THURSTON CONSTRUCTION MAPPING CLASSES WITH MINIMAL DILATATION

MARYAM CONTRACTOR AND OTTO REED

ABSTRACT. Given a pair of filling curves α, β on the surface of genus g with n punctures $\Sigma_{g,n}$, we explicitly compute the smallest dilatation mapping classes over all the pseudo-Anosov maps given by the Thurston construction. We do so by solving for the minimal dilatation in a congruence subgroup of $\mathrm{PSL}_2(\mathbb{Z})$. We apply this result to realized lower bounds on intersection number between α and β to give the minimal mapping class over any Thurston construction pA map given by a filling pair $\alpha \cup \beta$.

CONTENTS

1. INTRODUCTION

1.1. Preliminaries.

Given a surface of genus g with n punctures $\Sigma_{g,n}$, the *mapping class group* identifies homotopy classes of maps on $\Sigma_{g,n}$ which are not isotopic to the identity. Thurston classified these equivalence classes of maps into three categories:

Theorem 1.1 (Nielsen-Thurston Classification)**.** *Let $[f] \in \mathrm{Mod}(\Sigma_{g,n})$. Then there exists a representative homeomorphism $g \in [f]$ such that either*

(i) g is periodic: $g^d = I$ for some $0 \leq d < \infty$
(ii) g is reducible: there exists a (reducible) simple closed curve $\{\gamma_i\}$ such that $g(\{\gamma_i\}) = \{\gamma_i\}$.
(iii) g is unique and pseudo-Anosov

A pseudo-Anosov diffeomorphism f of a surface $\Sigma_{g,n}$ is characterized by the *dilatation* $\lambda > 1$ and two transverse directions, unstable and stable, along which f expands lengths by λ and contracts lengths by $1/\lambda$, respectively. Iterating f causes the lengths of curves in the unstable direction to tend to infinity and lengths of curves in the stable direction to tend to zero.

Thurston's classification gives that *every* mapping class in $\mathrm{Mod}(\Sigma_{g,n})$ which does not have finite order and is not reducible must be pA. To find such maps, one might first consider generators of the entire mapping class group. One example of such maps are *Dehn twists.*

Definition 1.2. Let S be an oriented surface and α a simple closed curve in S. The *Dehn twist* about α, $T_\alpha : S \to S$ is the homeomorphism given by cutting the surface along α, twisting a neighborhood of one of the boundary components by 2π and then regluing it. By convention we choose to twist counterclockwise, which gives a *left twist*.

We will typically take an isotopy class a of simple closed curves in S and write T_a to mean the Dehn twist about a representative $\alpha \in a$.

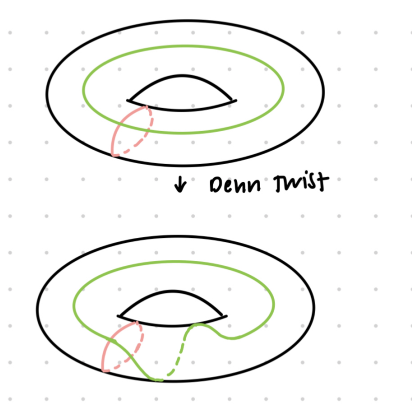

It is known that $\mathrm{Mod}(\Sigma_{g,n})$ is generated by finitely many Dehn twists (see [4], Theorem 4.1). However, Dehn twists about just two curves (specifically, two which *fill* $\Sigma_{g,n}$) generate some pA mapping classes.

Definition 1.3. If A and B are each unions of disjoint isotopy classes of simple closed curves on $\Sigma_{g,n}$, we say A and B *fill* $\Sigma_{g,n}$ if the complement $\Sigma_{g,n} \setminus (A \cup B)$ is a union of topological disks or punctured disks.

We can easily extend the definition of a Dehn twist to a multicurve $A = \{\alpha_1, \ldots, \alpha_k\}$ (i.e., a collection of pairwise disjoint simple closed curves) by defining the *multitwist* $T_A = \prod_{i=1}^{k} T_{\alpha_i}$.

For two multicurves $A = \{\alpha_1, \ldots, \alpha_k\}$, $B = \{\beta_1, \ldots, \beta_\ell\}$ which fill $\Sigma_{g,n}$, Thurston gives an explicit construction of pseudo-Anosov (pA) mapping classes using only information from these two curves ([4], Theorem 14.1).

Theorem 1.4 (Thurston's Construction)**.** *Suppose A and B are multicurves in $\Sigma_{g,n}$ so that $A \cup B$ fills $\Sigma_{g,n}$. Then there is a real number $\mu = \mu(\alpha, \beta)$ and a representation $\rho : \langle T_A, T_B \rangle \to \mathrm{PSL}_2(\mathbb{R})$ given by*

$$T_A \mapsto \begin{bmatrix} 1 & -\mu^{1/2} \\ 0 & 1 \end{bmatrix} \quad T_B \mapsto \begin{bmatrix} 1 & 0 \\ \mu^{1/2} & 1 \end{bmatrix}.$$

Moreover, ρ has the following properties:

(i) There is a bijective correspondence between periodic, reducible, and pA elements in $\langle T_A, T_B \rangle$ and elliptic, parabolic, and hyperbolic elements in $\mathrm{PSL}_2(\mathbb{R})$
(ii) Parabolic elements in $\rho(f)$ are exactly powers of T_A or T_B
(iii) If $\rho(f)$ is hyperbolic, then the dilatation of $[f] \in \mathrm{Mod}(\Sigma_g)$ is exactly the spectral radius of $\rho(f)$

We say $\langle T_A, T_B \rangle \subset \mathrm{Mod}(\Sigma_{g,n})$ are Thurston pA maps.

