

BOUNDS FOR THE NUMBER OF MOVES BETWEEN PANTS DECOMPOSITIONS, AND BETWEEN TRIANGULATIONS

MARC LACKENBY, MEHDI YAZDI

ABSTRACT. Given two pants decompositions of a compact orientable surface S , we give an upper bound for their distance in the pants graph that depends logarithmically on their intersection number and polynomially on the Euler characteristic of S . As a consequence, we find an upper bound on the volume of the convex core of a maximal cusp (which is a hyperbolic structures on $S \times \mathbb{R}$ where given pants decompositions of the conformal boundary are pinched to annular cusps). As a further application, we give an upper bound for the Weil–Petersson distance between two points in the Teichmüller space of S in terms of their corresponding short pants decompositions. Similarly, given two one-vertex triangulations of S , we give an upper bound for the number of flips and twist maps needed to convert one triangulation into the other. The proofs rely on using pre-triangulations, train tracks, and an algorithm of Agol, Hass, and Thurston.

1. INTRODUCTION

1.1. Pants decompositions. Let S be a compact orientable surface. An *essential* curve on S is a homotopically non-trivial and non-boundary-parallel simple closed curve. The *curve graph* $\mathcal{C}(S)$ of S , defined by Harvey [18], is a simplicial graph whose vertices are isotopy classes of essential simple closed curves on S , and with two vertices joined by an edge if their simple closed curves are disjoint up to isotopy. Hempel’s Lemma [20] states that the distance between two simple closed curves α and β in the curve graph of S is at most a logarithmic function of their geometric intersection number, namely $2 \log_2 i(\alpha, \beta) + 2$.

A *pants decomposition* of S is a maximal collection of disjoint, pairwise non-parallel, essential simple closed curves on S up to isotopy. The *pants graph* $\mathcal{P}(S)$ of S , defined by Hatcher and Thurston [19], is a simplicial graph whose vertices are pants decompositions of S , with two vertices joined by an edge if their pants decompositions are related by one *simple* or *associativity* move as in Figure 1. Hatcher and Thurston showed that the pants graph is connected. The pants graph is intimately related to hyperbolic geometry in dimensions two and three. For example, Brock [8] showed that the Teichmüller space of S equipped with the Weil–Petersson metric is quasi-isometric to the pants graph of S , with the quasi-isometry map sending a marked hyperbolic metric to one of its *short* pants decompositions (i.e. each curve in the pants decomposition being shorter than the *Bers constant*).

In light of Hempel’s Lemma, a natural question is whether there is an upper bound for the distance between two arbitrary pants decompositions P and P' in the pants graph that depends logarithmically on the geometric intersection number $i(P, P')$. The following example shows that this is too optimistic if we do not take into account the topological complexity of the underlying surface.

ML was partially supported by EPSRC grant EP/Y004256/1. For the purpose of open access, the authors have applied a CC BY public copyright licence to any author accepted manuscript arising from this submission.




FIGURE 1. The edges of the pants graph correspond to simple moves (left) and associativity moves (right). In each case, the red curve is replaced with the blue curve, and the remaining curves are unchanged.

Example 1.1. *The complexity of a compact orientable surface S of genus g with b boundary components is defined as $\xi(S) = 3g - 3 + b$. It is easy to see that $\xi(S)$ is the number of curves in a pants decomposition of S . Given an essential multicurve γ in S , define the subsurface S_γ of S to be the components of $S \setminus N(\gamma)$ that are not pairs of pants, where $N(\gamma)$ is a regular neighbourhood of γ . Define $\mathcal{P}_\gamma(S)$ to be the full subgraph of the pants graph $\mathcal{P}(S)$ consisting of pants decompositions containing γ . Taylor and Zupan [29] showed that if γ is an essential multicurve such that S_γ is a disjoint union of surfaces of complexity 1 then $\mathcal{P}_\gamma(S)$ is totally geodesic in $\mathcal{P}(S)$. Note that in this case $\mathcal{P}_\gamma(S)$ is a product of copies of the Farey graph.*

Now suppose that S is closed and take γ to be a multicurve such that S_γ is a disjoint union of n surfaces of complexity 1 where $n \geq g$. If P and P' are pants decompositions of S both containing γ , then the distance between them is equal to the sum over distances between the restrictions of P and P' to the components of S_γ . For example if the restriction of P and P' to each component of S_γ consists of two curves intersecting either once or twice then $i(P, P') \leq 2n$ while $d_{\mathcal{P}_S}(P, P') \geq n$. Recall that n is comparable to g since we assumed that $g \leq n$, and also for Euler characteristic reasons $n \leq 2g - 2$. This shows that there is no general upper bound of the form

$$A(\log i(P, P'))^B + C$$

for distance in the pants graph, where A, B, C are positive universal constants independent of g .

Given the above example, it is natural to ask if there is an upper bound for the distance in the pants graph that depends logarithmically on the geometric intersection number of the two pants decompositions, and polynomially on the Euler characteristic of the underlying surface. A first attempt would be to try to adapt the proof of Hempel's Lemma to the context of pants decompositions. The proof that we have in mind uses a surgery trick, due to Lickorish [24], to construct a curve γ such that both $i(\gamma, \alpha)$ and $i(\gamma, \beta)$ are at most one half of $i(\alpha, \beta)$, and then it proceeds inductively. The reader familiar with the proof would soon realise that if one starts with multicurves α and β instead, the produced multicurve γ might have more or less components. In particular, even if α and β are pants decompositions, there is no guarantee that γ would be a pants decomposition too. A second attempt would be to apply Hempel's Lemma successively and $|P|$ times, with $|P|$ the number of curves in the pants decomposition P , to convert one curve of P at a time to one curve of P' . The issue is that applying the proof of Hempel's Lemma to the first curve in P makes the remaining curves more complicated rapidly, and its successive application does not seem to produce a good bound for the distance between P and P' in the pants graph. Nevertheless, using different techniques, we show that such a bound exists. Throughout the paper, the implicit constants in the big O notation are universal; in particular they do not depend on the underlying surface. Note that if P and P' are two pants decompositions of S that intersect each other exactly once, then the distance

between P and P' in the pants graph is exactly 1. Therefore, the hypothesis $i(P, P') \geq 2$ in the following theorem is not restrictive.

Theorem 1.2. *Let S be a compact connected orientable surface, and P and P' be pants decompositions of S with $i(P, P') \geq 2$. The distance between P and P' in the pants graph is*

$$O(|\chi(S)|^2 \log(i(P, P'))).$$

Moreover, there is an algorithm that constructs a (not necessarily geodesic) path $P = P_0, P_1, \dots, P_n = P'$ of length $n = O(|\chi(S)|^2 \log(i(P, P')))$ connecting P and P' in the pants graph, in time that is a polynomial function of $|\chi(S)|$ and $\log(i(P, P'))$. Here, for the input, P is given as a union of pairs of pants with gluing instructions, and P' is given in normal form with respect to P . For the output, each P_i is given as a union of pairs of pants with gluing instructions together with the pants move from P_{i-1} to P_i .

Remark 1.3. Let $\chi = |\chi(S)|$. Example 1.1 shows that any general upper bound for distance in the pants graph should be at least of order

$$\frac{\chi}{\log \chi} \log i(P, P').$$

Therefore our bound and an optimal bound differ by a multiplicative factor of at most $\chi \log \chi$.

We now briefly explain what we mean by normal form for P' with respect to P and refer the reader to Section 2.4 and Figure 3 for precise definitions. The *Dehn coordinates* express the isotopy class of any multicurve γ in terms of

- (1) the geometric *intersection numbers* of γ with each curve in a fixed pants decomposition P , and
- (2) the *twist numbers* about the curves in P .

Both $i(P, P')$ and the distance in $\mathcal{P}(S)$ between P and P' are unchanged by the action of the mapping class group of $S - P$, which is generated by Dehn twists about the curves in P . Dehn twisting P' about a curve α of P does not change the intersection numbers, and changes the twist number about α by m_α , where m_α is the intersection number with α . Therefore, for the purpose of Theorem 1.2, it is enough to know the intersection numbers as in (1), and the twisting numbers modulo the corresponding intersection numbers. This is essentially the data that is referred to as normal form for P' with respect to P .

Our proof of Theorem 1.2 uses pre-triangulations (Definition 1.5), train tracks, as well as an algorithm of Agol, Hass, and Thurston [2] for counting the number of orbits for the action of a pseudogroup generated by *isometries* between subintervals of a (discrete) interval $[1, N] := \{1, \dots, N\}$. See the Background Section.

Convention 1.4. *To simplify the statements, throughout the paper we allow a 1-complex to have components that are simple closed curves with no vertices on them. Any such component would be considered as an edge of the 1-complex as well.*

Definition 1.5 (Pre-triangulation). A *pre-triangulation* of a surface S is an embedded finite 1-complex $\mathcal{T}$ in S such that

- (1) every complementary region is either a disc, or an annulus, or a pair of pants; and
- (2) each component of ∂S either lies inside $\mathcal{T}$ or is disjoint from $\mathcal{T}$.

With Convention 1.4, both triangulations and pants decompositions are special cases of pre-triangulations. The connection between pants decompositions and pre-triangulations is as follows; see Section 3 for the details. We describe a natural procedure that given a pre-triangulation $\mathcal{T}$ and an ordering O on the set of edges of $\mathcal{T}$, constructs a pants decomposition

$\mathbf{P}(\mathcal{T}, \mathcal{O})$ of S (Construction 3.10). Conversely, given a pants decomposition P and a choice of a suitable 1-complex γ in S , called a *nerve* (Definition 3.3), a connected pre-triangulation $\mathcal{T} = \mathbf{T}(P, \gamma)$ of S can be constructed (Construction 3.7). It is shown that for any pre-triangulation $\mathcal{T}$ and orderings $\mathcal{O}$ and $\mathcal{O}'$ on the edges of $\mathcal{T}$, the distance between the pants decompositions $\mathbf{P}(\mathcal{T}, \mathcal{O})$ and $\mathbf{P}(\mathcal{T}, \mathcal{O}')$ in the pants graph is $O(|\mathcal{T}|^2)$, where $|\mathcal{T}|$ is the number of edges of $\mathcal{T}$ (Lemma 3.19). Using this, we show in Lemma 3.20 that for any two pre-triangulations $\mathcal{T}$ and $\mathcal{T}'$ and orderings $\mathcal{O}$ and $\mathcal{O}'$ on their edges, the distance between the pants decompositions $\mathbf{P}(\mathcal{T}, \mathcal{O})$ and $\mathbf{P}(\mathcal{T}', \mathcal{O}')$ is $O(|\mathcal{T} \cup \mathcal{T}'|^2)$, where $|\mathcal{T} \cup \mathcal{T}'|$ is the number of edges of the pre-triangulation $\mathcal{T} \cup \mathcal{T}'$ obtained by superimposing $\mathcal{T}$ and $\mathcal{T}'$. In particular, this implies Corollary 3.21 stating that the distance between two pants decompositions P and P' in the pants graph is $O(i(P, P')^2)$. As far as we are aware, albeit much weaker than Theorem 1.2, even this quadratic upper bound for distance in the pants graph is new.

1.2. One-vertex triangulations. The idea of the proof of Theorem 1.2 can be used to give an upper bound for the number of *flip and twist moves* between two triangulations of a surface. The *flip graph* of a closed orientable surface S is a simplicial graph whose vertices are one-vertex triangulations of S up to isotopy, with two vertices joined by an edge if they differ by a *flip* of a diagonal in a simplicial square. It is well-known that the flip graph is connected; see for example [25, page 36]. The mapping class group of S acts by isometries on the flip graph, and the action is properly discontinuous and cocompact. By the Švarc–Milnor Lemma, the flip graph of S is quasi-isometric to the mapping class group of S , where the metric on the mapping class group is the word metric with respect to any fixed finite generating set. Given two one-vertex triangulations $\mathcal{T}$ and $\mathcal{T}'$ of S with the same vertex, one can convert $\mathcal{T}$ to $\mathcal{T}'$ using a sequence of at most $i(\mathcal{T}, \mathcal{T}')$ flips; see e.g. [13, Lemma 2.1] and below for the definition of $i(\mathcal{T}, \mathcal{T}')$. In general, the linear dependency on $i(\mathcal{T}, \mathcal{T}')$ cannot be improved, essentially because the flip graph is quasi-isometric to the mapping class group. For example, if $\mathcal{T}$ is a fixed one-vertex triangulation of S , α is a fixed essential simple closed curve in S , and $\mathcal{T}' := (T_\alpha)^n(\mathcal{T})$ where T_α is the Dehn twist about α , then

- $i(\mathcal{T}', \mathcal{T})$ is linear in n ; and
- $\mathcal{T}'$ and $\mathcal{T}$ have distance at least a constant multiple of n in the flip graph. This is because by the work of Farb, Lubotzky, and Minsky [17], the word length of $(T_\alpha)^n$ in the mapping class group is linear in n .

However, we show in Theorem 1.6 that if we are allowed to use powers of Dehn twists (*twist maps* for short) in addition to flips, then one can convert $\mathcal{T}$ to $\mathcal{T}'$ much more economically. In this case, the total number of flips and twist maps needed will be bounded above by a polynomial function of $\log(i(\mathcal{T}, \mathcal{T}'))$ and $|\chi(S)|$.

Given one-vertex (respectively ideal) triangulations $\mathcal{T}$ of a surface S with the same vertex set V , let $i(\mathcal{T}, \mathcal{T}')$ be the minimum number of intersections between $\mathcal{T}$ and $\mathcal{T}'$ outside of V , up to isotopies fixing V . Therefore $i(\mathcal{T}, \mathcal{T}') = 0$ if and only if $\mathcal{T} = \mathcal{T}'$ up to isotopies keeping the vertex set fixed. Furthermore, if $i(\mathcal{T}, \mathcal{T}') = 1$ then $\mathcal{T}$ and $\mathcal{T}'$ are related by a flip. Therefore in the following theorem the assumption $i(\mathcal{T}, \mathcal{T}') \geq 2$ is not restrictive.

Theorem 1.6. *Let S be a compact orientable surface. When S is closed (respectively, has non-empty boundary), let $\mathcal{T}$ and $\mathcal{T}'$ be one-vertex (respectively, ideal) triangulations of S . Assume that $i(\mathcal{T}, \mathcal{T}') \geq 2$. There is a sequence $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n = \mathcal{T}'$ of one-vertex (respectively, ideal) triangulations of S such that:*

- (1) *Each $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by either a flip or a power of a Dehn twist. When $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by Dehn twisting k times about a curve α , then α is a normal*

curve intersecting each edge of $\mathcal{T}_i$ at most three times, and the absolute value of k is bounded above by $i(\mathcal{T}, \mathcal{T}')$.

$$(2) \ n = O(|\chi(S)|^3) \log(i(\mathcal{T}, \mathcal{T}')).$$

Moreover, there is an algorithm that constructs the sequence $\mathcal{T}_1, \dots, \mathcal{T}_n$ in time that is a polynomial function of $\log(i(\mathcal{T}, \mathcal{T}'))$ and $|\chi(S)|$. Here we assume that $\mathcal{T}$ and $\mathcal{T}'$ have the same vertex (when S is closed), $\mathcal{T}$ is given as a union of (possibly ideal) triangles with gluing instructions, and $\mathcal{T}'$ is given in terms of its normal coordinates with respect to $\mathcal{T}$. For the output, each $\mathcal{T}_i$ is given as a union of (possibly ideal) triangles with gluing instructions together with the flip or twist move from $\mathcal{T}_i$ to $\mathcal{T}_{i+1}$.

See Definition 2.5 for the definition of the normal form for a triangulation with respect to another triangulation. Theorem 1.6 has been used by Baroni [3] and Lackenby [22] to study algorithmic problems about curves on surfaces and surface homeomorphisms, and it is likely to have further applications. Marc Bell and Richard Webb have communicated to us [5] that they, independently, were aware of an analogue of Theorem 1.6.

In Theorem 6.2, as a precursor to Theorem 1.6, we prove an analogous statement (Theorem 6.2) for a similar set of moves between two *polygonal decompositions* (Definition 2.7) of a compact orientable surface. In Theorem 6.9, we also prove a similar result for a set of moves between two spines of a compact orientable surface.

1.3. Volumes of hyperbolic 3-manifolds. A *maximal cusp* is a geometrically finite hyperbolic 3-manifold such that each component of its conformal boundary is a thrice-punctured sphere [11]. The *convex core* of a non-compact hyperbolic manifold is the smallest geodesically convex subset that contains every closed geodesic. Theorem 1.2 together with the work of Agol [1] implies the following.

Corollary 1.7. *Let Σ a closed orientable surface of genus $g \geq 2$, and P and P' be pants decompositions of Σ with no curve in common. Assume that M is a maximal cusp obtained from a quasi-Fuchsian 3-manifold homeomorphic to $\Sigma \times \mathbb{R}$ by pinching the multicurves P and P' to annular cusps on the two conformal boundary components of M . The volume of the convex core of M is*

$$O(g^2) \log(i(P, P')).$$

In Proposition 8.1 we give examples demonstrating that the bound in Corollary 1.7 is sharp up to a multiplicative factor of $g \log(g)$.

1.4. Weil–Petersson distance in the Teichmüller space. Let S be an orientable surface of finite type, and denote by $\text{Teich}(S)$ the Teichmüller space of S (in other words, the space of marked hyperbolic metrics possibly with cusps but with no boundary). Bers showed [7] that there is a constant $C = C(S)$ only depending on the topology of S such that for every $X \in \text{Teich}(S)$ there is a pants decomposition P_X in S such that each curve in P_X has length at most C with respect to the marked hyperbolic metric X . In other words every hyperbolic metric on S admits a ‘short’ pants decomposition. By Parlier’s [26] quantified version of Bers’s theorem, which is building on and improving the work of Buser [9] and of Buser and Seppälä [10], every hyperbolic metric X on S has a pants decomposition P in which every curve has length at most $2\pi|\chi(S)|$.

Denote the Weil–Petersson metric on $\text{Teich}(S)$ by d_{WP} . Using Theorem 1.2, the work of Wolpert [31], Cavendish and Parlier [12], and Parlier [26] we show the following.

Theorem 1.8. *Let S be an orientable surface of finite type, and let $X, Y \in \text{Teich}(S)$. Let P_X and P_Y be pants decompositions for X and Y respectively, in which each curve has length at most $2\pi|\chi(S)|$, which exist by a theorem of Parlier [26]. Then*

$$d_{\text{WP}}(X, Y) \leq O(|\chi(S)|^2) (1 + \log(i(P_X, P_Y) + 1)).$$

1.5. Idea of the proof of Theorem 1.2. Given pants decompositions P and P' , we first put P' in normal form with respect to P . Then we squeeze parallel normal arcs together to produce a train track τ such that P' is carried by τ . In other words, there is an integral weight μ on branches of τ such that the carried 1-complex $\mathcal{CC}(\tau, \mu)$ is equal to P' .

Given an integrally weighted train track (τ, μ) , and making some extra choices $\mathcal{I}$, we can define an *orbit counting problem* $\text{OCP}(\tau, \mu, \mathcal{I})$ (Definition 4.4). The extra data $\mathcal{I}$ is called an *initial interval identification* and can be safely ignored for now. Applying the Agol–Hass–Thurston algorithm, we can count the number of orbits $\text{OCP}(\tau, \mu, \mathcal{I})$. While the number of orbits of $\text{OCP}(\tau, \mu, \mathcal{I})$ is not any new information (it is always equal to the number of curves in P' , which is $(3|\chi(S)| - |\partial S|)/2$), we are interested in the way that the AHT algorithm calculates the number of orbits. In Proposition 4.13 and Lemma 4.17, we show that the AHT algorithm does this count in n steps with

$$n - 1 \leq E(1 + \log |\mu|),$$

where $E = O(|\chi(S)|)$ is the number of branches of τ and $|\mu| = i(P, P')$ is the total weight of all branches of τ . Moreover, each step of the the AHT algorithm can be seen as a geometric change on the underlying integrally weighted train track, namely as either a *splitting* or a *twirling* (Definition 4.6 and Figure 18). The algorithm proceeds until (τ, μ) is completely unwound to P' , so intuitively the algorithm does the reverse process of squeezing parallel arcs together, in a precise and efficient way. Hence, there is a sequence (τ_i, μ_i) of integrally weighted train tracks for $i = 1, \dots, n$ starting with $(\tau_1, \mu_1) = (\tau, \mu)$ and terminating at P' such that each (τ_{i+1}, μ_{i+1}) is obtained from (τ_i, μ_i) by splitting or twirling, and

$$(1) \quad n - 1 = O(|\chi(S)|)(1 + \log i(P, P')).$$

The underlying 1-complex of each τ_i is a pre-triangulation. By Lemma 4.10, since (τ_{i+1}, μ_{i+1}) is obtained from (τ_i, μ_i) by a simple geometric operation (splitting or twirling), the following holds: there are orderings O_i on the edges of τ_i such that for most i the distance between the pants decompositions $\mathbf{P}(\tau_i, O_i)$ and $\mathbf{P}(\tau_{i+1}, O_{i+1})$ is $O(|\chi(S)|)$. Moreover, by Lemma 3.20 the distance between the pants decompositions P and $\mathbf{P}(\tau_1, O_1)$ is $O(|\chi(S)|^2)$, because $i(\tau_1, P) = O(|\chi(S)|)$. Note also that $\mathbf{P}(\tau_n, O_n) = P'$. The triangle inequality now implies that the distance between P and P' is at most $O(|\chi(S)|^2) + (n - 1)O(|\chi(S)|)$, which combined with equation (1) gives the desired upper bound.

1.6. Previous related work. Several authors had previously used either train tracks and the Agol–Hass–Thurston algorithm, or the *tracing algorithm* of Erickson and Nayyeri [16], to simplify curves on surfaces efficiently.

- Motivated by the AHT algorithm, Dunfield and D. Thurston [14] adapted Brown’s algorithm in the context of splittings of train tracks to decide efficiently whether a random tunnel number one 3-manifold fibers over the circle.
- Dynnikov and Wiest [15] introduced a specialised version of the AHT algorithm for curves on a punctured disc, called the *transmission-relaxation algorithm*, to simplify a weighted train track associated with a braid in an n -times punctured disc. Using this, they proved that certain algebraic and geometric complexity measures for a braid

are comparable up to multiplicative constants only depending on n . They showed an improved running time for their algorithm compared to the general version of the AHT algorithm. In the proof of Proposition 4.13, we use their improved bound to analyse the time complexity of the AHT algorithm applied to a weighted train track.

- Erickson and Nayyeri [16] gave an algorithm to *trace* a curve γ with respect to a triangulation $\mathcal{T}$ in time that is a polynomial function of $\log(i(\mathcal{T}, \gamma))$ and $|\chi(S)|$.
- Bell [4] showed how to simplify a triangulation $\mathcal{T}$ of a surface S with respect to a *curve* γ using flips and twist maps, in time that is a polynomial function of $\log(i(\mathcal{T}, \gamma))$ but a super-exponential function of $|\chi(S)|$. He has communicated to us that the source code of his software program *flipper* does this simplification in time that is a polynomial function of both $\log(i(\mathcal{T}, \gamma))$ and $|\chi(S)|$. See also the work of Bell and Webb [6].

Another novel feature of our work is to show that the AHT algorithm can be used to establish non-algorithmic results such as Corollary 1.7 and Theorem 1.8.

1.7. Plan of the paper. In Section 2, we discuss the background on the algorithm of Agol, Hass, and Thurston. We also introduce various normal forms needed for the statements of the theorems. These include

- normal form for a pants decomposition with respect to another pants decomposition. This is a slight modification of Dehn coordinates, and is used in Theorem 1.2.
- normal form for a triangulation with respect to another triangulation with common vertices. This is a mild generalisation of the usual notion of normal curve with respect to a triangulation, and is used in Theorem 1.6.
- normal form for a 1-complex with respect to a polygonal decomposition. This is a more relaxed notion than the usual notion of a normal curve with respect to a triangulation, since the 1-complex can be quite general. This is used as a bookkeeping tool in Theorems 6.2 and 6.9.

In Section 3, we discuss pre-triangulations and their relation to pants decompositions. Section 4 discusses based weighted train tracks and orbit counting problems associated to them. The effect of the AHT algorithm applied to such orbit counting problems is seen as a geometric move on the underlying train track (Proposition 4.13). Section 5 contains the proof of Theorem 1.2. We first prove Theorems 6.2 and 6.9 in Section 6, and then use the latter to prove Theorem 1.6 in Section 7. Section 8 proves Corollary 1.7 using only the statement of Theorem 1.2 as well as the work of Agol [1]. Theorem 1.8 is proved in Section 9. Some open problems are discussed in Section 10.

