Thurston's fibered faces for non-orientable 3-manifolds and an application to minimal stretch factors

Sayantan Khan, Caleb Partin, and Becca Winarski

Abstract

We generalize Thurston norm and the related theory of fibered faces to the setting of non-orientable 3-manifolds. This lets us construct examples of pseudo-Anosov maps on non-orientable surfaces with small stretch factors. Using this, we prove that for a fixed number of punctures, the minimal stretch factor of a genus g non-orientable surface behaves like $\frac{1}{g}$, generalizing the techniques and a result of Yazdi.

Contents

1 Introduction **1**

2 Background **2**
2.1 Mapping classes on non-orientable surfaces 2
2.2 Thurston's fibered faces 5

3 Fibered face theory for non-orientable 3-manifolds **7**
3.1 The problem with homology in non-orientable 3-manifolds 7
3.2 Thurston norm for non-orientable 3-manifolds 9
3.3 Inverting the Poincaré duality map for embedded surfaces 10
3.4 Oriented sums of surfaces 13
3.5 Relating 1-forms and fibrations over S^1 14

4 Minimal stretch factors for non-orientable surfaces with marked points **17**

1 Introduction

Let M be a closed, fibered 3-manifold that fibers over S^1. It is a well known fact that these surface bundles over the circle all have a separate description, that of the mapping torus of some homeomorphism of a closed surface. Not only that, but a single 3-manifold can have many, possibly infinite, descriptions as a mapping torus. Thus when studying these 3-manifolds, an important question is whether one can understand or give a description of all these possible fibrations.

In 1986, Thurston gave a way to answer this question, a semi-norm on the second homology of an orientable 3-manifold that was able to "detect" when an embedded surface of the 3-manifold was the fiber of a fibration of said manifold. This *Thurston norm* is given a full treatment in [Thu86], in which Thurston shows that not only does this norm have unit ball which is a polyhedron, but the fibers of fibrations of one of these 3-manifolds M were in one-to-one correspondence with cones on open faces of this polyhedron unit ball. This almost combinatorial description of the fibers turns out to be a powerful tool in studying these fibered 3-manifolds.

All this also plays a key role in understanding mapping class groups of surfaces, since a surface and a mapping class on it is associated to a 3-manifold, namely, the mapping torus of the mapping class.

By relating the different ways a given mapping torus can fiber, one is able to relate mapping classes on different surfaces, and in doing so, construct mapping classes of interest.

However, there is a small issue with using Thurston norm based techniques: Thurston defines his norm on the second homology of an *orientable* 3-manifold. This leads to the question of whether the results that depend on the Thurston norm work in the non-orientable setting as well. While Thurston himself does comment on this in [Thu86], writing, "(m)ost of this paper works also for non-oriented manifolds, but for simplicity we deal only with the oriented case." It is the goal of this paper to deal with the other case, to see what works and what, if anything, possibly goes wrong when trying to extend the Thurston norm and its consequences for fibered 3-manifolds to the non-orientable setting.

One of the ways Thurston's results are used in the study of mapping classes is via the operation of *oriented sum*. Given the mapping torus M of some surface S and some homeomorphism φ, one can identify S with an embedded surface in M. By picking another embedded surface S' in an appropriate manner, one can perform local surgery to combine S and S': the resulting surface is called the oriented sum of S and S' and denoted $S + S'$. Under the appropriate hypothesis, the oriented sum is also the fiber of some other fibration, and thus has a mapping class φ' on it. It turns out one can relate the stretch factors of φ (which is a mapping class on S) and φ' (which is a mapping class on $S + S'$). We generalize this operation of oriented sum for non-orientable surfaces in Theorem 3.17.

To show that this paper isn't just generalization for the sake of generalization, we use this generalization of Thurston's results and oriented sums to study the asymptotic behaviour of the minimal stretch factor of punctured non-orientable surfaces. The result for orientable punctured surfaces was proven by Yazdi in 2019, and relied heavily on Thurston's result, among others. We are able to essentially "plug in" the non-orientable version of Thurston's results to get a non-orientable version of Yazdi's results. This is one of the main theorems of this paper.

Theorem 4.3. Let $\mathcal{N}_{g,n}$ be a non-orientable surface of genus g with n marked points, and let $l'_{g,n}$ be the smallest stretch factor of the pseudo-Anosov mapping classes acting on $\mathcal{N}_{g,n}$. Then for any fixed $n \in \mathbb{N}$, there are positive constants $B'_1 = B'_1(n)$ and $B'_2 = B'_2(n)$ such that for any $g \geq 2$, the stretch factor satisfies the following inequalities.

$$\frac{B'_1}{g} \leq l'_{g,n} \leq \frac{B'_2}{g}$$

2 Background

2.1 Mapping classes on non-orientable surfaces

Compact non-orientable surfaces with marked points are classified by their genus and number of marked points, just like compact orientable surfaces. A genus g non-orientable surface is the connect-sum of g copies of $\mathbb{RP}^2$, analogous to how a genus g surface is the connect-sum of g copies of a torus $S^1 \times S^1$. A common way of visualizing non-orientable surfaces is to think of them as orientable surfaces with *crosscaps* attached. Attaching a crosscap involves deleting a small open disc D, and gluing the boundary of that disc (on the surface) via the antipodal map. In pictures, this is often denoted by an X inscribed in a circle (see Figure 4 for an example of a surface with two crosscaps attached). A non-orientable surface $\mathcal{N}_{g,n}$ of genus g with n marked points has an orientable double cover, which is an orientable surface $\mathcal{S}_{g-1,2n}$ of genus $g-1$ and $2n$ marked points. Associated to this double cover is an orientation reversing deck transformation $\iota : \mathcal{S}_{g-1,2n} \to \mathcal{S}_{g-1,2n}$. The covering map from $\mathcal{S}_{g-1,2n}$ to $\mathcal{N}_{g,n}$ will be denoted by p. If the genus and the number of marked points is clear from the context, or unimportant, we will drop the subscripts, and just use $\mathcal{N}$ and $\mathcal{S}$ for the non-orientable surface and its orientable double cover respectively.

Given any homeomorphism $\varphi : \mathcal{N} \to \mathcal{N}$ on the non-orientable surface, it is always possible to lift it to a unique orientation preserving homeomorphism $\widetilde{\varphi} : \mathcal{S} \to \mathcal{S}$. This is an easy exercise in covering space theory, but we'll give a proof here for completeness.

Proposition 2.1. *For any homeomorphism $\varphi : \mathcal{N} \to \mathcal{N}$, there exists a unique orientation preserving lift $\widetilde{\varphi} : \mathcal{S} \to \mathcal{S}$. If the non-orientable surface has marked points fixed by φ, then the orientation preserving lift $\widetilde{\varphi}$ may not fix the marked points, but the lift of φ^2 will fix the marked points.*

Proof. One can always lift a homeomorphism $\varphi : \mathcal{N} \to \mathcal{N}$ if φ preserves the subgroup of $\pi_1(\mathcal{N})$ corresponding to the cover $\mathcal{S}$. This subgroup can be concretely described as the subgroup generated by the two sided curves in $\mathcal{N}$, i.e. the curves whose tubular neighbourhoods are cylinders, and not Möbius strips. Such a subgroup is clearly preserved by any homeomorphism φ, which means we always have a lift. There will be two choices for such a lift, since $\mathcal{S}$ is a two-sheeted cover. These two lifts $\widetilde{\varphi}_1$ and $\widetilde{\varphi}_2$ are related by the following identity.

$$\widetilde{\varphi}_1 = \iota \circ \widetilde{\varphi}_2$$

Since ι is orientation reversing, only one of $\widetilde{\varphi}_1$ or $\widetilde{\varphi}_2$ is orientation preserving, which gives us a unique choice.

If $\mathcal{N}$ has marked points that φ fixes, then the lift $\widetilde{\varphi}$ may or may not swap the pre-images of the marked points. But the square of the lift will definitely fix the pre-images as well, which proves the second part of the proposition. □

A consequence of the above proposition is that one can think of the mapping class group of a non-orientable surface as a subgroup of the mapping class group of the double cover. This inclusion also respects the Nielsen-Thurston classification of mapping classes, both qualitatively, and quantitatively, as the following proposition shows.

Proposition 2.2. *If φ is a self-homeomorphism of $\mathcal{N}$ and $\widetilde{\varphi}$ is its orientation preserving lift on $\mathcal{S}$, then the following statements are true.*

(i) φ is periodic, reducible, or pseudo-Anosov if and only if $\widetilde{\varphi}$ is periodic, reducible, or pseudo-Anosov respectively.

(ii) If φ is pseudo-Anosov with stretch factor k, then $\widetilde{\varphi}$ also has stretch factor k.

Proof. It's easy to see that if φ is periodic, so it $\widetilde{\varphi}$, and the other way round. If φ is reducible, that means it leaves some multi-curve on $\mathcal{N}$ invariant, which means $\widetilde{\varphi}$ leaves the lift of that multi-curve invariant as well. Conversely, if $\widetilde{\varphi}$ leaves some multi-curve γ invariant, so does $\iota \circ \widetilde{\varphi}$, since they commute. That means the union of γ and $\iota(\gamma)$ is also a multi-curve and thus descends to a multi-curve on $\mathcal{N}$ that is left invariant by φ. By the process of exclusion, any pseudo-Anosov on $\mathcal{N}$ must lift to a psuedo-Anosov on $\mathcal{N}$ and vice versa. This proves part (i) of the proposition.

Suppose now that φ is a psuedo-Anosov on $\mathcal{N}$ with stretch factor k and expanding and contracting foliations λ_e and λ_c respectively. Since φ is a pseudo-Anosov map, the following identity involving the intersection form i holds for all closed curves γ.

$$\begin{aligned} i(\varphi^{-1}\gamma, \lambda_e) &= i(\gamma, \varphi(\lambda_e)) && (1)\\ &= k \cdot i(\gamma, \lambda_e) && (2) \end{aligned}$$

A similar identity holds for λ_c.

$$\begin{aligned} i(\varphi^{-1}\gamma, \lambda_c) &= i(\gamma, \varphi(\lambda_c)) && (3)\\ &= \frac{1}{k} \cdot i(\gamma, \lambda_c) && (4) \end{aligned}$$

Note now that the foliations can be lifted to the double cover: call their lifts $\widetilde{\lambda}_e$ and $\widetilde{\lambda}_c$. For any closed curve $\widetilde{\gamma}$ on $\mathcal{S}$, consider its intersection number with the foliations. Observe that computing the intersection number is a local calculation. Start by picking an open cover U on $\mathcal{N}$ such that all the open sets in U are homeomorphic to the connected components of their pre-image in $\mathcal{S}$. By picking a partition of unity subordinate to this cover, one can compute the intersection number by restricting computation on each open set in the cover. This calculation lifts to the orientation double cover, giving us the following identity.

$$i(\widetilde{\gamma}, \widetilde{\lambda}_e) = i(\gamma, \lambda_e) \tag{5}$$

Combining identities (1) and (5), we get the following identity for intersection numbers on $\mathcal{S}$.

$$i(\widetilde{\varphi}^{-1}(\widetilde{\gamma}), \widetilde{\lambda}_e) = k \cdot i(\widetilde{\gamma}, \widetilde{\lambda}_e)$$

We get a similar expression for $\widetilde{\lambda}_c$, which proves that $\widetilde{\varphi}$ has the same stretch factor as f, thus proving the proposition. □

Finally, the last thing we need to know about mapping classes on non-orientable surfaces is how to construct examples of pseudo-Anosov maps. In the case of orientable surfaces, the Penner construction is used to construct pseudo-Anosov maps, as well compute their stretch factors. It turns out the Penner construction also works in the non-orientable setting, with some minor modifications. This construction is presented in detail in Section 2 of [Str17], but we'll give an outline of the key ideas.

The Penner construction in the orientable setting takes a pair of filling multicurves $A = \{a_1, \ldots, a_n\}$ and $B = \{b_1, \ldots, b_m\}$ and claims that a composition of Dehn twists T_{a_i} and $T_{b_i}^{-1}$ that uses all curves in the multicurves at least once will be pseudo-Anosov. The problem with making this work for non-orientable surfaces is that when defining Dehn twists about curves on a non-orientable surface, we don't have a well-defined notion of a left or right Dehn twist (i.e. which direction is the Dehn twist, and which direction is the inverse Dehn twist). The way we get around this for non-orientable surfaces is by having a collection of filling two-sided curves that are *marked inconsistently*.