Remark 1.5. In $\mathrm{PSL}_2(\mathbb{R})$, hyperbolic elements are characterized by two real eigenvalues $\lambda, \frac{1}{\lambda}$ where $\lambda > 1$. The correspondence given in Theorem 1.4(i) guarantees that dilatation of pA maps are real.

In this paper, we will prove the following:

Theorem 1.6. *The minimal dilatation over all Thurston pA mapping classes for all filling pairs in $\Sigma_{g,n}$, $g \neq 0, 2$ is given by*

$$\frac{1}{2}((2g-1)^2 + (2g-1)\sqrt{(2g-1)^2 - 4} - 2)$$

for $n = 0$ and

$$\frac{1}{2}((2g-1+n)^2 + (2g-1+n)\sqrt{(2g-1+n)^2 - 4} - 2)$$

for $n \geq 1$. Additionally, we have the following characterization:

Genus	*Punctures*	$i(\alpha, \beta)$	*Minimal Dilatation*
$g = 0$	$n \geq 4$ *even*	$n-2$	$\frac{1}{2}((n-2)^2 + (n-2)\sqrt{(n-2)^2 - 4} - 2)$
$g = 0$	n *odd*	$n-1$	$\frac{1}{2}((n-1)^2 + (n-1)\sqrt{(n-1)^2 - 4} - 2)$
$g = 2$	$n \leq 2$	4	$7 + 4\sqrt{3}$
$g = 2$	$n > 2$	$2g+n-2$	$\frac{1}{2}((2g+n-2)^2 + (2g+n-2)\sqrt{(2g+n-2)^2 - 4} - 2)$

In the special case where $A = \{\alpha\},\ B = \{\beta\}$, the number μ is equal to the square of the *geometric intersection number* $i(\alpha, \beta) = |\alpha \cap \beta|$; we assume the two curves are in *minimal position*, i.e., are represenatives of their respective isotopy classes that minimize this quantity.

Thus, our representation is given by

$$T_\alpha \mapsto \begin{bmatrix} 1 & -i(\alpha,\beta) \\ 0 & 1 \end{bmatrix} \quad T_\beta \mapsto \begin{bmatrix} 1 & 0 \\ i(\alpha,\beta) & 1 \end{bmatrix}.$$

Remark 1.7. Determining whether two curves are in minimal position is quite simple using Thurston's *bigon criterion* (see [4], Section 1.2.4): two simple closed curves α and β in $\Sigma_{g,n}$ are in minimal position if, and only if, they do not form any bigons–embedded disks in $\Sigma_{g,n}$ whose boundary is union of an arc of α and an arc of β intersecting at exactly two points (note that the disk cannot be punctured).

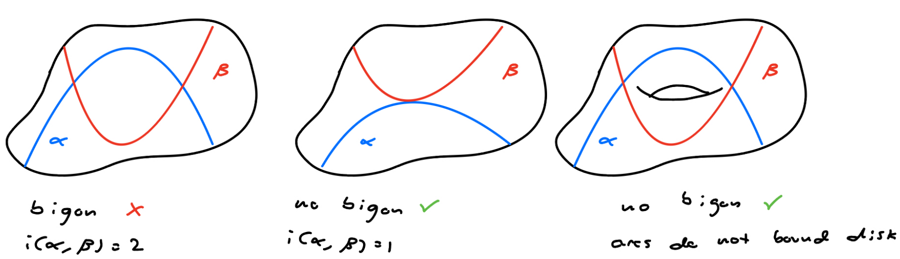

If α, β is a pair of filling curves on $\Sigma_{g,n}$ with $i(\alpha,\beta) \geq 2$, Section 14.1.2 of [4] gives that $\langle T_\alpha, T_\beta \rangle$ is free. Thus, the representation of T_α, T_β, given by $\left[\begin{smallmatrix} 1 & -i(\alpha,\beta) \\ 0 & 1 \end{smallmatrix}\right], \left[\begin{smallmatrix} 1 & 0 \\ i(\alpha,\beta) & 1 \end{smallmatrix}\right]$ generates a free subgroup in $\mathrm{PSL}_2(\mathbb{Z})$ of pseudo-Anosov mapping classes in $\mathrm{Mod}(\Sigma_{g,n})$. We use Thurston's construction to find the minimal dilatation maps in this subgroup of $\mathrm{Mod}(\Sigma_{g,n})$ by constructing minimally intersecting pairs of filling curves on $\Sigma_{g,n}$.

