

THE TRIANGULATION COMPLEXITY OF ELLIPTIC AND SOL 3-MANIFOLDS

MARC LACKENBY AND JESSICA S. PURCELL

ABSTRACT. The triangulation complexity of a compact 3-manifold M is the minimal number of tetrahedra in any triangulation of M . We compute the triangulation complexity of all elliptic 3-manifolds and all sol 3-manifolds, to within a universally bounded multiplicative error.

1. INTRODUCTION

The *triangulation complexity* $\Delta(M)$ of a compact 3-manifold M is the minimal number of tetrahedra in any triangulation of M . (In this paper, we use the definition of a triangulation that has become standard in low-dimensional topology: it is an expression of M as a union of 3-simplices with some of their faces identified in pairs via affine homeomorphisms.) Triangulation complexity is a very natural invariant, with some attractive properties. However, its precise value is known for only relatively small examples [17, 20] and for a few infinite families [21, 23, 8, 11, 12, 10, 7]. It bears an obvious resemblance to hyperbolic volume, and in fact the volume of a hyperbolic 3-manifold M forms a lower bound for $\Delta(M)$ via the inequality $\Delta(M) \geq \text{Vol}(M)/v_3$, due to Gromov and Thurston [26]. Here, $v_3 \simeq 1.01494$ is the volume of a regular hyperbolic ideal tetrahedron. But non-trivial lower bounds for manifolds with zero Gromov norm have been difficult to obtain. Jaco, Rubinstein and Tillmann [8, 11] were able to compute the triangulation complexity of lens spaces of the form $L(2n, 1)$ and $L(4n, 2n \pm 1)$. However, general lens spaces have remained out of reach. In this paper, we remedy this, by computing the triangulation complexity of all elliptic 3-manifolds and all sol 3-manifolds, to within a universally bounded multiplicative error. Our result about lens spaces confirms a conjecture of Jaco and Rubinstein [9] and Matveev [22, 20], up to a bounded multiplicative constant.

Theorem 1.1. *Let $L(p, q)$ be a lens space, where p and q are coprime integers satisfying $0 < q < p$. Let $[a_0, \dots, a_n]$ be the continued fraction expansion of p/q where each a_i is positive. Then there is a universal constant $k_{\text{lens}} > 0$ such that*

$$k_{\text{lens}} \sum_{i=0}^n a_i \leq \Delta(L(p, q)) \leq \sum_{i=0}^n a_i.$$

General elliptic 3-manifolds fall into three categories: lens spaces, prism manifolds and a third class that we call Platonic manifolds; for example see Scott [24] for a discussion of the classification of these manifolds. Recall that the prism manifold $P(p, q)$ is obtained from the orientable I -bundle over the Klein bottle, by attaching a solid torus, so that the meridian of the solid torus is identified with the p/q curve on the boundary torus. Here, a canonical framing of this boundary torus is used, so that the longitude and meridian are lifts of non-separating simple closed curves on the Klein bottle that are, respectively, orientation-reversing and orientation-preserving.

Theorem 1.2. *Let p and q be non-zero coprime integers and let $[a_0, \dots, a_n]$ denote the continued fraction expansion of p/q where a_i is positive for each $i > 0$. Then, $\Delta(P(p, q))$ is, to within a universally bounded multiplicative error, equal to $\sum_{i=0}^n a_i$.*

We say that an elliptic 3-manifold is *Platonic* if it admits a Seifert fibration where the base orbifold is the quotient of S^2 by the orientation-preserving symmetry group of a Platonic solid. These orbifolds have underlying space the 2-sphere and have three exceptional points with orders $(2, 3, 3)$, $(2, 3, 4)$ or $(2, 3, 5)$. The Seifert fibration is specified by the Seifert data, which describes the three singular fibres and includes the Euler number of the fibration. It turns out that the latter quantity controls the triangulation complexity.

Theorem 1.3. *Let M be a Platonic elliptic 3-manifold, and let e denote the Euler number of its Seifert fibration. Then, to within a universally bounded multiplicative error, $\Delta(M)$ is $|e|$.*

We also examine sol 3-manifolds. Recall that these are 3-manifolds of the form $(T^2 \times [0, 1])/(A(x, 1) \sim (x, 0))$ where A is an element of $\mathrm{SL}(2, \mathbb{Z})$ with $|\mathrm{tr}(A)| > 2$. Such a matrix A induces a homeomorphism of the torus that is known as *linear Anosov*. Let \bar{A} be the image of A in $\mathrm{PSL}(2, \mathbb{Z})$. Recall that $\mathrm{PSL}(2, \mathbb{Z})$ is isomorphic to $\mathbb{Z}_2 * \mathbb{Z}_3$ where the factors are generated by

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad T = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}.$$

Thus any element of $\mathrm{PSL}(2, \mathbb{Z})$ can be written uniquely as a word that is an alternating product of elements S and T or T^{-1} . The word is *cyclically reduced* if the first letter is neither the inverse of the final letter nor equal to the final letter. Any element of $\mathrm{PSL}(2, \mathbb{Z})$ is conjugate to a cyclically reduced word that is unique up to cyclic permutation. Our first theorem about sol manifolds relates the triangulation complexity of the manifold to the length of this cyclically reduced word.

Theorem 1.4. *Let A be an element of $\mathrm{SL}(2, \mathbb{Z})$ with $|\mathrm{tr}(A)| > 2$. Let M be the sol 3-manifold $(T^2 \times [0, 1])/(A(x, 1) \sim (x, 0))$. Let \bar{A} be the image of A in $\mathrm{PSL}(2, \mathbb{Z})$ and let $\ell(\bar{A})$ be the length of a cyclically reduced word in the generators S and $T^{\pm 1}$ that is conjugate to \bar{A} . Then, there is a universal constant $k_{\mathrm{sol}} > 0$ such that*

$$k_{\mathrm{sol}} \ell(\bar{A}) \leq \Delta(M) \leq (\ell(\bar{A})/2) + 6.$$

Note that this length $\ell(\bar{A})$ is readily calculable. For we may simplify A using row operations until it is the identity matrix. This writes A as a product of elementary matrices. The image of each elementary matrix in $\mathrm{PSL}(2, \mathbb{Z})$ is a word in the generators S and T . Thus, we obtain \bar{A} as a word in S and T . If the starting letter is equal to the inverse of the final letter or equal to the final letter, then we may conjugate by the inverse of this element to create a shorter word. Thus, eventually, we end with a cyclically reduced word, and $\ell(\bar{A})$ is its length.

Our second theorem relates the triangulation complexity of the 3-manifold M to the continued fraction expansion of $\sqrt{\mathrm{tr}(A)^2 - 4}$. As $\mathrm{tr}(A)^2 - 4$ is not a perfect square, the continued fraction expansion of $\sqrt{\mathrm{tr}(A)^2 - 4}$ does not terminate. Denote it by $[a_0, a_1, \dots]$. As $\sqrt{\mathrm{tr}(A)^2 - 4}$ is the square root of a positive integer, the continued fraction expansion is eventually periodic, in the sense that for some non-negative integer r and even positive integer t , $a_{i+t} = a_i$ for every $i \geq r$. The *periodic part*

of the continued fraction expansion is (a_r, \dots, a_{r+t-1}) , which is well-defined up to cyclic permutation.

Theorem 1.5. *Let A be an element of $\mathrm{SL}(2, \mathbb{Z})$ with $|\mathrm{tr}(A)| > 2$. Let \bar{A} be the image of A in $\mathrm{PSL}(2, \mathbb{Z})$. Suppose that \bar{A} is B^n for some positive integer n and some $B \in \mathrm{PSL}(2, \mathbb{Z})$ that cannot be expressed as a proper power. Let M be the sol 3-manifold $(T^2 \times [0, 1]) / (A(x, 1) \sim (x, 0))$. Let $[a_0, a_1, \dots]$ be the continued fraction expansion of $\sqrt{\mathrm{tr}(A)^2 - 4}$ where a_i is positive for each $i > 0$ and let (a_r, \dots, a_s) denote its periodic part. Then there is a universal constant $k'_{\mathrm{sol}} > 0$ such that*

$$k'_{\mathrm{sol}} n \sum_{i=r}^s a_i \leq \Delta(M) \leq 6 + n \sum_{i=r}^s a_i.$$

Crucial to our arguments is the analysis of triangulations of $T^2 \times [0, 1]$. This is because $T^2 \times [0, 1]$ arises when we cut a sol manifold along a torus fibre, or when we remove the core curves of the Heegaard solid tori of a lens space. Our result about these products is as follows.

Theorem 1.6. *Let \mathcal{T}_0 and \mathcal{T}_1 be 1-vertex triangulations of the torus T^2 . Let $\Delta(\mathcal{T}_0, \mathcal{T}_1)$ denote the minimal number of tetrahedra in any triangulation of $T^2 \times [0, 1]$ that equals \mathcal{T}_0 on $T^2 \times \{0\}$ and equals \mathcal{T}_1 on $T^2 \times \{1\}$. Then there is a universal constant $k_{\mathrm{prod}} > 0$ such that*

$$k_{\mathrm{prod}} d_{\mathrm{Tr}(T^2)}(\mathcal{T}_0, \mathcal{T}_1) \leq \Delta(\mathcal{T}_0, \mathcal{T}_1) \leq d_{\mathrm{Tr}(T^2)}(\mathcal{T}_0, \mathcal{T}_1) + 6.$$

Here, $\mathrm{Tr}(S)$ denotes the triangulation graph for a closed orientable surface S , defined to have a vertex for each isotopy class of 1-vertex triangulation of S , and where two vertices are joined by an edge if and only if the corresponding triangulations differ by a 2-2 Pachner move. Each edge is declared to have length 1, and this induces the metric $d_{\mathrm{Tr}(S)}(\cdot, \cdot)$. When S is the torus, this graph is in fact equal to the classical Farey tree and so distances in the graph can readily be computed using continued fractions.

This paper is a continuation of the work in [15], where we analysed the triangulation complexity of 3-manifolds M that fibre over the circle with fibre a closed orientable surface S with genus at least 2. We were able to estimate $\Delta(M)$, to within a bounded factor depending only on the genus of S , in the case where the monodromy ϕ of the fibration is pseudo-Anosov. Our theorem related $\Delta(M)$ to the translation length of the action of ϕ on various metric spaces.

Recall that if X is a space with metric d , and ϕ is an isometry of X , its *translation length* $\ell_X(\phi)$ is $\inf\{d(x, \phi(x)) : x \in X\}$. Its *stable translation length* $\bar{\ell}_X(\phi)$ is $\inf\{d(\phi^n(x), x)/n : n \in \mathbb{Z}_{>0}\}$, where $x \in X$ is chosen arbitrarily.

Each homeomorphism ϕ of S naturally induces an isometry of $\mathrm{Tr}(S)$. It also induces an isometry of the mapping class group $\mathrm{MCG}(S)$, where $\mathrm{MCG}(S)$ is given a word metric by making a fixed choice of some finite generating set. The homeomorphism ϕ also acts isometrically on Teichmüller space, with its Teichmüller or Weil-Petersson metrics. The *thick* part of Teichmüller space consists of those hyperbolic structures where every geodesic on the surface has length at least some suitably chosen $\epsilon > 0$. If $\epsilon > 0$ is sufficiently small, the thick part is path connected, and so may be given its path metric. The homeomorphism ϕ also induces an isometry of these metric spaces.

The following was the main theorem of [15]. The statement below combines the statement of [15, Theorem 1.3] with Theorems 3.5 and Proposition 2.7 of that paper.

Theorem 1.7. *Let S be a closed orientable connected surface with genus at least 2, and let $\phi: S \rightarrow S$ be a pseudo-Anosov homeomorphism. Then the following quantities are within bounded ratios of each other, where the bounds depend only on the genus of S and a choice of finite generating set for $\text{MCG}(S)$:*

- (1) *the triangulation complexity of $(S \times I)/\phi$;*
- (2) *the translation length (or stable translation length) of ϕ in the thick part of the Teichmüller space of S ;*
- (3) *the translation length (or stable translation length) of ϕ in the mapping class group of S ;*
- (4) *the translation length (or stable translation length) of ϕ in $\text{Tr}(S)$.*

We also analysed products.

Theorem 1.8 (Theorem 1.4 of [15]). *Let S be a closed orientable surface with genus at least 2 and let \mathcal{T}_0 and \mathcal{T}_1 be non-isotopic 1-vertex triangulations of S . Then the following are within a bounded ratio of each other, the bounds only depending on the genus of S :*

- (1) *the minimal number of tetrahedra in any triangulation of $S \times [0, 1]$ that equals \mathcal{T}_0 and \mathcal{T}_1 on $S \times \{0\}$ and $S \times \{1\}$ respectively;*
- (2) *the minimal number of 2-2 Pachner moves relating \mathcal{T}_0 and \mathcal{T}_1 .*

Theorems 1.4 and 1.6 are natural generalisations of these results to the case where S is a torus. For technical reasons, we were unable to deal with this case in [15]. In this paper, we deduce Theorems 1.4 and 1.6 from Theorem 1.8, by passing to a branched covering space.

We record here the analogue of Theorem 1.7 for sol manifolds.

Theorem 1.9. *Let $\phi: T^2 \rightarrow T^2$ be a linear Anosov homeomorphism. Then the following quantities are within universally bounded ratios of each other:*

- (1) *the triangulation complexity of $(T^2 \times I)/\phi$;*
- (2) *the translation length (or stable translation length) of ϕ in the thick part of the Teichmüller space of T^2 ;*
- (3) *the translation length (or stable translation length) of ϕ in the mapping class group of T^2 ;*
- (4) *the translation length (or stable translation length) of ϕ in $\text{Tr}(T^2)$.*

In (3), we metrize $\text{MCG}(T^2)$ by fixing the finite generating set

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

The structure of the paper is as follows. In Section 2, we recall some basic facts about handle structures, including the notion of a parallelity bundle. Section 3 contains the first substantial new result, Theorem 3.2. This asserts that for any triangulation \mathcal{T} of $S \times I$, the product of a closed orientable surface S and an interval, there is an arc isotopic to $* \times I$, for some point $* \in S$, that is simplicial in $\mathcal{T}^{(23)}$, the 23rd iterated barycentric subdivision. This is technically important, because it allows us to transfer a triangulation of $S \times I$ to a triangulation of a suitable branched cover. This is used later in Section 8, where Theorem 1.6 is proved using Theorem 1.8. A suitable finite branched cover of the torus T^2 over one point is used, which is a closed orientable surface S of genus greater than one. In order to compare translation lengths in $\text{Tr}(S)$ and $\text{Tr}(T^2)$, we develop some background

theory in Sections 4, 5, 6 and 7. In Section 4, we introduce $\text{Sp}(S)$, which is the space of spines for a closed orientable surface S . There is a quasi-isometry between $\text{Tr}(S)$ and $\text{Sp}(S)$ that is equivariant under the action of the mapping class group, but it is useful to consider both spaces. In Section 5, we recall the relationship between $\text{Tr}(T^2)$, the Farey graph and continued fractions. In Section 6, we recall results of Masur, Mosher and Schleimer [18], which use train tracks to estimate distances in $\text{Tr}(S)$. In Section 7, this is applied specifically in the case of the torus. Finally in Sections 9, 10, 11 and 12, we deal with products, sol manifolds, lens spaces, prism manifolds and Platonic manifolds.

2. HANDLE STRUCTURES AND PARALLELITY BUNDLES

Although the main results of this paper are on the complexity of triangulations of 3-manifolds, for our arguments it is often more convenient to work with handle structures, similarly to [15]. This section collects some of the definitions and results on handle structures that we will use.

Recall that a handle structure on a 3-manifold is a decomposition into i -handles $D^i \times D^{3-i}$, $i = 0, 1, 2, 3$, where we require:

- (1) Each i -handle intersects the handles of lower index in $\partial D^i \times D^{3-i}$.
- (2) Any two i -handles are disjoint.
- (3) The intersection of any 1-handle with any 2-handle is of the form:
 - $D^1 \times \alpha$ in the 1-handle $D^1 \times D^2$, where α is a collection of arcs in ∂D^2 ,
 - $\beta \times D^1$ in the 2-handle $D^2 \times D^1$, where β is a collection of arcs in ∂D^2 .
- (4) Any 2-handle runs over at least one 1-handle.

For example, given a triangulation of a 3-manifold M , there is an associated handle structure for M minus an open collar neighbourhood of ∂M , called the *dual handle structure*. This has 0-handles obtained by removing a thin regular open neighbourhood of the boundary of each tetrahedron, 1-handles obtained by taking a neighbourhood of each face not in ∂M and removing a thin regular open neighbourhood of each edge, 2-handles obtained by taking a neighbourhood of each edge not in ∂M with neighbourhoods of endpoints removed, and 3-handles consisting of a regular neighbourhood of each vertex not in ∂M .

One feature of a handle structure that does not hold for a triangulation is that it can be sliced along a normal surface to yield a new handle structure, whereas cutting a tetrahedron along a normal surface does not yield pieces that are tetrahedra, in general. Slicing along normal surfaces in this manner is important for our arguments; the definition below will help us investigate the handle structures that arise.

Definition 2.1. A handle structure of a 3-manifold M is *pre-tetrahedral* if the intersection between each 0-handle and the union of the 1-handles and 2-handles is one of the following possibilities:

- (1) *tetrahedral*, as shown in the far left of Figure 1;
- (2) *semi-tetrahedral*, as shown in the middle left of Figure 1;
- (3) *a product annulus of length 3*, as shown in the middle right of Figure 1;
- (4) *a parallelity annulus of length 4*, as shown in the far right of Figure 1.

In that figure, the shaded regions denote discs that can be components of intersection between the 0-handle and ∂M , or between the 0-handle and a 3-handle. The hashed regions denote components of intersection between the 0-handle and ∂M .




FIGURE 1. Far left: tetrahedral, Middle left: semi-tetrahedral, Middle right: product annulus of length 3, Far right: Parallelity annulus of length 4

Lemma 2.2. *Let \mathcal{T} be a triangulation of a compact orientable 3-manifold M . Let S be a normal surface properly embedded in M . Then $M \setminus S$ inherits a handle structure. Moreover when M is closed, the handle structure on $M \setminus S$ is pre-tetrahedral.*

Proof. This is explained in [15, Lemma 4.4] and so we just sketch the argument. We first form the handle structure \mathcal{H} dual to \mathcal{T} . The normal surface S intersects each i -handle in discs. When the i -handle is cut along these discs, the result is a collection of i -handles in the required handle structure for $M \setminus S$. \square

We will measure the size of a triangulation of a 3-manifold using the following quantity.

Definition 2.3. The *complexity* of a triangulation \mathcal{T} of a compact 3-manifold is the number of tetrahedra of \mathcal{T} and is denoted $\Delta(\mathcal{T})$.

The corresponding definition for a pre-tetrahedral handle structure is somewhat more complicated.

Definition 2.4. Let H_0 be a 0-handle of a pre-tetrahedral handle structure \mathcal{H} . Let α be the number of components of intersection between H_0 and the 3-handles. Define β as follows:

- (1) $\beta = 1/2$ if H_0 is tetrahedral;
- (2) $\beta = 1/4$ if H_0 is semi-tetrahedral;
- (3) $\beta = 0$ if H_0 is a product or parallelity annulus.