Each two-sided curve c on a non-orientable surface N has a neighborhood homeomorphic to an annulus A by a homeomorphism $\phi : A \to N$, called a *marking*. In this context, we can define the Dehn twist $T_{c,\phi}(x)$ around (c, ϕ) in the following manner.

$$T_{c,\phi}(x) = \begin{cases} \phi \circ T \circ \phi^{-1}(x) & \text{for } x \in \phi(A) \\ x & \text{for } x \in N - \phi(A) \end{cases}$$

Here T is the standard Dehn twist on A, i.e. $T(\theta, t) = (\theta + 2\pi t, t)$. If we fix an orientation of A, then we can pushforward this orientation to S. We say two marked curves (c, ϕ_c) and (d, ϕ_d) that intersect at a point p are marked inconsistently if the pushforward of the orientation of A by ϕ_c and ϕ_d disagree in a neighborhood of p. If all our curves are marked inconsistently and are filling, then once again a composition of Dehn twists around them that use all the curves at least once will be pseudo-Anosov.

The Penner construction not only promises that our map is pseudo-Anosov, but it also gives a way to compute the stretch factor of our map (see [Pen88]). The proof of the fact that the composition is pseudo-Anosov, and the computation of its stretch factor works the same is in the orientable setting. Given the collection of curves, one smooths out the intersections of the curves to obtain an *invariant train track*, an embedded graph on the surface that remains unchanged by the homeomorphism, which we'll call f. Let τ be the invariant train track and $\mathcal{C}$ the collection of curves on our surface that are used when defining f. Consider now the collection of transverse measures on our train track τ. For every curve (i.e. a single element subset of $\mathcal{C}$) $x \subset \mathcal{C}$, there is an associated transverse measure μ_x for τ that assigns 1 to all edges lying in x and 0 to everything else. Let V_τ be the cone of transverse measures on τ, and H the subspace of V_τ spanned by the transverse measure associated to curves in $\mathcal{C}$.

$$H = \text{span}(\{\mu_x \mid x \text{ is a connected curve in } \mathcal{C}\})$$

The μ_x are linearly independent and form the *standard basis* for H. This subspace H is invariant under the action of f on V_τ, thus f has a linear action on H. If we let A be the matrix representing this action in the standard basis, then the stretch factor of f, $\lambda(f)$, is the Perron-Frobenius eigenvalue of f.

2.2 Thurston's fibered faces

Given a surface S and a self homeomorphism $\varphi : S \to S$, one can construct a 3-manifold M via the *mapping torus* construction.

$$M := \frac{S \times [0,1]}{(x,1) \sim (\varphi(x),0)}$$

Inverting this construction is also a problem of interest: given a 3-manifold M, is it the mapping torus of some surface and self-homeomorphism (S,φ)? In how many ways can one express a 3-manifold as a mapping torus? Another way to think about these mapping tori is by observing that these correspond exactly to surface bundles over S^1, or *fibrations over* S^1. The surface S then is just the *fiber* of the fibration. Such a fibration also defines a flow on M, called the *suspension flow*, where (x,t_0) gets sent to (x,t_0+t), where x is a point in the fiber S, and t_0 is a point on S^1. It turns out that the Thurston norm is an extremely useful tool when studying fibrations over S^1 (and other related objects).

Given a orientable closed 3-manifold M, the Thurston norm is a semi-norm on its second homology group $H_1(M;\mathbb{R})$: to define the norm, we need to make some preliminary remarks and define a function. Let $\chi_-(S)$ denote the negative component of the Euler characteristic of a surface S, i.e. $\chi_-(S) = \max\{-\chi(S),0\}$, we call this the *complexity* of S. If a surface S has multiple components then its complexity is the sum of the χ_- for the individual components. Furthermore, for an *orientable* 3-manifold M, it is known that elements in $H_2(M;\mathbb{Z})$ can be represented by embedded surfaces inside of M. This lets us define a norm function x for every homology class a in $H_2(M;\mathbb{Z})$.

$$x(a) = \min\{\chi_-(S) \mid [S] = a \text{ and } S \text{ is compact, properly embedded and oriented}\}$$

One can then extend x to the rational points by claiming it must be linear on rays that go through the origin, and there is a unique way to extend a function defined on rational points to reals so that it is continuous. This defines an $\mathbb{R}$-valued function on $H_2(M;\mathbb{R})$ which is called the Thurston norm. It turns out that the unit ball in this norm is a convex polyhedron, and thus it makes sense to talk about the *faces* of the unit ball.

The way this ties back up with the question of realizing the 3-manifold M as a fibration over S^1 is via the following remarkable theorem of Thurston (the original statement appeared in [Thu86], but we use the restatement from [Yaz18] because it's cleaner).

Theorem 2.3 (Thurston)**.** *Let $\mathcal{F}$ be the set of homology classes in $H_2(M;\mathbb{R})$ that are representable by fibers of fibrations of M over the circle.*

- *(i) Elements of $\mathcal{F}$ are in one-to-one correspondence with (non-zero) lattice points inside some union of cones over open faces of the unit ball in the Thurston norm.*
- *(ii) If a surface F is transverse to the suspension flow associated to some fibration of $M \to S^1$ then $[F]$ lies in the closure of the corresponding cone in $H_2(M)$.*

Note that one needs to clarify here what the homology class $[F]$ actually is, since given an embedded surface F, one has two canonical choices of $[F]$ depending upon the orientation of F. The class $[F]$ referred to in the theorem comes from picking the orientation such that the positive flow direction is pointing outwards relative to the surface.

The open faces whose cones contain the fibers of fibrations are what are referred to as *fibered faces*. It turns out one can recover even more information about the self homeomorphism f of which M is the mapping torus. The following result of Thurston tells us precisely when the map f is psuedo-Anosov.

Theorem 2.4 (Thurston's Hyperbolization Theorem)**.** *If M is the mapping torus of (S,f), then M is hyperbolic if and only if f is pseudo-Anosov.*

In particular, a consequence of the above theorem is that if M is the mapping torus of (S, f) and (S', f') then f' is pseudo-Anosov if and only if f is pseudo-Anosov. From Theorem 2.3, we know that all such pairs (S, f) correspond to lattice points in some union of cones. A natural thing to do at this point would be to determine how the stretch factors of these pseudo-Anosov maps vary as we move around in the cone. To this end, we have the following two theorems, due to Fried-Matsumoto (see [Fri82], [Fri83], and [Mat87]) and Agol-Leininger-Margalit (see [ALM]).

Theorem 2.5 (Fried-Matsumoto)**.** *Let M be a hyperbolic 3-manifold and let $\mathcal{K}$ be the union of cones in $H_2(M;\mathbb{R})$ whose lattice points correspond to fibrations over S^1. There exists a strictly convex function $h : \mathcal{K} \to \mathbb{R}$ satisfying the following properties.*

(i) For all $t > 0$ and $u \in \mathcal{K}$, $h(tu) = \frac{1}{t}h(u)$.

(ii) For every primitive lattice point $u \in \mathcal{K}$, $h(u) = \log(k)$, where k is the stretch factor of the pseudo-Anosov map associated to this lattice point.

(iii) $h(u)$ goes to ∞ as u approaches $\partial\mathcal{K}$.

Theorem 2.6 (Agol-Leininger-Margalit)**.** *Let $\mathcal{K}$ be a fibered cone for a mapping torus M and let $\overline{\mathcal{K}}$ be its closure in $H_2(M;\mathbb{R})$. If $u \in \mathcal{K}$ and $v \in \overline{\mathcal{K}}$, then $h(u + v) < h(u)$.*

Theorems 2.5 and 2.6 make a quantitative link between the lattice points in the cones over the fibered faces and the stretch factor of the associated pseudo-Anosov maps, demonstrating the power of Thurston's fibered face theory.

It's also useful to know if an embedded surface S minimizes the Thurston norm in its homology class. That ends up happening if the embedded surface is incompressible.

Definition 2.7 (Incompressible surface)**.** A surface S (such that S is not a sphere) inside a 3-manifold M is said to be incompressible if there exists no embedded disc D in M such that $D \cap S = \partial D$, and D intersects S transversally.

The following theorem is due to Thurston.

Theorem 2.8. *A surface S minimizes the Thurston norm in its homology class iff it is incompressible.*

The following two families of examples of incompressible surfaces will be fairly important for us.

Example 2.9. Let M be a hyperbolic 3-manifold that fibers over S^1. Then an embedded surface S is incompressible if either of the two following conditions hold.

(i) S is the fiber of a fibration.

(ii) S has genus 2. In this case, S is incompressible, because if it weren't norm minimizing in its homology class, a genus 1 or genus 0 surface would be. But one can't have genus 0 or 1 incompressible surfaces in hyperbolic 3-manifolds.

The first example encompasses a large class of incompressible surfaces, as demonstrated by the following theorem.

Theorem 2.10 (Theorem 4 from [Thu86])**.** *Any incompressible surface S in the homology class of a fiber of a fibration is isotopic to the fiber.*

One of the primary goals of this paper to extend this definition of Thurston norm to non-orientable manifolds and be able to state the non-orientable version of the theorems above. Most of the work in doing so is concentrated in determining the right analog of Thurston norm for non-orientable surfaces, and then making Theorem 2.3 work with that definition. Once we have that, the rest of the theorems follow with relatively little work. That is what we will do in Section 3. Once we have the versions of the theorems for non-orientable surfaces, we'll generalize to non-orientable surfaces a trick due to McMullen that lets one construct pseudo-Anosov maps with small stretch factor. That will be used to prove bounds on the asymptotics of minimal stretch factors for non-orientable surfaces in Section 4 by adapting the methods in [Yaz18] for non-orientable surfaces.

3 Fibered face theory for non-orientable 3-manifolds

In what follows, we will restrict our attention our attention on two kinds of compact 3-manifolds: mapping tori of orientable surfaces with a pseudo-Anosov map (these will be the 3-manifolds we will be referring to when talking about orientable 3-manifolds), and the mapping tori of non-orientable 3-manifolds, again with a pseudo-Anosov map (these will be the manifolds we will be referring to when talking about non-orientable 3-manifolds). While a lot of our statements will hold more generally for compact non-orientable 3-manifolds, it will be easier to describe examples when working in this restricted setting; additionally, our application will only involve mapping tori of pseudo-Anosov maps.

3.1 The problem with homology in non-orientable 3-manifolds

A first attempt at defining the Thurston norm for a compact non-orientable 3-manifold might go as follows: for any embedded surface S (whether it's orientable or not), one defines the complexity function χ_-, much like in the case of orientable 3-manifolds, and then define the norm of a homology class $a \in H^2(M;\mathbb{Z})$ by minimizing $\chi_-(S')$ over all S' representing a. This may work, but is quite unsatisfying: this construction assigns 0 norm to all embedded non-orientable surfaces, since their fundamental classes are trivial, and thus map to 0 in $H_2(M;\mathbb{Z})$. This is bad, because we'd like incompressible surfaces in non-orientable 3-manifolds to have a positive norm. There are plenty of incompressible surfaces even in non-orientable 3-manifolds, namely fibers of fibrations over S^1.

It turns out that the fundamental problem with non-orientable 3-manifolds is that homology is a very coarse invariant: too coarse to detect embedded non-orientable surfaces. Our workaround will be to deal with the first cohomology $H^1(M)$ rather than the second homology $H_2(M)$. The two are same in the case of orientable 3-manifolds, by Poincaré duality, but that fails in the case of non-orientable 3-manifolds.

To see why it fails, consider an orientable 3-manifold M. We can explicitly work out the map from $H^1(M;\mathbb{Z})$ to $H_2(M;\mathbb{Z})$ given by Poincaré duality. To do so, we set up a correspondence between elements of $H^1(M;\mathbb{Z})$ and homotopy classes of maps from M to S^1. Given a cohomology class $[\alpha]$ in $H^1(M;\mathbb{Z})$, pick a representative 1-form α, and a basepoint x_0 in M. The associated map f_α is given by the following formula

$$f_\alpha(x) := \int_{x_0}^{x} \alpha \Big/ \mathbb{Z}$$

Changing the basepoint or the representative 1-form gives a different map to S^1 that is homotopic to the original map (see Section 5.1 of [Cal07] for the details). One can recover the 1-form α from the map f_α by pulling back the canonical volume form $d\theta$ on S^1 along f_α.

Consider now the surface $S = f_\alpha^{-1}(p)$ for some regular value $p \in S^1$. Note that we have S merely as a subset of M right now. To get a homology class, we pick an orientation on S by declaring that the outwards pointing normal vectors on S are the ones which get assigned a positive value by the form α. This defines an orientation of S because we already have an orientation on M, and thus defines a fundamental class $[S]$. We claim that $[S]$ is the Poincaré dual to α. More precisely, we claim the following.

Claim 3.1. *Let p and p' be two regular values of the function f_α and let S and S' be $f_\alpha^{-1}(p)$ and $f_\alpha^{-1}(p')$ respectively. Then for any closed 2-form ω on M, the following identity holds.*

$$\int_S \omega = \int_{S'} \omega$$

Furthermore, the following identity also holds.

$$\int_S \omega = \int_M \alpha \wedge \omega$$

In particular, the homology class of S is Poincaré dual to α.

