correction that we added — contains an extra term dealing with the parity of x and y. Since it is negative, forgetting it can only increase the result.

The observation here is that, except for what concerns the integer part operation in the second term, the above formula is linear in the variables x, x', y, y' on the four regions $\{x < x', y < y'\}$, $\{x < x', y > y'\}$, $\{x > x', y < y'\}$ and $\{x > x', y > y'\}$. By replacing $\lfloor (x-1)/2 \rfloor$ by $(x-1)/2$, thus slightly increasing the result, we obtain a formula that is linear is all four variables.

Therefore, in order to prove that $S_{4g+2}(V_{x,y}, V_{x',y'})$ is negative, we only have to evaluate the above formula on the extremal points of the four connected components of the domain that we are considering. These turn out to be 16 vertices of the cube $[1, 2g]^4$. Using symmetries, we can actually reduce the computation to six points, namely to $(1,1,1,1)$, $(1,1,1,2g)$, $(1,1,2g,2g)$, $(1,2g,1,2g)$, $(1,2g,2g,1)$ and $(2g,2g,2g,2g)$. It is then easy to check that the form is negative on these points. $\qquad\square$

Note that for all six points except $(1,1,2g,2g)$, the correction term provided by Lemma 6-18 is useless. However, at $(1,1,2g,2g)$, the uncorrected form is positive, while the corrected one is negative. This vertex corresponds to the linking number of two collections that go as right as possible, that is whose dynamical code is LR^{2g}. It is not a surprise that this vertex is where the form is the least negative, as the linking number of two (non-geodesic) collections whose dynamical code contains only R is positive (such collections are isotopic to a multiple of a fiber in $T^1\Sigma_{2,3,4g+2}$, and two such fibers are positively linked).

Proof of Proposition 6-14 (Case (b) of Theorem A) Since $T^1\Sigma_g$ is a finite cover of $T^1\Sigma_{2,3,4g+2}$, it is enough by Lemma 2-2 to show that the invariant lifts of the families γ and γ' have negative linking number in $T^1\Sigma_g$. By the construction of the template $\mathcal{B}_{4g+2}$ and by Theorem 4-9, these lifts are isotopic to two families $\hat{\gamma}, \hat{\gamma}'$ of periodic orbits of $\mathcal{B}_{4g+2}$. By Lemma 6-20, the reduced linear codes of $\hat{\gamma}, \hat{\gamma}'$ belong to the cone $C_{4g+2} \setminus \{0\}$, and, by Lemma 6-21, the form S_{4g+2} is negative on the pair formed by the two codes. By Lemma 6-18, the linking number between $\mathrm{Lk}(\hat{\gamma}, \hat{\gamma}')$ is then negative, and so is $\mathrm{Lk}(\gamma, \gamma')$. □

Thus the proof of Theorem A is complete.

7 Further questions

We conclude with a few remarks and questions about extensions of the above results. Here we shall both construct counter-examples showing some limitations for possible generalizations and discuss a few plausible conjectures.

7.1 Left-handed flows

We exhibited in Theorem A some hyperbolic orbifolds with no rational homology on which any two orbits of the geodesic flow have a negative linking number. It is natural to ask for further examples of orbifolds with the same property. One could even wonder whether the property could be true for every hyperbolic orbifold. This is *not* the case, and there exist counter-examples on every hyperbolic surface.

Proposition 7-1 *Let Σ_2 be a genus two hyperbolic surface. Then there exist two null-homologous collections γ, γ' of periodic orbits of Φ_{Σ_2} satisfying $\mathrm{Lk}(\gamma, \gamma') > 0$.*

Proof Let γ be the lift of the green collection and γ' be the lift of the orange collection in Figure 30.
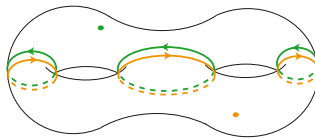

Figure 30

Then the lift of the green vector field is a surface whose boundary is the union of γ and twice the fiber of a point, and which does not intersect γ'. The same vector field on the other pair of pants connects γ' to twice another fiber. Then one checks that the linking number between two fibers is $+\frac{1}{2}$, and we thus obtain $\mathrm{Lk}(\gamma, \gamma') = +2$. □
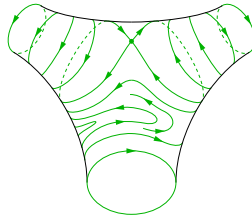
Figure 31

However, let us mention that such counter-examples are rare. Indeed, using the techniques of Section 6 and a computer, we have explored the possible linking numbers of periodic orbits of Φ_{Σ_2} and Φ_{Σ_3}. In a vast majority of cases, the linking number is negative, and the situation of Proposition 7-1 is exceptional. So far we have no explanation for this rarity.

