Computing Triangulations of Mapping Tori of Surface Homeomorphisms

Peter Brinkmann* and Saul Schleimer

Abstract

We present the mathematical background of a software package that computes triangulations of mapping tori of surface homeomorphisms, suitable for Jeff Weeks's program `SnapPea`. The package is an extension of the software described in [**?**]. It consists of two programs. `jmt` computes triangulations and prints them in a human-readable format. `jsnap` converts this format into `SnapPea`'s triangulation file format and may be of independent interest because it allows for quick and easy generation of input for `SnapPea`. As an application, we obtain a new solution to the restricted conjugacy problem in the mapping class group.

1 Introduction

In [**?**], the first author described a software package that provides an environment for computer experiments with automorphisms of surfaces with one puncture. The purpose of this paper is to present the mathematical background of an extension of this package that computes triangulations of mapping tori of such homeomorphisms, suitable for further analysis with Jeff Weeks's program `SnapPea` [**?**].[1]

*This research was partially conducted by the first author for the Clay Mathematics Institute.

2000 Mathematics Subject Classification. 57M27, 37E30.

Key words and phrases. Mapping tori of surface automorphisms, pseudo-Anosov automorphisms, mapping class group, conjugacy problem.

[1]Software available at `http://thames.northnet.org/weeks/index/SnapPea.html`

Pseudo-Anosov homeomorphisms are of particular interest because their mapping tori are hyperbolic 3-manifolds of finite volume [**?**]. The software described in [**?**] recognizes pseudo-Anosov homeomorphisms. Combining this with the programs discussed here, we obtain a powerful tool for generating and analyzing large numbers of hyperbolic 3-manifolds.

The software package described in [**?**] takes an automorphism ϕ of a surface S with one puncture (given as a sequence of Dehn twists) and computes the induced outer automorphism of the fundamental group of S, represented by a homotopy equivalence $f\colon G \to G$ of a finite graph $G \subset S$ homotopy equivalent to S, together with a loop σ in G homotopic to a loop around the puncture of S. The map f and the loop σ determine ϕ up to isotopy [**?**, Section 5.1].

In Section 2, we describe an effective algorithm for computing a triangulation of the mapping torus of $\phi\colon S \to S$, given only $f\colon G \to G$ and σ (Theorem 2.3). We also present an analysis of the complexity of this algorithm (Proposition 2.4). The first part of the software package is a program (called `jmt`) that implements this procedure. The program `jmt` prints its output in an intermediate human-readable format.

In Section 3, we explain how to use the software discussed here and the isometry checker of `SnapPea` to solve the restricted conjugacy problem in the mapping class group (i.e., the question of whether two pseudo-Anosov homeomorphisms are conjugate in the mapping class group). This problem was previously solved in [**?**] and [**?**]. One distinguishing feature of our solution is that much of it has already been implemented.

Appendix A discusses the second program in the software package (called `jsnap`), which converts the intermediate format of `jmt` into `SnapPea`'s triangulation file format. Since `SnapPea`'s format is rather complicated, it is not easy to generate input files for `SnapPea`, and `jsnap` may be of independent interest because it allows users to generate input for `SnapPea` without having to understand `SnapPea`'s file format.

Finally, in Appendix B, we present some sample computations that exhibit some of the capabilities of the combination of `SnapPea` and the software discussed here.

Immediate applications of the software described here include an experimental investigation of possible relationships between dynamical properties of pseudo-Anosov homeomorphisms (as computed by the first author's train track software) and topological properties of their mapping tori (as computed by `SnapPea`). For example, one might look for a relationship between growth

rate and volume. Another area where the package described in this paper has already been used is the study of slalom knots as introduced by Norbert A'Campo [**?**].

The software package is written in Java and should be universally portable. The programs `jmt` and `jsnap` are command line software and can be used to examine a large number of examples as a batch job. A graphical user interface with an online help feature is also available.

The package, including binary files, source code, complete online documentation, and a user manual, is available at `http://www.math.uiuc.edu/~brinkman/`.

We would like to thank Mladen Bestvina and John Stallings for many helpful discussions, as well as Jeff Weeks and Bill Floyd for explaining `SnapPea`'s intricacies. We would also like to express our gratitude to Kai-Uwe Bux for critiquing an early version of this paper.

2 Computing triangulations

Let $\phi\colon S \to S$ be an automorphism of a surface S with one puncture, represented by a homotopy equivalence $f\colon G \to G$ of a finite graph G and a loop σ in G representing a loop around the puncture of S (see Section 1). There is no loss in assuming that $f\colon G \to G$ maps vertices to vertices and that the restriction of f to the interior of each edge of G is an immersion.

In this section, we outline an effective procedure that computes a triangulation of the mapping torus of ϕ given only f and σ. To this end, we construct a simplicial 2-complex K and a face pairing e with the following properties.

1. The space $|K|$ is homeomorphic to a torus.
2. For each 2-simplex Δ of K, there exists a 2-simplex Δ' of K and an orientation reversing simplicial homeomorphism $e_\Delta\colon \Delta \to \Delta'$ such that $e_{\Delta'}^{-1} = e_\Delta$.
3. The space K/e is homotopy equivalent to the mapping torus of f.
4. If we let $M = (\text{cone over } K)/e$ and obtain M' from M by removing the cone point, then M' is a 3-manifold (in particular, the links of vertices in M' are 2-spheres).

In this situation, M' is homotopy equivalent to the mapping torus of f, which in turn is homotopy equivalent to the mapping torus M_ϕ of ϕ. As M' is a 3-manifold, M' is homeomorphic to M_ϕ [**?**, page 6].

The triangulation of K induces a triangulation of M, i.e., the tetrahedra of M are cones over the triangles of K. The vertices of K give rise to finite vertices of M, and the cone point is an ideal vertex corresponding to the torus cusp of M_ϕ. By computing the links of vertices, `SnapPea` recognizes finite vertices (whose links are 2-spheres) and ideal vertices (whose links are tori or Klein bottles).