2. BACKGROUND

2.1. Agol–Hass–Thurston (AHT) algorithm. By an *interval* $[M, N]$ with $M \leq N$ integers, we mean the set $\{M, M+1, \dots, N\}$. An *isometry* or *pairing* between two intervals is a bijection of the form $x \mapsto x+c$ or $x \mapsto -x+c$ for a constant integer c . Let $[1, N]$ be an interval and $\{g_i\}$, $1 \leq i \leq k$, be a collection of pairings between subintervals of $[1, N]$. Two pairings can be composed if the range of the first lies in the domain of the second. The pairings generate a pseudogroup $\mathcal{G}$ under the operations of composition where defined, inverses, and restriction to subintervals. Two points $x, y \in [1, N]$ are in the same orbit if there is a sequence $h_1, \dots, h_s$ of (not necessarily distinct) elements of $\mathcal{G}$ such that $h_s \circ \dots \circ h_1(x) = y$. This forms an equivalence relation on the interval $[1, N]$ and the equivalence classes are called *orbits*. Given an interval $[1, N]$ and isometric pairings $g_1, \dots, g_k$ as above, the associated *orbit counting problem* is the problem of counting the number of orbits for the action of the pseudogroup $\mathcal{G}$ (generated by $\{g_i\}$ for $1 \leq i \leq k$) on the interval $[1, N]$. Agol, Hass, and

Thurston gave an algorithm for counting the number of such orbits, in time that is a polynomial function of k and $\log(N)$. This logarithmic dependence on N will be crucial for the arguments in this paper.

An important example of an orbit counting problem comes from the theory of train tracks (see Definition 4.1). For example suppose S is a surface of finite type, $\tau \subset S$ is a train track, and γ is a multicurve carried by τ . Then γ can be described using integral weights on branches of τ . A question that arises in practice is to determine the number of connected components of γ given the weights associated to γ . Agol, Hass, and Thurston used their algorithm to count the number of connected components of γ , in time that is bounded above by a polynomial function of $E \log(N)$, where E is the number of branches of τ and N is the total weight of γ with respect to τ .

Similarly, Agol, Hass, and Thurston used their algorithm to count the number of connected components of a normal surface S in a triangulated 3-manifold M given by its normal coordinates with respect to the triangulation, and used it to verify an upper bound on the genus of a knot $K \subset M$ in non-deterministic polynomial time.

We now explain further terminology and the AHT algorithm briefly; the reader should see [2] for more details. A pairing $g: [a, b] \rightarrow [c, d]$ is *orientation-preserving* if $g(a) = c$ and $g(b) = d$, and otherwise *orientation-reversing*. If $a \leq c$ we refer to $[a, b]$ as the *domain* and $[c, d]$ as the *range* of the pairing. If the pairing preserves orientation, its *translation distance* t is defined as $t = c - a = d - b$. An interval is called *static* if it is in neither the domain nor the range of any pairing. Given a collection of pairings acting on $[1, N]$, a pairing is *maximal* if its range contains both N and the range of any other pairing containing N . A pairing $g: [a, b] \rightarrow [c, d]$ is *periodic* with period t if it is orientation-preserving with translation distance t and $a < c = a + t \leq b + 1$, so there is no gap between the domain and range. The combined interval $[a, d]$ is then called a *periodic interval* of period t .

The orbit-counting algorithm applies a series of the following modifications, in a particular order, to a collection of pairings.

- *Periodic merger*: Given two periodic pairings g_1 and g_2 with periodic intervals R_1 and R_2 and periods t_1 and t_2 such that $\text{width}(R_1 \cap R_2) \geq t_1 + t_2$, the periodic merger replaces g_1 and g_2 with a single periodic pairing of period $\text{GCD}(t_1, t_2)$. Here GCD stands for the greatest common divisor.

Remark 2.1. Assume that a periodic merger is done. Then before the merging, there is a point which lies in at least 3 domains or ranges of pairings, i.e.

$$\exists p \in [1, N] \text{ such that } \#\{i : p \in \text{domain}(g_i)\} + \#\{i : p \in \text{range}(g_i)\} \geq 3.$$

In our applications in this paper, every point will be in the domain or range of at most two pairings, hence periodic mergers will not occur. For this reason we will not elaborate on the properties of periodic merger further.

- *Contraction*: Contraction can be performed on a static interval $[r, s]$. We eliminate this interval, replace $[1, N]$ by $[1, N - (s - r + 1)]$, and change each g_j by replacing any point x in a domain or range which lies entirely to the right of s by $x - (s - r + 1)$. We will then decrease the number of orbits by $s - r + 1$, since the eliminated points are each unique representatives of an orbit.
- *Trimming*: The trimming operation simplifies an orientation-reversing pairing whose domain and range overlap.

Remark 2.2. In our applications in this paper, the orientability of the underlying surface implies that there will be no trimming.

- *Truncation*: Given an interval that lies in the domain or range of exactly one pairing, we can remove it without changing the orbit structure. This operation can be applied to remove points from the right of the interval $[1, N]$. Assume that there is a pairing $g: [a, b] \rightarrow [c, N]$ and a value N' with $c \leq N' + 1 \leq N$, such that all points in the interval $[N' + 1, N]$ are in the range of only g . Truncating g shortens the interval $[1, N]$ to the interval $[1, N']$, and similarly shortens the domain and range of g . If g is orientation-reversing, truncation is applied only when g has disjoint domain and range (i.e. after trimming).
- *Transmission*: In transmission, a pairing g_1 is used to shift the domain and range of a second pairing g_2 to the left as much as possible. Once the pairings are shifted leftwards, we can subsequently apply truncation. Transmissions allow significant simplifications of the orbit counting problem in one step.

If g_1 is orientation-reversing and has overlapping domain and range, then as a first step in transmission we trim g_1 . Now consider a pairing g_1 , and a second pairing g_2 whose range is contained in the range of g_1 . Consider two cases:

- If $\text{domain}(g_2) \not\subseteq \text{range}(g_1)$: then define the map $g'_2 = g_1^{-r} \circ g_2$, where $r = 1$ if g_1 is orientation-reversing and otherwise $r \geq 1$ is the largest integer such that $g_1^{-r+1}([c_2, d_2])$ is contained in the range of g_1 .
- If $\text{domain}(g_2) \subseteq \text{range}(g_1)$: then define the map $g'_2 = g_1^r \circ g_2 \circ g_1^s: g_1^{-s}([a_2, b_2]) \rightarrow g_1^r([c_2, d_2])$, where r is as above, $s = 1$ if g is orientation reversing, and otherwise $s \geq 1$ is the largest integer such that $g_1^{-s+1}([a_2, b_2])$ is contained in the range of g_1 . The operation of replacing g_2 by $g_1^{-r} \circ g_2 \circ g_1^s$ is called a transmission of g_2 by g_1 .

The AHT algorithm repeatedly applies the following steps (1)–(6) to the orbit counting problem.

- (1) Delete any pairings that are restrictions of the identity.
- (2) Make any possible contractions and, if any exist, increment the orbit counter by the sum of the number of points deleted by the contractions. If the number of pairings remaining is zero, output the number of orbits and stop.
- (3) Trim all orientation-reversing pairings whose domain and range overlap.
- (4) Search for pairs of periodic pairings g_i and g_j with periods t_1 and t_2 and with overlapping periodic intervals R_1 and R_2 such that $\text{width}(R_1 \cap R_2) \geq t_1 + t_2$. If any such pair exists, then perform a merger replacing g_i and g_j by a single periodic pairing. Repeat until no mergers can be performed.
- (5) Find a maximal g_i . For each $g_j \neq g_i$ whose range is contained in $[c_i, N']$, transmit g_j by g_i .
- (6) Find the smallest value of c such that the interval $[c, N']$ intersects the range of exactly one pairing. Truncate the pairing whose range contains the interval $[c, N']$.

A *cycle* of the AHT algorithm consists of applying steps (1)–(6) above.

Example 2.3 (Euclid’s algorithm for GCD). *Agol, Hass, and Thurston’s algorithm is in a sense a generalisation of Euclid’s algorithm for computing the greatest common divisor (GCD) of two natural numbers a and b . Let $N = a + b$, and consider the interval $[1, N] = [1, a + b]$ with pairings*

$$\begin{aligned} f(x) &= x + a, & f: [1, b] &\rightarrow [a + 1, a + b], \\ g(x) &= x + b, & g: [1, a] &\rightarrow [b + 1, a + b]. \end{aligned}$$

The number of orbits for the action of the pseudogroup generated by f and g is equal to $d := \text{GCD}(a, b)$. This will be a consequence of the upcoming discussion (or one can show it directly).

Let us apply one cycle of AHT algorithm to this orbit counting problem. Steps (1)–(4) are not applicable for the first cycle, and so we apply Step (5). Assume that $a < b$. Then f is maximal, and we transmit g by f . Let $b = na + r$ where n is an integer and $0 \leq r < a$. The transmission replaces g by $g' = f^{-n} \circ g$ where

$$g'(x) = x + b - na = x + r \quad g': [1, a] \rightarrow [r + 1, r + a].$$

Then Step (6) truncates f and replaces N with $N' = r + a$ and f with f' where

$$f'(x) = x + a, \quad f': [1, r] \rightarrow [a + 1, r + a].$$

Hence, after one cycle of the AHT algorithm we have replaced the pair (a, b) with the pair (r, a) . This is essentially reflecting the equality $\text{GCD}(a, b) = \text{GCD}(r, a)$ in Euclid's algorithm. Repeating this procedure, we obtain the number of orbits. Note that this inductively shows that the number of orbits is equal to $\text{GCD}(a, b)$, where the base of the induction corresponds to the case $a = 0$.

In Proposition 4.13, we will be interested in the effect of applying the AHT algorithm to certain orbit counting problems associated with train tracks. In particular, we would like to see the effect of one cycle as a geometric change in the underlying train track. In order to see a meaningful geometric change, we will shift the steps in one cycle of the AHT algorithm as follows.

Definition 2.4 (Shifted cycle of the AHT algorithm). The first *shifted cycle* of the AHT algorithm applies steps (1)–(4) of the orbit counting algorithm. Afterward, every shifted cycle applies steps (5)–(6) followed by steps (1)–(4). Clearly the AHT algorithm is the successive application of shifted cycles as well.

2.2. Normal form for a simple curve with respect to a triangulation. Let $\mathcal{T}$ be a triangulation of a compact surface. For a triangle Δ of $\mathcal{T}$, a *normal arc* is a simple arc with endpoints lying on the interior of edges of Δ and joining two different sides of Δ . Two normal arcs α and β in Δ are of the same *type* if there is an isotopy of Δ , preserving the edges and vertices of Δ throughout, taking one arc to the other. Each triangle contains 3 types of normal arcs, and so there are in total $3t$ types of normal arcs, where t is the number of triangles in $\mathcal{T}$; see Figure 2. Let γ be a properly embedded simple multicurve on S . We allow γ to consist of both simple arcs and simple closed curves. We say γ is a *normal curve* with respect to $\mathcal{T}$ if for each triangle Δ of $\mathcal{T}$, the intersection $\gamma \cap \Delta$ is a union of normal arcs. Fix a bijection between the set of normal arc types and $\{1, 2, \dots, 3t\}$. The *normal vector* or the *normal coordinates* $(\gamma) \in \mathbb{Z}^{3t}$ is the vector recording the number of normal arcs of each type for γ .

Any properly embedded simple multicurve on S , with no homotopically trivial simple closed curve component, can be isotoped through isotopies fixing ∂S setwise to a normal curve.

In the next few subsections, we extend the definition of normal form to other more general situations, which will be used in this paper.

2.3. Complexity of integral vectors. Given a vector $v = (v_1, \dots, v_n) \in \mathbb{Z}^n$ define the ℓ^1 -norm of v as

$$|v|_1 := \sum_{i=1}^n |v_i|.$$




FIGURE 2. The three types of normal arcs in a triangle.

Given an integer a , let $\text{dig}(a)$ be the number of digits of a in binary. In the case where a is negative, we view the minus sign in front as an extra digit. The *bit-sized complexity* of v is defined as

$$|v|_{\text{bit}} := \sum_{i=1}^n \text{dig}(v_i).$$

If γ is a normal curve with respect to a triangulation and $v = (\gamma)$ is the normal vector for γ , then we can speak of the ℓ^1 -norm of γ as $|(\gamma)|_1$, and of the *bit-sized complexity* of γ as $|(\gamma)|_{\text{bit}}$.

2.4. Normal form for pants decompositions. We define the Dehn parametrisation of isotopy classes of multicurves on a surface, following Penner and Harer [27, page 13]. We then modify Dehn coordinates slightly to define a normal form for a pants decomposition with respect to another pants decomposition. The motivation for this modification is that in Theorem 1.2, the isotopy class of P' is needed only up to the action of the mapping class group of $S - P$.

Let S be a compact orientable surface of genus g with b boundary components and let P be a pants decomposition of S . Denote the components of P by α_r where $1 \leq r \leq |P| = 3g - 3 + b$. For each pants curve α_i , choose a closed arc $w_i \subset \alpha_i$ called a *window*. For each pair of pants F in P , choose a collection of properly embedded arcs $\ell_{i,j}$ ($1 \leq i, j \leq 3$) in F with endpoints contained in the windows as in Figure 3. Let (m_r) be non-negative integers and (t_r) be integers for $1 \leq r \leq |P|$ such that

- if $m_r = 0$ then $t_r \geq 0$;
- for each embedded pair of pants F , the sum $m_{i_1} + m_{i_2} + m_{i_3}$ corresponding to the three boundary components of F is even; and
- the number m_r corresponding to a pants curve that bounds a torus minus a disc or a twice punctured disc is even.

Then we can construct a multicurve γ as follows: take a number of parallel copies of the arcs $\ell_{i,j}$ such that the intersection number with α_r is exactly m_r . These parallel arcs can be naturally pasted together to produce a multicurve, and this corresponds to the case of $t_r = 0$ and $m_r \neq 0$. If $m_r, t_r \neq 0$, then we paste these parallel arcs on the two sides of α_r by a $(\frac{t_r}{m_r})$ -fractional Dehn twist, where the twist is right-handed if $t_r > 0$. If $t_r \neq 0$ and $m_r = 0$ then we add t_r parallel copies of α_r to the multicurve. The numbers m_r are called the *intersection numbers*, and t_r are called the *twisting numbers*. The isotopy class of every multicurve, with all components essential, can be uniquely represented in this way. This is called the *Dehn coordinates* of the multicurve.

In our applications, we only need the twisting numbers t_r modulo the integers m_r , since a full twist (that is a Dehn twist) about α_r lies in the mapping class group of $S - P$. We say



FIGURE 3. The windows w_i and the arcs $\ell_{i,j}$.

FIGURE 4. Here ABC and $A'B'C'$ show two triangles of $\mathcal{T}$. Left: The three (blue dashed) arcs AB , AD , and EF of $\mathcal{T}'$ are normal with respect to $\mathcal{T}$. Right: The two (blue dashed) arcs $A'D'$ and $E'F'$ of $\mathcal{T}'$ are not normal with respect to $\mathcal{T}$.

that a pants decomposition P' is given in normal form with respect to a pants decomposition P if the coordinates $(m_r, t_r \bmod m_r)$ of P' with respect to P are given $(1 \leq r \leq |P|)$.

2.5. Normal form for triangulations with common vertices.

Definition 2.5 (Normal form for a triangulation). Let $\mathcal{T}$ and $\mathcal{T}'$ be triangulations of a compact surface S with the same set of vertices. We say that $\mathcal{T}'$ is in *normal form with respect to $\mathcal{T}$* if for each triangle Δ of $\mathcal{T}$, the intersection $\gamma \cap \Delta$ consist of a union of arcs of the following types

- an arc connecting different sides of Δ ; or
- an arc connecting a common vertex of Δ and $\mathcal{T}'$ to its opposite side; or
- a side of Δ whose endpoints are vertices of $\mathcal{T}'$.

See Figure 4 left for some examples of arcs that are normal, and Figure 4 right for some non-examples.

Proposition 2.6. *Let $\mathcal{T}$ and $\mathcal{T}'$ be triangulations of a compact surface S with the same set of vertices. Assume that every edge of $\mathcal{T}'$ with the same endpoints is homotopically essential. Then $\mathcal{T}'$ can be isotoped relative to its vertices to be in normal form with respect to $\mathcal{T}$.*




FIGURE 5. Here the two triangles belong to the polygonal decomposition $\mathcal{T}$, and the blue dashed 1-complex is the intersection of the 1-complex γ with the corresponding triangles. On the left the restriction of γ to the triangle is in normal form with respect to $\mathcal{T}$. On the right γ is not in normal form with respect to $\mathcal{T}$ because of either of the two shaded bigons between $\mathcal{T}$ and γ .

Proof. First, isotope $\mathcal{T}'$ relative to its vertices to be in general position with respect to $\mathcal{T}$. Isotope $\mathcal{T}'$ relative to its vertices and remove any bigons with one side in $\mathcal{T}$ and one side in $\mathcal{T}'$. This reduces the weight $|\mathcal{T}' \cap \mathcal{T}|$ until no bigons are left. At this point, for any triangle Δ of $\mathcal{T}$, any component of $\mathcal{T}' \cap \Delta$ is a normal arc or a simple closed curve lying entirely in Δ . The latter possibility is ruled out by the assumption that every edge of $\mathcal{T}'$ with the same endpoints is homotopically essential. $\square$

2.6. Normal form for a 1-complex with respect to a polygonal decomposition.

Definition 2.7 (Polygonal decomposition). A *polygonal decomposition* $\mathcal{D}$ of a compact surface S is an embedded connected finite 1-complex such that

- each complementary region adjacent to a boundary component of S is topologically either a disc or an annular neighbourhood of the boundary component; and
- every other complementary region is a topological disc.

The next definition is a generalisation of the usual notion of normal form for multicurves.

Definition 2.8 (Normal form for a 1-complex with respect to a polygonal decomposition). Let $\mathcal{T}$ be a polygonal decomposition of a compact surface S , and γ be a finite 1-complex embedded in S . We say that γ is in *normal form with respect to $\mathcal{T}$* if

- (1) γ has no homotopically trivial simple closed curve component that lies in a 2-cell of $\mathcal{T}$; and
- (2) there are no bigons between $\mathcal{T}$ and γ except possibly those bigons that were already present in γ ; i.e. there is no embedded disc $D \subset S$ whose interior is disjoint from $\gamma \cup \mathcal{T}$ and with $\partial D = \alpha \cup \beta$ where α and β are arcs with $\partial\alpha = \partial\beta$, and α lies in an edge of $\mathcal{T}$ but not in an edge of γ , and β lies in an edge of γ .

See Figure 5 left for an example of a 1-complex in normal form with respect to a polygonal decomposition, and Figure 5 right for a non-example.

Proposition 2.9. Let $\mathcal{T}$ be a polygonal decomposition of a compact surface S , and γ be a finite 1-complex embedded in S . Then γ can be isotoped relative to its vertices to be in normal form with respect to $\mathcal{T}$.

Proof. Denote the i -skeleton of $\mathcal{T}$ and γ by respectively $\mathcal{T}^i$ and γ^i for $i = 0, 1$. First, isotope γ relative to its vertices to be in general position with respect to $\mathcal{T}$; i.e. such that $\gamma^1 - \gamma^0$ is transverse to $\mathcal{T}^1 - \mathcal{T}^0$. All isotopes we consider will be relative to the vertices of γ . By a bigon we mean a bigon between $\mathcal{T}$ and γ . Denote by $i(\gamma, \mathcal{T})$ the number of transverse intersections




FIGURE 6. Left: the triangle belongs to the polygonal decomposition $\mathcal{T}$, and the dashed blue lines are the intersection of the 1-complex γ with $\mathcal{T}$. Right: the normal coordinates of γ with respect to $\mathcal{T}$ are shown.

of $\gamma^1 - \gamma^0$ with $\mathcal{T}^1 - \mathcal{T}^0$. If B is a bigon such that at least one vertex of B is not in γ^0 , then we can isotope γ relative to its vertices and remove B while reducing the intersection number $i(\gamma, \mathcal{T})$. Repeat until every bigon has both vertices in γ^0 . For any remaining bigon B with $\partial B = \alpha \cup \beta$, if one side of ∂B , say α , lies in an edge of $\mathcal{T}$ but not in an edge of γ , then isotope the other side β of ∂B into α , and increase the number of edges of γ that lie in an edge of $\mathcal{T}$. Repeating this process puts γ in normal form with respect to $\mathcal{T}$. $\square$

Definition 2.10 (Normal coordinates of a 1-complex with respect to a polygonal decomposition). Let $\mathcal{T}$ be a polygonal decomposition of a compact surface S , and let γ be a 1-complex that is in normal form (Definition 2.8) with respect to $\mathcal{T}$. Then $\mathcal{T}$ cuts γ into connected components, each of which is either

- (1) a graph with at least one vertex that is embedded in a closed 2-cell of $\mathcal{T}$; or
- (2) an arc properly embedded in a 2-cell of $\mathcal{T}$ whose endpoints lie on distinct edges of the 2-cell.

Let Δ be a 2-cell of $\mathcal{T}$ and a, b be two distinct sides of Δ . Given two disjoint simple properly embedded arcs $\gamma_1, \gamma_2 \subset \Delta$ as in (2) above and joining a and b , there is a rectangle with two opposite sides γ_1 and γ_2 and the other two sides lying on a and b . We say that γ_1 and γ_2 are of the same *type* if any other piece of γ in R is a properly embedded arc joining a and b as well. We can specify $\mathcal{T}$ by specifying a weighted 1-complex $\mathcal{N}$ in each 2-cell of $\mathcal{T}$ where

- $\mathcal{N}$ contains all components as in (1) above, and
- for each arc type appearing in (2), $\mathcal{N}$ contains one copy of the arc together with a positive integer weight that counts the number of arcs of that type.

The data of the weighted 1-complex $\mathcal{N}$ over all 2-cells of $\mathcal{T}$ is called the *normal coordinates of γ with respect to $\mathcal{T}$* .

Note that if γ is a multicurve, then only components as in (2) above (i.e. arcs) appear, and we recover the usual notion of normal coordinates for a multicurve.

Example 2.11. In Figure 6 left, the intersection of γ with a triangle of $\mathcal{T}$ is shown, which can be seen to be in normal form. On the right the normal coordinates of γ with respect to $\mathcal{T}$ are shown.

3. FROM PANTS DECOMPOSITIONS TO PRE-TRIANGULATIONS, AND VICE VERSA

3.1. From pants decompositions to pre-triangulations.

Definition 3.1 (Bipartite adjacency graph of a pants decomposition). Let P be a pants decomposition of a compact connected orientable surface S . Define the *bipartite adjacency graph* $\Gamma = \Gamma(P)$ with vertices $V = B \cup W$ partitioned into black B and white W colours

as follows: black vertices correspond to pairs of pants in $S \setminus P$, white vertices correspond to simple closed curves in P , and each black vertex is connected to the (not necessarily distinct) white vertices corresponding to the essential boundary components of the pair of pants.

Definition 3.2 (Projection of a path to the adjacency graph). Let S , P , and $\Gamma = \Gamma(P)$ be as above. Let $\gamma: [0, 1] \rightarrow S$ be a smooth path such that

- (1) γ is transverse to P , and
- (2) γ has no backtrack with respect to the pants curves in P , i.e. there is no embedded bigon D in S with $\partial D = \alpha \cup \beta$ where α is a subarc of γ and $\beta \subset P$ is an arc.

The *projection of γ to Γ* is a path defined as follows. Each component δ of $\gamma \setminus (\gamma \cap P)$ is an arc lying in a pair of pants P_δ such that one or both of its endpoints lie on ∂P_δ . If both endpoints of δ lie on ∂P_δ , then the projection of δ is the path of length 2 in Γ joining the white vertices corresponding to ∂P_δ . If δ is an arc with one endpoint on ∂P_δ and one endpoint in the interior of P_δ then the projection of δ is defined as the path of length 1 joining the white vertex corresponding to $\delta \cap \partial P_\delta$ to the black vertex corresponding to P_δ . Finally the projection of γ is obtained by concatenating the projections of components δ of $\gamma \setminus (\gamma \cap P)$. Note that the projection of γ is a path in the graph Γ , which does not depend on the parametrisation of γ .

Definition 3.3 (A nerve for a pants decomposition). Let S , P , and $\Gamma = \Gamma(P)$ be as above. Pick a base point $b \in S \setminus P$, and let b_0 be the black vertex of Γ corresponding to the pair of pants containing b . Let T be a maximal tree (also called spanning tree) for Γ . Denote the white vertices of Γ by w_i where $1 \leq i \leq m$, and let $\alpha_i \in P$ be the simple closed curve corresponding to w_i . For every vertex w_i , let p_i be the unique shortest path in T from b_0 to w_i .