Question 7-2 Let Σ_g be a genus-g hyperbolic surface. Characterize those pairs of collections of periodic orbits of Φ_{Σ_g} that have a positive linking number.

We note that the counter-examples of Proposition 7-1 involve parallel collections of geodesics. A more specific, and maybe more accessible question, could be:

Question 7-3 Let Σ_g be a genus-g hyperbolic surface. If γ, γ' are two collections of periodic orbits of Φ_{Σ_g} whose projections are not parallel and intersect, do we have $\mathrm{Lk}(\gamma, \gamma') \leq 0$?

In another direction, it is natural to wonder whether the assumption of a negative curvature can be dropped. Corollary 2-3 shows that the geodesic flow is also left handed on orbifolds with constant positive curvature, and, although their unit tangent bundle is not a homology sphere, orbifolds with constant zero curvature also yields flows that are left handed in some weak sense (see Theorem B and its corollaries). Nevertheless, one cannot hope for the geodesic flow on every sphere to be left handed.

Proposition 7-4 *If a surface Σ admits at least two separating geodesics that do not intersect, then the geodesic flow Φ_Σ is not left handed.*

Proof Figure 32 corresponds to the case of a sphere whose curvature has a nonconstant sign. The lifts of the two drawn curves are cohomologous, in the complement of the other curve, to a fiber and the opposite of a fiber respectively. Their linking number is $+\frac{1}{2}$. The argument is similar in the general case. $\qquad\square$
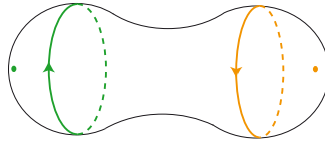
Figure 32

The situation of Proposition 7-4 cannot happen for a sphere with a positive curvature, and we propose:

Conjecture 7-5 Assume that Σ is a 2–sphere with a variable positive curvature. Then the geodesic flow Φ_Σ is left handed.

The particular case of an ellipsoid could be accessible as, in this case, the geodesic flow is integrable.

On the other hand, the counter-example of Proposition 7-1 heavily relies on the fact that the homology of Σ_2 is non-trivial. Therefore the conjecture of Ghys claiming that, if Σ is a hyperbolic 2–orbifold with $H_1(\Sigma, \mathbb{Q}) = 0$, then the geodesic flow Φ_Σ is left handed remains open and plausible.

So, in view of the known results and the above conjectures, the only cases for which the situation is totally unclear are those of orbifolds whose curvature has a non-constant sign and in which any two geodesics intersect, typically a pair of pants capped with three round hemispheres and slightly distorted so that the circles bounding the pants are not geodesic.

7.2 Template knots

The construction of Section 4 associates a (multi)-template with every regular tessellation of $\mathbb{H}^2$. Among all templates arising in this way, it is natural to pay special attention to those associated with the orbifolds $\Sigma_{p,q,\infty}$ of Section 5. In this case, the tiles exclusively are ideal polygons and, therefore, there exists a one-to-one correspondence between the periodic orbits of the template and the periodic geodesics on the orbifold. The knots appearing in this approach generalize Lorenz knots, which correspond to the special case $p = 2$, $q = 3$. Lorenz knots have many interesting properties, and one can wonder whether similar properties could be true for those knots that appear in the above more general setting.

Question 7-6 Which knots appear as periodic orbits of $\Phi_{\Sigma_{p,q,\infty}}$?

In this direction, Pinsky announced [14] that every periodic orbit of $\Phi_{\Sigma_{2,q,\infty}}$ is a prime knot (in the non-compact manifold $T^1\Sigma_{p,q,\infty}$). Also, our current results show that these knots are fibered in $T^1\Sigma_{p,q,\infty}$. So, in particular, all knots cannot appear in this way.

