

Minimally intersecting filling pairs on surfaces

TARIK AOUGAB
SHINNYIH HUANG

Let S_g denote the closed orientable surface of genus g . We construct exponentially many mapping class group orbits of pairs of simple closed curves which fill S_g and intersect minimally, by showing that such orbits are in correspondence with the solutions of a certain permutation equation in the symmetric group. Next, we demonstrate that minimally intersecting filling pairs are combinatorially optimal, in the sense that there are many simple closed curves intersecting the pair exactly once. We conclude by initiating the study of a topological Morse function \mathcal{F}_g over the moduli space of Riemann surfaces of genus g , which, given a hyperbolic metric σ , outputs the length of the shortest minimally intersecting filling pair for the metric σ . We completely characterize the global minima of \mathcal{F}_g and, using the exponentially many mapping class group orbits of minimally intersecting filling pairs that we construct in the first portion of the paper, we show that the number of such minima grows at least exponentially in g .

57M20, 57M50

1 Introduction

Let $S_{g,b}$ denote the orientable surface of genus g with b boundary components, and let $\text{Mod}(S_{g,0})$ denote the *mapping class group* of the closed surface $S_{g,0}$, the group of all orientation-preserving homeomorphisms of $S_{g,0}$ modulo isotopy (see Farb and Margalit [7] for background).

A pair of curves on S_g (shorthand for $S_{g,0}$) is said to *fill* the surface if the complement of their union is a disjoint union of topological disks.

Let $\mathcal{I}(S_g)$ denote the set of all isotopy classes of simple closed curves on S_g . Then $\text{Mod}(S_g)$ acts coordinatewise on the product $\mathcal{I}(S_g) \times \mathcal{I}(S_g)$, and this action descends to an action on the quotient

$$\mathcal{P}(S_g) := \mathcal{I}(S_g) \times \mathcal{I}(S_g) / [(\alpha, \beta) \sim (\beta, \alpha)].$$

Our first result establishes the exponential growth of the $\text{Mod}(S_g)$ -orbits of $\mathcal{P}(S_g)$ which fill and intersect minimally:

Theorem 1.1 *There exists a function $f(g) \sim 3^{g/2}/g^2$ satisfying the following: Let $N(g)$ denote the number of $\text{Mod}(S_g)$ -orbits of minimally intersecting filling pairs on S_g . Then:*

$$f(g) \leq N(g) \leq 2^{2g-2}(4g - 5)(2g - 3)!$$

In the statement of Theorem 1.1, \sim denotes the growth-rate relation

$$r(g) \sim s(g) \iff 0 < \lim_{g \rightarrow \infty} \frac{r(g)}{s(g)} < \infty.$$

We compare Theorem 1.1 to Anderson, Parlier and Pettet’s [2, Theorem 1], which says that if $\Omega = \{\alpha_1, \dots, \alpha_n\}$ is a filling set of simple closed curves on S_g , pairwise intersecting no more than K times, then

$$n^2 - n \geq \frac{4g - 2}{K}.$$

Thus, when $n = 2$, $K \geq 2g - 1$, and indeed we show in Section 2 that a minimally intersecting filling pair intersects $2g - 1$ times if and only if $g \neq 2$ (when $g = 2$, four intersections is minimal).

Given a pair of filling curves (α, β) , there are potentially many ways of defining a sort of combinatorial “complexity” for the pair. One such candidate is to count the number of simple closed curves intersecting the filling pair a small number of times. Heuristically, one expects a minimally intersecting filling pair to only “barely” fill S_g . Therefore, there should be many simple closed curves intersecting the union $\alpha \cup \beta$ only once. Indeed, we show the following:

Theorem 1.2 *Let (α, β) be a filling pair on S_g , $g > 2$, and define $T_1(\alpha, \beta) \in \mathbb{N}$ to be the number of simple closed curves intersecting $\alpha \cup \beta$ only once. Then $T_1(\alpha, \beta) \leq 4g - 2$, with equality if (α, β) is minimally intersecting.*

Let $\mathcal{M}(S_g)$ denote the *moduli space* of Riemann surfaces of genus g , which is identified with the space of all complete, finite-volume hyperbolic metrics on S_g up to isometry. The *systole* function $\mathcal{R}_g: \mathcal{M}(S_g) \rightarrow \mathbb{R}_+$ outputs the length of the shortest closed geodesic for a given hyperbolic metric $\sigma \in \mathcal{M}(S_g)$.

\mathcal{R}_g has been studied extensively (see Buser and Sarnak [6], Parlier [11], Balacheff and Parlier [3], and Schaller [12; 13; 14]) in part due to the fact that it is a *topological Morse function* (see Akrouf [1]), a generalization of the notion of a classical Morse function to functions which are not necessarily smooth, and for which versions of the Morse inequalities also hold. Thus a careful analysis of the critical points of \mathcal{R}_g can yield naturally arising cellular decompositions of $\mathcal{M}(S_g)$.

Motivated by this program, we initiate the study of another topological Morse function \mathcal{F}_g , the *filling pair* function, which outputs the length of the shortest minimally intersecting filling pair with respect to a metric $\sigma \in \mathcal{M}(S_g)$. Here, the length of a pair of curves is simply the sum of the individual lengths. We show the following.

Theorem 1.3 \mathcal{F}_g is proper and a topological Morse function. For any $\sigma \in \mathcal{M}_g$,

$$\mathcal{F}_g(\sigma) \geq \frac{1}{2}m_g,$$

where

$$m_g = (8g - 4) \cdot \cosh^{-1} \left(2 \left[\cos \left(\frac{2\pi}{8g-4} \right) + \frac{1}{2} \right] \right)$$

denotes the perimeter of a regular, right-angled $(8g-4)$ -gon. Furthermore, define

$$\mathcal{B}_g := \{ \sigma \in \mathcal{M}_g : \mathcal{F}_g(\sigma) = \frac{1}{2}m_g \}.$$

Then \mathcal{B}_g is finite and grows at least exponentially in g . If $\sigma \in \mathcal{B}_g$, the injectivity radius of σ is at least

$$\frac{1}{2} \cosh^{-1} \frac{9}{\sqrt{73}} \approx 0.3253.$$

Organization of paper In Section 2, we prove Theorem 1.1; in Section 3, we prove Theorem 1.2; and in Section 4, we prove Theorem 1.3.

Acknowledgements The authors would like to thank Yair Minsky for carefully reading through a draft of this project and for his many helpful suggestions. The first author would also like to thank Dan Margalit, Chris Aretines, Subhojoy Gupta, Babak Modami, and Jean Sun for many enlightening conversations regarding this work.

2 Counting minimally intersecting filling pairs

In this section we prove Theorem 1.1. We begin with some preliminary notions and notation. All curves will be simple and closed. Let \sim denote the free homotopy relation.

Let $\chi(S_g) = 2 - 2g$ denote the standard Euler characteristic.

If α, β are two curves on S_g , we define the *geometric intersection number* of the pair (α, β) , denoted $i(\alpha, \beta)$, by

$$i(\alpha, \beta) = \min_{\alpha' \sim \alpha} |\alpha' \cap \beta|,$$

where the minimum is taken over all curves α' freely homotopic to α .

If $|\alpha \cap \beta| = i(\alpha, \beta)$, we say that α and β are in *minimal position*.

Lemma 2.1 Suppose (α, β) is a pair of curves which fill S_g , $g \geq 1$. Then

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

Proof Assume α, β are in minimal position. That (α, β) fills S_g implies that the complement $S_g \setminus (\alpha \cup \beta)$ is a disjoint union D of topological disks. Thus we can view the 4-valent graph $\alpha \cup \beta$ as the 1-skeleton of a cellular decomposition of S_g .

The number of 0-cells is $i(\alpha, \beta)$, and 4-valency of $\alpha \cup \beta$ implies that the number of edges is $2i(\alpha, \beta)$. Therefore

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

Then since $|D| \geq 1$, we obtain the desired result. □

We remark that the proof of Lemma 2.1 implies that $i(\alpha, \beta) = 2g - 1$ exactly when $S_g \setminus (\alpha \cup \beta)$ consists of a single disk.

2.1 Filling permutations

Given a minimally intersecting filling pair (α, β) , choose orientations for α and β . Label the subarcs of α as $\alpha_1, \dots, \alpha_{2g-1}$, separated by the $2g - 1$ intersection points in accordance with the chosen orientation for α (and after choosing an initial arc for α), and similarly for $\beta_1, \dots, \beta_{2g-1}$. Once labeled, we call such a pair an *oriented filling pair*. A filling pair without the extra structure of an orientation is an *ordinary filling pair*.

The action of $\text{Mod}(S_g)$ on ordinary filling pairs lifts naturally to an action on the set of oriented filling pairs. Given an oriented filling pair, we obtain an ordinary one by forgetting the orientation.

Then associated to the oriented pair is a permutation $\sigma^{(\alpha, \beta)} \in \Sigma_{8g-4}$, the permutation group on $8g - 4$ symbols, which we define as follows.