A *nerve of (P, b) compatible with T* is a collection $\gamma = \cup \gamma_i$ of simple arcs $\gamma_i \subset S$ such that

- (1) γ_i starts at b and ends on α_i ;
- (2) γ_i is transverse to P and its projection to Γ is the path p_i ; and
- (3) if the last white vertex of the path $p_i \cap p_j$ is w_ℓ , then $\gamma_i \cap \gamma_j = \gamma_\ell$. In the special case that $p_i \cap p_j = \{b_0\}$, then $\gamma_i \cap \gamma_j = \{b\}$.

A *nerve of (P, b)* is a nerve of (P, b) compatible with some maximal tree $T \subset \Gamma(P)$.

Example 3.4. Let P be the pants decomposition of a surface of genus two as in Figure 7 top-left. Then the adjacency graph $\Gamma(P)$ is as in the top-right of the same figure. Let b be a base point lying in the interior of the left pair of pants, and T be the maximal tree as in bottom-right hand side of the figure. Then a nerve for (P, b) compatible with T is shown in the bottom-left side of the figure.

Lemma 3.5. Given S , P , T , and b as above, there is a nerve for (P, b) compatible with T .

Proof. Let d_i be the length (number of edges) of the path p_i . Define the arcs γ_i inductively based on the length d_i . The base case is when $d_i = 1$. There are either 1, 2, or 3 paths p_i with $d_i = 1$, and the corresponding curves α_i will be boundary components of the pair of pants P_0 containing b_0 . In this case we take γ_i to be simple arcs joining b to the corresponding curves in ∂P_0 that intersect only in b . Now assume that γ_i is not defined yet, and d_i is minimal with this property. Let w_j be the white vertex that is distance 2 from w_i along p_i . Let b_k be the common neighbour of w_i and w_j in p_i , and w_r ($r \neq i, j$) be the third neighbour of b_k . Then γ_j is defined by hypothesis. Note that $\gamma_j \subset \gamma_i$ should hold. Similarly if $p_j \subset p_r$, then $\gamma_j \subset \gamma_r$ should hold. Now if $p_j \subset p_r$ then we define $\gamma_i \setminus \gamma_j$ and $\gamma_r \setminus \gamma_j$ simultaneously such that $\gamma_i \cap \gamma_r = \gamma_j$ as in Figure 8 right, and otherwise define $\gamma_i \setminus \gamma_j$ as in Figure 8 left. $\square$




FIGURE 7. A pants decomposition P (top-left), its adjacency graph (top-right), a maximal tree T for the adjacency graph (bottom-right), a base point b and a nerve for (P, b) compatible with T (bottom-left).




FIGURE 8. Constructing a nerve

Remark 3.6. Assume that a triangulation $\mathcal{H}$ of S is given such that P is contained in the 1-skeleton of $\mathcal{H}$. Then we can construct a nerve γ for (P, b) compatible with T such that each arc γ_i passes through each face of $\mathcal{H}$ at most once, and it is disjoint from the vertices of the triangulation. To see this let q_i be a shortest path in the dual graph of $\mathcal{H}$ connecting $\gamma_j \cap \alpha_j$ to α_i in the above proof. Now if $p_j \not\subset p_r$, then define $\gamma_i \setminus \gamma_j$ such that its projection to the dual graph of $\mathcal{H}$ is equal to q_i . In particular, $\gamma_i \setminus \gamma_j$ passes through each face of $\mathcal{H}$ at most once. If $p_j \subset p_r$, define q_r similar to p_i . Then after possibly replacing q_r by another path of the same length and the same endpoints, q_i and q_r intersect each other in a connected interval; this follows from the shortest length hypothesis. Hence we can define $\gamma_i \setminus \gamma_j$ and $\gamma_r \setminus \gamma_j$ simultaneously such that their projections to the dual graph of $\mathcal{H}$ gives respectively q_i and q_r , and they only intersect in a single point.

A pants decomposition is a pre-triangulation as well. The next construction shows how to produce a *connected* pre-triangulation from a pants decomposition equipped with a nerve.

Construction 3.7 (From a pants decomposition to a connected pre-triangulation). *Let S be a compact connected orientable surface, and P be a pants decomposition of S . Let T be a maximal tree in the bipartite adjacency graph $\Gamma = \Gamma(P)$, and $b \in S \setminus P$ be a base point. Pick a nerve γ for (P, b) compatible with T . We construct a connected pre-triangulation $\mathcal{T} = \mathbf{T}(P, \gamma)$ of S . Assume the notations of Definitions 3.1 and 3.3. The only vertex of $\mathcal{T}$ is b . Moreover,*




FIGURE 9. From a pants decomposition to a connected pre-triangulation




FIGURE 10. The pants decomposition associated with the nerve shown in Figure 7 bottom-left.

$\mathcal{T}$ has an edge e_i for each curve α_i in the pants decomposition, constructed inductively as follows

- (1) First, for every white vertex w_i adjacent to b_0 , add an edge e_i that runs from b to α_i and remains close and almost parallel to γ_i , then goes around α_i , and finally comes back to b close and almost parallel to γ_i .
- (2) From now on at each step one or two edges are added. Pick a white vertex w_i for which γ_i is not defined yet and such that the distance between w_i and b_0 in T is minimal. Consider the shortest path p_i in T joining b_0 to w_i , and let w_j be the white vertex along p_i that is closest to w_i ($j \neq i$). By hypothesis, the edge e_j is already defined. Let b_k be the common neighbour of w_i and w_j in p_i , and let w_r ($r \neq i, j$) be the other neighbour of b_k in $\Gamma(P)$. If the edge joining b_k to w_r appears in p_r , then add e_i and e_r as in Figure 9 Right; otherwise add e_i as in Figure 9 Left. Again e_i runs from b to α_i remaining almost parallel to γ_i , then goes around α_i , and finally comes back to b almost parallel to γ_i .

Intuitively, e_i is obtained by pulling α_i to b along the path γ_i ; the inductive process for constructing e_i is to make sure that the edges e_i only intersect each other at the common vertex b .

Lemma 3.8. $\mathbf{T}(P, \gamma)$ is a pre-triangulation of S .

Proof. The nerve γ is an embedded tree in S , and there is a collapsing map $r: S \rightarrow S$ that collapses γ to the base point b , and the image of $P \cup \gamma$ under r is the 1-complex $\mathbf{T}(P, \gamma)$. Each complementary component of $\mathbf{T}(P, \gamma)$ in S is homeomorphic to a corresponding complementary component of $P \cup \gamma$ in S . Since $P \cup \gamma$ is a pre-triangulation, so is $\mathbf{T}(P, \gamma)$. $\square$

Example 3.9. Let P be the pants decomposition as in Figure 7 top-left, and γ be the nerve as in Figure 7 bottom-left. Then the pre-triangulation $\mathbf{T}(P, \gamma)$ is shown in Figure 10.




FIGURE 11. A pre-triangulation $\mathcal{T}$ with an ordering $O = \{e_1, e_2, e_3\}$ on its edges. The associated pants decomposition $\mathbf{P}(\mathcal{T}, O)$ is as in Figure 7 top-left.

3.2. From pre-triangulations to pants decompositions.

Construction 3.10 (From a pre-triangulation to a pants decomposition). *Let $\mathcal{T}$ be a pre-triangulation of a compact connected orientable surface S . Pick a forest T in the 1-skeleton $\mathcal{T}^1$ of $\mathcal{T}$ such that each component of T has at most one point of intersection with ∂S , and choose an ordering O on the edges of $\mathcal{T} \setminus (T \cup \partial S)$ as $e_1, \dots, e_\ell$. First define the subsurfaces N_i inductively: N_1 is a regular neighbourhood of $T \cup \partial S \cup e_1$ in S , and N_{i+1} is a regular neighbourhood of $N_i \cup e_{i+1}$ for each $1 \leq i < \ell$. Let F_i be the subsurface obtained from N_i by first capping off any boundary component that bounds a disc in $S \setminus N_i$ and then deleting any disc components. In particular, F_i is isotopic to the subsurface that $e_1 \cup \dots \cup e_i$ fills minus disc components. Define the multicurve P as the union of ∂F_i for $1 \leq i \leq \ell$ after discarding any boundary parallel or repeated simple closed curves. We will show in Lemma 3.13 that $P = \mathbf{P}(\mathcal{T}, T, O)$ is a pants decomposition. If $T = \emptyset$, we abbreviate $\mathbf{P}(\mathcal{T}, T, O)$ to $\mathbf{P}(\mathcal{T}, O)$.*

Example 3.11. *Let $\mathcal{T}$ be the pre-triangulation shown in Figure 11, and $O = \{e_1, e_2, e_3\}$ be the ordering on its edges. Then the associated pants decomposition $\mathbf{P}(\mathcal{T}, O)$ is the one in Figure 7 top-left.*

A connected embedded subsurface $F \subset S$ is called *essential* if it is π_1 -injective. An embedded subsurface is *essential* if each of its connected components is essential. When $F \subset S$ is an annulus or a pair of pants, then it is easy to see that F is essential if and only if each curve in ∂F is π_1 -injective in S .

Lemma 3.12. *Each surface F_j is obtained from F_{j-1} by thickening the boundary components to the outside and then either attaching a 1-handle, filling in a complementary annulus, or taking a disjoint union with at most one essential annulus ($2 \leq j \leq \ell$). In particular $F_j - F_{j-1}$ is a disjoint union of essential annuli and at most one essential pair of pants, and*

$$\chi(F_j - F_{j-1}) = \chi(F_j) - \chi(F_{j-1}) \geq -1.$$

Proof. The surface F_{j-1} is obtained by taking a regular neighbourhood N_{j-1} of $\partial S \cup T \cup e_1 \cup \dots \cup e_{j-1}$, capping off any complementary discs, and then discarding any disc components. Consider three cases:

- (1) The edge e_j is a loop disjoint from $\partial S \cup T \cup e_1 \cup \dots \cup e_{j-1}$. In this case N_j is the disjoint union of a regular neighbourhood of N_{j-1} and an annulus. If the core curve of the annulus is essential then F_j is the disjoint union of a regular neighbourhood of F_{j-1} and an essential annulus, and otherwise F_j is equal to a regular neighbourhood of F_{j-1} . Therefore $F_j - F_{j-1}$ is a disjoint union of essential annuli.
- (2) The edge e_j has only one endpoint on $\partial S \cup T \cup e_1 \cup \dots \cup e_{j-1}$. In this case N_j is a regular neighbourhood of N_{j-1} , and F_j is a regular neighbourhood of F_{j-1} .

- (3) The edge e_j has both endpoints on $\partial S \cup T \cup e_1 \cup \dots \cup e_{j-1}$. Then N_j is obtained by taking a regular neighbourhood of N_{j-1} and then attaching a band. In particular, $N_j - N_{j-1}$ is a union of annuli and one pair of pants. The pair of pants could be attached to N_{j-1} along either one or two of its boundary components, and N_j will have respectively either two or one new boundary components compared to N_{j-1} . Denote the boundary components of the attached pair of pants R by $\{\alpha, \beta, \gamma\}$. Now the lemma follows by considering various cases depending on the subset of $\{\alpha, \beta, \gamma\}$ along which R is attached to N_{j-1} , and the subset of $\{\alpha, \beta, \gamma\}$ that are homotopically trivial in S . For example if R is attached to N_{j-1} along α and β , and furthermore α and β are homotopically essential but γ is homotopically trivial, then the disc D bounding γ cannot contain α and β in its interior, and in this case F_j is obtained from F_{j-1} by attaching the annulus $R \cup D$. The remaining cases are similar and we leave them to the reader. $\square$

Lemma 3.13. $P = P(\mathcal{T}, T, O)$ is a pants decomposition.

Proof. Assume that α is an essential simple closed curve in $S \setminus P$. We need to show that α is isotopic to a curve in P . Since $\mathcal{T}$ is a pre-triangulation, F_ℓ is the entire surface S minus a (possibly empty) union of disjoint essential annuli and pairs of pants. Setting $F_0 = \emptyset$ and $F_{\ell+1} = S$, the surface S can be written as the disjoint union of $F_{i+1} - F_i$ for $0 \leq i \leq \ell$. Since α can be isotoped to be disjoint from $P = \bigcup_{i=1}^\ell \partial F_i$, there is an index i such that $\alpha \subset F_{i+1} - F_i$. However, by Lemma 3.13 every essential simple closed curve in $F_{i+1} - F_i$ is parallel to a component of ∂F_{i+1} or ∂F_i where $1 \leq i < \ell$. Moreover, since we assume that every component of T has at most one point of intersection with ∂S , it follows that F_1 is a union of annuli and at most one pair of pants and $\partial S \subset \partial F_1$. This implies that α is isotopic to a curve in $P = \bigcup \partial F_i$. $\square$

Remark 3.14. Let $O = (e_1, \dots, e_\ell)$ be an ordering on the edges of $\mathcal{T} \setminus T$. Let $e'_1, \dots, e'_m$ be any ordering on the edges of T . Define the ordering O' on the edges of $\mathcal{T}$ as $O' = (e'_1, \dots, e'_m, e_1, \dots, e_\ell)$. Then the pants decompositions $P(\mathcal{T}, T, O)$ and $P(\mathcal{T}, \emptyset, O)$ are equal.

Lemma 3.15 (Naturality). *Let S be a compact connected orientable surface, and P be a pants decomposition of S . Let $b \in S \setminus P$ be a base point, and γ be a nerve for (P, b) . Let $\mathcal{T}$ be a pre-triangulation obtained by possibly adding extra edges to the pre-triangulation $\mathbf{T}(P, \gamma)$. There is an ordering O on the edges of $\mathcal{T}$ such that $P(\mathcal{T}, O) = P$.*

Proof. A pants decomposition of S is a maximal collection of disjoint non-isotopic essential simple closed curves on S . Therefore it is enough to show that for a suitable choice of an ordering O we have $P \subset P(\mathcal{T}, O)$. Choose O such that the edges of $\mathbf{T}(P, \gamma)$ appear in the same order that they were constructed in the inductive recipe of Construction 3.7; if two (or three) edges were constructed simultaneously, the two (or three) edges appear consecutively in the ordering O . Finally, any edge of $\mathcal{T} \setminus \mathbf{T}(P, \gamma)$ appears after the edges of $\mathbf{T}(P, \gamma)$ in an arbitrary fashion. One can inductively see that $P \subset P(\mathcal{T}, O)$. $\square$

3.3. Modifying a pre-triangulation and its effect on the associated pants decomposition. In this subsection, we consider certain operations on pre-triangulations and their effect on the associated pants decompositions.

Lemma 3.16 (Effect of subdividing a pre-triangulation on the associated pants decomposition). *Let $\mathcal{T}$ be a pre-triangulation of a compact orientable surface S , and O be an ordering on the edges of $\mathcal{T}$. Consider any of the following two cases:*

- a) *Let $\mathcal{T}'$ be a pre-triangulation obtained by subdividing an edge e of $\mathcal{T}$. Define an ordering on the edges of $\mathcal{T}'$ as follows: if e appears as the i th element in the ordering $O = (e_1, e_2, \dots, e_n)$ and $e_i^1, \dots, e_i^k$ are the edges obtained by subdividing e_i , define $O' = (e_1, \dots, e_{i-1}, e_i^1, \dots, e_i^k, e_{i+1}, \dots, e_n)$.*
- b) *Let R be a complementary region to $\mathcal{T}$, and $\mathcal{T}'$ be the pre-triangulation obtained by adding a vertex v in R and connecting some of the vertices of R to v by new edges. Define an ordering on the edges of $\mathcal{T}'$ where the edges of $\mathcal{T}$ appear first and the newly added edges appear at the end in any order.*

Then the associated pants decompositions $\mathbf{P}(\mathcal{T}, O)$ and $\mathbf{P}(\mathcal{T}', O')$ are equal.

Proof. Part a) is clear. To see Part b), let $O' = (e_1, \dots, e_\ell, e'_1, \dots, e'_m)$ where e'_j are the newly added edges. Let N_i and F_i for $1 \leq i \leq \ell + m$ be the subsurfaces associated with $(\mathcal{T}', O')$ as in Construction 3.10. Part b) follows from the following three facts:

- i) $\mathbf{P}(\mathcal{T}, O)$ is obtained from $\cup_{i=1}^\ell \partial F_i$ by discarding repeated or boundary parallel curves. In particular the inclusion $\mathbf{P}(\mathcal{T}, O) \subset \mathbf{P}(\mathcal{T}', O')$ holds.
- ii) $\partial R \subset \partial N_\ell$.
- iii) Every essential simple closed curve in R is parallel to a component of ∂R . In particular, for $i > \ell$, the subsurface ∂F_i contains no essential curve that was not already present in ∂F_ℓ .

□

3.4. Changing the ordering and its effect on the associated pants decomposition.

Definition 3.17 (Consecutive transposition). Let $\mathcal{T}$ be a pre-triangulation of a compact orientable surface, and O and O' be two orderings on the edges of $\mathcal{T}$. We say that O' is obtained from O by a *consecutive transposition* if there are edges e_1 and e_2 of $\mathcal{T}$ that are consecutive in the ordering O and such that O' is obtained from O by swapping the order of e_1 and e_2 .

Lemma 3.18. *Let $\mathcal{T}$ be a pre-triangulation of a compact orientable surface S . Let O be an ordering on the edges of $\mathcal{T}$, and assume that O' is obtained from O by a consecutive transposition. The distance between the pants decompositions $\mathbf{P}(\mathcal{T}, O)$ and $\mathbf{P}(\mathcal{T}, O')$ in the pants graph is $O(1)$.*

Proof. Let $N_1, \dots, N_n$ (respectively $N'_1, \dots, N'_n$) be the sequence of surfaces associated with $(\mathcal{T}, O)$ (respectively $(\mathcal{T}, O')$) as in Construction 3.10. Define $F_1, \dots, F_n$ and $F'_1, \dots, F'_n$ similarly. Assume that O' is obtained from O by swapping the j -th and $(j+1)$ -th element. Then

$$F_k = F'_k \text{ if } |k - j| \geq 1.$$

Hence the pants decompositions $\mathbf{P}(\mathcal{T}, O')$ and $\mathbf{P}(\mathcal{T}, O)$ can differ only in curves that are supported on the subsurface $F_{j+1} \setminus F_{j-1}$. By Lemma 3.12, after discarding any disc or annulus components, $F_{j+1} \setminus F_{j-1}$ is an essential subsurface that has Euler characteristic at least -2 . Moreover, at most one component of $F_{j+1} \setminus F_{j-1}$ is not a pair of pants or annulus. Hence, after discarding all disc and annulus components, $F_{j+1} \setminus F_{j-1}$ is either an essential 4-times punctured sphere, an essential twice punctured torus, an essential once-punctured torus, or a disjoint union of at most two essential pairs of pants. The restrictions of $\mathbf{P}(\mathcal{T}, O')$ and




FIGURE 12. A schematic picture showing that the multicurves ∂N_j and ∂N_{j+1} (in red and blue respectively, and in different dashed lines) intersect in at most 4 points (solid dots). Here the horizontal rectangle is $N' \setminus N$, and the area below this rectangle is N . There are other possibilities: the bands are allowed to have twists, and the endpoints of the bands might not interlace or they might lie on different components of ∂N .

$\mathbf{P}(\mathcal{T}, \mathcal{O})$ to the subsurface $F_{j+1} \setminus F_{j-1}$, after discarding boundary-parallel curves, form two pants decompositions P and P' of $F_{j+1} \setminus F_{j-1}$. Since $F_{j+1} \setminus F_{j-1}$ is an essential subsurface of S , and the restrictions of $\mathbf{P}(\mathcal{T}, \mathcal{O}')$ and $\mathbf{P}(\mathcal{T}, \mathcal{O})$ to the exterior of $F_{j+1} \setminus F_{j-1}$ are equal, the distance between $\mathbf{P}(\mathcal{T}, \mathcal{O}')$ and $\mathbf{P}(\mathcal{T}, \mathcal{O})$ in the pants graph of S is bounded from above by the distance between P and P' in the pants graph of $F_{j+1} \setminus F_{j-1}$. We would like to give an upper bound for the distance between P and P' in the pants graph of $F_{j+1} \setminus F_{j-1}$. We claim that

$$i(\partial N_j, \partial N'_j) \leq 4.$$

If at least one of e_j and e_{j+1} is a simple closed curve, or at least one end of e_j or e_{j+1} is disjoint from $\partial S \cup e_1 \cup \dots \cup e_{j-1}$, then ∂N_j and $\partial N'_j$ can be isotoped to be disjoint. So assume that e_j and e_{j+1} have both ends on $\partial S \cup e_1 \cup \dots \cup e_{j-1}$. Choose regular neighbourhoods N and N' of $\partial S \cup e_1 \cup \dots \cup e_{j-1}$ in S such that $N \subset \text{int} N'$ and $N' \setminus N \cong \partial N \times [0, 1]$. Take a band $B_j \cong [0, 1] \times [0, 1]$ around e_j and connecting ∂N to itself such that

- $B_j \cap \partial N = [0, 1] \times \{0, 1\}$; and
- $B_j \cap (\overline{N' \setminus N})$ is identified with $[0, 1] \times ([0, \frac{1}{4}] \cup [\frac{3}{4}, 1])$ in B_j and with $(B_j \cap \partial N) \times [0, 1]$ in $N' \setminus N$.

Choose a band B_{j+1} around e_{j+1} , and disjoint from B_j , that connects $\partial N'$ to itself. Then ∂N_j is isotopic to $\partial(N \cup B_j)$, and ∂N_{j+1} is isotopic to $\partial(N' \cup B_{j+1})$. The only intersections of ∂N_j and ∂N_{j+1} come from the 4 intersection points between $\partial N'$ and ∂B_j , proving the desired inequality. See Figure 12. Hence

$$i(\partial F_j, \partial F'_j) \leq i(\partial N_j, \partial N'_j) \leq 4.$$

Therefore, the pants decompositions P and P' have distance $O(1)$ in the pants graph of $F_{j+1} \setminus F_{j-1}$, proving the lemma. $\square$

Corollary 3.19. *Let $\mathcal{T}$ be a pre-triangulation of a compact orientable surface S . Let $\mathcal{O}$ and $\mathcal{O}'$ be two orderings on the edges of $\mathcal{T}$. The distance between the pants decompositions $\mathbf{P}(\mathcal{T}, \mathcal{O})$ and $\mathbf{P}(\mathcal{T}, \mathcal{O}')$ in the pants graph is $O(n^2)$ where n is the number of edges of $\mathcal{T}$.*

Proof. If S_n is the symmetric group on n elements, then every permutation $\sigma \in S_n$ can be written as a product of at most

$$1 + 2 + \dots + (n - 1) = O(n^2)$$

consecutive transpositions of the form $\sigma_j = (j, j+1)$, where $1 \leq j \leq n-1$. Here σ_j swaps the j -th and $(j+1)$ -th elements. This can be seen by induction on n : assume $\sigma(n) = i$ and note that $\sigma_{n-1} \circ \dots \circ \sigma_{i+1} \circ \sigma_i \circ \sigma$ fixes n and hence can be considered as an element of S_{n-1} .

We can identify the symmetric group S_n with the group of permutations of the edges of $\mathcal{T}$ such that O corresponds to the identity. Hence there are orderings $O_1, \dots, O_k$ on the edges of $\mathcal{T}$ such that

- (1) $O = O_1, O' = O_k$;
- (2) each O_{i+1} is obtained from O_i by a single consecutive transposition; and
- (3) $k = O(n^2)$.

By Lemma 3.18, for each $1 \leq i < k$ the pants decompositions $\mathbf{P}(\mathcal{T}, O_{i+1})$ and $\mathbf{P}(\mathcal{T}, O_i)$ have distance $O(1)$ in the pants graph. The result now follows from the triangle inequality. $\square$

3.5. Upper bound for distance between pants decompositions coming from pre-triangulations. Given two pre-triangulations, and two orderings on their edges, the next lemma gives an upper bound for the distance between their associated pants decompositions.

Lemma 3.20. *Let $\mathcal{T}$ and $\mathcal{T}'$ be pre-triangulations of a compact orientable surface S . Assume that for any edge e of $\mathcal{T}$ and any edge e' of $\mathcal{T}'$, the intersection $e \cap e'$ is a disjoint union of subintervals of e and also of e' . Here we allow a subinterval to be a single point as well. Let O and O' be orderings on the edges of $\mathcal{T}$ and $\mathcal{T}'$. Then the pants decompositions $\mathbf{P}(\mathcal{T}, O)$ and $\mathbf{P}(\mathcal{T}', O')$ have distance at most $O(n^2)$, where n is the number of edges of the pre-triangulation obtained by superimposing $\mathcal{T}$ and $\mathcal{T}'$. In particular, if $\mathcal{T}$ and $\mathcal{T}'$ are transverse to each other, and $|\mathcal{T}|$ and $|\mathcal{T}'|$ are the number of edges of $\mathcal{T}$ and $\mathcal{T}'$, then $n \leq 2i(\mathcal{T}, \mathcal{T}') + |\mathcal{T}| + |\mathcal{T}'|$.*

Proof. Let $\mathcal{U}$ be the pre-triangulation obtained by superimposing $\mathcal{T}$ and $\mathcal{T}'$. Construct an ordering O_2 on the edges of $\mathcal{U}$ as follows.