Define the *complexity* of H_0 to be $(\alpha/8) + \beta$. Define the *complexity* $\Delta(\mathcal{H})$ of \mathcal{H} to be the sum of the complexities of its 0-handles.

The motivation for this definition is from the following result [15, Lemma 4.12].

Lemma 2.5. *Let \mathcal{T} be a triangulation of a closed orientable 3-manifold M . Let S be a normal surface embedded in M . Then the handle structure \mathcal{H} that $M \setminus S$ inherits, as in Lemma 2.2, satisfies $\Delta(\mathcal{H}) = \Delta(\mathcal{T})$.*

Lemma 2.6. *Let \mathcal{T} be a triangulation of a compact orientable 3-manifold M . Suppose that the intersection between any tetrahedron T of \mathcal{T} and ∂M is either empty, a vertex of T , an edge of T or a face of T . Then dual to \mathcal{T} is a pre-tetrahedral handle structure \mathcal{H} of M satisfying $\Delta(\mathcal{H}) \leq \Delta(\mathcal{T})$.*

Proof. The dual handle structure \mathcal{H} has an i -handle for each $(3-i)$ -simplex of \mathcal{T} that does not lie wholly in ∂M . When the intersection between a tetrahedron and ∂M is empty or a vertex, the dual 0-handle is tetrahedral. When the intersection between a tetrahedron and ∂M is an edge, the dual 0-handle is semi-tetrahedral.

When the intersection between a tetrahedron and ∂M is a face, the dual 0-handle is a product annulus of length 3. Since the number of 0-handles of \mathcal{H} is at most the number of tetrahedra of \mathcal{T} , and each 0-handle contributes at most 1 to $\Delta(\mathcal{H})$, we deduce that $\Delta(\mathcal{H}) \leq \Delta(\mathcal{T})$. \square

Remark 2.7. Given any triangulation \mathcal{T}' of a compact orientable 3-manifold M , we can form a triangulation \mathcal{T} of M satisfying the hypotheses of Lemma 2.6 with $\Delta(\mathcal{T}) \leq 33\Delta(\mathcal{T}')$. To do this, we attach a triangulation of $\partial M \times I$ to \mathcal{T}' . This triangulation is formed as follows. For each triangle of \mathcal{T}' in ∂M , form its product with I , which is a prism. Subdivide each of its square faces into two triangles. Then triangulate each prism by coning from a new vertex in its interior. Each prism is triangulated using 8 tetrahedra. Since the number of triangles of \mathcal{T}' in ∂M is at most $4\Delta(\mathcal{T}')$, the resulting triangulation \mathcal{T} satisfies $\Delta(\mathcal{T}) \leq 33\Delta(\mathcal{T}')$.

Definition 2.8. Let M be a compact 3-manifold with a handle structure \mathcal{H} , and let S be a subsurface of ∂M . When ∂M meets 0-handles in discs, ∂M inherits a handle structure from \mathcal{H} : an i -handle of ∂M is a component of intersection of ∂M with an $(i+1)$ -handle of \mathcal{H} . We say that \mathcal{H} is a *handle structure for the pair* (M, S) if the following all hold:

- (1) ∂S intersects each handle of ∂M in a collection of arcs;
- (2) ∂S misses the 2-handles of ∂M ;
- (3) ∂S respects the product structure of the 1-handles of ∂M .

Definition 2.9. Let \mathcal{H} be a handle structure for the pair (M, S) . A handle H of \mathcal{H} is a *parallelity handle* if it has a product structure $D^2 \times I$ such that

- (1) $D^2 \times \partial I = H \cap S$;
- (2) each component of intersection between H and any other handle is of the form $\beta \times I$ for some subset β of ∂D^2 .

For example, a product annulus of length 3 meeting S on its top and bottom, and a parallelity annulus of length 4 are both parallelity 0-handles. There will also be parallelity 1-handles and 2-handles.

Remark 2.10. The handle structure \mathcal{H} in Lemma 2.6 has no parallelity handles when viewed as a handle structure for the pair $(M, \partial M)$. To see this, note that any parallelity handle for $(M, \partial M)$ is adjacent to a parallelity 2-handle. However, a parallelity 2-handle of \mathcal{H} is dual to an edge of \mathcal{T} with both endpoints in ∂M but with interior in the interior of M . However, if there were such an edge of \mathcal{T} , then the intersection between any adjacent tetrahedron and ∂M would violate the hypothesis of Lemma 2.6.

Definition 2.11. The *parallelity bundle* for \mathcal{H} is the union of the parallelity handles.

It was shown in [13, Lemma 3.3] that the I -bundle structures on the parallelity handles can be chosen to patch together to form an I -bundle structure on the parallelity bundle. We therefore use the following standard terminology.

Definition 2.12. Let B be an I -bundle over a surface F . Its *horizontal boundary* $\partial_h B$ is the (∂I) -bundle over F . Its *vertical boundary* $\partial_v B$ is the I -bundle over ∂F .

It is often very useful to enlarge the parallelity bundle, forming the following structure.




FIGURE 2. An annular simplification removes P and replaces A with A' . Left: an annular simplification where P lies within a ball. Right: a cross-section of an annular simplification where P is a product region between A and A' .

Definition 2.13. Let M be a compact orientable 3-manifold and let S be a subsurface of ∂M . Let \mathcal{H} be a handle structure for (M, S) . A *generalised parallelity bundle* \mathcal{B} is a 3-dimensional submanifold of M such that

- (1) \mathcal{B} is an I -bundle over a compact surface;
- (2) the horizontal boundary $\partial_h \mathcal{B}$ of \mathcal{B} is the intersection between \mathcal{B} and S ;
- (3) \mathcal{B} is a union of handles of \mathcal{H} ;
- (4) any handle of \mathcal{B} that intersects $\partial_v \mathcal{B}$ is a parallelity handle, where the I -bundle structure on the parallelity handle agrees with the I -bundle structure of \mathcal{B} ;
- (5) whenever a handle of \mathcal{H} lies in \mathcal{B} then so do all incident handles of \mathcal{H} with higher index;
- (6) the intersection between $\partial_h \mathcal{B}$ and the non-parallelity handles lies in a union of disjoint discs in the interior of S .

Note that condition (6) is included in the definition given in [15] but is not in some earlier work [13].

The main reason why this is such a useful notion is the fact that frequently we may ensure that the horizontal boundary of a generalised parallelity bundle is incompressible.

Definition 2.14. Let M be a compact orientable irreducible 3-manifold and let S be a subsurface of ∂M . Let \mathcal{H} be a handle structure for (M, S) . Suppose M contains the following:

- (1) an annulus A' that is a vertical boundary component of a generalised parallelity bundle \mathcal{B} ;
- (2) an annulus A contained in S such that $\partial A = \partial A'$;
- (3) a 3-manifold P with $\partial P = A \cup A'$ such that P either lies in a 3-ball or is a product region between A and A' .

Suppose also that P is a union of handles of \mathcal{H} , that whenever a handle of \mathcal{H} lies in P , so do all incident handles with higher index, and that any parallelity handle of \mathcal{H} that intersects P lies in P . Finally, suppose that apart from the component of the generalised parallelity bundle incident to A' , all other components of \mathcal{B} in P are I -bundles over discs.

An *annular simplification* of the 3-manifold M is the manifold obtained by removing the interiors of P and A from M ; see Figure 2.

Lemma 2.15. *Let M be a compact orientable irreducible 3-manifold and let S be an incompressible subsurface of ∂M that is not a 2-sphere. Let \mathcal{H} be a handle structure for (M, S) . Let \mathcal{B} be a generalised parallelity bundle that is maximal, in the sense that it is not a proper subset of another generalised parallelity bundle. Suppose that \mathcal{H} admits no annular simplification. Then \mathcal{B} contains every parallelity handle of \mathcal{H} , and moreover, each component of \mathcal{B} :*

- (1) *has incompressible horizontal boundary, and*
- (2) *either has incompressible vertical boundary, or is an I -bundle over a disc.*

Proof. This is stated in [13, Corollary 5.7]. However, a slightly different definition of generalised parallelity bundle is used there that omits Condition 6 in Definition 2.13. This extra condition does not affect the argument there.

Alternatively, one can argue as follows. In [15, Theorem 6.18], we showed that \mathcal{B} contains every parallelity handle of \mathcal{H} , and every component either satisfies the conclusion of the lemma, or is a special case called boundary-trivial. In the boundary-trivial case, the component of \mathcal{B} lies within a 3-ball; the precise definition is [15, Definition 6.16]. However, [15, Lemma 6.17] implies that a boundary-trivial component admits an annular simplification. Thus we cannot have such components by hypothesis. \square

The *weight* of a surface properly embedded in a manifold M , in general position with respect to a triangulation \mathcal{T} , is defined to be the number of intersections between S and the edges of \mathcal{T} .

The following is [15, Lemma 6.15].

Theorem 2.16. *Let \mathcal{T} be a triangulation of a compact orientable irreducible 3-manifold M . Let S be an orientable incompressible normal surface properly embedded in M that has least weight, up to isotopy supported in the interior of M . Let \mathcal{H} be the handle structure that $M' = M \setminus S$ inherits, as in Lemma 2.2. Let $S' = \partial M' \setminus \partial M$. Then (M', S') admits no annular simplification. Hence, the parallelity bundle for \mathcal{H} extends to a maximal generalised parallelity bundle that has incompressible horizontal boundary.*

The following is [15, Lemma 8.14].

Lemma 2.17. *Let \mathcal{H} be a pre-tetrahedral handle structure of a pair (M, S) , and let \mathcal{B} be its parallelity bundle. Then the length of $\partial_v \mathcal{B}$, which is its number of 2-cells, is at most $56\Delta(\mathcal{H})$. Similarly if \mathcal{B}' is a maximal generalised parallelity bundle, then the length of $\partial_v \mathcal{B}'$ is at most $56\Delta(\mathcal{H})$.*

Definition 2.18. Let \mathcal{H} be a handle structure of a compact 3-manifold. Then the *associated cell structure* is obtained as follows:

- (1) each handle is a 3-cell;
- (2) each component of intersection between two handles or between a handle and ∂M is a 2-cell;
- (3) each component of intersection between three handles or between two handles and ∂M is a 1-cell;
- (4) each component of intersection between four handles or between three handles and ∂M is a 0-cell.

Lemma 2.19. *Let \mathcal{H} be a pre-tetrahedral handle structure of a compact orientable 3-manifold M . Suppose that \mathcal{H} has no parallelity 0-handles. Let \mathcal{C} be the associated*




FIGURE 3. Clipping a semi-tetrahedral 0-handle.

cell structure. Let \mathcal{T} be the triangulation obtained by placing a vertex in the interior of each 2-cell and coning off, and then placing a vertex in the interior of each 3-cell and coning off. Then $\Delta(\mathcal{T}) \leq 2304\Delta(\mathcal{H})$.

Proof. We first estimate how many triangles there are in the 2-skeleton of \mathcal{C} . Let \mathcal{H}^j denote the union of j handles of \mathcal{H} . Thus each 2-cell is a component of $\mathcal{H}^i \cap \mathcal{H}^j$ for $i \neq j \in \{0, 1, 2\}$ or of $\mathcal{H}^j \cap (\mathcal{H}^3 \cup \partial M)$ for $j < 3$.

There are as many triangles in $\mathcal{H}^2 \cap \mathcal{H}^1$ as in $\mathcal{H}^2 \cap \mathcal{H}^0$. Similarly, there are as many triangles in $\mathcal{H}^2 \cap (\mathcal{H}^3 \cup \partial M)$ as in $\mathcal{H}^2 \cap \mathcal{H}^0$. There are as many triangles in $\mathcal{H}^1 \cap \mathcal{H}^2$ as in $\mathcal{H}^1 \cap (\mathcal{H}^3 \cup \partial M)$. There are as many triangles in $\mathcal{H}^1 \cap \mathcal{H}^0$ as in $\mathcal{H}^1 \cap \mathcal{H}^2$. There are as many triangles in $\mathcal{H}^0 \cap (\mathcal{H}^3 \cup \partial M)$ as in $\mathcal{H}^0 \cap \mathcal{H}^2$. Each component of $\mathcal{H}^2 \cap \mathcal{H}^0$ is triangulated using 4 triangles. Hence, we see that the total number of triangles in the 2-skeleton of \mathcal{C} is $24|\mathcal{H}^2 \cap \mathcal{H}^0|$. Since \mathcal{H} is pre-tetrahedral, each 0-handle meets at most six 2-handles. Since \mathcal{H} has no parallelity 0-handles, each 0-handle contributes at least $1/8$ to $\Delta(\mathcal{H})$. Thus $24|\mathcal{H}^2 \cap \mathcal{H}^0|$ is at most $1152\Delta(\mathcal{H})$. Each tetrahedron of \mathcal{T} has a triangle in \mathcal{C} as a face, and each triangle in \mathcal{C} is a face of at most two tetrahedra. Hence, $\Delta(\mathcal{T}) \leq 2304\Delta(\mathcal{H})$. \square

In Section 9, one of our arguments will replace some semi-tetrahedral 0-handles in a pre-tetrahedral handle structure by 0-handles modified as follows. The boundary of a semi-tetrahedral 0-handle has two 1-handles that are bordered by exactly two 2-handles. We replace the union of one of these 1-handles and the adjacent 2-handles by a single 2-handle. For any semi-tetrahedral 0-handle, this replacement may be done on either one or both of its relevant 1-handles. We call the result a *clipped semi-tetrahedral 0-handle*. One clipped semi-tetrahedral 0-handle is shown in Figure 3.

We may define the complexity of a handle structure that is pre-tetrahedral aside from a finite number of clipped semi-tetrahedral 0-handles just as in Definition 2.4 by setting β to be $1/4$ for each clipped semi-tetrahedral 0-handle, and leaving the definition the same otherwise. Then we may modify Lemma 2.19 as follows.

Lemma 2.20. *Let \mathcal{H} be a handle structure of a compact orientable 3-manifold M . Suppose that \mathcal{H} is pre-tetrahedral, aside from a finite number of clipped semi-tetrahedral 0-handles. Suppose also that \mathcal{H} has no parallelity 0-handles. Let \mathcal{C} be the associated cell structure. Let \mathcal{T} be the triangulation obtained by placing a vertex in the interior of each 2-cell and coning off, and then placing a vertex in the interior of each 3-cell and coning off. Then $\Delta(\mathcal{T}) \leq 2304\Delta(\mathcal{H})$.*

Proof. The proof is identical to that of Lemma 2.19, since a clipped semi-tetrahedral 0-handle still meets at most six 2-handles, and contributes at least $1/4$ to $\Delta(\mathcal{H})$. \square

3. VERTICAL ARCS IN PRODUCTS

As discussed in the introduction, a central part of the paper will be an analysis of triangulations of $S \times [0, 1]$, where S is a closed orientable surface. We will want to transfer results about $S \times [0, 1]$ to results about $S' \times [0, 1]$ where S' is a branched cover of S . The branching locus will be an arc of the following form.

Definition 3.1. Let S be a closed surface. An arc properly embedded in $S \times [0, 1]$ is *vertical* if it is ambient isotopic to $\{*\} \times [0, 1]$ for some point $*$ in S .

The main result of this section is as follows.

Theorem 3.2. *Let S be a closed connected orientable surface. Let \mathcal{T} be a triangulation of $S \times [0, 1]$. Then the 23rd iterated barycentric subdivision $\mathcal{T}^{(23)}$ contains an arc in its 1-skeleton that is vertical.*

This will be proved using some normal surface theory. The following basic result in the theory is contained in [19, Proposition 3.3.24, Corollary 3.3.25].

Lemma 3.3. *Let M be a compact orientable irreducible 3-manifold with incompressible boundary, and let \mathcal{T} be a triangulation of M . Let S be an incompressible boundary-incompressible surface properly embedded in M , no component of which is a sphere or disc, and that is in general position with respect to \mathcal{T} . Then there is an ambient isotopy taking S to a normal surface with weight no greater than that of S . Moreover, if C is any boundary curve of S that is normal and intersects each edge of \mathcal{T} at most once, then the isotopy can be chosen to leave C fixed.*

Proposition 3.4. *Let \mathcal{T} be a triangulation of a compact 3-manifold M . Let S be a 2-sided normal surface properly embedded in M . Let S' be the copies of S in $M \setminus S$. Let \mathcal{B} be the parallelity bundle for the pair $(M \setminus S, S')$. Let α be an arc properly embedded in M with the following properties.*

- (1) *It lies within a copy of S in $M \setminus S$.*
- (2) *It is disjoint from the horizontal boundary of \mathcal{B} .*
- (3) *Its intersection with each normal triangle or square of S is either empty or a single properly embedded arc with endpoints on distinct edges of the triangle or square.*

Then α is simplicial in $\mathcal{T}^{(23)}$.

Proof. Within each tetrahedron of \mathcal{T} , the normal discs of S come in at most 5 types. Let D be the union of the outermost discs of each type. These discs within a single tetrahedron intersect each face of \mathcal{T} in at most 8 arcs. However, each face of \mathcal{T} might be adjacent to two tetrahedra of \mathcal{T} and there is no reason for the 8 arcs coming from the two adjacent tetrahedra to coincide. So, the intersection between D and any face of \mathcal{T} consists of at most 16 normal arcs. By [16, Lemma 6.7], the union of the arcs is simplicial in $\mathcal{T}^{(6)}$. Within each tetrahedron of \mathcal{T} , D consists of at most 10 normal discs. Hence, by [16, Lemma 6.11], we may use at most 16 further subdivisions to make these discs simplicial.

We now apply one further subdivision to the triangulation, forming $\mathcal{T}^{(23)}$. We may assume that the intersection between α and the 2-skeleton of \mathcal{T} is a union of vertices of $\mathcal{T}^{(23)}$. We may further isotope α so that it is simplicial. This follows from the general result that an arc in triangulated polygon may be isotoped to be simplicial in the barycentric subdivision. Moreover, if the endpoints of the arc

are already vertices of this subdivision, then the isotopy can keep these endpoints fixed. \square

Proof of Theorem 3.2. Suppose first that S is a 2-sphere. Pick any properly embedded simplicial arc in $\mathcal{T}^{(23)}$ joining $S \times \{0\}$ to $S \times \{1\}$. By the lightbulb trick, this is ambient isotopic to an arc of the form $\{*\} \times [0, 1]$, as required.

Thus, we may assume that S is not a 2-sphere. Hence, $S \times \{0\}$ contains an essential simple closed curve. Pick one, C , that is transverse to the 1-skeleton of \mathcal{T} on $S \times \{0\}$ and that intersects each edge of that 1-skeleton at most once. Then C is normal. It is not hard to prove that such a curve must exist; for example, we can take C to be non-trivial in $H_1(S; \mathbb{Z}/2\mathbb{Z})$ and with fewest points of intersection with the edges. Let A be the annulus $C \times [0, 1]$. By Lemma 3.3, this can be isotoped, without moving $C \times \{0\}$, to a normal surface. We pick A to have least weight among all annuli with one boundary component equal to $C \times \{0\}$ and the other boundary component on $S \times \{1\}$. Let $C \times \{1\}$ be the other boundary component of A , which is then a normal simple closed curve in $S \times \{1\}$.