Hence, we have reduced to problem of constructing a triangulation of the mapping torus of ϕ to the construction of the 2-complex K and face pairing e, given only the homotopy equivalence $f : G \to G$ and the loop σ. The construction of K and e is the purpose of the remainder of this section.

The construction of K and e proceeds in two steps. We construct the 2-torus T by gluing annuli using Stallings's folding construction [**?**]. Then we construct a triangulation and a face pairing for each of the annuli.

2.1 Step 1: Subdividing and folding

We review the notion of subdividing and folding [**?**, **?**]. Let G, G' be finite graphs, and let $f\colon G' \to G$ be a map that maps vertices to vertices and edges to edge paths.

If f fails to be an immersion, then there exist two distinct edges $a,\ b$ in G' emanating from the same vertex such that $f(a)$ and $f(b)$ have a nontrivial initial path in common. We construct a new graph G'_1 by subdividing a (resp. b) into two edges $a_1,\ a_2$ (resp. $b_1,\ b_2$).

Now f factors through G'_1, i.e., there are maps $s\colon G' \to G'_1$ and $g\colon G'_1 \to G$ such that $f = g \circ s$. Moreover, we can choose s and g such that $s(a) = a_1 a_2$, $s(b) = b_1 b_2$ and $g(a_1) = g(b_1)$. We obtain a new graph G'_2 from G'_1 by identifying the edges a_1 and b_1. Then g factors through G'_2, i.e., there is a map $h\colon G'_2 \to G$ such that $g = h \circ p$, where p is the natural projection $p\colon G'_1 \to G'_2$ (see Figure 1). We refer to this process as *folding* a_1 and b_1.

Remark 2.1. The notion of folds used in [**?**] differs slightly from that introduced in [**?**]. In [**?**], the authors consider homotopy equivalences $f : G \to G$, and folding changes *both* the domain and the range of f, whereas in [**?**], the author considers maps $f : G' \to G$, and folding only affects the domain G'. The notion of folds used in this paper is a slight modification of the folds in

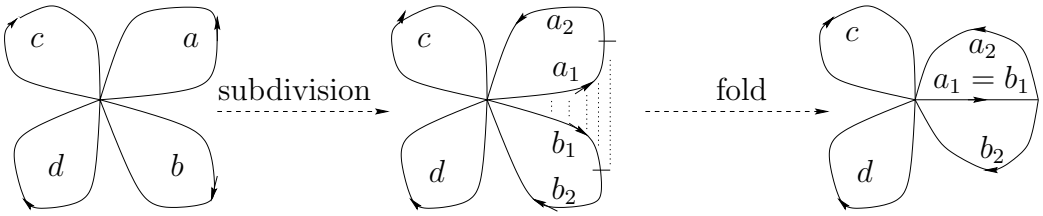

Figure 1: Subdividing and folding.

[**?**].

The remainder of this subsection details how to construct a sequence of graphs and maps

$$G = G_0 \xrightarrow{s_0} G_1 \xrightarrow{p_0} G_2 \ldots G_{2n-2} \xrightarrow{s_{n-1}} G_{2n-1} \xrightarrow{p_{n-1}} G_{2n} \xrightarrow{g_n} G$$

such that

$$f = g_n \circ p_{n-1} \circ s_{n-1} \circ \ldots \circ p_0 \circ s_0,$$

where $s_i \colon G_{2i} \to G_{2i+1}$ is a subdivision, $p_i \colon G_{2i+1} \to G_{2i+2}$ is a Stallings fold, and $g_n \colon G_{2n} \to G_0$ is an immersion. Since f is a homotopy equivalence, g_n will be onto, hence a homeomorphism. Moreover, for each $i = 0, \ldots, 2n$, we will construct a loop σ_i in G_i corresponding to a loop around the puncture of S.

Let $f \colon G \to G$ be induced by a homeomorphism $\phi \colon S \to S$, and let σ denote an edge loop in G corresponding to a loop around the puncture of S. Let $G = G_0$, $g_0 = f \colon G_0 \to G$, and $\sigma_0 = \sigma$.

Suppose that g_0 is not an immersion. Then there exist two edges $a,\ b$ emanating from the same vertex in G_0 such that $g_0(a)$ and $g_0(b)$ have a common initial segment. Since g_0 is induced by the homeomorphism $\phi \colon S \to S$, we can find a and b such that a and b are adjacent in the embedding of G_0 in S.

Since the loop σ_0 in G is homotopic to a loop around the puncture, a and b will be adjacent in the spelling of σ_0. Hence, we can detect a and b algorithmically by looking for cancellation between the images of adjacent edges in the spelling of σ_0.

We obtain G_1 from G_0 by subdividing a and b, and we obtain G_2 from G_1 by folding the initial segments of a and b. As above, we construct maps $s_0 \colon G_0 \to G_1$, $p_0 \colon G_1 \to G_2$, and $g_1 \colon G_2 \to G$ such that $g_0 = g_1 \circ p_0 \circ s_0$. Let $\sigma_1 = s_0(\sigma_0)$ and obtain σ_2 from $p_0(\sigma_1)$ by tightening. Since the edges a

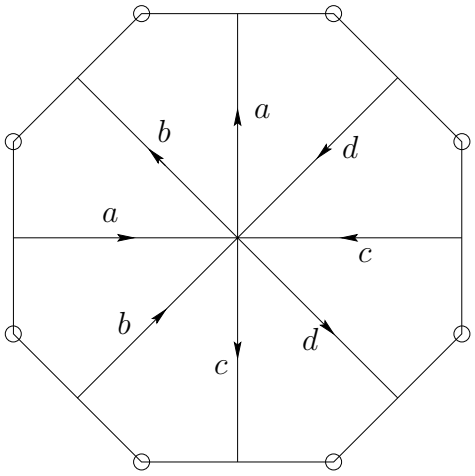

Figure 2: The graph G embedded in the surface S. The corners of the octagon correspond to the puncture of S. Two faces of the octagon are glued via an orientation-reversing map if the edge labels match up.

and b are adjacent in the embedding of G in S, the embedding of G_0 in S induces an embedding of G_1 and G_2 in S, and σ_1 and σ_2 are homotopic to σ_0 in S.