Proof. The first part of the claim follows from the fact that S and S' are homologous, i.e. $f_\alpha^{-1}([p,p'])$ is a singular 3-chain that has S and S' as boundaries. From Stokes' theorem, we get the following.

$$\begin{aligned}\int_{S-S'} \omega &= \int_{f_\alpha^{-1}([p,p'])} d\omega \\ &= 0\end{aligned}$$

To prove the second claim, observe that we can break up the second integral as a product integral.

$$\int_M \alpha \wedge \omega = \int_{S^1} \left(\int_{f_\alpha^{-1}(x)} \omega \right) d\theta$$

The above equation is true because α is the pullback of dx along the map f_α. Observe that the inner integral only makes sense when x is a regular value, but by Sard's theorem, almost every $x \in [0,1]$ is a regular value, so the right hand side is well-defined. By the first part of the claim, the inner integral is a constant function, as we vary over the x which are regular values of f_α, and the integral of $d\theta$ over S^1 is 1, giving us the identity we want.

$$\int_M \alpha \wedge \omega = \int_S \omega$$

□

What we have here is an explicit formula for the Poincaré duality map. For orientable 3-manifolds, this is an isomorphism, and more specifically the following theorem is true.

Theorem 3.2 (Poincaré duality for orientable 3-manifolds)**.** *Let M be an orientable 3-manifold, and let S be an oriented embedded surface. Then there exists a 1-form α such that $S = f_\alpha^{-1}(p)$ for some regular value p.*

Note however that this technique of getting from a 1-form to an embedded surface to the homology class of that surface still makes sense for a non-orientable 3-manifold M. In that case, one can see that the map from $H^1(M;\mathbb{Z})$ to $H_2(M;\mathbb{Z})$ has a kernel.

Example 3.3 (Failure of Poincaré duality for non-orientable 3-manifolds)**.** Let $\mathcal{N}$ be a non-orientable surface, and let φ be any self homeomorphism. Let M be the mapping torus of $(\mathcal{N}, f)$. We then have a map $f : M \to S^1$ given by mapping to the base of the mapping torus. Pulling back the form $d\theta$ along this map, we get a closed but not exact 1-form α on M. Observe that $f_\alpha = f$, because of how we constructed α. Furthermore $f_\alpha^{-1}(0)$ is $\mathcal{N}$ inside M. Thus, the "Poincaré duality map" for M maps a non-trivial element $\alpha \in H^1(M;\mathbb{Z})$ to the zero element $[\mathcal{N}] \in H_2(M;\mathbb{Z})$, since $\mathcal{N}$ is non-orientable. In particular, we end up losing information going from $H^1(M)$ to $H_2(M)$.

The above example also suggests an alternative method of defining the Thurston norm for non-orientable 3-manifolds: rather than working with $H_2(M;\mathbb{R})$, we can instead work with $H^1(M;\mathbb{R})$, because we don't want to lose information by going to $H_2(M;\mathbb{R})$. We'll also be interested in getting a partial inverse for this map. More specifically, given a non-orientable surface $\mathcal{N}$ inside a non-orientable 3-manifold M, we'd like to understand if M can be realized as the mapping torus of some homeomorphism of $\mathcal{N}$.

3.2 Thurston norm for non-orientable 3-manifolds

For this section, we'll use M to denote a non-orientable 3-manifold, and $\widetilde{M}$ to denote its orientation double cover. We will denote by ι the orientation reversing deck transformation of $\widetilde{M}$, and the covering map $\widetilde{M} \to M$ by p. If M is the mapping torus of the non-orientable surface $\mathcal{N}$ and a pseudo-Anosov map $\varphi : \mathcal{N} \to \mathcal{N}$, then $\widetilde{M}$ is the mapping torus of $(\mathcal{S}, \widetilde{\varphi})$, where $\mathcal{S}$ is the orientable double cover of $\mathcal{N}$, and $\widetilde{\varphi}$ is the orientation preserving lift of φ.

Since we've already concluded that the first cohomology is the "right" space to define the Thurston norm on, we need to relate $H^1(M;\mathbb{R})$ and $H^1(\widetilde{M};\mathbb{R})$. The obvious thing to do is to look at the pullback via p.

Lemma 3.4. *The pullback p^* maps $H^1(M;\mathbb{R})$ bijectively to the ι^*-invariant subspace of $H^1(\widetilde{M};\mathbb{R})$.*

Proof. Clearly, for any 1-form α on M, $p^*(\alpha)$ will be ι^*-invariant. This means that the image of p^* lands inside the ι^*-invariant subspace. To see that the map is injective at the level of H^1 (rather than at the level of 1-forms), consider a 1-form α on M such that $p^*\alpha$ is exact. We thus have a smooth function f on $\widetilde{M}$ such that the following relation holds.

$$df = p^*\alpha$$

But since $p^*\alpha$ is ι^*-invariant, we must have $df = \iota^* df$, and by pushing the ι^* inside, we get that $df = d(\iota^* f)$. That means f and $\iota^* f$ differ by a constant, but that constant must be 0 since ι is finite order. This shows that f descends to a function on M, and thus α is exact, which proves injectivity of p^*. Now we show surjectivity. Let $[\alpha]$ be an element in $H^1(\widetilde{M}, \mathbb{R})$ that is ι^*-invariant. That means if we pick a 1-form α in this equivalence class, the following identity holds for some smooth function f.

$$\alpha - \iota^*(\alpha) = df$$

Note that this means $\iota^* df = -df$. Using these two identities, it's easy to verify that the 1-form $\alpha - \frac{df}{2}$ is ι^*-invariant, and thus in the image of p^*. This proves surjectivity, and the lemma. ☐

The above lemma tells us that $H^1(M;\mathbb{R})$ is a subspace of $H^1(\widetilde{M};\mathbb{R})$ (and we know precisely what that subspace is), and thus the Thurston norm on $H^1(\widetilde{M};\mathbb{R})$ can be restricted to a norm on $H^1(M;\mathbb{R})$.

Definition 3.5 (Thurston norm for non-orientable 3-manifolds)**.** The Thurston norm x is a norm on $H^1(M;\mathbb{R})$, defined using the Thurston norm $\widetilde{x}$ on $H^1(\widetilde{M};\mathbb{R})$ (identified with $H_2(M;\mathbb{R})$ via Poincaré duality) in the following manner.

$$x(\alpha) := \widetilde{x}(p^*\alpha)$$

Now that we have a Thurston norm on $H^1(M;\mathbb{R})$, we need to describe some of its properties. All of these properties follow fairly easily from the orientable version.

Theorem 3.6. *The unit ball with respect to the dual Thurston norm on $\left(H^1(M;\mathbb{R})\right)^*$ is a polyhedron whose vertices are lattice points $\{\pm\beta_1, \dots \pm \beta_k\}$. The unit ball B_1 with respect to Thurston norm is a polyhedron given by the following inequalities.*

$$B_1 = \{a \mid |\beta_i(a)| \leq 1 \text{ for } 1 \leq i \leq k\}$$

Proof. The proof is identical to the original proof of Theorem 2 of Thurston in [Thu86]. The key ingredient of the proof is that the norm of any element in $H^1(M;\mathbb{Z})$ is an integer. That is true in our case because the norm of an element of $H^1(M;\mathbb{Z})$ is the Thurston norm of the corresponding element in $H^1(\widetilde{M};\mathbb{Z})$, which is an integer because it's the negative Euler characteristic of an embedded surface. The rest of the proof is just a matter of linear algebra, and works just as well in our setting. ☐

Observe that the way we defined the Thurston norm for non-orientable 3-manifolds is lacking in two ways. First of all, in the orientable case, the Thurston norm is a norm on the second homology, and thus also embedded surfaces. In other words, the Thurston norm tells us something about embedded surfaces in the manifold. We've already seen how working with second homology doesn't quite work, which is why we had to go to first cohomology instead. We would still like to talk about the norm of an embedded surface though, even if the homology class of that surface may be trivial. This is something we'll see in subsection 3.3.

The other issue is that when working with fibrations over S^1, the elements of $H^1(M;\mathbb{Z})$ are the elements of interest, rather than things in $H^1(M;\mathbb{R})$. Lemma 3.4 tells us that elements of $H^1(M;\mathbb{R})$ are precisely the ι^*-invariant elements of $H^1(\widetilde{M};\mathbb{R})$. That does not hold for $\mathbb{Z}$-coefficients: there are ι^*-invariant elements on $H^1(\widetilde{M};\mathbb{Z})$ that are not pullbacks of elements of $H^1(M;\mathbb{Z})$.

Example 3.7 (Failure of surjectivity)**.** Let $\mathcal{N}$ be a non-orientable surface, $\mathcal{S}$ its orientation double cover. Let γ be a one-sided curve on $\mathcal{N}$, i.e. a curve whose lift to $\mathcal{S}$ is an arc, and the pre-image is a single closed curve, which is twice as long. Call the pre-image $\widetilde{\gamma}$. Let the 3-manifolds M and $\widetilde{M}$ we're considering be the mapping tori of $\mathcal{N}$ and $\mathcal{S}$ with respect to some pseudo-Anosov map. We can then consider γ and $\widetilde{\gamma}$ as curves inside the 3-manifolds M and $\widetilde{M}$.

Pick a basis of $H_1(\widetilde{M};\mathbb{Z})$ containing $\widetilde{\gamma}$. Using this basis, we can construct an element of $H^1(\widetilde{M};\mathbb{Z})$ by simply assigning integer values to the basis elements. Pick an element α that assigns 0.5 to $\widetilde{\gamma}$ and an integer value to every other basis element. Consider the cohomology class $\alpha + \iota^*\alpha$. Since $\iota\widetilde{\gamma} = \widetilde{\gamma}$ (because $\widetilde{\gamma}$ is the pre-image of a curve of M), $\alpha + \iota^*\alpha$ is an ι^*-invariant element of $H^1(\widetilde{M};\mathbb{Z})$ that assigns 1 to $\widetilde{\gamma}$. Such a cohomology class cannot be a pullback of a class on M since the pullback of a cohomology class on M would assign an even value to $\widetilde{\gamma}$.

What the above example does show is that for any ι^*-invariant α in $H^1(\widetilde{M};\mathbb{Z})$, 2α definitely is a pullback of class in $H^1(M;\mathbb{Z})$.

3.3 Inverting the Poincaré duality map for embedded surfaces

In both the orientable and non-orientable setting, we have a way of assigning an embedded surface to an element α in $H^1(M;\mathbb{Z})$ by looking at $f_\alpha^{-1}(p)$, where p is a regular value. In the orientable setting, the homology class of this embedded surface is well defined, independent of the choice of representative 1-form in its cohomology class, as well as the choice of regular value. That's true in the non-orientable setting as well, but the homology class is not very useful. We'd like to invert this construction, i.e. given an embedded surface S, we want a closed 1-form α such that the surface S comes from α in the manner described above.

For an orientable 3-manifold M, this is just a consequence of Poincaré duality, so the only work we need to do is prove this for non-orientable manifolds. If we do so, we can talk about the norm of an embedded surface, and more generally, have an ad hoc version of Poincaré duality, associating embedded surfaces to 1-forms. However, we need a version of the orientability hypothesis to keep things well defined. For an embedded surface S in a non-orientable 3-manifold M, we define the notion of *relative orientability.*

Definition 3.8 (Relative orientability)**.** Let M be a 3-manifold, and S an embedded surface in M. S is said to be relatively oriented with respect to M if there is a nowhere vanishing normal vector field on S. Two such normal vector fields are said to induce the same orientation if locally they induce the same orientation after picking a local frame for S. A surface S is *relatively oriented* if both S and the choice of positive normal vector field are specified.

Note that relative orientability is a strictly weaker notion than orientability. If S and M are orientable, then S is relatively orientable with respect to M. But even if M is non-orientable, S may be relatively orientable with respect to M. For instance, if S is the fiber of a non-orientable mapping torus, then one can get a non-vanishing normal vector field by looking at the pre-image of a non-vanishing vector field on S^1. It is not the case that every embedded surface in a non-orientable 3-manifold is relatively orientable.

Example 3.9 (Relatively non-orientable surface)**.** Let S be the standard torus, i.e. $\mathbb{R}^2/\mathbb{Z}^2$, and let φ map (x, y) to $(-x, y)$. This is an orientation reversing homeomorphism, and thus the mapping torus of this homeomorphism will be a non-orientable 3-manifold M. Consider a vertical line in S preserved by φ, i.e. the line $x = 0$. The image of this closed curve in S under the suspension flow in M is a subsurface of M, which we'll call S'. The normal direction to S' when restricted to S is $\frac{\partial}{\partial x}$. But this can't be continuously extended to all of M, since that direction gets reversed by the suspension flow. This means that the surface S' is relatively non-orientable in M (despite being orientable itself.)