In the case where S is a torus, we need to be more precise about the choice of annulus A , as follows. Pick an oriented vertical arc in $S \times [0, 1]$ disjoint from A . Then the *winding number* of an oriented annulus, with boundary curves equal to ∂A , is the signed intersection number of the annulus with this vertical arc. For each winding number t , let $w(t)$ be the minimal weight of a normal annulus with boundary equal to ∂A and winding number t . If there is no normal annulus with a given winding number t , then we define $w(t)$ to be infinite. Note that $w(t)$ tends to infinity as $t \rightarrow \pm\infty$. Now, A has least weight among all normal annuli with the given boundary curves. Hence, $w(0)$ is a global minimum. However, there may be other values of t such that $w(t) = w(0)$. Choose t_0 to be maximal with this property. We replace A by a normal annulus, having the same weight and the same boundary curves, but with winding number t_0 . Call this new annulus A .

Let M be the 3-manifold $(S \times [0, 1]) \setminus A$. Let \tilde{A} be the two copies of A in ∂M . Let \mathcal{B} be the parallelity bundle for the pair (M, \tilde{A}) . This consists of the union of the regions between parallel normal discs of A . By choice of C , the curve $C \times \{0\}$ intersects each edge of \mathcal{T} at most once. Thus no normal disc of A incident to $C \times \{0\}$ is parallel to another normal disc of A . Hence, \mathcal{B} misses $S \times \{0\}$. By Theorem 2.16, \mathcal{B} extends to a maximal generalised parallelity bundle \mathcal{B}_+ that has incompressible horizontal boundary. Its vertical boundary is a union of vertical boundary components of \mathcal{B} , and hence it also misses $S \times \{0\}$.

Since $\partial_h \mathcal{B}_+$ is an incompressible subsurface of the annuli \tilde{A} , it is a collection of annuli and discs. Hence, each component of \mathcal{B}_+ is an I -bundle over a disc, annulus or Möbius band.

Claim 1. No component of \mathcal{B}_+ is an I -bundle over a Möbius band.

The I -bundle over a core curve of this Möbius band would be a Möbius band embedded in $S \times [0, 1]$ with boundary in \tilde{A} . We could then attach an annulus to its boundary, to create a Möbius band embedded in $S \times [0, 1]$ with boundary in $S \times \{0\}$. We could then double $S \times [0, 1]$ along $S \times \{0\}$ to create another copy of $S \times [0, 1]$ containing a Klein bottle. We could then embed this in the 3-sphere, which is well known to be impossible.

Claim 2. No component of \mathcal{B}_+ is an I -bundle over an annulus that intersects both components of \tilde{A} .

Let B be such a component. Since $\partial_h B$ consists of incompressible annuli, and because these annuli are disjoint from $S \times \{0\}$, at least one boundary component of $\partial_h B$ is disjoint from $\partial \tilde{A}$. It is a core curve of a component of \tilde{A} . Let V be the vertical boundary component of \mathcal{B}_+ incident to this core curve; recall V lies in the parallelity bundle \mathcal{B} . Then V is disjoint from $S \times \{0, 1\}$ and ∂V consists of core curves disjoint from $\partial \tilde{A}$. Note that V specifies a free homotopy between the two boundary curves of $S \setminus C$. Hence, we deduce in this case that S is a torus.

Let A_1 and A_2 be the two components of \tilde{A} . Each A_i is divided into smaller annuli A'_i and A''_i by ∂V , where A'_i is the component intersecting $S \times \{0\}$. We can construct two annuli $A'_1 \cup V \cup A''_2$ and $A'_2 \cup V \cup A''_1$. Using a small isotopy supported in the interior of $S \times [0, 1]$, these annuli can be made normal. Both of these have the same boundary curves as A . One has winding number one less than A , the other has winding number one more than A . The one with winding number greater than A has, by our choice of A , weight strictly greater than A . But the sum of the weights of $A'_1 \cup V \cup A''_2$ and $A'_2 \cup V \cup A''_1$ is twice the weight of A . Hence, the other annulus has weight less than that of A . But A was chosen to have minimal weight, which is a contradiction.

Claim 3. Each annular component of $\partial_h \mathcal{B}_+$ is disjoint from $\partial \tilde{A}$.

Let B be any component of \mathcal{B}_+ that is an I -bundle over an annulus. By Claim 2, its two horizontal boundary components both lie in the same component of \tilde{A} . Call this component A_1 . Note each component of $\partial_h B$ contains a core curve of A_1 , because $\partial_h B$ is essential. Suppose that one component of $\partial_h B$ is an annulus H intersecting $\partial \tilde{A}$. Because \mathcal{B} misses $S \times \{0\}$, H must meet $S \times \{1\}$. Then $\partial_v B$, which lies in the parallelity bundle, intersects $S \times \{1\}$. It follows that the other component of $\partial_h B$ also intersects $S \times \{1\}$. Since H intersects $S \times \{1\}$ by assumption, the other components of $A_1 \setminus H$ are an annulus incident to $S \times \{0\}$ and possibly discs incident to $S \times \{1\}$. But the other component of $\partial_h B$ also intersects $S \times \{1\}$, and so it must lie in one of these discs. But it cannot then contain a core curve of A_1 , which is a contradiction.

Claim 4. There are no annular components of $\partial_h \mathcal{B}_+$.

Let B be any component of \mathcal{B}_+ that is an I -bundle over an annulus. By Claim 3, $\partial_h B$ is disjoint from $S \times \{0, 1\}$ and by Claim 2, it lies in a single component of \tilde{A} . Let V be any vertical boundary component of B . Then ∂V cobounds an annulus A' in A . If we remove A' from A and replace it by V , the result is an annulus with the same boundary as A but with smaller weight. By Lemma 3.3, we may isotope this to a normal annulus without increasing its weight and without moving its intersection curve with $S \times \{0\}$. This contradicts our choice of A .

We are now in a position to prove the theorem. Since \mathcal{B}_+ consists only of I -bundles over discs, we may find an arc α in \tilde{A} running from $S \times \{0\}$ to $S \times \{1\}$ and that avoids \mathcal{B}_+ . We can choose α with the property that it intersects each triangle or square of \tilde{A} in a single properly embedded arc with endpoints on distinct edges of the triangle or square. Thus, α satisfies the hypotheses of Proposition 3.4. It therefore is simplicial in $\mathcal{T}^{(23)}$. It is the required vertical arc. \square

The following lemma will be useful when modifying a given triangulation of $S \times [0, 1]$.

Lemma 3.5. *Let \mathcal{T} be a triangulation of $S \times [0, 1]$ and let α be a simplicial arc that is vertical in $S \times [0, 1]$. Let \mathcal{T}' be a triangulation obtained from \mathcal{T} by attaching a*

tetrahedron to $S \times [0, 1]$ to realise a Pachner move of the boundary triangulation. Then α extends to a simplicial arc α' in \mathcal{T}' that is also vertical.

Proof. The attachment of the tetrahedron realises a Pachner move on the boundary that has type 1-3, 2-2 or 3-1. In the cases of a 1-3 move and a 2-2 move, the arc α remains properly embedded and vertical, and so in these cases, we set α' to be α . In the case of a 3-1 Pachner move, the new tetrahedron is incident to three triangles that meet at a vertex. If α does not end at that vertex, then we again set α' to be α . If α does end at that vertex, then we form α' by adding one of the edges that is incident to two of the triangles. This is vertical. \square

4. SPINES, TRIANGULATIONS AND MAPPING CLASS GROUPS

In this section, we define a graph associated with a closed orientable surface S , the spine graph $\text{Sp}(S)$ on S . We show $\text{Sp}(S)$ is quasi-isometric to the triangulation graph $\text{Tr}(S)$ defined in the introduction. We also obtain properties of spines and methods of modifying them that we will use in future arguments.

Recall the triangulation graph $\text{Tr}(S)$ defined in the introduction. Related to the triangulation graph is the spine graph, defined as follows.

Definition 4.1. A *spine* for a closed orientable surface S is a graph Γ embedded in S that has no vertices of degree 1 or 2 and where $S \setminus \Gamma$ is a disc.

Definition 4.2. In an *edge contraction* on a spine Γ , one collapses an edge that joins distinct vertices, thereby amalgamating these vertices into a single vertex. An *edge expansion* is the reverse of this operation.

Definition 4.3. The *spine graph* $\text{Sp}(S)$ for a closed orientable surface S is a graph defined as follows. It has a vertex for each spine of S , up to isotopy of S . Two vertices are joined by an edge if and only if their spines differ by an edge contraction or expansion.

We wish to compare the spine graph and triangulation graph. Dual to each 1-vertex triangulation is a spine. Each 2-2 Pachner move on a 1-vertex triangulation has the following effect on the dual spines: contract an edge and then expand. Thus, each edge in $\text{Tr}(S)$ maps to a concatenation of two edges in $\text{Sp}(S)$. We therefore get a map $\text{Tr}(S) \rightarrow \text{Sp}(S)$.

It will also be useful to recall the following variant of the triangulation graph [15, Definition 2.6].

Definition 4.4. Let S be a closed orientable surface and let n be a positive integer. Then $\text{Tr}(S; n)$ denotes the space of triangulations with at most n vertices. This is a graph with a vertex for each isotopy class of such triangulations, and with an edge for each 2-2, 3-1, or 1-3 Pachner move between them.

There is an obvious inclusion $\text{Tr}(S) \rightarrow \text{Tr}(S; n)$ for any positive integer n . Note also that the mapping class group of S acts on $\text{Tr}(S)$, $\text{Tr}(S; n)$ and $\text{Sp}(S)$ by isometries. Moreover, this action is properly discontinuous and cocompact. Hence, the mapping class group of S is quasi-isometric to each of $\text{Tr}(S)$, $\text{Tr}(S; n)$ and $\text{Sp}(S)$, via an application of the Milnor-Švarc lemma ([2, Proposition 8.19]). In fact, we obtain the following result.

Lemma 4.5. *The maps $\text{Tr}(S) \rightarrow \text{Tr}(S; n)$ and $\text{Tr}(S) \rightarrow \text{Sp}(S)$ are quasi-isometries.*

Proof. Pick a 1-vertex triangulation T for S . By the Milnor-Švarc lemma, the map $\text{MCG}(S) \rightarrow \text{Tr}(S)$ sending $g \in \text{MCG}(S)$ to gT is a quasi-isometry; see, for example [2, Proposition 8.19]. A quasi-inverse is given as follows. For T fixed and any point p in $\text{Tr}(S)$, pick a triangulation of the form gT that is closest to p . Then the quasi-inverse sends p to g . The composition of this quasi-inverse with $\text{MCG}(S) \rightarrow \text{Tr}(S; n)$ is a quasi-isometry $\text{Tr}(S) \rightarrow \text{Tr}(S; n)$. There is a uniform upper bound to the distance between the image of a triangulation under this map $\text{Tr}(S) \rightarrow \text{Tr}(S; n)$ and its image under the inclusion map. Hence, the inclusion map is also a quasi-isometry as required.

The argument for $\text{Tr}(S) \rightarrow \text{Sp}(S)$ is identical. \square

A modification that one can make to a spine that is slightly more substantial than an edge contraction or expansion is as follows.

Definition 4.6. Let Γ be a spine for a closed surface S . Let e_1 be an arc properly embedded in the disc $S \setminus \Gamma$. Let e_2 be an edge of the graph $\Gamma \cup e_1$ that has distinct components of $S \setminus (\Gamma \cup e_1)$ on either side of it. Then the result of removing e_2 from Γ and adding e_1 is a new spine Γ' for S . We say that Γ and Γ' are related by an *edge swap*.

The following is [15, Lemma 8.3].

Lemma 4.7. *Let S be a closed orientable surface. Let Γ be a spine for S . Then an edge swap can be realised by a sequence of at most $24g(S)$ edge expansions and contractions.*

Definition 4.8. Let S be a closed orientable surface with a cell structure. A spine for S is *cellular* if it is a subcomplex of the 1-skeleton of the cell complex. The *length* of this spine is the number of 1-cells that it contains.

The following is [15, Corollary 8.8].

Lemma 4.9. *Let S be a closed orientable surface with a cell structure \mathcal{C} , and with a cellular spine Γ . Let D_1, \dots, D_m be cellular subsets of S , each of which is an embedded disc, and with disjoint interiors. Let ℓ be the sum of the lengths of $\partial D_1, \dots, \partial D_m$. Then there is a sequence of at most $6mg(S) + 2\ell$ edge swaps taking Γ to a cellular spine Γ' that is disjoint from the interior of D_1, \dots, D_m .*

Remark 4.10. A slight strengthening of the lemma remains true, with the same proof. Instead of D_1, \dots, D_m being embedded discs, we can allow them to be the images of immersed discs in S , where the restriction of the immersion to $\text{int}(D_1) \cup \dots \cup \text{int}(D_m)$ is an embedding. In other words, we allow the boundaries of the discs to self-intersect and to intersect each other.

The following is a version of Lemma 4.9 dealing with both discs and annuli.

Lemma 4.11. *Let S be a closed orientable surface with a cell structure \mathcal{C} , and with a cellular spine Γ . Let A_1, \dots, A_m be cellular subsets of S , each of which is the image of an immersed disc or annulus, and where the restriction of the immersion to the interior of these discs and annuli is an embedding. Let ℓ be the sum of the lengths of $\partial A_1, \dots, \partial A_m$. Then there is a sequence of at most $6mg(S) + 16g(S) + 2m + 2\ell$ edge swaps taking Γ to a cellular spine Γ' that is disjoint from the interior of the disc components of A_1, \dots, A_m and that intersects the interior of each annular component in at most one essential embedded arc. Moreover, this arc is a subset of the original spine Γ .*

Proof. In each essential annular component, there must be an essential properly embedded arc that is a subset of Γ , as otherwise the disc $S \setminus \Gamma$ would contain a core curve of this annulus. Pick one such arc in each essential annular component. Let α be the union of these arcs. If A_i is an essential annulus, then define D_i to be $A_i \setminus \alpha$. If A_i is an inessential annulus, then let D_i be the disc in S containing A_i that has boundary equal to a component of ∂A_i . If A_i is a disc, let D_i be A_i . Some of these discs may be nested, in which case discard the smaller disc. Thus, D_1, \dots, D_m is a collection of discs as in Remark 4.10. So, there is a sequence of edge swaps taking Γ to a cellular spine that is disjoint from the interior of D_1, \dots, D_m . It must intersect the interior of each essential annulus A_i in a single arc.

Unfortunately, the number of these edge swaps is bounded above by a linear function of the total length of the boundary of D_1, \dots, D_m , which depends not just on ℓ but also on the length of the arcs α . To deal with this, we define a new cell structure \mathcal{C}' on S , as follows. Away from $A_1 \cup \dots \cup A_m$, this agrees with \mathcal{C} , but within each A_i , the 2-cells are the components of $A_i \setminus \Gamma$. Note that Γ is still cellular with respect to \mathcal{C}' . We can now bound the length of $\partial D_1, \dots, \partial D_m$ in \mathcal{C}' in terms of ℓ . This length is at most the length of $\partial A_1, \dots, \partial A_m$ plus twice the length of α . The length of α with respect to \mathcal{C}' is at most the number of vertices of Γ plus the number of essential annular components of A_1, \dots, A_m . By an Euler characteristic argument, using the fact that each vertex has degree at least three, the number of vertices of Γ is at most $4g(S) - 2$. The number of essential annular components of A_1, \dots, A_m is at most m . So, the length of $\partial D_1, \dots, \partial D_m$ in \mathcal{C}' is at most $\ell + 8g(S) + 2m$. Now apply Lemma 4.9 and Remark 4.10 to turn Γ into a spine that is cellular in \mathcal{C}' and that intersects the interior of $A_1 \cup \dots \cup A_m$ in the arcs α . It is then cellular with respect to \mathcal{C} . \square

Lemma 4.12. *Let \mathcal{T} and \mathcal{T}' be triangulations of a closed surface S that differ by a sequence of n Pachner moves. Let Γ be a subcomplex of \mathcal{T} that is a spine of S . Then there is a sequence of at most n edge swaps and some isotopies taking Γ to a spine that is a subcomplex of \mathcal{T}' .*

Proof. It suffices to consider the case where $n = 1$, and so \mathcal{T} and \mathcal{T}' differ by a single Pachner move. Some terminology: throughout this proof, edges will refer to edges in a spine, with endpoints on vertices of the spine of valence at least 3. We refer to edges of the triangulation \mathcal{T} , which are not necessarily edges of Γ even when they lie in Γ , by 1-cells.

If the Pachner move is a 1-3 move, then there is nothing to prove as the 1-skeleton of \mathcal{T} is then a subcomplex of \mathcal{T}' .

Suppose it is a 2-2 move, removing a 1-cell e and inserting a new 1-cell e' . If e is not part of Γ , then Γ is a subcomplex of \mathcal{T}' and so no edge swaps are required. So suppose that e is contained in Γ . Since $S \setminus \Gamma$ is a disc, there is an arc α running from the midpoint of e back to the midpoint of e but on the other side of e and that is otherwise disjoint from Γ . We may assume that α is disjoint from the vertices of \mathcal{T} and intersects each 1-cell of \mathcal{T} at most once. It must intersect at least one 1-cell e'' in the boundary of the square that is the union of the two triangles involved in the Pachner move. If both endpoints of e'' lie in Γ , let $e''' = e''$. Otherwise, let e''' be the union of e'' and the third 1-cell of the triangle formed by e'' and e . We perform the edge swap that adds e''' to Γ and removes the edge of Γ containing e .

We now consider a 3-1 move. Let e_1, e_2 and e_3 be the three 1-cells of \mathcal{T} that are removed. If none of these are part of Γ , then we leave the spine unchanged. There

cannot be just one of these 1-cells in Γ , since no vertex of Γ has degree 1. Suppose Γ runs over exactly two 1-cells in $\{e_1, e_2, e_3\}$, say e_1 and e_2 . These are two 1-cells of a triangle of \mathcal{T} . The third 1-cell e' of this triangle cannot lie in Γ , as $S \setminus \Gamma$ is a single disc. Hence, we may isotope $e_1 \cup e_2$ across the triangle to e' . Suppose finally that all three of e_1, e_2 and e_3 are part of Γ . Let e be the other 1-cell of the triangle formed by e_1 and e_2 . Adding e to Γ and removing e_1 is an edge swap. We then remove e_2 and e_3 and add the third 1-cell of the triangle that they span. This is realised by an isotopy of the spine and so no edge swap is required. \square

Lemma 4.13. *Let \mathcal{T} be a triangulation of a torus with v vertices. Then there is a sequence of at most $4v$ Pachner moves taking \mathcal{T} to a 1-vertex triangulation.*

Proof. This is contained in the proof of [14, Proposition 10.3], and so we only sketch the argument. Suppose $v > 1$, as otherwise we are done. The strategy is to apply at most 4 Pachner moves to the triangulation, after which the number of vertices is reduced.

If there is an edge of the triangulation with the same triangle on both sides, then one endpoint of the edge is a vertex with valence 1. It is possible to apply two 2-2 Pachner moves to increase this valence to 3. Then one can apply a 3-1 Pachner move to remove this vertex.