Note that the *size* of g_1, i.e., the sum of the lengths of the images under g_1 of the edges in G_2, is strictly smaller than the size of g_0. Hence, after repeating this construction finitely many times, we reach a map $g_n\colon G_{2n} \to G$ that cannot be folded and thus has to be an immersion. We have found the desired sequence of subdivisions and folds.

Example 2.2. Let G be the graph with one vertex and four edges, labeled $a, \ldots, d$, embedded in a punctured surface S of genus 2 as shown in Figure 2. We consider the following homotopy equivalence $f\colon G \to G$ induced by an automorphism of S.

$$\begin{aligned} f(a) &= a\underline{cd\bar{c}\bar{b}} \\ f(b) &= bc\bar{d}\bar{c}b\underline{cd\bar{c}\bar{b}} \\ f(c) &= bc\bar{d}\bar{d} \\ f(d) &= dd\bar{c}\bar{b}d \\ \sigma &= a\bar{b}\bar{a}bc\bar{d}\bar{c}d \end{aligned}$$

In $f(\sigma)$, cancellation occurs between the underlined parts of $f(a)$ and $f(\bar{b})$, and we subdivide a and b in preparation for folding, which gives us the maps $s_0\colon G \to G_1$ and $g_1\colon G_1 \to G$ (see Figure 1). In order to reduce notational complexity, we only change the labels of those edges that are subdivided.

$$\begin{aligned} s_0(a) &= a_1a_2 \\ s_0(b) &= b_1b_2 \\ s_0(c) &= c \\ s_0(d) &= d \\ \sigma_1 &= a_1a_2\bar{b}_2\bar{b}_1\bar{a}_2\bar{a}_1b_1b_2c\bar{d}\bar{c}d \end{aligned}$$

Now we fold the edges $\bar{a_2}$ and $\bar{b_2}$.

$$\begin{aligned} p_0(a_1) &= a_1 \\ p_0(a_2) &= b_2 \\ p_0(b_1) &= b_1 \\ p_0(b_2) &= b_2 \\ p_0(c) &= c \\ p_0(d) &= d \end{aligned}$$

Finally, we compute the map $g_1\colon G_2 \to G$.

$$\begin{aligned} g_1(a_1) &= a \\ g_1(b_1) &= bc\bar{d}\bar{c}b \\ g_1(b_2) &= cd\bar{c}\bar{b} \\ g_1(c) &= bc\bar{d}\bar{d} \\ g_1(d) &= dd\bar{c}\bar{b}d \\ \sigma_2 &= a_1\bar{b}_1\bar{b}_2\bar{a}_1b_1b_2c\bar{d}\bar{c}d \end{aligned}$$

2.2 Step 2: Triangulating annuli

We can interpret the loops σ_i as immersions $\sigma_i\colon S^1 \to G_i$. The preimage of the vertex set of G_i subdivides S^1 into intervals, and the restriction of σ_i

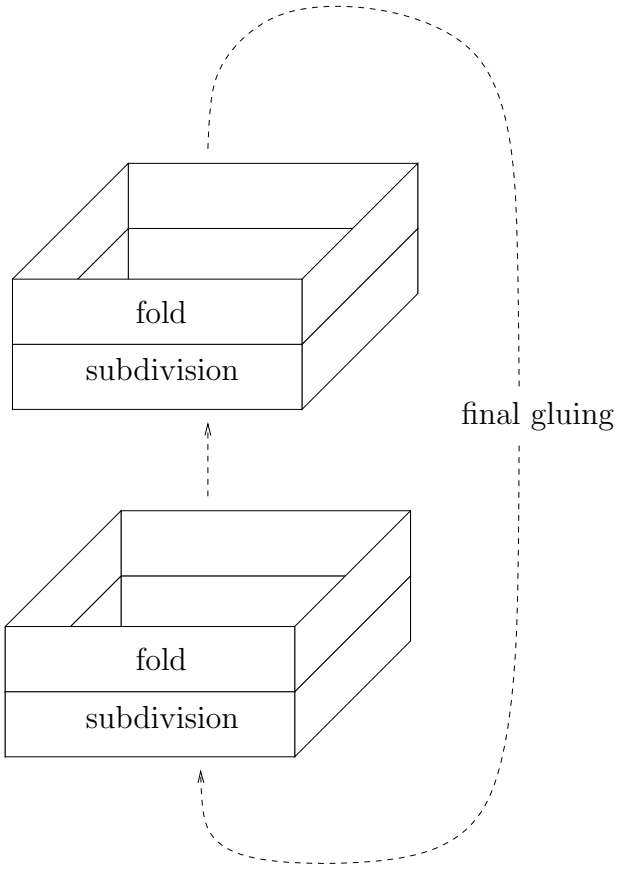

Figure 3: Decomposition into annuli.

to such an interval is a homeomorphism onto the interior of an edge in G_i. Hence, we can label each interval with the corresponding edge in G_i. We refer to this construction as *spelling* σ_i *along* S^1.

Now, for each $i \in \{0, \dots, 2n\}$, we take an annulus A_i and spell the word σ_i along one boundary component and σ_{i+1} along the other. We orient the two boundary components of A_i such that they are freely homotopic as oriented loops.

This labeling defines a gluing of A_i and A_{i+1}, and the homeomorphism $g_n \colon G_{2n} \to G = G_0$ induces a gluing of A_{2n-1} and A_0 (which we refer to as the *final gluing*), giving us the desired torus T (see Figure 3). Figure 4 shows the gluing of A_0 and A_1 for Example 2.2.