We will identify Σ_{8g-4} as the permutations of the ordered set

$$A(g) = \{\alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_{2g-1}, \beta_{2g-1}, \alpha_1^{-1}, \beta_1^{-1}, \dots, \alpha_{2g-1}^{-1}, \beta_{2g-1}^{-1}\}.$$

Cutting along $\alpha \cup \beta$ produces an $(8g-4)$ -gon $P(\alpha, \beta)$, whose edges correspond to elements of $A(g)$. Pick an initial edge of $P(\alpha, \beta)$ and orient it clockwise. Then $\sigma^{(\alpha, \beta)}$ is defined by

$$\begin{aligned} \sigma^{(\alpha, \beta)}(j) &:= k \\ &\iff \text{the } j^{\text{th}} \text{ element of } A(g) \text{ coincides with the } (k - 1)^{\text{st}} \text{ edge of } P(\alpha, \beta). \end{aligned}$$

Note that $\sigma^{(\alpha,\beta)}$ must be an $(8g-4)$ -cycle, and corresponds to the permutation of $A(g)$ obtained by rotating $P(\alpha, \beta)$ counterclockwise by $2\pi/(8g-4)$.

Let $Q_g \in \Sigma_{8g-4}$ denote the standard ordered $(8g-4)$ -cycle, which in cycle notation is

$$Q_g = (1, 2, 3, \dots, 8g-4).$$

Note also that Q_g^{4g-2} sends an edge to itself, oriented in the opposite direction. Let $\tau_g \in \Sigma_{8g-4}$ denote the permutation

$$\tau_g = (1, 3, 5, \dots, 4g-3)(2, 4, 6, \dots, 4g-2) \cdot (8g-5, 8g-7, \dots, 4g-1)(8g-4, 8g-6, \dots, 4g),$$

which corresponds to moving forward by one subarc in each of α, β .

In what follows, in order to simplify the relevant permutation equations, we will suppress the reference to g and (α, β) when referring to a permutation (eg Q_g will be written as Q and $\sigma^{(\alpha,\beta)}$ as σ).

We say that a permutation γ respects parity, or is parity-respecting, if

$$i \equiv j \pmod{2} \iff \gamma(i) \equiv \gamma(j) \pmod{2}.$$

Lemma 2.2 *Let (α, β) be a minimally intersecting filling pair on S_g . Then the permutation $\sigma = \sigma^{(\alpha,\beta)}$ satisfies the equation*

$$\sigma Q^{4g-2} \sigma = \tau.$$

Conversely, if $\sigma \in \Sigma_{8g-4}$ is a parity-respecting $(8g-4)$ -cycle which solves this equation, then σ defines a minimally intersecting filling pair on S_g .

Proof Let $k \in \{1, \dots, 8g-4\}$. Then $\sigma(k)$ is the labeling of the edge of $P(g)$ immediately following the edge labeled with the k^{th} element of $A(g)$. The ordering of $A(g)$ has been chosen so that $Q^{4g-2}(j)$ is precisely the element of $A(g)$ identified with the inverse arc to j . Thus $\sigma Q^{4g-2} \sigma$ sends k to the next arc along whichever curve k is a part of, with respect to the chosen orientation of that curve, which is equal to $\tau(k)$.

Conversely, let σ satisfy the above permutation equation, and suppose further that σ is an $(8g-4)$ -cycle and is parity-respecting. Then associated to σ is a clockwise labeling of the edges of an $(8g-4)$ -gon (up to cyclic permutation) from the set $\{1, \dots, 8g-4\}$, such that edges alternate between receiving even and odd labels. This is simply the labeling obtained by choosing an initial edge of the polygon, and labeling edges in order from the initial edge in accordance with the permutation σ written in cycle notation.

This produces a labeling of the edges by the elements of $A(g)$, which we in turn interpret as a gluing pattern; each edge glues to its inverse. Thus we obtain a closed surface S , which is orientable because each edge is glued to one which is oppositely oriented. It remains to check that the Euler characteristic of the surface is $2 - 2g$, and that the boundary of the polygon projects to a minimally intersecting filling pair.

The equality $\chi(S) = 2 - 2g$ is equivalent to there being $2g - 1$ equivalence classes of vertices under the gluing. Indeed, the permutation equation ensures that each equivalence class contains precisely four vertices; as seen in Figure 1, the vertex v separating the edges labeled α_k and β_{j+1} glues to the vertex separating the edges labeled β_{j+1}^{-1} and α_{k+1} .




Figure 1: The permutation equation implies that if α_k and β_{j+1} share a vertex in common along the boundary of the polygon $P(g)$, then β_{j+1}^{-1} and α_{k+1} must also share a vertex in common, and these two vertices are glued together to obtain S . Iterating this argument shows that each gluing class contains four vertices.

Similarly, the vertex separating β_{j+1}^{-1} and α_{k+1} is glued to the vertex separating α_{k+1}^{-1} and β_j^{-1} , which in turn glues to the vertex separating β_j and α_k^{-1} . Finally, this fourth vertex glues to v and the cycle is complete.

By the same analysis, the terminal vertex of α_k is identified to the initial vertex of α_{k+1} , and therefore the α -arcs are concatenated in the quotient to produce a single simple closed curve α , and similarly for the β -arcs. □

We also note that if a permutation in Σ_{8g-4} is an $(8g - 4)$ -cycle and parity-respecting, it must send each even number to an odd number.

Henceforth, any permutation satisfying the requirements of Lemma 2.2 (being an $(8g - 4)$ -cycle, parity-respecting, and satisfying the permutation equation) will be called a *filling permutation*.

Filling permutations parametrize oriented filling pairs. Therefore, in order to count ordinary filling pairs, it suffices to count filling permutations, and divide by the number of distinct orientations that a fixed ordinary filling pair admits. Lemma 2.3 characterizes which filling permutations correspond to the same ordinary filling pair.

Lemma 2.3 *If $\Gamma_1 = (\alpha_1, \beta_1), \Gamma_2 = (\alpha_2, \beta_2)$ are two minimally intersecting filling pairs on S_g in the same $\text{Mod}(S_g)$ -orbit, then $\sigma^{(\alpha_1, \beta_1)} = \sigma^{(\alpha_2, \beta_2)}$, modulo conjugation by permutations of the form*

$$\mu_g^l \kappa_g^k \delta_g^j \eta_g^i \quad \text{with } l, i \in \{0, 1\}, j, k \in \{0, 1, \dots, 2g - 2\},$$

where

$$\begin{aligned} \mu_g &= (1, 2)(3, 4) \cdots (8g - 5, 8g - 4), \\ \kappa_g &= (1, 3, 5, \dots, 4g - 3)(4g - 1, 4g + 1, \dots, 8g - 7, 8g - 5), \\ \delta_g &= (2, 4, 6, \dots, 4g - 2)(4g, 4g + 2, \dots, 8g - 4), \\ \eta_g &= (1, 4g - 1)(3, 4g + 1) \cdots (4g - 3, 8g - 5). \end{aligned}$$

Proof If Γ_1 and Γ_2 are oriented filling pairs in the same $\text{Mod}(S_g)$ -orbit, there is a homeomorphism of S_g taking Γ_1 to Γ_2 , which lifts to a homeomorphism of the $(8g - 4)$ -gon that is simplicial in the sense that it fixes the boundary and sends edges to edges and vertices to vertices.

The only option for such a map is a dihedral symmetry of the regular Euclidean $(8g - 4)$ -gon; since the map must also be orientation-preserving, this leaves rotation as the only option. Thus the filling permutation for Γ_2 , written in cycle notation, is obtained from that of Γ_1 by a cyclic rotation, and thus determines the same underlying permutation of Σ_{8g-4} .

However, in general there will exist oriented filling pairs which are not in the same $\text{Mod}(S_g)$ -orbit, but whose underlying ordinary filling pairs are. Nevertheless, the fiber of the map from oriented pairs to ordinary pairs is in one-to-one correspondence with the set of distinct orientations with which one can equip an ordinary pair. Fixing an ordinary filling pair, any choice of orientation can be obtained from any other by simply choosing a different initial arc along α (corresponding to conjugating by a power of κ_g) or β (conjugating by δ_g), reversing the direction of α (conjugating by η_g), and interchanging α and β (conjugating by μ_g). □

We call permutations of the form $\mu_g^l \kappa_g^k \delta_g^j \eta_g^i$ *twisting* permutations. In particular, we have the following bound for such permutations.

Remark 2.4 In light of Lemma 2.3, to obtain a lower bound on $N(g)$, it suffices to find a lower bound on the set of filling permutations, and divide by an upper bound on the number of twisting permutations. Since κ_g and δ_g both have order $2g - 1$ and η_g, μ_g are involutions, there are at most $4 \cdot (2g - 1)^2$ such permutations.

2.2 Proof of the lower bound

Let (α, β) be an oriented minimally intersecting filling pair on S_g , and label the intersections v_1, \dots, v_{2g-1} in accordance with the chosen direction along α . We first describe a way of “extending” (α, β) to a minimally intersecting pair on S_{g+2} (see Remark 2.6 for an explanation of why we do not extend to S_{g+1}).

Fix some $k, 1 \leq k \leq 2g - 1$. To create a minimally intersecting filling pair on S_{g+2} , we first excise a small disk around v_k , yielding $S_{g,1}$. Doing so turns α and β into arcs $\tilde{\alpha}, \tilde{\beta}$ which intersect $2g - 2$ times and which fill $S_{g,1}$. To complete the construction, we will glue a decorated copy Z of $S_{2,1}$ to $S_{g,1}$ along its boundary component; by “decorated”, we mean that Z will come equipped with a special pair of oriented arcs (a, b) , such that if we concatenate $\tilde{\alpha}$ with a , and $\tilde{\beta}$ with b , we obtain a minimally intersecting filling pair on S_{g+2} .