The dilatation of pA mapping class is related to the the length of closed geodesics in the *Moduli space* of $\Sigma_{g,n}$. Roughly speaking, $\mathcal{M}_{g,n}$ is a set of hyperbolic *structures* on $\Sigma_{g,n}$ up to actions of the mapping class group. To define it concretely, we must begin with Teichmuller space.

Definition 1.8. We say (X, φ) is a *marking* of $\Sigma_{g,n}$ if there is a homeomorphism $\varphi : \Sigma_{g,n} \to X$ where X has a hyperbolic metric.

Definition 1.9. The *Teichmuller space* of $\Sigma_{g,n}$ is this set of markings modulo homotopy (defined as isometry between two spaces X and Y which respects composition of markings).

There is a natural interpretation of the Teichmuller space of $\Sigma_{g,n}$. If all of the isotopy classes of essential simple closed curves on Σ_g are denoted as $\mathcal{S}$, then for any $\mathcal{X} \in \mathrm{Teich}(\Sigma_g)$ we may define a length function $\ell_{\mathcal{X}} : \mathcal{S} \to \mathbb{R}_{\geq 0}$. Each point $\mathcal{X} \in \mathrm{Teich}(\Sigma_g)$ corresponds to an isotopy class of a marking (X, φ). Then $\ell_{\mathcal{X}}(c)$ for some $c \in \mathcal{S}$ would be given by the length of the geodesic in X of the isotopy class $\varphi(c)$. Taking the pullback of the geodesic for an isotopy class from any fixed marking gives a new hyperbolic metric on our original surface $\Sigma_{g,n}$.

Moduli space generalizes this space of structures further by considering the action of the mapping class group on Teichmuller space. Formally, $\mathcal{M}_{g,n}$ is given by

$$\mathcal{M}_{g,n} = \mathrm{Teich}(\Sigma_{g,n}) / \mathrm{Mod}(\Sigma_{g,n})$$

Each point (equivalence class) in moduli space corresponds to a hyperbolic metric on $\Sigma_{g,n}$, and there is a bijective correspondance between $\mathcal{M}_{g,n}$ and $\mathrm{Mod}(\Sigma_{g,n})$ given as follows:

$$\{\text{length spectrum of } \mathcal{M}(\Sigma_{g,n})\} \leftrightarrow \{\text{set of dilatations in } \mathrm{Mod}(\Sigma_{g,n})\}$$

where the *length spectrum in* $\mathcal{M}_{g,n}$ is the length of closed geodesics. Thus, finding minimal dilatation Thurston pA maps minimizes the length of the geodesic between two hyperbolic metrics in moduli space.

1.2. **The case of the torus.** We begin the proof of Theorem 1.6 with a simple case: the torus. The result is well known–we simply include it for the sake of completeness.

Theorem 2.5 in [4] gives that $\mathrm{Mod}(T^2) \simeq \mathrm{SL}_2(\mathbb{Z})$. Thus, the proof of Theorem 1.6 amounts to finding the matrix in $SL_2(\mathbb{Z})$ with the smallest spectral radius exceeding 2, which is equivalent to minimizing the roots of the characteristic polynomial equation

$$x^2 - \mathrm{tr}(\alpha)x + 1.$$

In $\mathrm{SL}_2(\mathbb{Z})$, eigenvalues grow monotonically as a function of trace; the smallest magnitude trace is 3, so we have

$$x^2 - 3x + 1 = 0 \implies \lambda = \frac{3 + \sqrt{5}}{2}$$

Now, finding $\alpha = \left[\begin{smallmatrix} w & x \\ y & z \end{smallmatrix}\right]$ follows immediately from the conditions $w + z = 3, wz - xy = 1$: the solution is given by $\alpha = \left[\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right]$. Furthermore, α has two distinct real eigenvalues, so this solution is unique up to conjugacy.

2. Minimal dilatation in Λ_n

To solve the general problem, we find the minimal dilatation matrices in the subgroup of $\mathrm{SL}_2(\mathbb{Z})$ given by

$$\Lambda_n := \left\langle \begin{bmatrix} 1 & n \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ n & 1 \end{bmatrix} \right\rangle.$$

We define the smallest dilatation function λ in Λ_n to be

$$\lambda(\Lambda_n) := \inf\{|\lambda(\alpha)| : |\lambda(\alpha)| > 2, \alpha \in \Lambda_n\}$$

where $\lambda(\alpha)$ is the spectral radius (i.e. magnitude of the larger eigenvalue). Since Λ_n is discrete, this infimum must be realized.

For the case $n = 1$, we use the fact that $\Lambda_1 \simeq SL_2(\mathbb{Z})$. Then, from the previous section, we know that $SL_2(\mathbb{Z}) \simeq \mathrm{Mod}(T^2)$, so the minimal dilatation pA map is simply $\left[\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right]$.