Fix some $i \in \{0, \dots, 2n\}$ and spell σ_i along S^1. Notice that each label on

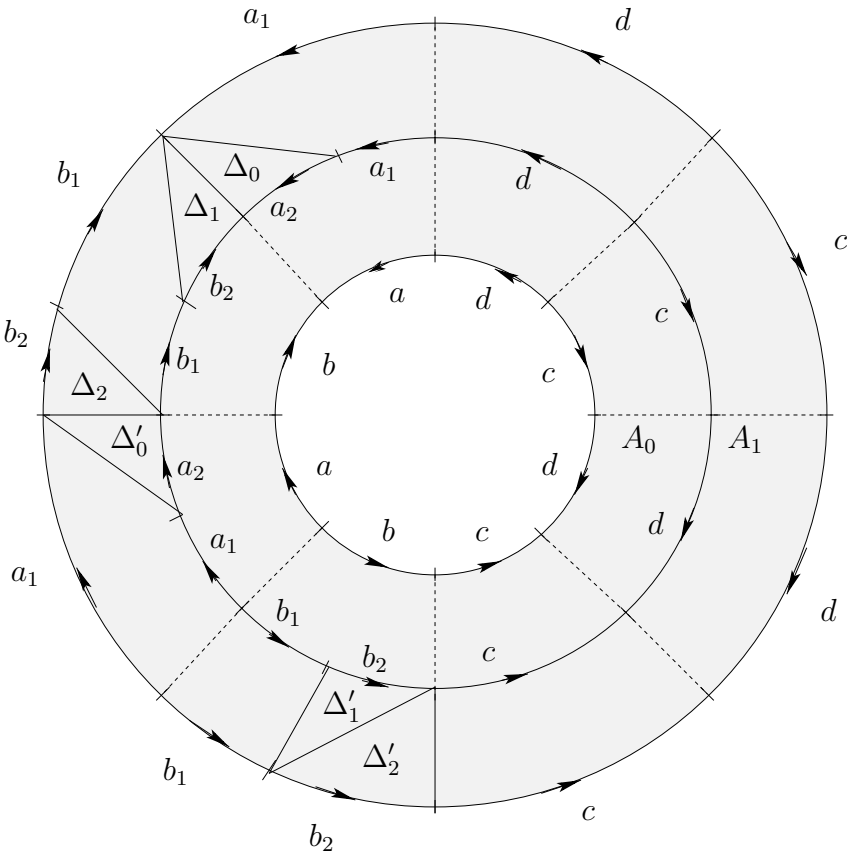

Figure 4: The gluing of A_0 and A_1 for Example 2.2.

S^1 occurs twice, once for each direction (see Example 2.2 and Figure 2). If we identify corresponding intervals, we obtain the graph G_i, and if we take the cone over S^1, remove the cone point, and identify intervals with identical labels, then we obtain the surface S [**?**, Section 5.1].

Hence, we only need to extend the edge pairing on the boundary of the annulus A_i to an appropriate face pairing of a triangulation of all of A_i. Recall that we want to choose the face pairing in such a way that the cone over T (with the cone point removed) becomes a 3-manifold when we identify corresponding triangles.

We first find a suitable triangulation of an annulus A_{2i} corresponding to a subdivision s_i. We decompose A_{2i} into rectangles corresponding to edges

that are not subdivided and pentagons corresponding to those edges that are subdivided (see Figure 4). The edge pairing on the boundary of A_{2i} induces a pairing of rectangles (resp. pentagons). Finding a triangulation of A_{2i} that is compatible with this pairing is straightforward.

We now construct a suitable triangulation of an annulus A_{2i+1} corresponding to a fold p_i. Edges that are not involved in the fold give rise to paired rectangles contained in A_{2i+1} (see Figure 4), and as above, we easily find a triangulation of these rectangles that is compatible with the pairing.

Let a, b denote the two edges involved in the fold, i.e., we have $p_i(a) = p_i(b) = b'$. By exchanging a and b or reversing the orientation of a and b as necessary, we may assume that the loop σ_{2i+1} has a subpath of the form $w = a\bar{b}u\bar{a}$ or $w = \bar{a}\bar{b}ub$, where u is a path that contains neither a nor b. For concreteness, we focus on the case $w = a\bar{b}u\bar{a}$. The first fold of Example 2.2 falls into this case, with $w = a_2\bar{b}_2\bar{b}_1\bar{a}_2$. The construction in the remaining case is similar.

The loop σ_{2i+2} has a corresponding subpath of the form $w' = u'\bar{b}'$. Let Δ_0 be the triangle spanned by the initial endpoint of w' and the occurrence of a in w. We pair Δ_0 with the triangle Δ_0' spanned by the terminal endpoint of w' and the occurrence of $\bar{a}$ in w (see Figure 4).

Let Δ_1 be the triangle spanned by the occurrence of $\bar{b}$ in w and the initial endpoint of w', and let Δ_2 be the triangle spanned by the occurrence of $\bar{b}'$ in w' and the terminal endpoint of u. Observe that after identifying paired edges, Δ_1 and Δ_2 have a side in common, so we can think of Δ_1 and Δ_2 as spanning a rectangle between $\bar{b}$ and $\bar{b}'$, which induces a triangulation of the rectangle spanned by the occurrence of b in σ_{2i+1} and the occurrence of $\bar{b}$ in σ_{2i+2} (see Figure 4). This completes the triangulation and face pairing of A_{2i+1}, which completes our construction.

Given the triangulation of A_{2i+1} constructed above, we can think of the annulus A_{2i+1} as interpolating between the graph G_{2i+1} and the graph G_{2i+2}. In other words, we think of the fold as occurring continuously, by identifying larger and larger segments of the edges a and b (see Figure 5). Moreover, this continuous folding process is compatible with the embedding of the graphs in the surface S. This observation shows that the complex K and its triangulation have the desired properties, in particular Property 4.

