

Small dilatation pseudo-Anosovs and 3-manifolds

Benson Farb, Christopher J. Leininger, and Dan Margalit *

May 6, 2009

Abstract

The main result of this paper is a universal finiteness theorem for the set of all small dilatation pseudo-Anosov homeomorphisms $\phi : S \rightarrow S$, ranging over all surfaces S . More precisely, we consider pseudo-Anosovs $\phi : S \rightarrow S$ with $|\chi(S)| \log(\lambda(\phi))$ bounded above by some constant, and we prove that, after puncturing the surfaces at the singular points of the stable foliations, the resulting set of mapping tori is finite. Said differently, there is a finite set of fibered hyperbolic 3-manifolds so that all small dilatation pseudo-Anosovs occur as the monodromy of a Dehn filling on one of the 3-manifolds in the finite list, where the filling is on the boundary slope of a fiber.

1 Introduction

Given a pseudo-Anosov homeomorphism ϕ of a surface S , let $\lambda(\phi)$ denote its dilatation. For any $P \geq 1$, we define

$$\Psi_P = \left\{ \phi : S \rightarrow S \mid \chi(S) < 0, \phi \text{ pseudo-Anosov, and } \lambda(\phi) \leq P^{\frac{1}{|\chi(S)|}} \right\}.$$

It follows from work of Penner [Pe] that for P sufficiently large, and for each closed surface S_g of genus $g \geq 2$, there exists $\phi_g : S_g \rightarrow S_g$ so that

$$\{\phi_g : S_g \rightarrow S_g\}_{g=2}^{\infty} \subset \Psi_P.$$

We refer to Ψ_P as the set of *small dilatation pseudo-Anosov homeomorphisms*.

Given a pseudo-Anosov homeomorphism $\phi : S \rightarrow S$, let $S_\phi^\circ = S^\circ \subset S$ be the surface obtained by removing the singularities of the stable and unstable foliations for ϕ and let $\phi|_{S^\circ} : S^\circ \rightarrow S^\circ$ denote the restriction. The set of pseudo-Anosov homeomorphisms

$$\Psi_P^\circ = \{\phi|_{S^\circ} : S^\circ \rightarrow S^\circ \mid (\phi : S \rightarrow S) \in \Psi_P\}$$

is therefore also infinite.

The main discovery contained in this paper is a universal finiteness phenomenon for all small dilatation pseudo-Anosov homeomorphisms: they are, in an explicit sense described below, all “generated” by a finite number of examples. To give a first statement, let $\mathcal{T}(\Psi_P^\circ)$ denote the homeomorphism classes of mapping tori of elements of Ψ_P° .

Theorem 1.1. *The set $\mathcal{T}(\Psi_P^\circ)$ is finite.*

We will begin by putting Theorem 1.1 in context, and by explaining some of its restatements and corollaries.

*The authors gratefully acknowledge support from the National Science Foundation.

1.1 Dynamics and geometry of pseudo-Anosov homeomorphisms

For a pseudo-Anosov homeomorphism $\phi : S \rightarrow S$, the number $\lambda(\phi)$ gives a quantitative measure of several different dynamical properties of ϕ . For example, given any fixed metric on the surface, $\lambda(\phi)$ gives the growth rate of lengths of a geodesic under iteration by ϕ [FLP, Exposé 9, Proposition 19]. The number $\log(\lambda(\phi))$ gives the minimal topological entropy of any homeomorphism in the isotopy class of ϕ [FLP, Exposé 10, §IV]. Moreover, a pseudo-Anosov homeomorphism is essentially the unique minimizer in its isotopy class—it is unique up to conjugacy by a homeomorphism isotopic to the identity [FLP, Exposé 12, Théorème III].

From a more global perspective, recall that the set of isotopy classes of orientation preserving homeomorphisms $\phi : S \rightarrow S$ forms a group called the mapping class group, denoted $\text{Mod}(S)$. This group acts properly discontinuously by isometries on the Teichmüller space $\text{Teich}(S)$ with quotient the moduli space $\mathcal{M}(S)$ of Riemann surfaces homeomorphic to S . The closed geodesics in the orbifold $\mathcal{M}(S)$ correspond precisely to the conjugacy classes of mapping classes represented by pseudo-Anosov homeomorphisms, and moreover, the length of a geodesic associated to a pseudo-Anosov $\phi : S \rightarrow S$ is $\log(\lambda(\phi))$. Thus, the length spectrum of $\mathcal{M}(S)$ is the set

$$\text{spec}(\text{Mod}(S)) = \{\log(\lambda(\phi)) : \phi : S \rightarrow S \text{ is pseudo-Anosov}\} \subset (0, \infty).$$

Arnoux–Yoccoz [AY] and Ivanov [Iv] proved that $\text{spec}(\text{Mod}(S))$ is a closed discrete subset of \mathbb{R} . It follows that $\text{spec}(\text{Mod}(S))$ has, for each S , a least element, which we shall denote by $L(S)$. We can think of $L(S)$ as the *systole* of $\mathcal{M}(S)$.

1.2 Small dilatations

Penner proved that there exists constants $0 < c_0 < c_1$ so that for all closed surfaces S with $\chi(S) < 0$, one has

$$c_0 \leq L(S)|\chi(S)| \leq c_1. \tag{1}$$

The proof of the lower bound comes from a spectral estimate for Perron–Frobenius matrices, with $c_0 > \log(2)/6$ (see [Pe] and [Mc2]). As such, this lower bound is valid for all surfaces S with $\chi(S) < 0$, including punctured surfaces. The upper bound is proven by constructing pseudo-Anosov homeomorphisms $\phi_g : S_g \rightarrow S_g$ on each closed surface of genus $g \geq 2$ so that $\lambda(\phi_g) \leq e^{c_1/(2g-2)}$; see also [Ba].

The best known upper bound for $\{L(S_g)|\chi(S_g)|\}$ is due to Hironaka–Kin [HK] and independently Minakawa [Mk], and is $2 \log(2 + \sqrt{3})$. The situation for punctured surfaces is more mysterious. There is a constant c_1 so that the upper bound of (1) holds for punctured spheres and punctured tori; see [HK, Ve, Ts]. However, Tsai has shown that for a surface $S_{g,p}$ of fixed genus $g \geq 2$ and variable number of punctures p , there is no upper bound c_1 for $L(S_{g,p})|\chi(S_{g,p})|$, and in fact, this number grows like $\log(p)$ as p tends to infinity; see [Ts].

One construction for small dilatation pseudo-Anosov homeomorphisms is due to McMullen [Mc2]. The construction, described in the next section, uses 3-manifolds and is the motivation for our results.

1.3 3-manifolds

Given $\phi : S \rightarrow S$, the mapping torus $M = M_\phi$ is the total space of a fiber bundle $f : M \rightarrow S^1$, so that the fiber is the surface S . If $H^1(M; \mathbb{R})$ has dimension at least 2, then one can perturb the fibration to another fibration $f' : M \rightarrow S^1$ whose dual cohomology class is projectively close to the dual of $f : M \rightarrow S^1$; see [Ti]. In fact, work of Thurston [Th2] implies that there is an open cone in $H^1(M; \mathbb{R})$ with the property that every integral class in this cone is dual to a fiber in a fibration. Moreover, the absolute value of the euler characteristic of these fibers extends to a linear function on this cone which we denote $|\chi(\cdot)|$.

Fried [Fr] proved that the dilatation of the monodromy extends to a continuous function $\lambda(\cdot)$ on this cone, such that the reciprocal of its logarithm is homogeneous. Therefore, the product $|\chi(\cdot)| \log(\lambda(\cdot))$ depends only on the projective class and varies continuously. McMullen [Mc2] refined the analysis of the extension $\lambda(\cdot)$, producing a polynomial invariant whose roots determine $\lambda(\cdot)$, and which provides a much richer structure.

In the course of his analysis, McMullen observed that for a fixed 3-manifold $M = M_\phi$, if $f_k : M \rightarrow S^1$ is a sequence of distinct fibrations of M with fiber $S(k)$ and monodromy $\phi_k : S(k) \rightarrow S(k)$, for which the dual cohomology classes converge projectively to the dual of the original fibration $f : M \rightarrow S^1$, then $|\chi(S(k))| \rightarrow \infty$, and

$$|\chi(S(k))| \log(\lambda(\phi_k)) \rightarrow |\chi(S)| \log(\lambda(\phi)).$$

In particular, $|\chi(S(k))| \log(\lambda(\phi_k))$ is uniformly bounded, independently of k , and therefore $\{\phi_k : S(k) \rightarrow S(k)\} \subset \Psi_P$ for some P .

There is a mild generalization of this construction obtained by replacing $\phi : S \rightarrow S$ with $\phi|_{S^\circ} : S^\circ \rightarrow S^\circ$ (as defined above) and considering the resulting mapping torus $M^\circ = M_{\phi|_{S^\circ}}$. McMullen's construction produces pseudo-Anosov homeomorphisms of punctured surfaces from the monodromies of various fibrations of $M^\circ \rightarrow S^1$. By filling in any invariant set of punctures (for example, all of them), we obtain homeomorphisms that may or may not be pseudo-Anosov (after filling in the punctures, there may be a 1-pronged singularity in the stable and unstable foliation, and this is not allowed for a pseudo-Anosov). In any case, we can say that all of the pseudo-Anosov homeomorphisms obtained in this way are *generated* by $\phi : S \rightarrow S$. Theorem 1.1 can be restated as follows.

Corollary 1.2. *Given $P > 1$, there exists a finite set of pseudo-Anosov homeomorphisms that generate all pseudo-Anosov homeomorphisms $\phi \in \Psi_P$ in the sense above.*

Corollary 1.2 can be viewed as a kind of converse of McMullen's construction. In view of Corollary 1.2, we have the following natural question.

Question 1.3. *For any given P , can one explicitly find a finite set of pseudo-Anosov homeomorphisms that generate Ψ_P ?*

1.4 Dehn filling

Puncturing the surface at the singularities is necessary for the validity of Theorem 1.1. Indeed, the set of all pseudo-Anosov homeomorphisms that occur as the monodromy for a fibration of a fixed 3-manifold have a uniform upper bound for the number of prongs at

any singularity of the stable foliation. On the other hand, Penner's original construction produces a sequence of pseudo-Anosov homeomorphisms $\phi_g : S_g \rightarrow S_g$ in which the number of prongs at a singularity tends to infinity with g . So, the set $\mathcal{T}(\Psi_P)$ of homeomorphism classes of mapping tori of elements of Ψ_P is an infinite set for P sufficiently large.

Removing the singularities from the stable and unstable foliations of ϕ , then taking the mapping torus, is the same as *drilling out* the closed trajectories through the singular points of the suspension flow in M_ϕ . Thus, we can reconstruct $M = M_\phi$ from $M^\circ = M_{\phi|_{S^\circ}}$ by *Dehn filling*; see [Th1]. The Dehn-filled solid tori in M are regular neighborhoods of the closed trajectories through the singular points and the *Dehn filling slopes* are the (minimal transverse) intersections of S with the boundaries of these neighborhoods. Back in M° , the boundaries of the neighborhoods are tori that bound product neighborhoods of the ends of M° homeomorphic to $[0, \infty) \times T^2$. Here the filling slopes are described as the intersections of S° with these tori and are called the *boundary slopes of the fiber S°* .

Moreover, since all of the manifolds in $\mathcal{T}(\Psi_P^\circ)$ are mapping tori for pseudo-Anosov homeomorphisms, they all admit a complete finite volume hyperbolic metric by Thurston's Geometrization Theorem for Fibered 3-manifolds (see, e.g. [Mc1], [Ot]). As another corollary of Theorem 1.1, we obtain the following.

Corollary 1.4. *For each $P > 1$ there exists $r \geq 1$ with the following property. There are finitely many complete, noncompact, hyperbolic 3-manifolds M_1, \dots, M_r fibering over S^1 , with the property that any $\phi \in \Psi_P$ occurs as the monodromy of some bundle obtained by Dehn filling one of the M_i along boundary slopes of a fiber.*

1.5 Volumes

Because hyperbolic volume decreases after Dehn filling [NZ, Th1], as a corollary of Corollary 1.4, we have the following.

Corollary 1.5. *The set of hyperbolic volumes of 3-manifolds in $\mathcal{T}(\Psi_P)$ is bounded by a constant depending only on P .*

It follows that we can define a function $\mathcal{V} : \mathbb{R} \rightarrow \mathbb{R}$ by the formula:

$$\mathcal{V}(\log P) = \sup\{\text{vol}(M_\phi) \mid \phi \in \mathcal{T}(\Psi_P)\}.$$

Given a pseudo-Anosov $\phi : S \rightarrow S$, we have $\phi \in \Psi_P$ for $\log P = |\chi(S)| \log(\lambda(\phi))$, and so as a consequence of Corollary 1.5 we have the following.

Corollary 1.6. *For any pseudo-Anosov $\phi : S \rightarrow S$, we have*

$$\text{vol}(M_\phi) \leq \mathcal{V}(|\chi(S)| \log(\lambda(\phi))).$$

Using the relationship between hyperbolic volume and simplicial volume [Th2], a careful analysis of our proof of Theorem 1.1 shows that \mathcal{V} is bounded above by an exponential function.

For a homeomorphism ϕ of S , let $\tau_{WP}(\phi)$ denote the translation length of ϕ , thought of as an isometry of $\text{Teich}(S)$ with the Weil–Petersson metric. Brock [Br] has proven that

the volume of M_ϕ and $\tau_{WP}(\phi)$ satisfy a bilipschitz relation (see also Agol [Ag]), and in particular

$$\text{vol}(M_\phi) \leq c\tau_{WP}(\phi).$$

Moreover, there is a relation between the Weil–Peterson translation length and the Teichmüller translation length $\tau_{\text{Teich}}(\phi) = \log(\lambda(\phi))$ (see [Li]), which implies

$$\tau_{WP}(\phi) \leq \sqrt{2\pi|\chi(S)|} \log(\lambda(\phi)).$$

However, Brock’s constant $c = c(S)$ depends on the surface S , and moreover $c(S) \geq |\chi(S)|$ when $|\chi(S)|$ is sufficiently large. In particular, Corollary 1.5 does not follow from these estimates. See also [KKT] for a discussion relating volume to dilatation for a fixed surface.