For the general case, let $A = \left[\begin{smallmatrix} 1 & n \\ 0 & 1 \end{smallmatrix}\right]$ and $B = \left[\begin{smallmatrix} 1 & 0 \\ n & 1 \end{smallmatrix}\right]$. Since we are considering matrices in the subgroup generated by A and B, we can write any element M of the subgroup as a word $M = A^{k_1}B^{\ell_1}A^{k_2}B^{\ell_2}\cdots A^{k_n}$ for integers k_i, ℓ_i (we will ignore infinite words, since a simple argument shows they cannot attain the minimal dilatation). Since the subgroup is free, M corresponds to a unique *reduced word*, i.e., where we have $k_i, \ell_i \neq 0$ for all i. This gives a well-defined *word length* for each element of the subgroup.

We will consider a surface $\Sigma_{g,n}$ and assume $i(\alpha, \beta) \neq 2$ for any filling pair α, β; later (Remark 3.4) we show that Λ_2 is not the representation given by the Thurston construction for any number of genus or punctures.

Theorem 2.1. *The minimal dilatation $\lambda(\Lambda_n)$, $n > 2$, is given by $\frac{1}{2}(n^2 + n\sqrt{n^2-4} - 2)$, which is achieved by the matrix $\left[\begin{smallmatrix} 1-n^2 & -n \\ n & 1 \end{smallmatrix}\right]$.*

Fix $n > 2$. In $\mathrm{PSL}_2(\mathbb{Z})$, the spectral radius of a matrix α is given by the larger root of the characteristic polynomial

$$x^2 - \mathrm{tr}(\alpha)x + 1 = 0$$

Explicitly, these solutions are

$$x = \frac{\mathrm{tr}(\alpha) \pm \sqrt{(\mathrm{tr}(\alpha))^2 - 4}}{2}$$

We wish to minimize spectral radius over hyperbolic matrices, so we assume also that $|\mathrm{tr}(\alpha)| > 2$. Thus the spectral radius λ is monotonically increasing as a function of the magnitude of the trace; thus minimizing spectral radius is equivalent to minimizing trace magnitude. Thus we minimize over trace and then compute the corresponding dilatation.

To begin, we show the following, which was originally stated, but not proved in [3], Theorem 4a:

Proposition 2.2. *Let $\alpha \in \Lambda_n$, $n > 2$. Then α has the form*

$$\begin{bmatrix} 1 + k_1 n^2 & k_2 n \\ k_3 n & 1 + k_4 n^2 \end{bmatrix} \quad k_i \in \mathbb{Z}$$

Proof. For simplicity, we say a matrix γ is *congruent* (denoted $\gamma \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix}$) if γ takes on the form $\begin{bmatrix} 1 + k_1 n^2 & k_2 n \\ k_3 n & 1 + k_4 n^2 \end{bmatrix}$. We induct on the length of α. Say α has the form X_1X_2 for $X_1, X_2 \in \{A, A^{-1}, B, B^{-1}\}$. Assume without loss of generality that $X_2 = A$. Then we have the following cases:

$$A^{-1}A = \mathrm{id} \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix}$$
$$AA = \begin{bmatrix} 1 & 2n \\ 0 & 1 \end{bmatrix} \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix}$$
$$BA = \begin{bmatrix} 1 & n \\ n & 1+n^2 \end{bmatrix} \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix}$$
$$B^{-1}A = \begin{bmatrix} 1 & n \\ -n & 1-n^2 \end{bmatrix} \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix},$$

For $X_2 = A^{-1}, X_2 = B^{\pm 1}$, the cases follow similarly, so the base case is proven. Now assume for some k that $\prod_{i=1}^k X_i \cong \left[\begin{smallmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{smallmatrix}\right]$ for $X_i \in \{A, A^{-1}, B, B^{-1}\}$. Say $\prod_{i=1}^k X_i = \begin{bmatrix} a_i & b_i \\ c_i & d_i \end{bmatrix}$ for $a_i, b_i, c_i, d_i \in \mathbb{Z}[n]$. Then we have

$$A^{\pm 1} \prod_{i=1}^k X_i = \begin{bmatrix} a_i \pm c_i n & b_i \pm d_i n \\ c_i & d_i \end{bmatrix}$$

By assumption, $c_i \equiv 0 \mod n$, so $c_i n \equiv 0 \mod n^2$. Also, $d_i \equiv 0 \mod n^2$ so $d_i n \equiv 0 \mod n$. It follows that $A^{\pm 1} \prod_{i=1}^k X_i \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix}$.

Similarly,

$$B^{\pm 1} \prod_{i=1}^k X_i = \begin{bmatrix} a_i & b_i \\ c_i \pm a_i n & d_i \pm b_i n \end{bmatrix}$$

Using the same reasoning as above, $B^{\pm 1} \prod_{i=1}^k X_i \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix}$. Thus in all possible cases for X_{k+1}, we know $\prod_{i=1}^{k+1} X_i \cong \begin{bmatrix} 1 \mod n^2 & 0 \mod n \\ 0 \mod n & 1 \mod n^2 \end{bmatrix}$. □

Proof of Theorem 2.1. To prove Theorem 2.1, by Proposition 2.2 it suffices to minimize trace over all matrices of the form

$$\alpha = \begin{bmatrix} k_1 n^2 + 1 & k_2 n \\ k_3 n & k_4 n^2 + 1 \end{bmatrix}, \quad k_i \in \mathbb{Z} \text{ such that } (k_1 n^2 + 1)(k_4 n^2 + 1) - k_2 k_3 n^2 = 1 \tag{2.0.1}$$