We do however get relative orientability for free when working with only non-orientable surfaces and 3-manifolds.

Proposition 3.10. *Let M be a non-orientable 3-manifold, and let S be an embedded connected non-orientable surface in M. Then S is relatively orientable with respect to M.*

Proof. Let $\widetilde{M}$ be the orientation double cover of M, and $\widetilde{S}$ be the pre-image of S. Then the restriction of the orientation reversing deck transformation ι to $\widetilde{S}$ is orientation reversing on $\widetilde{S}$ as well, since the quotient S is non-orientable. That means S leaves the outward pointing normal direction from $\widetilde{S}$ invariant, and that descends to an outward pointing normal direction on S. This shows that S is relatively orientable with respect to M. $\square$

We care about relatively orientable surfaces because for these surfaces can be mapped to cohomology classes.

Theorem 3.11 (Poincaré duality for non-orientable 3-manifolds)**.** *Let M be a non-orientable 3-manifold, and let S be a relatively oriented embedded surface. Then there exists a cohomology class $[\alpha]$ in $H^1(M;\mathbb{Z})$ such that for some representative α, S is $f_\alpha^{-1}(p)$ for some regular value $p \in S^1$. Furthermore, α assigns positive values to the positively oriented normal vector field on S.*

The idea of the proof of this theorem is fairly straightforward. Starting with the embedded surface S in M, we look at the pre-image $\widetilde{S}$ in the orientation double cover $\widetilde{M}$. We show that the Poincaré dual to $\widetilde{S}$ is ι^*-invariant.

Lemma 3.12. *Let S be a relatively oriented embedded surface in M, and $\widetilde{S}$ its pre-image in $\widetilde{M}$. Then the Poincaré dual to $[\widetilde{S}]$ is ι^*-invariant.*

Proof. Observe that if S comes with a relative orientation, then $\widetilde{S}$ inherits that relative orientation. Since $\widetilde{S}$ and $\widetilde{M}$ are orientable, this defines an orientation on $\widetilde{S}$, and thus the homology class $[\widetilde{S}]$ is well defined.

Observe now that the deck transformation ι reverses the orientation on $\widetilde{S}$. To see why this is the case, pick a local frame (v_1, v_2, v_3) around some point in $\widetilde{S}$ such that v_3 is the outwards pointing normal vector field. Since the outwards pointing normal vector field descends to the quotient by the orientation reversing map ι, that means $\iota(v_3)$ must also be outwards pointing (and not inwards pointing). If ι has to reverse the orientation on $\widetilde{M}$, it must do so by reversing the orientation on the sub-basis (v_1, v_2). In particular, that means ι reverses the orientation on $\widetilde{S}$.

This means $[\widetilde{S}]$ is in the -1-eigenspace of the ι_* action on $H_2(\widetilde{M};\mathbb{R})$. Therefore the Poincaré dual to $[\widetilde{S}]$, which we'll call $\widetilde{\alpha}$, is in the 1-eigenspace, i.e. ι^*-invariant. This just follows from the following chain of equalities which hold for all closed 2-forms ω.

$$\begin{aligned}
\int_{\iota_*\widetilde{S}} \omega &= \int_{\widetilde{S}} \iota^*\omega && \text{(By a change of variables)} \\
&= \int_{\widetilde{M}} \widetilde{\alpha} \wedge \iota^*\omega && \text{(Poincaré duality)} \\
&= \int_{\widetilde{M}} \iota^*\left(\iota^*\widetilde{\alpha} \wedge \omega\right) && \\
&= \int_{\widetilde{M}} -\left(\iota^*\widetilde{\alpha} \wedge \omega\right) && \text{(ι is orientation reversing)}
\end{aligned}$$

On the other hand, the following equalities follow from the fact that $\iota_*[\widetilde{S}] = -[\widetilde{S}]$.

$$\begin{aligned}
\int_{\iota_*\widetilde{S}} \omega &= -\int_{\widetilde{S}} \omega \\
&= -\int_{\widetilde{M}} \widetilde{\alpha} \wedge \omega
\end{aligned}$$

Comparing the right hand side of the two equations, it follows that $\widetilde{\alpha}$ is ι^*-invariant. ☐

We now have an ι^*-invariant 1-form $\widetilde{\alpha}$ such that $\widetilde{S}$ is $f_{\widetilde{\alpha}}^{-1}(p)$ for some regular value p on S^1. The next claim we want to make is that the map $f_{\widetilde{\alpha}} : \widetilde{M} \to S^1$ factors through the quotient M.

Lemma 3.13. *The map $f_{\widetilde{\alpha}}$ factors through M, i.e. for all points $x \in \widetilde{M}$, $f_{\widetilde{\alpha}}(x) = f_{\widetilde{\alpha}}(\iota(x))$.*

Proof. Recall that $f_{\widetilde{\alpha}}(x)$ is given by the following integral formula.

$$f_{\widetilde{\alpha}}(x) = \int_{x_0}^{x} \widetilde{\alpha} \Big/ \mathbb{Z}$$

Here x_0 is some arbitrarily chosen basepoint on $\widetilde{M}$. For $f_{\widetilde{\alpha}}(x)$ to equal $f_{\widetilde{\alpha}}(\iota(x))$ for all x, we must have the following.

$$\left(\int_{x_0}^{x} \widetilde{\alpha} - \int_{x_0}^{\iota(x)} \widetilde{\alpha}\right) \in \mathbb{Z}$$

By a change of variables, and using the ι^*-invariance of $\widetilde{\alpha}$, the left hand side of the above condition can be transformed, giving us the following condition.

$$\left(\int_{x_0}^{\iota(x_0)} \widetilde{\alpha}\right) \in \mathbb{Z} \tag{6}$$

In other words, we want the integral of $\widetilde{\alpha}$ along any curve γ from x_0 to $\iota(x_0)$ to be an integer. Equivalently, it will suffice to show that the integral of $\widetilde{\alpha}$ along δ is an even integer, where δ is the closed curve obtained by taking the union of γ and $\iota(\gamma)$.

Recall now that the 2-parity of $\displaystyle\int_{\delta} \widetilde{\alpha}$ is precisely the 2-parity of the intersection number of δ and $\widetilde{S}$ as long as all the intersections are transversal. Furthermore, both δ and $\widetilde{S}$ are lifts of a curve and surface from M. Which means the number of intersections they have in $\widetilde{M}$ is twice the number of intersections have in M. But the latter number must be an integer, and thus the former number must be an even integer, showing that condition (6) holds. In particular, this shows that the map $f_{\widetilde{\alpha}}$ factors through, proving the lemma. ☐

We now have everything we need to finish proving Theorem 3.11.

Proof of Theorem 3.11. Starting with a relatively oriented surface S in M, we look at its pre-image $\widetilde{S}$ in $\widetilde{M}$. The relative orientation of the pre-image gives us the homology class $[\widetilde{S}]$, and we get a 1-form $\widetilde{\alpha}$, which is Poincaré dual to the homology class of $\widetilde{S}$. More specifically, we have that $\widetilde{S}$ is $f_{\widetilde{\alpha}}^{-1}(p)$ for some regular value p. By lemma 3.12, $\widetilde{\alpha}$ is ι^*-invariant, and by lemma 3.13, we have that $f_{\widetilde{\alpha}}$ factors through M, which means $f_{\alpha}^{-1}(p) = S$. Furthermore, we also have that α lies in $H^1(M;\mathbb{Z})$. That just follows from the fact that α is the pullback of the form $d\theta$ on S^1 along the map f_α. This proves the result. $\square$

3.4 Oriented sums of surfaces

We now have a way of going from an embedded surface to an element of $H^1(M;\mathbb{Z})$. To make this mapping even more useful, we'll describe a way of adding two surfaces via the operation of taking *oriented sums*: this will be additive in two senses: the Euler characteristic, and the cohomology class of the dual. This operation is well-known in the case of orientable 3-manifolds (along with orientable embedded surfaces), but we will sketch out the relevant details for completeness. The same construction works for relatively orientable surfaces; one just needs to verify consistency.

Oriented sum for oriented manifolds Let S and S' be oriented embedded surfaces in an oriented manifold M. Assume that S and S' intersect trasversally. Thus, $S \cap S'$ is a disjoint union of copies of S^1. Pick a small neighbourhood of each component of the intersection such that the cross section looks like Figure 1.

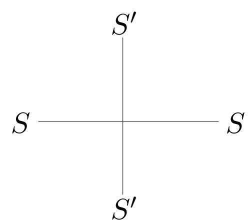

Figure 1: Cross section of intersection of S and S'.

We then perform a local surgery such that each S leaf joins an S' leaf. We have two possible choices: we could join the left S leaf to either the top or the bottom S' leaf. Since both S and S' are oriented submanifolds of M, there is an outward pointing normal vector field on S and S'. Suppose the outward normal vector field on S points upwards and the outward normal vector on S' points to the right. In that case, we'd glue the left S leaf to the bottom S' leaf to maintain a consistent outward normal vector field. See Figure 2 to see how the choice affects orientability.

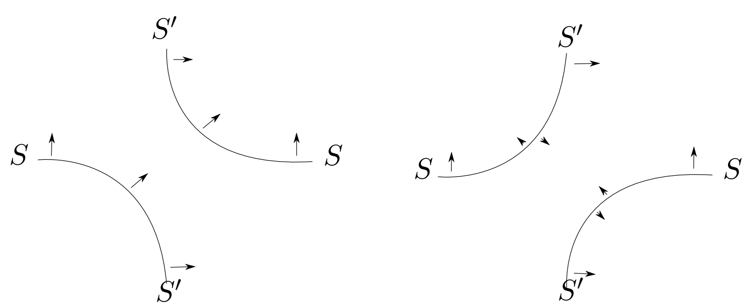

Figure 2: On the left, the normal vectors on S and S' are consistent. On the right, they aren't.

By performing this surgery at all the intersections, we get a new submanifold S'' (which may have multiple components). This new submanifold S'' is the oriented sum of S and S'. The operation of taking oriented sums adds the Euler characteristic, as well as the homology classes (and thus the cohomology classes of their Poincaré duals).

$$\chi(S'') = \chi(S) + \chi(S')$$
$$[S''] = [S] + [S']$$

Oriented sum for non-orientable manifolds Observe that in order to canonically choose the right leaves to join, all we needed was a relative orientation for both S and S'. That suggests that the same construction ought to work. Like in the case of an orientable ambient manifold, at every transversal intersection, we perform surgery based on the outwards pointing normal vector field. We just need to verify that this construction is consistent with the covering map: i.e. taking the oriented sum of S and S' is the same as taking the oriented sum of $\widetilde{S}$ and $\widetilde{S'}$ and then quotienting with ι.

It boils down to verifying that if we glue a leaf of $\widetilde{S}$ and $\widetilde{S'}$, then their images under the orientation reversing deck transformation ι also get glued together. Consider Figure 3, which shows the outward point normal vectors to $\widetilde{S}$ and $\widetilde{S'}$, which dictate which leaves are glued together.

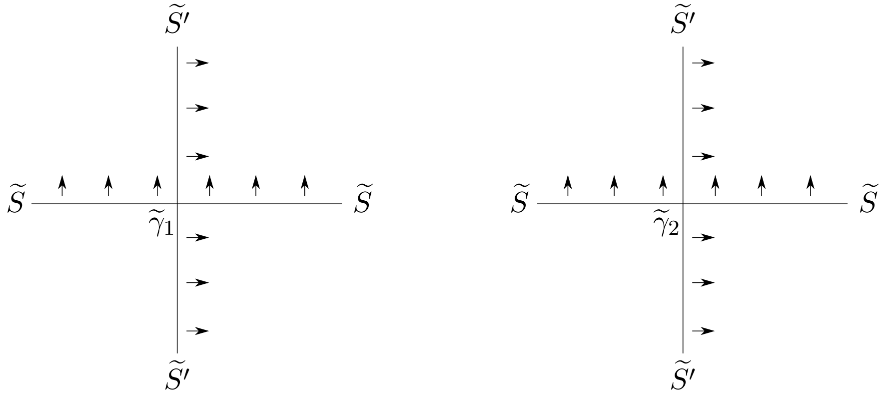

Figure 3: Neighbourhoods of $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$, with the outward pointing normal vector field.