Summing up, we have obtained the following main result of this paper.

Theorem 2.3. *Let f be a homotopy equivalence of a finite graph G, representing a homeomorphism ϕ of a once-punctured surface. Then the following*

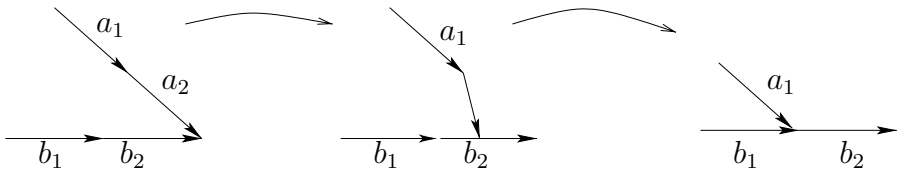

Figure 5: Continuous fold for Example 2.2.

is an effective procedure for computing a triangulation of the mapping torus of ϕ.

1. *Decompose the homotopy equivalence $f\colon G \to G$ into a sequence of subdivisions and folds, followed by a homeomorphism (Section 2.1).*

2. *Obtain the torus K as a gluing of one annulus for each subdivision and fold in the above decomposition (Figure 3).*

3. *Triangulate the individual annuli and construct a face pairing (Section 2.2).*

4. *Construct a triangulation of the mapping torus of the surface homeomorphism ϕ by taking the cone of K and glue tetrahedra according to the face pairing.*

□

The program `jmt` is an implementation of this procedure.

2.3 Complexity analysis

The purpose of this subsection is to obtain an estimate on the number of tetrahedra in the triangulation that we have constructed.

Recall that we defined the size of a map $h\colon G' \to G$ to be the sum of the lengths (in the usual path metric) of the images of the edges of G'. For example, the size of the first map in Example 2.2 is 23. Let $S(h)$ denote the size of h.

We say that a map $g\colon G' \to G$ is *tight* if

1. for every edge e of G', the restriction of g to the interior of e is an immersion, and

2. for every vertex v of G', there are two edges emanating from v that cannot be folded (not even after a subdivision).

Note that tightness can always be achieved by homotopy.

Proposition 2.4. *Let $f : G \to G$ be a tight homotopy equivalence representing a homeomorphism ϕ of a once-punctured surface of genus g, and assume that G has no vertices of valence less than three. Then the number of tetrahedra in the triangulation of M_ϕ is bounded by $16(5g-2)S(f)$.*

Proof. Since folding reduces size and annuli come in pairs (a subdivision annulus followed by a folding annulus), the number of annuli is bounded above by $2S(f)$. A simple application of Euler characteristics shows that G has no more than $6g-3$ edges because the valence of each vertex is at least three. Subdividing and folding, however, may create additional edges as well as vertices of valence one or two, so we need to understand the effect of subdividing and folding on the number of edges.

To this end, we introduce the notion of *partial folds*, i.e., folds where both participating edges have to be subdivided, and *full folds*, i.e., folds where at least one of the participating edges is not subdivided [**?**]. Clearly, a subdivision followed by a full fold does not increase the number of edges, whereas a subdivision followed by a partial fold increases the number of edges by one.

A partial fold reduces the number of possible folds by one because the map resulting from it is an immersion around the new vertex created by the fold, so the only folds that are possible after a partial fold are those that were available before.

Similarly, a full fold does not increase the number of possible folds. This means that the number of partial folds that occurs in our construction is bounded by the number of folds that the map $f\colon G \to G$ admits. Since f is tight, the number of folds at one vertex v is bounded by $\mathrm{val}(v)-2$. Summing up, we see that the number of possible folds is bounded by

$$\sum_{v \in G^{(0)}} (\mathrm{val}(v) - 2) = -2\chi(G) = 4g - 2.$$

Hence, the number of edges after a fold is bounded by $6g-3+4g-2 = 10g-5$. Subdivisions increase the number of edges, so the number of edges at any point in our construction is bounded by $10g-4$.

The number of tetrahedra belonging to one annulus is bounded above by $8(5g-2)$ (four tetrahedra per edge), which gives us a theoretical upper bound of $16(5g-2)S(f)$ on the number of tetrahedra in our triangulation of M_ϕ. □

Similar arguments show that for fixed genus, the time it takes to compute a triangulation is linear in the size of the input. We note that in practice, partial folds seldom occur, and triangulations tend to be considerably smaller than the bound given in Proposition 2.4.

3 Solving the conjugacy problem

The algorithm from Section 2 and `SnapPea`'s isometry checker provide a practical way of testing a necessary condition for two pseudo-Anosov homeomorphisms ϕ, γ to be conjugate in the mapping class group. Namely, we can compute the mapping tori of ϕ and γ as in Section 2, and then `SnapPea`'s isometry checker will determine whether the two mapping tori are isometric. If they are not isometric, we can immediately conclude that the ϕ and γ are not conjugate. However, if the mapping tori are isometric, we cannot yet conclude that they are conjugate. The purpose of this section is to provide an effective sufficient criterion for conjugacy.

Problem 3.1. Let γ and ϕ be automorphisms of a surface S with one puncture. Both automorphisms are assumed to be presented as products of Dehn twists on S.

The *conjugacy problem* asks for a decision procedure to determine whether or not γ and ϕ are conjugate in the mapping class group of S. That is, does there exist an automorphism, φ, of S such that $\gamma = \varphi^{-1}\phi\varphi$?

The *restricted conjugacy problem* is slightly easier in that it assumes that ϕ is in fact pseudo-Anosov. As this is not a large restriction from now on we will assume that ϕ is pseudo-Anosov.

A complete (albeit impractical) solution for the conjugacy problem has been given by Hemion [**?**]. The restricted case has also been solved by Mosher [**?**].