Note the second constraint comes from the fact that $\alpha \in \mathrm{SL}_2(\mathbb{Z})$, so its determinant $(k_1 n^2 + 1)(k_4 n^2 + 1) - k_2 k_3 n^2$ is 1. Rearranging the determinant equation gives $k_2 k_3 = k_1 k_4 n^2 + (k_1 + k_4) \in \mathbb{Z}$. Thus given any fixed $k_1, k_4 \in \mathbb{Z}$, there always exists k_2, k_3 such that the matrix $\left[\begin{smallmatrix} 1+k_1 n^2 & k_2 n \\ k_3 n & 1+k_4 n^2 \end{smallmatrix}\right]$ is in $\mathrm{SL}_2(\mathbb{Z})$.

For any α given by 2.0.1, $|\operatorname{tr}(\alpha)|$ is given by

$$|2 + n^2(k_1 + k_4)|$$

which is the quantity which we want to minimize, subject to our constraints. Of course, this quantity is the smallest when $k_1 + k_4 = 0$–but this would imply $\operatorname{tr}(\alpha) = 2$, i.e., α is not hyperbolic. Thus, $k_1 + k_4 \neq 0$.

Consider the case where $|k_1+k_4| = 1$. If $k_1+k_4 = -1$, then $|\operatorname{tr}(\alpha)| = 2-n^2$ and if $k_1+k_4 = 1$, then $|\operatorname{tr}(\alpha)| = 2+n^2 > n^2-2$. Finally, for $|k_1+k_4| > 1$, $|2+n^2(k_1+k_4)| \in \{(k_1+k_4)n^2-2, (k_1+k_4)n^2+2\}$ and in either case is greater in magnitude than n^2-2.

It is left to show that a matrix in Λ_n achieves the minimum trace of n^2-2. Choosing $k_1 = -1, k_4 = 0$ gives the matrix $\begin{bmatrix} 1-n^2 & k_2 n \\ k_3 n & 1 \end{bmatrix}$, which implies $k_2 = -n, k_3 = n$. But this matrix is equal to AB, given by $\begin{bmatrix} 1-n^2 & -n \\ n & 1 \end{bmatrix}$. Thus both AB and BA (which are conjugate in Λ_n) achieve the minimum dilatation of $\frac{1}{2}(n^2 + n\sqrt{n^2-4} - 2)$. □

3. Construction of Filling Curves

The goal of this section is to obtain a lower bound for the intersection number of a pair of filling curves and subsequently construct examples achieving these minima. We will use the filling permutations of Aougab & Huang [1] and Aougab & Taylor [2] and the generalized filling permutations of Jeffreys [5], which gives us an algebraic way to describe "gluing patterns" of polygons. The idea is to construct polygons whose sides are identified in such a way that, once glued, they form the surface Σ_g with the glued sides becoming the filling curves α, β. Each polygon will correspond to a disk in the complement of $\alpha \cup \beta$ on Σ_g, so we can retroactively puncture the polygons to form $\Sigma_{g,n}$. Since we will "place" the punctures, our convention will be to treat them as marked points and thus exclude them from the Euler characteristic.

We begin with a general lower bound for the intersection number on any surface $\Sigma_{g,n}$ from Aougab and Huang ([1], Lemma 2.1).

Lemma 3.1. *Fix $g \geq 1, n \geq 0$. If α, β fill $\Sigma_{g,n}$, then*

$$i(\alpha, \beta) \geq 2g - 1$$

where i denotes geometric intersection number.

Proof. We model α, β as a 4-valent graph G (where vertices v are intersection points) since the complement $\Sigma_g \setminus (\alpha, \beta)$ is a union of topological discs D. The Euler characteristic of the graph must match that of $\Sigma_{g,n}$. We know

$$\sum_{v \in G} \deg_v(G) = 2|E| = 4|V| = 2i(\alpha, \beta)$$

Then we obtain

$$\chi(\Sigma_g) = 2 - 2g = |D| - 2i(\alpha, \beta) + i(\alpha, \beta)$$

and since $|D| \geq 1$, we have the result. □

This bound is only realized in the case when $n = 0$. For punctured surfaces, however, we can come very close. To construct an explicit example where equality is realized, we now introduce the notion of *filling permutations* from [1] and [5].

Fix a surface $\Sigma_{g,n}$ for $g, n \geq 0$ and let $i_{g,n}$ denote the minimal intersection number for a pair of filling curves on the surface. Suppose α, β fill $\Sigma_{g,n}$ and intersect $i(\alpha, \beta) = m \geq i_{g,n}$ times. Choose orientations for α, β and label the subarcs $\alpha_1, \ldots, \alpha_m$ where each subarc corresponds to arcs beginning and ending at intersection points. Similarly, label subarcs of β as $\beta_1, \ldots, \beta_m$.