So we may suppose that every edge of the triangulation has distinct triangles on both sides. Using the fact that the Euler characteristic of the torus is zero, there is a vertex with valence at most 6. Suitably chosen 2-2 Pachner moves then reduce this to 3. A 3-1 Pachner move can then be used to remove the vertex. \square

Lemma 4.14. *Let \mathcal{T} be a triangulation of a compact surface S with t triangles. Then the barycentric subdivision $\mathcal{T}^{(1)}$ is obtained from \mathcal{T} by $4t$ Pachner moves and an isotopy.*

Proof. First perform a 1-3 Pachner move to each triangle of \mathcal{T} . Each original edge of \mathcal{T} is then adjacent to two new triangles. Choose one of the two, and perform a 1-3 Pachner move in that triangle. Then perform the 2-2 move that removes the edge. The resulting triangulation is isotopic to $\mathcal{T}^{(1)}$. In total, we have performed $4t$ Pachner moves. \square

Lemma 4.15. *Let T be a torus equipped with a cell structure. Let Γ be a cellular spine for T . Let C be a cellular essential simple closed curve with length ℓ . Then there exists a spine for T that is obtained from Γ by at most $24 + 4\ell$ edge swaps and that contains C .*

Proof. The annulus $T \setminus C$ has boundary length 2ℓ . Apply Lemma 4.11 to turn Γ into a spine Γ' that intersects the interior of $T \setminus C$ in a single arc, using at most $24 + 4\ell$ edge swaps. The spine Γ' must contain all of C , as otherwise $S \setminus \Gamma'$ would contain an essential simple closed curve. \square

5. TRIANGULATIONS OF A TORUS

In this section, we recall a description of the space $\text{Tr}(T^2)$ of all 1-vertex triangulations of a torus.

The *Farey graph* is a graph with vertex set $\mathbb{Q} \cup \{\infty\}$, and where two vertices p/q and r/s are joined by an edge if and only if $|ps - qr| = 1$. Here, we assume that the fractions are in their lowest terms and that $\infty = 1/0$. Now, $\mathbb{Q} \cup \{\infty\}$ is a subset




FIGURE 4. The Farey graph and the dual Farey tree

of $\mathbb{R} \cup \{\infty\}$, which is the circle at infinity of the upper-half plane. We can realise each edge of the Farey graph as an infinite geodesic in the hyperbolic plane; see Figure 4. The edges of the Farey graph form the edges in a tessellation of \mathbb{H}^2 by ideal triangles. We call this the *Farey tessellation*. Each triangle has three points on the circle at infinity, and these correspond to three slopes on the torus, with the property that any two of these slopes intersect once. Given three such slopes, we can realise them as Euclidean geodesics in the torus, which we think of as $\mathbb{R}^2/\mathbb{Z}^2$. We can arrange that these geodesics each go through the image of the origin, and hence all intersect at this point. Thus, this forms a 1-vertex triangulation of the torus. Conversely, given any 1-vertex triangulation of the torus, we may isotope the vertex to the origin, and then isotope each of the edges to Euclidean geodesics. Thus, we see that there is a 1-1 correspondence between 1-vertex triangulations of the torus, up to isotopy, and ideal triangles in the Farey tessellation.

When a 2-2 Pachner move is performed, this removes one of the edges of the triangulation, forming a square, and then inserts the other diagonal of the square. The remaining two edges of the triangulation are preserved, and these correspond to an edge of the Farey graph. Thus we see that two triangulations differ by a 2-2 Pachner move if and only if their corresponding ideal triangles in the Farey tessellation share an edge of the Farey graph.

It is natural to form the dual of the Farey tessellation, which is the *Farey tree*. This has a vertex for each ideal triangle of the Farey tessellation, and two vertices of the Farey tree are joined by an edge if and only if the dual triangles share an edge. As the name suggests, this is a tree. The above discussion has the following immediate consequence.

Theorem 5.1. *The graph $\text{Tr}(T^2)$ is isomorphic to the Farey tree.* □

5.1. A Cayley graph for $\text{PSL}(2, \mathbb{Z})$. The mapping class group of the torus is isomorphic to $\text{SL}(2, \mathbb{Z})$. It was shown by Serre [25] that $\text{SL}(2, \mathbb{Z})$ is isomorphic to the amalgamated free product of \mathbb{Z}_4 and \mathbb{Z}_6 , amalgamated over the subgroups of order 2. The non-trivial element in the amalgamating subgroup is the matrix $-I$, which is central. If we quotient by this subgroup, the result is $\text{PSL}(2, \mathbb{Z})$, which is




FIGURE 5. The Cayley graph of $\text{PSL}(2, \mathbb{Z})$ with respect to the generators S and T

isomorphic to $\mathbb{Z}_2 * \mathbb{Z}_3$. The factors are generated by

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad T = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}.$$

The Farey tree is closely related to the Cayley graph for $\text{PSL}(2, \mathbb{Z})$ with respect to these generators. This group acts on upper half space by isometries. It preserves $\mathbb{Q} \cup \{\infty\}$ in the circle at infinity, and hence it preserves the Farey tessellation and the dual Farey tree. The Cayley graph for $\text{PSL}(2, \mathbb{Z})$ with respect to these generators embeds in upper half space as follows. We set the vertex v corresponding to the identity element of $\text{PSL}(2, \mathbb{Z})$ to lie at $\epsilon + i$ for some small real $\epsilon > 0$. The images of this point v under the action of $\text{PSL}(2, \mathbb{Z})$ form the vertices of the Cayley graph. Emanating from the vertex v there are oriented edges, joining v to Sv and Tv . The images of these edges under the action of $\text{PSL}(2, \mathbb{Z})$ form the edges of the Cayley graph. The graph is shown in Figure 5. Note that v lies in the triangle with corners $0, 1$ and ∞ . The stabiliser of this triangle in $\text{PSL}(2, \mathbb{Z})$ is $\{T^i : i = 0, 1, 2\}$. Hence, there are three vertices of the Cayley graph in this triangle that are connected by edges labelled by T . The edge joining v to Sv intersects the geodesic joining 0 and ∞ and is disjoint from the remaining edges of the Farey graph. Hence, each S -labelled edge of the Cayley graph is dual to an edge of the Farey graph. Each edge of the Farey graph is associated with two such S -labelled edges, which join the same pair of the vertices.

Thus, in summary, the Cayley graph is obtained from the Farey tree as follows. Replace each vertex of the Farey tree by a little triangle, with edges labelled by T . Replace each edge of the Farey tree by two edges of the Cayley graph labelled by S .

5.2. Continued fractions. There is a well-known connection between continued fractions and the Farey tessellation.

A *continued fraction* for a rational number r is an expression

$$r = a_0 + \frac{1}{a_1 + \frac{1}{a_2 + \frac{1}{\dots + \frac{1}{a_n}}}}$$

where each a_i is an integer. This is written $r = [a_0, a_1, \dots, a_n]$.

One can also consider an irrational number r , which also has a *continued fraction expansion* $[a_0, a_1, \dots]$. This means that if $r_n = [a_0, a_1, \dots, a_n]$, then $r_n \rightarrow r$ as $n \rightarrow \infty$. We will focus on the case where a_i is positive for each $i > 0$. Subject to this condition, every real number r has a unique continued fraction expansion, which we shall call *the* continued fraction expansion for r .

The continued fraction expansion of r is *periodic* if there is a non-negative integer k and an even positive integer t such that $a_{i+t} = a_i$ for every $i \geq k$. The smallest such t is the *length* of the periodic part. The following is a well-known result of Lagrange; see, for example [4].

Lemma 5.2 (Lagrange). *The continued fraction expansion of a real number r is periodic exactly when $\mathbb{Q}(r)$ is a quadratic extension of \mathbb{Q} . This happens exactly when $r = p + q\sqrt{d}$ for some square free integer $d > 1$ and some rational numbers p and q where $q \neq 0$. Moreover, for fixed d , two real numbers $p + q\sqrt{d}$ and $p' + q'\sqrt{d}$, for $p, p' \in \mathbb{Q}$ and $q, q' \in \mathbb{Q} \setminus \{0\}$ have the same periodic part. That is, if $[a_0, a_1, a_2, \dots]$ and $[a'_0, a'_1, a'_2, \dots]$ are their continued fraction expansions, then there are integers l and m such that $a'_i = a_{i+m}$ for all $i \geq l$.*

One can read off the continued fraction expansion of a positive real number r from the Farey tessellation, as follows. Consider any hyperbolic geodesic γ starting in the hyperbolic plane on the imaginary axis, and ending at r on the circle at infinity. It intersects each triangle of the Farey tessellation in at most one arc; see Figure 6. As one travels along γ and one enters such a triangle, it either goes to the right or the left in this triangle, except when r is rational and γ lands on r . So, when r is irrational, one reads off the *cutting sequence* of γ , which is a sequence of lefts and rights, written as $L^{a_0}R^{a_1}L^{a_2}\dots$. Then the continued fraction expansion of r is $[a_0, a_1, a_2, \dots]$. When r is rational, we also get a cutting sequence, but we must be careful about the final triangle that γ runs through. Here, γ goes neither left nor right, but instead straight on towards r . We view this final triangle as giving a final L or R to the cutting sequence, where L or R is chosen to be the same as the previous letter. Thus, we obtain a sequence $L^{a_0}R^{a_1}L^{a_2}\dots L^{a_n}$ or $L^{a_0}R^{a_1}L^{a_2}\dots R^{a_n}$ with $a_n \geq 2$. Then $[a_0, a_1, a_2, \dots, a_n]$ is the continued fraction expansion of r .

The following lines in the Farey tree will play an important role in our analysis of lens spaces.

Definition 5.3. For $p/q \in \mathbb{Q} \cup \{\infty\}$, the line $L(p/q)$ in the Farey tree is the union of edges that are dual to an edge of the Farey graph emanating from p/q .

Lemma 5.4. *When $0 < q < p$, the distance in the Farey tree between the lines $L(q/p)$ and $L(\infty)$ is $(\sum_{i=0}^n a_i) - 1$ where $[a_0, \dots, a_n]$ is the continued fraction of expansion of p/q with each a_i positive.*




FIGURE 6. The cutting sequence determined by a geodesic starting on the imaginary axis and ending at r




FIGURE 7. The line $L(0)$ in the Farey tree

Proof. The line $L(\infty)$ runs parallel to the horocycle $\{(x, y) : y = 1\}$ in upper half-space. The line $L(q/p)$ forms a loop starting and ending at q/p . Let γ be the vertical geodesic in the half plane running from ∞ to the point q/p . As it comes from infinity, it hits $L(\infty)$, then it intersects various edges in the Farey graph, and then it hits $L(q/p)$. This determines a path in the Farey tree from $L(\infty)$ to $L(q/p)$. There is no shorter path, because each edge e of the Farey tessellation crossed by γ separates $L(\infty)$ from $L(q/p)$, and so any path in the Farey tree from $L(\infty)$ to $L(q/p)$ must run along the edge dual to e .

A closely related geodesic γ' determines the continued fraction expansion for q/p . This starts on the imaginary axis and ends at q/p . But because $0 < q/p < 1$, γ and γ' hit the same edges of the Farey graph (except the edge that forms the imaginary axis). Note that the continued fraction expansion of q/p is $[0, a_0, \dots, a_n]$. Hence, $(\sum_{i=0}^n a_i) - 1$ is exactly the length of the path in the Farey tree joining $L(\infty)$ to $L(q/p)$. \square




FIGURE 8. Splits and slide applied to a train track

6. TRAIN TRACK SPLITTING SEQUENCES

We will estimate distances in triangulation graphs using train tracks. We start by recalling some terminology.

A *pre-track* is a graph τ smoothly embedded in the interior of a surface S such that at each vertex v , the following hold:

- (1) there are three edges coming into v ;
- (2) these edges all have non-zero derivative at v , all of which lie in the same tangent line;
- (3) one edge approaches v along this line from one direction, and the other two edges approach from the other direction.

The vertices of τ are called *switches* and the edges are called *branches*.

Each component R of $S \setminus \tau$ is a surface, but its boundary is not necessarily smooth. Its boundary is composed of a union of arcs, one for each edge of τ . When two of these arcs cannot be combined into a single smooth arc, their point of intersection is a *cuspl*. The *index* of R is equal to $\chi(R)$ minus half the number of cusps of ∂R . We say that τ is a *train track* if each component of $S \setminus \tau$ has negative index. If we add up the index of the components of $S \setminus \tau$, the result is $\chi(S)$. Hence, we deduce that the number of complementary regions of a train track is at most $-2\chi(S)$.

A train track is *filling* if each component of $S \setminus \tau$ is either a disc or an annular neighbourhood of a component of ∂S . When τ is filling, it is dual to a triangulation of S , possibly with some ideal vertices.

Two train tracks differ by a *split* or a *slide* if one is obtained from the other by one of the modifications shown in Figure 8.

Pick a point in the interior of each branch. Cutting the branch at such a point creates two intervals, which are called *half-branches*. A half-branch is called *small* if at the cusp at its endpoint, there is another half-branch coming in from the same direction. If a half-branch is not small, it is *large*.

For any train track τ , its regular neighbourhood $N(\tau)$ is naturally a union of intervals called *fibres*, and there is a collapsing map $N(\tau) \rightarrow \tau$ that collapses each

fibre to a point. A curve C is said to be *carried* by τ if C is embedded in $N(\tau)$ and is transverse to all the fibres.

A train track τ is *transversely recurrent* if for each branch of τ , there is a simple closed curve C transverse to τ that intersects this branch and such that $\tau \cup C$ does not have a complementary region that is a bigon. By a *bigon*, we mean a component of $S \setminus (C \cup \tau)$ that is a disc with boundary consisting of the union of two arcs, one lying in C , the other lying in τ and having no cusps. The train track τ is *recurrent* if for each branch of τ , there is a simple closed curve carried by τ that runs over the branch. It is *birecurrent* if it is both recurrent and transversely recurrent. We say that *the set of curves carried by τ fills S* if, for every essential simple closed curve C in S , there is a curve carried by τ that cannot be isotoped off C . The train track is then said to be *filling*.

The following is essentially due to Masur, Mosher and Schleimer [18].

Theorem 6.1. *Let τ and τ' be filling birecurrent train tracks in a closed orientable surface S of genus at least 2. Suppose that there is a sequence of splits and slides taking τ to τ' . Let \mathcal{T} and \mathcal{T}' be the triangulations dual to τ and τ' . Let $n = -2\chi(S)$. Then, the distance in $\text{Tr}(S; n)$ between \mathcal{T} and \mathcal{T}' is, up to a bounded multiplicative error, equal to the number of splits. This bound only depends on the Euler characteristic of S .*

Proof. In [18, Section 6.1], the *marking graph* $\mathcal{M}(S)$ is defined. This is quasi-isometric to the mapping class group of S . In [18, Section 6.1], a map from filling birecurrent train tracks to $\mathcal{M}(S)$ is defined. Composing this with the quasi-isometries $\mathcal{M}(S) \rightarrow \text{MCG}(S) \rightarrow \text{Tr}(S; n)$ we obtain a map from filling birecurrent train tracks to $\text{Tr}(S; n)$. This is a bounded distance from the map that sends each filling birecurrent train track to its dual triangulation.

We are supposing that there is a sequence of splits and slides taking τ to τ' . Then by [18, Theorem 6.1], a sequence of such splits and slides is sent to a quasi-geodesic in $\mathcal{M}(S)$, with quasi-geodesic constants depending only on S . The length of this quasi-geodesic is the number of splits in the sequence. We compose this with the quasi-isometries $\mathcal{M}(S) \rightarrow \text{MCG}(S) \rightarrow \text{Tr}(S; n)$, and we obtain a quasi-geodesic in $\text{Tr}(S; n)$. Its start and end vertices are a bounded distance from \mathcal{T} and \mathcal{T}' , the bound depending only on S . Hence, the distance in $\text{Tr}(S; n)$ between \mathcal{T} and \mathcal{T}' is, up to a bounded multiplicative error, equal to the number of splits. \square

Now suppose that the train track τ has a transverse measure μ . (We refer to [6] for the definition of transverse measures and their relationship with pseudo-Anosov homeomorphisms.) At any large branch of τ , one may split τ in three possible ways, but only one of these ways is compatible with μ . The result is a measured train track (τ', μ') .

A *maximal split* on a measured train track (τ, μ) is obtained by performing a measured split at each large branch of τ . The following was proved by Agol [1, Theorem 3.5].

Theorem 6.2. *Let S be a compact orientable surface. Let $\phi: S \rightarrow S$ be a pseudo-Anosov homeomorphism, and let (τ, μ) be a measured train track that carries its stable measured lamination. Let λ be its dilatation. For each positive integer i , let (τ_i, μ_i) be the result of performing a sequence of i maximal splits to (τ, μ) . Then there are integers $n > 0$ and $m \geq 0$ such that, for each $i \geq m$, $\tau_{n+i} = \phi(\tau_i)$ and $\mu_{n+i} = \lambda^{-1}\mu_i$.*

7. HOMEOMORPHISMS OF THE TORUS

The famous classification of orientation-preserving homeomorphisms of closed orientable surfaces into periodic, reducible and pseudo-Anosov [27] is a generalisation of the special case of the torus. In this case, the third category is known as *linear Anosov*, which we can define to be isotopic to a linear map with determinant 1, and with irrational real eigenvalues.

An orientation-preserving homeomorphism of the torus induces an action on the homology of the torus and hence gives an element of $\mathrm{SL}(2, \mathbb{Z})$. There is a homomorphism $\mathrm{SL}(2, \mathbb{Z})$ to the isometry group of the hyperbolic plane. Thus, our homeomorphism of the torus induces an isometry of the hyperbolic plane that preserves the Farey tessellation, and hence is an isometry of the Farey tree. An alternative way of viewing this action is to note that the Farey tree is $\mathrm{Tr}(T^2)$ and any homeomorphism of the torus naturally induces an isometry of $\mathrm{Tr}(T^2)$.

Recall that any isometry of a tree either has a fixed point or has an *invariant axis*. This is a subset of the tree isometric to the real line, such that the isometry acts as non-trivial translation upon this line. Any subset of the Farey tree isometric to the real line has two well-defined endpoints on the circle at infinity, although these endpoints need not be distinct. We can now see the classification of orientation-preserving homeomorphisms of the torus in terms of the action on the Farey tree:

- (1) A homeomorphism is periodic if and only if its action on the Farey tree has a fixed point in the interior.
- (2) A homeomorphism is reducible and not periodic if and only if its action on the Farey tree has an invariant axis, but the endpoints of this axis are the same point on the circle at infinity.
- (3) A homeomorphism is linear Anosov if and only if its action on the Farey tree has an invariant axis, and the endpoints of the axis are distinct points on the circle at infinity.

The following is well-known (see for example [3, Section 0]).

Lemma 7.1. *A matrix $A \in \mathrm{SL}(2, \mathbb{Z})$ acts on the torus as a linear Anosov homeomorphism if and only if its trace $\mathrm{tr}(A)$ satisfies $|\mathrm{tr}(A)| > 2$.*

Proof. The matrix A projects to an element of $\mathrm{PSL}(2, \mathbb{Z})$, which is a subgroup of orientation-preserving isometries of the hyperbolic plane. This isometry induces the action of A on the Farey tessellation and hence on the Farey tree.