We now describe the arcs (a, b) . Consider the 3–component multicurve m on $S_{2,0}$ pictured in red in Figure 2. The curve w intersects each component of m ; let $T_m(w)$ denote the right Dehn twist of w around m .

Lemma 2.5 *The pair $(T_m(w), w)$ fills $S_{2,0}$, and the complement of the pair consists of four disks. Furthermore, when $(T_m(w), w)$ are drawn in minimal position, there exists an intersection point which is on the boundary of all of these complementary disks.*




Figure 2: The simple closed curve w intersects each component of the red pants decomposition.

Proof The proof is completely constructive. $T_m(w)$ and w are pictured in Figure 3 ($T_m(w)$ in red, w in black). The representatives drawn in Figure 3 intersect six times;

thus both curves are subdivided into six arcs: $\{x_1, \dots, x_6\}$ for w and $\{y_1, \dots, y_6\}$ for $T_m(w)$, ordered cyclically with respect to some chosen orientations of w and $T_m(w)$, respectively.

To examine the complement $S_{2,0} \setminus (T_m(w) \cup w)$, simply cut along the union of these two curves to obtain the four labeled polygons in Figure 4. $S_{2,0}$ is obtained by gluing these four polygons back together along the oriented edge labelings. The four green vertices in Figure 4 are all identified together in this gluing. \square

To complete the construction, we puncture $S_{2,0}$ at this green point by removing a small disk centered there to obtain Z , a copy of $S_{2,1}$; define the arc a to be what remains of w , and what remains of $T_m(w)$ is defined to be the second arc b . We then glue Z to $S_{g,1}$ along the boundary of the disk we removed centered at the intersection point v_k , matching the endpoints of a with those of $\tilde{\alpha}$ and the endpoints of b with $\partial\tilde{\beta}$, obtaining two simple closed curves (α_k, β_k) on $S_{g+2,0}$. Here, the subscript k denotes the fact that we obtained this pair of curves by excising a disk around the intersection point v_k .




Figure 3: w in black and $T_m(w)$ in red. The green point is the desired intersection point, which is on the boundary of all four complementary regions.

We claim that (α_k, β_k) fill $S_{g+2,0}$. Assuming this, they must intersect minimally because, by construction, $i(\alpha_k, \beta_k) = 2g + 3 = 2(g + 2) - 1$. Indeed, because the green point of intersection was on the boundary of all four complementary regions of $w \cup T_m(w)$, the number of complementary regions of $S_{g+2} \setminus (\alpha_k \cup \beta_k)$ is equal to the number of complementary regions of $S_g \setminus (\alpha \cup \beta)$. Therefore $S_{g+2} \setminus (\alpha_k \cup \beta_k)$ is connected, and by the same basic Euler characteristic argument used in the proof of Lemma 2.1, the complement of $\alpha_k \cup \beta_k$ must be simply connected.

Remark 2.6 It is unclear how to attempt such a construction by extending a minimally intersecting filling pair on S_g to one on S_{g+1} . Our construction works by finding a pair of simple closed curves on S_2 intersecting six times, and such that any other




Figure 4: Cutting along $T_m(w) \cup w$ produces four complementary simply connected regions. When glued back up according to the oriented edge labelings, the four green vertices are identified.

simple closed curve intersects their union at least twice. The key point here is that the arcs on $S_{2,1}$ fill in the *arc and curve sense*, meaning that no essential arc or simple closed curve is disjoint from their union.

To extend a minimally intersecting filling pair on S_g to one on S_{g+1} in this fashion, we would need a pair of simple closed curves on S_1 intersecting four times such that any other simple closed curve intersects their union at least twice. One way of guaranteeing this is to choose a pair of curves that are distance 3 in the Farey graph. However, if a curve is distance 3 or more from the $(1, 0)$ curve, it must intersect $(1, 0)$ at least five times; since $SL(2, \mathbb{Z})$ acts transitively on the vertex set of the Farey graph, the same is true for any vertex.

It remains to show that by choosing different intersection points v_j to excise and glue a copy of Z along the resulting boundary component, and by choosing different initial filling pairs on S_g and carrying out this construction, we obtain distinct filling permutations.

After gluing a copy of Z to S_g along a minimally intersecting filling pair (α, β) in the above fashion, we obtain a filling pair (α', β') on S_{g+2} which contains a copy of the arcs a, b as subarcs of α' and β' , respectively.

That is, there exists an open subarc a' of α' whose closure has boundary consisting of two intersection points of $\alpha' \cup \beta'$, such that

- (1) it consists of six consecutive arcs $\alpha'_k, \dots, \alpha'_{k+5}$ (since a' is open, we define α'_k, α'_{k+5} to be open, and the other four arcs to be closed),
- (2) there is a subarc b' of β' , composed of six consecutive β' -arcs, such that $|a' \cap b'| = 5$, and the combinatorics of these intersections completely coincide with those displayed between the arcs a and b in Figure 3.

That is to say, the j^{th} intersection point counted along a' coincides with the k^{th} intersection point along b' if and only if this is the case for the arcs a and b on Z . We call b' the *companion arc* to a' . The union $a' \cup b'$ is called a Z -piece. When convenient, we will alternate between thinking of a Z -piece as being a surface with boundary, or as being contained within a particular filling pair by identifying it with the arcs $a' \cup b'$.

Remark 2.7 Given a Z -piece within a filling pair (α, β) , with α and β arcs labeled x_1, x_2, \dots, x_6 and y_1, \dots, y_6 , respectively, as in Figure 3, note that none of the following four points can coincide on the surface:

- (1) The initial point of x_1 .
- (2) The initial point of y_1 .
- (3) The terminal point of y_6 .
- (4) The terminal point of x_6 .

This observation is trivial, but it will play an important role in the proof of Lemma 2.9.

Note that there exists a filling pair $(\alpha(1), \beta(1))$ on S_1 intersecting $2g - 1 = 1$ times, and at the end of this section we will demonstrate a filling pair $(\alpha(4), \beta(4))$ on S_4 intersecting seven times. If (α, β) on S_g is a minimally intersecting filling pair obtained from one of these two “seed” pairs by successively gluing on copies of Z , whether (α, β) is obtained from $(\alpha(1), \beta(1))$ or $(\alpha(4), \beta(4))$ in this fashion is completely determined by the parity of g . In the case that $g > 4$ is even, we call (α, β) a $(g - 4)/2$ -child of the ancestor $(\alpha(4), \beta(4))$, and similarly if $g \geq 3$ is odd.

The purpose of the next lemma is to show that Z -pieces are symmetric with respect to interchanging the roles of the arcs a and b in Figure 3:

Lemma 2.8 *There exists a homeomorphism of a Z -piece interchanging the arcs a and b and preserving the labeling, in the sense that subarc x_k along a , $1 \leq k \leq 6$, maps to subarc y_k along b .*

Proof Consider the polygonal decomposition in Figure 4; if, for each k , one relabels x_k as y'_k and y_k as x'_k , then the concatenation $x'_1 * x'_2 * \dots * x'_6$ displays the exact intersection combinatorics with the arc $y'_1 * y_2 * \dots * y'_6$ as the oriented arc a does with b . □

Given Z_1 a Z -piece, define $Z_1^{(\alpha)}$ to be the interior arcs of α contained in Z_1 ; that is, $Z_1^{(\alpha)}$ consists of the subarcs $\{x_2, x_3, x_4, x_5\}$. Define $Z_1^{(\beta)}$ similarly.

If Z_1, Z_2 are two Z -pieces in (α, β) , let $(Z_1, Z_2)^\cap$ denote the set of arcs

$$(Z_1, Z_2)^\cap = (Z_1^{(\alpha)} \cap Z_2^{(\alpha)}) \cup (Z_1^{(\beta)} \cap Z_2^{(\beta)}).$$

For $1 \leq k \leq 6$, let $x_k(Z_1)$ denote the terminal point of the subarc x_k in a in the Z -piece Z_1 , and define $y_k(Z_1)$ similarly.

The purpose of the next lemma is to show that distinct Z -pieces cannot share any α -arcs.

Lemma 2.9 *Let $Z_1, Z_2 \subset \alpha \cup \beta$ be two Z -pieces on a k -child (α, β) for any $k \in \mathbb{N}$. Then*

$$Z_1^{(\alpha)} \cap Z_2^{(\alpha)} \neq \emptyset \implies Z_1 = Z_2.$$

Proof Without loss of generality, Z_1 starts before Z_2 along α and terminates before Z_2 along α . We claim that it suffices to show

$$Z_1 \cap \alpha = Z_2 \cap \alpha;$$

indeed, the arcs of α contained in a Z -piece determine uniquely the arcs of β composing the companion arc, and therefore

$$Z_1 \cap \alpha = Z_2 \cap \alpha \implies Z_1 \cap \beta = Z_2 \cap \beta \implies Z_1 = Z_2.$$

Identify Z_2 with the Z -piece pictured in Figure 3; that is, the α -arcs within Z_2 are labeled x_1, x_2, \dots, x_6 , and the β -arcs are labeled y_1, \dots, y_6 in that order, respectively.