3.1 Notation

In order to describe our solution of the restricted conjugacy problem, we need to introduce some notation. Let $\gamma\colon\ S \to S$ and $\phi\colon\ S \to S$ be pseudo-Anosov homeomorphisms with isometric mapping tori. `SnapPea` will detect this and compute an isometry $h\colon\ M_\gamma \to M_\phi$.

Let $i_\gamma\colon\ S \to M_\gamma$ and i_ϕ be the two inclusion maps that realize S as a fiber of the induced fiber structures $\mathcal{F}_\gamma$ and $\mathcal{F}_\phi$ on M_γ and M_ϕ. Set $F_\gamma = i_\gamma(S)$ and define F_ϕ in a similar fashion.

Let $p_\gamma\colon\ M_\gamma \to S^1$ be the map from M_γ to the circle induced by $\mathcal{F}_\gamma$ and define p_ϕ similarly.

Let $\sigma \in \mathrm{Mod}(M_\phi)$ denote a typical element of the isometry group of M_ϕ. As M_ϕ is hyperbolic, $\mathrm{Mod}(M_\phi)$ is a finite group and can be computed by `SnapPea`. Set $G_\sigma = \sigma h(F_\gamma)$.

Finally, pick any $g \in \pi_1(M_\gamma)$ with the property that $(p_\gamma)_* g = 1 \in \mathbb{Z}$. We say that such a loop *represents the S^1-orientation.* If $(p_\phi \sigma h)_* g$ equals 1 (-1) then we say that *σh preserves (reverses) S^1-orientation.*

3.2 Retriangulation and the fundamental group

Before solving the restricted conjugacy problem we will need a pair of subroutines which determine the images of elements of $\pi_1(M_\gamma)$ under the map $(p_\phi \sigma h)_*$.

First, we need an algorithm which decides whether $(p_\phi)_*\colon\ \pi_1(G_\sigma) \to \mathbb{Z}$ has trivial image. The idea here is to keep track of a set of generators for $\pi_1(S)$ under the maps i_γ, h, σ and p_ϕ. Unfortunately, these homeomorphisms do not all respect common simplicial structures on M_γ and M_ϕ. To fix this problem one must find an appropriate set of generators in each retriangulation of M_γ. Then, once `SnapPea` finds a geometric triangulation of M_γ, we can push the generators onto the one-skeleton of the Ford domain. The isometry σh takes them to edge paths in the one-skeleton of M_ϕ where we can reverse the process. Finally, p_ϕ projects the generators of $\pi_1(G_\sigma)$ to the circle where it is easy to check whether or not they are all contractible.

Second, we will need an algorithm which decides whether σh preserves S^1-orientation. To do this, construct any loop $g \in \pi_1(M_\gamma)$ which represents the S^1-orientation. As in the previous algorithm take the image of g under the map $(p_\phi \sigma h)_*$. Check that this image is the positive generator of $\pi_1(S^1)$.

The bookkeeping problem of keeping track of surface subgroups of $\pi_1(M^3)$

under retriangulation does not yet have an implemented solution. We would be very interested in the work of any reader who is willing to write such a program. It should be remarked that the subroutines above do a little more work than is strictly necessary. It would suffice to keep track of a two-chain representing the fiber of M_γ. Again, this is a straightforward problem which does not yet have an implemented solution.

3.3 The algorithm

Begin with two homeomorphisms $\gamma, \phi\colon S \to S$. If one of them fails to be pseudo-Anosov, then the software described in [**?**] will detect this. Otherwise, construct the mapping tori of γ and ϕ, M_γ and M_ϕ. If `SnapPea` reports that M_γ and M_ϕ are not isometric, we conclude that γ cannot be conjugate to ϕ. This resolves the issue for a vast majority of possible pairs of γ and ϕ.

Now, suppose `SnapPea` reports the two mapping tori are isometric. We cannot yet conclude that the two automorphisms are conjugate. It may be that γ and ϕ have homeomorphic mapping tori but are not conjugate because they give rise to distinct fiber structures in the resulting three manifold. Also, it may happen that γ and ϕ induce identical fiber structures while reversing S^1-orientation. In this case we show that γ is conjugate to ϕ^{-1}.

If $\pi_1(G_\sigma)$ has nontrivial image under $(p_\phi)_*$ then clearly G_σ is not isotopic to F_ϕ. We conclude that if $\pi_1(G_\sigma)$ has nontrivial image in $\pi_1(S^1) = \mathbb{Z}$ for every σ in $\mathrm{Mod}(M_\phi)$ then γ is not conjugate to ϕ.

We claim that if there exists some σ such that σh preserves S^1-orientation and $\pi_1(G_\sigma)$ is contained in the kernel of the natural projection, then γ is conjugate to ϕ, which completes our solution of the restricted conjugacy problem. The rest of this section is devoted to a proof of this claim.

3.4 Isotoping G_σ

Assume now that $\pi_1(G_\sigma)$ has trivial image in $(p_\phi)_*$ for some fixed $\sigma \in \mathrm{Mod}(M_\phi)$. At this point we need a weak form of Theorem 4 from [**?**]:

Theorem 3.2. *G_σ is isotopic to a properly embedded surface that is either a leaf of $\mathcal{F}_\phi$, or has only saddle singularities for the induced singular foliation of G_σ. The boundary component of the isotoped G_σ is either a leaf of $\mathcal{F}_\phi|\partial M_\phi$ or is transverse to $\mathcal{F}_\phi|\partial M_\phi$.* □

We use this as follows:

Corollary 3.3. *If $\pi_1(G_\sigma)$ is in the kernel of $(p_\phi)_*$ then G_σ is isotopic to F_ϕ.*

Proof. Suppose, to obtain a contradiction, that G_σ is not isotopic to F_ϕ. By Thurston's theorem we may isotope G_σ so that the induced foliation has only saddle singularities.