The normal vector field tells us that the left $\widetilde{S}$ leaf gets glued to the bottom $\widetilde{S'}$ leaf near $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$. Consider now the deck transformation ι. Note that ι is an orientation reversing self map for $\widetilde{M}$, $\widetilde{S}$ and $\widetilde{S'}$. We've already seen that ι preserves the relative orientation, and thus leaves both the outwards normal vector fields invariant. This means the gluing is ι-invariant, i.e. two leaves get glued iff their images under ι get glued. This shows that the gluing descends to M, and we can thus define an oriented sum operation on embedded surfaces in M that is consistent with the oriented sum on the orientation double cover.

By the consistency of the oriented sum in M and $\widetilde{M}$, it easily follows that the oriented sum is additive in Euler characteristic, as well as in terms of Poincaré dual, since the Poincaré dual was also defined by going to the orientation double cover.

3.5 Relating 1-forms and fibrations over S^1

While we have informally described what a fibration over S^1 is prior to this section, it will be useful to formally define a fibration at this stage.

Definition 3.14 (Fibration over S^1)**.** Given a 3-manifold M, a fibration (or a fiber bundle) over S^1 is a map $f : M \to S^1$ such that the derivative of f has full rank at all points in M. The pre-image of every point in S^1 is an relatively oriented embedded surface in M, where the positive normal direction is the pre-image of the positive direction in S^1. This surface is called the fiber of the fibration.

Note that any 3-manifold that admits a fibration over S^1 is a mapping torus of the fiber (which is a surface), along with the homeomorphism that comes from the transition map when changing coordinate charts on S^1. We can instead look at homotopy classes of fibrations, and every equivalence class will correspond to a homotopy class of a homeomorphism of the fiber, i.e. a mapping class. Since we're mostly interested in mapping tori of mapping classes rather than mapping classes of specific homeomorphisms in those mapping classes, we'll be focusing on homotopy classes of fibrations.

It turns out that the "right" way to think about fibrations $f : M \to S^1$ is by studying a specific kind of 1-form on M, namely *non-singular integer* 1*-forms*.

Definition 3.15. A non-singular integer 1-form on a 3-manifold M is a smooth nowhere vanishing 1-form α on M such that for any closed loop γ, the integral of α along γ lies in $\mathbb{Z}$.

$$\int_\gamma \alpha \in \mathbb{Z}$$

Given a non-singular integer 1-form, there is a canonical way of getting a fibration f_α over S^1, as we've seen before. The map f_α is given by the following formula.

$$f_\alpha(x) := \int_{x_0}^{x} \alpha \Big/ \mathbb{Z}$$

Conversely, given a fibration $f : M \to S^1$, one can obtain a non-singular integer 1-form by pulling back a 1-form along the map f. The correct 1-form to pull back is $d\theta$, i.e. the non-vanishing 1-form on S^1 such that $\int_{S^1} d\theta = 1$ (note that despite the notation, this is not an exact form). These two constructions are inverses of each other, which is fairly easy to verify. Furthermore, if we change α to $\alpha + df$, where df is an exact form, then the associated map to S^1 is not the same, but homotopic to the original map. Conversely, if we pull back $d\theta$ along a map homotopic to f rather than f, we get a form that differs from the original form by an exact form (see Section 5.2.1 of [Cal07] for the details). The takeaway here is that if we only care about the mapping torus structure of the mapping classes, we can focus our attention to the elements of $H^1(M;\mathbb{Z})$ that admit a non-singular 1-form representative.

We now have all we need to prove a version of Theorem 2.3 for non-orientable 3-manifolds.

Theorem 3.16. *Let M be a non-orientable* 3*-manifold, and let $\mathcal{F}$ be the set of all possible ways M fibers over S^1 (up to homotopy). Then the following results hold for $\mathcal{F}$.*

(i) Elements of $\mathcal{F}$ are in a one-to-one correspondence with (non-zero) lattice points inside some union of cones over open faces of the unit ball with respect to the Thurston norm in $H^1(M;\mathbb{R})$.

(ii) If a embedded relatively oriented surface S is transverse to the suspension flow associated to some fibration f such that the flow direction is the outwards normal direction, then the Poincaré dual to S lies in the closure of the cone corresponding to f.

Proof. We've already done most of the work in reducing the proof of this result to the orientable version. To get the union of cones $\mathcal{K}$ in $H^1(M;\mathbb{R})$ corresponding to fibrations, we look at the corresponding union of cones $\widetilde{\mathcal{K}}$ in $H^1(\widetilde{M};\mathbb{R})$ (i.e. the first cohomology of the orientable double cover). Recall that $H^1(M;\mathbb{R})$ bijectively maps into $H^1(\widetilde{M};\mathbb{R})$ as a subspace. We define $\mathcal{K}$ to be the restriction of $\widetilde{\mathcal{K}}$ to the subspace corresponding to $H^1(M;\mathbb{R})$.

Consider now any fibration $f : M \to S^1$. By composing it with the covering map $p : \widetilde{M} \to M$, we get a fibration of $\widetilde{M}$: $f \circ p : \widetilde{M} \to S^1$. By pulling back $d\theta$ along this map, we get an element of $H^1(\widetilde{M};\mathbb{Z})$, which lies in $\widetilde{\mathcal{K}}$. This element is a pullback of $f^*(d\theta) \in H^1(M;\mathbb{Z})$, and therefore lies in the subspace corresponding to $H^1(M;\mathbb{Z})$. This shows that every homotopy class of fibrations injectively maps into $\mathcal{K}$. Conversely, suppose we have some element α in $\mathcal{K}$. Its pullback $\widetilde{\alpha}$ lies in $\widetilde{\mathcal{K}}$, and therefore corresponds to a fibration $f_{\widetilde{\alpha}} : \widetilde{M} \to S^1$. We would like to pushforward this map to a map from M to S^1. Recall the proof of Lemma 3.13 in which we saw that for this to happen, $\widetilde{\alpha}$ must satisfy the following condition for any basepoint x_0 in $\widetilde{M}$.

$$\int_{x_0}^{\iota(x_0)} \widetilde{\alpha} \in \mathbb{Z}$$

But observe that since any path from x_0 to $\iota(x_0)$ is a lift of a closed curve γ on M from $p(x_0)$ to $p(\iota(x_0)) = p(x_0)$, and $\widetilde{\alpha}$ is the pullback of the $\alpha \in H^1(M;\mathbb{Z})$, the above integral is equal to an integral on M.

$$\int_{x_0}^{\iota(x_0)} \widetilde{\alpha} = \int_{\gamma} \alpha$$

The right hand side term is clearly an integer, since α is an integer 1-form. This shows that the map $f_{\widetilde{\alpha}}$ descends to a map on M, and therefore α corresponds to a fibration. This proves part (i) of the theorem.

Let α be the Poincaré dual to S. Because of the way we defined α, the pullback $\widetilde{\alpha}$ is the Poincaré dual to the pre-image $\widetilde{S}$ of S. Furthermore, since S is transverse to the suspension flow direction, with the flow direction pointing outwards, the same must hold for $\widetilde{S}$. Therefore the Poincaré dual to $\widetilde{S}$ lies in the closure of the corresponding cone. But since the dual also is a pullback of α, it lies in the subspace corresponding to $H^1(M;\mathbb{R})$. This shows that α lies in the closure of the cone restricted to $H^1(M;\mathbb{R})$ and proves part (ii) of the result. □

Part (ii) of the above theorem (and Theorem 2.3) is especially useful when one is trying to decompose a 3-manifold into a mapping torus. Suppose one starts off with a mapping torus $M = (\mathcal{N}, \varphi)$. One then constructs another relatively oriented surface $\mathcal{N}'$ inside M such that $\mathcal{N}'$ is transverse to the suspension flow direction. By the above theorem, the Poincaré dual to $\mathcal{N}$, which we'll call α, lies in a cone with other 1-forms also coming from fibrations. Furthermore, the Poincaré dual to $\mathcal{N}'$, which we'll call α', lies in the closure of said cone. By taking positive integer linear combinations of α and α', we can get other elements of the cone. At the level of surfaces, that corresponds to taking oriented sums of $\mathcal{N}$ and $\mathcal{N}'$ to get new relatively oriented surfaces in M. Under reasonably mild conditions on $\mathcal{N}$ and $\mathcal{N}'$, we can actually realize their oriented sums as fibers of fibrations.

Theorem 3.17. *Let M, $\mathcal{N}$ and $\mathcal{N}'$ be as described above. Suppose that $\mathcal{N}'$ is incompressible, and the oriented sum of $\mathcal{N}$ and $\mathcal{N}'$, which we'll denote by $\mathcal{N} + \mathcal{N}'$, is connected. Then $\mathcal{N} + \mathcal{N}'$ is isotopic to the fiber of the fibration given by $\alpha + \alpha'$ (the Poincaré duals to $\mathcal{N}$ and $\mathcal{N}'$).*

Proof. The first step is to observe that in this case, one can compute the Thurston norm of α and α' using $\mathcal{N}$ and $\mathcal{N}'$. For α and α', the Thurston norms are $2\chi_-(\mathcal{N})$ and $2\chi_-(\mathcal{N}')$ respectively. This follows by passing to the orientation double cover, and noting that incompressible surfaces minimize the Thurston norm in their homology class. The first surface $\mathcal{N}$ is incompressible by virtue of being the fiber of a fibration, and the second surface $\mathcal{N}'$ is incompressible by hypothesis.

Since α and α' lie in a cone over a fibered face, the Thurston norm x is linear. We can thus compute the Thurston norm of $\alpha + \alpha'$ in terms of $\mathcal{N} + \mathcal{N}'$.

$$\begin{aligned} x(\alpha + \alpha') &= x(\alpha) + x(\alpha') \\ &= 2\chi_-(\mathcal{N}) + 2\chi_-(\mathcal{N}') \\ &= 2\chi_-(\mathcal{N} + \mathcal{N}') \end{aligned}$$

The last equality follows from the properties of the oriented sum. This tells us that the pre-image of $\mathcal{N} + \mathcal{N}'$ in the double cover must be Thurston norm minimizing, and thus incompressible.

By Theorem 3.16, we have that $\alpha + \alpha'$ corresponds to some other fibration of M. Since M is non-orientable, there are two kinds of fiber possible.

(i) The fibration is the mapping torus of a non-orientable surface along with a homeomorphism.

(ii) The fibration is the mapping torus of an orientable surface along with an orientation reversing homeomorphism.

In the first case, the fiber is a non-orientable surface homologous to $\mathcal{N}+\mathcal{N}'$. By passing to the double cover, we get two homologous orientable surfaces, both of which minimize the Thurston norm. By Theorem 2.10, we have that the double cover of $\mathcal{N}+\mathcal{N}'$ is isotopic to the fiber, which means $\mathcal{N}+\mathcal{N}'$ is isotopic to the fiber in M, and M can be realized as the mapping torus of some homeomorphism on $\mathcal{N}+\mathcal{N}'$.

The second case can be ruled out using a similar argument. In the case that M is the mapping torus of an orientable surface S with an orientation reversing homeomorphism, we can pass to the double cover. The fiber of a point in the double cover is two disjoint copies of S. But that is homologous to the double cover of $\mathcal{N}+\mathcal{N}'$ which will have a single component, since $\mathcal{N}+\mathcal{N}'$ is non-orientable. Theorem 3.16 says these two surfaces must be isotopic, but that can't happen since they have a different number of connected components. □

The non-orientable versions of Theorems 2.5 and 2.6 follow in a straightforward manner from the orientable versions.

Theorem 3.18. *Let M be a non-orientable hyperbolic 3-manifold and let $\mathcal{K}$ be the union of cones in $H^1(M;\mathbb{R})$ whose lattice points correspond to fibrations over S^1. There exists a strictly convex function $h:\mathcal{K}\to\mathbb{R}$ satisfying the following properties.*

(i) For all $t>0$ and $u\in\mathcal{K}$, $h(tu)=\frac{1}{t}h(u)$.

(ii) For every primitive lattice point $u\in\mathcal{K}$, $h(u)=\log(k)$, where k is the stretch factor of the pseudo-Anosov map associated to this lattice point.

(iii) $h(u)$ goes to ∞ as u approaches $\partial\mathcal{K}$.

Proof. We already have such a function $\widetilde{h}$ on $H^1(\widetilde{M};\mathbb{R})$. Restricting that function to the subspace corresponding to $H^1(M;\mathbb{R})$, we get a convex function satisfying properties (i) and (iii). To verify property (ii), we need verify that the stretch factor of a pseudo-Anosov map on a non-orientable surface is the same as the stretch factor of the unique lift to its double cover. This follows from Proposition 2.2. □

The exact statement of Theorem 2.6 holds for the non-orientable setting too: one just restricts the function h on $H^1(\widetilde{M};\mathbb{R})$ to the subspace corresponding to $H^1(M;\mathbb{R})$.