Case 1: $x_6(Z_1) = x_2(Z_2)$ Note that, in Figure 3, the concatenation of arcs $y_5 * x_2^{-1}$ forms a loop.

Therefore, if $x_6(Z_1) = x_2(Z_2)$, the concatenation $y_3 * x_4$ must also be a loop. By inspection, it is not.

Case 2: $x_6(Z_1) = x_3(Z_2)$ As seen in Figure 3, the concatenation $y_3 * x_4 * x_5$ is a loop.

Then, if $x_6(Z_1) = x_3(Z_2)$, the concatenation of y_6 with the next two subarcs of α also forms a loop, but this would imply that $y_6(Z_2)$ coincides with the initial point of x_1 , contradicting Remark 2.7.

Case 3: $x_6(Z_1) = x_4(Z_2)$ Note that the concatenation of arcs $y_2^{-1} * x_5$ forms a loop.

Then, as in the previous two cases, if $x_6(Z_1) = x_4(Z_2)$, the concatenation $y_3^{-1} * x_6$ must be a loop, but as seen in Figure 3, it is not.

Case 4: $x_6(Z_1) = x_5(Z_2)$ The concatenation C of y_1^{-1} with the forward direction of the next α subarc cannot be a loop, by Remark 2.7. However, $y_2^{-1} * x_5$ is a loop, and therefore $x_6(Z_1) = x_5(Z_2)$ implies that C is a loop.

Therefore, if $Z_1^{(\alpha)} \cap Z_2^{(\alpha)} \neq \emptyset$, $Z_1^{(\alpha)} = Z_2^{(\alpha)}$, and thus $Z_1 = Z_2$. □

Since, by Lemma 2.8, Z -pieces are symmetric with respect to the arcs a and b , Lemma 2.9 immediately implies the following corollary:

Corollary 2.10 *Let $Z_1, Z_2 \subset \alpha \cup \beta$ be two Z -pieces on a k -child (α, β) for any $k \in \mathbb{N}$. Then*

$$Z_1^{(\beta)} \cap Z_2^{(\beta)} \neq \emptyset \implies Z_1 = Z_2.$$

Proof If, on the filling pair (α, β) , Z_1 and Z_2 share β -arcs, then by Lemma 2.8, on the pair $(\alpha', \beta') := (\beta, \alpha)$, Z_1 and Z_2 share α' -arcs; apply Lemma 2.9. □

Putting Lemma 2.9 and Corollary 2.10 together, we obtain the following important fact, which states that distinct Z -pieces cannot overlap:

Lemma 2.11 *Let Z_1, Z_2 be Z -pieces on a k -child (α, β) for any $k \in \mathbb{N}$. Then*

$$(Z_1, Z_2)^\cap \neq \emptyset \implies Z_1 = Z_2.$$

Proof Immediate from Lemma 2.9, Corollary 2.10 and the definition of $(Z_1, Z_2)^\cap$. □

Henceforth, we assume that $g \geq 3$ is odd. The even case is completely analogous, and will therefore follow from the construction of a minimally intersecting filling pair on S_4 at the end of this section.

Given an odd integer $g \geq 3$, we define the set L_g of positive integer sequences by

$$L_g := \{(s_1, s_2, \dots, s_{(g-1)/2}) : 1 = s_1 < s_2 < \dots < s_{(g-1)/2}, s_i \leq 4i - 3\}.$$

Let W_g denote the set of all oriented minimally intersecting filling pairs on S_g .

Lemma 2.12 $|W_g| \geq |L_g|$

Proof Fix g and a positive sequence $(s) = (s_1, \dots, s_{(g-1)/2})$ in L_g . We show the existence of an injection $f = f_g: L_g \hookrightarrow W_g$. Specifically, we construct a corresponding oriented minimally intersecting filling pair in a unique way.

Start with the filling pair $(\alpha(1), \beta(1))$ on S_1 , and attach a Z -piece along the only intersection point to obtain an oriented minimally intersecting filling pair $(\alpha(3), \beta(3))$ on S_3 . As in the original construction at the beginning of this section, label the $2g - 1 = 5$ intersection points v_1, v_2, \dots, v_5 by choosing an initial intersection point and counting with respect to the orientation along $\alpha(3)$. Note that $s_2 \leq 5$; then attach a Z -piece by excising a small disk centered around v_{s_2} and gluing to obtain an oriented pair $(\alpha(5), \beta(5))$ on S_5 .

Relabel the $2(5) - 1 = 9$ intersection points as follows: any of the original intersection points v_k coming from $\alpha(3) \cup \beta(3)$ with $k < s_2$ are labeled v_k . Intersection points v_j coming from $\alpha(3) \cup \beta(3)$ with $j > s_2$ are relabeled as v_{j+4} . The five intersection points coming from the Z -piece are labeled, in order along the arc a , $v_{s_2}, v_{s_2+1}, v_{s_2+2}, v_{s_2+3}, v_{s_2+4}$. This produces a labeling of all nine intersection points which respects the orientation of $\alpha(5)$.

We then glue on an additional Z -piece by excising the intersection point v_{s_3} , obtaining the oriented pair $(\alpha(7), \beta(7))$ on S_7 , and we relabel the now 13 intersection points as above.

We repeat this process for each $i \leq \frac{1}{2}(g - 1)$, each time excising the intersection point labeled v_{s_i} and gluing on another Z -piece. This process terminates with an oriented minimally intersecting filling pair in W_g , which we define to be $f((s))$.

Remark 2.13 It is unclear how many Z -pieces the filling pair $f((s))$ contains. However, it necessarily contains at least one Z -piece: the one associated to the last index $s_{(g-1)/2}$.

It remains to show that f is injective. Assume $f((s)) = f((s'))$. Then we claim that

$$s_{(g-1)/2} = s'_{(g-1)/2}.$$

Suppose not; then by Remark 2.13, the filling pair $f((\alpha)) = f((\alpha'))$ contains two distinct Z -pieces Z, Z' , associated to $s_{(g-1)/2}, s'_{(g-1)/2}$, respectively. Without loss of generality, $s_{(g-1)/2} < s'_{(g-1)/2}$. By Lemma 2.11, Z and Z' cannot overlap.

Therefore, if we remove Z by cutting along the separating curve in S_g associated to its boundary and gluing in a disk, we obtain an oriented minimally intersecting

filling pair on S_{g-2} , which still contains Z' . Then since (α) is a monotonically increasing sequence, by the same argument the filling pair obtained by successively cutting along the Z -pieces associated to smaller and smaller indices of (α) will still contain Z' . Eventually, we conclude that the filling pair $(\alpha(1), \beta(1))$ on S_1 contains Z' , a contradiction. Hence $s_{(g-1)/2} = s'_{(g-1)/2}$.

Applying the exact same argument inductively, we have $s_k = s'_k$ for each $k < \frac{1}{2}(g-1)$ as well. Hence f is injective. □

Corollary 2.14 For g odd,

$$N(g) > \frac{\prod_{k=1}^{\frac{g-1}{2}-1} 3k + 1}{4 \cdot (2g - 1)^2 \cdot (\frac{g-1}{2} - 1)!}$$

Proof By Lemma 2.12, it suffices to show that

$$(2-1) \quad |L_g| > \frac{\prod_{k=1}^{\frac{g-1}{2}-1} 3k + 1}{(\frac{g-1}{2} - 1)!}$$

The lower bound on $N(g)$ follows by dividing by an upper bound on the number of twisting permutations.

Moreover, we bound $|L_g|$ below by counting the number of length $\frac{1}{2}(g-1) - 1$ integer sequences $(s)_{i \geq 2}$ such that no two entries are the same, and $1 < \alpha_i \leq 4i - 3$. Simple calculation then leads to (2-1). □

The following lemma then implies the exponential growth of $N(g)$:

Lemma 2.15
$$\frac{\prod_{k=1}^{\frac{g-1}{2}-1} 3k + 1}{4 \cdot (2g - 1)^2 \cdot (\frac{g-1}{2} - 1)!} \sim \frac{3^{g/2}}{g^2}$$

Proof Begin by factoring out a copy of $3^{\frac{g-1}{2}-1}$ from the numerator. What remains after pulling out this exponential term is the product

$$(1 + \frac{1}{3})(2 + \frac{1}{3}) \cdots (\frac{g-1}{2} - 1 + \frac{1}{3}),$$

which, after expanding, we interpret as a polynomial in $\frac{1}{3}$, the constant term being $(\frac{1}{2}[g-1] - 1)!$, which will cancel with the factorial in the denominator. All of the other coefficients will be dominated by the factorial term in the denominator, and therefore all go to 0 as $g \rightarrow \infty$. □

As mentioned above, the proof of the lower bound for g even, $g > 2$, is exactly analogous. Hence, to complete the proof of the lower bound, it suffices to construct a minimally intersecting filling pair on S_4 ; equivalently, we demonstrate the existence of a filling permutation in Σ_{28} :

$$(1, 10, 7, 18, 5, 24, 3, 4, 9, 20, 21, 12, 13, 28, 25, 8, 15, 2, 17, 22, 23, 6, 11, 26, 19, 16, 27, 14) \in \Sigma_{28}.$$

Recall that the proof of Lemma 2.12 relied on the assumption that there are no Z -pieces in this filling pair. In Section 2.4 below, we show that there are no filling pairs on S_2 whose complement is connected. Thus, if there was a Z -piece within the filling pair on S_4 associated to the permutation above in Σ_{28} , we could excise a copy of Z from S_4 to obtain a filling pair on S_2 whose complement is connected, a contradiction.