1.6 Minimizers

Very little is known about the actual values of $L(S)$. It is easy to prove $L(S_1) = \log(\frac{3+\sqrt{5}}{2})$. The number $L(S)$ is also known when $|\chi(S)|$ is relatively small; see [CH, HS, HK, So, SKL, Zh, LT]. However, the exact value is not known for any surface of genus greater than 2; see [LT] for some partial results for surfaces of genus less than 9. In spite of the fact that elements in the set $\{L(S)\}$ seem very difficult to determine, the following corollary shows that infinitely many of these numbers are generated by a single example.

Corollary 1.7. *There exists a complete, noncompact, finite volume, hyperbolic 3–manifold M with the following property: there exist Dehn fillings of M giving an infinite sequence of fiberings over S^1 , with closed fibers S_{g_i} of genus $g_i \geq 2$ with $g_i \rightarrow \infty$, and with monodromy ϕ_i , so that*

$$L(S_{g_i}) = \log(\lambda(\phi_i)).$$

We do not know of an explicit example of a hyperbolic 3–manifold as in Corollary 1.7; the work of [KKT, KT, Ve] gives one candidate.

We now ask if the Dehn filling is necessary in the statement of Corollary 1.7:

Question 1.8. *Does there exist a single 3–manifold that contains infinitely many minimizers? That is, does there exist a hyperbolic 3–manifold M that fibers over the circle in infinitely many different ways $\phi_k : M \rightarrow S^1$, so that the monodromies $f_k : S_{g_k, p_k} \rightarrow S_{g_k, p_k}$ satisfy $\log(\lambda(\phi_k)) = L(S_{g_k, p_k})$?*

1.7 Outline of the proof

To explain the motivation for the proof, we again consider McMullen’s construction. In a fixed fibered 3–manifold $M \rightarrow S^1$, Oertel proved that all of the fibers $S(k)$ for which the dual cohomology classes lie in the open cone described above are carried by a finite number of *branched surfaces* transverse to the suspension flow. A branched surface is a 2–dimensional analogue in a 3–manifold of a train track on a surface; see [Oe].

Basically, an infinite sequence of fibers in M whose dual cohomology classes are projectively converging are building up larger and larger “product regions” because they all live in a single 3–manifold. These product regions can be collapsed down, and the images of the fibers under this collapse define a finite number of branched surfaces.

Our proof follows this idea by trying to find “large product regions” in the mapping torus that can be collapsed, so that the fiber projects onto a branched surface of uniformly bounded complexity. In our case, we do not know that we are in a fixed 3–manifold (indeed, we may not be), so the product regions must come from a different source than in the single 3–manifold setting; this is where the small dilatation assumption is used. Moreover, the collapsing construction we describe does not in general produce a branched surface in a 3–manifold. However, the failure occurs only along the closed trajectories through the singular points, and after removing the singularities, the result of collapsing is indeed a 3–manifold.

While this is only a heuristic (for example, the reader never actually needs to know what a branched surface is), it is helpful to keep in mind while reading the proof, which we now outline. First we replace all punctures by marked points, so S is a closed surface with marked points. Let $M = M_\phi$ be the mapping torus, which is now a compact 3–manifold. We associate to any small dilatation pseudo-Anosov homeomorphism $\phi : S \rightarrow S$ a Markov partition \mathcal{R} with a “small” number of rectangles [Section 4].

Step 1: We prove that all but a uniformly bounded number of the rectangles of \mathcal{R} are essentially permuted [Lemma 4.6]. This follows from the relationship with Perron–Frobenius matrices and their adjacency graphs together with an application of a result of Ham and Song [HS] [Lemma 3.1]. Moreover, we prove that any rectangle meets a uniformly bounded number of other rectangles along its sides [Lemma 5.1]. These rectangles, and their images, are used to construct a cell structure Y on S [Section 5].

Step 2: When a rectangle $R \in \mathcal{R}$ and all of the rectangles “adjacent” to it are taken homeomorphically onto other rectangles of the Markov partition by both ϕ and ϕ^{-1} , then we declare R and $\phi(R)$ to be “ Y –equivalent” [Section 6]. This generates an equivalence relation on rectangles with a uniformly bounded number of equivalence classes [Corollary 6.5]. Moreover, if two different rectangles are equivalent by a power of ϕ , then that power of ϕ is cellular on those rectangles with respect to the cell structure Y on S [Proposition 6.6].

Step 3: In the mapping torus M_ϕ , the suspension flow ϕ_t applied to the rectangle for $0 \leq t \leq 1$ defines a “box” in the 3–manifold, and applying this to all rectangles in \mathcal{R} produces a decomposition of M into boxes which can be turned into a cell structure \hat{Y} on M so that $S \subset M$ with its cell structure Y is a subcomplex [Section 7].

Step 4: If a rectangle R is Y –equivalent to $\phi(R)$, then we call the box constructed from R a “filled box”. The filled boxes stack end-to-end into “prisms” [Section 8]. Each of these prisms admits a product structure of an interval times a rectangle, via the suspension flow. We construct a quotient N of M by collapsing out the flow direction in each prism. The resulting 3–complex N has a uniformly bounded number of cells in each dimension 0, 1, 2, 3, and with a uniform bound on the complexity of the attaching maps [Proposition 8.4]. There are finitely many such 3–complexes [Proposition 2.1]. The uniform boundedness is a result of careful analysis of the cell structures Y and \hat{Y} on S and M , respectively.

Step 5: The flow lines through the singular and marked points are closed trajectories in M that become 1–subcomplexes in the quotient N . We remove all of these closed trajectories from M to produce $M^\circ \in \mathcal{T}(\Psi_P^\circ)$ and we remove their image 1–subcomplexes from N to produce N° . We then prove that the associated quotient $M^\circ \rightarrow N^\circ$ is homotopic to a homeomorphism [Theorem 9.1]. Thus, $N^\circ \cong M^\circ$ is obtained from a finite list of 3–complexes

by removing a finite 1-subcomplex, which completes the proof.

Acknowledgements. We would like to thank Mladen Bestvina, Jeff Brock, Dick Canary, and Nathan Dunfield for helpful conversations.

2 Finiteness for CW-complexes

Our goal is to prove that the set of homeomorphism types of the 3-manifolds in $\mathcal{T}(\Psi_P^\circ)$ is finite. We will accomplish this by finding a finite list of compact 3-dimensional CW-complexes $\overline{\mathcal{T}(\Psi_P)}$ with the property that any 3-manifold in $\mathcal{T}(\Psi_P^\circ)$ is obtained by removing a finite subcomplex of the 1-skeleton from one of the 3-complexes in $\overline{\mathcal{T}(\Psi_P)}$.

To prove that the set of 3-complexes in $\overline{\mathcal{T}(\Psi_P)}$ is finite, we will first find a constant $K = K(P)$ so that each complex in $\overline{\mathcal{T}(\Psi_P)}$ is built from at most K cells in each dimension. This alone is not enough to conclude finiteness. For example, one can construct infinitely many homeomorphism types of 2-complexes using one 0-cell, one 1-cell, and one 2-cell. To conclude finiteness, we must therefore also impose a bound to the complexity of our attaching maps.

For an integer $K \geq 0$, we define the notion of *K -bounded complexity* for an n -cell of a CW-complex X in the following inductive manner. All 0-cells have K -bounded complexity for all $K \geq 0$. Having defined K -bounded complexity for $(n-1)$ -cells, we say that an n -cell has K -bounded complexity if the attaching map $\partial\mathbb{D}^n \rightarrow X^{(n-1)}$ has the following properties:

1. The domain $\partial\mathbb{D}^n$ can be given the structure of an $(n-1)$ -complex with at most K cells in each dimension so that each cell has K -bounded complexity, and
2. The attaching map is a homeomorphism on the interior of each cell.

Observe that a 1-cell has K -bounded complexity for all $K \geq 2$. A 2-cell has K -bounded complexity if the boundary is subdivided into at most K arcs, and the attaching map sends the interior of these arcs homeomorphically onto the interiors of 1-cells in $X^{(1)}$.

A finite cell complex has *K -bounded complexity* if each cell has K -bounded complexity and there are at most K cells in each dimension. Such a cell complex is special case of a *combinatorial complex*; see [BrH, I.8 Appendix].

Proposition 2.1. *Fix integers $K, n \geq 0$. The set of CW-homeomorphism types of n -complexes with K -bounded complexity is finite.*

Proof. Let X be an n -complex with K -bounded complexity. We may subdivide each cell of X in order to obtain a complex X' where each cell is a simplex and each attaching map is a homeomorphism on the interior of each cell; in the language of [Ha], such a complex is called a Δ -complex. What is more, we may choose X' so that the number of cells is bounded above by a constant that only depends on K and n . One way to construct X' from X is to subdivide inductively by skeleta as follows.

Since any graph is a Δ -complex, there is nothing to do for the 1-skeleton. To make the 2-skeleton into a Δ -complex, it suffices to add one vertex to the interior of each 2-cell,

and an edge connecting the new vertex to each vertex of the boundary of the 2-cell; in other words, we cone off the boundary of each 2-cell. We then cone off the 3-cells, then 4-cells, and so on, until at last we cone off all the n -cells. The number of simplices in any dimension is bounded by a function of K and the dimension, so in particular, there is a uniformly bounded number of simplices, depending only on K and n .

The second barycentric subdivision X'' of X' is a simplicial complex (this is true for any Δ -complex), and the number of vertices of X'' is bounded above by some constant K' that only depends on K and n . Observe that X is homeomorphic to X'' . In fact, the CW-homeomorphism type of X is determined by the simplicial isomorphism type of X'' , up to finite ambiguity: there are only finitely many CW-homeomorphism types that subdivide to the simplicial isomorphism type of X'' . We also note that X'' is determined up to simplicial isomorphism by specifying which subsets of the K' vertices span a simplex. Therefore, there are only finitely many simplicial isomorphism types of X'' and so only finitely many CW-homeomorphism types of X . \square

3 The adjacency graph for a Perron–Frobenius matrix

Our proof of Theorem 1.1 requires a few facts about Perron–Frobenius matrices and their associated adjacency graphs. After recalling these definitions, we give a result of Ham and Song [HS] that bounds the complexity of an adjacency graph in terms of the spectral radius of the associated matrix.




Figure 1: The adjacency graph for the 3×3 matrix A .

Let A be an $n \times n$ *Perron–Frobenius matrix*, that is, A has nonnegative integer entries $a_{ij} \geq 0$ and there exists a positive integer k so that each entry of A^k is positive. Associated to any Perron–Frobenius matrix A is its *adjacency graph* Γ_A . This is a directed graph with n vertices labeled $\{1, \dots, n\}$ and a_{ij} directed edges from vertex i to vertex j . For example,

Figure 1 shows the adjacency graph for the matrix:

$$A = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 2 & 1 & 0 \end{pmatrix}$$

The (i, j) -entry of A^m is the number of distinct combinatorial directed paths of length m from vertex i to vertex j of Γ_A . In particular, it follows that if A is a Perron–Frobenius matrix, then there is a directed path from any vertex to any other vertex.

Remark. A Perron–Frobenius matrix is sometimes defined as a matrix A with nonnegative integer entries with the property that, for each pair (i, j) , there is a k —depending on (i, j) —so that the (i, j) -entry of A^k is positive. As a consequence of our stronger definition, every positive power of a Perron–Frobenius matrix is also Perron–Frobenius.

The Perron–Frobenius Theorem (see e.g. [Ga, §III.2]) tells us that A has a positive eigenvalue $\lambda(A)$, called the *spectral radius*, which is strictly greater than the modulus of every other eigenvalue and which has a 1-dimensional eigenspace spanned by a vector with all positive entries. Moreover, one can bound $\lambda(A)$ from above and below by the maximal and minimal row sums, respectively:

$$\min_i \sum_j a_{ij} \leq \lambda(A) \leq \max_i \sum_j a_{ij}. \quad (2)$$

For a directed graph Γ and a vertex v of Γ , let $\deg_{\text{out}}(v)$ and $\deg_{\text{in}}(v)$ denote the number of edges beginning and ending at v , respectively. Since each edge has exactly one initial endpoint and one terminal endpoint, it follows that the number of edges of Γ is precisely

$$\sum_{v \in \Gamma^{(0)}} \deg_{\text{out}}(v) = \sum_{v \in \Gamma^{(0)}} \deg_{\text{in}}(v).$$

The following fact is due to Ham–Song [HS, Lemma 3.1]

Lemma 3.1. *Let A be an $n \times n$ Perron–Frobenius matrix with adjacency graph Γ_A . We have*

$$1 + \sum_{v \in \Gamma_A^{(0)}} (\deg_{\text{out}}(v) - 1) = 1 + \sum_{v \in \Gamma_A^{(0)}} (\deg_{\text{in}}(v) - 1) \leq \lambda(A)^n. \quad (3)$$

Both sums in the statement of Lemma 3.1 are equal to $-\chi(\Gamma_A)$. In particular, Lemma 3.1 bounds the number of homeomorphism types of graphs Γ_A in terms of $\lambda(A)$, since Γ_A cannot have any valence 1 vertices.

As Lemma 3.1 is central to our paper, we give the proof due to Ham and Song.

Proof. Let T be a spanning tree for Γ_A . Since $T^{(0)} = \Gamma_A^{(0)}$, we have that T has exactly n vertices. Since T is a tree, we have $\chi(T) = 1$, and so T has $n - 1$ edges.

We claim that the number of directed paths of length n starting from a vertex v of Γ_A is greater than or equal to the number of edges of Γ_A not contained in T . Indeed, each path

of length n must leave T , and so there is a surjective set map from the set of directed paths of length n based at v to the set of edges of Γ_A not contained in T : for each such path, take the first edge in the path that is not an edge of T .

Let v_0 be the vertex of Γ_A corresponding to the row of A^n that has the smallest row sum. We have:

$$\begin{aligned}
\lambda(A)^n &\geq \text{smallest row sum for } A^n \\
&= \text{number of directed paths of length } n \text{ starting from } v_0 \\
&\geq \text{number of edges of } \Gamma_A \text{ not contained in } T \\
&= \left(\sum_{v \in \Gamma_A^{(0)}} \text{deg}_{\text{out}}(v) \right) - (n-1) \\
&= 1 + \sum_{v \in \Gamma_A^{(0)}} (\text{deg}_{\text{out}}(v) - 1).
\end{aligned}$$

Since $\sum \text{deg}_{\text{out}}(v) = \sum \text{deg}_{\text{in}}(v)$, we are done. □

4 Pseudo-Anosov homeomorphisms and Markov partitions

Let S be a surface of genus g with p marked points (marked points are more convenient for us than punctures in what follows). We still write $\chi(S) = 2 - 2g - p$ and assume $\chi(S) < 0$.