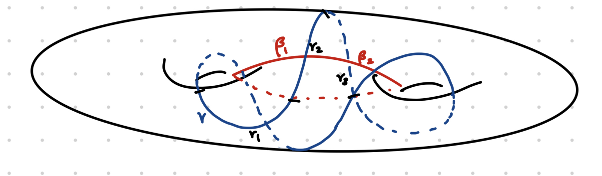

Let $Q = Q_{\alpha,\beta} \in S_{4m}$ be given as $Q = (1, 2, \ldots, 4m)^{2m}$. We note that Q changes nothing except the orientation of every edge, i.e., it sends j to k if and only if the jth and kth elements of A are inverses of each other. Finally, define $\tau = \tau_{\alpha,\beta} \in S_{4m}$ as

$$\tau = (1, 3, 5, \ldots, 2m-1)(2, 4, 6, \ldots, 2m)(4m-1, 4m-3, \ldots, 2n+1)(4m, 4m-2, \ldots, 2m+2).$$

The first cycle represents sending α_i to α_{i+1}, the second β_i to β_{i+1}, the third α_k^{-1} to α_{k+1}^{-1}, and the fourth β_k^{-1} to β_{k+1}^{-1}. In other words, τ moves each arc in α to the next arc of α with the same orientation, and similarly for β.

We will say that a permutation is parity-respecting if it sends even numbers to even numbers and odd numbers to odd numbers and parity-reversing if it sends even numbers to odd numbers and odd numbers to even numbers.

The following lemma from Jeffreys ([5], Lemma 2.3) gives the conditions necessary to define a filling permutation on a surface $\Sigma_{g,n}$; we will subsequently construct the filling curves by finding a permutation that satisfies these hypotheses.

Lemma 3.2. *Let α, β be a filling pair on $\Sigma_{g,n}$ with $i(\alpha, \beta) = m \geq i_{g,n}$. Then, $\sigma = \sigma_{\alpha,\beta}$ satisfies $\sigma Q \sigma = \tau$. Conversely, a parity-reversing permutation $\sigma \in S_{4m}$ consisting of $m+2-2g$ cycles and no more than n 2-cycles that satisfies the above relation defines a filling pair on $\Sigma_{g,n}$ with intersection number m.*

Proof. Take $j \in \{1, 2, \ldots, 4m\}$. The edge labelled as the jth element of A is followed by the edge labelled by the $\sigma(j)$th element of A. Additionally, $Q(\sigma(j))$ is the inverse of $\sigma(j)$ and the edge labelled by the $\sigma(Q(\sigma(j)))$th element of A follows that labelled by the $Q(\sigma(j))$th element. The latter is exactly the edge labelled by the $j+1$th element of A, so $\sigma Q \sigma = \tau$.

Conversely, assume $\sigma \in S_{4m}$ satisfies the hypothesis. Since σ is parity-reversing, it consists of cycles of even length, which we associate to polygons with a corresponding number of sides. Puncture every 2-gon and any remaining polygons at most once until all punctures have been placed. Note that this is possible because a filling pair on $\Sigma_{g,n}$ intersecting $i_{g,n}$ times gives a graph with $2 - 2g + i_{g,n}$ faces, which must be at least equal to the number of punctured discs. Hence, $m \geq i_{g,n}$ implies $m + 2 - 2g \geq n$. There are at most n 2-cycles, so we can puncture all of the bigons to ensure α and β are in minimal position.

Label each polygon clockwise with the elements of the corresponding cycle and glue polygons together according to the elements of A identified with these labels. Each edge occurs eactly once with each orientation, so the resulting closed surface has n punctures.

Now we show that the Euler characteristic of this surface is given by $2 - 2g$; then, by the classification of surfaces, this surface would be homeomorphic to our original $\Sigma_{g,n}$. We know that there are $m + 2 - 2g$ faces and $2m$ edges, so $V - E + F = 2 - 2g$ if, and only if, there are m equivalence classes of vertices under gluing.

Each equivalence class of vertices contains exactly 4 vertices. The one between α_k and β_{j+1} glues to that between β_{j+1}^{-1} and α_{k+1}; similarly the vertex between β_{j+1}^{-1} and α_{k+1} glues to the one in between α_{k+1}^{-1} and β_j^{-1}. The one between α_{k+1}^{-1} and β_j^{-1} glues to the one between β_j, α_k^{-1}, which glues to the original vertex between α_k, β_{j+1}. This creates a 4-cycle and thus the equivalence class under the gluing consists of exactly 4 vertices. Hence we have $4m/4 = m$ total equivalence classes, as desired.

Finally, since the last vertex of α_k is identified to the first of α_{k+1}, it follows that the α-arcs form a simple closed curve α. Performing the same for the β-arcs, we acquire a surface of genus g filled by two simple closed curves α, β which intersect exactly m times. □

Now we have the necessary ingredients to compute the minimal realized number of intersection points on $\Sigma_{g,n}$; we closely follow the proof in [2].