- (1) First subdivide the edges of $\mathcal{T}$ by introducing the new vertices that are in the intersection of $\mathcal{T}$ and $\mathcal{U}$. Let $\mathcal{T}_1$ be the new pre-triangulation, and O_1 a new ordering on the edges of $\mathcal{T}_1$ constructed by repeated application of Lemma 3.16 a).
- (2) Secondly, starting with $(\mathcal{T}_1, O_1)$, add the vertices of $\mathcal{U}$ that are disjoint from $\mathcal{T}_1$ together with their edges in $\mathcal{U}$ to obtain $\mathcal{U}$. Let O_2 be an ordering on the edges of $\mathcal{U}$ constructed by repeated application of Lemma 3.16 b).

By Lemma 3.16, the pants decompositions $\mathbf{P}(\mathcal{T}, O)$ and $\mathbf{P}(\mathcal{U}, O_2)$ are equal. Similarly, there is an ordering O'_2 on the edges of $\mathcal{U}$ such that the pants decompositions $\mathbf{P}(\mathcal{T}', O')$ and $\mathbf{P}(\mathcal{U}, O'_2)$ are equal. Now by Lemma 3.19 the distance between the pants decompositions $\mathbf{P}(\mathcal{U}, O_2)$ and $\mathbf{P}(\mathcal{U}, O'_2)$ is $O(n^2)$, proving the first part of the lemma. For the second part we have

$$\begin{aligned} 2n &= \text{sum of the degrees of vertices in } \mathcal{U} \\ &\leq 4i(\mathcal{T}, \mathcal{T}') + \text{sum of degrees in } \mathcal{T} + \text{sum of degrees in } \mathcal{T}' \\ &= 4i(\mathcal{T}, \mathcal{T}') + 2|\mathcal{T}| + 2|\mathcal{T}'|. \end{aligned}$$

$\square$

Corollary 3.21. *Let P and P' be pants decompositions of a compact orientable surface S . The distance between P and P' in the pants graph is $O(i(P, P')^2)$.*

Proof. Let $P_\cap$ be the union of simple closed curves that appear in both P and P' . Let S_1 be the surface obtained by cutting S along $P_\cap$, and let P_1 and P'_1 be the corresponding pants decompositions of S_1 . Hence P_1 is the image of $P \setminus P_\cap$ in S_1 , and P'_1 is defined similarly.

The pants decompositions P_1 and P'_1 are (disconnected) pre-triangulations of S_1 . Isotope P'_1 to intersect P_1 minimally. Note that for any ordering O on P_1 we have $\mathbf{P}(P_1, O) = P_1$; and similarly $\mathbf{P}(P'_1, O') = P'_1$. Hence the distance between the pants decompositions $P_1 = \mathbf{P}(P_1, O)$ and $P'_1 = \mathbf{P}(P'_1, O')$ is $O(n^2)$, where n is the number of edges of the pre-triangulation $P_1 \cup P'_1$ obtained by superimposing P_1 and P'_1 . In particular, using the degree sum formula for the 4-regular graph $P_1 \cup P'_1$ we have

$$n = \text{number of edges of } P_1 \cup P'_1 = \frac{1}{2} \text{ degree sum} = 2i(P_1, P'_1) = 2i(P, P').$$

□

4. TRAIN TRACKS, AND AGOL–HASS–THURSTON ALGORITHM

Agol, Hass, and Thurston used their algorithm to count the number of connected components of a simple closed multicurve γ carried by a train track τ , in time that is bounded above by a polynomial function of $E \log(N)$, where E is the number of branches of τ and N is the total weight of γ with respect to τ . We are interested in the way that the algorithm does this count; this will be via a sequence of *splitting* and *twirling* (Definition 4.6, and Figure 18) the weighted train track until it completely unwinds to the original curve γ . We will apply the AHT algorithm to another similar setting and show in Proposition 4.13 that if one starts with say a one-vertex triangulation $\mathcal{T}$ carried by a *based train track* τ (Definition 4.1), and set up an appropriate orbit counting problem (Definition 4.4), then the AHT algorithm does the count via splitting and twirling the based weighted train track until it completely unwinds it to $\mathcal{T}$. Moreover the number of splitting and twirling is bounded above by a polynomial function of $E \log(N)$, with E equal to the number of branches of the based train track and N equal to the total weight of τ .

4.1. Based integrally weighted train tracks, and orbit counting problems. Train tracks were introduced by Thurston to study simple closed multicurves on a surface. Here we will work with 1-complexes, and a variant of a train track, which we call a based train track, will be useful.

Definition 4.1 (Based integrally weighted train track). A *based train track* (τ, V) in a surface S is an embedded finite 1-complex with a distinguished subset V of its vertices called *base vertices* such that

- the embedding is C^1 in the interior of each edge of τ ;
- at every vertex $w \notin V$ there is a well-defined tangent line, and there are edges entering w from both directions. See Figure 13, top-left.

An edge of τ is called a *branch*. Fix a point p_e in the interior of each branch e . Each component of $e \setminus p_e$ is called a *half-branch*. Therefore the half-branches at a vertex $w \notin V$ can be partitioned into two non-empty subsets, which we can locally think of as incoming vs outgoing; the choice of which subset is incoming and which is outgoing is arbitrary. An *integral weight* μ on τ is an assignment of *positive integers* to branches of τ such that at each vertex $w \notin V$ the *switch condition* is satisfied; i.e. sum of the weights of incoming half-branches is equal to that of outgoing half-branches. We denote this common number by μ_w . A *based integrally weighted train track* (τ, V, μ) is a based train track (τ, V) together with an integral weight μ on its branches. When V is empty, we obtain the usual notions of a *train track* τ and an *integrally weighted train track* (τ, μ) .




FIGURE 13. Left: on top, a non-base vertex of the based integrally weighted train track (τ, V, μ) is shown; a base vertex is depicted in the bottom. The local picture for the associated 1-complex $\mathcal{CC}(\tau, V, \mu)$ is shown on the right.

Remark 4.2. Note that unlike some texts, we do not require the complementary regions to a train track to have negative (or non-positive) *index* (a variant of the Euler characteristic).

The following is analogous to the multicurve carried by an integrally weighted train track.

Definition 4.3 (The 1-complex carried by a based integrally weighted train track). Given a based integrally weighted train track (τ, V, μ) on a surface S , define the 1-complex $\mathcal{CC}(\tau, V, \mu)$ embedded in S as follows: given a branch e of τ , not adjacent to any vertex in V , replace e by μ_e parallel segments where μ_e is the weight of e . Given a branch e that is adjacent to a vertex $v \in V$, replace e with μ_e parallel segments and identify one ends of these segments with v . At each vertex $w \notin V$ of τ , glue the endpoints of the incoming segments adjacent to w to those of the outgoing segments in an order-preserving way. See Figure 13. We call $\mathcal{CC}(\tau, V, \mu)$ the 1-complex carried by (τ, V, μ) .

Given a based integrally weighted train track (τ, V, μ) on a surface S , define a *branched neighbourhood* $N(\tau, V, \mu)$ of (τ, V, μ) in S as follows: For each branch e , take a rectangle $R_e \cong e \times [0, \mu_e]$ around e equipped with the horizontal foliation whose leaves consist of $e \times \text{point}$. At each vertex $w \notin V$ glue the rectangles R_e adjacent to w along their vertical boundary $\partial e \times [0, \mu_e]$ according to their adjacency and the length of their vertical boundary; this is possible by the switch condition. See Figure 14. For any base vertex $v \in V$, let $\deg(v)$ be the number of half-branches adjacent to v . Take a $2\deg(v)$ -gon P_v around v and identify half of its edges with the vertical boundaries of the adjacent rectangles in an alternating fashion; see Figure 15. The *free sides* of P_v are those that are not identified with the vertical boundary components of adjacent rectangles. Note that $N(\tau, V, \mu) \setminus (\cup_{v \in V} P_v)$ comes equipped with a *horizontal foliation* obtained by gluing together the horizontal foliations of all rectangles R_e .

Each rectangle R_e has an I -bundle structure given by fibres $\text{point} \times [0, \mu_e]$. The *tie interval above a point* in τ is defined as follows:

- For each point in the interior of an edge e , the *tie interval* above the point is the I -fibre above that, see the (blue) vertical solid lines in Figure 14.
- For a vertex $w \notin V$, the *tie interval above w* in $N(\tau, \mu)$ is the union of the I -fibres over w coming from the half-branches adjacent to w ; see Figure 14. Moreover, a point c on the tie interval above $w \notin V$ is called a *cusp* if a neighbourhood of c in the leaf of the horizontal foliation through it is not homeomorphic to $\mathbb{R}$; see Figure 14.
- For a base vertex $v \in V$, the *tie interval above v* is the disjoint union of I -fibres over v coming from the half-branches adjacent to v , see Figure 15.




FIGURE 14. A non-base vertex w of a based integrally weighted train track (τ, V, μ) on the left, and the local picture for the branched neighbourhood $N(\tau, V, \mu)$ on the right. The tie interval above w is shown by dashed lines, and the cusp points c_i are depicted as dots on it. Other tie intervals are shown with (blue) vertical solid lines.




FIGURE 15. A base vertex v of a based integrally weighted train track (τ, V, μ) on the left, and the local picture for the branched neighbourhood $N(\tau, V, \mu)$ on the right. The tie intervals are shown with solid arcs.

Note that (τ, V, μ) is determined uniquely, up to isotopy, by the branched neighbourhood $N(\tau, V, \mu)$.

Just as in splittings of (unweighted) train tracks, it makes sense to talk about a splitting of an (unweighted) based train track (τ, V) along a branch e . Let e be an oriented edge with the initial vertex $w \notin V$ and the terminal vertex w' . A splitting of (τ, V) is a combinatorial local move as in Figure 16.

A splitting of (τ, V, μ) is defined just like a splitting of a weighted train track except we require the support of the splitting to be disjoint from V . Namely let $N(\tau, V, \mu)$ be a branched neighbourhood of (τ, V, μ) . Consider a vertex $w \notin V$, and the tie interval t above w . Let p be a cusp point of t , and l be the leaf of the horizontal foliation on $N(\tau, V, \mu) \setminus (\cup_{v \in V} P_v)$ passing through p and oriented starting at p . Denote by e the first branch of τ travelled by l , and give e the orientation from l . If $l \neq p$, let q be the next point of intersection of l with the union of the tie intervals above the vertices of τ , it could be that $q = p$ if l is a simple closed curve. Denote the segment of l between p and q by $l_0 = [p, q)$ where l_0 contains p but not q . If $l = p$, set $l_0 = p$. Then the *splitting of $N(\tau, V, \mu)$ along the cusp point p (or the leaf segment l_0)* is defined as the closed complement $N(\tau, V, \mu) \setminus l_0$ with one exception: if q lies on a tie interval above a base vertex $v \in V$ then we split $N(\tau, V, \mu)$ along l_0 and then modify the polygon P_v by adding a free side around q ; see Figures 16 and 17. Sometimes we refer to this operation as a splitting of $N(\tau, V, \mu)$ along the branch e , if we do not need to stress the choice of the cusp point p . We say that (τ', V, μ') is obtained by *splitting* (τ, V, μ) if $N(\tau', V, \mu')$ is obtained by splitting $N(\tau, V, \mu)$; note that the set of base vertices V is not changed during the splitting.

Definition 4.4 (Orbit counting problem associated with a based integrally weighted train track). Let (τ, V, μ) be a based integrally weighted train track. For each vertex $w \notin V$, define μ_w as sum of the weights of the incoming (or outgoing) edges at w , and pick an interval




FIGURE 16. Splitting an unweighted train track along an edge or a vertex. The picture on the left is replaced with the one on the right, after splitting along the middle vertex or edge.




FIGURE 17. Splitting of a branched neighbourhood. The shaded regions are the rectangles over various edges. The dashed line is the leaf segment $l = [p, q)$, where q lies on P_v here.

$t_w \subset \mathbb{N}$ of length μ_w . Visually we think of t_w as the intersection of the tie interval above w with the carried 1-complex $\mathcal{CC}(\tau, V, \mu)$. Similarly, let $e_1, \dots, e_r$ be the half-branches adjacent to base vertices, and for each e_j consider an interval $t_j \subset \mathbb{N}$ of length μ_{e_j} . Again we think of t_j as the intersection of a component of the tie interval above a base vertex v with the 1-complex $\mathcal{CC}(\tau, V, \mu)$. By abuse of notation, we call the intervals t_w and t_j the tie intervals. Pick intervals $I_j \subset \mathbb{N}$ where the index j varies in $\mathcal{J} := \{1, \dots, r\} \cup \{w | w \notin V \text{ is a vertex}\}$ such that

- (1) the interval I_j has the same length as t_j for $j \in \mathcal{J}$;
- (2) the intervals I_j for $j \in \mathcal{J}$ are pairwise disjoint and their union is $[1, N]$ for some positive integer N ; and
- (3) the union of I_j where $j \in \{1, \dots, r\}$ is $[1, M]$ for some positive integer M .

Pick isometric identifications $i_j: t_j \rightarrow I_j$; there are two such choices for each index j . The data $\mathcal{I} := (t_j, I_j, i_j)_{j \in \mathcal{J}}$ is called an *initial interval identification*. Define an orbit counting problem $\text{OCP}(\tau, V, \mu, \mathcal{I})$ as follows. Let e be an edge of τ connecting vertices w and w' , and assume that $I_{w'}$ lies to the right of I_w as subsets of $\mathbb{N}$. Define an isometric pairing g_e by following the segments of the 1-complex $\mathcal{CC}(\tau, V, \mu)$ along the rectangle R_e where $\text{domain}(g_e) \subset I_w$ and $\text{range}(g_e) \subset I_{w'}$. Then the isometric pairings $\{g_e | e \text{ is a branch}\}$ form the orbit counting problem $\text{OCP}(\tau, V, \mu, \mathcal{I})$ with total interval $[1, N]$.

Remark 4.5. By construction, every point in $[1, N]$ appears in the domain and range of at most two pairings of $\text{OCP}(\tau, V, \mu, \mathcal{I})$.

The following geometric operation on a based integrally weighted train track appears naturally when applying the AHT algorithm to the orbit counting problem $\text{OCP}(\tau, V, \mu, \mathcal{I})$.




FIGURE 18. A twirling. A neighbourhood of the branch e and the underlying train track is shown on the right. Here $k = 2$, $m = 2$, and $n = 1$. In the left side, the two A should be identified with each other to construct the rectangle R_e . The solid points from top to bottom are p_1^0 , p_2^0 , p_1^1 , p_2^1 , p_1^2 , p_2^2 and p_1^3 . The numbers **a**, **b**, and **c** are the weights of e and the two edges coming into it from right.

Definition 4.6 (Twirling). Let (τ, V, μ) be a based integrally weighted train track on an orientable surface S , and let $N = N(\tau, V, \mu)$ be a branched neighbourhood of it. Let $w \notin V$ be a vertex, and e be a branch connecting w to itself. Denote the tie interval above w by t , and fix an identification of t with the interval $[0, \mu_w]$ (there are two choices for such identification). Call $0 \in [0, \mu_w] \cong t$ the *top* point of t . Let $e \times [0, \mu_e] \cong R_e \subset N$ be the rectangle around e . Write $\partial e = \partial_- e \cup \partial_+ e$, and orient e from $\partial_- e$ to $\partial_+ e$. Call $\partial_- e \times [0, \mu_e]$ the negative vertical boundary of R_e , similarly define the positive vertical boundary. Assume that R_e starts and ends on different sides of t , and that the images of the negative and positive vertical boundary components of R_e in N intersect each other. (The operation of twirling is only defined under this assumption). Let the positive side of t be the side of $\partial_+ e$; similarly for the negative side. Assume that e is the top branch coming into w on the negative side of t . See Figure 18.

Let $p_1, \dots, p_k$ be the cusp points on t that lie above $\partial_+ e \times (0, \mu_e)$. Let l_i be the leaf of the horizontal foliation through p_i , and s_i be connected component of $l_i \cap R_e$ that contains p_i . Split N along all s_i . More precisely, set $p_i^0 = p_i$, and if $l_i \neq \{p_i^0\}$ then define p_i^1 to be the first intersection point of l_i with the positive vertical boundary of R_e ; and p_i^1 is not defined otherwise. Then by the orientability of S , the relative order of p_i^1 on t is the same as that of p_i ; i.e. p_i^1 is on top of p_j^1 if $i < j$ as long as both are defined. Now if all p_i^1 are defined and lie in the negative vertical boundary of R_e , we can repeat the same process starting with p_i^1 to obtain points p_i^2 and so on. Hence there is an m such that at least one of the points $p_1^m, \dots, p_k^m$ is either not defined or does not lie in the negative vertical boundary of R_e ; say only $p_1^m, \dots, p_n^m$ are defined and lie in the negative vertical boundary of R_e for some $0 \leq n < k$. Then we split one further time along $p_1^m, \dots, p_n^m$. The resulting based integrally weighted train track is called a *twirling* of (τ, V, μ) along the oriented branch e .

Remark 4.7. If (τ', V, μ') is obtained from (τ, V, μ) via a splitting, then the number of branches of τ' is no more than that of τ . Since every twirling is a concatenation of a number of splits, the same holds if (τ', V, μ') is obtained from (τ, V, μ) via a twirling.

Lemma 4.8 (Effect of twirling on the underlying train track). *Let (τ, V, μ) be a based integrally weighted train track on an orientable surface S . Assume that (τ', V, μ') is obtained from (τ, V, μ) by twirling along the branch e . Let α be the simple closed curve that is the closure*

of e , and m be as in Definition 4.6. Denote the Dehn twist along α by T_α . Then τ' is obtained from τ via twisting by $(T_\alpha)^{m-1}$ followed by splitting the (unweighted) based train track $(T_\alpha)^{m-1}(\tau, V)$ a number of times along e .

Proof. Define the points p_i^r as in Definition 4.6. Let l_i^r be the segment of l_i between p_i^r and p_i^{r+1} , if both p_i^r and p_i^{r+1} are defined. Then repeatedly splitting along $\cup_{i=1}^k l_i^r$ for $r = 0, 1, \dots, m-2$ replaces τ with $(T_\alpha)^{m-1}(\tau)$. Then splitting along $\cup_{i=1}^k p_i^{m-1}$ followed by splitting further along $\cup_{i=1}^m p_i^m$ is a sequence of splits along the oriented branch e . See Figure 18. $\square$

Lemma 4.9. *Let (τ, V) be a based train track on a compact orientable surface S , and that (τ', V) is obtained from (τ, V) by a splitting. Assume that the underlying 1-complex of τ' is a polygonal decomposition (respectively a pre-triangulation) of S . Then the underlying 1-complex of τ is also a polygonal decomposition (respectively a pre-triangulation). The corresponding statement holds if based train tracks are replaced with based integrally weighted train tracks.*

Proof. The Euler characteristics of the complementary components of a train track will not increase as a result of splitting, but they might decrease if some complementary components merge together. Now each complementary component of τ' has Euler characteristic at least 1 (respectively -1). It follows that the complementary components of τ' have Euler characteristic at least 1 (respectively -1) as well. It remains to show that no complementary component of τ is topologically a torus minus a disc. Observe that if τ' is obtained by splitting τ , and α is an essential (i.e. homotopically non-trivial and non-boundary-parallel) curve in a complementary component R of τ , then some complementary component of τ' also contains an essential curve. It follows that no complementary component R of τ is a torus minus a disc; otherwise R contains an essential curve which implies that some complementary component of τ' also contains an essential curve, a contradiction.

A splitting of a based integrally weighted train track is in particular a splitting of the underlying unweighted based train track. Therefore, the same proof applies to the weighted case as well. $\square$

Lemma 4.10. *Let (τ, V, μ) be a based integrally weighted train track with E edges on a compact orientable surface S . Assume that (τ', V, μ') is obtained from (τ, V, μ) by either a split or a twirl. Assume that the underlying 1-complex of τ is a pre-triangulation. Set $k = \chi(\tau') - \chi(\tau) \geq 0$, where $\chi(\tau)$ is the Euler characteristic of τ as a 1-complex. For every ordering O on the edges of τ , there is an ordering O' on the edges of τ' such that the pants decompositions $P(\tau, O)$ and $P(\tau', O')$ have distance $O((k+1)E)$ in the pants graph.*

Moreover, given the train track τ , the ordering O , and the split or twirl move from τ to τ' , there is an algorithm that constructs such an ordering O' together with a sequence of $O((k+1)E)$ consecutive transpositions taking O to O' . The algorithm runs in time that is a polynomial function of E . Here the train track is given as follows: the complementary components of τ are cusped polygons, annuli, or pairs of pants; these are given together with their side identifications. Furthermore, the twirl or split move is given by specifying an oriented edge.

Proof. We first show how to reduce to the case where τ' is obtained from τ by splitting a number of times along the same oriented edge. Assume that (τ', V, μ') is obtained from (τ, V, μ) by a twirl along a branch e , and let α be the simple closed curve that is the closure of e . By Lemma 4.8, τ' is obtained from τ by some power $(T_\alpha)^\ell$ of the Dehn twist along α followed by splitting $(T_\alpha)^\ell(\tau, V)$ a number of times along e . Given the ordering O on τ , we can apply at most E consecutive transpositions to obtain O_1 such that e appears as the first

edge in O_1 . The distance between the pants decompositions $\mathbf{P}(\tau, O)$ and $\mathbf{P}(\tau, O_1)$ is $O(E)$ by Lemma 3.18. Since $(T_\alpha)^\ell$ is a mapping class, it induces an ordering $((T_\alpha)^\ell)_*(O_1)$ on $(T_\alpha)^\ell(\tau)$. We claim that the pants decompositions $\mathbf{P}(\tau, O_1)$ and $\mathbf{P}((T_\alpha)^\ell(\tau), ((T_\alpha)^\ell)_*(O_1))$ are the same. To see this, note that e appears as the first edge in O_1 , and so Construction 3.10 produces α as the first curve in $\mathbf{P}(\tau, O_1)$. Hence

$$\mathbf{P}((T_\alpha)^\ell(\tau), ((T_\alpha)^\ell)_*(O_1)) = (T_\alpha)^\ell(\mathbf{P}(\tau, O_1)) = \mathbf{P}(\tau, O_1).$$

The last equality holds because Dehn twisting a pants decomposition along one of its curves does not change the pants decomposition. This shows that at the cost of distance $O(E)$ in the pants graph, we may replace (τ, O) with $((T_\alpha)^\ell(\tau), ((T_\alpha)^\ell)_*(O_1))$, completing the reduction.

Now assume that τ' is obtained from τ by splitting a number of times along the same oriented edge e . Again at the cost of distance E in the pants graph, we may assume that e appears as the first edge in O . Assume that $k \leq 1$; the general case follows from this. Since τ splits to τ' , the train track τ' is carried by τ . This means that there is a neighbourhood $N(\tau)$ of τ foliated by interval fibers, a projection map $\pi: N(\tau) \rightarrow \tau$ collapsing each fiber to a point, and a map $f: \tau' \rightarrow N(\tau)$ transverse to the fibers. In particular, the composition $\pi \circ f$ gives a map $\pi \circ f: \tau' \rightarrow \tau$, which we abbreviate to $\pi: \tau' \rightarrow \tau$. Consider the following cases:

Case 1): If $k = 0$ and e has distinct endpoints. Here $T := \pi^{-1}(e)$ is a tree. Contract e in τ to a point to obtain a pre-triangulation τ_1 . Let O_1 be the ordering on τ_1 induced from the ordering O on τ (by removing e). By Construction 3.10, since e appears as the first edge in O and the endpoints of e are distinct, we have

$$\mathbf{P}(\tau, O) = \mathbf{P}(\tau_1, O_1).$$

For each point in $\pi^{-1}(\partial e)$, add it as a vertex (possibly of degree two) in τ' . Define an ordering O' on τ' , where the restriction of O' to $\tau' - T$ agrees with the restriction of O to $\tau - e$ and such that the edges of τ' lying in T appear first (in any order) in O' . Contract the tree T in τ' to obtain τ'_1 , let O'_1 be the ordering on τ'_1 obtained from the ordering O' on τ' by deleting the edges that lie in T . We have

$$\mathbf{P}(\tau'_1, O'_1) = \mathbf{P}(\tau', O').$$

Note that (τ_1, O_1) is equal to (τ'_1, O'_1) up to isotopy. Hence

$$\mathbf{P}(\tau, O) = \mathbf{P}(\tau_1, O_1) = \mathbf{P}(\tau'_1, O'_1) = \mathbf{P}(\tau', O').$$

Case 2): If $k = 1$ and e has distinct endpoints. Here $\pi^{-1}(e)$ has two connected components, each of which is a tree. Contract e in τ to obtain (τ_1, O_1) . For each point in $\pi^{-1}(\partial e)$ that is not a vertex of τ' , add it as a vertex of degree two to obtain $\hat{\tau}'$. Attach an extra edge e' (joining the two components of T) to $\hat{\tau}'$ to obtain a 1-complex τ'_1 such that if we contract $e' \cup \pi^{-1}(e)$ in τ'_1 the result is isotopic to τ_1 . See Figure 19.