4.1 Pseudo-Anosov homeomorphisms

First recall that one can describe a complex structure and integrable holomorphic quadratic differential q on S by a Euclidean cone metric with the following properties:

1. Each cone angle has the form $k\pi$ for some $k \in \mathbb{Z}_+$, with $k \geq 2$ at any unmarked point $z \in S$.
2. There is an orthogonal pair of singular foliations \mathcal{F}_h and \mathcal{F}_v on S , called the *horizontal* and *vertical foliations*, respectively, with all singularities at the cone points, and with all leaves geodesic.

Such a metric has an atlas of *preferred coordinates* on the complement of the singularities, which are local isometries to \mathbb{R}^2 and for which the vertical and horizontal foliations are sent to the vertical and horizontal foliations of \mathbb{R}^2 .

A homeomorphism $\phi : S \rightarrow S$ is pseudo-Anosov if and only if there exists a complex structure on S , a quadratic differential q on S , and a $\lambda > 1$, so that in any preferred coordinate chart for q , the map ϕ is given by

$$(x, y) \mapsto (\lambda x + c, \frac{y}{\lambda} + c')$$

for some $c, c' \in \mathbb{R}$. In particular, observe that ϕ preserves \mathcal{F}_h and \mathcal{F}_v . The horizontal foliation is called the *stable foliation*, and the leaves are all stretched by λ ; the vertical

foliation is called the *unstable foliation*, and the leaves of this foliation are contracted. The number λ is nothing other than the dilatation $\lambda = \lambda(\phi)$.

4.2 Markov partitions

The main technical tool we will use for studying pseudo-Anosov homeomorphisms is the theory of Markov partitions, which we now describe. Let ϕ be a pseudo-Anosov homeomorphism, and let q be a quadratic differential as in our description of a pseudo-Anosov homeomorphism. A *rectangle* (for q) is an immersion

$$\rho : [0, 1] \times [0, 1] \rightarrow S$$

that is affine with respect to the preferred coordinates, and that satisfies the following properties:

1. ρ maps the interior injectively onto an open set in S ,
2. $\rho(\{x\} \times [0, 1])$ is contained in a leaf of \mathcal{F}_v and $\rho([0, 1] \times \{y\})$ is contained in a leaf of \mathcal{F}_h , for all $x \in (0, 1)$ and $y \in (0, 1)$, and
3. $\rho(\partial([0, 1] \times [0, 1]))$ is a union of arcs of leaves and singularities of \mathcal{F}_v and \mathcal{F}_h .

As is common practice, we abuse notation by confusing a rectangle and its image $R = \rho([0, 1] \times [0, 1])$. The *interior* of a rectangle $\text{int}(R)$ is the ρ -image of its interior.

For any rectangle R define

$$\partial_v R = \rho(\{0, 1\} \times [0, 1]) \quad \text{and} \quad \partial_h R = \rho([0, 1] \times \{0, 1\}).$$

We thus have $\partial R = \partial_v R \cup \partial_h R$.

A *Markov partition*¹ for a pseudo-Anosov homeomorphism $\phi : S \rightarrow S$ is a finite set of rectangles $\mathcal{R} = \{R_i\}_{i=1}^n$ satisfying the following:

1. $\text{int}(R_i) \cap \text{int}(R_j) = \emptyset$ for $i \neq j$,
2. $S = R_1 \cup \dots \cup R_n$,
3. $\phi(\cup \partial_v R_i) \subseteq \cup \partial_v R_i$,
4. $\phi^{-1}(\cup \partial_h R_i) \subseteq \cup \partial_h R_i$, and
5. each marked point of S lies on the boundary of some R_i .

We say that a Markov partition $\mathcal{R} = \{R_i\}_{i=1}^n$ for ϕ is *small* if

$$|\mathcal{R}| \leq 9|\chi(S)|.$$

The following lemma appears in the work of Bestvina–Handel; see §3.4, §4.4, and §5 of [BH].

Lemma 4.1. *Suppose $\chi(S) < 0$. Every pseudo-Anosov homeomorphism of S admits a small Markov partition.*

¹What we call a “Markov partition” is sometimes called a “pre-Markov partition” in the literature; see, e.g., Exposé 9, Section 5 of [FLP].




Figure 2: A rectangle in a Markov partition is stretched horizontally and compressed vertically, then mapped over other rectangles, sending vertical sides to vertical sides. In this local picture there are 9 rectangles in one part of the surface being mapped over 8 rectangles in some other part of the surface. The colored rectangles are the ones that are mixed, while the other rectangles are unmixed.

4.3 The adjacency graph of a Markov partition

Given a Markov partition $\mathcal{R} = \{R_i\}_{i=1}^n$ for $\phi : S \rightarrow S$, there is an associated nonnegative $n \times n$ integral matrix $A = A(\phi, \mathcal{R})$, called the *transition matrix* for (ϕ, \mathcal{R}) , whose (i, j) -entry is

$$a_{ij} = |\phi^{-1}(\text{int}(R_j)) \cap R_i|,$$

where the absolute value sign denotes the number of components. That is, A records how many times the i^{th} rectangle “maps over” the j^{th} rectangle after applying ϕ .

The relation between the Perron–Frobenius theory and the pseudo-Anosov theory is provided by the following; see e.g. [FLP, Exposé 10] or [CB, pp. 101–102].

Proposition 4.2. *If $\phi : S \rightarrow S$ is pseudo-Anosov and \mathcal{R} is a Markov partition for ϕ with transition matrix A , then A is a Perron–Frobenius matrix and $\lambda(A) = \lambda(\phi)$.*

Let $\phi : S \rightarrow S$ be a pseudo-Anosov homeomorphism and let \mathcal{R} be a Markov partition for ϕ . We say that a rectangle R of \mathcal{R} is *unmixed* by ϕ if ϕ maps R homeomorphically onto a rectangle $R' \in \mathcal{R}$, and we say that it is *mixed* by ϕ otherwise.

Given $R \in \mathcal{R}$, any rectangle $R' \in \mathcal{R}$ for which $\phi^{-1}(\text{int } R') \cap R \neq \emptyset$ is called a *target rectangle* of R for ϕ . The *degree* of $R \in \mathcal{R}$ for ϕ is defined as

$$\deg(R, \phi) = \left| \phi^{-1} \left(\bigcup_{R' \in \mathcal{R}} \text{int}(R') \right) \cap R \right|$$

Informally, $\deg(R, \phi)$ records the total number of rectangles over which R maps (counted with multiplicity). The *codegree* of a rectangle $R \in \mathcal{R}$ for ϕ is the sum

$$\text{codeg}(R, \phi) = \sum_{R'} \deg(R', \phi)$$

where the sum is taken over all rectangles $R' \in \mathcal{R}$ having R as a target rectangle.

We will omit the dependence on ϕ when it is clear from context, and will simply write $\deg(R)$ and $\text{codeg}(R)$. However, it will sometimes be important to make this distinction since a Markov partition for ϕ is also a Markov partition for every nontrivial power of ϕ .

It is immediate from the definitions that if Γ is the adjacency matrix associated to (ϕ, \mathcal{R}) , and v_R is the vertex associated to the rectangle $R \in \mathcal{R}$, then $\deg(R, \phi) = \deg_{\text{out}}(v_R)$ and $\text{codeg}(R, \phi) = \deg_{\text{in}}(v_R)$. Translating other standard properties of adjacency graphs for Perron–Frobenius matrices in terms of Markov partitions, we obtain the following.

Proposition 4.3. *Suppose $\phi : S \rightarrow S$ is pseudo-Anosov, \mathcal{R} is a Markov partition for ϕ , and Γ is the adjacency graph for the associated transition matrix. Let $R \in \mathcal{R}$, and let v_R be the associated vertex of Γ . The number of distinct combinatorial directed paths in Γ that have length k and that emanate from v_R is precisely $\deg(R, \phi^k)$. The number of distinct combinatorial paths of length k leading to v_R is precisely $\text{codeg}(R, \phi^k)$.*

Likewise, the following is immediate from the definitions and the related properties of Perron–Frobenius matrices.

Proposition 4.4. *Let $\phi : S \rightarrow S$ be a pseudo-Anosov homeomorphism and let \mathcal{R} be a Markov partition for ϕ . A rectangle is unmixed by ϕ if and only if $\deg(R, \phi) = 1$ and $\text{codeg}(R', \phi) = 1$ for the unique target rectangle R' of R for ϕ .*

For a positive integer $k > 0$, if $\deg(R, \phi^k) > 1$, then $\deg(R, \phi^j) > 1$ for all $j \geq k$. Similarly, for any $k > 0$, if $\text{codeg}(R', \phi^k) > 1$ for some target rectangle R' of R by ϕ^k , then the same is true for some target rectangle of R by ϕ^j for all $j \geq k$.

Proposition 4.4 immediately implies the following.

Corollary 4.5. *Suppose $\phi : S \rightarrow S$ is pseudo-Anosov and \mathcal{R} is a Markov partition. Then if $R \in \mathcal{R}$ is mixed by ϕ^k , then it is mixed by ϕ^j for all $j \geq k$.*

We can also translate Lemma 3.1 into the context of Markov partitions for small dilatation pseudo-Anosov homeomorphisms.

Lemma 4.6. *There is an integer $C = C(P) > 0$, depending only on P , so that if $(\phi : S \rightarrow S) \in \Psi_P$ and if $\mathcal{R} = \{R_i\}$ is a small Markov partition for ϕ , then the number of rectangles of \mathcal{R} that are mixed by ϕ is at most C . Moreover, the sum of the degrees of the mixed rectangles is at most C , and the sum of the codegrees of all targets of mixed rectangles is at most C .*

Proof. By the definition of a small Markov partition, we have $|\mathcal{R}| \leq 9|\chi(S)|$, and so

$$P^{\frac{1}{|\chi(S)|}} \leq P^{\frac{9}{|\mathcal{R}|}}.$$

Then, by the definition of Ψ_P , we have $\lambda(\phi) \leq P^{\frac{1}{|\chi(S)|}} \leq P^{\frac{9}{|\mathcal{R}|}}$.

Let A be the transition matrix for the pair (ϕ, \mathcal{R}) , and let Γ be the corresponding adjacency graph. We denote by v_R the vertex of Γ associated to a rectangle $R \in \mathcal{R}$. Since $\lambda(A) = \lambda(\phi) \leq P^{\frac{9}{|\mathcal{R}|}}$, Lemma 3.1 implies that

$$\sum_{v \in \Gamma^{(0)}} (\deg_{\text{out}}(v) - 1) = \sum_{v \in \Gamma^{(0)}} (\deg_{\text{in}}(v) - 1) \leq P^9 \quad (4)$$

from which it follows that Γ has at most P^9 vertices v_R with $\deg(R, \phi) = \deg_{\text{out}}(v_R) > 1$ and at most P^9 vertices v_R with $\text{codeg}(R, \phi) = \deg_{\text{in}}(v_R) > 1$. That is, $\deg(R, \phi) = 1$ for all but at most P^9 rectangles, and, among the rectangles R with $\deg(R, \phi) = 1$, all but at most P^9 of these have $\text{codeg}(R', \phi) = 1$ for their unique target rectangle R' . Therefore, by Proposition 4.4 there are at most $2P^9$ mixed rectangles.

Finally, observe that for any $R \in \mathcal{R}$ we have

$$\deg(R, \phi) = \deg_{\text{out}}(v_R) = (\deg_{\text{out}}(v_R) - 1) + 1 \leq P^9 + 1$$

and similarly $\text{codeg}(R, \phi) \leq P^9 + 1$. Since there are at most $2P^9$ rectangles that are mixed, the sum of the degrees of the mixed rectangles is at most $2P^9(P^9 + 1)$, and similarly the sums of the codegrees of target rectangles of mixed rectangles is at most $2P^9(P^9 + 1)$.

Setting $C = 2P^9(P^9 + 1)$ completes the proof. \square

Given a rectangle R of a Markov partition, we let $\ell(R)$ denote its *length*, which is the length of either side of $\partial_h R$ with respect to the cone metric associated to q . Similarly, we let $w(R)$ denote its *width*, which is the length of either side of $\partial_v R$ with respect to q .

Lemma 4.7. *Let $(\phi : S \rightarrow S) \in \Psi_P$ and let \mathcal{R} be a small Markov partition for ϕ . If R and R' are any two rectangles of \mathcal{R} , then*

$$P^{-9} \leq \frac{\ell(R)}{\ell(R')} \leq P^9 \quad \text{and} \quad P^{-9} \leq \frac{w(R)}{w(R')} \leq P^9.$$

Proof. It suffices to show that $\ell(R)/\ell(R') \leq P^9$. By replacing ϕ with ϕ^{-1} and switching the roles of R and R' , the other inequalities will follow.

Let A be the transition matrix for (ϕ, \mathcal{R}) and let Γ be its adjacency graph. For some $k \leq |\mathcal{R}|$ the (i, j) -entry of A^k is positive, which means that R_j is a target rectangle of R_i for ϕ^k . That is, ϕ^k stretches R_i over R_j . In particular

$$\ell(\phi^{|\mathcal{R}|}(R_i)) \geq \ell(\phi^k(R_i)) \geq \ell(R_j)$$

and so it suffices to show that $\ell(\phi^{|\mathcal{R}|}(R_i)) \leq P^9 \ell(R_i)$. Indeed:

$$\ell(\phi^{|\mathcal{R}|}(R_i)) = \lambda(\phi)^{|\mathcal{R}|} \ell(R_i) \leq P^{\frac{|\mathcal{R}|}{|\chi(S)|}} \ell(R_i) \leq P^9 \ell(R_i).$$

The equality uses the definition of a Markov partition, the first inequality follows from the definition of Ψ_P , and the second inequality comes from the definition of a small Markov partition. \square

5 Cell structures on S

Let ϕ be a pseudo-Anosov homeomorphism of S , and let $\mathcal{R} = \{R_i\}$ be a Markov partition for ϕ . The Markov partition determines a cell structure on S , which we denote by $X = X(\mathcal{R})$, as follows. The vertices of X are the corners of the rectangles R_i together with the marked points and the singular points of the stable (and unstable) foliation. The vertices divide the boundary of each rectangle of \mathcal{R} into a union of arcs, which we declare to be the 1-cells of our complex. Finally, the 2-cells are simply the rectangles themselves. We refer to 1-cells as either *horizontal* or *vertical* according to which foliation they lie in.