Consider first the case that the action has a fixed point in the interior. If A has rows (a, b) and (c, d) , a fixed point is an element x with positive imaginary part such that $(ax + b)/(cx + d) = x$. Solving for x , this gives a quadratic polynomial with discriminant $\sqrt{\mathrm{tr}(A)^2 - 4}$ and with highest order term cx^2 . Thus there is a fixed point with positive imaginary part if and only if $|\mathrm{tr}(A)| < 2$ and $c \neq 0$. Note however that if $c = 0$, then the condition that A has determinant 1 forces $\mathrm{tr}(A)$ to be equal to ± 2 . Thus there is a fixed point with positive imaginary part if and only if $|\mathrm{tr}(A)| < 2$.

Consider next the eigenvectors of A . Specifically, $(u, 1)^T$ is an eigenvector of A if and only if $u/1$ is an endpoint of an invariant axis of the action of A on the Farey tree. Now, the characteristic polynomial for A is $x^2 - \mathrm{tr}(A)x + 1$, with roots $(\mathrm{tr}(A) \pm \sqrt{\mathrm{tr}(A)^2 - 4})/2$. Thus A has distinct real eigenvalues if and only

if $|\operatorname{tr}(A)| > 2$; this is the case A induces a linear Anosov. The remaining case, $|\operatorname{tr}(A)| = 2$, corresponds to the case A is reducible and not periodic. \square

When an isometry ϕ of a tree has an invariant axis, then this is also the invariant axis for any non-zero power of ϕ . Hence, we have the following result.

Lemma 7.2. *Let A be a homeomorphism of the torus. Then for any integer n , $\ell_{\operatorname{Tr}(T^2)}(A^n) = n \ell_{\operatorname{Tr}(T^2)}(A)$. Hence, the stable translation length satisfies $\ell_{\operatorname{Tr}(T^2)}(A) = \ell_{\operatorname{Tr}(T^2)}(A)$.*

7.1. Translation length of a linear Anosov. The following well known proposition gives the length of the translation in the Farey tree of an Anosov in terms of continued fractions.

Proposition 7.3. *Let $A \in \operatorname{SL}(2, \mathbb{Z})$ act as a linear Anosov homeomorphism on the torus. Let \bar{A} be the image of A in $\operatorname{PSL}(2, \mathbb{Z})$. Suppose that \bar{A} is B^n for some positive integer n and some matrix $B \in \operatorname{PSL}(2, \mathbb{Z})$ that is not a proper power. Let (a_r, \dots, a_s) denote the periodic part of the continued fraction expansion of $\sqrt{\operatorname{tr}(A)^2 - 4}$. Then the translation distance of A in the Farey tree is $n \sum_{i=r}^s a_i$.*

Proof. As in the proof of Lemma 7.1, the matrix A corresponds to a linear Anosov homeomorphism when $|\operatorname{tr}(A)| > 2$, with eigenvalues

$$\frac{\operatorname{tr}(A) \pm \sqrt{\operatorname{tr}(A)^2 - 4}}{2}.$$

Let λ be either of these eigenvalues. Then the determinant of the matrix $A - \lambda I$ is zero, and hence the two rows are multiples of each other. Let (a, b) be one of its rows. Suppose $(u, 1)^T$ is an eigenvector for A , so $u/1$ is an endpoint of the invariant axis γ for A . Then $au + b = 0$ and so $u = -b/a$. Thus, we deduce that u lies in $\mathbb{Q}(\lambda) = \mathbb{Q}(\sqrt{\operatorname{tr}(A)^2 - 4})$. So, the periodic part of the continued fraction of u is equal to the periodic part of the continued fraction expansion of $\sqrt{\operatorname{tr}(A)^2 - 4}$ by Lemma 5.2.

Let γ' be a geodesic starting at a point in the hyperbolic plane on the imaginary axis and ending at u on the circle at infinity. The edges of the Farey graph that it crosses determines the cutting sequence $L^{a_0} R^{a_1} L^{a_2} \dots$ for γ' and hence the continued fraction expansion $[a_0, a_1, \dots]$ for u . This cutting sequence is eventually the same as that of the invariant axis γ . In particular, they have the same periodic parts. Now, as γ is the axis of A , its cutting sequence is periodic. However, the length of the corresponding path in the Farey tree may be a multiple of this period. This happens precisely when \bar{A} is B^n for some integer n and some matrix $B \in \operatorname{PSL}(2, \mathbb{Z})$ that is not a proper power. Hence, translation distance of A in the Farey tree is $n \sum_{i=r}^s a_i$. \square

7.2. Moving between vertices in the Farey tree.

Lemma 7.4. *Given any two ideal triangles $a_1 b_1 c_1$ and $a_2 b_2 c_2$ of the Farey tessellation, there is a homeomorphism ϕ of the torus such that $\phi(a_1) = a_2$, $\phi(b_1) = b_2$, $\phi(c_1) = c_2$. This is unique up to isotopy and composition by the map $-\operatorname{id}$. It is orientation-preserving if and only if ϕ preserves the cyclic ordering of the vertices around the circle at infinity.*

Proof. Since the slopes a_1 and b_1 have intersection number 1, we may choose a basis for the first homology of the torus so that $a_1 = (1, 0)$ and $b_1 = (0, 1)$, when

these slopes are oriented in some way. The function $a_1 \mapsto a_2$ and $b_1 \mapsto b_2$ may be realised by an element of $\mathrm{GL}(2, \mathbb{Z})$. Note that the determinant of this map is indeed ± 1 , since a_2 and b_2 are Farey neighbours. This linear map sends c_1 to a slope that has intersection number one with both a_2 and b_2 . If this slope is not c_2 , then pre-compose the linear map by $(1, 0) \mapsto (-1, 0)$ and $(0, 1) \mapsto (0, 1)$. This gives the required homeomorphism ϕ .

To establish uniqueness, it suffices to check that if $\phi(a_1) = a_1$, $\phi(b_1) = b_1$ and $\phi(c_1) = c_1$ then ϕ is isotopic to $\pm \mathrm{id}$. But if the linear map ϕ sends $(1, 0)$ to $\pm(1, 0)$, sends $(0, 1)$ to $\pm(0, 1)$ and sends $(1, 1)$ to $\pm(1, 1)$, then ϕ is $\pm \mathrm{id}$.

We now show that ϕ is orientation-preserving if and only if ϕ preserves the cyclic ordering of the vertices around the circle at infinity. We established above that ϕ is either an element of $\mathrm{SL}(2, \mathbb{Z})$ or a composition of an element of $\mathrm{SL}(2, \mathbb{Z})$ with a reflection. In the former case, ϕ is orientation-preserving and realised by a Möbius transformation of upper half-space, which therefore preserves the cyclic ordering of triples in the circle at infinity. In the latter case, ϕ is orientation-reversing and reverses the cyclic ordering of the vertices. \square

Lemma 7.5. *Suppose $a_1b_1c_1$ and $a_2b_2c_2$ are the vertices of distinct ideal triangles of the Farey tessellation. Let α be the unique embedded path in the Farey tree joining the centre of $a_1b_1c_1$ to the centre of $a_2b_2c_2$. Let e be the first edge of the Farey graph that α crosses. Let ϕ be as in Lemma 7.4. Then ϕ acts on the Farey tree, and so sends α to an arc $\phi(\alpha)$. Suppose that $\phi(\alpha)$ leaves the triangle $a_2b_2c_2$ by a different edge from the one α came in through. Suppose also that $\phi(e)$ and e do not share a vertex. Then ϕ is linear Anosov. Moreover, the axis of ϕ in the Farey tree runs through the vertices dual to $a_1b_1c_1$ and $a_2b_2c_2$.*

Proof. The infinite line $\bigcup_{n=-\infty}^{\infty} \phi^n(\alpha)$ forms an invariant axis. The endpoints of this axis on the circle at infinity are distinct, because they are separated by the endpoints of e and $\phi(e)$. Hence, as discussed above, ϕ is linear Anosov. \square

Proposition 7.6. *Let \mathcal{T}_1 and \mathcal{T}_2 be distinct 1-vertex triangulations of the torus that do not differ by a 2-2 Pachner move. Then there is a linear Anosov homeomorphism ϕ such that $\phi(\mathcal{T}_1) = \mathcal{T}_2$. Moreover, the axis of ϕ in the Farey tree runs through the vertices dual to \mathcal{T}_1 and \mathcal{T}_2 .*

Proof. Let a_1, b_1, c_1 and a_2, b_2, c_2 in $\mathbb{Q} \cup \{1/0\}$ correspond to the slopes of \mathcal{T}_1 and \mathcal{T}_2 in the Farey tessellation, where the edge b_1c_1 and b_2c_2 are closest to each other. Choose the labelling so that both a_1, b_1, c_1 and a_2, b_2, c_2 appear in a clockwise fashion around the circle at infinity. Since \mathcal{T}_1 and \mathcal{T}_2 do not differ by a 2-2 Pachner move, $b_2 \neq c_1$ or $c_2 \neq b_1$, say $b_2 \neq c_1$. By Lemma 7.4, there is a homeomorphism ϕ such that $\phi(a_1) = c_2$, $\phi(b_1) = a_2$ and $\phi(c_1) = b_2$. It is orientation-preserving, since it preserves the cyclic ordering of the vertices.

Let α be the arc in the Farey tree joining the centre of $a_1b_1c_1$ to the centre of $a_2b_2c_2$. The first edge e of the Farey graph that it crosses is b_1c_1 . Then $\phi(e)$ is a_2b_2 . This is different from the edge b_2c_2 that α crosses. Note also that e and $\phi(e)$ do not share a vertex. So by Lemma 7.5, it is linear Anosov, with axis as claimed. \square

7.3. Splitting sequences between two ideal triangulations.

Theorem 7.7. *Let \mathcal{T} and \mathcal{T}' be distinct ideal triangulations of the once-punctured torus. Then there are (filling) train tracks τ and τ' dual to \mathcal{T} and \mathcal{T}' and a sequence of splits taking τ to τ' of length $d_{\mathrm{Tr}(\mathcal{T}^2)}(\mathcal{T}, \mathcal{T}')$. These train tracks are birecurrent,*




FIGURE 9. Left: A simple closed curve on a triangulated torus. Right: The associated measured train track

and the set of curves carried by τ fill the once-punctured torus, as do the set of curves carried by τ' .

Proof. First observe that if \mathcal{T} and \mathcal{T}' differ by a 2-2 Pachner move, then it is straightforward to realise their dual trees as train tracks τ and τ' in the once-punctured torus that differ by a single split.

So we assume that \mathcal{T} and \mathcal{T}' do not differ by a 2-2 Pachner move. The triangulations \mathcal{T} and \mathcal{T}' correspond to vertices v and v' of the Farey tree. By Proposition 7.6, there is a linear Anosov homeomorphism ϕ of the torus taking v to v' . Moreover, the axis of ϕ goes through v and v' . This has a stable lamination \mathcal{L} with a transverse measure. We may isotope \mathcal{L} so that it intersects each triangle of \mathcal{T} in normal arcs. The edges of \mathcal{T} then inherit a transverse measure. There are three possible normal arc types in each triangle, and some triangle must be missing an arc type, as otherwise \mathcal{L} would contain a simple closed curve encircling the puncture. Thus, in that triangle, the three edges have measures a , b and $a + b$ for some non-negative real numbers a and b . As this is the torus, these are the three edges of the other triangle of \mathcal{T} , and hence this triangle is also missing an arc type. Now in fact, a and b must both be positive, as otherwise \mathcal{L} would be a thickened simple closed curve. So, the weights on the edges determine a train track $\tau = \tau_0$ that is dual to \mathcal{T} and that carries \mathcal{L} . See Figure 9. Let μ_0 be the transverse measure on τ_0 .

We can view ϕ as specifying a pseudo-Anosov homeomorphism of the once punctured torus, where its stable lamination is again \mathcal{L} . We now apply Agol's result, Theorem 6.2, which provides a splitting sequence, giving a sequence of transversely measured train tracks (τ_i, μ_i) starting at (τ_0, μ_0) . A split does not increase the number of complementary regions of a train track. Hence, each train track τ_i has a single complementary region that is an annular neighbourhood of the puncture. It is therefore dual to an ideal triangulation \mathcal{T}_i . Thus, this sequence of train tracks produces an injective path in the Farey tree starting at v . Since it is eventually periodic, at some point, this path must land on the axis of ϕ and follow this axis from then onwards. However, v is already on the axis of ϕ . Thus, this path just follows the axis. The axis goes through v' , and so when the path reaches v' , the result is a train track dual τ' to \mathcal{T}' . Thus, the required splitting sequence has been produced.

We now show that τ and τ' are birecurrent. Let τ'' be the train track that is dual to the ideal triangulation with edges having slopes $1/0$, $0/1$ and $1/1$, and where the latter is dual to the large branch. There is a homeomorphism of the once-punctured

torus taking the ideal triangulation dual to τ to the one with edges $1/0$, $0/1$ and $1/1$, and taking the edge dual to the large branch of τ to $1/1$. Thus, there is a homeomorphism taking τ to τ'' . Similarly, there is a homeomorphism taking τ' to τ'' . So it suffices to show that τ'' is birecurrent. But the simple closed curves with slopes $1/0$, $0/1$ and $1/1$ can be arranged to intersect the branches of τ'' in the required way, thereby establishing that τ'' is transversely recurrent. Also, τ'' is recurrent, since for each of its three branches, there is an obvious simple closed curve carried by τ'' that runs over this branch. Hence, τ'' is birecurrent, as required.

Finally, the curves carried by τ'' fill the once-punctured torus, since they include $1/0$ and $0/1$. Hence, the curves carried by τ also fill the once-punctured torus, as do the curves carried by τ' . \square

8. BRANCHED COVERS OF THE TORUS

We will consider branched covering maps $p: S \rightarrow T$, where T is the torus and S is a closed orientable surface. We will require that there is a single branch point b in T . Our first result says that any 1-vertex triangulation \mathcal{T} of T has a well-defined lift to S .

Lemma 8.1. *Let $p: S \rightarrow T$ be a branched cover of the torus T , branched over a single point b in T . Let \mathcal{T} and \mathcal{T}' be isotopic 1-vertex triangulations of the torus T , with their vertices both equal to b . Then their inverse images $\tilde{\mathcal{T}}$ and $\tilde{\mathcal{T}}'$ in S are isotopic.*

Proof. The triangulations \mathcal{T} and \mathcal{T}' are isotopic, but the isotopy is not assumed to preserve basepoints. The isotopy is a 1-parameter family of homeomorphisms, the final one being a homeomorphism $h: (T, b) \rightarrow (T, b)$. The Birman exact sequence [5, Section 4.2.3] for the torus gives that the natural map $\text{MCG}(T, 1 \text{ point}) \rightarrow \text{MCG}(T)$ is an isomorphism. Here, $\text{MCG}(T, 1 \text{ point})$ denotes the group of orientation-preserving homeomorphisms of T that fix a specific point, up to isotopies that fix this point throughout. Hence, the homeomorphism h is isotopic to the identity, via an isotopy that keeps b fixed throughout. This isotopy lifts to an isotopy of S that keeps $p^{-1}(b)$ fixed throughout. This isotopy takes $\tilde{\mathcal{T}}$ to $\tilde{\mathcal{T}}'$. \square

As a consequence of the above lemma, it makes sense to compare distances in $\text{Tr}(T)$ with distances in suitable triangulation graphs for S .

Theorem 8.2. *Let $p: S \rightarrow T$ be a branched cover of the torus, branched over a single point b in T , with finite degree $\deg(p)$. Suppose that the branching index around each point in $p^{-1}(b)$ is greater than 1. Let \mathcal{T}_1 and \mathcal{T}_2 be 1-vertex triangulations of the torus T with vertex at b , and let $\tilde{\mathcal{T}}_1$ and $\tilde{\mathcal{T}}_2$ be their inverse images in S . Then there are constants $k_1, k_2 > 0$, depending only on S , such that*

$$k_1 d_{\text{Tr}(T)}(\mathcal{T}_1, \mathcal{T}_2) - k_2 \leq d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2) \leq \deg(p) d_{\text{Tr}(T)}(\mathcal{T}_1, \mathcal{T}_2).$$

Proof. The upper bound follows immediately from the fact that a 2-2 Pachner move on a triangulation of T with vertex at b induces $\deg(p)$ 2-2 Pachner moves on the corresponding triangulation of S .

So, we focus on the other inequality. We may assume that \mathcal{T}_1 and \mathcal{T}_2 do not differ by a 2-2 Pachner move, for otherwise $d_{\text{Tr}(T)}(\mathcal{T}_1, \mathcal{T}_2) = 1$, and the inequality is trivial.

We view \mathcal{T}_1 and \mathcal{T}_2 as ideal triangulations of the once-punctured torus. By Theorem 7.7, they are dual to filling birecurrent train tracks τ_1 and τ_2 , and there is a splitting sequence taking τ_1 to τ_2 such that the length of this sequence is $d_{\text{Tr}(T)}(\mathcal{T}_1, \mathcal{T}_2)$. View these as pre-tracks in the torus disjoint from the branch point b . Their inverse images $\tilde{\tau}_1$ and $\tilde{\tau}_2$ in S are pre-tracks. In fact, they are train tracks, because each complementary region is a regular neighbourhood of a point of $p^{-1}(b)$. Since the branching index around this point is greater than 1, the inverse image of the two cusps around b is at least four cusps. Then the index of the complementary region is at most $\chi(\text{disc}) - \frac{1}{2}(4) \leq -1$. Hence the number of complementary regions of $\tilde{\tau}_1$ and $\tilde{\tau}_2$ is at most $-\chi(S)$. So their duals $\tilde{\mathcal{T}}_1$ and $\tilde{\mathcal{T}}_2$ each have at most $-\chi(S)$ vertices.

The splitting sequence from τ_1 to τ_2 lifts to a splitting sequence from $\tilde{\tau}_1$ to $\tilde{\tau}_2$, of length $\deg(p)d_{\text{Tr}(T^2)}(\mathcal{T}_1, \mathcal{T}_2)$. We wish to use Theorem 6.1 to show that the number of splits gives a lower bound on $d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2)$, up to multiplicative error depending only on the Euler characteristic of S . We need to check the hypotheses of Theorem 6.1.

We first show that $\tilde{\tau}_1$ is transversely recurrent. Consider any branch \tilde{e} of $\tilde{\tau}_1$. It projects to a branch e of τ_1 . Since τ_1 is transversely recurrent, there is a simple closed curve C through e that intersects τ_1 transversely and where $\tau \cup C$ has no bigon complementary region. Let \tilde{C} be the inverse image of C in S . The component of \tilde{C} going through \tilde{e} establishes the transverse recurrence of $\tilde{\tau}_1$. The same argument establishes that $\tilde{\tau}_2$ is transversely recurrent.

We now show that $\tilde{\tau}_1$ is recurrent. For each branch \tilde{e} of $\tilde{\tau}_1$, let e be its image branch in τ_1 . There is a curve C carried by τ_1 running over e . Its inverse image in S is a collection \tilde{C} of curves carried by $\tilde{\tau}_1$. One of these runs over \tilde{e} as required.

We now show that the curves carried by $\tilde{\tau}_1$ fill S . Let C be a finite collection of curves carried by τ_1 that fill the punctured torus. Place the curves of C in minimal position with respect to each other, in the sense that no two of them have a bigon complementary region. Then the complement of these curves in T is a union of discs. Let \tilde{C} be the inverse image of these curves in S . These are carried by $\tilde{\tau}_1$. They are in minimal position. Their complement in S is a union of discs. Hence, \tilde{C} fills S .