Let O'_1 be the ordering on τ'_1 such that the edges in $e' \cup \pi^{-1}(e)$ appear first, and otherwise it agrees with O_1 . We have

$$\mathbf{P}(\tau, O) = \mathbf{P}(\tau_1, O_1) = \mathbf{P}(\tau'_1, O'_1).$$

Let O'_2 be the ordering on τ'_1 such that e' appears last, and otherwise the ordering agrees with O'_1 . Then O'_2 can be obtained from O'_1 by $O(E)$ consecutive transpositions. Therefore, the distance between $\mathbf{P}(\tau'_1, O'_2)$ and $\mathbf{P}(\tau'_1, O'_1)$ is $O(E)$. Let $\hat{O}'$ be the ordering on $\hat{\tau}'$ obtained by removing e' (the last edge) from τ'_1 . We claim that $\mathbf{P}(\hat{\tau}', \hat{O}')$ is equal to $\mathbf{P}(\tau'_1, O'_2)$. This is because, e' appears last in O'_2 , and given that τ' and hence $\hat{\tau}'$ is a pre-triangulation, adding e to $\hat{\tau}'$ (to obtain τ'_1) does not create any new curve in Construction 3.10. Hence the distance




FIGURE 19. The train track τ' on the right is obtained by splitting τ (left). The added edge e' is shown with a dotted segment. The edge e and the preimage $\pi^{-1}(e)$ are shown in solid blue segments.

between $\mathbf{P}(\tau, \mathcal{O}) = \mathbf{P}(\tau'_1, \mathcal{O}'_1)$ and $\mathbf{P}(\hat{\tau}', \hat{\mathcal{O}}') = \mathbf{P}(\tau'_1, \mathcal{O}'_2)$ is $O(E)$. On the other hand, the ordering $\hat{\mathcal{O}}'$ on $\hat{\tau}'$ induces an ordering $\mathcal{O}'$ on τ' , simply by deleting the vertices $\pi^{-1}(e)$ and we have $\mathbf{P}(\hat{\tau}', \hat{\mathcal{O}}') = \mathbf{P}(\tau', \mathcal{O}')$.

Case 3: If $k \leq 1$ and the endpoints of e coincide. In this case we can add a vertex of degree 2 in the middle of e to obtain edges e_1 and e_2 such that $e_1 \cup e_2 = e$, each e_i has distinct endpoints, and the splittings along e can be obtained by splittings along e_1 followed by splittings along e_2 . This reduces the problem to the previous two cases.

It is straightforward to see that the above construction of $\mathcal{O}'$ can be done algorithmically in time that is a polynomial function of E , $|\chi(S)|$, and k . Since τ is a pre-triangulation, E is at least linear in $|\chi(S)|$. Similarly, k is at most E . Therefore the algorithm runs in polynomial time in E . □

4.2. Simplifying a based integrally weighted train track using the AHT algorithm.

Given a based integrally weighted train track (τ, V, μ) , one can always simplify it to its carried 1-complex using successive splittings. The number of splittings could be large, and one can instead simplify (τ, V, μ) much faster by using both splitting and twirling moves. This section gives an upper bound for the numbers of splitting and twirling moves needed to simplify a based integrally weighted train track to its carried 1-complex, using the algorithm of Agol, Hass, and Thurston. The arguments are technical and could be skipped in a first reading, but see Remark 4.16. Before giving the technical details we show the underlying principle using a simple example related to Euclid's algorithm.

Example 4.11 (Euclid's algorithm revisited). *Let (τ, μ) be the weighted train track on the torus shown in Figure 20 top-left, with weights a , b , and $a+b$. This train track has two vertices v and w and no base vertices. Denote the standard meridian and longitude of the torus by m and ℓ respectively. The carried 1-complex by (τ, μ) is a multicurve γ whose homology class is $am + b\ell$. Therefore the number of connected components of γ is $\text{GCD}(a, b)$. In Figure 20 top-right, a branched neighbourhood N of (τ, μ) is shown, obtained by gluing together 3 rectangles together one for each branch of τ . Identify the tie interval above the vertex v (shown with a red bracket interval) with $[1, a+b]$ such that the top part of the bracket in the figure is identified with $1 \in [1, a+b]$. Similarly identify the tie interval above w with $[a+b+1, 2a+2b]$ such that the top part of the bracket in the figure is identified with $a+b+1 \in [a+b+1, 2a+2b]$. This is our initial interval identification $\mathcal{I}$. The orbit counting problem associated to $(\tau, \mu, \mathcal{I})$ has*




FIGURE 20. Top: a weighted train track (left) where the weights are $a, b, a+b$ and the vertices are v, w , and a branched neighbourhood of it (right). In Example 4.11 we apply the AHT algorithm to unwind this train track, and see how the effect of each shifted cycle of the algorithm can be seen as a splitting or twirling on the train track. Bottom: the weighted train track obtained by splitting τ (left), and a branched neighbourhood of it (right).

the following pairings:

$$\begin{aligned} f(x) &= x + a + 2b, & f: [1, a] &\rightarrow [a + 2b + 1, 2a + 2b], \\ g(x) &= x + b, & g: [a + 1, a + b] &\rightarrow [a + b + 1, a + 2b], \\ h(x) &= x + a + b, & h: [1, a + b] &\rightarrow [a + b + 1, 2a + 2b]. \end{aligned}$$

Here f, g and h are the pairings corresponding to the edges with weights respectively a, b and $a + b$.

We saw in Example 2.3 that the AHT algorithm counts the number of orbits, which we know is equal to $\text{GCD}(a, b)$. Let us see the effect of each shifted cycle of AHT algorithm on the weighted train track. The first shifted cycle consists of Steps (1)–(4), which are not applicable in our case, so we start with the second shifted cycle. Recall that a shifted cycle applies Steps (5) and (6), followed by Steps (1)–(4).

The pairing h is maximal. Step (4) transmits g by h and replaces g with $g' = h^{-1} \circ g$ where

$$g'(x) = x - a, \quad g': [a + 1, a + b] \rightarrow [1, b].$$

Equivalently we can use the inverse $(g')^{-1}$ instead of g' where

$$(g')^{-1}(x) = x + a, \quad (g')^{-1}: [1, b] \rightarrow [a + 1, a + b].$$

Similarly transmitting f by h replaces f with $f' = h^{-1} \circ f$ where

$$f'(x) = x + b, \quad f': [1, a] \rightarrow [1 + b, a + b].$$

Step (5) then truncates h to h' , and $N = 2a + 2b$ to $N' = a + b$, where

$$h'(x) = x, \quad h': [1, a + b] \rightarrow [1, a + b].$$

Step (1) then deletes the pairing h' , and Steps (2)–(4) are not applicable. Therefore the second shifted cycle of the AHT algorithm replaces $\{f, g, h\}$ with $\{f', (g')^{-1}\}$. This is the orbit counting problem associated to the train track τ' shown in Figure 20 bottom-left. Note that τ' is obtained by splitting τ .

Now we apply the shifted cycles of AHT algorithm to the weighted train track τ' . Assume that $a < b$, the other case is similar. Let $b = na + r$, where n is an integer and $0 \leq r < a$. The pairing $(g')^{-1}$ is maximal now. Step (4) transmits f' by $(g')^{-1}$ and replaces f' by $f'' = (g')^n \circ f'$ where

$$f''(x) = x + b - na = x + r, \quad f'': [1, a] \rightarrow [1 + r, r + a].$$

Then Step (5) truncates $(g')^{-1}$ to g'' , and replaces $N' = a + b$ with $N'' = a + r$, where

$$g''(x) = x + a, \quad g'': [1, r] \rightarrow [1 + a, r + a].$$

Assume first that $r \neq 0$. Then Steps (1)–(4) are not applicable. Hence the shifted cycle replaces $\{f', (g')^{-1}\}$ with $\{f'', g''\}$. This is the orbit counting problem associated with the train track τ'' obtained by twirling the edge of τ' that has weight a , see Figure 20 bottom-left. If $r = 0$, then Step (1) deletes the identity pairing f'' and the pairing g'' has empty domain, so at this point the algorithm terminates and outputs a . This corresponds to the weighted train track getting completely unwound to the multicurve γ by the corresponding twirling.

The following is a restatement of a definition due to Dynnikov and Wiest [15, Definition 2.8 and Remark 2.2]. It will allow us, as in the case of Dynnikov and Wiest, to obtain an improved estimate for the running time of the AHT algorithm applied to the orbit counting problem associated with an integrally weighted train track.

Definition 4.12 (AHT-complexity). Let S be a compact orientable surface, and (τ, V, μ) be a based integrally weighted train track on S . Let $\mathcal{P} = \text{OCP}(\tau, V, \mu, \mathcal{I})$ be the orbit counting problem for some initial interval identification $\mathcal{I}$. Denote the pairings of $\mathcal{P}$ by g_e where e ranges over the branches of τ . Denote by $z_1, \dots, z_r$ the length of the intervals $\text{domain}(g_e)$; i.e. z_i are the measures associated to the branches of τ . Let N be the maximum natural number appearing in the intervals $\text{domain}(g_e) \cup \text{range}(g_e)$. Denote by $\tilde{z}$ the length of the narrower interval attached to N ; i.e. if N is an endpoint of a tie interval t_w then $\tilde{z}$ is the smaller measure for the two branches that come into t_w on the side of N . Define the *AHT-complexity* as

$$c_{\text{AHT}}(\mathcal{P}) := r + \sum_{i=1}^r \log(z_i) - \frac{1}{2} \log(\tilde{z}).$$

The next proposition studies the effect of running one shifted cycle of the AHT algorithm on a based integrally weighted track (τ, V, μ) , or more precisely on the associated orbit counting problem $\text{OCP}(\tau, V, \mu, \mathcal{I})$. Roughly speaking, it states that running one shifted cycle of the AHT algorithm on $\text{OCP}(\tau, V, \mu, \mathcal{I})$ gives $\text{OCP}(\tau', V', \mu', \mathcal{I}')$ for a new based integrally weighted train track (τ', V', μ') that is obtained from (τ, V, μ) via splitting or twirling. In what follows, we write (τ, V, μ) as a disjoint union of the *active part* (τ_a, V_a, μ_a) and the *stationary part* (τ_s, V_s, μ_s) . The stationary part is completely unwound, so that τ_s has no vertices beside its base vertices V_s . We apply the AHT algorithm to the active part, and as a result the active and stationary parts are updated. While this notation is cumbersome, it will be useful to keep track of the stationary part. See Remark 4.14 in the proof of the proposition.

Proposition 4.13 (Applying the AHT algorithm to a weighted train track). *Let S be a compact orientable surface, and (τ, V, μ) be a based integrally weighted track on S having E branches. Assume that (τ, V, μ) is a disjoint union of (τ_a, V_a, μ_a) and (τ_s, V_s, μ_s) , where τ_a (the active part of τ) has at least one vertex besides its base vertices V_a , and τ_s (the stationary part of τ) has no vertices besides V_s . Let $\mathcal{I}_a$ be an initial interval identification for (τ_a, V_a, μ_a) . There is a based integrally weighted train track (τ', V', μ') on S , which is a disjoint union of active (τ'_a, V'_a, μ'_a) and stationary (τ'_s, V'_s, μ'_s) parts, and an initial interval identification $\mathcal{I}'_a$ on the active part of τ' with the following properties:*

- (1) *Applying one shifted cycle of the AHT algorithm to $\text{OCP}(\tau_a, V_a, \mu_a, \mathcal{I}_a)$ gives $\text{OCP}(\tau'_a, V'_a, \mu'_a, \mathcal{I}'_a)$.*
- (2) *(τ', V', μ') is obtained from (τ, V, μ) by either*
 - *at most E splits along the same oriented branch of τ ; or*
 - *a twirling.*
- (3) *The AHT-complexity of $\text{OCP}(\tau'_a, V'_a, \mu'_a, \mathcal{I}'_a)$ is at most that of $\text{OCP}(\tau_a, V_a, \mu_a, \mathcal{I}_a)$ minus one.*

Moreover, there is an algorithm that constructs (τ', V', μ') together with its partition into active and stationary parts and the splits or twirl move from (τ, V, μ) to (τ', V', μ') , in time that is a polynomial function of E and $\log |\mu|$.

Proof. We first prove items (1)–(2). Recall the various steps in the orbit counting algorithm:

- (1) Delete any pairings that are restrictions of the identity.
- (2) Make any possible contractions.
- (3) Trim all orientation-reversing pairings whose domain and range overlap.
- (4) Perform periodic mergers.
- (5) Find a maximal g_i . For each $g_j \neq g_i$ whose range is contained in $[c_i, N']$, transmit g_j by g_i .
- (6) Find the smallest value of c such that the interval $[c, N']$ intersects the range of exactly one pairing. Truncate the pairing whose range contains the interval $[c, N']$.

Remark 4.14. Before starting the proof, we note that the reason for updating the stationary part of the weighted train track is that step (1) above may delete a pairing. This results in loss of data regarding some part of τ . To remedy this, whenever step (1) deletes a pairing, we preserve the data by adding the corresponding part of the train track to the stationary part. For example, if the carried 1-complex of (τ, μ) is a pants decomposition P , then by successively applying the AHT algorithm all components of P will eventually be removed, and so we keep a record of them by defining τ'_s .

Consider the first shifted cycle of the AHT algorithm applied to $\text{OCP}(\tau_a, V_a, \mu_a, \mathcal{I}_a)$. This consists of applying steps (1)–(4). First consider the combined effect of steps (1)–(2). In step (2), there could be no contractions unless we had deleted some pairing in step (1) that was the restriction of the identity. This happens only when τ has a component that is a simple closed curve. If this happens, then we add the resulting weighted simple closed curve to (τ_s, V_s, μ_s) . Note that steps (3)–(4) are not applicable: there is no trimming since S is orientable, and periodic mergers are not possible since every point in $[1, N']$ appears in the domain and range of at most 2 pairings. Hence the result holds for the first shifted cycle of the AHT algorithm.

Now consider a shifted cycle of the AHT algorithm consisting of steps (5)–(6) followed by steps (1)–(4). As argued before, steps (3)–(4) are not applicable.

First we show that step (5) includes a transmission. The union of the tie intervals above the base vertices V_a is $[1, M]$ for some integer M by the definition of the initial interval

identification; see Definition 4.4. Furthermore $M \neq N'$ by the assumption that τ_a has at least one vertex besides the base vertices V_a . Hence, there is a vertex $w' \notin V_a$ such that the tie interval t' above w' includes N' . Let the top of t' be the side containing N' . Then at least one of the two half-branches entering $t' \subset N(\tau_a, V_a, \mu_a)$ on the top side corresponds to a maximal pairing, allowing a transmission.

Let g_i be a maximal pairing, and $g_{j_1}, \dots, g_{j_k}$ be the pairings whose range is contained in $[c_i, N']$. After re-indexing, we may assume that $\{j_1, \dots, j_k\} = \{1, \dots, k\}$. Let $g_i = g_e$ for a branch e with initial vertex w and terminal vertex $w' \notin V_a$. Let e_- and e_+ be the half-branches of e adjacent to w and w' . Let $p_1, \dots, p_\ell$ be the cusp points on $\text{range}(g_e) \subset t'$. See e.g. Figure 18 left, where there are two such cusp points. Note that $\ell \geq k$ with equality if and only if for each $1 \leq j \leq k$ we have $\text{domain}(g_j) \not\subset \text{range}(g_e)$.

For $1 \leq j \leq \ell$, denote by l_j the leaf of the horizontal foliation on $N(\tau_a, V_a, \mu_a)$ containing p_j . Consider the following possibilities:

Case a: If the domain and range of g_e are disjoint.

Denote by R_e the rectangle corresponding to e in the branched neighbourhood $N(\tau_a, V_a, \mu_a)$. Let s_j be the connected component of $l_j|_{R_e}$ containing p_j . Let $N(\tau'_a, V'_a, \mu'_a)$ be obtained by splitting $N(\tau_a, V_a, \mu_a)$ along $\cup_{j=1}^\ell s_j$. Then applying one shifted cycle of the AHT algorithm to $\text{OCP}(\tau_a, V_a, \mu_a, \mathcal{I}_a)$ will result in $\text{OCP}(\tau'_a, V'_a, \mu'_a, \mathcal{I}'_a)$ for some $\mathcal{I}'_a$. More precisely, for each $1 \leq j \leq k$, step (5) replaces g_j with $g'_j := g_e^{-1} \circ g_j \circ g_e^s$, where $s \in \{0, 1\}$ is equal to 1 if and only if $\text{domain}(g_j) \subset \text{range}(g_e)$. Let $c := \min_{j=1}^\ell p_j \in [c_0, N'] = t$. Hence, after step (5) the points in the interval $[c, N']$ lie in the range of only g_e , and c is minimal with this property. Subsequently, step (6) will truncate the pairing g_e . After step (6) there are no pairings that are restrictions of the identity, and so step (1) will have no effect. Finally step (2) will contract the static interval $[c, N']$.

In this case (τ'_a, V'_a, μ'_a) is obtained from (τ_a, V_a, μ_a) via ℓ splits along the same branch e . Note that ℓ is at most twice the number of branches of τ .

Remark 4.15. Note that item (2) in the definition of an initial interval identification is used here. This was to ensure that N' lies on a tie interval above a vertex $w \notin V_a$; otherwise step (2) could have deleted a branch of the train track adjacent to a base vertex.

Case b: If the domain and range of g_e have non-empty intersection.

In this case $w = w'$. Let s_j be the connected component of $l_j|_{R_e}$ containing p_j . If $l_j \neq \{p_j\}$, then s_j starts at p_j , runs horizontally across R_e until it reaches the other vertical boundary of R_e , then it possibly enters R_e again, and continues this for a finite number of times until it hits a tie interval in a point not included in $\text{range}(g_e)$. Let $N(\tau'_a, V'_a, \mu'_a)$ be obtained by splitting $N(\tau_a, V_a, \mu_a)$ along $\cup_{j=1}^\ell s_j$. Then applying Steps (5)–(6) of the AHT algorithm to $\text{OCP}(\tau_a, V_a, \mu_a, \mathcal{I}_a)$ will result in $\text{OCP}(\tau'_a, V'_a, \mu'_a, \mathcal{I}'_a)$ for some $\mathcal{I}'_a$. Next if we apply Steps (1)–(2), the entire collection of pairings corresponding to some components of τ might be deleted; any such component is necessarily a simple closed curve, because of item (2) in the definition of an initial interval identification. We add any such components to the stationary part.

Let α be the simple closed curve that is the image of e in τ . Then (τ'_a, V'_a, μ'_a) is obtained from (τ_a, V_a, μ_a) via twirling along α , and then possibly transferring some simple closed curve components to the stationary part. See Definition 4.6 and Figure 18. This finishes the proof of items (1)–(2).

Now we prove item (3) following the proof of [15, Lemma 2.9, part a]. Let N be the maximum natural number appearing in the orbit counting problem $\text{OCP}(\tau_a, V_a, \mu_a, \mathcal{I}_a)$; hence N corresponds to an endpoint of a tie interval above a vertex $w \notin V_a$. There are exactly two pairings whose domain or range include N ; let $z \geq \tilde{z}$ be the length of the intervals associated

with these pairings; in particular z is the length of $\text{range}(g_e)$. Define $z' \geq \tilde{z}'$ similarly for $\text{OCP}(\tau'_a, V'_a, \mu'_a, \mathcal{I}'_a)$. Let z_i be the measure of branches of τ . Define z'_i similarly. Recall that in both Case a and Case b, $N(\tau'_a, V'_a, \mu'_a)$ is obtained by splitting $N(\tau_a, V_a, \mu_a)$ along $\cup_{j=1}^\ell s_j$, and then possibly transferring some components to the stationary part. Assume that p_ℓ is the lowest point between all p_j on the tie interval containing it; i.e. p_ℓ is the furthest from the end point N . There are two cases to consider:

- i) If p_ℓ is the minimum of $\text{range}(g_e)$, then the number of pairings will be reduced by one as a result of splitting. Hence

$$\begin{aligned} c_{\text{AHT}}^{\text{old}} - c_{\text{AHT}}^{\text{new}} &= \left(r + \sum_{i=1}^r \log(z_i) - \frac{1}{2} \log(\tilde{z})\right) - \left(r - 1 + \sum_{i=1}^{r-1} \log(z'_i) - \frac{1}{2} \log(\tilde{z}')\right) \\ &= 1 + \log(z) - \frac{1}{2} \log(\tilde{z}) + \frac{1}{2} \log(\tilde{z}') \geq 1. \end{aligned}$$

- ii) If p_ℓ is not the minimum of $\text{range}(g_e)$, then the number of pairings will not increase after splitting, and it decreases only if some simple closed curve component of τ is transferred to the stationary part. Moreover, $\tilde{z}'$ will be equal to the length of $\text{range}(g_e)$ minus a whole positive multiple of the length of $[p_\ell, N]$, where $[p_\ell, N]$ is the closed segment between p_ℓ and N along the tie interval I_w . Hence $z \geq \tilde{z} + \tilde{z}'$. We have

$$\begin{aligned} c_{\text{AHT}}^{\text{old}} - c_{\text{AHT}}^{\text{new}} &\geq \left(r + \sum_{i=1}^r \log(z_i) - \frac{1}{2} \log(\tilde{z})\right) - \left(r + \sum_{i=1}^r \log(z'_i) - \frac{1}{2} \log(\tilde{z}')\right) \\ &= \left(r + \log(z) + \sum_{z_i \neq z} \log(z_i) - \frac{1}{2} \log(\tilde{z})\right) - \left(r + \sum_{z'_i \neq \tilde{z}'} \log(z'_i) + \frac{1}{2} \log(\tilde{z}')\right) \\ &\geq \log(z) - \frac{1}{2} \log(\tilde{z}) - \frac{1}{2} \log(\tilde{z}') = \frac{1}{2} \log\left(\frac{z^2}{\tilde{z}\tilde{z}'}\right) \geq 1, \end{aligned}$$

where we used the inequality $(a + b)^2 \geq 4ab$.

This completes the proof of (3). Since the AHT-complexity is reduced by at least 1 at each step, the above algorithm terminates in time that is a polynomial function of E and $\log |\mu|$. $\square$

Remark 4.16. The operations of splitting and twirling do not change the 1-complex carried by a based integrally weighted train track. Hence, if (τ', V', μ') is obtained from (τ, V, μ) via applying a shifted cycle of the AHT algorithm as in Proposition 4.13, then the carried 1-complexes $\mathcal{CC}(\tau', V', \mu')$ and $\mathcal{CC}(\tau, V, \mu)$ are equal. Therefore, if τ' has no vertices besides the base vertices V and no complementary component to $\mathcal{CC}(\tau, V, \mu)$ is a bigon, then $\tau' = \mathcal{CC}(\tau, V, \mu)$.

The reduction of AHT-complexity in part (3) of Proposition 4.13 allows us to control the number of shifted cycles needed to completely simplify a based integrally weighted train track.