For the next lemma, recall the notion of K -bounded complexity, which is defined in Section 2.

Lemma 5.1. *There is an integer $D_1 = D_1(P) > 0$ so that for any $\phi \in \Psi_P$, and for any small Markov partition \mathcal{R} for ϕ , each of the cells of $X(\mathcal{R})$ has D_1 -bounded complexity*

Proof. Let $R \in \mathcal{R}$. The first observation is that each of the four sides of R contains at most one marked point or one singularity of the stable foliation for ϕ . Indeed, if a side of ∂R were to contain more than one singularity or marked point, then the vertical or horizontal interval of ∂R connecting these two points would be a leaf of either the stable or unstable foliation. This is a contradiction since some power of ϕ would have to take this segment to itself, but ϕ nontrivially stretches or shrinks all leaves of the given foliation.

Now, for each 1-cell e in the boundary of R , at least one of the following holds:

- (i) e is an entire edge of the boundary of a rectangle, or
- (ii) e has a corner of R as an endpoint, or
- (iii) e has a marked point or singularity of the stable foliation for ϕ as an endpoint.

Along each side of R there are at most two 1-cells of type (ii) and, since there is at most one singularity or marked point, there are at most two 1-cells of type (iii). By Lemma 4.7, the length of any rectangle of \mathcal{R} is at least $P^{-9}\ell(R)$ and the width is at least $P^{-9}w(R)$. Thus, there are at most P^9 1-cells of type (i) along each side of R . Summarizing, there are at most $(P^9 + 4)$ 1-cells along each side of R , and hence at most $(4P^9 + 16)$ 1-cells in ∂R . Thus, the constant $D_1 = 4P^9 + 16$ satisfies the conclusion of the lemma (this is trivial for 0- and 1-cells). \square

For any rectangle $R \in \mathcal{R}$, its image $\phi(R)$ is a union of subrectangles of rectangles in \mathcal{R} . Thus, if we define

$$\phi(\mathcal{R}) = \{\phi(R) \mid R \in \mathcal{R}\}$$

then $\mathcal{R} \cup \phi(\mathcal{R})$ determines a rectangle decomposition of S , with all rectangles being subrectangles of rectangles in \mathcal{R} . We define another cell structure $Y = Y(\phi, \mathcal{R})$ on S , in exactly the same way as above, using this new rectangle decomposition; see Figure 3.

We can briefly explain the reason for introducing the cell structure Y . In Section 7, we will construct a cell structure on the mapping torus for ϕ . Each 3-cell will come from crossing a rectangle of \mathcal{R} with the unit interval. The gluing map for the one face of this “box” is given by ϕ , and hence would not be cellular in the X -structure.




Figure 3: The image of the mixed rectangles are colored (see Figure 2) and they determine the subdivision of the original rectangles into subrectangles.

The following gives bounds on the complexity of the cells of Y and the “relative complexity” of X with respect to Y .

Lemma 5.2. *There is an integer $D_2 = D_2(P)$ so that if $\phi \in \Psi_P$ and $\mathcal{R} = \{R_i\}$ is a small Markov partition for ϕ , then the following hold.*

1. *Each 1-cell of $X = X(\mathcal{R})$ is a union of at most D_2 1-cells of Y .*
2. *Each 2-cell of X is a union of at most D_2 2-cells of Y .*
3. *For each 2-cell e of X , $\phi(e)$ is a union of at most D_2 2-cells of Y .*
4. *For each 1-cell e of X , $\phi(e)$ is a union of at most D_2 1-cells of Y .*
5. *Each of the cells of $Y(\phi, \mathcal{R})$ has D_2 -bounded complexity.*

Proof. By the definition of a Markov partition, the cell structure $Y = Y(\phi, \mathcal{R})$ is obtained from the cell structure $X = X(\mathcal{R})$ by subdividing each rectangle of X into subrectangles along horizontal arcs. Given $R \in \mathcal{R}$, the number of rectangles in the subdivision is precisely $\text{codeg}(R)$, and is thus bounded by $C = C(P)$ according to Lemma 4.6. Therefore, part 2 holds for $D_2 \geq C$. Similarly, given any rectangle R , $\phi(R)$ is a union of $\text{deg}(R) \leq C$ 2-cells of Y , and so part 3 holds for $D_2 \geq C$.

The subdivision of R is obtained by first subdividing the vertical sides of R by adding at most $2(\text{codeg}(R) - 1)$ new 0-cells, then adding $\text{codeg}(R) - 1$ horizontal 1-cells from the left side of R to the right. Therefore, since each vertical 1-cell of X lies in exactly two rectangles R and R' , it is subdivided into at most $(\text{codeg}(R) + \text{codeg}(R')) \leq 2C$ 1-cells of Y by Lemma 4.6. Every horizontal 1-cell of X is still a horizontal 1-cell of Y , so part 1 holds for $D_2 \geq 2C$.

For part 5, observe that any 2-cell e of Y is a subrectangle of some rectangle $R \in \mathcal{R}$. Since each rectangle R is a 2-cell of X , it has at most D_1 1-cells of X in its boundary. By the previous paragraph, each of these 1-cells is subdivided into at most $2C$ 1-cells of Y . Therefore, there are at most $2CD_1$ 1-cells of Y in the boundary of R , and thus no more than this number in the boundary of e . Part 5 holds for any $D_2 \geq 2CD_1$ (this is trivial for 0- and 1-cells).

For part 4, observe that if e is a vertical 1-cell of X , then $\phi(e)$ is contained in the vertical boundary of some rectangle $R \in \mathcal{R}$. As mentioned in the previous paragraph, the boundary of any rectangle has at most $2CD_1$ 1-cells of Y in its boundary, so $\phi(e)$ is a union of at most $2CD_1$ 1-cells of Y . If e is a horizontal 1-cell, then e is contained in the horizontal boundary of some rectangle $R \in \mathcal{R}$, and so $\phi(e)$ is contained in the union of at most $\deg(R) \leq C$ horizontal boundaries of rectangles of the subdivision. As already mentioned, the horizontal boundary of each of these rectangles is either one of the added 1-cells, or else a horizontal boundary of some rectangle of \mathcal{R} , which contains at most D_1 1-cells of X (which are also the 1-cells of Y as they are contained in the horizontal edges of a rectangle). Therefore, $\phi(e)$ contains at most CD_1 1-cells of Y . Therefore, part 4 holds for any $D_2 \geq 2CD_1$.

The lemma now follows by setting $D_2 = 2CD_1$. \square

6 Equivalence relations on rectangles

Given a pseudo-Anosov $(\phi : S \rightarrow S) \in \Psi_P$, let \mathcal{R} be a small Markov partition for ϕ . In this section we will use this data to construct a quotient space $\pi : S \rightarrow T$, and a cell structure W on T for which π is cellular with respect to Y . Moreover, we will prove that W has uniformly bounded complexity.

The quotient T will be obtained by gluing together rectangles of \mathcal{R} via a certain equivalence relation. This relation has the property that if a rectangle R is equivalent to R' , then there is a (unique) power $\phi^{\beta(R,R')}$ that takes R homeomorphically onto the rectangle R' . Moreover, $\phi^{\beta(R,R')}$ will be cellular with respect to Y -cell structures on R and R' , respectively.

This equivalence relation is most easily constructed from simpler equivalence relations. Our approach will be to define a first approximation to the equivalence relation we are searching for, and then refine it (twice) to achieve the equivalence relation with the required properties. Along the way we verify other properties that will be needed later.

6.1 The first approximation: h -equivalence

Define an equivalence relation \sim^h on \mathcal{R} by declaring

$$R \sim^h R'$$

if there exists $\beta \in \mathbb{Z}$ so that ϕ^β takes R homeomorphically onto R' . The ‘h’ stands for “homeomorphically”. As no rectangle can be taken homeomorphically to itself by a non-trivial power of ϕ , it follows that if $R \sim^h R'$ then there exists a unique integer $\beta(R, R')$ for which $\phi^{\beta(R,R')}(R) = R'$.

Proposition 6.1. *Let $\phi \in \Psi_P$, and let \mathcal{R} be a small Markov partition for ϕ . The h -equivalence classes in \mathcal{R} have the form*

$$\{\phi^j(R)\}_{j=0}^k$$

where each $k = k(R)$ is an integer. There is a constant $E_h = E_h(P)$ so that there are at most E_h h -equivalence classes.

Proof. We first observe that $R \stackrel{h}{\sim} R'$ if and only if R is unmixed for ϕ^j and $\phi^j(R) = R'$. Proposition 4.4 then implies that $R \stackrel{h}{\sim} R'$ if and only if $\deg(R, \phi^j) = 1$, R' is the unique target rectangle of R and $\text{codeg}(R', \phi^j) = 1$. Corollary 4.5 now implies that the h -equivalence classes have the required form.

To prove the second statement of the proposition, we observe that, according to the description of the h -equivalence relation provided in the previous paragraph, the last rectangle R in any given equivalence class (which is well-defined by the first part of the proposition) satisfies one of the following two conditions:

(i) $\deg(R, \phi) > 1$.

(ii) $\deg(R, \phi) = 1$ and $\text{codeg}(R', \phi) > 1$, where R' is the unique target rectangle of R .

The number of rectangles R of type (i) is bounded above by the constant $C = C(P)$ from Lemma 4.6. The number of type (ii) is bounded above by the sum of the codegrees of rectangles with codegree greater than 1. Again, by Lemma 4.6, this is bounded by C . Thus, we may take E_h to be $2C$. \square

Given $(\phi : S \rightarrow S) \in \Psi_P$ and \mathcal{R} a small Markov partition for ϕ , we index the rectangles of \mathcal{R} as $\{R_i\}_{i=1}^n$ so that

1. The h -equivalence classes all have the form $\{R_i, R_{i+1}, \dots, R_{i+k}\}$.
2. If $\{R_i, R_{i+1}, \dots, R_{i+k}\}$ is an equivalence class, then $R_{i+j} = \phi^j(R_i)$.

That this is possible follows from Proposition 6.1.

6.2 The second approximation: N -equivalence

Say that two rectangles $R, R' \in \mathcal{R}$ are *adjacent* if $R \cap R'$ contains at least one point that is not a singular point or a marked point. Observe that, by definition, a rectangle is adjacent to itself. We let $N_1(R) \subset \mathcal{R}$ denote the set of rectangles adjacent to R . We think of $N_1(R)$ as the “1-neighborhood” of R .

Lemma 6.2. *There exists a constant $D_3 = D_3(P)$ so that if $(\phi : S \rightarrow S) \in \Psi_P$ and \mathcal{R} is a small Markov partition for ϕ , then the number of rectangles in $N_1(R)$ is at most D_3 for any $R \in \mathcal{R}$.*

Proof. The number of rectangles in $N_1(R)$ is at most the number of 1-cells of X in the boundary of R plus 5. This is because each 1-cell contributes at most one new adjacent rectangle to $N_1(R)$, each corner vertex of R contributes at most one more rectangle, and R itself contributes 1. From Lemma 5.2 it follows that we can take $D_3 = D_2 + 5$. \square

Now we define a refinement of h -equivalence, called N -equivalence, by declaring

$$R \stackrel{N}{\sim} R'$$

if $\phi^{\beta(R,R')}$ does not mix R'' for any $R'' \in N_1(R)$ (the ‘ N ’ stands for 1-neighborhood). That is, $\phi^{\beta(R,R')}$ maps the entire 1-neighborhood $N_1(R)$ to the 1-neighborhood of R' , taking each rectangle homeomorphically onto another rectangle. We leave it to the reader to check that N -equivalence is a refinement of h -equivalence.

Figure 4 shows the local picture of the surface after applying ϕ^j , for $j = 0, 1, 2, 3$. We have labeled a few of the rectangles:

$$R_i, R_{i+1} = \phi(R_i), R_{i+2} = \phi^2(R_i), R_{i+3} = \phi^3(R_i) \quad \text{and} \\ R_j, R_{j+1} = \phi(R_j), R_{j+2} = \phi^2(R_j), R_{j+3} \supsetneq \phi^3(R_j).$$

The rectangles are related as follows:

$$R_i \stackrel{h}{\sim} R_{i+1} \stackrel{h}{\sim} R_{i+2} \stackrel{h}{\sim} R_{i+3} \\ R_i \not\stackrel{N}{\sim} R_{i+1} \stackrel{N}{\sim} R_{i+2} \not\stackrel{N}{\sim} R_{i+3}$$

(also $R_i \not\stackrel{N}{\sim} R_{i+3}$), and

$$R_j \stackrel{h}{\sim} R_{j+1} \stackrel{h}{\sim} R_{j+2} \not\stackrel{h}{\sim} R_{j+3} \\ R_j \stackrel{N}{\sim} R_{j+1} \stackrel{N}{\sim} R_{j+2} \not\stackrel{N}{\sim} R_{j+3}.$$

Of course, each rectangle is equivalent to itself with respect to either relation.

Proposition 6.3. *Let $\phi \in \Psi_P$, and let \mathcal{R} be a small Markov partition for ϕ . Each N -equivalence class has the form*

$$\{R_i, R_{i+1}, \dots, R_{i+s}\} = \{R_i, \phi(R_i), \dots, \phi^s(R_i)\}$$

for some $s = s(R_i) \in \mathbb{Z}$. There is a constant $E_N = E_N(P)$ so that \mathcal{R} has at most E_N N -equivalence classes.