2.3 Proof of the upper bound

The upper bound from Theorem 1.1 is obtained by bounding the number of filling permutations from above.

Lemma 2.16 *Let $\sigma \in \Sigma_{8g-4}$ be a filling permutation. Then $\sigma = Q^{4g-2}C$, where C is a square root of $Q^{4g-2}\tau$. Conversely, if C is a square root of $Q^{4g-2}\tau$ such that $\sigma = Q^{4g-2}C$ is a parity-respecting $(8g - 4)$ -cycle, then σ is a filling permutation.*

Proof Suppose σ is a filling permutation. Then

$$(Q^{4g-2}\sigma)^2 = Q^{4g-2}(\sigma Q^{4g-2}\sigma) = Q^{4g-2}\tau.$$

Conversely, if C is a square root of $Q^{4g-2}\tau$,

$$(Q^{4g-2}C)Q^{4g-2}(Q^{4g-2}C) = Q^{4g-2}C^2 = \tau. \quad \square$$

Therefore, $N(g)$ is bounded above by the number of square roots of $Q^{4g-2}\tau$; the upper bound of Theorem 1.1 is obtained by identifying a collection of such square roots C such that $Q^{4g-2}C$ is not an $(8g - 4)$ -cycle, and subtracting this from the total number of square roots.

Note that $Q^{4g-2}\tau$ is a composition of disjoint transpositions:

$$Q^{4g-2}\tau = (1, 4g + 1)(2, 4g + 2)(3, 4g + 3) \cdots (4g - 4, 8g - 4)(4g - 3, 4g - 1)(4g - 2, 4g).$$

The square roots of $Q^{4g-2}\tau$ are obtained by partitioning these transpositions into pairs and “interleaving” the pairs. For instance, one such root is obtained by pairing

off $(1, 4g + 1)$ with $(2, 4g + 2)$, and in general for each $i \leq 4g - 3$ pairing off the i^{th} transposition above with the $(i + 1)^{\text{st}}$ and converting each such pair into a single 4-cycle by interleaving the four numbers:

$$(1, 2, 4g + 1, 4g + 2)(3, 4, 4g + 3, 4g + 4) \cdots (4g - 3, 4g - 2, 4g - 1, 4g).$$

Note that for a particular pair of transpositions, there are two 4-cycles we can obtain by interleaving. Furthermore, because Q^{4g-2} is parity-respecting and sends even numbers to even numbers, the only way $Q^{4g-2}C$ can be a parity-respecting $(8g - 4)$ -cycle is if C sends even numbers to odd ones. Therefore there are exactly $2^{2g-1}(2g - 1)!$ square roots C of $Q^{4g-2}\tau$ such that $Q^{4g-2}C$ is parity-respecting.

We remark that the fact that $2^{2g-1}(2g - 1)!$ is an upper bound for $N(g)$ is immediate without the help of Lemma 2.16 (label the intersection points from 1 to $2g - 1$ according to their order along α and comparing this labeling to an analogous one obtained by ordering along β produces a permutation in Σ_{2g-1} ; specifying the sign of each intersection accounts for the factor of 2^{2g-1}).

However, using Lemma 2.16 and the structure of the square roots of $Q^{4g-2}\tau$, we can obtain strictly smaller lower bounds. In particular, consider the following $2^{2g-2}(2g - 1)(2g - 3)!$ permutations:

Pair the transposition $(1, 4g + 1)$ with an arbitrary transposition (k, j) of even numbers in the disjoint cycle notation for $Q^{4g-2}\tau$; there are $2 \cdot (2g - 1)$ ways of doing this. If C is any square root of $Q^{4g-2}\tau$ with $(1, k, 4g + 1, j)$ as one of its 4-cycles, then $Q^{4g-2}C$ maps 1 to $(k + 4g - 2) \pmod{8g - 4}$ (or $8g - 4$ if $k = 4g - 2$), the number representing the inverse direction of the arc corresponding to k . Note that, by inspection of $Q^{4g-2}\tau$,

$$(k + 4g - 2) \not\equiv j \pmod{8g - 4},$$

and therefore we are free to pair the transposition containing $(k + 4g - 2) \pmod{8g - 4}$ with the transposition containing $4g - 1$, and we choose the interleaving of these transpositions so that the resulting 4-cycle sends $(k + 4g - 2) \pmod{8g - 4}$ to $4g - 1$.

Then we can obtain a parity-respecting square root C by pairing off the remaining transpositions in any way, evens with odds. However, no square root C obtained in this fashion has the property that $Q^{4g-2}C$ is an $(8g - 4)$ -cycle, because, by construction,

$$(Q^{4g-2}C)^2(1) = 1.$$

Therefore, the number of admissible square roots is bounded above by

$$2^{2g-1}(2g - 1)! - 2^{2g-2}(2g - 1)(2g - 3)! = 2^{2g-2} \frac{(4g - 5)(2g - 1)!}{2g - 2}.$$

Note that the number of twisting permutations is bounded below by $2g - 1$, since κ_g has order $2g - 1$. This completes the proof of the upper bound of Theorem 1.1. \square

2.4 Genus 2

In this section, we show that there are no filling pairs intersecting three times on S_2 . For a filling pair with four intersections on S_2 , see Farb and Margalit [7, page 41].

Theorem 2.17 *Suppose (α, β) is a filling pair on S_2 . Then $i(\alpha, \beta) \geq 4$.*

Remark 2.18 Dan Margalit has pointed out the following alternative argument for a proof of Theorem 2.17:

Suppose (α, β) is a pair of simple closed curves on S_2 such that $i(\alpha, \beta) = 3$. Then consider the hyperelliptic involution $\pi: S_2 \rightarrow S_{0,6}$, which is a double covering map branched over the six punctures. It follows that $\pi(\alpha), \pi(\beta)$ must be a pair of simple arcs, intersecting once in their interior and sharing an endpoint (each interior intersection accounts for two intersections in the cover).

Up to homeomorphism, there are only a handful of such configurations, and in each case, one can find an arc on $S_{0,6}$ disjoint from the union $\pi(\alpha) \cup \pi(\beta)$ which lifts to a simple closed curve disjoint from $\alpha \cup \beta$. Therefore (α, β) cannot be a filling pair on S_2 .

Proof Suppose $i(\alpha, \beta) = 3$. We can schematically represent $\alpha \cup \beta$ as a partially directed graph $G = (V, E)$ with six vertices, three undirected edges and six directed edges as follows. To construct G , begin with two disjoint 3-cycles, one of which is oriented clockwise and the other counterclockwise. Label the vertices on each cycle from 1 to 3 in cyclic order. If α is oriented on the surface by decorating α with an arrow, then one of the cycles of G represents α as viewed from its left, and the other cycle represents α as viewed from its right. The three intersection points between α and β correspond to the three vertices on each cycle.

The remaining three undirected edges of G correspond to the three arcs of β . If (α, β) fill, their complement is a single disk, and this corresponds to the existence of a certain kind of Eulerian cycle. Concretely, there is a cycle c on G beginning with the oriented edge on the clockwise cycle whose initial point is vertex 1, satisfying the following:

- (1) The edges along c alternate between being directed and undirected.
- (2) Every directed edge appears along c exactly once.
- (3) Every undirected edge appears along c exactly once in each direction.

As demonstrated below in Figure 5, no undirected edge can connect vertices on the same 3-cycle, or else α and β cannot be in minimal position.

Therefore, without loss of generality we can assume that the first undirected edge connects vertex 1 on the clockwise cycle to vertex 3 on the counterclockwise cycle. This will then determine G , because vertex 3 on the clockwise cycle cannot then be connected to vertex 1 on the counterclockwise cycle, or else β will close up without ever passing through vertex 2.

Therefore, G must be as pictured in Figure 6. Note then that G cannot satisfy the Eulerian condition described above; starting at vertex 1 and traveling along the clockwise cycle, one arrives at vertex 2. Then the undirected edge incident to vertex 2 on the clockwise cycle is connected to vertex 1 on the counterclockwise cycle at its other end. Then we must travel counterclockwise to vertex 3 on the counterclockwise cycle, and finally the undirected edge at this vertex connects back to vertex 1 on the clockwise cycle. Therefore, no single cycle can satisfy the three necessary requirements. \square




Figure 5: If one of the undirected edges connects vertices on the same cycle, there must be at least two undirected edges with this property. This will imply the existence of two bigons labeled D_1, D_2 above in the complement of $\alpha \cup \beta$.




Figure 6: When undirected edges connect vertices on opposite cycles, G cannot satisfy the Eulerian condition described above.