Proposition 3.3. *If α, β are minimally intersecting filling curves on $\Sigma_{g,n}$, $g \neq 0, 2$, $i(\alpha, \beta) = 2g-1$ if $n = 0$ and $i(\alpha, \beta) = 2g + n - 2$ if $n \geq 1$.*

Proof. Using the same argument as in Lemma 3.1, we have that $i(\alpha, \beta) = 2g + n - 2 + |D|$ where D is the disks in the complement of α, β. Thus we have the lower bounds and it is left to show that these bounds are realized. The first case is given explicitly by Lemma 3.2; for the second, we induct on n. When $n = 1$, $2g - 1 = 2g + n - 2$. Thus the filling curves given in Proposition 4.2 which have a single disk D in their complement still fill $\Sigma_{g,1}$, obtained by puncturing D once. When $g = 1$ the formula for intersection number on the torus (that being a (p, q) and (r, s) intersecting precisely $ps - qr$ times, [4]) gives a simple way to find two curves intersecting exactly p times for $n \geq 1$. The complement of these two curves is n topological disks, and puncturing each gives $2g + n - 2$ intersections on $\Sigma_{1,n}$.

Now we describe a method to give a filling pair on $\Sigma_{g,n+2}$ given a pair on $\Sigma_{g,n}$ which gives 2 more intersection points. As before, let α, β be a filling pair on $\Sigma_{g,n}$, and orient and label them into subarcs $\alpha_1, \dots, \alpha_{i(\alpha,\beta)}$ and $\beta_1, \dots, \beta_{i(\alpha,\beta)}$ according to intersection points. Suppose α_1 and $\beta_{i(\alpha,\beta)}$ cross each other. Then pushing α_1 across $\beta_{i(\alpha,\beta)}$ and back over forms 2 bigons. Puncturing them both gives the same pair of filling curves on $\Sigma_{g,n+2}$ with intersection number $i(\alpha, \beta) + 2$.

If $n = 2k + 1$ is odd and $g > 2$, take a filling pair whose complement is connected. Puncturing this region gives a filling pair on $\Sigma_{g,1}$. Then performing the above double bigon construction k times gives a pair of filling curves on $\Sigma_{g,2k+1} = \Sigma_{g,n}$ which intersects $2g + n - 2$ times. For n even, the argument generalizes if there exists a filling pair α, β on $\Sigma_{g,0}$ intersecting $2g$ times. Puncturing both gives α, β which fill and intersect $2g + p - 2$ times. Thus it suffices to construct these filling curves on $\Sigma_{g,0}$ for $g > 2$ intersecting $2g$ times.

Consider the following two curves on $\Sigma_{2,0}$:

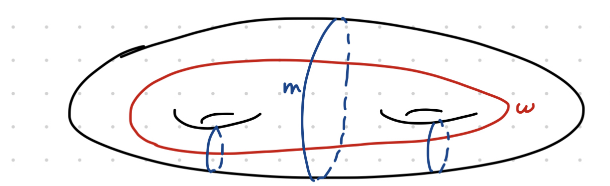

The right Dehn twist of w about m, coupled with w (call these two curves (x, y)) fill $\Sigma_{2,0}$ and its complement is 4 topological disks:

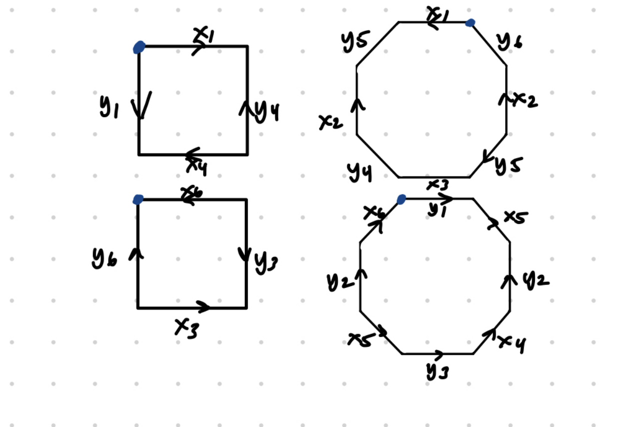

Then x, y intersect 6 times (given there are 6 subarcs of x, y.) Take $\Sigma_{1,0}$ with the filling pair above intersecting twice. Cut a small disk about each intersection point to get $\Sigma'_{1,0}$. This produces a torus with boundary component and modified curves $\overline{\alpha_1}, \overline{\alpha_2}$. Now take $\Sigma_{2,0}$ with (x, y) as above and cut a small disk about the blue dot (which, in the gluing scheme, is identified as one point on $\Sigma_{2,0}$). Cutting this disk gives $\Sigma'_{2,0}$ which is of genus two with $\overline{x}, \overline{y}$. Glue $\Sigma'_{1,0}$ and $\Sigma'_{2,0}$, concatenating $\overline{\alpha_1}$ with $\overline{x}$ and $\overline{\alpha_2}$ with $\overline{y}$. We claim this gives a filling pair α, β on $\Sigma_{3,0}$ intersecting 6 $(= 2g)$ times.

Proving this amounts to taking γ a simple closed curve and assuming it is disjoint from α, β. Then the projection of γ to the curve graph of the subsurface $\Sigma'_{2,0}$ $\pi_{\Sigma'_{2,0}}(\gamma)$ is disjoint from $\overline{x}, \overline{y}$. Thus it must be homotopic to $\partial\Sigma'_{2,0}$ since $\overline{x}, \overline{y}$ fill. Thus it is homotopic into $\Sigma'_{1,0}$. Since $\overline{\alpha_2}, \overline{\alpha_2}$ fill this subsurface, it follows that γ cannot be disjoint from both.