4 Minimal stretch factors for non-orientable surfaces with marked points

Recall that associated to every pseudo-Anosov homeomorphism f there is a number $\lambda(f)$, the *dilatation* or *stretch factor*, the amount that the stable and unstable foliations of the pseudo-Anosov change by. Given a surface S, it is natural to ask what we can say about the set of all possible stretch factors, i.e.

$$\{\log(\lambda(f)) \mid f\in\mathrm{Mod}(S) \text{ is pseudo-Anosov}\}$$

We call this set the *spectrum* of S. A first step at understanding this set of all stretch factors associated to a surface is considering the following quantity

$$l_{g,n}=\min\{\log(\lambda(f)) \mid f\in\mathrm{Mod}(\mathcal{S}_{g,n}) \text{ is pseudo-Anosov}\}.$$

The study of this minimal stretch factor $l_{g,n}$ was initiated by Penner in his work [Pen91]. In this paper Penner studied the asymptotic behavior of minimal stretch factors of orientable surfaces without punctures, i.e. the behavior of $l_{g,0}$. He showed that there exist positive constants A_1 and A_2 such that the following inequalities held for any $g \geq 2$.

$$\frac{A_1}{g} \leq l_{g,0} \leq \frac{A_2}{g}$$

This showed that asymptotically $l_{g,0}$ behaves like $\frac{1}{g}$ for $g \geq 2$. It turns out something similar is true when one starts adding punctures. Yazdi showed that for any fixed n, $l_{g,n}$ behaves like $\frac{1}{g}$. More precisely, he proved the following theorem.

Theorem 4.1 (Theorem 1.2 of [Yaz18])**.** *For any fixed $n \in \mathbb{N}$, there are positive constants $B_1 = B_1(n)$ and $B_2 = B_2(n)$ such that the following inequalities hold for any $g \geq 2$.*

$$\frac{B_1}{g} \leq l_{g,n} \leq \frac{B_2}{g}.$$

Yazdi's result is one of many recent results in studying the asymptotics of $l_{g,n}$ for different subsets of the (g, n) plane. See the introductions of [Yaz18] and [Tsa09] for more examples of results of this form. Yazdi proves an additional result along these lines for a large subset of the (g, n) plane, one containing balls of arbitrary large radii.

Theorem 4.2 (Yazdi)**.** *There exists positive constants A, B and C such that for any $n \geq 1$ and $g \geq Cn\log^2(n)$ such that the following inequalities hold.*

$$\frac{B}{g} \leq l_{g,n} \leq \frac{A}{g}$$

A key tool in Yazdi's proof of both these theorems was the fibered face theory of Thurston. With the non-orientable analog of Thurston's fibered face theory we developed in the previous section, it's possible to prove an analogous theorem for non-orientable punctured surfaces. Let $\mathcal{N}_{g,n}$ be the genus g non-orientable surface with n punctures and let $l'_{g,n}$ be the minimal stretch factor of $\mathcal{N}_{g,n}$.

$$l'_{g,n} = \min \{\log(\lambda(f)) \mid f \in \mathrm{Mod}(\mathcal{N}_{g,n}) \text{ is pseudo-Anosov}\}$$

Then we have the following result, analogous to Yazdi's results.

Theorem 4.3. *For any fixed $n \in \mathbb{N}$, there are positive constants $B'_1 = B'_1(n)$ and $B'_2 = B'_2(n)$ such that for any $g \geq 2$, the stretch factor satisfies the following inequalities.*

$$\frac{B'_1}{g} \leq l'_{g,n} \leq \frac{B'_2}{g}$$

Observe that the lower bound for the non-orientable case follows easily from the lower bound for the orientable case. Let f be the pseudo-Anosov map with the minimal stretch factor on $\mathcal{N}_{g,n}$. Then, by Proposition 2.2, this map lifts to a map $\widetilde{f}$ on $\mathcal{S}_{g-1,2n}$ (possibly after squaring). Furthermore, $\widetilde{f}$ has the same stretch factor as f. The former is bounded below by $\frac{B}{g}$, and thus the stretch factor of f is bounded below as well. The more challenging part of the proof is showing the upper bound holds. This will be done by explicitly constructing pseudo-Anosov maps with small stretch factors, adapting Yazdi's techniques to the non-orientable setting.

The construction of Yazdi proceeds in five steps: in steps 1 and 2, a family of small dilatation psuedo-Anosov maps is constructed on $\mathcal{S}_{g_i,n}$, where $\{g_i\}$ is a sequence of genera going off to infinity, by not containing every element of $\mathbb{N}$, i.e. there are plenty of gaps. Steps 3 and 5 deal with filling in the gaps, i.e. constructing small dilatation pseudo-Anosov maps on the missing surfaces, and this is where Thurston's fibered face theory enters the picture. In this section, we will adapt the steps to work for non-orientable surfaces.

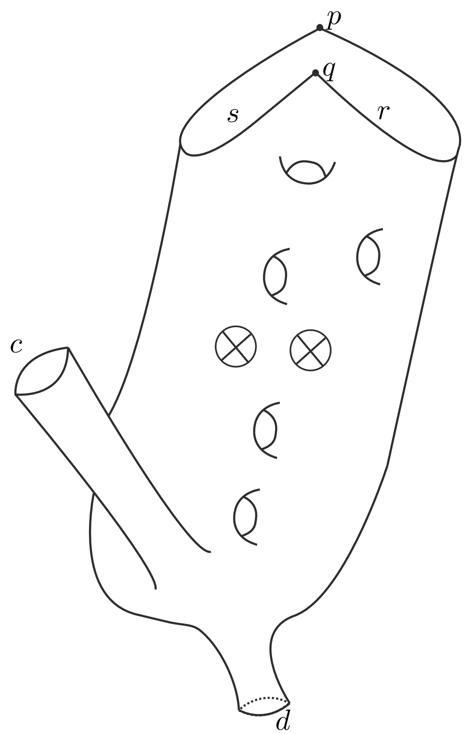

Figure 4: The surface T, which will be the building block of the construction.

Step 1: Constructing the surfaces The first step in the construction is defining a family of surfaces that exhibit a sort of rotational symmetry. Using this symmetry, if one shows that a power of some homeomorphism is pseudo-Ansov, then so is the original homeomorphism. Yazdi cites this insight as being due to Penner in his construction in [Pen91].

Note that we will try to follow Yazdi's notation as close as we can, in order to make it clear to the reader how our construction replicates his.

We begin by defining a family of surfaces $P_{n,k}$. Let T be an orientable surface of genus 5 with 3 boundary components c, d and e. Orient T and give c, d and e the induced orientations from this orientation. Now add two cross-caps to T but keep the boundaries of T oriented. Let p (respectively q) be a puncture (respectively a marked point) on the boundary component e of T, with oriented arcs r and s connecting them on ∂T. See Figure 4 for a picture of T.

Let $T_{i,j}$ be copies of the surface T, where $i, j \in \mathbb{Z}$. We will use similar notation to refer to the boundary components of $T_{i,j}$. Define an infinite surface S_∞ as the following quotient:

$$S_\infty \coloneqq \left(\bigcup T_{i,j}\right) \Big/ \sim$$

Here, i and j are integers. The gluing $\sim$ is given by the following two families of identifications.

$$c_{i,j} \sim d_{i+1,j}$$
$$r_{i,j} \sim s_{i,j+1}$$

Furthermore, the boundary components are glued by an orientation-reversing homeomorphisms. We have two natural shift maps $\overline{\rho_1}, \overline{\rho_2} : T_\infty \to T_\infty$ that act in the following manner.

$$\overline{\rho_1} : T_{i,j} \mapsto T_{i+1,j}$$
$$\overline{\rho_2} : T_{i,j} \mapsto T_{i,j+1}$$

Note that these maps commute. Define the surface $P_{n,k}$ as the quotient of the surface T_∞ by the covering action of the group generated by $(\overline{\rho_1})^n$ and $(\overline{\rho_2})^k$. Therefore, $\overline{\rho_1}$ and $\overline{\rho_2}$ induce maps on the surface $P_{n,k}$, which we denote by ρ_1 and ρ_2.

A natural question at this point is why did we choose the surface T for our building block? It comes down to two main problems:

- The combinatorics of the curves make the associated matrix we get from the Penner construction satisfy the conditions of Lemma 4.6. This is used to prove our family of pseudo-Anosov maps have stretch factors bounded above by the quantity we desire.
- Having a curve γ such that it and its image under our map form the boundary of an embedded $\mathbb{R}P^2$ with two boundary components in the mapping torus, which will come into play when extending our family of surfaces in Step 3.

Lemma 4.4. *Define the sequence $g_{n,k}$ in the following manner for $n \geq 1$ and $k \geq 3$.*

$$g_{n,k} = (14k - 2)n + 2$$

The genus of $P_{n,k}$ is $g_{n,k}$.

Proof. Consider the subsurface $U \subset P_{n,k}$ defined as

$$U = \left(\bigcup_{i=0}^{k-1} T_{0,i} \right) \Big/ \sim$$

Then U is a compact, non-orientable surface of genus $12k$ with $2k$ boundary components, and forms a fundamental domain for the covering action of $\overline{\rho_1}$ on T_∞. We have a formula for the Euler characteristic of U.

$$\begin{aligned} \chi(U) &= 2 - 12k - 2k \\ &= 2 - 14k \end{aligned}$$

This gives us a formula for the Euler characteristic of $P_{n,k}$.

$$\begin{aligned} \chi(P_{n,k}) &= n \cdot \chi(U) \\ &= -n(14k - 2) \end{aligned}$$

since $P_{n,k}$ is formed by gluing n copies of U together along circle boundary components. By the relation between genus and Euler characteristic, we have the claimed formula for genus.

$$g_{n,k} = n(14k - 2) + 2$$

□

Step 2: Constructing the maps We now construct maps $f_{n,k} : P_{n,k} \to P_{n,k}$ that are defined as a composition of specific Dehn twists followed by a finite order mapping class. The key insight is that a power of this map will be a composition of Dehn twists that satisfy the criteria to be a Penner construction and thus pseudo-Anosov. This is how we take advantage of the rotational symmetry of the $P_{n,k}$.

Recall that for non-orientable surfaces, we don't initially have a well-defined notion of a positive or negative Dehn twist. As we saw in Section 2.1, in order to perform the Penner construction, we need to ensure that the curves we are working with are marked inconsistently. Note that our labeling of the curves already gives us an inconsistent marking. For any alpha curve α_i, we let the marking ϕ_{α_i} be orientation

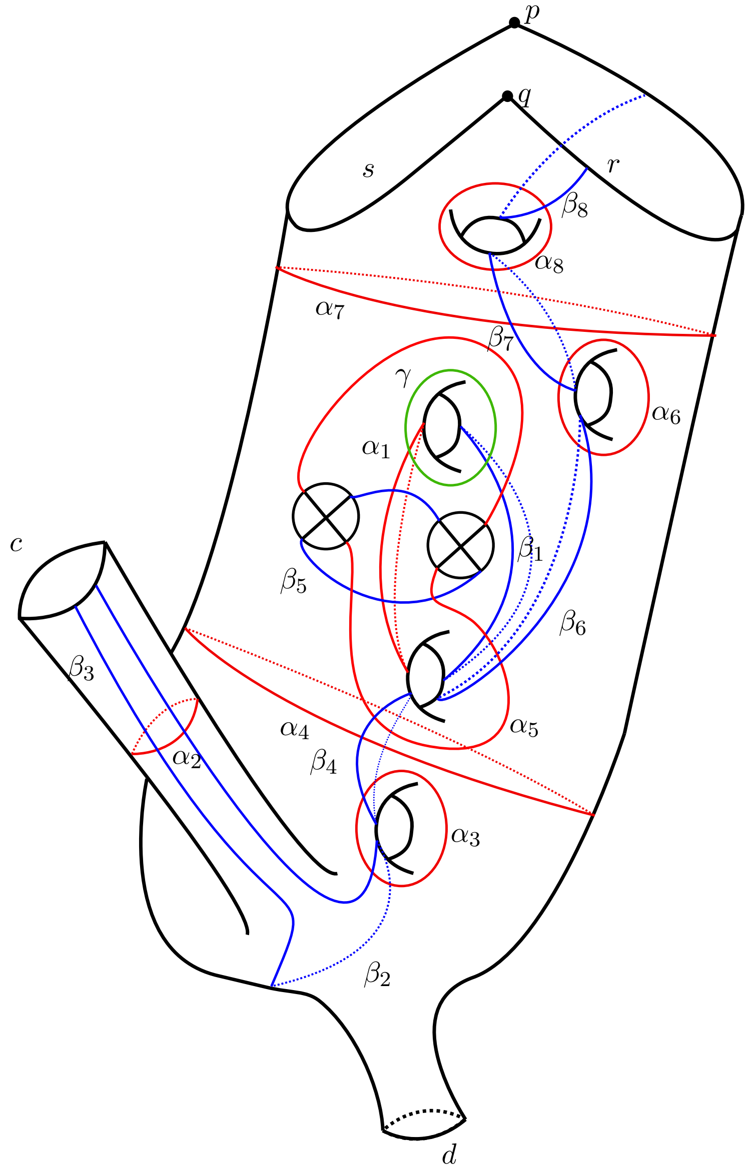

Figure 5: The curves that lie on $T_{0,0}$

preserving and for beta curves β_j let ϕ_{β_j} be orientation reversing. Since alpha curves only intersect beta curves (and vice versa), we have an inconsistent marking at each point of intersection.