Lemma 4.17 (Bound on the number of moves for simplifying a weighted train track using the AHT algorithm). *Let (τ, V, μ) be as in Proposition 4.13. Assume that no complementary component to the carried 1-complex $\mathcal{CC}(\tau, V, \mu)$ is a bigon. Apply shifted cycles of the AHT algorithm repeatedly to obtain a sequence of based integrally weighted train tracks (τ_i, V, μ_i) for $i = 1, \dots, n$ with $(\tau_1, V, \mu_1) = (\tau, V, \mu)$ and $\tau_n = \mathcal{CC}(\tau, V, \mu)$. Let $|\mu|$ be the sum of the weights of all branches of τ , and E be the number of branches of τ . Then the number n of the train tracks is at most*

$$1 + E + E \log |\mu|.$$

Proof. The number $n - 1$ is bounded above by the number of shifted cycles of the AHT algorithm applied to the orbit counting problem $\text{OCP}(\tau, V, \mu, \mathcal{I})$. Let c_i be the AHT-complexity of $\text{OCP}(\tau_i, V, \mu_i, \mathcal{I}_i)$; see Definition 4.12. By part 3) of Proposition 4.13 we have $0 \leq c_{i+1} \leq c_i - 1$ for each i , hence the number of shifted cycles is at most the value of c_1 . We have

$$c_1 = r + \sum_{i=1}^r \log(z_i) - \frac{1}{2} \log(\tilde{z}) \leq E + E \log |\mu|.$$

□

5. UPPER BOUND FOR DISTANCE IN THE PANTS GRAPH

Theorem 1.2. *Let S be a compact connected orientable surface, and P and P' be pants decompositions of S with $i(P, P') \geq 2$. The distance between P and P' in the pants graph is*

$$O(|\chi(S)|^2 \log(i(P, P'))).$$

Moreover, there is an algorithm that constructs a (not necessarily geodesic) path $P = P_0, P_1, \dots, P_n = P'$ of length $n = O(|\chi(S)|^2 \log(i(P, P')))$ connecting P and P' in the pants graph, in time that is a polynomial function of $|\chi(S)|$ and $\log(i(P, P'))$. Here, for the input, P is given as a union of pairs of pants with gluing instructions, and P' is given in normal form with respect to P . For the output, each P_i is given as a union of pairs of pants with gluing instructions together with the pants move from P_{i-1} to P_i .

Proof. Step 1: Assume that P' is in normal form with respect to P . By squeezing parallel (in $S - P$) arcs of P' together, we obtain an integrally weighted train track (τ, μ) with vertices lying on P and with the carried 1-complex $\mathcal{CC}(\tau, \mu)$ isotopic to P' . Each component of $S - P$ contains at most 3 branches of τ , and so the number of branches of τ is at most $3|\chi(S)|$. Moreover, the total weight $|\mu|$ of all branches of τ is equal to $i(P, P')$.

Step 2: Pick any initial interval identification $\mathcal{I}$ for (τ, μ) , and repeatedly apply shifted cycles of the AHT algorithm to the orbit counting problem $\text{OCP}(\tau, \mu, \mathcal{I})$. Proposition 4.13 describes the geometric effect of each shifted cycle of the AHT algorithm on an integrally weighted train track: this is either at most $O(|\chi(S)|)$ splittings along the same oriented branch or a twirling. We apply shifted cycles of the AHT algorithm until we obtain the carried complex $\mathcal{CC}(\tau, \mu)$, see Remark 4.16 to see why the output of the algorithm is the carried 1-complex. This is possible since the AHT-complexity is reduced by at least 1 after applying each shifted cycle; see Proposition 4.13, item (3). The algorithm returns the carried 1-complex $\mathcal{CC}(\tau, \mu)$ which we know is equal to P' , but we are interested in the way that the algorithm computes the output rather than the output itself. By Proposition 4.13, there is a sequence of integrally weighted train tracks $(\tau_1, \mu_1), \dots, (\tau_n, \mu_n)$ with the following properties.

- (1) $(\tau_1, \mu_1) = (\tau, \mu)$, and $\tau_n = P'$.
- (2) The underlying 1-complex of each τ_i is a pre-triangulation. This can be seen as follows. By Proposition 4.13, each (τ_{i+1}, μ_{i+1}) is obtained from (τ_i, μ_i) by either a number of splittings and/or a twirling. By Lemma 4.9, the property of being a pre-triangulation pulls back under splitting. By Lemma 4.8, the effect of twirling on the underlying train track is applying a power of a Dehn twist, which keeps the property of being a pre-triangulation invariant, and a number of splittings. Therefore, we have shown that if τ_{i+1} is a pre-triangulation, then so is τ_i . Since $\tau_n = P'$ is a pre-triangulation, it follows inductively that each τ_i is a pre-triangulation.

Set $k_i = \chi(\tau_{i+1}) - \chi(\tau_i) \geq 0$. By Proposition 4.10, there are orderings O_i on the edges of τ_i such that for each i the pants decompositions $\mathbf{P}(\tau_i, O_i)$ and $\mathbf{P}(\tau_{i+1}, O_{i+1})$ have distance $O((k_i + 1)|\chi(S)|)$ in the pants graph. Note that

$$\sum_{i=1}^{n-1} k_i = \chi(\tau_n) - \chi(\tau_1) = O(|\chi(S)|),$$

and so if $d(\cdot, \cdot)$ denotes the distance in the pants graph, then

$$\sum_{i=1}^{n-1} d(\mathbf{P}(\tau_i, O_i), \mathbf{P}(\tau_{i+1}, O_{i+1})) = (n-1)O(|\chi(S)|) + O(|\chi(S)|^2).$$

(3) By Lemma 4.17, the number n of the train tracks is

$$1 + O(|\chi(S)|) \log(i(P, P')).$$

Step 3: Recall that τ_n has no vertices and so $\mathbf{P}(\tau_n, O_n) = \mathcal{CC}(\tau_n, \mu_n) = P'$. Applying the triangle inequality we have the following

$$\begin{aligned} d(P, P') &\leq d(P, \mathbf{P}(\tau_1, O_1)) + \sum_{i=1}^{n-1} d(\mathbf{P}(\tau_i, O_i), \mathbf{P}(\tau_{i+1}, O_{i+1})) \\ &\leq d(P, \mathbf{P}(\tau_1, O_1)) + (n-1)O(|\chi(S)|) + O(|\chi(S)|^2). \end{aligned}$$

Note that by Lemma 3.20, the pants decompositions $\mathbf{P}(\tau_1, O_1)$ and P have distance $O(|\chi(S)|^2)$, since $i(P, \tau_1) = O(|\chi(S)|)$. This gives an upper bound for $d(P, P')$ of the form

$$O(|\chi(S)|^2) + O(|\chi(S)|^2) \log(i(P, P')).$$

Since we assumed that $i(P, P') \geq 2$, we have

$$|\chi(S)|^2 \log(i(P, P')) \geq \log 2 \cdot |\chi(S)|^2.$$

Therefore, the upper bound for $d(P, P')$ can equivalently be written as

$$O(|\chi(S)|^2) \log(i(P, P')).$$

We now discuss the algorithmic and computational part of the statement. By Proposition 4.13, the sequence (τ_i, μ_i) for $1 \leq i \leq n$ can be constructed in time that is a polynomial function of $|\chi(S)|$ and $\log(i(P, P'))$. By Proposition 4.10, given O_i one can construct O_{i+1} together with a sequence of at most $O((k_i + 1)|\chi(S)|)$ consecutive transpositions taking O_i to O_{i+1} , in time that is a polynomial function of $|\chi(S)|$. For each pair (τ_i, O_i) , the pants decomposition $\mathbf{P}(\tau_i, O_i)$ can be constructed in time that is polynomial in $|\chi(S)|$. Moreover, given two pairs (τ_i, O_i) and (τ_{i+1}, O_{i+1}) , one can construct a sequence of pants moves of length $O((k_i + 1)|\chi(S)|)$ taking $\mathbf{P}(\tau_i, O_i)$ to $\mathbf{P}(\tau_{i+1}, O_{i+1})$, in time that is polynomial in $|\chi(S)|$. This completes the proof of the algorithmic part of the statement. $\square$

6. POLYGONAL DECOMPOSITIONS

In this section, we prove Theorem 6.2, which will be used to relate two polygonal decompositions of a surface.

Definition 6.1. Let S be a compact surface, and let $\mathcal{T}$ be a polygonal decomposition that is disjoint from ∂S . Let e be an arc embedded in the interior of S with both endpoints being vertices of $\mathcal{T}$ and $e \cap \mathcal{T} = \partial e$. Then the polygonal decomposition $\mathcal{T} \cup e$ is obtained from $\mathcal{T}$ by *adding the diagonal* e . We say that $\mathcal{T}$ is obtained from $\mathcal{T} \cup e$ by *deleting the diagonal* e .

Theorem 6.2. *Let S be a compact orientable surface, and $\mathcal{T}$ and $\mathcal{T}'$ be polygonal decompositions of S disjoint from ∂S with at most E edges. There is a sequence $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n = \mathcal{T}'$ of polygonal decompositions of S such that:*

- (1) *Each $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by contracting or expanding an embedded edge, or adding or deleting a diagonal, or a power of a Dehn twist. When $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by Dehn twisting k times about a curve α , then α is the closure of an edge (necessarily with the same endpoints) of $\mathcal{T}_i$, and the absolute value of k is bounded above by $i(\mathcal{T}, \mathcal{T}')$.*
- (2) *Each $\mathcal{T}_i$ is disjoint from ∂S .*
- (3) *The number of vertices of each $\mathcal{T}_i$ is $O(E)$.*
- (4) *The number n of steps in this sequence is $O(E^2 \log(i(\mathcal{T}, \mathcal{T}')) + E^2)$.*

Moreover, there is an algorithm that constructs the sequence $\mathcal{T}_1, \dots, \mathcal{T}_n$ in time that is a polynomial function of $\log(i(\mathcal{T}, \mathcal{T}'))$ and E . For the input, $\mathcal{T}$ is given as a union of polygonal discs and annuli with gluing instructions, and $\mathcal{T}'$ is given by its normal coordinates with respect to $\mathcal{T}$. For the output, each $\mathcal{T}_i$ is given as a union of polygonal discs and annuli with gluing instructions together with the move from $\mathcal{T}_i$ to $\mathcal{T}_{i+1}$.

Remark 6.3. Note that contracting or expanding edges and Dehn twists are not enough to go between two arbitrary polygonal decompositions of S , since they preserve the number of complementary regions. Similarly, adding or deleting diagonals and Dehn twists preserve the number of vertices.

Lemma 6.4. *Let $\mathcal{T}$ and $\mathcal{T}'$ be polygonal decompositions of a compact surface S that are disjoint from ∂S . Let $|\mathcal{T} \cup \mathcal{T}'|$ be the number of edges of the polygonal decomposition $\mathcal{T} \cup \mathcal{T}'$ obtained by superimposing $\mathcal{T}$ and $\mathcal{T}'$. Then $\mathcal{T}$ can be transformed into $\mathcal{T}'$ using $O(|\mathcal{T} \cup \mathcal{T}'|)$ contracting or expanding embedded edges, and $O(|\mathcal{T} \cup \mathcal{T}'|)$ adding or deleting diagonals.*

Moreover, there is an algorithm that produces such sequence of moves in time that is polynomial in $|\mathcal{T} \cup \mathcal{T}'|$. Here $\mathcal{T}$ is given as a union of polygonal discs and annuli with gluing instructions, and $\mathcal{T}'$ is given in normal form with respect to $\mathcal{T}$. The output gives a sequence of polygonal decompositions (described as a union of polygonal discs and annuli with gluing instructions) interpolating between $\mathcal{T}$ and $\mathcal{T}'$ together with the moves between consecutive polygonal decompositions.

Proof. It is enough to show that $\mathcal{T} \cup \mathcal{T}'$ can be transformed into each of $\mathcal{T}$ and $\mathcal{T}'$ using $O(|\mathcal{T} \cup \mathcal{T}'|)$ such moves. For any complementary region R of $\mathcal{T}$, remove all edges of $\mathcal{T} \cup \mathcal{T}'$ that lie in the interior of R by contracting embedded edges and deleting diagonals. Then remove any remaining extra vertices of $\mathcal{T} \cup \mathcal{T}'$ that lie on $\mathcal{T}$ by contracting embedded edges. It is easy to see that this can be done in polynomial time as a function of $|\mathcal{T} \cup \mathcal{T}'|$. $\square$

Proof of Theorem 6.2. Let V be the set of vertices of $\mathcal{T}'$. First we construct a based integrally weighted train track (τ, V, μ) such that the carried 1-complex $\mathcal{CC}(\tau, V, \mu)$ is equal to $\mathcal{T}'$. To do this, for each edge of $\mathcal{T}$ squeeze together all points of intersection with $\mathcal{T}'$ to a single switch of τ , and then for each set of parallel normal arcs of $\mathcal{T}'$ squeeze them together to a single edge of τ ; this gives the based integrally weighted train track (τ, V, μ) . See Figure 21 for an example, where hollow dots show the vertices of $\mathcal{T}'$ which in turn will correspond to base vertices of the associated train track.

Apply shifted cycles of the AHT algorithm to unwind (τ, V, μ) to $\mathcal{T}'$. By Proposition 4.13 we obtain a sequence of based integrally weighted train tracks (τ_i, V, μ_i) for $1 \leq i \leq n$ such that

- (1) $(\tau_1, V, \mu_1) = (\tau, V, \mu)$, and $(\tau_n, V, \mu_n) = (\mathcal{T}', V, \mathbf{1})$ where $\mathbf{1}$ denotes the weight assigning 1 to every branch.




FIGURE 21. Constructing a based integrally weighted train track by squeezing parallel normal arcs together. On the left the intersection of the polygonal decomposition $\mathcal{T}'$ with a face of the polygonal decomposition $\mathcal{T}$ is shown. On the right, the associated based integrally weighted track is shown, any edge with no weight written on it has weight 1. The midpoints of the polygon will then correspond to non-base vertices of the train track.

- (2) Each $(\tau_{i+1}, V, \mu_{i+1})$ is obtained from (τ_i, V, μ_i) by either splits along the same oriented branch, or by a twirling.
- (3) The number n is $O(E \log(i(\mathcal{T}, \mathcal{T}')) + E)$. This follows from Lemma 4.17 and the following two observations:
 - $|\mu|$ is at most $i(\mathcal{T}, \mathcal{T}') + |E|$.
 - The number of edges of $\mathcal{T}$ is at most E , and so the number of branches of the based train track (τ, V) is $O(E)$ as well.

Lemma 4.17 then implies that n is at most $O(E) + O(E) \log(i(\mathcal{T}, \mathcal{T}') + |E|)$. When $i(\mathcal{T}, \mathcal{T}') \geq |E|$, then $i(\mathcal{T}, \mathcal{T}') + |E| \leq 2i(\mathcal{T}, \mathcal{T}')$, and so the upper bound on n becomes $O(E \log(i(\mathcal{T}, \mathcal{T}')) + E)$, as required. On the other hand, when $i(\mathcal{T}, \mathcal{T}') < |E|$, Lemma 6.4 implies that we can find a sequence of polygonal decompositions with length $O(E)$, which is at most the required bound.

- (4) The sequence (τ_i, V, μ_i) for $1 \leq i \leq n$ can be constructed in time that is a polynomial function of $\log(i(\mathcal{T}, \mathcal{T}'))$ and E .

By Lemma 6.4, there is a sequence of $O(E)$ addition or deletion of diagonals that takes τ_1 to $\mathcal{T}$ and this sequence can be constructed in time that is polynomial in E . This is because $i(\mathcal{T}, \tau_1)$ is $O(E)$ and so superimposing τ_1 and $\mathcal{T}$ will produce a 1-complex with $O(E)$ edges.

Additionally, each τ_{i+1} is obtained from τ_i by $O(E)$ contractions or expansions of embedded edges, $O(E)$ addition or deletion of diagonals, and at most one twist map. This again follows from Lemma 6.4 since if τ_{i+1} is obtained from τ_i by splits along the same oriented branch, then after a suitable isotopy the 1-complex $\tau_i \cup \tau_{i+1}$ obtained by superimposing them has $O(E)$ edges. In this case, by Lemma 6.4 such a sequence of moves from τ_i to τ_{i+1} can be constructed in time that is polynomial in E . When τ_{i+1} is obtained from τ_i by a twist map, then the curve that we twist about and the power of the Dehn twist are already given to us by Proposition 4.13.

Combining the sequence of moves

$$\mathcal{T} \rightarrow \tau_1 \rightarrow \tau_2 \rightarrow \cdots \rightarrow \tau_n = \mathcal{T}'$$

gives the desired sequence of polygonal decompositions from $\mathcal{T}$ to $\mathcal{T}'$ of length

$$n = O(E^2 \log(i(\mathcal{T}, \mathcal{T}')) + E^2).$$

□

Notation 6.5. Let S be a surface and let Γ be an embedded finite 1-complex in S . Then $S \setminus \setminus \Gamma$ denotes the surface obtained by cutting S along Γ . In other words, equipping S with a




FIGURE 22. An example of a spine for a torus (left), a surface of genus one with one boundary component (middle), and a surface of genus one with two boundary components.




FIGURE 23. An edge swap on a spine Γ of S . Here the polygon represents $S \setminus \Gamma$.

Riemannian metric, then $S \setminus \Gamma$ is the metric completion of $S - \Gamma$ with respect to the induced metric on $S - \Gamma$.

We now consider an alternative version of Theorem 6.2 for the following important type of polygonal decompositions.

Definition 6.6. A *spine* for a closed surface S is an embedded 1-complex $\Gamma \subset S$ such that the surface $S \setminus \Gamma$ obtained by cutting S along Γ is a disc. A *spine* for a compact surface S with non-empty boundary is a 1-complex Γ embedded in the interior of S such that the surface $S \setminus \Gamma$ is a regular neighbourhood of ∂S . See Figure 22.

The following is one natural way of modifying spines.

Definition 6.7 (Edge swap [23]). Let Γ be a spine for a compact surface S . Let e' be an arc properly embedded in $S \setminus \Gamma$ with endpoints on Γ . Let e be an edge of the graph $\Gamma \cup e'$ that has distinct components of $S \setminus (\Gamma \cup e')$ on either side of it, at least one of which is a disc. Then the result of removing e from Γ and adding e' is a new spine Γ' for S . We say that Γ and Γ' are related by an *edge swap*. See Figure 23.

Example 6.8. Let Γ be a spine for the torus shown in Figure 24 left. Here Γ (shown in blue) is the union of all edges except for e' (shown in red), and the vertices of Γ are shown with solid dots. The complement of Γ is a 6-gon, and so Γ is a spine. Swapping the edge e with e' results in the spine shown on the right hand side of the same Figure, whose complement is an 8-gon in this case. Note that in an edge swap, the endpoints of e' need not be vertices of Γ , as seen in this example.

Theorem 6.9. Let S be a compact orientable surface, and let Γ and Γ' be spines of S . There is a sequence $\Gamma = \Gamma_0, \Gamma_1, \dots, \Gamma_n = \Gamma'$ of spines for S such that:

- (1) Each Γ_{i+1} is obtained from Γ_i by an edge swap, an expansion or contraction of an embedded edge, or a power of a Dehn twist. When Γ_{i+1} is obtained from Γ_i by Dehn




FIGURE 24. Edge swap: the edge e of the spine (in blue) is replaced with the edge e' (in red). The new spine is shown on the right.

twisting k times about a curve α , then $\alpha \cap \Gamma_i$ is a vertex of Γ_i , and the absolute value of k is bounded above by $i(\Gamma, \Gamma')$.

(2) The number n of steps in this sequence is $O(\chi(S)^2 \log(i(\Gamma, \Gamma')) + \chi(S)^2)$.

Moreover, there is an algorithm that constructs the sequence $\Gamma_0, \dots, \Gamma_n$ in time that is a polynomial function of $\log(i(\Gamma, \Gamma'))$ and $\chi(S)$. For the input, Γ is given as a polygon or a collection of annuli with gluing instructions, and Γ' is given by its normal coordinates with respect to Γ . For the output, each Γ_i is given as a polygon or annuli with gluing instructions together with the move from Γ_i to Γ_{i+1} .

Proof. By Theorem 6.2, there is a sequence of polygonal decompositions

$$\Gamma = \mathcal{P}_0, \mathcal{P}_1, \dots, \mathcal{P}_m = \Gamma'$$

such that:

- (1) Each $\mathcal{P}_{i+1}$ is obtained from $\mathcal{P}_i$ by contracting or expanding an embedded edge, adding or deleting a diagonal, or a power of a Dehn twist. When $\mathcal{P}_{i+1}$ is obtained from $\mathcal{P}_i$ by Dehn twisting k times about a curve α , then α is the closure of an edge of $\mathcal{P}_i$, and the absolute value of k is bounded above by $i(\Gamma, \Gamma')$.
- (2) Each $\mathcal{P}_i$ is disjoint from ∂S .
- (3) The number of vertices of each $\mathcal{P}_i$ is $O(|\chi(S)|)$.
- (4)

$$m = O(\chi(S)^2 \log(i(\Gamma, \Gamma')) + \chi(S)^2).$$

- (5) There is an algorithm that constructs the sequence $\mathcal{P}_1, \dots, \mathcal{P}_m$ together with the moves from $\mathcal{P}_i$ to $\mathcal{P}_{i+1}$, in time that is a polynomial function of $\log(i(\Gamma, \Gamma'))$ and $\chi(S)$.

For each polygonal decomposition $\mathcal{P}_i$, we will pick a collection of edges X_i in $\mathcal{P}_i$ so that $\mathcal{P}_i \setminus X_i$ is a spine Γ_i . Initially we will assume that S is closed.

We will construct X_i recursively. Initially, $X_0 = \emptyset$. To define X_{i+1} from X_i , we consider various cases:

- (i) If $\mathcal{P}_i \rightarrow \mathcal{P}_{i+1}$ is the insertion of a diagonal e , we define X_{i+1} to be $X_i \cup \{e\}$. So $\Gamma_{i+1} = \Gamma_i$.
- (ii) Suppose that $\mathcal{P}_i \rightarrow \mathcal{P}_{i+1}$ is the removal of a diagonal e . If e is in X_i , then set X_{i+1} to be $X_i \setminus \{e\}$. Then $\Gamma_i = \mathcal{P}_i \setminus X_i = \mathcal{P}_{i+1} \setminus X_{i+1} = \Gamma_{i+1}$. So we consider when e does not lie in X_i . Since we are assuming that S is closed, $S \setminus \Gamma_i$ is a single disc, and so there is an arc in this disc region joining the two sides of e . However, in $S \setminus \mathcal{P}_i$, the two sides of e lie in distinct disc regions. So there is an arc α in X_i , consisting of a union of edges, so that α lies in the boundaries of both these disc regions. See Figure 25. Let X_{i+1} equal X_i minus the edges in α . Then Γ_{i+1} is obtained from Γ_i by removing e and adding this arc α . In other words, Γ_{i+1} is obtained from Γ_i by performing an edge swap.




FIGURE 25. Proof of Theorem 6.9, case (ii): The edge e is a diagonal of the disc $S \setminus \Gamma_i$, and $\alpha \subset X_i$ is a union of edges.

- (iii) If $\mathcal{P}_i \rightarrow \mathcal{P}_{i+1}$ is the expansion of an edge e , we set X_{i+1} to be X_i , unless this creates a 1-valent vertex of Γ_{i+1} , in which case we set $X_{i+1} = X_i \cup \{e\}$. So, Γ_{i+1} is obtained from Γ_i by expanding the edge e , or is a copy of Γ_i .
- (iv) Suppose that $\mathcal{P}_i \rightarrow \mathcal{P}_{i+1}$ is the contraction of an embedded edge e . If e is not in X_i , then we set $X_{i+1} = X_i$ and we note Γ_{i+1} is obtained from Γ_i by contracting e . So suppose that e is in X_i . If we were to add e to Γ_i , the complementary disc $S \setminus \Gamma_i$ would be divided into two discs, A and B say. Moreover, the complementary regions incident to one endpoint of e are not all in A or B . Hence, there is some edge e' of Γ_i emanating from this endpoint of e that has A on one side and B on the other. Let $X'_i = X_i \cup \{e'\} \setminus \{e\}$. Let $\Gamma'_i = \mathcal{P}_i \setminus X'_i$. Then Γ'_i is obtained from Γ_i by an edge swap. Moreover, we may set X_{i+1} to be X'_i . Then Γ_{i+1} is obtained from Γ'_i by contracting e .
- (v) Suppose $\mathcal{P}_i \rightarrow \mathcal{P}_{i+1}$ is obtained by Dehn twisting about a curve α that is the closure of an embedded edge. Then we may perturb α so that it intersects $\mathcal{P}_i$ in the vertex at the endpoints of α . Hence, α intersects Γ_i in exactly this vertex. We let X_{i+1} be the image of X_i under these Dehn twists.

Suppose now that ∂S is non-empty. Cases (i), (iii) and (v) are identical to the situation where S is closed.