Thus, the hypotheses of Theorem 6.1 hold and so $d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2)$ is at least the number of splits in a sequence taking $\tilde{\tau}_1$ to $\tilde{\tau}_2$, up to bounded multiplicative error with bound depending only on S . We know from above that this number of splits is $\deg(p)d_{\text{Tr}(T)}(\mathcal{T}_1, \mathcal{T}_2)$. It follows that $d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2)$ is at least $k_1 d_{\text{Tr}(T)}(\mathcal{T}_1, \mathcal{T}_2) - k_2$, for some constants $k_1, k_2 > 0$ depending only on S . \square

Suppose $p: S \rightarrow T$ is a branched cover of the torus, branched over a single point b with degree $\deg(p)$. Suppose Γ is a spine for T disjoint from b . Then observe that the inverse image of Γ in S might not be a spine, as its complement consists of at most $\deg(p)$ discs. However, a spine can be formed by removing at most $\deg(p) - 1$ edges.

Corollary 8.3. *Let $p: S \rightarrow T$ be a branched cover of the torus, branched over a single point b , with finite degree $\deg(p)$. Suppose that the branching index around each point in $p^{-1}(b)$ is greater than 1. Let Γ_1 and Γ_2 be spines for T that are disjoint from b , and let $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$ be their inverse images in S . Remove at most $\deg(p) - 1$ edges from each of $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$ to form spines $\tilde{\Gamma}'_1$ and $\tilde{\Gamma}'_2$ for S . Then*

there are constants $c_1, C_1, c_2, C_2 > 0$, depending only on p , such that

$$c_1 d_{\mathrm{Sp}(T)}(\Gamma_1, \Gamma_2) - c_2 \leq d_{\mathrm{Sp}(S)}(\tilde{\Gamma}'_1, \tilde{\Gamma}'_2) \leq C_1 d_{\mathrm{Sp}(T)}(\Gamma_1, \Gamma_2) + C_2.$$

Proof. Note first that it does not matter which edges of $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$ that we remove. For suppose that $\tilde{\Gamma}'_1$ and $\tilde{\Gamma}''_1$ are other spines also obtained from $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$ by removing at most $\deg(p) - 1$ edges from each. Then, $\tilde{\Gamma}'_1$ and $\tilde{\Gamma}''_1$ differ by at most $24(\deg(p) - 1)$ edge contractions and expansions by Lemma 4.7, and similarly so do $\tilde{\Gamma}'_2$ and $\tilde{\Gamma}''_2$. Thus the difference can be picked up by the constants.

Define a map $\mathrm{Tr}(T) \rightarrow \mathrm{Tr}(S; -\chi(S))$ first on the vertices: This sends a vertex in $\mathrm{Tr}(T)$, corresponding to a 1-vertex triangulation of T with vertex at b , to the vertex in $\mathrm{Tr}(S; -\chi(S))$ corresponding to the triangulation that is the inverse image of S . Each edge in T corresponds to a 2-2 Pachner move and this lifts to $\deg(p)$ 2-2 Pachner moves in S . Hence, the map $\mathrm{Tr}(T) \rightarrow \mathrm{Tr}(S; -\chi(S))$ can also be defined on edges and is continuous. Theorem 8.2 implies that this map is a quasi-isometry. By Lemma 4.5, the maps $\mathrm{Tr}(T) \rightarrow \mathrm{Sp}(T)$ and $\mathrm{Tr}(S) \rightarrow \mathrm{Sp}(S)$ are quasi-isometries, and the inclusion $\mathrm{Tr}(S) \rightarrow \mathrm{Tr}(S; n)$ is a quasi-isometry for any positive integer n . A quasi-inverse is given by taking any triangulation with at most n vertices, dualising it to form a trivalent graph, then removing edges to form a spine, and then dualising to form a 1-vertex triangulation. This construction can be chosen to be invariant under the mapping class group, and hence forms a quasi-isometry $\mathrm{Tr}(S; n) \rightarrow \mathrm{Tr}(S)$. The composition

$$\mathrm{Sp}(T) \rightarrow \mathrm{Tr}(T) \rightarrow \mathrm{Tr}(S; -\chi(S)) \rightarrow \mathrm{Tr}(S) \rightarrow \mathrm{Sp}(S)$$

is a quasi-isometry. Thus, we obtain the required inequalities. \square

Remark 8.4. Theorem 8.2 and Corollary 8.3 require a branched cover of the torus with a single branch point b and where each point of $p^{-1}(b)$ has branching index at least two. One such branched cover is obtained by the following process. First, take the $\mathbb{Z}/2 \times \mathbb{Z}/2$ cover of T^2 arising from the natural homomorphism $\pi_1(T^2) \rightarrow H_1(T^2; \mathbb{Z}/2)$. Then restrict this to a cover $F \rightarrow T^2 - \{b\}$, where F is a four-times punctured torus. Next form the cover of F arising from $\pi_1(F) \rightarrow H_1(F; \mathbb{Z}/2)$. Finally complete this to form the required branched cover S of T^2 .

9. TRIANGULATIONS AND HANDLE STRUCTURES OF PRODUCTS

Theorem 1.6. *Let \mathcal{T}_0 and \mathcal{T}_1 be 1-vertex triangulations of the torus T^2 . Let $\Delta(\mathcal{T}_0, \mathcal{T}_1)$ denote the minimal number of tetrahedra in any triangulation of $T^2 \times [0, 1]$ that equals \mathcal{T}_0 on $T^2 \times \{0\}$ and equals \mathcal{T}_1 on $T^2 \times \{1\}$. Then there is a universal constant $k_{\mathrm{prod}} > 0$ such that*

$$k_{\mathrm{prod}} d_{\mathrm{Tr}(T^2)}(\mathcal{T}_0, \mathcal{T}_1) \leq \Delta(\mathcal{T}_0, \mathcal{T}_1) \leq d_{\mathrm{Tr}(T^2)}(\mathcal{T}_0, \mathcal{T}_1) + 6.$$

Proof. The upper bound is straightforward: Let α be a shortest path in $\mathrm{Tr}(T^2)$ from \mathcal{T}_0 to \mathcal{T}_1 . This determines a sequence of 1-vertex triangulations, starting at \mathcal{T}_0 and ending at \mathcal{T}_1 . We use this to build a triangulation of $T^2 \times [0, 1]$, as follows. Start with $T^2 \times \{0\}$ triangulated using \mathcal{T}_0 . Then take the product of this with $[0, 1]$ and triangulate each of the resulting prisms using 3 tetrahedra. If chosen correctly, these patch together to form a triangulation of $T^2 \times [0, 1]$ with 6 tetrahedra, where $T^2 \times \{0\}$ and $T^2 \times \{1\}$ are both triangulated using \mathcal{T}_0 . Then layer onto $T^2 \times \{1\}$ a sequence of tetrahedra, specified by the sequence of Pachner moves, until we reach \mathcal{T}_1 . This gives the upper bound.

As for the lower bound, let $p: S \rightarrow T^2$ be a branched cover, with single branch point b and where each point of $p^{-1}(b)$ has branching index at least two. For example, take the branched cover of Remark 8.4.

Now let \mathcal{T} be a triangulation of $T^2 \times [0, 1]$ that equals \mathcal{T}_0 on $T^2 \times \{0\}$ and equals \mathcal{T}_1 on $T^2 \times \{1\}$ and that realises $\Delta(\mathcal{T}_0, \mathcal{T}_1)$. By Theorem 3.2, the 23rd iterated barycentric subdivision $\mathcal{T}^{(23)}$ contains an arc in its 1-skeleton that is vertical. Note $\mathcal{T}^{(23)}$ consists of $(24)^{23} \Delta(\mathcal{T}_0, \mathcal{T}_1)$ tetrahedra. By Lemma 4.14, there is a sequence of at most $4(1 + 6 + \dots + 6^{22}) < 6^{23}$ Pachner moves taking \mathcal{T}_0 to $\mathcal{T}_0^{(23)}$. Perform the reverse of this sequence, and realise each Pachner move on the boundary of $S \times [0, 1]$ by attaching a tetrahedron to $S \times \{0\}$. Do the same for $S \times \{1\}$. Let \mathcal{T}_+ be the resulting triangulation of $S \times [0, 1]$. It has at most $(24)^{23} \Delta(\mathcal{T}_0, \mathcal{T}_1) + 2 \cdot 6^{23}$ tetrahedra and it equals \mathcal{T}_0 and \mathcal{T}_1 on its boundary. By Lemma 3.5, it also contains an arc in its 1-skeleton that is vertical. Hence, the inverse image of \mathcal{T}_+ under the branched covering map is a triangulation of $S \times [0, 1]$. It has at most $(24)^{23} \deg(p) \Delta(\mathcal{T}_0, \mathcal{T}_1) + 2 \cdot 6^{23} \deg(p)$ tetrahedra.

Let $\tilde{\mathcal{T}}_0$ denote the restriction of the triangulation $p^{-1}(\mathcal{T}_+)$ to $S \times \{0\}$, and let $\tilde{\mathcal{T}}_1$ denote the triangulation on $S \times \{1\}$. By Theorem 8.2, there are constants $k_1, k_2 > 0$, depending only on p , such that

$$(1) \quad k_1 d_{\text{Tr}(T)}(\mathcal{T}_0, \mathcal{T}_1) - k_2 \leq d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}_0, \tilde{\mathcal{T}}_1).$$

By Lemma 4.13, there is a sequence of at most $4 \deg(p)$ Pachner moves taking $\tilde{\mathcal{T}}_0$ to a 1-vertex triangulation $\tilde{\mathcal{T}}'_0$, and a sequence of at most $4 \deg(p)$ Pachner moves taking $\tilde{\mathcal{T}}_1$ to a 1-vertex triangulation $\tilde{\mathcal{T}}'_1$. Considering the reverse moves, we obtain

$$(2) \quad d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}_0, \tilde{\mathcal{T}}_1) \leq d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}'_0, \tilde{\mathcal{T}}'_1) + 8 \deg(p).$$

Each Pachner move corresponds to the addition of a 3-simplex to the triangulation of $S \times [0, 1]$. So, we obtain a triangulation \mathcal{T}' of $S \times [0, 1]$ that equals $\tilde{\mathcal{T}}'_0$ on $S \times \{0\}$, equals $\tilde{\mathcal{T}}'_1$ on $S \times \{1\}$ and has at most $(24)^{23} \deg(p) \Delta(\mathcal{T}_0, \mathcal{T}_1) + (2 \cdot 6^{23} + 8) \deg(p)$ tetrahedra.

Let $\Delta(\mathcal{T}')$ denote the minimal number of tetrahedra in any triangulation of $S \times [0, 1]$ that equals \mathcal{T}'_0 on $S \times \{0\}$ and \mathcal{T}'_1 on $S \times \{1\}$. By the above observation,

$$(3) \quad \Delta(\mathcal{T}') \leq (24)^{23} \deg(p) \Delta(\mathcal{T}_0, \mathcal{T}_1) + (2 \cdot 6^{23} + 8) \deg(p).$$

By Theorem 1.8, there is a constant $k_3 > 0$ depending only on S such that

$$(4) \quad \Delta(\mathcal{T}') \geq k_3 d_{\text{Tr}(S)}(\tilde{\mathcal{T}}'_0, \tilde{\mathcal{T}}'_1).$$

By Lemma 4.5, there are constants $k_4, k_5 > 0$ depending only on S and $\deg(p)$ such that

$$(5) \quad d_{\text{Tr}(S)}(\tilde{\mathcal{T}}'_0, \tilde{\mathcal{T}}'_1) \geq k_4 d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}'_0, \tilde{\mathcal{T}}'_1) - k_5.$$

Putting this all together, we obtain

$$(24)^{23} \deg(p) \Delta(\mathcal{T}_0, \mathcal{T}_1) + (2 \cdot 6^{23} + 8) \deg(p) \geq \Delta(\mathcal{T}') \quad \text{by (3)}$$

$$\geq k_3 d_{\text{Tr}(S)}(\tilde{\mathcal{T}}'_0, \tilde{\mathcal{T}}'_1) \quad \text{by (4)}$$

$$\geq k_3 (k_4 d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}'_0, \tilde{\mathcal{T}}'_1) - k_5) \quad \text{by (5)}$$

$$\geq k_3 k_4 (d_{\text{Tr}(S; -\chi(S))}(\tilde{\mathcal{T}}_0, \tilde{\mathcal{T}}_1) - 8 \deg(p)) - k_3 k_5 \quad \text{by (2)}$$

$$\geq k_3 k_4 k_1 d_{\text{Tr}(T)}(\mathcal{T}_0, \mathcal{T}_1) - k_3 k_4 k_2 - k_3 k_5 - 8 k_3 k_4 \deg(p) \quad \text{by (1)}$$

This gives a linear lower bound on $\Delta(\mathcal{T}_0, \mathcal{T}_1)$ in terms of $d_{\text{Tr}(T)}(\mathcal{T}_0, \mathcal{T}_1)$. For all but at most finitely many positive values of $d_{\text{Tr}(T)}(\mathcal{T}_0, \mathcal{T}_1)$, this lower bound will be positive and implies that there exists $k_{\text{prod}} > 0$ such that

$$\Delta(\mathcal{T}_0, \mathcal{T}_1) \geq k_{\text{prod}} d_{\text{Tr}(T)}(\mathcal{T}_0, \mathcal{T}_1).$$

For the remaining values, $\Delta(\mathcal{T}_0, \mathcal{T}_1)$ is positive, and so a universal $k_{\text{prod}} > 0$ can be chosen appropriately. \square

In [15], given a pre-tetrahedral handle structure \mathcal{H} on $S \times [0, 1]$, we considered the number of edge swaps required to transfer a cellular spine on $S \times \{0\}$ to a cellular spine on $S \times \{1\}$. The main technical theorem of that paper, [15, Theorem 9.11], gives a linear lower bound on the number of edge swaps in terms of $\Delta(\mathcal{H})$ under appropriate hypotheses. We will prove an analogous theorem for tori.

Theorem 9.1. *Let \mathcal{H} be a pre-tetrahedral handle structure for $T^2 \times [0, 1]$ that admits no annular simplification. Let Γ_0 be a cellular spine in $T^2 \times \{0\}$. Then there is a sequence of at most $k_{\text{hand}} \Delta(\mathcal{H})$ edge swaps taking Γ_0 to a spine Γ_1 that is cellular with respect to $T^2 \times \{1\}$. Here, k_{hand} is a universal constant.*

Proof. For ease of notation, set $\Gamma_0 = \Gamma^{(0)}$. In the course of the proof, we will obtain spines $\Gamma^{(0)}, \Gamma^{(1)}, \dots, \Gamma^{(n)}$, with each $\Gamma^{(j)}$ obtained from $\Gamma^{(j-1)}$ by at most $k^{(j)} \Delta(\mathcal{H})$ edge swaps, where $k^{(j)}$ is a universal constant. The number of spines in the sequence will be universally bounded (by $n = 6$), so the result follows.

First, let \mathcal{B} be a maximal generalised parallelity bundle for \mathcal{H} . By Lemma 2.15, each component of \mathcal{B} has incompressible horizontal boundary, and is either an I -bundle over a disc or has incompressible vertical boundary. Since $T^2 \times I$ does not contain a properly embedded Möbius band or Klein bottle, the only possibilities are that \mathcal{B} consists of a union of I -bundles over discs and annuli, or \mathcal{B} is all of $T^2 \times I$. We also claim that each component of \mathcal{B} that is an I -bundle over an annulus has one horizontal boundary component in $T^2 \times \{0\}$ and one horizontal boundary component in $T^2 \times \{1\}$. Suppose that on the contrary, there is an I -bundle over an annulus with both horizontal boundary components in the same component of $T^2 \times \partial I$. Then its two vertical boundary components are boundary-parallel. Pick such a component of \mathcal{B} that is outermost in $T^2 \times I$. Let A' be its vertical boundary component that is not outermost. Then $\partial A'$ bounds an annulus A in $T^2 \times \partial I$, and $A \cup A'$ bounds a product region P . Hence, \mathcal{H} admits an annular simplification, contrary to hypothesis.

Let $\mathcal{B}' \subset \mathcal{B}$ denote the parallelity bundle for \mathcal{H} . Recall from Definition 2.13 (6), the definition of a generalised parallelity bundle, that the intersection between $\partial_h \mathcal{B}'$ and each essential component A'' of $\partial_h \mathcal{B}$ contains a component that is equal to A'' with some discs D removed from its interior. As a first step, we will adjust $\Gamma_0 = \Gamma^{(0)}$ by sliding it off all such discs, using Lemma 4.9.

By Lemma 2.17, the length of $\partial_v \mathcal{B}'$ is at most $56\Delta(\mathcal{H})$. Hence, the length of ∂D is at most $112\Delta(\mathcal{H})$. In particular, the number of components of D is at most $112\Delta(\mathcal{H})$. Then Lemma 4.9 implies that after at most $6 \cdot 112\Delta(\mathcal{H}) + 2 \cdot 112\Delta(\mathcal{H}) = 896\Delta(\mathcal{H})$ edge swaps, we obtain a spine $\Gamma^{(1)}$ whose intersection with \mathcal{B} lies entirely within the parallelity bundle \mathcal{B}' : $\Gamma^{(1)} \cap \mathcal{B} \subset \mathcal{B}'$.

Case 1. \mathcal{B} is all of $T^2 \times [0, 1]$.

Then some component of the parallelity bundle \mathcal{B}' is of the form $(T^2 \setminus D) \times [0, 1]$, where D is a union of disjoint discs in T^2 . In particular, $\Gamma^{(1)}$ lies in parallelity

handles that run from $T^2 \times \{0\}$ to $T^2 \times \{1\}$ and that respect the product structure of $T^2 \times [0, 1]$. Transfer $\Gamma^{(1)}$ to $T^2 \times \{1\}$ using the product structure on \mathcal{B}' , obtaining a spine $\Gamma^{(2)}$ in $T^2 \times \{1\}$ without any additional edge swaps. Observe this is cellular in the cell structure associated with \mathcal{H} , so set $\Gamma_1 = \Gamma^{(2)}$. The proof is complete in this case.

Case 2. \mathcal{B} consists of I -bundles over discs and annuli.

In this case, we cannot ensure that all of $\Gamma^{(1)}$ lies only in the parallelity bundle, and so we cannot transfer as simply as in the previous case. Instead, we will obtain a triangulation from a simplified handle structure and apply Theorem 1.6.

First we adjust $\Gamma^{(1)}$ further. Again Lemma 2.17 implies that the length of $\partial_v \mathcal{B}$ is at most $56\Delta(\mathcal{H})$, and there are at most $112\Delta(\mathcal{H})$ components of $\partial_h \mathcal{B}$. We will now apply Lemma 4.11. Adjust $\Gamma^{(1)}$ to a new spine $\Gamma^{(2)}$ that is disjoint from the interior of the disc components of $\partial_h \mathcal{B}$, and intersects the interior of each annular component of $\partial_h \mathcal{B}$ in at most one arc. Moreover, this arc is a subset of $\Gamma^{(1)}$ and so the arc lies in the parallelity bundle \mathcal{B}' . By Lemma 4.11, $\Gamma^{(2)}$ is obtained from $\Gamma^{(1)}$ from a number of edge swaps bounded by

$$6 \cdot 112\Delta(\mathcal{H}) + 16 + 2 \cdot 112\Delta(\mathcal{H}) + 2 \cdot 112\Delta(\mathcal{H}) \leq 10 \cdot 112\Delta(\mathcal{H}) + 16 \leq 1248\Delta(\mathcal{H}).$$

Here we used that $\Delta(\mathcal{H}) \geq 1/8$, since otherwise $\Delta(\mathcal{H}) = 0$ and the parallelity bundle for \mathcal{H} is then all of \mathcal{H} , which is dealt with in Case 1.