Let $\mathcal{B}$ be the union of all β curves except β_1 in $T_{0,0} \cup T_{0,1} \cup T_{1,0}$ (see figures below). Let $\rho_1(\mathcal{B})$ be the image of $\mathcal{B}$ under ρ_1. Define ϕ_b as the composition of Dehn twists along all the curves in the set $\overline{\mathcal{B}} \coloneqq \mathcal{B} \cup \rho_1(\mathcal{B}) \cup \cdots \cup \rho_1^{n-1}(\mathcal{B})$. Since the curves in $\overline{\mathcal{B}}$ are disjoint, Dehn twists along them commute and therefore it is not necessary to specify the order in which we compose these Dehn twists in ϕ_b. Let $\mathcal{R}$ be the union of all α curves except α_1 in $T_{0,0}$. Define $\overline{\mathcal{R}}$ and ϕ_r in the exact same way.

Let $\alpha_1, \beta_1 \subset T_{0,0}$ be the curves in Figure 5. Let ϕ be the composition of Dehn twists along all the curves $\alpha_1, \rho_1(\alpha_1), \ldots, \rho_1^{n-1}(\alpha_1)$ followed by Dehn twists along all the curves $\beta_1, \rho_1(\beta_1), \ldots, \rho_1^{n-1}(\beta_1)$. Define the map $f_{n,k}$ in the following manner.

$$f_{n,k} \coloneqq \rho_2 \circ \phi \circ \phi_b \circ \phi_r$$

It follows from the Penner construction that $(f_{n,k})^k$ is pseudo-Anosov. Hence $f_{n,k}$ itself is pseudo-Anosov and an invariant train track $\tau_{n,k}$ for $f_{n,k}$ can be obtained from Penner's construction that we described

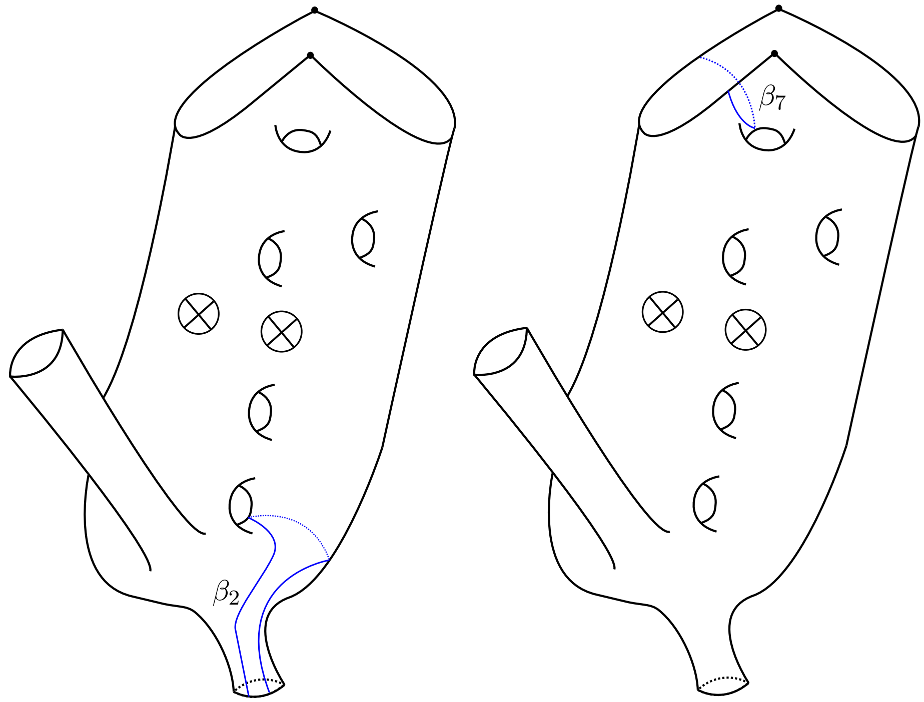

Figure 6: The parts of curves β_2 and β_7 on $T_{0,1}$ and $T_{1,0}$

in Section 2.1.

Step 3: The Mapping Torus We have now constructed an infinite family of non-orientable surfaces and pseudo-Anosov maps, but this isn't enough. Looking back at Lemma 4.4, the genera of this family of surfaces do not include every positive integer. In fact, it misses infinitely many integers. We will use our extension of the Thurston's fibered face theory to fill in the gaps, by constructing fibers for fibrations of the mapping tori of the pseudo-Anosov maps we defined above.

Let $M_{n,k}$ be the mapping torus of $f_{n,k}$. Likewise, let $\mathcal{K}_{n,k}$ denote the fibered cone of $H^1(M_{n,k}, \mathbb{R})$ corresponding to the map $f_{n,k}$. We will show that $M_{n,k}$ contains a closed, relatively orientable, incompressible surface homeomorphic to $\mathcal{N}_3$ that is transverse to the suspension flow direction. This will allow us to apply Theorem 3.17 to construct new fibrations of $M_{n,k}$.

Lemma 4.5. *There is a relatively orientable incompressible surface $F_{n,k}$ in $M_{n,k}$ that is homeomorphic to $\mathcal{N}_3$. Moreover it is transverse to the suspension flow direction given by $f_{n,k}$ and its Poincaré dual is in the closure $\overline{\mathcal{K}_{n,k}}$.*

Proof. Let $\gamma \subset T_{0,0}$ be the curve as shown in Figure 5. Note that as we mentioned in Step 1, γ was specifically chosen so that γ and $\phi(\gamma)$ bound a non-orientable surface $\hat{F}$ of genus 1 with boundary, i.e. cutting along them creates an an annulus with a cross-cap. For convenience, we will denote $\phi(\gamma)$ by $\hat{\gamma}$. We are going to follow the image of γ under iterations of our pseudo-Anosov map $f_{n,k}$, this will allow us to attach tubes (annuli) to the boundary of $\hat{F}$ to get a closed $\mathcal{N}_3$. Following γ under $f_{n,k}$ gives us the following.

$$\begin{aligned} f_{n,k}(\gamma) &= \rho_2 \circ \phi \circ \phi_b \circ \phi_r(\gamma) \\ &= \rho_2 \circ \phi(\gamma) \\ &= \rho_2(\hat{\gamma}) \end{aligned}$$

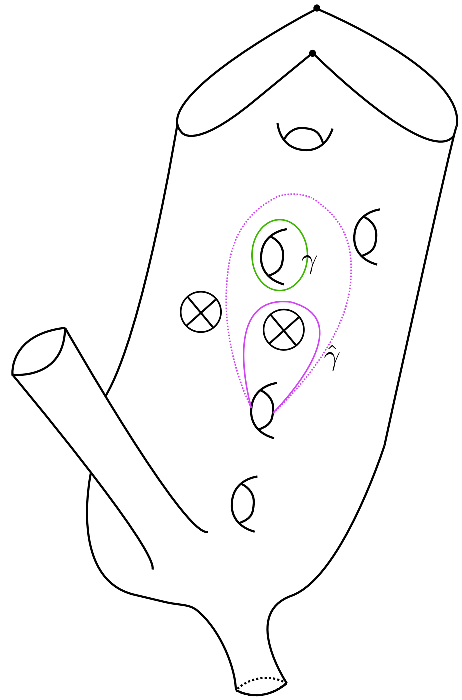

Figure 7: The curves γ and $\hat{\gamma}$ bound an a non-orientable surface of genus 1.

Applying $f_{n,k}$ repeatedly to γ gives us the following.

$$\begin{aligned} f_{n,k}^2(\gamma) &= \rho_2^2(\hat{\gamma}) \\ &\vdots \\ f_{n,k}^k(\gamma) &= \rho_2^k(\hat{\gamma}) \\ &= \hat{\gamma} \end{aligned}$$

Let T_i be a tube (annulus) that connects $f_{n,k}^{i-1}(\gamma)$ to $f_{n,k}^{i}(\gamma)$ (which as we saw above will just be γ or copies of $\hat{\gamma}$ on a ρ_2 rotation of $T_{0,0}$) in the mapping torus $M_{n,k}$, for $1 \leq i \leq k$. We obtain these tubes by following the suspension flow of $f_{n,k}$ around $M_{n,k}$. We can now obtain our embedded surface $F_{n,k}$ by taking the union of $T_1, T_2, \ldots, T_k$ and $\hat{F}$. Since we are "adding an orientable genus" to a non-orientable surface of genus 1, we get that $F_{n,k}$ is homeomorphic to $\mathcal{N}_3$.

Since the resulting surface is an embedded non-orientable surface in a non-orientable 3-manifold, we have relative orientability by Proposition 3.10.

The proof of the fact that $F_{n,k}$ can be isotoped to be transverse to the suspension flow is the same as the proof in [Yaz18], which in turn follows the proof in [LM13]. The proof goes through even in this setting essentially because of the local nature of the proof.

Finally, $F_{n,k}$ is incompressible in $M_{n,k}$ because $M_{n,k}$ is hyperbolic, and $F_{n,k}$ is genus 3, the lowest possible genus for a hyperbolic non-orientable surface. ☐

Step 4: Bounding the Stretch Factor In [Yaz18], Yazdi shows that the family of pseudo-Anosov maps that we have constructed all have the log of their stretch factor bounded above by a similar factor. In order to to do this, recall in Section 2 we saw that pseudo-Anosov maps give rise to matrices whose Perron-Frobenius eigenvalue is our stretch factor. So a way to find an upper bound of the stretch factor of the maps we have constructed is to bound the spectral radius of the associated matrices. The following lemma by Yazdi does just this for a specific class of matrices that our examples are based off of.

Lemma 4.6 (Lemma 2.3 of [Yaz18])**.** *Let A be a non-negative integral matrix, Γ be the adjacency graph of A, and $V(\Gamma)$ the set of vertices of Γ. For each $v \in V(\Gamma)$, define v^+ to be the set of vertices u such that there is an oriented edge from v to u. Let D and k be fixed natural numbers. Assume the following conditions hold for Γ.*

(i) For each $v \in V(\Gamma)$ we have $\deg_{out}(v) \leq D$.

(ii) There is a partition $V(\Gamma) = V_1 \cup \cdots \cup V_k$ such that for each $v \in V_i$ we have $v^+ \subset V_{i+1}$, for any $1 \leq i \leq k$ except possibly when $i = 1$ or 3 (indices are mod k).

(iii) For each $v \in V_1$, we have $v^+ \subset V_2 \cup V_3$.

(iv) For each $v \in V_3$ we have $v^+ \subset V_3 \cup V_4$, and for $u \in v^+ \cap V_3$ we have $u^+ \subset V_4$.

(v) For all $3 < j \leq k$ and each $v \in V_j$, the set v^+ consists of a single element.

With this result in hand, we can now show that the stretch factors for our main family of examples are all bounded above in the way we hope.

Lemma 4.7. *Let $\lambda_{n,k}$ be the stretch factor of $f_{n,k}$. Then there exists a universal positive constant C' and such that for every $n \geq 1$ and $k \geq 3$, we have the following upper bound on $\log(\lambda_{n,k})$.*

$$\log(\lambda_{n,k}) \leq C' \frac{n}{g_{n,k}}$$

Proof. We deliberately constructed our examples so our curves are in the same "general form" as the ones in [Yaz18] and thus they will still satisfy the criteria of Lemma 4.6. All intersections between the curves happen inside the building block T, except for the intersections between building blocks given by the beta curves β_3 and β_8. Though we still will explicitly show that it is the case our set of curves will satisfy the above lemma by Yazdi as well.

We define the following multi-curves.

$$\begin{aligned}
\mathcal{A} &\coloneqq \mathcal{B} \cup \mathcal{R} \cup \{\alpha_1, \beta_1\} \\
\overline{\mathcal{A}} &\coloneqq \mathcal{A} \cup \rho_1(\mathcal{A}) \cup \cdots \cup \rho_1^{n-1}(\mathcal{A}) \\
\widehat{\mathcal{A}} &\coloneqq \overline{\mathcal{A}} \cup \rho_2\left(\overline{\mathcal{A}}\right) \cup \cdots \cup \rho_2^{k-1}\left(\overline{\mathcal{A}}\right).
\end{aligned}$$

Thus $\hat{\mathcal{A}}$ is all the curves on our surface we are Dehn twisting around to get $f_{n,k}$.