Let $M_\mathbb{Z} = F_\phi \times \mathbb{R}$ be the infinite cyclic cover of M coming from $\mathcal{F}_\phi$. By assumption we may lift G_σ to $M_\mathbb{Z}$. Note that projection onto the second factor $M_\mathbb{Z} \to \mathbb{R}$ gives a Morse function when restricted to G_σ. However this Morse function must have a maximum. The induced foliation of G_σ has a singularity at this maximum, but this singularity cannot possibly be a saddle singularity. This is a contradiction.

□

Thus we may isotope σh so as to obtain $G_\sigma = F_\phi$. Cutting along F_γ and F_ϕ we obtain a map $h' \colon S \times I \to S \times I$ that takes $S \times 0$ and $S \times 1$ to $S \times 0$ and $S \times 1$, but not necessarily in that order. It may be that σh reverses the S^1-orientation.

It follows that h', and hence σh, is isotopic to a map that preserves fibers. (See, for example, Lemma 3.5 of [**?**].) Letting $h_0 = (\sigma h)|F_\gamma$ we find that either $\gamma = h_0^{-1}\phi h_0$ or $\gamma = h_0^{-1}\phi^{-1}h_0$ depending on whether σh preserves or reverses S^1-orientation. This completes the proof of the claim and thus shows the correctness of our algorithm, which we sum up in the following theorem.

Theorem 3.4. *Let γ and ϕ be pseudo-Anosov homeomorphisms of a once punctured surface, S. The following is a procedure to decide whether the two mappings are conjugate in the mapping class group of S, if* `SnapPea` *is allowed as a subroutine.*

1. *Apply Theorem 2.3 to obtain the mapping tori, M_γ and M_ϕ.*
2. *Using* `jsnap` *and* `SnapPea`*, determine whether or not M_γ and M_ϕ are isometric. If not, then γ is not conjugate to ϕ.*
3. *Using* `SnapPea`*, enumerate all isometries between M_γ and M_ϕ.*
4. *Determine whether any of these are S^1-orientation and fibre preserving. If none are, then γ is not conjugate to ϕ. Otherwise, γ is conjugate to ϕ.*

□

The complexity of this algorithm depends on the complexity of `SnapPea`, which we treat as a black box in this paper.

A Generating input for `SnapPea`

In order to generate input for `SnapPea`, the output of `jmt` has to be translated into `SnapPea`'s triangulation file format by a second program (called `jsnap`). The purpose of this section is to discuss the intermediate format, which allows for quick and easy generation of input for `SnapPea`. We illustrate this format with a familiar example from [**?**].

The intermediate format admits two types of input lines, for tetrahedra and for gluings. An input line defining a tetrahedron has the form "`T` v_1 v_2 v_3 v_4," where $v_1, \ldots, v_4$ are distinct labels of the vertices. Tetrahedra are glued implicitly if they have three vertices in common, and they can be glued explicitly by entering a line of the form "`G` v_1 v_2 v_3 w_1 w_2 w_3," where v_1, v_2, v_3 and w_1, w_2, w_3 are the labels of two faces of tetrahedra. In this gluing, the side $[v_1, v_2]$ is glued to the side $[w_1, w_2]$, the side $[v_2, v_3]$ is glued to the side $[w_2, w_3]$, etc. Empty lines and comments beginning with '//' are also allowed.

Although the four vertex labels of a tetrahedron have to be distinct, two or more vertices of a tetrahedron may be identified after gluing. Distinct vertex labels are only needed in order to uniquely specify the sides of tetrahedra.

Example A.1 (Figure-eight knot). Consider the familiar figure-eight knot (see Figure 6). An ideal triangulation of the complement of this knot can be expressed as a gluing of two tetrahedra [**?**, Chapter 1] (see Figure 7).

We encode the gluing as follows.

```
T a b c d
T b c d e
G b e d a c d
G c b e a b d
G c e d a c b
```

Feeding the above five lines into `jsnap` yields an encoding of the triangulation in `SnapPea`'s file format. According to `SnapPea`, the fundamental group of the resulting manifold M has the following presentation, which we

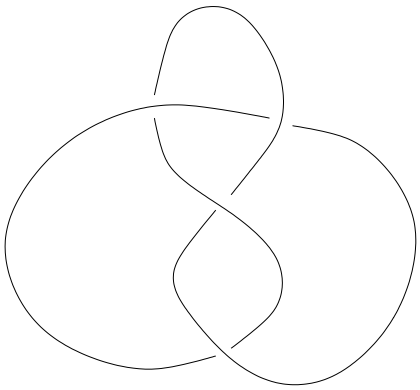

Figure 6: The figure-eight knot.

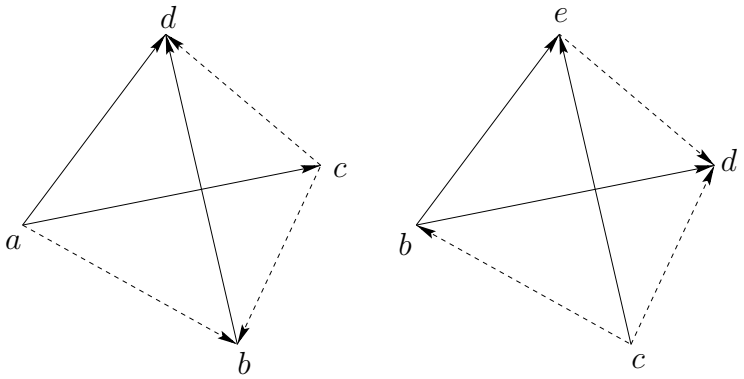

Figure 7: The figure-eight knot complement, expressed as a gluing of two ideal tetrahedra. Triangles are glued such that the line style and the direction of arrows are matched [**?**, Chapter 1].

modify by a sequence of Tietze transformations [**?**, Chapter II.2].

$$\begin{aligned}
\pi_1(M) &\cong < x,\ y \mid \bar{y}\bar{x}\bar{x}\bar{x}\bar{y}xyyx = 1 > \qquad (1)\\
&\cong < x,\ y\ ,\ z \mid \bar{y}\bar{x}\bar{x}\bar{x}\bar{y}xyyx = 1,\ y = z\bar{x} >\\
&\cong < x,\ z \mid x\bar{z}\bar{x}\bar{x}\bar{z}xz\bar{x}z = 1 >\\
&\cong < x,\ z \mid \bar{x}zx\bar{z}\bar{x}\bar{x}\bar{z}xz = 1 > \qquad (2)
\end{aligned}$$

Presentation (2) is the presentation of the fundamental group of the complement of the figure-eight knot given in [**?**, Example 3.8].