Proof. Similarly to the proof of Proposition 6.1, the first statement follows from Corollary 4.5. Now, consider an arbitrary h -equivalence class:

$$\{R_i, R_{i+1}, \dots, R_{i+k}\} = \{R_i, \phi(R_i), \dots, \phi^k(R_i)\}.$$

This is partitioned into its N -equivalence classes as follows. The class divides at R_{i+j} (that is, R_{i+j+1} begins a new N -equivalence class) if and only if $N_1(R_{i+j})$ contains a rectangle that is mixed by ϕ . By Lemma 4.6, at most $C = C(P)$ rectangles of \mathcal{R} are mixed. Moreover, according to Lemma 6.2, each rectangle—in particular, each mixed rectangle—is a 1-neighbor to at most $D_3 = D_3(P)$ rectangles. Therefore, there are at most CD_3 rectangles that are 1-neighbors of mixed rectangles. Thus, each h -equivalence class can be subdivided into at most $(CD_3 + 1)$ N -equivalence classes. According to Proposition 6.1 the number of h -equivalence classes is at most E_h , and thus the proposition follows if we take $E_N = E_h(CD_3 + 1)$. \square

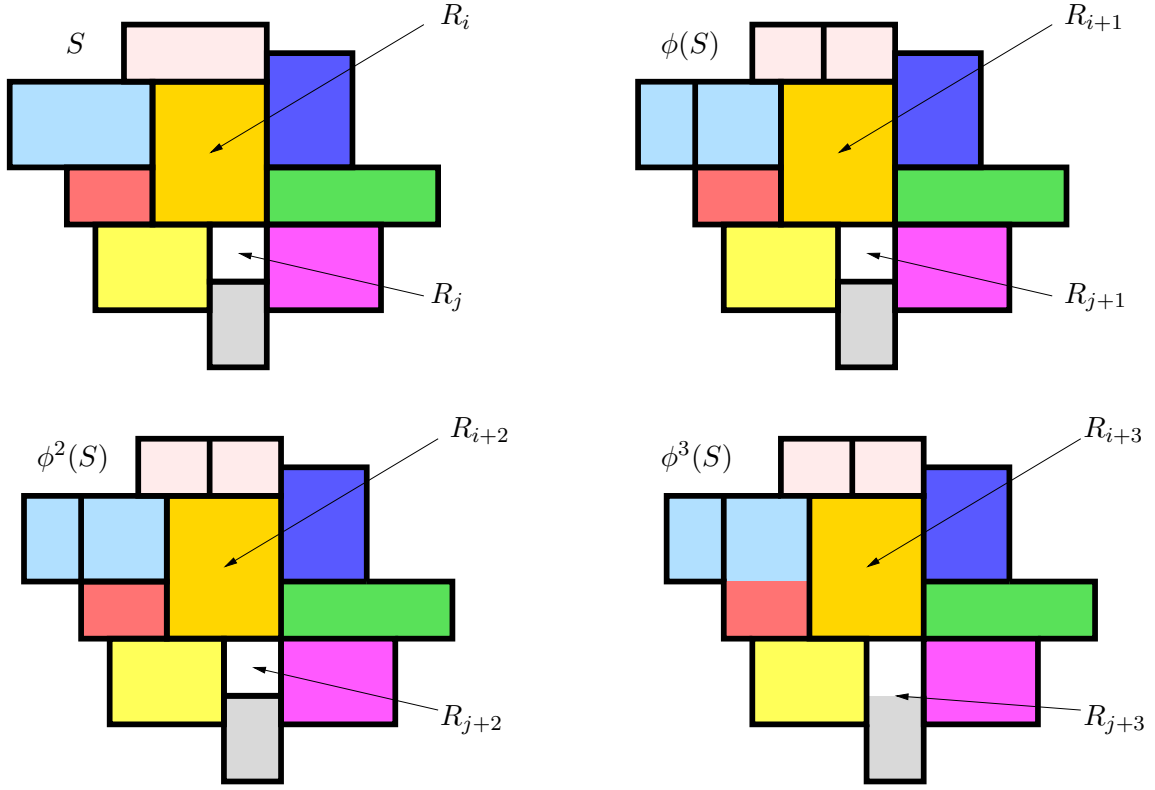


Figure 4: A local picture of the surface after applying ϕ^j , $j = 0, 1, 2, 3$.

The next proposition explains one advantage of the N -equivalence relation over the h -equivalence relation.

Proposition 6.4. *Let $(\phi : S \rightarrow S) \in \Psi_P$ and let \mathcal{R} be a small Markov partition for ϕ . If $R \stackrel{N}{\sim} R'$, then $\phi^{\beta(R,R')}|_R : R \rightarrow R'$ is cellular with respect to X .*

Proof. The cell structure X is defined using the rectangles of \mathcal{R} , together with the way adjacent rectangles intersect one another, and the singular and marked points. Since ϕ preserves the singular and marked points, and since $R \stackrel{N}{\sim} R'$ implies $\phi^{\beta(R,R')}$ maps all rectangles of $N_1(R)$ homeomorphically onto a rectangle in $N_1(R')$, the result follows. \square

6.3 The “right” relation on rectangles: Y -equivalence

We define a refinement of N -equivalence, called Y -equivalence, by dividing each N -equivalence class $\{R_i, R_{i+1}, \dots, R_{i+k}\}$ with more than one element into two Y -equivalence classes $\{R_i\}$ and $\{R_{i+1}, R_{i+2}, \dots, R_{i+k}\}$. That is, we split off the initial element of each N -equivalence class into its own Y -equivalence class. The Y -equivalence classes are also consecutive with

respect to the indices and so we can refer to the initial and terminal rectangles of a Y -equivalence class. We write $R \overset{Y}{\sim} R'$ if R and R' are Y -equivalent.

The new feature of Y -equivalence is that any time two distinct rectangles are Y -equivalent, they both have codegree one. In fact, all rectangles in the 1-neighborhood are ϕ -images of unmixed rectangles. It follows that this equivalence relation behaves nicely with respect to the cell structure Y , hence the terminology; see Proposition 6.6 below.

If we take $E_Y = 2E_N$, we immediately obtain the following consequence of Proposition 6.3.

Corollary 6.5. *There is a constant $E_Y = E_Y(P)$ so that for any $\phi \in \Psi_P$, any small Markov partition \mathcal{R} for ϕ has at most E_Y Y -equivalence classes.*

The next proposition is the analogue of Proposition 6.4 for Y -equivalence and the cell structure Y .

Proposition 6.6. *Let $(\phi : S \rightarrow S) \in \Psi_P$ and let \mathcal{R} be a small Markov partition for ϕ . If the Y -equivalence class of R contains more than one element, and $R \overset{Y}{\sim} R'$, then the X - and Y -cell structures on R agree, and likewise for R' . Moreover, $\phi^{\beta(R,R')}|_R : R \rightarrow R'$ is cellular with respect to Y .*

The assumption that the Y -equivalence class of R contains more than one element is necessary for the first statement, since otherwise it would follow that the X and Y cell structures coincide on all of S . This would imply that ϕ is cellular with respect to X , and hence finite order, which is absurd.

Proof. Since the Y -equivalence class of R and R' contains more than one element, neither R nor R' is the initial rectangle of the N -equivalence class they lie in. Thus, $\phi^{-1}(R)$ and $\phi^{-1}(R')$ are both elements of \mathcal{R} and

$$\phi^{-1}(R) \overset{N}{\sim} R \overset{N}{\sim} \phi^{-1}(R') \overset{N}{\sim} R'.$$

By Proposition 6.4, the maps

$$\phi|_{\phi^{-1}(R)} : \phi^{-1}(R) \rightarrow R \quad \text{and} \quad \phi|_{\phi^{-1}(R')} : \phi^{-1}(R') \rightarrow R'$$

are both cellular with respect to X . Thus, by the definition of Y , it follows that the Y -structures on R and R' are exactly the same as the X -structures on R and R' , respectively. As $R \overset{N}{\sim} R'$, Proposition 6.4 guarantees that

$$\phi^{\beta(R,R')}|_R : R \rightarrow R'$$

is cellular with respect to X , and hence also with respect to Y , as required. \square

6.4 From rectangles to the surface

We will use the Y -equivalence relation to construct a quotient of S by gluing R to R' by $\phi^{\beta(R,R')}$ whenever $R \overset{Y}{\sim} R'$. To better understand this quotient, we will study the equivalence relation on S that this determines. This is most easily achieved by breaking the equivalence relation up into simpler relations as follows.

We write $(x, R) \leftrightarrow (x', R')$ to mean that the following conditions hold:

1. $x \in R$
2. $x' \in R'$
3. $R \overset{Y}{\sim} R'$
4. $\phi^{\beta(R,R')}(x) = x'$

Now write $x \leftrightarrow x'$ and say that x is related to x' if $(x, R) \leftrightarrow (x', R')$ for some $R, R' \in \mathcal{R}$.

The relation \leftrightarrow is easily seen to be symmetric and reflexive, but it may not be transitive. This is because if x' lies in two distinct rectangles R' and R'' , it may be that $(x, R) \leftrightarrow (x', R')$ and $(x', R'') \leftrightarrow (x'', R''')$. We let \sim denote the equivalence relation on S obtained from the transitive closure of the relation \leftrightarrow on S .

Lemma 6.7. *The \sim -equivalence class of x , which is contained in $\{\phi^k(x)\}_{k \in \mathbb{Z}}$, consists of consecutive ϕ -iterates of x . What is more, if $x \sim x'$, then after possibly interchanging the roles of x and x' , there exists $k \in \mathbb{Z}_{\geq 0}$, and a sequence of rectangles $R_{i_0}, \dots, R_{i_{k-1}} \in \mathcal{R}$, so that*

$$x \leftrightarrow \phi(x) \leftrightarrow \dots \leftrightarrow \phi^{k-1}(x) \leftrightarrow \phi^k(x) = x'$$

with $(\phi^j(x), R_{i_j}) \leftrightarrow (\phi^{j+1}(x), \phi(R_{i_j}))$ for $j = 0, \dots, k-1$.

The second part of the lemma says that when $x \sim x'$, after possibly interchanging x and x' , we can get from x to x' moving forward through consecutive elements of the ϕ -orbit by applying the relation \leftrightarrow . We caution the reader that it may be necessary to interchange the roles of x and x' , even when they lie in a periodic orbit.

Remark. For a periodic point, the orbit $\{\phi^k(x)\}_{k \in \mathbb{Z}}$ is a finite set, in which case the ordering is a cyclic ordering. However, it still makes sense to say that a set consists of consecutive ϕ -iterates of x .

Proof of Lemma 6.7. Note that if $(x, R) \leftrightarrow (x', R')$ with $\beta(R, R') \geq 0$ then we also have

$$(x, R) \leftrightarrow (\phi^j(x), \phi^j(R))$$

for $j = 0, \dots, \beta(R, R')$, since $\{R, \phi(R), \dots, \phi^{\beta(R,R')}(R)\}$ is contained in a single Y -equivalence class (that is, the Y -equivalence classes of rectangles are consecutive). Since \sim is the transitive closure of \leftrightarrow , the full equivalence class is obtained by stringing together these sets whenever they intersect, and it follows that the equivalence class of x consists of consecutive ϕ -iterates of x , and further that they are related as in statement of the lemma. \square

6.5 The cell structure Y and \sim -equivalence

The main purpose of this section is to describe the structure of the \sim -equivalence classes and how they relate to the Y cell structure.

Given a point $x \in S$, set

$$k_-(x) = \inf\{k \in \mathbb{Z} \mid \phi^j(x) \sim x \text{ for all } k \leq j \leq 0\}$$

and

$$k_+(x) = \sup\{k \in \mathbb{Z} \mid \phi^j(x) \sim x \text{ for all } 0 \leq j \leq k\}.$$

That is, we look at all consecutive points in the orbit of x that are equivalent, and take the infimum and supremum, respectively, of the consecutive exponents (beginning at 0) that occur. According to Lemma 6.7, the set $\{\phi^j(x)\}_{j=k_-(x)}^{k_+(x)}$ is the \sim -equivalence class of x . Observe that if x is a fixed point of ϕ then $k_-(x) = -\infty$ and $k_+(x) = +\infty$.

Lemma 6.8. *If x is not a marked point or singular point, then $k_-(x), k_+(x) \in \mathbb{Z}$. In particular, it cannot be the case that the entire ϕ -orbit of x lies in the same \sim -equivalence class.*

Proof. First suppose that x is a periodic point of order n . The only way that $k_+(x) = \infty$ or $k_-(x) = -\infty$ is if any two points of the orbit are \sim -equivalent. Up to changing x within its orbit, we may assume that there is a rectangle R containing x that is mixed by ϕ ; indeed, otherwise the pseudo-Anosov map ϕ^{kn} would preserve the collection of rectangles containing x for all k which is impossible. We now prove $x \not\sim \phi(x)$, which will complete the proof in the case of a periodic point.

The proof is by contradiction, so assume that $x \sim \phi(x)$. According to Lemma 6.7, there are two cases to consider depending on whether or not the roles of x and $\phi(x)$ must be interchanged. Thus, we either have a rectangle R' for which

$$(x, R') \leftrightarrow (\phi(x), \phi(R'))$$

or else there exists a sequence of rectangles $R_{i_0}, \dots, R_{i_{n-2}}$ so that

$$(\phi^j(\phi(x)), R_{i_j}) \leftrightarrow (\phi^{j+1}(\phi(x)), \phi(R_{i_j}))$$

for $j = 0, \dots, n-2$. In particular, in the second case we have $(\phi(x), R_{i_0}) \leftrightarrow (\phi^2(x), \phi(R_{i_0}))$.

The first case is clearly impossible since R' is adjacent to R which is mixed by ϕ , and hence $R' \overset{Y}{\not\sim} \phi(R')$. In particular, it follows that $(x, R') \not\leftrightarrow (\phi(x), \phi(R'))$, which is a contradiction. In the second case, we have $R_{i_0} \overset{Y}{\sim} \phi(R_{i_0})$ which implies $\phi^{-1}(R_{i_0}) \overset{N}{\sim} R_{i_0}$. Since $\phi^{-1}(R_{i_0})$ is a rectangle containing x , it is adjacent to (the mixed rectangle) R and therefore $\phi^{-1}(R_{i_0}) \overset{N}{\not\sim} R_{i_0}$, another contradiction.

Therefore, it must be the case that $x \not\sim \phi(x)$, as required.