To obtain such a construction on an odd genus, we can iterate by choosing a pair intersecting $2(2(k-1)+1)$ times on $\Sigma_{2(k-1)+1,0}$, cutting a disk about an intersection point and gluing to $\Sigma'_{2,0}$. $\square$

A similar application of the double bigon construction gives minimal intersection numbers for $\Sigma_{g,n}$ for $g = 0, 2$ (see [2], Lemma 3.1 and [5] Theorem 3.3). We summarize the results as follows:

Genus	Punctures	$i(\alpha, \beta)$
$g = 0$	$n \geq 4$ even	$n - 2$
$g = 0$	$n \geq 4$ odd	$n - 1$
$g = 2$	$n \leq 2$	4
$g = 2$	$n > 2$	$2g + n - 2$

Remark 3.4. The case $g = 0$, $n < 4$ is not considered because the filling curves have intersection number zero: if there are two or fewer punctures then a single curve fills and if there are exactly three punctures then the filling pair does not intersect.

We now see that the proof of Theorem 1.6 immediately follows from the Thurston construction, Proposition 3.3, and Theorem 2.1.

Remark 3.5. We note that the value of $n = 2$ is never realized for any $\Sigma_{g,n}$ justifying the exclusion of this value in Proposition 3.3.

4. Future Directions

Revisiting the Thurston construction, we recall that for two multicurves A and B which fill $\Sigma_{g,n}$ we obtain a representation $\rho : \langle T_A, T_B \rangle \to \mathrm{PSL}_2(\mathbb{R})$ given by

$$T_\alpha \mapsto \begin{bmatrix} 1 & -\mu^{1/2} \\ 0 & 1 \end{bmatrix} \quad T_\beta \mapsto \begin{bmatrix} 1 & 0 \\ \mu^{1/2} & 1 \end{bmatrix}.$$

Throughout this paper we exclusively explored the case where A and B are single curves α and β, respectively, but the problem of finding the minimal dilatation Thurston pA map extends to the general case of multicurves $A = \{\alpha_1, \dots, \alpha_k\},\ B = \{\beta_1, \dots, \beta_\ell\}$. Since we have two families of curves the number μ is no longer just the intersection number between two curves; instead, we form the $k \times \ell$ matrix N whose (n, m) entry is given by

$$N_{n,m} = i(\alpha_n, \beta_m)$$

and take μ to be the *Perron-Frobenius eigenvalue* of $N^\top N$ (note that we must work with this matrix instead of N since the latter is not necessarily square). We defer the reader to [4], Section 14.1.2 for some background on Perron-Frobenius theory.

The advantage to multicurves is that it is much easier to construct a pair of filling multicurves than simply a filling pair, but they also add a considerable element of complexity. The number μ now depends on all entries of the matrix N and not in a trivial way, since there is not a clear-cut relationship between how the eigenvalues of matrix change with its entries. Thus, even determining the minimal dilatation map for families of two curves, $A = \{\alpha_1, \alpha_2\},\ B = \{\beta_1, \beta_2\}$ would likely require a large leap in understanding.

Another interesting angle to considering is that of most efficient mixing, i.e., the Thurston pA map with "highest" dilatation. Of course, there is no maximum dilatation in a literal sense, since for any map f we can let $g = f^n$ to obtain a higher dilatation. However, we can define a "most efficient" map by fixing generators, say A and B and maximizing dilatation over reduced word length. It is not clear for which generators (if any) this supremum will exist or be attained, but inducting on word length and analyzing the growth of the entries of the word would likely be a good start. The problem for multicurves is again more difficult in this context for the same reasons as finding the minimal dilatation, namely the slippery relationship between μ and the matrix N.

Acknowledgements

The authors would first like to thank Professor Benson Farb and Faye Jackson for their support and guidance throughout the summer. The authors are also grateful for the help of Noah Caplinger and Peter Huxford, whose insights were essential to this paper. The second author would also like to thank Professor Peter May for organizing the 2024 Chicago REU and the first author for her patience and algebra knowledge. The first author thanks the second author for his patience and topology knowledge.

References

[1] Tarik Aougab and Shinnyih Huang. Minimally intersecting filling pairs on surfaces. *Algebraic & Geometric Topology*, 15(2):903–932, 2015.

[2] Tarik Aougab and Samuel J. Taylor. Small intersection numbers in the curve graph. *Bulletin of the London Mathematical Society*, 46(5):989–1002, 2014.

[3] Anastasiia Chorna, Katherine Geller, and Vladimir Shpilrain. On two-generator subgroups in sl2(z), sl2(q), and sl2(r). *Journal of Algebra*, 478:367–381, 2017.

[4] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups (PMS-49)*. Princeton University Press, 2012.

[5] Luke Jeffreys. Minimally intersecting filling pairs on the punctured surface of genus two. *Topology and its Applications*, 254:101–106, 2019.