We now explain how (ii) is modified in the case where ∂S is non-empty. Again suppose that $\mathcal{P}_i \rightarrow \mathcal{P}_{i+1}$ is the removal of a diagonal e . Again, the difficult case is where e does not lie in X_i , and hence lies in Γ_i . If the two sides of e lie in the same annular region of $S \setminus \Gamma_i$, then there must be an arc α consisting of a union of edges of X_i in this annulus separating these two copies of e in the boundary of the annulus. In that case, we set X_{i+1} to be equal to X_i minus the edges of α , and then Γ_{i+1} is obtained from Γ_i by an edge swap. So suppose that the two sides of e lie in distinct annular regions of $S \setminus \Gamma_i$. One side of e lies in a disc component D of $S \setminus \Gamma_i$. Say that this lies in a component A of $S \setminus \Gamma_i$. Then the edges in X_i separate D from $\partial S \cap A$. Hence, there is an arc α properly embedded in A consisting of a union of edges of X_i that separates D from $\partial S \cap A$. Again set X_{i+1} to be equal to X_i minus the edges of α , and again Γ_{i+1} is obtained from Γ_i by an edge swap.

The argument in case (iv) is very similar to the case when S is closed. Suppose that $\mathcal{P}_i \rightarrow \mathcal{P}_{i+1}$ is the contraction of an embedded edge e . The difficult situation is where e is in X_i . Then e is an arc properly embedded in an annulus component of $S \setminus \Gamma_i$. It is disjoint from ∂S , and hence it separates the annulus into an annulus and a disc D . Emanating from the vertex at one endpoint of e , there is an edge e' that has D on one side and an annulus of $S \setminus (\Gamma_i \cup e)$ on the other. Let $X'_i = X_i \cup \{e'\} \setminus \{e\}$ and let $\Gamma'_i = \mathcal{P}_i \setminus X'_i$. Then Γ'_i is obtained

from Γ_i by an edge swap. Setting X_{i+1} to be X'_i , Γ_{i+1} is obtained from Γ_i by the contraction of e . □

7. ONE-VERTEX AND IDEAL TRIANGULATIONS

In this section, we improve Theorem 6.2 by showing that one can stay within the class of one-vertex triangulations or ideal triangulations.

Theorem 1.6. *Let S be a compact orientable surface. When S is closed (respectively, has non-empty boundary), let $\mathcal{T}$ and $\mathcal{T}'$ be one-vertex (respectively, ideal) triangulations of S . Assume that $i(\mathcal{T}, \mathcal{T}') \geq 2$. There is a sequence $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n = \mathcal{T}'$ of one-vertex (respectively, ideal) triangulations of S such that:*

- (1) *Each $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by either a flip or a power of a Dehn twist. When $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by Dehn twisting k times about a curve α , then α is a normal curve intersecting each edge of $\mathcal{T}_i$ at most three times, and the absolute value of k is bounded above by $i(\mathcal{T}, \mathcal{T}')$.*
- (2) *$n = O(|\chi(S)|^3 \log(i(\mathcal{T}, \mathcal{T}'))$.*

Moreover, there is an algorithm that constructs the sequence $\mathcal{T}_1, \dots, \mathcal{T}_n$ in time that is a polynomial function of $\log(i(\mathcal{T}, \mathcal{T}'))$ and $|\chi(S)|$. Here we assume that $\mathcal{T}$ and $\mathcal{T}'$ have the same vertex (when S is closed), $\mathcal{T}$ is given as a union of (possibly ideal) triangles with gluing instructions, and $\mathcal{T}'$ is given in terms of its normal coordinates with respect to $\mathcal{T}$. For the output, each $\mathcal{T}_i$ is given as a union of (possibly ideal) triangles with gluing instructions together with the flip or twist move from $\mathcal{T}_i$ to $\mathcal{T}_{i+1}$.

We will prove this by dualising the spines of Theorem 6.9. Recall that the *dual* of a spine Γ is a polygonal decomposition that has a complementary disc for each vertex of Γ and an edge dual to each edge of Γ . When the surface is closed, the dual of Γ has a single vertex in the disc region $S \setminus \Gamma$. When the surface has non-empty boundary, the vertices of the dual of Γ are all 1-valent and lie on ∂S .

The dual 1-complex to a one-vertex triangulation of a closed surface S or an ideal triangulation of a surface with boundary is a *trivalent spine*; i.e. a spine in which every vertex has degree 3. A flip on a triangulation or ideal triangulation corresponds to a *Whitehead move* on its dual spine. Therefore, Theorem 1.6 can be rephrased in terms of two given trivalent spines and the existence of a ‘short’ sequence of Whitehead moves and twist maps taking one to the other.

Each edge swap between trivalent spines can be written as a composition of a controlled number of Whitehead moves. This is the content of the next lemma which is an analogue of Lemma 8.3 of [23] for *trivalent* spines.

Lemma 7.1. *Let Γ and Γ' be trivalent spines for a compact orientable surface S such that Γ' is obtained from Γ by an edge swap. Then Γ' can be obtained from Γ by $O(|\chi(S)|)$ Whitehead moves.*

Moreover, there is an algorithm that constructs such a sequence of Whitehead moves in time that is a polynomial function of $|\chi(S)|$.

Proof. We follow the proof of Lemma 8.3 in [23]. Assume that Γ' is obtained from Γ by deleting the edge e and adding $e' \neq e$. Note that by definition of an edge swap, e and e' can only intersect at endpoints if at all. Let A be the surface obtained by cutting S along $\Gamma \setminus \{e\}$. This is an annulus (when S is closed) or a three-times punctures sphere or once-punctured

annulus (when $\partial S \neq \emptyset$). Then e and e' are two essential properly embedded arcs in A , and hence they are isotopic in A . Indeed, there is a disc component of $A \setminus (e \cup e')$ with boundary equal to the concatenation of an arc in ∂A , the edge e , another arc in ∂A and the other edge e' . First we move one endpoint of e across this disc to one endpoint of e' . Then we move the other endpoint of e across this disc to the other endpoint of e' . Finally we isotope e to e' keeping its endpoints fixed; this third step only requires an isotopy and no Whitehead moves. It is enough to show that the first step can be done with $O(|\chi(S)|)$ Whitehead moves, as the second step is similar. There are $O(|\chi(S)|)$ vertices of Γ between the first endpoints of e and e' along ∂A . Passing the endpoint of e across any one of these vertices can be seen as a Whitehead move. Therefore, the total number of Whitehead moves needed is $O(|\chi(S)|)$. $\square$

Lemma 7.2. *Let D be a polygon with n sides. Let $\mathcal{T}$ and $\mathcal{T}'$ be triangulations of D where each side of D is an edge and with no vertices in the interior of D . Then $\mathcal{T}$ and $\mathcal{T}'$ differ by a sequence of at most $2n - 6$ flips. Moreover, there is an algorithm that constructs such a sequence of flips in time that is polynomial in n .*

We term triangulations $\mathcal{T}$ and $\mathcal{T}'$ as above *diagonal subdivisions* of D .

Proof. The proof is as in [28, Lemma 2] and we repeat it to make the algorithmic part of the statement clear. Given a diagonal subdivision of D and a vertex x , if the degree $\deg(x)$ is not equal to $n - 3$, then we can increase $\deg(x)$ by performing a flip. Hence after $n - 3 - \deg(x)$ flips, we can convert the subdivision into a new subdivision where all diagonals have one endpoint at x . Hence, $\mathcal{T}$ can be converted to $\mathcal{T}'$ using at most $2n - 6$ flips. $\square$

The dual of a triangulation $\mathcal{T}$ as in the above lemma is a tree embedded within the disc D that has 1-valent vertices on ∂D and trivalent vertices in the interior of D . We can view the lemma as providing a sequence of Whitehead moves between any two such trees.

Proof of Theorem 1.6. We are given 1-vertex or ideal triangulations $\mathcal{T}$ and $\mathcal{T}'$. Let Γ and Γ' be the spines dual to $\mathcal{T}$. By Theorem 6.9, there is a sequence $\Gamma = \Gamma_0, \Gamma_1, \dots, \Gamma_n = \Gamma'$ of spines for S such that:

- (1) Each Γ_{i+1} is obtained from Γ_i by an edge swap, an expansion or contraction of an embedded edge, or a power of a Dehn twist. When Γ_{i+1} is obtained from Γ_i by Dehn twisting k times about a curve α , then $\alpha \cap \Gamma_i$ is a vertex of Γ_i , and the absolute value of k is bounded above by $i(\Gamma, \Gamma')$.
- (2) The number n of steps in this sequence is $O(\chi(S)^2 \log(i(\mathcal{T}, \mathcal{T}')) + \chi(S)^2)$.

We will remove a small regular neighbourhood of each vertex of Γ_i and replace it by a tree. Each vertex of this tree has either degree 1 or 3. The 1-valent vertices lie where the tree meets the remnants of the edges of Γ_i , and the remaining vertices of the tree are trivalent, see Figure 26. Let F_i be the union of these trees. Set $\mathcal{Q}_i$ to be the resulting trivalent spine.

The spine Γ_0 is dual to a one-vertex or ideal triangulation, so any small neighbourhood of one of its vertices is of the right form. The procedure can be initiated by taking F_0 to be the union of a small neighbourhood of each vertex and then $\mathcal{Q}_0 = \Gamma_0$. Initially, each component of F_0 has a single trivalent vertex and three 1-valent vertices. To define F_{i+1} , we consider various cases:

- (1) Suppose that $\Gamma_i \rightarrow \Gamma_{i+1}$ is an edge swap, removing an edge e and inserting an edge e' . Then e corresponds to an edge (also called e) of $\mathcal{Q}_i$. The edge e' may have one or both of its endpoints on a vertex of Γ_i , in which case when we remove a regular neighbourhood of these vertices, we also remove the end segments of e' , but we can then extend the remnant of e' to an edge e'' with one or both endpoints on the interior




FIGURE 26. Proof of Theorem 1.6: a small regular neighbourhood of each vertex of Γ_i is replaced with a tree.




FIGURE 27. Proof of Theorem 1.6, case (1): the edge e' has an endpoint on a vertex of Γ_i , see the left side of the figure. We then construct e'' as in right such that its corresponding endpoint lies in the interior of an edge of $\mathcal{Q}_i$.




FIGURE 28. Proof of Theorem 1.6, case (2): The new spine is obtained by contracting an edge e , see the top. The corresponding trivalent spine is isotopic to the previous one, see the bottom.

of an edge of F_i . See Figure 27. If we remove e from $\mathcal{Q}_i$ and attach on e'' , the result is a trivalent spine $\mathcal{Q}_{i+1}$. The forest F_{i+1} is defined to be the intersection between $\mathcal{Q}_{i+1}$ and the regular neighbourhood of the vertices of Γ_{i+1} . By construction, $\mathcal{Q}_{i+1}$ is obtained from $\mathcal{Q}_i$ by an edge swap. Hence by Lemma 7.1, $\mathcal{Q}_i$ and $\mathcal{Q}_{i+1}$ differ by a sequence of $O(|\chi(S)|)$ Whitehead moves.

- (2) Suppose $\Gamma_i \rightarrow \Gamma_{i+1}$ is the contraction of an edge e . In $\mathcal{Q}_i$, there is a copy of e and at its endpoints there are two components of F_i . We amalgamate them into a single tree by attaching the edge e , and we declare that this is a component of F_{i+1} . The remaining components of F_i become components of F_{i+1} . In this way, $\mathcal{Q}_{i+1}$ is isotopic to $\mathcal{Q}_i$. See Figure 28.
- (3) Now consider the case where $\Gamma_i \rightarrow \Gamma_{i+1}$ is the expansion of an edge e from a vertex v . Let T be the component of F_i in a regular neighbourhood of v . Let v_1 and v_2 be the vertices at the endpoint of e . Pick trees of the required form T_1 and T_2 for these




FIGURE 29. Left: A surface of genus two is shown as a polygon with side identifications. The spine Γ (in red solid lines) is the union of solid curves passing through the center of the polygon. The normal curve α (in dashed blue lines) is the union of the two dashed lines passing through the center of the polygon and five other normal arcs going close to the vertex of the polygon. Right: The spine Γ is perturbed to a trivalent spine $\mathcal{Q}$ (in red solid lines), which is dual to the shown triangulation $\mathcal{T}$ of the surface. The curve α is isotoped to be a concatenation of normal arcs (with respect to $\mathcal{T}$) together with a connected union of two edges of $\mathcal{Q}$. It can be seen in this example that the isotoped α could have 3 normal arcs in some triangles, including 2 normal arcs of the same type.

regular neighbourhoods to be components of F_{i+1} . We can view $T_1 \cup e \cup T_2$ to be a tree lying in a regular neighbourhood of v . Using Lemma 7.2, T can be transformed into $T_1 \cup e \cup T_2$ using $O(|\chi(S)|)$ Whitehead moves. Hence, $\mathcal{Q}_i$ and $\mathcal{Q}_{i+1}$ differ by a sequence of $O(|\chi(S)|)$ Whitehead moves.

- (4) Finally suppose that $\Gamma_i \rightarrow \Gamma_{i+1}$ is a power of a Dehn twist along a curve α that intersects Γ_i in a vertex. In $\mathcal{Q}_i$, this vertex is replaced by a tree, and α can be arranged to intersect this tree in a connected union of edges or a single vertex. We set F_{i+1} to be the image of F_i under this power of a Dehn twist.

We now dualise this sequence of trivalent spines to form a sequence of 1-vertex or ideal triangulations of S . The dual of each Whitehead move is a flip. Hence, we need at most $O(|\chi(S)|^3 \log(i(\mathcal{T}, \mathcal{T}')) + |\chi(S)|^3)$ flips and powers of Dehn twists. Since we assumed that $i(\mathcal{T}, \mathcal{T}') \geq 2$, this can be equivalently written as $O(|\chi(S)|^3 \log(i(\mathcal{T}, \mathcal{T}')))$. When we Dehn twist, the curves that we twist along intersects the spine in a connected union of edges or a single vertex. When we dualise, this curve is a concatenation of normal arcs, together with a part that runs into a (possibly ideal) vertex of the triangulation. Push the curve off the (possibly ideal) vertex, and we obtain a normal curve that intersects each edge of the triangulation at most three times. To see this note that when we push the curve off the vertex, it skirts around the vertex, and in doing so it picks up at most one new normal arc of each type in each triangle. Since such a triangle may already have had a normal arc in it, we get at most four normal arcs, and these may intersect each edge at most three times. See Figure 29 for an example.

It is clear that this sequence of 1-vertex and ideal triangulations is constructible in polynomial time as a function of $|\chi(S)|$ and $\log(i(\mathcal{T}, \mathcal{T}'))$. □

Lemma 7.3. *Let S be a compact connected orientable surface with non-empty boundary, and let A be a collection of disjoint arcs properly embedded in S . Let V be a finite collection of*

points on ∂S disjoint from A , with at least one point of V on each component of ∂S . In the case where S is a disc, suppose also $|V| \geq 3$. Then V is the vertex set of a triangulation $\mathcal{T}$ of S with the property that each edge of $\mathcal{T}$ intersects each component of A at most twice. Moreover, given a triangulated surface S , a normal multi-arc A , and a set of points V as above, there is an algorithm that constructs the triangulation $\mathcal{T}$. The algorithm runs in time that is a polynomial function of the number of triangles in some input triangulation of S , the cardinality of V , and the ℓ^1 -norm of the normal coordinates of A .

Proof. Consider first the case where S is a disc. Note that each component of $\partial S \setminus V$ intersects each component of A at most twice, in a subset of the endpoints of that component. These components of $\partial S \setminus V$ will be edges of $\mathcal{T}$, and so we have verified the required condition for these edges. If $|V| = 3$, then we set $\mathcal{T}$ to be a single triangle. So suppose $|V| > 3$. Pick three vertices in V that are consecutive around ∂S . Join the outermost two by an edge that runs parallel to ∂S . This will be an edge of $\mathcal{T}$. It intersects each component of A at most twice. These three vertices now span a triangle. Removing this triangle from S gives a disc with one fewer vertices in its boundary. Hence, by induction, S has the required triangulation.

Now suppose that S is not a disc. Suppose also that some component of ∂S contains more than one vertex in V . Pick three consecutive vertices on this component of ∂S (where the outermost two may be equal) and join the outermost two by an edge and then remove a triangle, as above. In this way, we may suppose that each component of ∂S contains a single vertex. We now modify A to a new set of arcs A' as follows:

- (1) remove any inessential arcs;
- (2) replace parallel essential arcs by a single arc;
- (3) if any complementary region is not a disc or is a disc with more than three arcs in its boundary, then add in an essential arc not parallel to a previous one, and avoiding A . Repeat this as much as possible. This step turns the arc system into a hexagonal decomposition of the surface where for each hexagon the edges alternately lie in ∂S and in A' .

The resulting arcs form the 1-skeleton of an ideal triangulation of S . By construction, each either is equal to a component of A or is disjoint from A . We now add further arcs to A' , one for each component of ∂S . Consider any component C of ∂S and the vertex v in V that it contains. Pick some orientation on C . Let p_1 and p_2 be the endpoints of A' that are adjacent to v , where the orientation on C runs from p_1 to v . Say that p_i lies in the arc a_i in A' . We add the following arc to A' : it starts at the end of a_1 that is not p_1 , it runs along a_1 and then along the sub-arc of ∂S containing v up to p_2 . This sub-arc of ∂S may contain endpoints of inessential arcs of A , in which case we modify this new arc so that it avoids these inessential arcs of A . We repeat this for each component of ∂S . Let A'' be the union of A' and these new arcs, perturbed a little so that they are disjoint from each other. By construction, they are disjoint from A .

We now slide the endpoints of A'' along ∂S , using the chosen orientations on the components of ∂S . We stop when all the endpoints of A'' lie in V . The result is the 1-skeleton of the required triangulation $\mathcal{T}$. Note that this sliding operation may introduce points of intersection between the edges of $\mathcal{T}$ and A , but each edge of $\mathcal{T}$ intersects each component of A at most twice, near the endpoints of that component of A . $\square$

Theorem 7.4. *Let S be a closed orientable surface, and $\mathcal{T}$ be a one-vertex triangulation of S . Let γ be an essential simple closed normal curve given by its normal vector (γ) with respect to $\mathcal{T}$, and denote the bit-sized complexity of (γ) by $|(\gamma)|_{\text{bit}}$ and its ℓ^1 -norm by $|(\gamma)|_1$. There is*

an algorithm that constructs a sequence of one-vertex triangulations $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n$ of S and a sequence of curves $\gamma = \gamma_0, \gamma_1, \dots, \gamma_n$ such that

- (1) γ_i is isotopic to γ for every i .
- (2) γ_i is in normal form with respect to $\mathcal{T}_i$ for every $i < n$.
- (3) γ_n lies in the 1-skeleton of $\mathcal{T}_n$.
- (4) Each $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by either a flip or a power of a Dehn twist. When $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by Dehn twisting k times about a curve α , then α intersects each edge of $\mathcal{T}_i$ at most three times, and the absolute value of k is bounded above by a polynomial function of $|(\gamma)|_1$ and $|\chi(S)|$.
- (5) The algorithm runs in time that is a polynomial function of $|(\gamma)|_{\text{bit}}$ and $|\chi(S)|$.

For the output, each $\mathcal{T}_i$ is given as a union of triangles with gluing instructions together with the flip or twist move from $\mathcal{T}_i$ to $\mathcal{T}_{i+1}$, and γ_i is given by its normal coordinates with respect to $\mathcal{T}_i$.

Proof. The idea is to extend γ to a one-vertex triangulation $\mathcal{T}'$ of S , and then repeat the proof of Theorem 1.6. More precisely, we show that there is an algorithm that extends γ to a one-vertex triangulation $\mathcal{T}'$ of S such that:

- (1) $\mathcal{T}'$ is in normal form with respect to $\mathcal{T}$.
- (2) The algorithm runs in time that is a polynomial function of $|(\gamma)|_{\text{bit}}$ and $|\chi(S)|$. In particular the bit-sized complexity of the normal coordinates of $\mathcal{T}'$ with respect to $\mathcal{T}$ are bounded from above by such a polynomial function.

Step 1: Construction of $\mathcal{T}'$.

The normal arcs of γ decompose triangles of $\mathcal{T}$ into several 0-handles (or 2-cells). A 0-handle of $S \setminus (\mathcal{T} \cup \gamma)$ is called a *parallelity handle* if it is a 4-gon with two of its opposite sides being parallel normal arcs of γ and the other two sides lying in the edges of $\mathcal{T}$. Each parallelity 0-handle comes with the structure of an I -bundle over an interval, where the interval base is parallel to a normal arc of γ . See Figure 30. The I -bundle structures on parallelity 0-handles glue together to form an I -bundle $\mathcal{B}$, called the *parallelity bundle*, whose base B is a possibly disconnected compact 1-manifold. Moreover, the base B can be considered as a normal (possibly not closed) multicurve. Since we assumed that γ is connected, the base B is a finite union of closed intervals. The *vertical boundary* $\partial_v \mathcal{B}$ of $\mathcal{B}$ is defined as the restriction of the I -bundle $\mathcal{B}$ to ∂B . The *horizontal boundary* $\partial_h \mathcal{B}$ of $\mathcal{B}$ is the (∂I) -bundle over B , obtained by the restriction of the I -bundle. Therefore, the boundary of $\mathcal{B}$ is the union of its vertical boundary and horizontal boundary. Similarly, for each component of $\mathcal{B}$, we can speak of its vertical and horizontal boundary.

In each triangle of $\mathcal{T}$, there are at most four 0-handles that are not parallelity handles, see Figure 30. The union of the 0-handles that are not parallelity handles forms a 2-complex called the *gut region*. Therefore, the number of 0-handles of the gut region is at most $4t$, where t is the number of triangles in $\mathcal{T}$.

Denote the vertex of $\mathcal{T}$ by v . Let Δ be a triangle 0-handle of the gut region, and x be the side of Δ opposite the vertex v . Place a vertex w on x , and isotope γ by dragging the vertex w to v along a straight line in Δ . After this isotopy and normalisation, γ is a simple closed normal curve with one vertex on it that coincides with the vertex of $\mathcal{T}$. See Figure 31. We will apply the Agol–Hass–Thurston algorithm to compute the following data about the parallelity bundle: Denote the components of $\mathcal{B}$ by $\mathcal{B}_1, \dots, \mathcal{B}_k$. For each $\mathcal{B}_i$, we compute the normal coordinates for its base, together with the attachment of the vertical boundary of $\mathcal{B}_i$ to the gut region, and the relative I -direction for the two components of the vertical boundary of $\mathcal{B}$.




FIGURE 30. Normal arcs decompose a triangle into a union of 0-handles (or 2-cells). On the left the 0-handles that are not parallelity handles are shaded. Note that in general there are at most four 0-handles that are not parallelity handles. The fibers of the I -bundle structure of a parallelity handle (well-defined up to isotopy) are shown on the right hand side.




FIGURE 31. Proof of Theorem 7.4, isotoping γ such that it becomes a normal curve passing through the vertex v of the triangulation: Some of the triangles adjacent to v are shown. We first drag the point w of γ to the vertex v inside Δ , and then normalise the resulting curve. The normalised curve passing through the vertex v is shown in dashed lines.

Therefore, we have a handle decomposition $\mathcal{H}$ of $X = S \setminus \setminus \gamma$ into 0-handles, where each 0-handle is either a 0-handle of the gut region, or it is a 4-gon that is equal to a component of $\mathcal{B}$. This handle decomposition $\mathcal{H}$ of X has $O(t)$ 0-handles. There is a natural immersion $i: X \rightarrow S$ whose restriction to the interior of X is an embedding, and such that γ lies in the image of the boundary of X under the map i . Note that X is a compact orientable surface with two or one connected components according to whether γ is separating in S or not. Let V be the copies of the vertex v in X . By Lemma 7.3, X admits a triangulation $\mathcal{T}'_X$ with vertex set V and where each edge intersects each component of $\partial_v \mathcal{B}$ at most twice.

Define the one-vertex triangulation $\mathcal{T}'$ of S as the image of the triangulation $\mathcal{T}'_X$ of $S \setminus \setminus \gamma$ under the map i . We can now read off the normal coordinate with respect to $\mathcal{T}$ of each edge e of $\mathcal{T}'$. To see this, consider two cases:

- i) For each part of e in the gut region, we can read its normal coordinate with respect to $\mathcal{T}$.
- ii) Let H be a 0-handle of $\mathcal{H}$ that forms a component of the parallelity bundle $\mathcal{B}$. We previously computed the base of H as a normal arc with respect to $\mathcal{T}$, using the Agol–Hass–Thurston algorithm. Each time e runs through H , it enters and exits H via

$\partial_v H$ and so each component of $e \cap H$ is normally parallel to the base of H . So, we can read off the normal coordinate of $e \cap H$ with respect to $\mathcal{T}$.