We now consider the case where x is not a periodic point for ϕ . By Lemma 6.7, the equivalence class of x consists of consecutive ϕ -iterates. Let j be a positive integer so that one of the rectangles containing $\phi^j(x)$, say R , is mixed by ϕ . As in the periodic case, it follows that $\phi^j(x) \not\leftrightarrow \phi^{j+1}(x)$. Since x is aperiodic, it follows from Lemma 6.7 that $x \not\sim \phi^{j+1}(x)$, so $k_+(x) \leq j + 1$.

Similarly, note that there is some $j > 0$ so that $\phi^{-j}(x)$ is contained in a rectangle R that is mixed by ϕ . So, $\phi^{-j}(x) \not\leftrightarrow \phi^{-j+1}(x)$, and again aperiodicity of x together with Lemma 6.7 implies $\phi^{-j}(x) \not\sim x$, so $k_-(x) \geq -j + 1$. □

We now prove that $k_{\pm}(x)$ depends only on the cell containing x in its interior.

Proposition 6.9 (Structure of \sim). *Let e be a cell of Y that is not a singular point or a marked point. If $x, y \in \text{int}(e)$, then $k_{\pm}(x) = k_{\pm}(y)$. If $x \in \text{int}(e)$ and $y \in \partial e$, then*

$$k_-(y) \leq k_-(x) \leq k_+(x) \leq k_+(y).$$

In particular, we can define $k_{\pm}(e) = k_{\pm}(x)$ for any $x \in \text{int}(e)$ (independently of the choice of $x \in \text{int}(e)$), and for every integer $\alpha \in [k_-(e), k_+(e)]$, the map $\phi^\alpha|_e$ is cellular with respect to Y .

Proof. We suppose first that $x \in \text{int}(e)$ and $y \in e$ and we prove that $k_+(x) \leq k_+(y)$. If $k_+(x) = 0$, there is nothing to prove so suppose $k = k_+(x) > 0$. Combining Lemmas 6.7 and 6.8, there exists a sequence of rectangles $R_{i_0}, \dots, R_{i_{k-1}}$, so that

$$(\phi^j(x), R_{i_j}) \leftrightarrow (\phi^{j+1}(x), \phi(R_{i_j}))$$

for $j = 0, \dots, k-1$. For each j , Proposition 6.6 implies that $\phi|_{R_{i_j}}$ is cellular with respect to Y . Since e is a cell in R_{i_0} , by induction and Proposition 6.6, we see that $\phi^{j+1}(e)$ is a cell in both $\phi(R_{i_j})$ and $R_{i_{j+1}}$. Hence $\phi^j(y) \in R_{i_j}$ and $\phi^{j+1}(y) \in \phi(R_{i_j})$, and therefore

$$(\phi^j(y), R_{i_j}) \leftrightarrow (\phi^{j+1}(y), \phi(R_{i_j}))$$

for every $j = 0, \dots, k-1$. It follows that $y \sim \phi^j(y)$ for $j = 1, \dots, k$ and thus $k_+(y) \geq k = k_+(x)$.

Observe that if $y \in \text{int}(e)$, we can reverse the roles of x and y to obtain $k_+(y) \leq k_+(x)$, and hence $k_+(x) = k_+(y)$. Otherwise, $y \in \partial e$, and we only know $k_+(x) \leq k_+(y)$.

The proof that $k_-(x) = k_-(y)$ if $y \in \text{int}(e)$ and $k_-(y) \leq k_-(x)$ if $x \in \partial e$ follows a similar argument.

That we can define $k_{\pm}(e) = k_{\pm}(x)$ for any $x \in \text{int}(e)$ now follows. Finally, the fact that $\phi^\alpha|_e$ is cellular is proven in the course of the proof above. \square

6.6 The quotient of S

Denote S/\sim by T and let $\pi : S \rightarrow T$ be the quotient map.

Proposition 6.10. *Let $(\phi : S \rightarrow S) \in \Psi_P$ and let \mathcal{R} be a small Markov partition for ϕ . There is a cell structure $W = W(\phi, \mathcal{R})$ on T , so that π is cellular with respect to the cell structure $Y = Y(\phi, \mathcal{R})$ on S . Moreover, for each cell e of Y , π restricts to a homeomorphism from the interior of e onto the interior of a cell $\pi(e)$ of W . Finally, there is a $D = D(P)$ so that W has D -bounded complexity.*

Proof. It follows from Proposition 6.9 that \sim defines an equivalence relation on the cells of Y , which we also call \sim , by declaring

$$e \sim \phi^\alpha(e) \quad \text{for all } k_-(e) \leq \alpha \leq k_+(e).$$

If e is a cell in Y that is neither a singular point nor a marked point, and $e \sim e'$, then appealing to Lemma 6.7, Lemma 6.8, and Proposition 6.9, there is a unique integer $\alpha(e, e') \in$

$[k_-(e), k_+(e)]$ so that $\phi^{\alpha(e,e')}(e) = e'$. Moreover, by Proposition 6.9, $\phi^{\alpha(e,e')}|_e : e \rightarrow e'$ is cellular with respect to Y .

An exercise in CW-topology shows that T admits a cell structure W with one cell for each equivalence class of cells in Y . Moreover, the characteristic maps for these cells can be taken to be the characteristic maps for cells of Y , composed with π . Indeed, if $e \sim e'$ in Y , and e and e' are p -cells with $p \geq 1$, then we may assume that the characteristic maps ψ_e and $\psi_{e'}$ for e and e' , respectively, are related by

$$\phi^{\alpha(e,e')} \circ \psi_e = \psi_{e'}.$$

Of course, for 0-cells, the characteristic maps are canonical. It follows that π is cellular and is a homeomorphism on the interior of any cell.

Since π is cellular and is a homeomorphism on the interior of each cell, Lemma 5.2 implies that each 2-cell of W has D -bounded complexity for any $D \geq D_2 = D_2(P)$. Recall that the 0- and 1-cells trivially have D -bounded complexity for all $D \geq 2$. Lemma 5.2 also implies that every rectangle $R \in \mathcal{R}$ is subdivided into at most D_2 2-cells of Y . By Corollary 6.5 there are at most E_Y Y -equivalence classes of rectangles, and therefore there are at most $D_2 E_Y$ 2-cells of W . Since each 2-cell has D_2 -bounded complexity, there are at most $D_2^2 E_Y$ 0- and 1-cells. Thus, W has D -bounded complexity for $D = D_2^2 E_Y$. \square

7 Cell structure for the mapping torus

Let $\phi : S \rightarrow S$ be a pseudo-Anosov homeomorphism, and let $M = M_\phi$ denote the mapping torus. This is the 3-manifold obtained as a quotient of $S \times [0, 1]$ by identifying $(x, 1)$ with $(\phi(x), 0)$ for all $x \in S$. We view S as embedded in M as follows:

$$S \rightarrow S \times \{0\} \rightarrow S \times [0, 1] \rightarrow M.$$

Let $\{\phi_t \mid t \in \mathbb{R}\}$ denote the suspension flow on M ; this is the flow on M determined by the local flow $\tilde{\phi}_t(x, s) = (x, s + t)$ on $S \times [0, 1]$. The time-one map restricted to S is the first return map, that is, $\phi_1|_S = \phi$. From this it follows that $\phi_k(x) = \phi^k(x)$ for every integer $k \in \mathbb{Z}$ and $x \in S$.

We also have the punctured surface version. The punctured surface S° embeds in S and is ϕ -invariant. Hence, we may view M° as $M_{\phi|_{S^\circ}}$, the mapping torus of $\phi|_{S^\circ} : S^\circ \rightarrow S^\circ$, embedded in $M = M_\phi$.

7.1 Boxes

Let $(\phi : S \rightarrow S) \in \Psi_P$ and let $\mathcal{R} = \{R_i\}$ be a small Markov partition for ϕ . For each rectangle R_i , there is an associated *box*, defined by

$$B_i = \bigcup_{0 \leq t \leq 1} \phi_t(R_i).$$

More precisely, if $\rho_i : [0, 1] \times [0, 1] \rightarrow R_i$ is the parameterized rectangle, then the box is parameterized as

$$\hat{\rho}_i : ([0, 1] \times [0, 1]) \times [0, 1] \rightarrow B_i$$

where

$$\hat{\rho}_i(\mathbf{u}, t) = \phi_t(\rho_i(\mathbf{u})).$$

As with rectangles, the map $\hat{\rho}_i$ is only an embedding when restricted to the interior. And as in the case of rectangles, we abuse notation, and refer to B_i as a subset of M , although it is formally a map into M .

7.2 The cell structure

We impose on M a cell structure $\hat{Y} = \hat{Y}(\phi, \mathcal{R})$ defined by the cell structures X and Y on S as follows. The 0-cells of \hat{Y} are the 0-cells of Y (recall that $S \subset M$). The 1-cells of \hat{Y} are the 1-cells of Y , called *surface 1-cells*, together with suspensions of 0-cells of X :

$$\bigcup_{0 \leq t \leq 1} \phi_t(v) \text{ for } v \in X^{(0)}.$$

Because $X^{(0)} \subset Y^{(0)}$ and since $\phi(X^{(0)}) \subset Y^{(0)}$, the boundary of each of these 1-cells is contained in $\hat{Y}^{(0)} = Y^{(0)}$, as required. We call these the *suspension 1-cells*.

The 2-cells of \hat{Y} are the 2-cells of Y , called *surface 2-cells*, together with suspensions of 1-cells of X :

$$\bigcup_{0 \leq t \leq 1} \phi_t(e) \text{ for } e \text{ a 1-cell of } X.$$

Observe that for each 1-cell e of X , both e and $\phi(e)$ are 1-subcomplexes of Y . Furthermore, since the boundary of e is contained in $X^{(0)}$, we see that the boundary of the 2-cell defined by the above suspension is contained in $\hat{Y}^{(1)}$, as required. We call these 2-cells *suspension 2-cells*.

Finally, the 3-cells of \hat{Y} are the boxes. All 3-cells can be thought of as a *suspension 3-cells*, since they are suspensions of surface 2-cells.

Proposition 7.1. *There exists a positive integer $K_1 = K_1(P)$ with the following property. If $(\phi : S \rightarrow S) \in \Psi_P$ and if \mathcal{R} is a small Markov partition for ϕ , then each cell of $\hat{Y} = \hat{Y}(\phi, \mathcal{R})$ has K_1 -bounded complexity.*

Proof. All 0-cells and 1-cells trivially have K_1 -bounded complexity for all $K_1 \geq 2$.

Since S is embedded in M as the subcomplex $Y \subset \hat{Y}$, it follows from part 5 of Lemma 5.2 that the surface 2-cells of \hat{Y} have K_1 -bounded complexity for any $K_1 \geq D_2 = D_2(P)$.

If e is a suspension 2-cell obtained by suspending a 1-cell e_0 of X , then the boundary of e consists of 2 suspension 1-cells, and by parts 1 and 4 of Lemma 5.2, at most $2D_2$ surface 1-cells. Therefore, these 2-cells have K_1 -bounded complexity for any $K_1 \geq 2D_2 + 2$.

The boundary of each 3-cell is a union 6 rectangles: a “bottom” and a “top” (which are rectangles of \mathcal{R} and $\phi(\mathcal{R})$, respectively), and four “suspension sides”. The number of 0-cells is just the number of 0-cells in the top and bottom rectangles (since all vertices lie in S). Each of the top and bottom rectangles is a union of at most D_2 surface 2-cells by parts 2 and 3 of Lemma 5.2. Each surface 2-cell has D_2 -bounded complexity, so has at most D_2 vertices in its boundary. It follows that each 3-cell has at most $2D_2^2$ vertices in its boundary.

The number of suspension 2-cells in the boundary of a 3-cell B_i is the number of 1-cells of X in the boundary of $R_i \in \mathcal{R}$. Since the cells of X have $D_1 = D_1(P)$ -bounded complexity by Lemma 5.1, and the 2-cells of X are precisely the rectangles of \mathcal{R} , it follows that the boundary of a 3-cell has at most D_1 suspension 2-cells in its boundary. Combining this with the bounds from the previous paragraph on the number of surface 2-cells, it follows that the boundary of a 3-cell has at most $D_1 + 2D_2$ 2-cells in its boundary. Finally, since the boundary of a 3-cell is a 2-sphere, the Euler characteristic tells us that the number of 1-cells in the boundary is 2 less than the sum of the numbers of 0-cells and 2-cells, and so is at most $D_1 + 2D_2 + 2D_2^2 \leq D_1 + 4D_2^2$. It follows that if $K_1 \geq D_1 + 4D_2^2$, then each 3-cell has K_1 -bounded complexity.

Therefore, setting $K_1 = D_1 + 4D_2^2$ completes the proof. \square

A subset of the suspension 1-cells are the *singular* and *marked* 1-cells; these are the 1-cells of the form

$$\bigcup_{0 \leq t \leq 1} \phi_t(v)$$

for v either a singular point or marked point, respectively. These 1-cells, together with their vertices, form 1-dimensional subcomplexes called the *singular* and *marked subcomplexes*. These subcomplexes are unions of circles in M , and M° is obtained from M by removing them.

8 Quotient spaces I: bounded complexity 3-complexes

In this section we use the notion of Y -equivalence to produce a quotient of the compact 3-manifold $M = M_\phi$; the quotient we obtain will be compact, but might not be a manifold. We will prove that this quotient admits a 3-dimensional cell structure with uniformly bounded complexity for which the quotient map is cellular. The singular and marked subcomplexes of \hat{Y} will define subcomplexes of the quotient, and in the next section we will prove that the space obtained by removing these subcomplexes from the quotient is homeomorphic to M° , the corresponding mapping torus of the punctured surface S° . First, we define and analyze the quotients.

Let $(\phi : S \rightarrow S) \in \Psi_P$ be a pseudo-Anosov homeomorphism and let $\mathcal{R} = \{R_i\}_{i=1}^n$ be a small Markov partition for ϕ , indexed so that each Y -equivalence class has the form $\{R_i, \dots, R_{i+k}\}$ with $R_{i+j} = \phi^j(R_i)$; see Section 6. Let ϕ_t be the suspension flow on M . As in Section 7, we obtain collection of boxes $\{B_i\}_{i=1}^n$. Let \hat{Y} be the cell structure on M as described in Section 7.