Form a new handle structure \mathcal{H}' as follows. Replace each component of \mathcal{B} that is an I -bundle over a disc by a single 2-handle. Replace each component of \mathcal{B} that is an I -bundle over an annulus by a 1-handle and a 2-handle, arranged such that the intersection of these two handles contains an arc of $\Gamma^{(2)}$ within the annulus, if there is one. This is possible because each component of \mathcal{B} that is an I -bundle over an annulus intersects $T^2 \times \{0\}$ in a single component. Then $\Gamma^{(2)}$ remains a cellular spine in this new handle structure \mathcal{H}' . Note \mathcal{H}' has no parallelity 0-handles.

However, note that \mathcal{H}' may no longer be pre-tetrahedral. For each component of \mathcal{B} that is not a 2-handle, its vertical boundary consists of alternating components of intersection with 1-handles and 2-handles, and these must be adjacent to semi-tetrahedral 0-handles, attached where a 1-handle is bounded by exactly two 2-handles. The process of replacing a $D^2 \times I$ component of \mathcal{B} replaces one or both instance of such a 1-handle and its adjacent 2-handles with a single 2-handle. This gives a clipped semi-tetrahedral 0-handle as in Figure 3. Similarly clipped semi-tetrahedral 0-handles may arise when replacing components that are I -bundles over an annulus, but these are the only adjustments that need to be made.

Recall that we may still define the complexity of a handle structure that is pre-tetrahedral aside from a finite number of clipped semi-tetrahedral 0-handles. Since clipped semi-tetrahedral handles contribute the same as semi-tetrahedral handles to complexity, we have $\Delta(\mathcal{H}') \leq \Delta(\mathcal{H})$.

Let \mathcal{T} be the triangulation obtained from \mathcal{H}' as in Lemma 2.20. This satisfies $\Delta(\mathcal{T}) \leq 1152\Delta(\mathcal{H}') \leq 1152\Delta(\mathcal{H})$. Let \mathcal{T}_0 and \mathcal{T}_1 be the induced triangulations of $T^2 \times \{0\}$ and $T^2 \times \{1\}$. Observe that our spine $\Gamma^{(2)}$ is a subcomplex of \mathcal{T}_0 .

By Lemma 4.13, there is a sequence of at most $16\Delta(\mathcal{T})$ Pachner moves taking \mathcal{T}_0 and \mathcal{T}_1 to 1-vertex triangulations \mathcal{T}'_0 and \mathcal{T}'_1 . By Lemma 4.12, there is a sequence of at most $16\Delta(\mathcal{T}) \leq 16 \cdot 1152\Delta(\mathcal{H})$ edge swaps taking $\Gamma^{(2)}$ to a subcomplex $\Gamma^{(3)}$ of \mathcal{T}'_0 .

Realise the Pachner moves as tetrahedra, to form a triangulation \mathcal{T}' of $T^2 \times [0, 1]$ with triangulations \mathcal{T}'_0 and \mathcal{T}'_1 on $T^2 \times \{0\}$ and $T^2 \times \{1\}$, respectively, and with $\Delta(\mathcal{T}') \leq 17\Delta(\mathcal{T})$.

By Theorem 1.6, $d_{\text{Tr}(T^2)}(\mathcal{T}'_0, \mathcal{T}'_1)$ is at most $k_{\text{prod}} \Delta(\mathcal{T}')$ for some universal constant $k_{\text{prod}} > 0$. That is, there is a sequence of at most $k_{\text{prod}} \Delta(\mathcal{T}') \leq 17k_{\text{prod}} \Delta(\mathcal{T})$ Pachner moves taking \mathcal{T}'_0 to \mathcal{T}'_1 . Follow this by at most $16\Delta(\mathcal{T})$ Pachner moves taking \mathcal{T}'_1 to \mathcal{T}_1 , for a total of $(17k_{\text{prod}} + 16)\Delta(\mathcal{T})$ Pachner moves taking \mathcal{T}'_0 to \mathcal{T}_1 . Again Lemma 4.12 implies there is a sequence of at most $(17k_{\text{prod}} + 16)\Delta(\mathcal{T}') \leq k\Delta(\mathcal{H})$ edge swaps taking the spine $\Gamma^{(3)}$ to a subcomplex $\Gamma^{(4)}$ of \mathcal{T}_1 , where $k > 0$ is a universal constant.

The 1-skeleton of the cell structure associated with \mathcal{H}' on $T^2 \times \{1\}$ is a subcomplex of \mathcal{T}_1 , so the next step is to adjust $\Gamma^{(4)}$ to lie only in this subcomplex. The 2-cells of the cell structure are discs. The total length of their boundary is at most the total length of the 1-skeleton of the tetrahedra of \mathcal{T} , which is at most $6\Delta(\mathcal{T})$. The number of discs is at most the number of triangles in \mathcal{T} , which is at most $4\Delta(\mathcal{T})$. Hence by Lemma 4.9, we may modify $\Gamma^{(4)}$ to a spine $\Gamma^{(5)}$ that is cellular with respect to \mathcal{H}' using at most $36\Delta(\mathcal{T})$ edge swaps.

The cell structure that $T^2 \times \{1\}$ inherits from \mathcal{H}' agrees with that of \mathcal{H} away from $\partial_h \mathcal{B}$. Within the interior of each essential annular component of $\partial_h \mathcal{B}$, the structure of \mathcal{H}' consists of two 1-cells and two 2-cells. One of these 1-cells is the arc of intersection with $\Gamma^{(2)}$, which we arranged to be part of the parallelity bundle. Hence, it is also cellular in the cell structure that inherits from \mathcal{H} . If necessary, a single edge swap takes $\Gamma^{(5)}$ to a spine intersecting the interior of this annulus just in this arc. As there are at most $56\Delta(\mathcal{H})$ annular components of $\partial_h \mathcal{B}$ in $T^2 \times \{1\}$, this can be done for all annuli with at most $56\Delta(\mathcal{H})$ additional edge swaps, obtaining a spine $\Gamma^{(6)}$. Now $\Gamma^{(6)}$ is cellular with respect to \mathcal{H} , so we set this equal to Γ_1 .

In summary, starting with the spine $\Gamma_0 = \Gamma^{(0)}$, we have found a sequence of edge swaps taking the spine $\Gamma^{(i)}$ to a spine $\Gamma^{(i+1)}$, for $i = 0, \dots, 5$, where $\Gamma^{(6)} = \Gamma_1$ is the desired cellular spine on $T^2 \times \{1\}$, such that the number of edge swaps required to take $\Gamma^{(i)}$ to $\Gamma^{(i+1)}$ is bounded by a uniform constant times $\Delta(\mathcal{H})$, for each $i = 1, \dots, 5$. \square

10. TRIANGULATIONS OF SOL MANIFOLDS

In this section, we prove Theorems 1.4, 1.5 and 1.9.

Lemma 10.1. *Let $\phi: T^2 \rightarrow T^2$ be a linear Anosov homeomorphism. Then the triangulation complexity satisfies*

$$\Delta((T^2 \times I)/\phi) \leq 6 + \ell_{\text{Tr}(T^2)}(\phi).$$

Proof. There is a 1-vertex triangulation t of the torus such that the distance in the Farey tree between t and $\phi(t)$ realises the translation length $\ell_{\text{Tr}(T^2)}(\phi)$. As in the proof of Theorem 1.6, we triangulate $T^2 \times [0, 1]$ using $6 + \ell_{\text{Tr}(T^2)}(\phi)$ tetrahedra, with $T^2 \times \{0\}$ triangulated using t and $T^2 \times \{1\}$ triangulated using $\phi(t)$. Then glue top to bottom using ϕ . The result is a triangulation of $(T^2 \times I)/\phi$ with $6 + \ell_{\text{Tr}(T^2)}(\phi)$ tetrahedra. \square

Lemma 10.2. *Let $\phi: T^2 \rightarrow T^2$ be a linear Anosov homeomorphism. Then there exists a universal constant $k'_{\text{sol}} > 0$ such that*

$$k'_{\text{sol}} \ell_{\text{Tr}(T^2)}(\phi) \leq \Delta((T^2 \times I)/\phi).$$

Proof. Let $M = (T^2 \times I)/\phi$. We will show that $\Delta(M)$ is at least a constant times the translation distance of ϕ in the spine graph $\text{Sp}(T^2)$. As $\text{Sp}(T^2)$ and $\text{Tr}(T^2)$ are quasi-isometric, by Lemma 4.5, this will prove the result.

Consider a triangulation \mathcal{T} for M with $\Delta(\mathcal{T}) = \Delta(M)$. Let S be a normal fibre in M with least weight. This corresponds to a surface, also called S , that intersects each handle in the dual handle structure in a collection of properly embedded discs. This surface S inherits a cell structure in which each of these discs is a 2-cell. Pick some spine Γ for S that is cellular. Let \mathcal{H} be the handle structure that results from cutting M along S . Then by Lemma 2.5, $\Delta(\mathcal{H}) = \Delta(\mathcal{T})$.

Since S has least weight in its isotopy class, \mathcal{H} does not admit any annular simplifications, by Theorem 2.16. So, by Theorem 9.1, there is a sequence of at most $k_{\text{hand}}\Delta(\mathcal{H})$ edge contractions and expansions taking Γ in $T^2 \times \{0\}$ to a spine Γ_1 in $T^2 \times \{1\}$ that is cellular. Now apply the gluing map ϕ between $T^2 \times \{1\}$ and $T^2 \times \{0\}$ to get the spine $\phi(\Gamma_1)$ in $T^2 \times \{0\}$. Next apply Theorem 9.1 again, to obtain a sequence of at most $k_{\text{hand}}\Delta(\mathcal{H})$ edge contractions and expansions taking $\phi(\Gamma_1)$ in $T^2 \times \{0\}$ to a cellular spine Γ_2 in $T^2 \times \{1\}$. Keep repeating this process, giving a sequence of spines Γ_i that are cellular in $T^2 \times \{1\}$. Thus, the distance in $\text{Sp}(T^2)$ between $\phi(\Gamma_i)$ and Γ_{i+1} is at most $k_{\text{hand}}\Delta(\mathcal{H})$. There are only finitely many cellular spines in $T^2 \times \{1\}$ and so there are integers $r < s$ such that $\Gamma_r = \Gamma_s$. By relabelling, we may assume that $r = 0$ and $s = n$, say. Thus, with respect to the metric d on $\text{Sp}(T^2)$, we have the following inequalities:

$$\begin{aligned} d(\phi^n \Gamma_0, \Gamma_0) &= d(\phi^n \Gamma_0, \Gamma_n) \\ &\leq d(\phi^n \Gamma_0, \phi^{n-1} \Gamma_1) + d(\phi^{n-1} \Gamma_1, \phi^{n-2} \Gamma_2) + \cdots + d(\phi \Gamma_{n-1}, \Gamma_n) \\ &= d(\phi \Gamma_0, \Gamma_1) + d(\phi \Gamma_1, \Gamma_2) + \cdots + d(\phi \Gamma_{n-1}, \Gamma_n) \\ &\leq k_{\text{hand}} n \Delta(\mathcal{H}). \end{aligned}$$

So, the translation length $\ell_{\text{Sp}(T^2)}(\phi^n)$ of ϕ^n is at most $k_{\text{hand}} n \Delta(\mathcal{H})$. But $\ell_{\text{Sp}(T^2)}(\phi^n)$ is n times the translation length $\ell_{\text{Sp}(T^2)}(\phi)$, since ϕ acts on the tree $\text{Sp}(T^2)$ by translation along an axis; see Lemma 7.2. Therefore, the translation length of ϕ is at most $k_{\text{hand}} \Delta(\mathcal{H}) = k_{\text{hand}} \Delta(M)$. \square

Recall from Section 5.1 that $\text{PSL}(2, \mathbb{Z})$ is isomorphic to $\mathbb{Z}_2 * \mathbb{Z}_3$ where the factors are generated by

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad T = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}.$$

Theorem 1.4. *Let A be an element of $\text{SL}(2, \mathbb{Z})$ with $|\text{tr}(A)| > 2$. Let M be the sol 3-manifold $(T^2 \times [0, 1])/(A(x, 1) \sim (x, 0))$. Let \bar{A} be the image of A in $\text{PSL}(2, \mathbb{Z})$ and let $\ell(\bar{A})$ be the length of a cyclically reduced word in the generators S and $T^{\pm 1}$ that is conjugate to \bar{A} . Then, there is a universal constant $k_{\text{sol}} > 0$ such that*

$$k_{\text{sol}} \ell(\bar{A}) \leq \Delta(M) \leq (\ell(\bar{A})/2) + 6.$$

Proof. As explained in Section 5.1, the Farey tree is closely related to the Cayley graph of $\mathbb{Z}_2 * \mathbb{Z}_3$ with respect to the generators S and T . Specifically this Cayley graph is obtained from the Farey tree as follows: replace each vertex of the tree by a triangle, with each edge oriented and labelled by T ; replace each edge of the Farey tree by two edges, both labelled by S and pointing in opposite directions. The element A in $\text{SL}(2, \mathbb{Z})$ acts on this Cayley graph as it does on the Farey tree, and the translation lengths of these two actions differ by a factor of 2. Indeed,

the invariant axis in the Farey tree can be used to produce an invariant geodesic in the Cayley graph, and the translation length along this geodesic is twice that of the translation length along the axis in the Farey tree. As one travels along this geodesic, one reads off a word in S , T and T^{-1} which is a cyclically reduced representative for a conjugate of \bar{A} . Thus its length $\ell(\bar{A})$ is twice the translation length of the action of A on the Farey tree. By Lemmas 10.2 and 10.1, this is, up to a bounded multiplicative factor, the triangulation complexity of M . \square

Theorem 1.5. *Let A be an element of $\mathrm{SL}(2, \mathbb{Z})$ with $|\mathrm{tr}(A)| > 2$. Let \bar{A} be the image of A in $\mathrm{PSL}(2, \mathbb{Z})$. Suppose that \bar{A} is B^n for some positive integer n and some $B \in \mathrm{PSL}(2, \mathbb{Z})$ that cannot be expressed as a proper power. Let M be the sol 3-manifold $(T^2 \times [0, 1]) / (A(x, 1) \sim (x, 0))$. Let $[a_0, a_1, \dots]$ be the continued fraction expansion of $\sqrt{\mathrm{tr}(A)^2 - 4}$ where a_i is positive for each $i > 0$ and let (a_r, \dots, a_s) denote its periodic part. Then there is a universal constant $k'_{\mathrm{sol}} > 0$ such that*

$$k'_{\mathrm{sol}} n \sum_{i=r}^s a_i \leq \Delta(M) \leq 6 + n \sum_{i=r}^s a_i.$$

Proof. Let $\phi: T^2 \rightarrow T^2$ be the homeomorphism determined by A . By Lemma 7.1, ϕ is linear Anosov, so by Lemmas 10.1 and 10.2, we have

$$k'_{\mathrm{sol}} \ell_{\mathrm{Tr}(T^2)}(\phi) \leq \Delta(M) \leq 6 + \ell_{\mathrm{Tr}(T^2)}(\phi).$$

Now by Theorem 5.1, $\ell_{\mathrm{Tr}(T^2)}(\phi)$ is the translation length of ϕ in the Farey tree. By Proposition 7.3, this translation length is $n \sum_{i=r}^s a_i$. \square

Theorem 1.9. *Let $\phi: T^2 \rightarrow T^2$ be a linear Anosov homeomorphism. Then the following quantities are within bounded ratios of each other:*

- (1) the triangulation complexity of $(T^2 \times I) / \phi$;
- (2) the translation distance (or stable translation distance) of ϕ in the thick part of the Teichmüller space of T^2 ;
- (3) the translation distance (or stable translation distance) of ϕ in the mapping class group of T^2 ;
- (4) the translation distance (or stable translation distance) of ϕ in $\mathrm{Tr}(T^2)$.

In (3), we metrize $\mathrm{MCG}(T^2)$ by fixing the finite generating set

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

Proof. The fact that the quantities (2), (3) and (4) are within a bounded ratio of each other is a rapid consequence of the Milnor-Svarc lemma, as explained in [15, Section 1.2]. The relationship between (1) and (4) now follows immediately from Lemmas 10.1 and 10.2. \square

11. TRIANGULATIONS OF LENS SPACES

The following was one of the main theorems of [16].

Theorem 11.1. *Let M be a lens space other than a prism manifold $L(4p, 2p \pm 1)$ or \mathbb{RP}^3 . Let \mathcal{T} be any triangulation of M . Then the iterated barycentric subdivision $\mathcal{T}^{(139)}$ contains in its 1-skeleton the union of the two core curves.*

We will use this to prove our main result about lens spaces.

Theorem 1.1. *Let $L(p, q)$ be a lens space, where p and q are coprime integers satisfying $0 < q < p$. Let $[a_0, \dots, a_n]$ be the continued fraction expansion of p/q where each a_i is positive. Then there is a universal constant $k_{\text{lens}} > 0$ such that*

$$k_{\text{lens}} \sum_{i=0}^n a_i \leq \Delta(L(p, q)) \leq \sum_{i=0}^n a_i.$$

Proof. The upper bound on $\Delta(L(p, q))$ is fairly straightforward. Indeed, Jaco and Rubinstein [9] provide a triangulation with $(\sum_{i=0}^n a_i) - 3$ tetrahedra for $p > 3$ and they conjecture that this is equal to $\Delta(L(p, q))$. For $p = 2$ or 3 , the lens space is \mathbb{RP}^3 or $L(3, 1)$, which both satisfy $\Delta(L(p, q)) = 2$. Note $\sum_{i=0}^n a_i \geq 2$, so the inequality holds in these cases.

We now focus on the lower bound for $\Delta(L(p, q))$. The triangulation complexity of $L(4p', 2p' \pm 1)$ was shown by Jaco, Rubinstein and Tillmann [11] to be p' for $p' \geq 2$. So we now assume that the lens space is not of this form and also is not \mathbb{RP}^3 or $L(4, 1)$.

Let \mathcal{T} be a triangulation of $L(p, q)$. Our goal is to show that $\Delta(\mathcal{T}) \geq k_{\text{lens}} \sum_{i=0}^n a_i$ for some universal constant $k_{\text{lens}} > 0$. The approach that we will take is as follows. We will drill from (a subdivision of) the triangulation \mathcal{T} two core curves C and C' for the solid tori making up $L(p, q)$, with meridians μ and μ' . This gives a manifold homeomorphic to $T^2 \times [0, 1]$. Using the triangulation \mathcal{T} , we will obtain spines for $T^2 \times \{0\}$ and $T^2 \times \{1\}$ containing μ and μ' , respectively; these correspond to points on $L(\mu)$ and $L(\mu')$ in the Farey tree. Theorem 9.1 and Lemma 4.15 then build a path in the Farey tree from $L(\mu)$ to $L(\mu')$, with length bounded by a constant times $\Delta(\mathcal{T})$. This must be at least as long as the shortest path from $L(\mu)$ to $L(\mu')$, which is $\sum a_i$. Thus our constructed path will give a bound of the form $\sum a_i \leq (1/k_{\text{lens}})\Delta(\mathcal{T})$ for some constant k_{lens} . This completes the outline of the proof of the theorem.