Summing these coordinates over each 0-handle of $\mathcal{H}$ gives the normal vector of e with respect to $\mathcal{T}$. Finally, the normal coordinate of $\mathcal{T}'$ with respect to $\mathcal{T}$ can be obtained by summing up the normal coordinates of its edges. Note that by construction we have

$$(2) \quad i(\mathcal{T}, \mathcal{T}') \leq O(|\chi(S)|) \cdot |(\gamma)|_1 + O(\chi(S)^2).$$

To see this note that the base of the parallelity bundle is a normal multi-arc of ℓ^1 -norm at most $|(\gamma)|_1$. Moreover, $\mathcal{T}'$ has $O(|\chi(S)|)$ edges and each edge of $\mathcal{T}'$ passes through each component of $\mathcal{B}$ at most twice. Therefore the intersection of $\mathcal{T}'$ with $\mathcal{T} \cap \mathcal{B}$ contributes at most $O(|\chi(S)|) \cdot |(\gamma)|_1$ intersection points. Additionally, each edge of $\mathcal{T}'$ intersects the gut region of $\mathcal{H}$ at most $O(|\chi(S)|)$ times, and so $\mathcal{T}'$ intersects the restriction of $\mathcal{T}$ to the gut region at most $O(\chi(S)^2)$ times.

Step 2: Construction of $\mathcal{T}_i$ and γ_i .

Let

$$\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n = \mathcal{T}'$$

be the sequence of one-vertex triangulations given by Theorem 1.6. Set $\gamma_0 := \gamma$. Each $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by either a flip or a twist map. Given the normal curve γ_i with respect to $\mathcal{T}_i$, we put γ_i in normal form with respect to $\mathcal{T}_{i+1}$ and define γ_{i+1} as this normal representative. Consider two cases:

- a) If $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by a flip, then it is easy to put γ_i in normal form with respect to $\mathcal{T}_{i+1}$ and find its normal coordinates.
- b) If $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by a twist map $(T_\alpha)^k$, then the normal coordinates of γ_i with respect to $\mathcal{T}_{i+1}$ are equal to the normal coordinates of $(T_\alpha)^{-k}(\gamma_i)$ with respect to $\mathcal{T}_i$. The normal coordinates of $(T_\alpha)^{-k}(\gamma_i)$ with respect to $\mathcal{T}_i$ can be read off from the normal coordinates of γ_i with respect to $\mathcal{T}_i$, the normal coordinates of α with respect to $\mathcal{T}_i$, and the value of k , and all this information is given to us by Theorem 1.6.

By construction, γ_i satisfy conditions (1)–(3) of the statement of the theorem. By Theorem 1.6 and equation (2), conditions (4)–(5) of the statement are satisfied as well. $\square$

We can similarly prove a version of Theorem 7.4 for ideal triangulations of surfaces with boundary.

Theorem 7.5. *Let S be a compact orientable surface with non-empty boundary, and $\mathcal{T}$ be an ideal triangulation of S . Let γ be an essential simple closed normal curve or an essential simple normal arc given by its normal vector (γ) with respect to $\mathcal{T}$, and denote the bit-sized complexity of (γ) by $|(\gamma)|_{\text{bit}}$ and its ℓ^1 -norm by $|(\gamma)|_1$. There is an algorithm that constructs a sequence of ideal triangulations $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n$ of S and a sequence of curves $\gamma = \gamma_0, \gamma_1, \dots, \gamma_n$ such that*

- (1) γ_i is isotopic to γ for every i .
- (2) γ_i is in normal form with respect to $\mathcal{T}_i$ for every $i < n$.
- (3) γ_n intersects each edge of $\mathcal{T}_n$ at most twice.
- (4) Each $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by either a flip or a power of a Dehn twist. When $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by Dehn twisting k times about a curve α , then α intersects

each edge of $\mathcal{T}_i$ at most three times, and the absolute value of k is bounded above by a polynomial function of $|(\gamma)|_1$ and $|\chi(S)|$.

(5) The algorithm runs in time that is a polynomial function of $|(\gamma)|_{\text{bit}}$ and $|\chi(S)|$.

For the output, each $\mathcal{T}_i$ is given as a union of ideal triangles with gluing instructions together with the flip or twist move from $\mathcal{T}_i$ to $\mathcal{T}_{i+1}$, and γ_i is given by its normal coordinates with respect to $\mathcal{T}_i$.

Proof. The proof is similar to the proof of Theorem 7.4. When γ is a closed curve, we first isotope it to create an arc that passes through an ideal vertex v of $\mathcal{T}$. We then repeat the argument as in the proof of Theorem 7.4 to construct a sequence of ideal triangulations $\mathcal{T} = \mathcal{T}_0, \dots, \mathcal{T}_n$ and arcs $\gamma_0, \dots, \gamma_n$ such that γ_n is an edge of $\mathcal{T}_n$. Finally when γ is a closed curve, we perturb γ_n off the ideal vertex v to create a curve that intersects each edge of $\mathcal{T}'$ at most twice. When γ is an arc, we perturb γ_n so that it is normal and disjoint from the edges of $\mathcal{T}'$.

Note that if γ is a separating curve and $\mathcal{T}$ has exactly one ideal vertex, the geometric intersection number between γ_n and an edge of $\mathcal{T}_n$ is even. Therefore, assuming further that γ is essential, there is an edge e of $\mathcal{T}_n$ such that γ intersects e at least twice. $\square$

8. APPLICATION TO VOLUMES OF HYPERBOLIC 3-MANIFOLDS

Theorem 1.2 together with Agol's explicit construction of hyperbolic structures in [1] has the following corollary.

Corollary 1.7. *Let Σ a closed orientable surface of genus $g \geq 2$, and P and P' be pants decompositions of Σ with no curve in common. Assume that M is a maximal cusp obtained from a quasi-Fuchsian 3-manifold homeomorphic to $\Sigma \times \mathbb{R}$ by pinching the multicurves P and P' to annular cusps on the two conformal boundary components of M . The volume of the convex core of M is*

$$O(g^2) \log(i(P, P')).$$

Proof. Since P and P' have no curve in common, each curve in P intersects P' at least once. Therefore $i(P, P') \geq |P| \geq 2$, where $|P| = 3g - 3$ is the number of curves in P . By Theorem 1.2, there is a path C consisting of $P_0 = P, P_1, \dots, P_m = P'$ in the pants graph of Σ with length $m = O(g^2) \log(i(P, P'))$.

We recall Agol's work from [1]; Agol's proof is written for mapping tori but works with minor changes for product manifolds too as we will see below. Define the sequence of circles $\beta_1, \dots, \beta_m$ in Σ where β_{i+1} is the circle in P_{i+1} replacing a circle in P_i . Also set $\beta_0 = P_0$. Define the subset $\mathcal{A}$ of $\{\beta_1, \dots, \beta_m\}$ as follows: β_i belongs to $\mathcal{A}$ whenever β_i is isotopic to a curve in P' and for no $j > i$, β_j is isotopic to β_i . For $0 \leq i \leq m$, define the multicurves B_i in $\Sigma \times [0, 1]$ as follows: $B_i = \beta_i \times \{\frac{i}{m}\}$ if $\beta_i \notin \mathcal{A}$, and $B_i = \beta_i \times \{1\}$ if $\beta_i \in \mathcal{A}$. Consider the complement $M_C := \Sigma \times [0, 1] \setminus N(\cup B_i)$ where $N(\cup B_i)$ is a regular neighbourhood of $\cup B_i$. In other words, we drill out β_i at successive heights from $\Sigma \times [0, 1]$, except that the curves in P' are all drilled out from the level $\Sigma \times \{1\}$. Likewise, all the curves in P_0 are drilled from the level $\Sigma \times \{0\}$, since $\beta_0 \notin \mathcal{A}$. Let M be the maximal cusp obtained from $\Sigma \times \mathbb{R}$ by pinching the multicurves P and P' , and N be the convex core of M . Then N is obtained from M_C by Dehn filling boundary components corresponding to those curves B_i that lie in the interior of $\Sigma \times [0, 1]$. Agol [1, Lemma 2.3] constructed a complete hyperbolic structure with totally geodesic boundary and with rank-1 and rank-2 cusps on M_C by gluing together 'model pieces'. The rank-1 cusps correspond to the curves drilled out from $\Sigma \times \{0, 1\}$, which are $P \times \{0\}$ and $P' \times \{1\}$. The rank-2 cusps correspond to the curves drilled out from the

interior of $\Sigma \times [0, 1]$. Finally the totally geodesic boundary corresponds to the union of pair of pants in $\Sigma \times \{1\} \setminus P' \times \{1\}$ and $\Sigma \times \{0\} \setminus P \times \{0\}$. The explicit construction of the hyperbolic structure shows that

$$\text{Vol}(M_C) = (2A + S)V_{\text{oct}} \leq 2V_{\text{oct}}m,$$

where Vol denotes the volume of the hyperbolic structure, A and S indicate the number of *associativity/simple* moves in the path C , and V_{oct} is the volume of the regular ideal octahedron; see the proof of [1, Corollary 2.4]. Let N be the manifold obtained by filling in the boundary components of M_C corresponding to those B_i that lie in the interior of $\Sigma \times [0, 1]$. Then N is a hyperbolic 3-manifold with totally geodesic boundary and rank-1 cusps and

$$\text{Vol}(N) < \text{Vol}(M_C).$$

To see this note that by Thurston's hyperbolisation theorem for Haken manifolds, N is hyperbolic since it is Haken, atoroidal, and anannular. By Thurston [30], the volume decreases under Dehn filling and so

$$\text{Vol}(N) < \text{Vol}(M_C).$$

Alternatively we can first double each of M_C and N along the subset $(\Sigma \times \{0\} \setminus P \times \{0\}) \cup (\Sigma \times \{1\} \setminus P' \times \{1\})$ of boundary, and work with the doubled manifold to avoid rank-1 cusps. Hence we have

$$\text{Vol}(N) < \text{Vol}(M_C) \leq 2V_{\text{oct}}m = O(g^2) \log(i(P, P')).$$

□

The next result shows that our bound in Corollary 1.7 is sharp up to a multiplicative factor of $g \log(g)$.

Proposition 8.1. *There is a universal positive constant C such that for any $g \geq 3$, there are maximal cusps M obtained from $\Sigma_g \times \mathbb{R}$ by pinching the multicurves P and P' to annular cusps such that the volume of the convex core of M is greater than*

$$C \frac{g}{\log(g)} \log(i(P, P')).$$

Proof. Let P_1 be a pants decomposition of a surface Σ_2 of genus 2, and denote the simple closed curves in P_1 by $\{\alpha_1, \alpha_2, \alpha_3\}$. Pick a base point b on Σ_2 , and let $\pi_1(\Sigma_2, b)$ be the fundamental group. Let $f: \Sigma_2 \rightarrow \Sigma_2$ be a pseudo-Anosov mapping class that acts trivially on the first homology group $H_1(\Sigma_2; \mathbb{Z})$. After an isotopy, we can assume that f fixes the base point b , and so it acts on $\pi_1(\Sigma_2, b)$. Define the pants decomposition $P'_1 := f(P_1)$. Since f is pseudo-Anosov, it does not fix the isotopy class of any essential simple closed curve on Σ_2 . So, after possibly replacing f by a power of itself, we can assume that $f(\alpha_i)$ is not isotopic to α_j for any $1 \leq i, j \leq 3$.

Picking a tree connecting α_i to the base point b in Σ_2 , we can identify α_i with elements of $\pi_1(\Sigma_2, b)$. Let $\phi: \pi_1(\Sigma_2, b) \rightarrow \mathbb{Z}$ be a surjective homomorphism such that $\phi(\alpha_i) = 0$ for $1 \leq i \leq 3$, for example ϕ can be taken to be the algebraic intersection pairing with some α_j that is non-separating. Denote by ϕ_k the composition of ϕ with the reduction map $\mathbb{Z} \rightarrow \mathbb{Z}/k\mathbb{Z}$ modulo k for $k \geq 2$, and define G as the kernel of ϕ_k . Therefore, G is an index k subgroup of $\pi_1(\Sigma, b)$ that contains all α_i for $1 \leq i \leq 3$. The image of ϕ_k is abelian and so it factors through the abelianisation map $\pi_1(\Sigma_2, b) \rightarrow H_1(\Sigma_2; \mathbb{Z})$. Since f acts trivially on homology, we conclude that the curves in $P'_1 = f(P_1)$ also lie in $G = \ker(\phi_k)$.

Let M_1 be the maximal cusp obtained from $\Sigma_2 \times \mathbb{R}$ by pinching the pants decompositions P_1 and P'_1 to annular cusps, here we use the fact that P_1 and P'_1 have no curve in common.

Let M be the k -sheeted cover of M_1 corresponding to the subgroup $G < \pi_1(\Sigma_2, b) \cong \pi_1(M_1)$. The hyperbolic structure on M_1 lifts to a hyperbolic structure on M , and so there exists g such that M is obtained from $\Sigma_g \times \mathbb{R}$ by pinching the lifts $P := \tilde{P}_1$ and $P' := \tilde{P}'_1$ of respectively P_1 and P'_1 to annular cusps. Since the curves in P_1 and P'_1 lie in $G = \ker(\phi_k)$ by construction, we deduce that $P := \tilde{P}_1$ and $P' := \tilde{P}'_1$ are pants decompositions (that is, they cut Σ into pairs of pants). By comparing Euler characteristics we have

$$\chi(\Sigma_g) = k \cdot \chi(\Sigma_2) \implies g - 1 = k.$$

Since $k \geq 2$, we have that $g \geq 3$. Denote the convex core of a hyperbolic manifold N by $\text{CC}(N)$, and its hyperbolic volume by $\text{Vol}(N)$. Since M is a k -sheeted cover of M_1

$$\text{Vol}(\text{CC}(M)) = k \cdot \text{Vol}(\text{CC}(M_1)) = (g - 1)\text{Vol}(\text{CC}(M_1)).$$

To see this, let Γ_1 be a discrete group of isometries of hyperbolic three-space $\mathbb{H}^3$ such that $M_1 = \mathbb{H}^3/\Gamma_1$, and $\Gamma < \Gamma_1$ be a subgroup of index k with $M = \mathbb{H}^3/\Gamma$. Denote the limit sets of Γ_1 and Γ by respectively $\Lambda(\Gamma_1)$ and $\Lambda(\Gamma)$. It is easy to see that $\Lambda(\Gamma_1) = \Lambda(\Gamma)$, since $\Gamma < \Gamma_1$ is of finite index. Denote the convex hull of $\Lambda(\Gamma_1)$ by $\text{CH}(\Lambda(\Gamma_1))$, and define $\text{CH}(\Lambda(\Gamma))$ similarly. Then the convex core of the hyperbolic manifold $\mathbb{H}^3/\Gamma_1$ is the image of $\text{CH}(\Lambda(\Gamma_1))$ under the projection $p_1: \mathbb{H}^3 \rightarrow \mathbb{H}^3/\Gamma_1$, see [30, Chapter 8.3]. Similarly, the convex core of $\mathbb{H}^3/\Gamma$ is the image of $\text{CH}(\Lambda(\Gamma))$ under the projection $p: \mathbb{H}^3 \rightarrow \mathbb{H}^3/\Gamma$. It follows from $\Lambda(\Gamma_1) = \Lambda(\Gamma)$ that $\text{Vol}(\text{CH}(\Lambda(\Gamma_1))/\Gamma_1) = k \cdot \text{Vol}(\text{CH}(\Lambda(\Gamma))/\Gamma)$, proving the claim.

We also have

$$i(P, P') = k \cdot i(P_1, P'_1) = (g - 1)i(P_1, P'_1).$$

Fix the pants decompositions P_1, P'_1 on Σ_2 and the homomorphism ϕ , and allow k to vary. To simplify the notation set $A = \text{Vol}(\text{CC}(M_1))$ and $B = i(P_1, P'_1)$, and note that $A, B > 0$ are universal constants. Therefore, for $C > 0$ sufficiently small we have

$$\begin{aligned} \text{Vol}(\text{CC}(M)) &= (g - 1)\text{Vol}(\text{CC}(M_1)) = A(g - 1) \geq \\ C \frac{g}{\log(g)} \log(B(g - 1)) &= C \frac{g}{\log(g)} \log(i(P, P')). \end{aligned}$$

For example if $C = \frac{A}{3 \log(B)}$, then the above inequality is satisfied because

$$\begin{aligned} C \frac{g}{\log(g)} \log(B(g - 1)) &= C \frac{g}{\log(g)} (\log B + \log(g - 1)) < \\ C \frac{g}{\log(g)} (\log B + \log g) &= C \log(B) \frac{g}{\log(g)} + Cg < \\ \frac{A}{3} \frac{g}{\log(g)} + \frac{A}{3 \log(B)} g &< \\ \frac{A}{3} g + \frac{A}{3} g &= \frac{2A}{3} g \leq A(g - 1). \end{aligned}$$

In the penultimate line above we used the fact that $B = i(P_1, P'_1) \geq 3$ since every curve in P_1 intersects P'_1 at least once, and so $\log(B) > 1$. In the last line we used that $g \geq 3$ and hence $2g \leq 3(g - 1)$. □

9. APPLICATION TO TEICHMÜLLER GEOMETRY

Our next application is to Teichmüller space, endowed with the Weil-Petersson metric. Denote by $\text{Teich}(S)$ the Teichmüller space of an orientable surface S of finite type (in other

words, the space of marked hyperbolic metrics possibly with cusps but with no boundary). Denote its Weil–Petersson metric by d_{WP} .

Theorem 1.8. *Let S be an orientable surface of finite type, and let $X, Y \in \text{Teich}(S)$. Let P_X and P_Y be pants decompositions for X and Y respectively, in which each curve has length at most $2\pi|\chi(S)|$, which exist by a theorem of Parlier [26]. Then*

$$d_{\text{WP}}(X, Y) \leq O(|\chi(S)|^2) (1 + \log(i(P_X, P_Y) + 1)).$$

Proof. Set $L = 2\pi|\chi(S)|$. By Parlier’s quantified version of Bers’s theorem [26], which is building on and improving the work of Buser and Seppälä [10], every hyperbolic metric X on S has a pants decomposition P in which every curve has length at most L . Let $\ell_X(P)$ denote the sum of the lengths of the curves in P . Hence, $\ell_X(P) \leq 3|\chi(S)|L$. Let $N(P)$ be the nodal surface where each curve in P has length pinched to zero. We may view $N(P)$ as a point in the metric completion of $\text{Teich}(S)$ with the Weil–Petersson distance. Wolpert showed (Corollary 4.10 in [31]) that the Weil–Petersson distance from X to $N(P)$ is at most

$$\sqrt{2\pi\ell_X(P)} \leq \sqrt{12\pi}|\chi(S)|.$$

Therefore, $d_{\text{WP}}(X, Y) \leq d_{\text{WP}}(N(P_X), N(P_Y)) + O(|\chi(S)|)$.

Cavendish and Parlier [12] introduced a metric graph which they term the *cubical pants graph* $\mathcal{CP}(S)$. This is obtained from the pants graph $\mathcal{P}(S)$ by adding edges (which may have length greater than one). Hence, for the two pants decompositions P_X and P_Y , their distance in $\mathcal{CP}(S)$ is at most their distance in $\mathcal{P}(S)$. In Lemma 4.1 of [12], it is shown that there is an absolute constant C such that

$$d_{\text{WP}}(N(P_X), N(P_Y)) \leq C d_{\mathcal{CP}(S)}(P_X, P_Y).$$

So,

$$\begin{aligned} d_{\text{WP}}(X, Y) &\leq d_{\text{WP}}(N(P_X), N(P_Y)) + O(|\chi(S)|) \\ &\leq C d_{\mathcal{CP}(S)}(P_X, P_Y) + O(|\chi(S)|) \\ &\leq C d_{\mathcal{P}(S)}(P_X, P_Y) + O(|\chi(S)|) \\ &\leq O(|\chi(S)|^2) (1 + \log(i(P_X, P_Y))), \end{aligned}$$

where the latter inequality is from Theorem 1.2. □

10. QUESTIONS

Question 10.1. *Is there an algorithm that takes as input a compact connected orientable surface S and two pants decompositions P and P' , and computes the distance $d_{\mathcal{P}}(P, P')$ in the pants graph, in time that is a polynomial function of $|\chi(S)|$ and $\log(1 + i(P, P'))$?*

It follows from the work of Irmer [21] that there is an algorithm that computes the distance in the pants graph.

A simple argument shows that the distance in the flip graph of one-vertex triangulations is bounded from above by the intersection number; see [13, Lemma 2.1]. This prompts the following question.

Question 10.2. *Is there a universal constant $C > 0$ such that for every compact orientable surface S and pants decompositions P and P' of S , the distance between P and P' in the pants graph is at most $Ci(P, P')$?*

REFERENCES

- [1] I. Agol. Small 3-manifolds of large genus. *Geometriae Dedicata*, 102(1):53–64, 2003.
- [2] I. Agol, J. Hass, and W. Thurston. The computational complexity of knot genus and spanning area. *Transactions of the American Mathematical Society*, 358(9):3821–3850, 2006.
- [3] F. Baroni. Uniformly polynomial-time classification of surface homeomorphisms. *arXiv preprint arXiv:2402.00231*, 2024.
- [4] M. C. Bell. Simplifying triangulations. *Discrete & Computational Geometry*, 66(1):1–11, 2021.
- [5] M. C. Bell and R. Webb. Personal communication.
- [6] M. C. Bell and R. C. Webb. Applications of fast triangulation simplification. *arXiv preprint arXiv:1605.03514*, 2016.
- [7] L. Bers. An inequality for riemann surfaces. *Differential Geometry and Complex Analysis: A Volume Dedicated to the Memory of Harry Ernest Rauch*, pages 87–93, 1985.
- [8] J. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *Journal of the American Mathematical Society*, 16(3):495–535, 2003.
- [9] P. Buser. *Riemannsche Flächen und Längenspektrum vom trigonometrischen Standpunkt aus*. PhD thesis, Mathematisches Institut der Univ., 1980.
- [10] P. Buser and M. Seppälä. Symmetric pants decompositions of Riemann surfaces. 1992.
- [11] R. D. Canary, M. Culler, S. Hersonsky, and P. B. Shalen. Approximation by maximal cusps in boundaries of deformation spaces of Kleinian groups. *Journal of Differential Geometry*, 64(1):57–109, 2003.
- [12] W. Cavendish and H. Parlier. Growth of the Weil-Petersson diameter of moduli space. *Duke Math. J.*, 161(1):139–171, 2012.
- [13] V. Disarlo and H. Parlier. The geometry of flip graphs and mapping class groups. *Transactions of the American Mathematical Society*, 372(6):3809–3844, 2019.
- [14] N. M. Dunfield and D. P. Thurston. A random tunnel number one 3-manifold does not fiber over the circle. *Geometry & Topology*, 10(4):2431–2499, 2006.
- [15] I. Dynnikov and B. Wiest. On the complexity of braids. *Journal of the European Mathematical Society*, 9(4):801–840, 2007.
- [16] J. Erickson and A. Nayyeri. Tracing compressed curves in triangulated surfaces. *Discrete Comput Geom*, 49:823–863, 2013.
- [17] B. Farb, A. Lubotzky, and Y. Minsky. Rank-1 phenomena for mapping class groups. *Duke Mathematical Journal*, 106(3):581 – 597, 2001.
- [18] W. J. Harvey. Boundary structure of the modular group. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, NY, 1978)*, volume 97, pages 245–251, 1981.
- [19] A. Hatcher and W. Thurston. A presentation for the mapping class group of a closed orientable surface. *Topology*, 19(3):221–237, 1980.
- [20] J. Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40(3):631–657, 2001.
- [21] I. Irmer. Stable lengths on the pants graph are rational. *New York J. Math.*, 21:1153–1168, 2015.
- [22] M. Lackenby. Some fast algorithms for curves in surfaces. *arXiv preprint arXiv:2401.16056*, 2024.
- [23] M. Lackenby and J. S. Purcell. The triangulation complexity of fibred 3-manifolds. *arXiv preprint arXiv:1910.10914*, 2019.
- [24] W. R. Lickorish. A representation of orientable combinatorial 3-manifolds. *Annals of Mathematics*, pages 531–540, 1962.
- [25] L. Mosher. Tiling the projective foliation space of a punctured surface. *Transactions of the American Mathematical Society*, pages 1–70, 1988.
- [26] H. Parlier. A shorter note on shorter pants. *arXiv preprint arXiv:2304.06973*, 2023.
- [27] R. C. Penner and J. L. Harer. *Combinatorics of train tracks*. Number 125. Princeton University Press, 1992.
- [28] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Journal of the American Mathematical Society*, 1(3):647–681, 1988.
- [29] S. J. Taylor and A. Zupan. Products of farey graphs are totally geodesic in the pants graph. *Journal of Topology and Analysis*, 8(02):287–311, 2016.
- [30] W. P. Thurston. *The geometry and topology of three-manifolds*. Princeton University Princeton, NJ, 1979.
- [31] S. A. Wolpert. Behavior of geodesic-length functions on Teichmüller space. *J. Differential Geom.*, 79(2):277–334, 2008.