8.1 Prisms

To each Y -equivalence class of rectangles, we will associate a *prism* as follows. If $\{R_i, \dots, R_{i+k}\}$ is a Y -equivalence class, then the associated prism is

$$P_i = \bigcup_{0 \leq t \leq k} \phi_t(R_i).$$

The rectangles R_i and R_{i+k} are the *bottom* and *top* of the prism, respectively. Alternatively, if $\{R_i, R_{i+1}, \dots, R_{i+k}\}$ is a Y -equivalence class with more than one element, then the associated prism is

$$P_i = B_i \cup B_{i+1} \cup \dots \cup B_{i+k-1}.$$

On the other hand if $\{R_i\}$ is a Y -equivalence class consisting of the single rectangle R_i , then $P_i = R_i \subset S \subset M$. If a box is contained in a prism, then we will say it is a *filled box*. See Figure 5.

Let L be the union of the prisms in M . For clarification, we note that L is the union of S together with the set of filled boxes, and is a subcomplex of \hat{Y} .




Figure 5: The figure shows two prisms (green and yellow) in M_ϕ . The prisms stop near the top of the picture as the rectangles at the tops of the prisms are mixed: they are both mapped by ϕ into the topmost rectangle in the picture. At the bottom of the picture, both rectangles are unmixed, but only one of the boxes (the yellow one) is filled. This means that the rectangle directly below the green prism is not Y -equivalent to its image by ϕ , the bottom of the green prism (although it is h -equivalent).

The prisms define an equivalence relation \approx on M that is a “continuous version” of \sim on S . More precisely, we declare $y \approx z$ if and only if either $y = z$ or else there exists an interval $I = [0, r]$ or $I = [r, 0]$ in \mathbb{R} so that $\phi_r(z) = y$ and

$$\bigcup_{t \in I} \phi_t(z) \subset L.$$

That is, $y \approx z$ if there is an arc of a flow line containing y and z that is contained in L .

We will be particularly interested in the flow lines in M through points in S . We denote the flow line in M through $x \in S$ by

$$\ell_x = \bigcup_{t \in \mathbb{R}} \phi_t(x).$$

Proposition 8.1. *The equivalence relation \approx restricted to S is precisely the equivalence relation \sim .*

Proof. According to Lemma 6.7, if $x \sim x'$, then after reversing the roles of x and x' , we have rectangles $R_{i_0}, \dots, R_{i_{k-1}} \in \mathcal{R}$ so that

$$x \leftrightarrow \phi(x) \leftrightarrow \phi^2(x) \leftrightarrow \dots \leftrightarrow \phi^{k-1}(x) \leftrightarrow \phi^k(x) = x'$$

with $(\phi^j(x), R_{i_j}) \leftrightarrow (\phi^{j+1}(x), \phi(R_{i_j}))$ for $j = 0, \dots, k-1$. Therefore,

$$\phi^j(x), \phi^{j+1}(x) \in \bigcup_{0 \leq t \leq 1} \phi_t(\phi^j(x)) = \ell_{\phi^j(x)} \cap B_{i_j}.$$

By the definition of the \leftrightarrow relation, we have that $R_{i_j} \overset{Y}{\sim} \phi(R_{i_j})$, and so B_{i_j} is a filled box. Thus $\phi^j(x) \approx \phi^{j+1}(x)$ for each $j = 0, \dots, k-1$. By transitivity, $x \approx x'$.

For the other direction, suppose $x \approx x'$. We have that $x' \in \ell_x$ and the arc $\ell_x^0 \subset \ell_x$ from x to x' is contained in L . Say that $x' = \phi^j(x)$ for $j \geq 0$ (reverse the roles of x and x' if necessary). We can write ℓ_x^0 as

$$\ell_x^0 = \bigcup_{0 \leq t \leq j} \phi_t(x) = \bigcup_{i=0}^{j-1} \bigcup_{0 \leq t \leq 1} \phi_t(\phi^i(x)).$$

Since each arc

$$\bigcup_{0 \leq t \leq 1} \phi_t(\phi^i(x))$$

lies in L , it is contained in a filled box. It follows that $\phi^i(x) \leftrightarrow \phi_1(\phi^i(x)) = \phi^{i+1}(x)$ (recall that a Y -equivalence between rectangles induces \leftrightarrow -relations between all pairs of corresponding points). Therefore

$$x \leftrightarrow \phi(x) \leftrightarrow \phi^2(x) \leftrightarrow \dots \leftrightarrow \phi^j(x) = x'.$$

Since \sim is the transitive closure of \leftrightarrow , it follows that $x \sim x'$. □

8.2 Flowing out of L

The next proposition provides a bound to how long a flow line can stay in L .

Proposition 8.2. *Let $(\phi : S \rightarrow S) \in \Psi_P$ and let \mathcal{R} be a small Markov partition for ϕ . If $x \in S$ is not a singular point or a marked point, then $\ell_x \cap L$ is a union of compact arcs and points in ℓ_x . Moreover, there exists an integer $Q = Q(\phi, \mathcal{R})$ so that each component of $\ell_x \cap L$ has length strictly less than Q (with respect to the flow parameter).*

Proof. Fix $x \in S$ to be a nonsingular, unmarked point. Since the set of such points is invariant under ϕ , we have that $\ell_x \cap S = \{\phi^j(x)\}_{j=-\infty}^{\infty}$ contains no singular or marked points.

We need only consider the component $\ell_x^0 \subset \ell_x \cap L$ containing x (since every component has this form for some x). Recall that Lemma 6.8 implies that the \sim -equivalence class of x is precisely

$$\{\phi^{k_-(x)}(x), \dots, \phi^{k_+(x)}(x)\}$$

for integers $k_-(x) \leq 0 \leq k_+(x)$. By Proposition 8.1, it follows that ℓ_x^0 is equal to the compact arc

$$\bigcup_{k_-(x) \leq t \leq k_+(x)} \phi_t(x) = \ell_x^0$$

of length $k_+(x) - k_-(x)$ (if $k_-(x) = 0 = k_+(x)$ the arc is a point).

On the other hand, Proposition 6.9 implies $k_{\pm}(x) = k_{\pm}(e)$, where e is the unique cell containing x in its interior. Therefore, the length of the longest arc of intersection of a flow line not passing through a marked point or singular point is

$$Q' = \max \{k_+(e) - k_-(e) \mid e \text{ is an unmarked, nonsingular cell of } Y\},$$

which is finite since there are only finitely many cells in Y . Setting $Q = Q' + 1$ completes the proof. \square

8.3 Product structures

In the next proposition, we identify the box B_i with $R_i \times [0, 1]$ via the suspension flow as described in Section 7.

Given a space V , a subspace $U \subset V$ and a cell structure Z on V , if U happens to be a subcomplex with respect to the cell structure Z , then we will refer to the induced cell structure on U as the *restriction* of Z to U , and write it as $Z|U$.

Proposition 8.3. *For each filled box B_i , the restriction $\hat{Y}|B_i$ agrees with the product cell structure on $R_i \times [0, 1]$ coming from $Y|R_i$ and the cell structure on $[0, 1]$ that has a single 1-cell.*

Proof. In a product cell structure, the cells all have the form $e_0 \times e_1$, where e_0 and e_1 are cells of the first and second factors, respectively. Thus, the cells of the product $Y|R_i \times [0, 1]$ are of the form (1) $e_0 \times [0, 1]$, (2) $e_0 \times \{0\}$, and (3) $e_0 \times \{1\}$. So the proposition is saying that for each cell e of \hat{Y} in a filled box B_i , there is a cell e_0 from $Y|R_i$ so that one of the following holds:

$$(1) \ e = \bigcup_{0 \leq t \leq 1} \phi_t(e_0), \quad (2) \ e = e_0, \quad \text{or} \quad (3) \ e = \phi_1(e_0).$$

So suppose

$$B_i = \bigcup_{0 \leq t \leq 1} \phi_t(R_i)$$

is a filled box. Then $R_i \stackrel{Y}{\sim} R_{i+1} = \phi(R_i)$, and, according to Proposition 6.6, we have that $\phi|_{R_i} : R_i \rightarrow R_{i+1}$ is cellular with respect to $Y|R_i = X|R_i$ and $Y|R_{i+1} = X|R_{i+1}$.

Since the suspension cells are precisely the suspensions of cells of X , which in R_i are exactly the cells of Y , it follows that all suspension cells are of the form (1) above. All other cells are surface cells in R_i and R_{i+1} . The cells in R_i are of course of the form (2). Since, again, $\phi|_{R_i}$ is cellular with respect to Y , every cell of $Y|R_{i+1}$ is of the form $\phi_1(e_0)$ for some cell e_0 of $Y|R_i$, and thus has the form (3). \square

8.4 Cell structure on the quotient

Let N denote M/\approx and let $p : M \rightarrow N$ be the quotient map. According to Proposition 8.1, the restriction to S of the relation \approx is precisely the relation \sim and so the inclusion of $S \rightarrow M$ descends to an inclusion $T \rightarrow N$, and we have the following commutative diagram:

$$\begin{array}{ccc} S & \xrightarrow{\pi} & T \\ \downarrow & & \downarrow \\ M & \xrightarrow{p} & N \end{array}$$

Proposition 8.4. *The quotient N admits a cell structure $\hat{W} = \hat{W}(\phi, \mathcal{R})$ so that p is cellular with respect to \hat{Y} on M . Moreover, there exists $K = K(P)$ so that \hat{W} has K -bounded complexity.*

Proof. We view T as a subset of N . First, recall that L is the union of the prisms and that S is contained in L , and observe that all of L is mapped (surjectively) to T . Indeed, on any filled box B_i the map to N is obtained by first projecting onto R_i using its product structure coming from the flow, then projecting R_i to T by π . By Proposition 8.3, this is a cellular map from $\hat{Y}|L$ to W , the cell structure on T from Proposition 6.10.

Let B_{i_1}, \dots, B_{i_k} denote the set of unfilled boxes. The 3-manifold M , with cell structure \hat{Y} , is the union of two subcomplexes:

$$L \quad \text{and} \quad L^c = S \cup B_{i_1} \cup \dots \cup B_{i_k}.$$

Note that L and L^c are not complementary, as both contain S . Since both S and L map onto T , it follows that L^c maps onto N :

$$N = p(L^c) = T \cup p(B_{i_1}) \cup \dots \cup p(B_{i_k}).$$

We now construct a cell structure \hat{W} on N for which $p|_{L^c} : L^c \rightarrow N$ is cellular with respect to $\hat{Y}|L^c$ and for which T is a subcomplex with $\hat{W}|T = W$. If we can do this, then since $p|_L : L \rightarrow N$ is already cellular, and M is the union of the two subcomplexes, it will follow that p is cellular.

The above description of N as the image of S and the unfilled boxes indicates the way to build the cell structure \hat{W} on N . Namely, we start by giving T the cell structure W and then describe the remaining cells as images of certain cells in L^c .

First suppose that e is a suspension 1-cell that is not contained in L . Then since any suspension 1-cell is contained in the intersection of some set of boxes, it must be that the

1-cell meets L only in its endpoints. In particular, p is injective when restricted to the interior of e , and the endpoints map into $W^{(0)}$. We take $p(e)$ as a 1-cell in \hat{W} .

In addition to the 2-cells of W , we also construct a new 2-cell e' for each suspension 2-cell e that is not contained in L , as follows. The suspension 2-cells are contained in intersections of boxes, and so since e is not contained in L , it meets L only in its boundary. Therefore, p is injective on the interior of e . Moreover, the boundary contains two suspension 1-cells and two arcs that are unions of surface 1-cells of \hat{Y} . The map p sends all the surface 1-cells of ∂e cellularly into T , and is injective on the interior of each cell. Each of the two suspension 1-cells is either sent to a 1-cell constructed in the previous section (hence injectively on the interior), or is collapsed to a 0-cell. The cell e' is essentially $p(e)$, the only difference being the cell structure on $\partial e'$ has collapsed those suspension 1-cells in ∂e that are collapsed by p in N .

Finally, the 3-cells come from the unfilled boxes B_{i_1}, \dots, B_{i_k} . Each unfilled box has a graph on its boundary coming from the cell structure \hat{Y} . Moreover, by Proposition 7.1, this graph has at most $K_1 = K_1(P)$ vertices, edges, and complementary regions, all of which are disks. These disks are suspension and surface 2-cells in $\hat{Y}|L^c$. The 3-cells in \hat{W} are obtained by first collapsing any suspension 1-cells and 2-cells in the boundary of the box (using the product structure from the flow) that are contained in L , then projecting by p to N . This collapsing produces a new graph in the boundary of the 3-cell, but the new graph has no more edges, vertices or complementary regions than the original one. Moreover, all complementary components are still disks; see Figure 6 for a picture of the 3-cell and the new cell structure on the boundary. With this new cell structure on the boundary, the restriction of p to the interior of each cell is injective. We also observe that each of these 3-cells has K_1 -bounded complexity.




Figure 6: On the left is a box in M with the suspension cells contained in L shown in red. On the right is a picture of the 3-cell after collapsing the suspension cells of the boundary contained in L .

It is straightforward to check that this is the desired cell structure; p is cellular on L^c , and by construction T is a subcomplex with $\hat{W}|T = W$. Therefore, p is cellular on all of $N = L \cup L^c$ as discussed above, and we have proven the first part of the proposition.

Now we find a $K = K(P) > 0$ so that \hat{W} has K -bounded complexity. First, the number of 3-cells is precisely the number of unfilled boxes. Each unfilled box is given by

$$\bigcup_{0 \leq t \leq 1} \phi_t(R)$$

where R is a terminal rectangle in its Y -equivalence class. Therefore, the unfilled boxes

correspond precisely to the Y -equivalence classes of rectangles. By Corollary 6.5, there are at most $E_Y = E_Y(P)$ unfilled boxes, and so at most E_Y 3-cells. As mentioned above, each of these 3-cells has K_1 -bounded complexity.

The cell structure W has $D = D(P)$ -bounded complexity by Proposition 6.10, so there are at most D 1-cells in W . Each 1-cell of \hat{W} not in W is contained in the boundary of some 3-cell. Since these have K_1 -bounded complexity, and since there are at most E_Y of these, it follows that there are at most $E_Y K_1$ 1-cells in \hat{W} that are not in W . Therefore, there are at most $(D + E_Y K_1)$ 1-cells in \hat{W} .