3 Minimally intersecting filling pairs as combinatorial optimizers

In this section, we consider a potential way of measuring the combinatorial efficiency of a filling pair. Namely, given a filling pair (α, β) , let $T_k(\alpha, \beta) \in \mathbb{N}$ be the number of simple closed curves (up to isotopy) which intersect the union $\alpha \cup \beta$ no more than k times. We expect an “optimal” filling pair to have large values of T_k for small values of k ; in other words, there should be many simple closed curves which only intersect the union a small number of times.

Specifically, we prove Theorem 1.2, which shows that minimally intersecting filling pairs optimize the value of T_1 :

Theorem 1.2 *Let (α, β) be a filling pair on S_g , $g > 2$ and define $T_1(\alpha, \beta) \in \mathbb{N}$ to be the number of simple closed curves intersecting $\alpha \cup \beta$ only once. Then $T_1(\alpha, \beta) \leq 4g - 2$, with equality if and only if (α, β) is minimally intersecting.*

Proof Suppose first that (α, β) is a minimally intersecting filling pair. Then

$$T_1(\alpha, \beta) = 4g - 2,$$

because for each of the $2g - 1$ subarcs α_k of α (resp. β), there is a single simple closed curve which intersects α_k and no other subarc. This simple closed curve corresponds to the arc in the $(8g - 4)$ -gon obtained by cutting along $\alpha \cup \beta$ which connects the two edges which project down to the subarc α_k .

Now suppose that (α, β) do not intersect minimally, so that $S_g \setminus (\alpha \cup \beta)$ has multiple connected components. The number of components is equal to

$$i(\alpha, \beta) - 2g + 2;$$

thus if we cut along $\alpha \cup \beta$ we obtain a disjoint union of $i(\alpha, \beta) - 2g + 2$ even-sided polygons $P_1, \dots, P_{i(\alpha, \beta) - 2g + 2}$, such that the total number of edges over all polygons is $4 \cdot i(\alpha, \beta)$. $T_1(\alpha, \beta)$ is then equal to the number of pairs of edges which project to the same arc in S_g and which belong to the same polygon. For each j , there must be at least one edge of P_j which has the property that the inverse edge is not on P_j (or else P_j will be disconnected from all of the other polygons after glued up). Thus, for each j there must be at least two edges of P_j whose inverse edges belong to other polygons, since P_j is even-sided for each j .

Suppose there exists $j \in \{1, \dots, i(\alpha, \beta) - 2g + 2\}$ such that there are exactly two edges of P_j whose inverse edge is not on P_j . Then one must project to a subarc α_k of α ,

and the other to a subarc β_h of β . Consider the arc β_l which immediately precedes α_k along P_j , and assume $\beta_l \neq \beta_h$ (otherwise, the arc immediately preceding β_h cannot be α_k , so in this case we just switch the roles of β_h and α_k).

By assumption, (α, β) is not minimally intersecting, but there is still a map M acting on the edges of $P_1 \sqcup \cdots \sqcup P_{i(\alpha, \beta) - 2g + 2}$ playing the same role as the permutation C discussed in Section 2.3 and in the proof of Lemma 2.2 in Section 2.1. Namely, send an edge e to the edge e' immediately following it in clockwise order along some P_i , and then send e' to the edge which projects down to the same arc as e' does after being glued up (ie the inverse edge for e').

This map M has order 4, because it corresponds to a “rotation” by $\pi/2$ about an intersection point. Note that $M(\beta_l) = \alpha_k^{-1}$, which by assumption is not on P_j . Then $M(\alpha_k^{-1}) = \beta_{l+1}^{-1}$ (by β_{l+1} , we simply mean the subarc immediately following β_l in the direction of β_l ; see Figure 7).




Figure 7: Schematic of an intersection point of $\alpha \cup \beta$. Arrows indicate the action of M on each oriented edge.

We claim that $\beta_{l+1} = \beta_h^{-1}$, for, if not, then $M(\alpha_k^{-1}) = \beta_{l+1}^{-1}$ is not on P_j . β_{l+1}^{-1} is immediately followed by α_{k-1}^{-1} . Thus $M(\beta_{l+1}^{-1}) = \alpha_{k-1}$, which is also not on P_j , because otherwise both α_k and α_{k-1} would have the property that their inverse edges are not on P_j . However, M must have order 4, which implies that β_l^{-1} immediately follows α_{k-1} (see Figure 7) and that, in particular, the inverse of β_l is not on P_j , which contradicts our choice of β_l .

Therefore, $\beta_{l+1} = \beta_h^{-1}$. This in particular implies that α_k^{-1} immediately precedes β_h^{-1} on P_m for some m . We claim that β_h must immediately follow α_k on P_j (see Figure 8). If not, then let β_q denote the edge immediately preceding α_k^{-1} on P_m . Then $M(\beta_q) = \alpha_k$ on P_j , and $M(\alpha_k) = \beta_r$ for some r , also on P_j . Thus $M(\beta_r) = \alpha_{k+1}^{-1}$, also on P_j , and since M has order 4, β_q^{-1} immediately follows α_{k+1}^{-1} on P_j , and therefore $\beta_q^{-1} = \beta_h$. But this is a contradiction, because β_q and β_h^{-1} represent distinct edges on P_m ; see Figure 8.




Figure 8: Using the fact that M has order 4, we conclude that $\beta_q^{-1} = \beta_h$, but this is impossible because β_q and β_h^{-1} are distinct edges of P_m by assumption.

This implies that a sufficiently small neighborhood in S_g around the vertex on P_j separating α_k and β_h contains only two edges of $\alpha \cup \beta$. This is a contradiction, because $\alpha \cup \beta$ is 4-valent. Since j was arbitrary, we conclude that each polygon $P_1, \dots, P_{i(\alpha, \beta) - 2g + 2}$ must have at least four edges whose inverses are not on the same polygon. This leaves at most

$$4i(\alpha, \beta) - 4(i(\alpha, \beta) - 2g + 2) = 8g - 8$$

edges whose inverse can potentially be on the same polygon, and therefore

$$T_1(\alpha, \beta) \leq 4g - 4. \quad \square$$

It would be interesting to analyze the statistics of $T_k(\alpha, \beta)$ as k goes to ∞ . We pose the following question:

Question 1 What is the minimal k such that T_k is not maximized by a minimally intersecting filling pair?

4 Minimal length filling pairs on hyperbolic surfaces

In this section, we prove Theorem 1.3. In all that follows, let l_h denote length in the hyperbolic plane \mathbb{H}^2 . As in the introduction, define \mathcal{M}_g to be the *moduli space* of S_g , the space of all complete hyperbolic metrics on S_g up to isometry. Then $\mathcal{F}_g : \mathcal{M}(S_g) \rightarrow \mathbb{R}$ is the function which outputs the length of the shortest minimally intersecting filling pair on a hyperbolic surface σ , where the length of a filling pair is the sum of the two individual lengths of the unique geodesic representatives of the simple closed curves in the pair.

In recalling the definition of a topological Morse function, we follow [13]:

Let $f: \mathcal{M}_g \rightarrow \mathbb{R}$ be a continuous function. Then $x \in \mathcal{M}_g$ is called an *ordinary point* of f if there exists an open neighborhood U of x , a chart $\phi: U \rightarrow \mathbb{R}^{6g-6}$ inducing real-valued coordinates y_1, \dots, y_{6g-6} , and an integer j , $1 \leq j \leq 6g - 6$, such that for any $z \in U$,

$$f(z) = y_j.$$

Otherwise, x is a *critical point*. If x is a critical point, x is called a *non-degenerate critical point* of f if there exists an open neighborhood U of x , a chart $\phi: U \rightarrow \mathbb{R}^{6g-6}$ inducing real-valued coordinates y_1, \dots, y_{6g-6} , and an integer j , $1 \leq j \leq 6g - 6$ such that for any $z \in U$,

$$f(z) - f(x) = \sum_{i=1}^j y_i^2 - \sum_{i=j+1}^{6g-6} y_i^2.$$

We call j the *index* of the non-degenerate critical point x . For our purposes, a *topological Morse function* is a real-valued continuous function f on \mathcal{M}_g such that f has only finitely many critical points, all of which are non-degenerate.

Theorem 1.3 says that \mathcal{F}_g is a topological Morse function, and that the number of critical points of index 0 (ie the global minima) grow at least exponentially in genus:

Theorem 1.3 *Let $g \geq 3$. \mathcal{F}_g is proper and a topological Morse function. For any $\sigma \in \mathcal{M}_g$,*

$$\mathcal{F}_g(\sigma) \geq \frac{1}{2}m_g,$$

where

$$m_g = (8g - 4) \cdot \cosh^{-1} \left(2 \left[\cos \left(\frac{2\pi}{8g-4} \right) + \frac{1}{2} \right] \right)$$

denotes the perimeter of a regular, right-angled $(8g-4)$ -gon. Furthermore, define

$$\mathcal{B}_g := \{ \sigma \in \mathcal{M}_g : \mathcal{F}_g(\sigma) = \frac{1}{2}m_g \}.$$

Then \mathcal{B}_g is finite and grows exponentially in g . If $\sigma \in \mathcal{B}_g$, the injectivity radius of σ is at least $\frac{1}{2} \cosh^{-1}(9/\sqrt{73})$.