By Theorem 11.1, $\mathcal{T}^{(139)}$ contains in its 1-skeleton the union of the two core curves C and C' . To build our spines on $T^2 \times \{0\}$ and $T^2 \times \{1\}$, we not only need C and C' to be simplicial, but also a regular neighbourhood of these curves must be simplicial, and a meridian of the simplicial neighbourhood must also be simplicial. This can be obtained by further subdivision. However, as we subdivide, we will need a bound on the length of the simplicial meridian to apply Lemma 4.15. Again this can be obtained by careful subdivision, as follows. Beginning with $\mathcal{T}^{(139)}$, take a further two barycentric subdivisions, so that a regular neighbourhood $N(C \cup C')$ for $C \cup C'$ is simplicial in $\mathcal{T}^{(141)}$. Observe that $\partial N(C \cup C')$ consists of faces, edges, and vertices that belong to $\mathcal{T}^{(141)}$ but do not lie in faces, edges and vertices (respectively) of $\mathcal{T}^{(139)}$. Observe that barycentric subdivision adds a new central vertex to each tetrahedron of $\mathcal{T}^{(140)}$, and this vertex meets exactly 24 tetrahedra in $\mathcal{T}^{(141)}$. Observe also that an edge running from this central vertex to one of the new vertices on a face will meet 30 tetrahedra: 24 at one endpoint, and an additional six at the other endpoint.

We may find two curves \bar{C} and \bar{C}' that are simplicial on $\partial N(C \cup C')$, that are parallel copies of C and C' , respectively, and that are made up of edges each meeting at most 30 tetrahedra. We may take a further two barycentric subdivisions, creating $\mathcal{T}^{(143)}$, so that $N(\bar{C} \cup \bar{C}')$ is simplicial. Then $N(\bar{C} \cup \bar{C}')$ consists of those simplices in $\mathcal{T}^{(143)}$ that have non-empty intersection with \bar{C} and \bar{C}' . We remove the interior of this regular neighbourhood from $L(p, q)$, and thereby obtain a triangulation of

$T^2 \times [0, 1]$. Because each barycentric subdivision increased the number of tetrahedra by a factor of 24, this triangulation has complexity at most $(24)^{143} \Delta(\mathcal{T})$.

Consider any vertex v of $\mathcal{T}^{(141)}$ lying on \overline{C} . The union of the simplices in $\mathcal{T}^{(143)}$ incident to this vertex is a simplicial 3-ball B . The intersection between ∂B and \overline{C} consists of two points. The union of the simplices in ∂B incident to one of these points is a disc. The boundary of this disc is a meridian curve μ for $N(\overline{C})$. Since at most 30 tetrahedra of $\mathcal{T}^{(141)}$ are incident to v , we deduce that there is a universal upper bound (120 in fact) for the length of μ in $\mathcal{T}^{(143)}$. Each vertex of μ lies in the interior of a 3-simplex or 2-simplex of $\mathcal{T}^{(141)}$. Hence, it is incident to at most two 3-simplices of $\mathcal{T}^{(141)}$. These 3-simplices contain at most $2 \times (24)^2$ tetrahedra of $\mathcal{T}^{(143)}$ and hence at most $12 \times (24)^2 = 6912$ edges of $\mathcal{T}^{(143)}$. We deduce that each vertex of μ is incident to at most 6912 edges in $\partial N(\overline{C})$. Similarly, on $\partial N(\overline{C}')$, there is a meridian curve μ' with length at most 120 and again with the property that each vertex that it runs through is incident to at most 6912 edges in $\partial N(\overline{C}')$.

As in Remark 2.7, we attach a triangulation of $\partial N(C \cup C')$ onto the triangulation of $T^2 \times [0, 1]$, to form a new triangulation \mathcal{T}' of $T^2 \times [0, 1]$. This satisfies $\Delta(\mathcal{T}') \leq 33(24)^{143} \Delta(\mathcal{T})$. By Lemma 2.6, the dual handle structure \mathcal{H} is pre-tetrahedral and satisfies $\Delta(\mathcal{H}) \leq 33(24)^{143} \Delta(\mathcal{T})$. As explained in Remark 2.10, it has no parallelity handles. In particular, it admits no annular simplifications.

There is a copy of the meridian curve μ in \mathcal{T}' . This is a sequence of vertices and edges, and hence it corresponds to a sequence of 2-handles and 1-handles in the handle structure of $T^2 \times \{0, 1\}$. The boundary of each of these 2-handles has length at most 6912×2 in the cell structure, and the boundary of each 1-handle has length 4. The union of these 1-handles and 2-handles is an annulus. Each of the boundary components of the annulus is cellular and has length at most $120 \times ((6912 \times 2) + 4) = 1659360$ in \mathcal{H} . Pick one of these boundary components and extend it to a cellular spine Γ in $T^2 \times \{0\}$. Similarly, there is a cellular curve in $T^2 \times \{1\}$ that is parallel to μ' and that has length at most 1659360.

By Theorem 9.1 and Lemma 4.7, there is a universal constant $k_{\text{hand}} > 0$ and a sequence of at most $24k_{\text{hand}} \Delta(\mathcal{H})$ edge contractions and expansions taking Γ to a cellular spine Γ' in $T^2 \times \{1\}$. By Lemma 4.15, there is a further sequence of at most $24 + 4 \cdot 1659360 = 6637464$ edge swaps taking Γ' to a spine Γ'' that contains μ' . By Lemma 4.7, this is realised by at most $24 \cdot 6637464 < 10^8$ edge contractions and expansions. Since Γ and Γ'' contain meridians as subsets, they correspond to vertices in the Farey tree that lie on the lines $L(\mu)$ and $L(\mu')$, as in Definition 5.3. The distance in the spine graph is exactly twice the distance in the Farey tree, since a 2-2 Pachner move in a triangulation is realised by an edge contraction then expansion. Hence, the distance between these lines is at most $24k_{\text{hand}} \Delta(\mathcal{H}) + 10^8 \leq 33(24)^{144} k_{\text{hand}} \Delta(\mathcal{T}) + 10^8$ in the Farey tree. By Lemma 5.4, this distance is $(\sum_{i=0}^n a_i) - 1$, where $[a_0, \dots, a_n]$ is the continued fraction expansion of p/q . Hence, $\Delta(L(p, q))$ is at least a linear function of $\sum_{i=0}^n a_i$. Since $\sum_{i=0}^n a_i \geq 1$, the additive part of this linear function can be eliminated, at the possible cost of changing the multiplicative constant. So, $\Delta(L(p, q))$ is at least $k_{\text{lens}} \sum_{i=0}^n a_i$ for some universal $k_{\text{lens}} > 0$. \square

12. PRISM MANIFOLDS AND PLATONIC MANIFOLDS

We start by considering the prism manifold $P(p, q)$. This is obtained by gluing together the solid torus and $K^2 \times I$, the orientable I -bundle over the Klein bottle, via

a homeomorphism between their boundaries. The resulting manifold is determined, up to homeomorphism, by the slope on the boundary of $K^2 \times I$ to which a meridian disc of the solid torus is attached. Now, the boundary of $K^2 \times I$ has a canonical framing, as follows. There are only two non-separating simple closed curves λ and μ on the Klein bottle, where λ is orientation-reversing and μ is orientation-preserving. The inverse images of these in the boundary of $K^2 \times I$ are curves with slopes λ and $\tilde{\mu}$. The prism manifold $P(p, q)$ is obtained by attaching the meridian disc of the solid torus along a curve with slope $p\tilde{\lambda} + q\tilde{\mu}$, when these slopes are given some choice of orientation.

Theorem 1.2. *Let p and q be non-zero coprime integers and let $[a_0, \dots, a_n]$ denote the continued fraction expansion of p/q where a_i is positive for each $i > 0$. Then, $\Delta(P(p, q))$ is, to within a universally bounded multiplicative error, equal to $\sum_{i=0}^n a_i$.*

Proof. As usual, the upper bound on $\Delta(P(p, q))$ is fairly straightforward. Start with a 1-vertex triangulation of the Klein bottle in which λ and μ are edges. The inverse image in $K^2 \times I$ of the three edges in K^2 is three squares, which we can triangulate using two triangles each. If we cut $K^2 \times I$ along these three squares, the result is two prisms, which can be triangulated using eight tetrahedra. Now attach onto the boundary of $K^2 \times I$ some tetrahedra, so that the resulting boundary has a 1-vertex triangulation, and where two of its edges have slopes $\tilde{\lambda}$ and $\tilde{\mu}$. Now apply 2-2 Pachner moves taking this triangulation to one that includes p/q as an edge. Then glue on a triangulation of the solid torus with boundary triangulation containing a meridian as an edge. Thus the resulting number of tetrahedra in the triangulation of $P(p, q)$ is at most $\sum_{i=0}^n a_i$ plus a constant.

For the lower bound, note that $P(p, q)$ is double covered by a lens space L . Hence, $\Delta(P(p, q)) \geq \Delta(L)/2$. The inverse image of $K^2 \times I$ in this double cover is a copy of $T^2 \times [0, 1]$. The inverse image of the solid torus is two solid tori, one attached along the slope p/q and the other attached along the slope $-p/q$. According to Theorem 1.1 and Lemma 5.4, $\Delta(L)$ is at least a constant times the distance in the Farey graph between the lines $L(p/q)$ and $L(-p/q)$. To compute this distance, consider the hyperbolic geodesic joining p/q to $-p/q$. This is divided into two half-infinite geodesics by the imaginary axis. As we travel along one of these geodesics, starting at the imaginary axis and ending at p/q , we recover the splitting sequence for p/q . Hence, this corresponds to a path in the Farey tree with length $\sum_{i=0}^n a_i$. The path running from the imaginary axis to $L(-p/q)$ has the same length. Hence, $\Delta(L)$ is at least a constant times $\sum_{i=0}^n a_i$. \square

A Platonic manifold is an elliptic manifold M that admits a Seifert fibration with base space that is the quotient of the 2-sphere by the orientation-preserving symmetry group of a Platonic solid. In other words, the base space Σ is a 2-sphere with three cone points of orders $(2, 3, 3)$, $(2, 3, 4)$ or $(2, 3, 5)$.

If we remove the three singular fibres from M , the result is a circle bundle over the three-holed sphere. Thus, M is obtained from this circle bundle by attaching three solid tori. For convenience, we also remove one regular fibre, and the resulting manifold M' is a circle bundle over the four-holed sphere S . Now orientable circle bundles over orientable surfaces with non-empty boundary are trivial. Thus, M' is just a copy of $S \times S^1$. We fix a meridian and longitude for each boundary component of $S \times S^1$, by declaring the longitude to be of the form $C \times \{*\}$ for the relevant boundary component C of S , and declaring the meridian to be $\{*\} \times S^1$. Thus,

the Dehn filling slopes are given by four fractions p_0/q_0 , p_1/q_1 , p_2/q_2 and p_3/q_3 , where $q_0 = 1$, $q_1 = 2$, $q_2 = 3$ and $q_3 = 3, 4$ or 5 . The Euler number is just the sum $p_0/q_0 + p_1/q_1 + p_2/q_2 + p_3/q_3$. Without changing the manifold or its Seifert fibration, we can adjust these slopes by adding an integer to one and subtracting an integer from another. In this way, we can arrange p_1/q_1 , p_2/q_2 and p_3/q_3 all to lie strictly between 0 and 1.

Theorem 1.3. *Let M be a Platonic elliptic 3-manifold, and let e denote the Euler number of its Seifert fibration. Then, to within a universally bounded multiplicative error, $\Delta(M)$ is $|e|$.*

Proof. The upper bound is straightforward. We can form a triangulation of $S \times S^1$ with a fixed number of tetrahedra, and where the longitudes and meridians are all simplicial. We can also arrange that the triangulation of each boundary component has a single vertex. Since p_1/q_1 , p_2/q_2 and p_3/q_3 take only finitely many possible values, we can attach triangulated solid tori so that the meridian disc is attached to these slopes, using a universally bounded number of tetrahedra. The final slope $p_0/1$ is integral. Hence, using at most $|p_0|$ many 2-2 Pachner moves, we may arrange that this slope $p_0/1$ is simplicial. We can then attach a triangulated solid torus to form a triangulation of M . The difference between p_0 and the Euler number e is bounded above by 3, since p_1/q_1 , p_2/q_2 and p_3/q_3 all lie between 0 and 1. So, the number of tetrahedra is at most $|e| + c$ for some universal constant c . This is at most a multiple of $|e|$ as e cannot be zero, since M would then contain an embedded non-separating orientable surface, which is impossible in a rational homology 3-sphere.

We now establish the lower bound on $\Delta(M)$. The Seifert fibration $M \rightarrow \Sigma$ induces a surjective homomorphism $\pi_1(M) \rightarrow \pi_1(\Sigma)$, where the latter group is the orbifold fundamental group of the base space Σ . The kernel of this homomorphism has index at most 60. Let \tilde{M} be the corresponding finite cover of M . Then $\Delta(\tilde{M}) \leq 60\Delta(M)$. The Seifert fibration on M lifts to a Seifert fibration on \tilde{M} . The Euler number \tilde{e} of \tilde{M} is related to the Euler number of M as follows, using [24, Theorem 3.6]. If d_1 is the degree of the covering between the base orbifolds, and d_2 is the degree of the coverings between regular fibres, then $\tilde{e} = ed_1/d_2$. In particular, $|\tilde{e}| \geq |e|/60$. The Seifert fibration of \tilde{M} has base space a 2-sphere and has no singular fibres, and therefore \tilde{M} is the lens space $L(\tilde{e}, 1)$. Therefore, by Theorem 1.1,

$$\Delta(M) \geq \Delta(\tilde{M})/60 \geq (k_{\text{lens}}/60)|\tilde{e}| \geq (k_{\text{lens}}/3600)|e|. \quad \square$$

REFERENCES

1. Ian Agol, *Ideal triangulations of pseudo-Anosov mapping tori*, Topology and geometry in dimension three, Contemp. Math., vol. 560, Amer. Math. Soc., Providence, RI, 2011, pp. 1–17. MR 2866919 [23]
2. Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR 1744486 [14, 15]
3. Andrew J. Casson and Steven A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts, vol. 9, Cambridge University Press, Cambridge, 1988. MR 964685 [24]
4. H. Davenport, *The higher arithmetic*, eighth ed., Cambridge University Press, Cambridge, 2008, An introduction to the theory of numbers, With editing and additional material by James H. Davenport. MR 2462408 [20]
5. Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012. MR 2850125 [28]

6. Albert Fathi, François Laudenbach, and Valentin Poénaru, *Thurston's work on surfaces*, Mathematical Notes, vol. 48, Princeton University Press, Princeton, NJ, 2012, Translated from the 1979 French original by Djun M. Kim and Dan Margalit. MR 3053012 [23]
7. William Jaco, Hyam Rubinstein, Jonathan Spreer, and Stephan Tillmann, *On minimal ideal triangulations of cusped hyperbolic 3-manifolds*, J. Topol. **13** (2020), no. 1, 308–342. MR 4138740 [1]
8. William Jaco, Hyam Rubinstein, and Stephan Tillmann, *Minimal triangulations for an infinite family of lens spaces*, J. Topol. **2** (2009), no. 1, 157–180. MR 2499441 [1]
9. William Jaco and J. Hyam Rubinstein, *Layered triangulations of 3-manifolds*, arXiv:math/0603601, 2006. [1, 37]
10. William Jaco, J. Hyam Rubinstein, Jonathan Spreer, and Stephan Tillmann, *\mathbb{Z}_2 -Thurston norm and complexity of 3-manifolds, II*, Algebr. Geom. Topol. **20** (2020), no. 1, 503–529. MR 4071382 [1]
11. William Jaco, J. Hyam Rubinstein, and Stephan Tillmann, *Coverings and minimal triangulations of 3-manifolds*, Algebr. Geom. Topol. **11** (2011), no. 3, 1257–1265. MR 2801418 [1, 37]
12. ———, *\mathbb{Z}_2 -Thurston norm and complexity of 3-manifolds*, Math. Ann. **356** (2013), no. 1, 1–22. MR 3038119 [1]
13. Marc Lackenby, *The crossing number of composite knots*, J. Topol. **2** (2009), no. 4, 747–768. MR 2574742 [7, 8, 9]
14. ———, *The efficient certification of knottedness and Thurston norm*, Adv. Math. **387** (2021), Paper No. 107796, 142. MR 4274879 [17]
15. Marc Lackenby and Jessica Purcell, *The triangulation complexity of fibred 3-manifolds*, Geom. Topol., to appear. arXiv:1910.10914, 2019. [3, 4, 5, 6, 8, 9, 14, 15, 32, 36]
16. Marc Lackenby and Saul Schleimer, *Recognising elliptic manifolds*, arXiv:2205.08802, 2022. [11, 36]
17. Bruno Martelli and Carlo Petronio, *Three-manifolds having complexity at most 9*, Experiment. Math. **10** (2001), no. 2, 207–236. MR 1837672 [1]
18. Howard Masur, Lee Mosher, and Saul Schleimer, *On train-track splitting sequences*, Duke Math. J. **161** (2012), no. 9, 1613–1656. MR 2942790 [5, 23]
19. S. Matveev, *Algorithmic topology and classification of 3-manifolds*, second ed., Algorithms and Computation in Mathematics, vol. 9, Springer, Berlin, 2007. MR 2341532 [11]
20. S. V. Matveev, *Tabulation of three-dimensional manifolds*, Uspekhi Mat. Nauk **60** (2005), no. 4(364), 97–122. MR 2190925 [1]
21. Sergei Matveev, Carlo Petronio, and Andrei Vesnin, *Two-sided asymptotic bounds for the complexity of some closed hyperbolic three-manifolds*, J. Aust. Math. Soc. **86** (2009), no. 2, 205–219. MR 2507595 [1]
22. S.V. Matveev, *Tables of 3-manifolds up to complexity 6*, Max-Planck Institute Preprint, MPI, 1998, pp. 1–50. [1]
23. Carlo Petronio and Andrei Vesnin, *Two-sided bounds for the complexity of cyclic branched coverings of two-bridge links*, Osaka J. Math. **46** (2009), no. 4, 1077–1095. MR 2604922 [1]
24. Peter Scott, *The geometries of 3-manifolds*, Bull. London Math. Soc. **15** (1983), no. 5, 401–487. MR 705527 [1, 40]
25. Jean-Pierre Serre, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation. MR 1954121 [18]
26. William P. Thurston, *The geometry and topology of three-manifolds*, Princeton Univ. Math. Dept. Notes, 1979, Available at <http://www.msri.org/communications/books/gt3m>. [1]
27. ———, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) **19** (1988), no. 2, 417–431. MR 956596 [24]

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, OXFORD, OX2 6GG, UK

SCHOOL OF MATHEMATICS, MONASH UNIVERSITY, VIC 3800, AUSTRALIA