While `SnapPea`'s format is an extremely efficient representation of triangulations, understanding it requires some effort. The program `jsnap` acts as

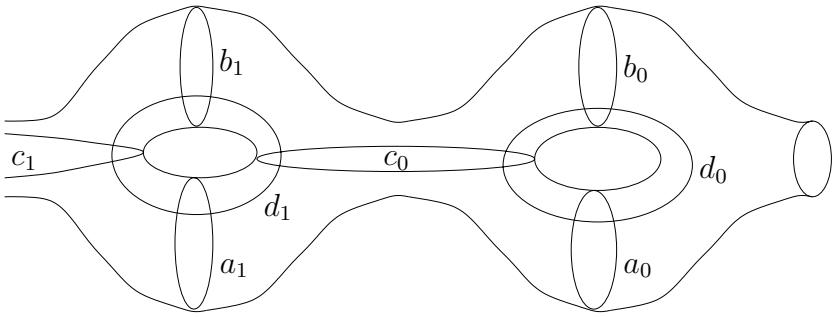

Figure 8: A set of generators of the mapping class group.

an interface between `SnapPea` and the human user. If you can draw or visualize a triangulation of a 3-manifold, then you can also enter it into `jsnap`.

B Sample Computations

We present some sample computations that illustrate the power of the software discussed here. We define surface automorphisms as compositions of Dehn twists with respect to the set of curves shown in Figure 8. We equip the surface with a normal vector field. When twisting with respect to a curve c, we turn right whenever we hit c (resp. left for inverse twists). The set of these Dehn twists generates the mapping class group [**?**].

We first present an example for which we can easily verify correctness.

Example B.1 (Figure-eight knot revisited). Let S be a punctured torus, and let $\phi\colon\ S \to S$ be given by

$$\phi = D_{a_0}^{-1} D_{d_0}.$$

The complement of the figure-eight knot (see Figure 6) is homeomorphic to the mapping torus of ϕ (we will verify this soon). Given this composition of Dehn twists, the train track software described in [**?**] determines that ϕ is pseudo-Anosov with growth rate $\lambda \approx 2.61803399$, and `SnapPea` determines that the mapping torus M_ϕ is a hyperbolic 3-manifold of volume $V \approx 2.02988321$ with one torus cusp.

Moreover, the train track software computes the following topological representative $f\colon\ G \to G$ of ϕ, where G is a graph with one vertex and two

edges.

$$\begin{aligned} f(a) &= ba \\ f(b) &= bba \\ \sigma &= a\bar{b}\bar{a}b \end{aligned}$$

This yields the following presentation of the fundamental group of the mapping torus M_ϕ of ϕ.

$$\pi_1(M_\phi) \cong < a,\ b,\ t \mid \bar{t}at = ba,\ \bar{t}bt = bba >$$

A few Tietze transformations show that $\pi_1(M_\phi)$ is the fundamental group of the figure-eight knot complement.

$$\begin{aligned} \pi_1(M_\phi) &\cong < a,\ b,\ t \mid \bar{t}at = ba,\ \bar{t}bt = bba > \\ &\cong < a,\ t \mid t\bar{a}\bar{a}ta\bar{t}\bar{a}\bar{t}a = 1 > \\ &\cong < a,\ c \mid \bar{a}ca\bar{c}\bar{a}\bar{a}\bar{c}ac = 1 > \end{aligned}$$

As in Example A.1, we have obtained the presentation given in [**?**, Example 3.8]. Finally, the presentation of $\pi_1(M_\phi)$ computed by `SnapPea` is

$$\pi_1(M_\phi) \cong < x,\ y \mid \bar{x}\bar{y}\bar{y}\bar{y}\bar{x}yxxy = 1 >,$$

which agrees with Presentation (1) in Example A.1.

Alternatively, we can run `SnapPea`'s isometry checker on Example A.1 and Example B.1 in order to see that we get the same hyperbolic 3-manifold in both examples.

Example B.2 (Genus 3). Let S be a surface of genus 3 with one puncture, and let $\phi\colon\ S \to S$ be given by

$$\phi = D_{d_0} D_{c_0} D_{d_1} D_{c_1} D_{d_2} D_{a_2}^{-1}.$$

The train track software identifies ϕ as a pseudo-Anosov homeomorphism with growth rate $\lambda \approx 2.04249053$, and `SnapPea` determines that the mapping torus M_ϕ is a hyperbolic 3-manifold of volume $V \approx 4.93524268$ with one torus cusp.

These applications only show a small part of all the possibilities. The train track software computes a plethora of information about surface homeomorphisms [**?**], and `SnapPea` allows for a detailed analysis of (hyperbolic) 3-manifolds. We believe that the combination of the two packages may become a valuable tool for topologists.

Peter Brinkmann
Department of Mathematics
273 Altgeld Hall
1409 W. Green Street
Urbana, IL 61801, USA
E-mail: brinkman@math.uiuc.edu

Saul Schleimer
Department of Mathematics
UC Berkeley
Berkeley, CA 94720, USA
E-mail: saul@math.berkeley.edu