Proof Let $C \subset \mathbb{R}$ be a compact subset. Then we claim that there exists $\epsilon > 0$ such that $\mathcal{F}_g^{-1}(C)$ is contained in \mathcal{M}_g^ϵ , the ϵ -thick part of the moduli space, defined to be the subset of \mathcal{M}_g consisting of hyperbolic surfaces with injectivity radius at least ϵ . To see this, note that this is implied by the statement $\mathcal{F}_g \rightarrow \infty$ as the injectivity radius approaches 0.

Lemma 4.1 *For $k \in \mathbb{N}$, there exists $\epsilon = \epsilon(k)$ such that any hyperbolic surface σ with injectivity radius $< \epsilon$ satisfies $\mathcal{F}_g(\sigma) > k$.*

Proof By the collar lemma (see [5]), there exists a function $c: \mathbb{R} \rightarrow \mathbb{R}$ with

$$\lim_{x \rightarrow 0} c(x) = \infty$$

such that if σ is a hyperbolic surface with a closed geodesic γ of length $< 2\epsilon$, there exists an embedded tubular neighborhood around γ whose width is at least $c(\epsilon)$. Since any filling pair must intersect γ at least once, the result follows. \square

Therefore $\mathcal{F}_g^{-1}(C)$ is contained in the thick part of \mathcal{M}_g , which is compact by a result of Mumford [10]; then continuity of \mathcal{F}_g^{-1} implies properness.

To show that \mathcal{F}_g is a topological Morse function, we use the following criterion, due to Akrouit [1]:

Let $(l_u)_{u \in \mathcal{C}}$ be a collection of real-valued smooth functions defined over \mathcal{M}_g and indexed by some countable set \mathcal{C} . Given $x \in \mathcal{M}_g$ and $\lambda \in \mathbb{R}$, define

$$\mathcal{C}_x^{\leq \lambda} := \{u \in \mathcal{C} : l_u(x) \leq \lambda\}.$$

Then the function $\rho := \inf_{u \in \mathcal{C}} l_u$ is called a *generalized systole function* if, for each $x \in \mathcal{M}_g$ and each $\lambda \in \mathbb{R}$, there exists a neighborhood W of x such that $\bigcup_{q \in W} \mathcal{C}_q^{\leq \lambda}$ is a finite subset of \mathcal{C} .

Akrouit shows that any generalized systole function is a topological Morse function. Define \mathcal{C} to be the collection of minimally intersecting filling pairs (not taken up to homeomorphism), and for each such pair $u \in \mathcal{C}$, let $l_u(\sigma)$ denote the length of the filling pair u on the hyperbolic surface σ . Then $\mathcal{F}_g = \inf_{u \in \mathcal{C}} l_u$.

Fix $x \in \mathcal{M}_g$ and $\lambda > 0$; then by the collar lemma, only finitely many pairs in \mathcal{C} can have length less than λ on x . Let s be one of the shortest simple closed geodesics on x , and let $l_s(x)$ denote the length of s on x . Let U be a sufficiently small neighborhood around x so that for any $y \in U$, $|l_s(y) - l_s(x)| < \epsilon$, for some small ϵ .

Thus, any hyperbolic surface in U has the property that s has length at most $l_s(x) + \epsilon$, so by another application of the collar lemma, there can only be finitely many filling pairs having the property that there exists a point in U on which the pair has length at most λ . Therefore \mathcal{F}_g is a topological Morse function.

Next, we identify the minima \mathcal{B}_g of \mathcal{F}_g . Recall that, given a minimally intersecting filling pair (α, β) on a hyperbolic surface σ , cutting along $\alpha \cup \beta$ produces a single hyperbolic $(8g-4)$ -gon P_σ .

Fix $n \in \mathbb{N}$, $n \geq 3$, and let $\lambda > 0$ be such that there exists a hyperbolic n -gon with area λ (that is, λ is less than the area of an ideal n -gon). Then by a result of Bezdek [4],

the hyperbolic n -gon enclosing an area of λ with the smallest perimeter is the regular n -gon with area λ . Therefore, the perimeter of P is at least as large as the perimeter of a regular $(8g-4)$ -gon enclosing an area of $2\pi(2g-2)$ (the area of σ , by Gauss-Bonnet). The regular $(8g-4)$ -gon with area $2\pi(2g-2)$ is right-angled. The perimeter of such a polygon is m_g , and its perimeter is twice the length of the filling pair.

Now we demonstrate a uniform lower bound on the injectivity radius for any $\sigma \in \mathcal{B}_g$. We identify a simple closed geodesic γ on σ with its lift $\tilde{\gamma}$ in the right-angled regular $(8g-4)$ -gon P_σ that one obtains by cutting along the minimally intersecting filling pair (α, β) of minimal length $\frac{1}{2}m_g$ on σ ; $\tilde{\gamma}$ is a disjoint collection of geodesic arcs connecting edges of P_σ .

Assume first that every arc in $\tilde{\gamma}$ connects adjacent edges of P_σ . Since γ, α, β are all geodesics, they are in pairwise minimal position, and therefore no two of them form a bigon on the surface.

If $\tilde{\gamma}$ contains some arc x with one endpoint on some edge a of P_σ belonging to α , and the other endpoint on an adjacent edge b belonging to β , there must be some arc y of $\tilde{\gamma}$ with an endpoint on the edge of P_σ corresponding to b^{-1} . By the previous paragraph, the other endpoint of y cannot be located on the edge c belonging to α immediately following a (see Figure 9).




Figure 9: If the right endpoint of y is on c , γ and α cannot be in minimal position.

But then the arc of γ obtained by concatenating x and y is longer than the edge b , because P_σ is right-angled. All edges of P_σ have length

$$\frac{m_g}{8g-4} = \cosh^{-1}\left(2\left[\cos\left(\frac{2\pi}{8g-4}\right) + \frac{1}{2}\right]\right).$$

For $g \geq 3$, this is an increasing function of g , and therefore

$$\frac{m_g}{8g-4} > \frac{m_3}{20} = \cosh^{-1}\left(2\left[\frac{1}{2} + \sqrt{\frac{5}{8} + \frac{\sqrt{5}}{8}}\right]\right) > \cosh^{-1}\frac{9}{\sqrt{73}}.$$

Therefore, we can assume that there exists an arc r in $\tilde{\gamma}$ whose endpoints are not on adjacent edges of P_σ . Let e be an edge of P_σ containing an endpoint of r . Let u, v




Figure 10: r must cross the geodesic connecting u to v .

denote the endpoints of the two edges adjacent to e , which are not shared endpoints with e , and let x denote the geodesic segment in P_σ orthogonal to both e and the geodesic connecting u to v (see Figure 10).

We claim that the geodesic segment g connecting u to v must separate e and its two neighboring edges from the rest of P_σ . Indeed, if any part of ∂P_σ crosses g , then g is not contained in P_σ , which contradicts the fact that P_σ is convex. Thus, since by assumption the other endpoint of r is not on an edge adjacent to e , r must cross g , and therefore its length is at least the length of x .

Computing the length of x , which we denote by $\lambda(g)$, is an exercise in elementary hyperbolic geometry; see [5, Theorem 2.3.1]:

$$\lambda(g) = \cosh^{-1} \frac{1 + 2 \cos \frac{\pi}{2-4g}}{\left(4 \cos \frac{\pi}{2-4g} \left(1 + \cos \frac{\pi}{2-4g}\right) + \left(1 + 2 \cos \frac{\pi}{2-4g}\right)^{-2}\right)^{1/2}}.$$

This is a decreasing function of g , and

$$\lim_{g \rightarrow \infty} \lambda(g) = \cosh^{-1} \frac{9}{\sqrt{73}}.$$

This proves that any simple closed geodesic on σ has length at least $\cosh^{-1}(9/\sqrt{73})$, and therefore the injectivity radius is at least half of this quantity.

Finally, we address the growth of $|\mathcal{B}_g|$. Properness of \mathcal{F}_g implies that all minima are contained in a compact subset of \mathcal{M}_g , and the fact that \mathcal{F}_g is a topological Morse function implies that the minima are isolated. Therefore $|\mathcal{B}_g| < \infty$.

It remains to show that \mathcal{B}_g grows exponentially as a function of g . Each $\text{Mod}(S_g)$ -orbit of minimally intersecting filling pairs can be associated to an element of \mathcal{B}_g by simply gluing the edges of a regular right-angled $(8g-4)$ -gon together in accordance with the chosen orbit. However, this association need not be injective; a priori, it may be the case that many minimal-length, minimally intersecting filling pairs occur on the same hyperbolic surface.

Suppose σ is a hyperbolic surface admitting r minimal-length, minimally intersecting filling pairs. Thus there are r immersed right-angled $(8g-4)$ -gons on σ , and developing σ into \mathbb{H}^2 , these polygons lift to r distinct tilings T_1, \dots, T_r of \mathbb{H}^2 by regular, right-angled $(8g-4)$ -gons.

Let Γ denote the holonomy of this developing map; ie Γ is the discrete subgroup of $\text{PSL}_2(\mathbb{R})$ identified with the fundamental group of σ . Recall that the *commensurator* of Γ , denoted $\text{comm}(\Gamma)$, is the subgroup of $\text{Isom}(\mathbb{H}^2)$ consisting of elements g such that $g\Gamma g^{-1}$ is *commensurable* with Γ , which is to say that $g\Gamma g^{-1} \cap \Gamma$ is a finite-index subgroup of Γ and $g\Gamma g^{-1}$. Note that $\Gamma < \text{comm}(\Gamma)$.