7.3 Gauss linking forms

Let M be 3–manifold. A Gauss linking form on M is a differential $(1,1)$–form whose integral along every pair of null-homologous curves equals their linking number. Gauss linking forms exist on arbitrary 3–manifolds, but explicit formulas are known in very few cases: essentially, the only known examples are the those of [6] for the cases of $\mathbb{S}^3, \mathbb{R}^3$ and $\mathbb{H}^3$.

Now, Ghys' theorem [10] states that a flow is left handed if and only if there exists a Gauss linking form that is negative on the flow. Therefore, Theorem A implies the existence, for the considered orbifolds Σ, of a Gauss linking form in $T^1\Sigma$ that is negative along Φ_Σ. However, our proof of Theorem A gives no indication about the involved Gauss linking forms.

Question 7-7 Are there explicit formulas for the Gauss linking forms implicitly involved in Theorem A?

More generally, better understanding Gauss linking forms appears as a plausible way to address Question 1-2 and Conjecture 7-5.

References

[1] **V I Arnol′d**, *The asymptotic Hopf invariant and its applications*, Selecta Math. Soviet. 5 (1986) 327–345 MR891881

[2] **S Baader**, **J Marché**, *Asymptotic Vassiliev invariants for vector fields*, Bull. Soc. Math. France 140 (2012) 569–582 MR3059851

[3] **G D Birkhoff**, *Dynamical systems with two degrees of freedom*, Trans. Amer. Math. Soc. 18 (1917) 199–300 MR1501070

[4] **J S Birman**, **R F Williams**, *Knotted periodic orbits in dynamical systems, I: Lorenz's equations*, Topology 22 (1983) 47–82 MR682059

[5] **P Dehornoy**, *Invariants topologiques des orbites périodiques d'un champ de vecteurs*, PhD thesis, ÉNS de Lyon (2011)

[6] **D DeTurck**, **H Gluck**, *The Gauss linking integral on the 3–sphere and in hyperbolic 3–space* arXiv:math/0406276

[7] **J-M Gambaudo**, **É Ghys**, *Signature asymptotique d'un champ de vecteurs en dimension 3*, Duke Math. J. 106 (2001) 41–79 MR1810366

[8] **R W Ghrist**, **P J Holmes**, **M C Sullivan**, *Knots and links in three-dimensional flows*, Lecture Notes in Mathematics 1654, Springer, Berlin (1997) MR1480169

[9] **É Ghys**, *Knots and dynamics*, from: "International Congress of Mathematicians, I", (M Sanz-Solé, J Soria, J L Varona, J Verdera, editors), Eur. Math. Soc., Zürich (2007) 247–277 MR2334193

[10] **É Ghys**, *Right-handed vector fields & the Lorenz attractor*, Jpn. J. Math. 4 (2009) 47–61 MR2491282

[11] **U Kaiser**, *Link theory in manifolds*, Lecture Notes in Mathematics 1669, Springer, Berlin (1997) MR1479639

[12] **J M Montesinos**, *Classical tessellations and three-manifolds*, Springer, Berlin (1987) MR915761

[13] **A Ogg**, *Modular forms and Dirichlet series*, Benjamin, New York (1969) MR0256993

[14] **T Pinsky**, *Templates for geodesic flows*, Ergodic Theory Dynam. Systems 34 (2014) 211–235 MR3163031

[15] **S Schwartzman**, *Asymptotic cycles*, Ann. of Math. 66 (1957) 270–284 MR0088720

[16] **W P Thurston**, *The geometry and topology of* 3*–manifolds*, lecture notes, Princeton University (1978–1981) Available at http://library.msri.org/books/gt3m

[17] **J Van Horn-Morris**, *Constructions of open book decompositions*, PhD thesis, University of Texas at Austin (2007)

Institut Fourier, 100 rue des maths
BP74, 38402 St. Martin-d'Hères, France

pierre.dehornoy@ujf-grenoble.fr

http://www-fourier.ujf-grenoble.fr/~dehornop/