Recall from above that we stated we need to find the eigenvalue of the matrix that represents the action of $f_{n,k}$ on the subspace of the cone of transverse measures that is spanned by the measures assigning 1 to single curves in $\hat{\mathcal{A}}$ and 0 to everything else. Let A be said matrix and Γ the adjacency graph of A. In order to bound the spectral radius of A, we need to show that Γ satisfies the criteria of Lemma 1. To do this we first need to partition the vertices of γ, which is equivalent to a partition of the curves in $\hat{\mathcal{A}}$:

$$\mathcal{A} = \bigcup_{i=1}^{k} \rho_2^{i-2}(\overline{\mathcal{A}}).$$

Then define V_i for $1 \leq i \leq k$ as the vertices of Γ corresponding to elements in $\rho_2^{i-2}(\overline{\mathcal{A}})$.

We can now check the conditions of Lemma 1, based on the combinatorics of the curves on our surface:

(i) From the way that the curves are constructed, there will exist a constant D', independent of n and k, such that for every connected curve (single element subset) $c \subset \hat{A}$, the geometric intersection number between c and $\overline{A}$ is at most D'. Recall from Section 2.1 that we refer to the linear action of

$f_{n,k}$ on the subspace of the cone of transverse measures on our invariant train track corresponding to connected curves in $\hat{\mathcal{A}}$ as A. Following Yazdi we express A as the following product.

$$A = M_4 M_3 M_2 M_1$$

Here each M_i is the linear action of ρ_2, ϕ, ϕ_b and ϕ_r respectively. For a connected curve $x \in \hat{\mathcal{A}}$, the L^1-norm of $A(\mu_x)$ is bounded above by the geometric intersection of $f_{n,k}(x)$ with the curves in $\overline{A}$, thus each of M_1, M_2 and M_3 will change the norm by a factor of at most $(1 + D')$. Since ρ_2 won't change intersection numbers, M_4 will preserve the L^1-norm. If we let $D = (1 + D')^3$, then the outward degree of each vertex in Γ is at most D.

(ii) As above, we now have a partition of our vertices where $V_i := \rho_2^{i-2}(\overline{\mathcal{A}})$. So suppose that $v \in V_i$, $i \neq 1, 3$, is a vertex that corresponds to μ_c for a curve $c \in \hat{\mathcal{A}}$. By the partitioning c must be a curve in $\rho_2^{i-2}(\overline{\mathcal{A}})$, for $i \neq 1, 3$. Recall that $f_{n,k}$ is defined as $f_{n,k} = \rho_2 \circ \phi \circ \phi_b \circ \phi_r$. The action of $\phi \circ \phi_b \circ \phi_r$ will send μ_c to a sum of μ_y where y corresponds to elements of V_i, since $\phi \circ \phi_b \circ \phi_r$ will send curves in these partitions to curves that just intersect curves in the same partition. Then ρ_2 will rotate all the curves to the next partition, thus sending μ_y to μ_z, where z corresponds to an element of V_{i+1}.

(iii) We need to see which vertices in $v \in V_1$ have $v^+ \not\subset V_2$. As in Yazdi, this is precisely the vertices corresponding to the following set.

$$X = \{\mu_y \mid \exists\, i \;\text{ s.t. } y \text{ is a connected curve in } \rho_1(\rho_2^{-1}(\beta_8))\}$$

One can see from the picture of the curves that elements corresponding to X will have $v^+ \subset V_2 \cup V_3$.

(iv) The elements $v \in V_3$ such that $v^+ \not\subset V_4$ are the ones that correspond to the elements of the following set.

$$Y = \{\mu_y \mid \exists\, i \text{ s.t. } y \text{ is a connected curve in } \rho_1^i(\rho_2(\alpha_8))\}$$

This is due to the intersection of the curves $\rho_1^i(\rho_2(\alpha_8))$ with the curves $\rho_1^i(\beta_8)$. Moreover, for any element $v \in V_3$ corresponding to Y and any $u \in v^+ \cap V_3$, u no longer corresponds to Y and hence $u^+ \subset V_4$.

(v) All the curves corresponding to an element of V_j, $3 < j \leq k$ are disjoint from all the curves in $\overline{A}$. Thus $f_{n,k}$ just acts by rotation.

Setting $\lambda = \lambda_{n,k}$, Lemma 4.6 implies the following.

$$\lambda^{k-1} = \rho(A)^{k-1} = \rho(A^{k-1}) \leq 4D^4 \implies (k-1) \cdot \log(\lambda) \leq \log(4D^4)$$
$$\implies \frac{k}{2} \log(\lambda) \leq (k-1) \log(\lambda) \leq \log(4D^4)$$

On the other hand, we know $g_{n,k} = (14k - 2)n + 2 \leq 14kn$. Therefore

$$\log(\lambda) \leq 2 \log(4D^4) \cdot \frac{1}{k} \leq 2 \log(4D^4) \cdot \frac{14n}{g_{n,k}} = C' \frac{n}{g_{n,k}}$$

where $C' := 28 \log(4D^4)$. □

Step 5: Filling in the Gaps We now want to use the mapping tori $M_{n,k}$ of our maps $f_{n,k}$ to construct pseudo-Anosov maps with small stretch factors on the surfaces of the genera we our missing from our family $P_{n,k}$. To do this we consider the following surfaces: $P^r_{n,k} = P_{n,k} + r(F_{n,k})$, that is taking the oriented sum of $P^r_{n,k}$ and $F_{n,k}$ r times as defined in Theorem 3.17.

Lemma 4.8. *The surface $P^r_{n,k}$ have genus equal to $g^r_{n,k} = g_{n,k} + r$. In particular as r varies between 0 and $6n$, the genera of $P^r_{n,k}$ cover the range between $g_{n,k}$ and $g_{n,k+1}$. Moreover, $P^r_{n,k}$ is isotopic to a fiber of a fibration of $M_{n,k}$ with pseudo-Anosov monodromy that fixes $2n$ of the singularities of its invariant foliation.*

Proof. We know that the Euler characteristic of an oriented sum is the sum of the Euler characteristics of the summands.

$$\begin{aligned}\chi(P^r_{n,k}) &= \chi(P_{n,k}) + r \cdot \chi(F_{n,k}) \\ &= (-2g_{n,k} + 2) - 2r \\ &= -2(g_{n,k} - 1) + 2\end{aligned}$$

This proves the identity for the genus of $P^r_{n,k}$.

From Theorem 3.17, we know that $P^r_{n,k}$ is isotopic to a fiber of a fibration of $M_{n,k}$ since we showed that $F_{n,k}$ is incompressible and transverse to the suspension flow in Lemma 4.5. If we now let $f^r_{n,k}$ be the first return map of this new fibration, since one monodromy of $M_{n,k}$ is pseudo-Anosov, all monodromies are and $f^r_{n,k}$ is a pseudo-Anosov map.

As in the proof of Lemma 3.5 of [Yaz18], the singularities of the stable foliation of $f_{n,k}$ that are fixed are the $2n$ intersection points of the axis of ρ_1 with $P_{n,k}$. Furthermore we have already seen that the surface $F_{n,k}$ can be isotoped to be transverse to the suspension flow and disjoint from the orbit of the $2n$ singularities of $f_{n,k}$, hence the monodromy $f^r_{n,k}$ still fixes the corresponding $2n$ singularities on $P^r_{n,k}$. ☐

We can now prove the non-orientable version of Lemma 3.6 of [Yaz18].

Lemma 4.9. *Let $\lambda^r_{n,k}$ be the stretch factor of $f^r_{n,k}$.Then there exists a constant $C > 0$ such that for every $n \geq 1$, $k \geq 3$, and $0 \leq r \leq 6n$ we have the following upper bound on $\log(\lambda^r_{n,k})$.*

$$\log(\lambda^r_{n,k}) \leq C \frac{n}{g^r_{n,k}}$$

Proof. Let $\mathcal{K} = \mathcal{K}_{n,k}$ be our fibered faces and $h : \mathcal{K} \to \mathbb{R}$ the function described in Theorem 3.18. Note that we have the following bounds on $g^r_{n,k}$.

$$\begin{aligned}g^r_{n,k} &= g_{n,k} + r \\ &\leq g_{n,k} + 14n \\ &< 2g_{n,k}\end{aligned}$$

Now if we let ω be the Poincaré dual of $P^r_{n,k}$ and α the Poincaré dual of $P_{n,k}$, we have the following bounds by the convexity of the entropy function h.

$$\begin{aligned}h([\omega]) &< h([\alpha]) \\ &\leq C' \frac{n}{g_{n,k}} \\ &\leq 2C' \frac{n}{g^r_{n,k}}\end{aligned}$$

☐

So now we have that our surfaces $P^r_{n,k}$ are isotopic to fibers of fibrations of $M_{n,k}$ with pseudo-Anosov monodromies with bounded stretch factors.

We can now think of $f^r_{n,k}$ as a map on a non-orientable surface of genus $g^r_{n,k}$. Note from above we know that $g^r_{n,k}$ covers all natural numbers between $g_{n,k}$ and $g_{n,k+1}$, thus this set of genera for all r covers all natural numbers larger than $g_{n,3} = 40n+2$. Recall that all of these surfaces will have $2n$ singularities, so we can either puncture n or $n+1$ to account for all possible number of punctures.

We can now give a proof of Theorem 4.3.

Theorem 4.3. For any fixed $n \in \mathbb{N}$, there are positive constants $B'_1 = B'_1(n)$ and $B'_2 = B'_2(n)$ such that for any $g \geq 2$, the stretch factor satisfies the following inequalities.

$$\frac{B'_1}{g} \leq l'_{g,n} \leq \frac{B'_2}{g}$$

Proof. We begin by proving the upper bound. By Lemma 4.9 and above, we have that there exists a number $C' > 0$ such that for $g \geq 40n+2$, $l'_{g,n} \leq 2C'\frac{n}{g}$. So we take $B'_2(n)$ to be the following quantity.

$$B'_2(n) = \max\{4C'n, l'_{1,n}, 2l'_{2,n}, \ldots, (40n+1)l'_{40n+1,n}\}$$

The lower bound follows easily from the lower bound in the orientable setting, as demonstrated in the discussion following Theorem 4.3. For ease of reading, we replicate that here. Let f be the pseudo-Anosov map with the minimal stretch factor on $\mathcal{N}_{g,n}$. Then, by Proposition 2.2, this map lifts to a map $\widetilde{f}$ on $\mathcal{S}_{g-1,2n}$ (possibly after squaring). Furthermore, $\widetilde{f}$ has the same stretch factor as f. The former is bounded below by $\frac{B}{g}$, and thus the stretch factor of f is bounded below as well. $\square$

References

[ALM] I Agol, CJ Leininger, and D Margalit. Pseudo-anosov stretch factors and homology of mapping tori, preprint (2014). *arXiv preprint arXiv:1409.6983.*

[Cal07] Danny Calegari. *Foliations and the geometry of 3-manifolds.* Oxford University Press on Demand, 2007.

[Fri82] David Fried. Flow equivalence, hyperbolic systems and a new zeta function for flows. *Commentarii Mathematici Helvetici*, 57(1):237–259, 1982.

[Fri83] David Fried. Transitive anosov flows and pseudo-anosov maps. *Topology*, 22(3):299–303, 1983.

[LM13] Christopher J Leininger and Dan Margalit. On the number and location of short geodesics in moduli space. *Journal of Topology*, 6(1):30–48, 2013.

[Mat87] Shigenori Matsumoto. Topological entropy and thurston's norm of atoroidal surface bundles over the circle. *J. Fac. Sci. Univ. Tokyo Sect. IA Math*, 34(3):763–778, 1987.

[Pen88] Robert C Penner. A construction of pseudo-anosov homeomorphisms. *Transactions of the American Mathematical Society*, 310(1):179–197, 1988.

[Pen91] Robert C Penner. Bounds on least dilatations. *Proceedings of the American Mathematical Society*, 113(2):443–450, 1991.

[Str17] Balázs Strenner. Algebraic degrees of pseudo-anosov stretch factors. *Geometric and Functional Analysis*, 27(6):1497–1539, Nov 2017.

[Thu86] William P Thurston. A norm for the homology of 3-manifolds. *Memoirs of the American Mathematical Society*, 59(339):99–130, 1986.

[Tsa09] Chia-Yen Tsai. The asymptotic behavior of least pseudo-anosov dilatations. *Geometry & Topology*, 13(4):2253–2278, 2009.

[Yaz18] Mehdi Yazdi. Pseudo-anosov maps with small stretch factors on punctured surfaces. *arXiv preprint arXiv:1801.01754*, 2018.