Let Λ_i denote the full isometry group of the tiling T_i ; note that $\Gamma < \Lambda_i$ for each i , and Γ acts freely and transitively on the set of tiles. Thus Γ has finite-index in Λ_i , and therefore $\text{comm}(\Gamma) = \text{comm}(\Lambda_i)$.

Lemma 4.2 $r \leq [\text{comm}(\Lambda_1) : \Lambda_1]$

Proof For each $k \in \{1, \dots, r\}$, let $g_k \in \text{Isom}(\mathbb{H}^2)$ be an element sending T_1 to T_k . Then $g_k \Lambda_1 g_k^{-1}$ is the group of isometries of T_k , and by assumption this is also Λ_k ; $\Lambda_k \cap \Lambda_1$ contains Γ , which is finite-index in both, and therefore $g_k \in \text{comm}(\Lambda_1)$.

Suppose for $j \neq k$ that g_k, g_j are in the same coset of Λ_1 in $\text{comm}(\Lambda_1)$. Then $g_j^{-1} g_k$ fixes T_1 . Thus g_j^{-1} sends T_k to T_1 , but by definition it also sends T_j to T_1 , a contradiction unless $T_k = T_j$ or $g_j = g_k$, either of which implies the other, and that $j = k$. □

Lemma 4.3 *For g sufficiently large, $\text{comm}(\Lambda_1)$ is a discrete subgroup of $\text{Isom}(\mathbb{H}^2)$.*

Proof A fundamental polygon for the tiling T_1 can be subdivided into $16g-8$ copies of a hyperbolic triangle with angles $2\pi/(16g-8), \pi/4, \pi/2$. As a consequence, Λ_1 is commensurable with the triangle group of signature $(16g-8, 8, 4)$, the subgroup of $\text{Isom}(\mathbb{H}^2)$ generated by reflections in the sides of one of these triangles, which has the presentation

$$\langle a, b, c \mid a^2 = b^2 = c^2 = (ab)^8 = (bc)^4 = (ca)^{16g-8} = 1 \rangle.$$

Denote this group by Δ_g . By a result of Long, Machlachlan and Reid [8] (also Takeuchi [15]), there are only finitely many arithmetic triangle groups. By a theorem of Margulis [9], a subgroup of $\text{Isom}(\mathbb{H}^2)$ is arithmetic if and only if its commensurator is not discrete. Therefore, for all sufficiently large g , $\text{comm}(\Delta_g)$ is discrete.

Since Λ_1 is commensurable with Δ_g , $\text{comm}(\Lambda_1) = \text{comm}(\Delta_g)$, and is therefore also discrete. □

Using Lemma 4.3, for sufficiently large g , we can assume that the quotient space $\mathbb{H}^2/(\text{comm}(\Lambda_1))$ is a hyperbolic orbifold \mathcal{O}_g , and therefore the area of \mathcal{O}_g is at least $\pi/21$ [7]. \mathcal{O}_g is covered by the orbifold \mathbb{H}^2/Λ_1 , and the degree of the covering map is equal to the index $[\text{comm}(\Lambda_1) : \Lambda_1]$, and is also equal to the ratio of the area of \mathbb{H}^2/Λ_1 to the area of \mathcal{O}_g . The area of a fundamental domain for the action of Λ_1 is bounded above by the area of σ , and therefore

$$r \leq \frac{2\pi(2g - 2)}{\pi/21} = 42(2g - 2).$$

Thus, at most $42(2g - 2)$ minimal-length, minimally intersecting filling pairs can coincide on the same hyperbolic surface, and hence there are at least $N(g)/(42(2g - 2))$ surfaces containing such a pair. Since $N(g)$ grows exponentially by Theorem 1.1, so must the number of such surfaces. □

Remark 4.4 More generally, we can also consider the function $\mathcal{F}_g^{(k)}: \mathcal{M}_g \rightarrow \mathbb{R}$, which, given a hyperbolic metric σ on S_g , outputs the length of the shortest filling pair whose complement has k or fewer complementary regions. In this notation, $\mathcal{F}_g = \mathcal{F}_g^{(1)}$.

We claim that $\mathcal{F}_g^{(2)}$ has the same minima as $\mathcal{F}_g^{(1)}$. To see this, let (α, β) be a filling pair of geodesics on σ intersecting $2g$ times, so that cutting along $\alpha \cup \beta$ produces a pair P_1, P_2 of hyperbolic polygons. Between both polygons, there is a total of $8g$ edges.




Figure 11: A schematic picture of the formation of $P_{1,2}$ from P_1 and P_2 . We lose a total of four edges in the creation of $P_{1,2}$: e, e^{-1} , and a and b concatenate smoothly because they represent adjacent arcs on the same geodesic, and therefore they are replaced by a single edge. By the same argument, so are c and d .

Then choose an edge e on P_1 which glues to an edge e^{-1} on P_2 as in Figure 9. Since e and e^{-1} have the same hyperbolic length, we can glue P_1 and P_2 together along e to produce a new polygon $P_{1,2}$, which must have four fewer sides than the total number of sides in $P_1 \sqcup P_2$, for we lose the two edges e, e^{-1} which lie in the interior of $P_{1,2}$, and the edges a, b concatenate together smoothly in $P_{1,2}$ to form a single long edge since they correspond to adjacent arcs along the same geodesic on σ . The same is true for the edges c, d .

Thus, the total perimeter of $P_1 \sqcup P_2$ is strictly larger than the perimeter of a regular, right-angled $(8g-4)$ -gon, and therefore we've shown:

Corollary 4.5 *For all $\sigma \in \mathcal{M}_g$, $\mathcal{F}_g^{(2)} \geq \frac{1}{2}m_g$, and the same hyperbolic surfaces minimize both $\mathcal{F}_g^{(1)}$ and $\mathcal{F}_g^{(2)}$. Therefore, the entirety of Theorem 1.3 also holds for $\mathcal{F}_g^{(2)}$.*

Motivated by Corollary 4.5, we make the following conjecture:

Conjecture 4.6 *Define $\mathcal{Y}_g: \mathcal{M}_g \rightarrow \mathbb{R}$ to be the “filling pair systole” function, which outputs the length of the shortest filling pair. Then $\mathcal{Y}_g \geq \frac{1}{2}m_g$.*

References

- [1] **H Akrouf**, *Singularités topologiques des systoles généralisées*, *Topology* 42 (2003) 291–308 MR1941437
- [2] **J W Anderson, H Parlier, A Pettet**, *Small filling sets of curves on a surface*, *Topology Appl.* 158 (2011) 84–92 MR2734699
- [3] **F Balacheff, H Parlier**, *Bers’ constants for punctured spheres and hyperelliptic surfaces*, *J. Topol. Anal.* 4 (2012) 271–296 MR2982444
- [4] **K Bezdek**, *Ein elementarer Beweis für die isoperimetrische Ungleichung in der Euklidischen und hyperbolischen Ebene*, *Ann. Univ. Sci. Budapest. Eötvös Sect. Math.* 27 (1984) 107–112 MR823098
- [5] **P Buser**, *Geometry and spectra of compact Riemann surfaces*, *Progress in Mathematics* 106, Birkhäuser, Boston (1992) MR1183224
- [6] **P Buser, P Sarnak**, *On the period matrix of a Riemann surface of large genus*, *Invent. Math.* 117 (1994) 27–56 MR1269424
- [7] **B Farb, D Margalit**, *A primer on mapping class groups*, *Princeton Mathematical Series* 49, Princeton Univ. Press (2012) MR2850125
- [8] **D D Long, C Maclachlan, A W Reid**, *Arithmetic Fuchsian groups of genus zero*, *Pure Appl. Math. Q.* 2 (2006) 569–599 MR2251482
- [9] **G A Margulis**, *Discrete subgroups of semisimple Lie groups*, *Ergeb. Math. Grenzgeb.* 17, Springer, Berlin (1991) MR1090825
- [10] **D Mumford**, *A remark on Mahler’s compactness theorem*, *Proc. Amer. Math. Soc.* 28 (1971) 289–294 MR0276410
- [11] **H Parlier**, *The homology systole of hyperbolic Riemann surfaces*, *Geom. Dedicata* 157 (2012) 331–338 MR2893492
- [12] **P Schmutz**, *Riemann surfaces with shortest geodesic of maximal length*, *Geom. Funct. Anal.* 3 (1993) 564–631 MR1250756

- [13] **P Schmutz Schaller**, *Geometry of Riemann surfaces based on closed geodesics*, Bull. Amer. Math. Soc. 35 (1998) 193–214 MR1609636
- [14] **P Schmutz Schaller**, *Systoles and topological Morse functions for Riemann surfaces*, J. Differential Geom. 52 (1999) 407–452 MR1761080
- [15] **K Takeuchi**, *Commensurability classes of arithmetic triangle groups*, J. Fac. Sci. Univ. Tokyo Sect. IA Math. 24 (1977) 201–212 MR0463116

*Department of Mathematics, Yale University
10 Hillhouse Avenue, New Haven, CT 06510, USA*

tarik.aougab@yale.edu, shinnyih.huang@yale.edu

Received: 26 January 2014 Revised: 4 July 2014