Similarly, the number of 2-cells is bounded by the number of 2-cells in W plus the number of 2-cells in each of the unfilled boxes. A count as in the previous paragraph implies that there are at most $D + E_Y K_1$ 2-cells in \hat{W} . Finally, the 2-cells of W have D -bounded complexity by Proposition 6.10, and each of the suspension 2-cells has K_1 -bounded complexity since it is in the boundary of a 3-cell, which has K_1 -bounded complexity.

Therefore, setting

$$K = K(P) = D + E_Y K_1,$$

it follows that \hat{W} has K -bounded complexity. \square

Let

$$\overline{\mathcal{T}(\Psi_P)} = \{\hat{W}(\phi, \mathcal{R}) \mid \phi \in \Psi_P \text{ and } \mathcal{R} \text{ a small Markov partition for } \phi\} / \text{CW-homeomorphism}.$$

The following is immediate from Propositions 2.1 and 8.4.

Corollary 8.5. *The set $\overline{\mathcal{T}(\Psi_P)}$ is finite.*

9 Quotient spaces II: Finitely many 3-manifolds

We continue with the notation from the previous section: $M_{\phi|_{S^\circ}} = M^\circ \subset M = M_\phi$, L is the union of S and the prisms, and $p : M \rightarrow N$ is the quotient defined by collapsing arcs of flow lines in L to points. Let N° denote $p(M^\circ)$, the complement in N of the singular and marked subcomplexes. We also write $p : M^\circ \rightarrow N^\circ$ for the restriction. The main goal of this section is to prove the following.

Theorem 9.1. *The map $p : M^\circ \rightarrow N^\circ$ is homotopic to a homeomorphism.*

Recall that a 3-manifold is *irreducible* if every 2-sphere in the 3-manifold bounds a 3-ball. We will deduce Theorem 9.1 from Waldhausen's Theorem (the version we will use can be found in [He, Corollary 13.7]):

Theorem 9.2 (Waldhausen). *Suppose M_1 and M_2 are compact, orientable irreducible 3-manifolds with nonempty boundary. If $f : (M_1, \partial M_1) \rightarrow (M_2, \partial M_2)$ is a map such that $f_* : \pi_1(M_1) \rightarrow \pi_1(M_2)$ is an isomorphism and $f_* : \pi_1(\partial_0 M_1) \rightarrow \pi_1(\partial_0 M_2)$ is an injection for each component $\partial_0 M_1 \subset \partial M_1$ and the component $\partial_0 M_2 \subset \partial M_2$ containing $f(\partial_0 M_1)$, then f is homotopic to a homeomorphism.*

Assuming Theorem 9.1, we can prove Theorem 1.1, which states that the set $\mathcal{T}(\Psi_P^\circ)$ is finite.

Proof of Theorem 1.1. By Corollary 8.5, $\overline{\mathcal{T}(\Psi_P)}$ is a finite set of CW–homeomorphism types of cell complexes. Therefore, since each N° is obtained by removing a 1–subcomplex from a compact 3–complex $N \in \overline{\mathcal{T}(\Psi_P)}$, it follows that the set

$$\{N^\circ \mid N \text{ represents an element of } \overline{\mathcal{T}(\Psi_P)}\} / \text{homeomorphism}$$

is finite. But by Theorem 9.1, this set is equal to $\mathcal{T}(\Psi_P^\circ)$. \square

So, we are left to prove Theorem 9.1. This will occupy the remainder of the section.

Let $\rho : M^\circ \rightarrow S^1$ denote the fibration coming from the mapping torus description of M° . Let $\omega \in \Omega^1(S^1)$ be a volume form so that $\int_{S^1} \omega = 1$, and let $\mu = \rho^*(\omega) \in \Omega^1(M^\circ)$ be the pullback (which on $S^\circ \times [0, 1]$ we can take to be dt). This 1–form μ represents the Poincaré dual of the class represented by $S^\circ \subset M^\circ$. Integration of μ defines an epimorphism $\int \mu : \pi_1(M^\circ) \rightarrow \mathbb{Z}$, and we let

$$q : \widetilde{M}^\circ \rightarrow M^\circ$$

denote the infinite cyclic cover of M° determined by $\int \mu$.

Alternatively, q is simply the cover corresponding to $\pi_1(S^\circ)$, and so is obtained by unwrapping the bundle. Therefore, we have

$$\widetilde{M}^\circ \cong S^\circ \times \mathbb{R}.$$

We make this diffeomorphism explicit as follows. First, choose a lift $S^\circ \subset \widetilde{M}^\circ$ of the embedding $S^\circ \subset M^\circ$. The flow ϕ_t lifts to a flow $\widetilde{\phi}_t$ on \widetilde{M}° . We now define a map

$$S^\circ \times \mathbb{R} \rightarrow \widetilde{M}^\circ$$

by

$$(x, t) \mapsto \widetilde{\phi}_t(x).$$

We use this diffeomorphism to identify \widetilde{M}° with $S^\circ \times \mathbb{R}$.

The covering group for $q : \widetilde{M}^\circ \rightarrow M^\circ$ is infinite cyclic. Let δ be the generator that induces a translation by $+1$ in the second factor of $\widetilde{M}^\circ \cong S^\circ \times \mathbb{R}$. With respect to the product structure, we have

$$\delta(x, t) = (\phi^{-1}(x), t + 1).$$

Writing the last formula in terms of the flow $\widetilde{\phi}_t$, we obtain

$$\delta(\widetilde{\phi}_t(x)) = \delta(x, t) = (\phi^{-1}(x), t + 1) = \widetilde{\phi}_{t+1}(\phi^{-1}(x)),$$

and iterating, we obtain

$$\delta^k(\widetilde{\phi}_t(x)) = \widetilde{\phi}_{t+k}(\phi^{-k}(x)). \tag{5}$$

We will study the quotient N° of M° by looking at the associated quotient of \widetilde{M}° . Let $L^\circ = L \cap M^\circ$, and let $\widetilde{L}^\circ = q^{-1}(L^\circ)$. For any $x \in S^\circ$, we consider the flow line

$$\widetilde{\ell}_x = \bigcup_{t \in \mathbb{R}} \widetilde{\phi}_t(x).$$

Since $\tilde{\ell}_x$ is a lift of the flow line ℓ_x , Proposition 8.2 implies that $\tilde{\ell}_x \cap \tilde{L}^\circ$ is a union of compact arcs (and points) of length less than $Q = Q(\phi, \mathcal{R})$ (measured with respect to the flow parameter). For $y, z \in \tilde{M}^\circ$ define $y \equiv z$ if there exists $x \in S^\circ$ so that y, z are in the same component of $\tilde{\ell}_x \cap \tilde{L}^\circ$. Set

$$\tilde{N}^\circ = \tilde{M}^\circ / \equiv$$

and let

$$\tilde{p} : \tilde{M}^\circ \rightarrow \tilde{N}^\circ$$

be the quotient map.

We will now give an alternative description of \tilde{N}° as an open subset of $S^\circ \times \mathbb{R}$; we will realize it as the image of a smooth map $f : \tilde{M}^\circ \rightarrow S^\circ \times \mathbb{R}$. Let $h \in C^\infty(M^\circ)$ be a smooth function that is positive on $M^\circ - L^\circ$ and is identically zero on L° . Let $\nu = h\mu$, and $\tilde{\nu} = q^*(\nu)$. Since $\tilde{\nu}$ is a pullback, it is invariant by δ :

$$\delta^*(\tilde{\nu}) = \tilde{\nu}.$$

Let dt denote the 1-form on $\tilde{M}^\circ \cong S^\circ \times \mathbb{R}$ obtained by pulling back the volume form on \mathbb{R} , and observe that on $S^\circ \times \mathbb{R}$ we have

$$\tilde{\nu} = (h \circ q) dt.$$

Now define $f : S^\circ \times \mathbb{R} \rightarrow S^\circ \times \mathbb{R}$ by

$$f(x, t) = \left(x, \int_0^t \tilde{\nu}_x \right)$$

where the notation in the integral means the integral of $\tilde{\nu}$ over the path $\phi_s(x)$ as s runs from 0 to t . Since $\tilde{\nu}$ is smooth, the map f is smooth. Set

$$V = f(\tilde{M}^\circ).$$

We remark that V may or may not be all of $S^\circ \times \mathbb{R}$.

Lemma 9.3. *The fibers of f are precisely the \equiv -equivalence classes. That is, f and \tilde{p} have the same fibers.*

Proof. First observe that if $x \neq x'$, then $f(\tilde{\ell}_x) \cap f(\tilde{\ell}_{x'}) = \emptyset$ since $f(\tilde{\ell}_x) = \tilde{f}(\{x\} \times \mathbb{R}) \subset \{x\} \times \mathbb{R}$. It follows that we can only have $y \equiv z$ if y and z both lie on $\tilde{\ell}_x$ for some x . Therefore, it suffices to show that for any $x \in S^\circ$ and for all $y, z \in \tilde{\ell}_x$, we have $f(y) = f(z)$ if and only if $y \equiv z$.

To prove this, we observe that for any flow line $\tilde{\ell}_x$, the restriction $\tilde{\nu}_x$ of $\tilde{\nu}$ to $\tilde{\ell}_x$ is zero precisely on the intersection $\tilde{\ell}_x \cap \tilde{L}^\circ$. Therefore, $f|_{\tilde{\ell}_x}$ is constant on each component of $\tilde{\ell}_x \cap L^\circ$ and monotone (increasing) on the complementary intervals. It follows that the fibers of $f|_{\tilde{\ell}_x}$ are precisely the \equiv -equivalence classes. \square

Because $\tilde{\nu}$ is invariant by δ , the map f semiconjugates δ to a map $\hat{\delta} : V \rightarrow V$. That is, $\hat{\delta} \circ f = f \circ \delta$. An explicit formula for $\hat{\delta}$ is given by

$$\hat{\delta}(x, s) = \left(\phi^{-1}(x), s + \int_0^1 \tilde{\nu}_{\phi^{-1}(x)} \right).$$

Proposition 9.4. *The image $V = f(S^\circ \times \mathbb{R})$ is an open subset of $S^\circ \times \mathbb{R}$, and is homeomorphic to $S^\circ \times \mathbb{R}$. Moreover, $\langle \hat{\delta} \rangle$ acts properly discontinuously and freely on V , and with respect to this product structure, $\hat{\delta}$ is given by ϕ^{-1} on the first factor.*

Proof. Let $Q = Q(\phi, \mathcal{R})$ be the constant from Proposition 8.2. As mentioned above (immediately after the definition of $\tilde{\ell}_x$), the integer Q is greater than the length of the component of any flow line $\tilde{\ell}_x^0 \subset \tilde{\ell}_x \cap \tilde{L}^\circ$. In particular, it follows that

$$\int_0^Q \tilde{\nu}_x > 0$$

for all $x \in S$.

It follows that $f(S^\circ \times \{0\}) \cap f(S^\circ \times \{Q\}) = \emptyset$. On the other hand, we have

$$\hat{\delta}^Q(f(S^\circ \times \{0\})) = f(\delta^Q(S^\circ \times \{0\})) = f(S^\circ \times \{Q\}).$$

Since $f(S^\circ) = S^\circ = S^\circ \times \{0\} \subset V$, this implies

$$S^\circ \cap \hat{\delta}^Q(S^\circ) = \emptyset$$

in V .

The region between S° and $\hat{\delta}^Q(S^\circ)$ in $S^\circ \times \mathbb{R}$ is homeomorphic to a product. Indeed, it is exactly the region between the graphs of the zero function on S° and the smooth positive function

$$x \mapsto \int_0^Q \tilde{\nu}_x.$$

We denote this region by Δ . Observe that

$$\Delta = \left\{ \left(x, \int_0^t \tilde{\nu}_x \right) \mid x \in S^\circ \text{ and } 0 \leq t \leq Q \right\} = f(S^\circ \times [0, Q])$$

and so $\Delta \subset V$.

The region Δ is a fundamental domain for the action of $\langle \hat{\delta}^Q \rangle$ on V , and V is equal to

$$\bigcup_{j=-\infty}^{\infty} \hat{\delta}^{jQ}(\Delta) \tag{6}$$

with $\langle \hat{\delta}^Q \rangle$ acting properly discontinuously. Since $\langle \hat{\delta}^Q \rangle$ is contained in $\langle \hat{\delta} \rangle$ with finite index, it follows that $\langle \hat{\delta} \rangle$ also acts properly discontinuously. Because $\langle \hat{\delta} \rangle$ is torsion free, its action on V is also free.

Any homeomorphism from Δ to $S^\circ \times [0, Q]$ that is the identity on the surface factor uniquely extends to a homeomorphism $V \rightarrow S^\circ \times \mathbb{R}$ that conjugates the action of $\langle \hat{\delta}^Q \rangle$ to

the action of $\langle \delta^Q \rangle$. As the homeomorphism is the identity on the first factor, the conjugate of $\hat{\delta}$ (which is different from δ) restricts to ϕ^{-1} on the first factor.

From (6), we also see that V is an open set: any point is either in the interior of a $\langle \hat{\delta}^Q \rangle$ -translate of Δ , or the union of two consecutive translates. \square

Finally, we prove Theorem 9.1, which states that the map $p : M^\circ \rightarrow N^\circ$ is homotopic to a homeomorphism.

Proof of Theorem 9.1. We are going to make connections between the various spaces and maps we have constructed (and some we have yet to construct). Throughout, the following diagram will serve as a guide.

$$\begin{array}{ccccc}
 & & V & \xrightarrow{\cong} & S^\circ \times \mathbb{R} \\
 & \nearrow f & \downarrow \eta & & \\
 \widetilde{M}^\circ & \xrightarrow{\tilde{p}} & \widetilde{N}^\circ & & \\
 \downarrow q & & \searrow r & & \\
 M^\circ & \xrightarrow{p} & N^\circ & \xrightarrow{\cong} & \widetilde{N}^\circ / \langle \hat{\delta} \rangle
 \end{array}$$

The map $f : \widetilde{M}^\circ \rightarrow V \cong S^\circ \times \mathbb{R}$ is a proper continuous surjection, and hence is a quotient map. Also, by definition, $\tilde{p} : \widetilde{M}^\circ \rightarrow \widetilde{N}^\circ$ is a quotient map. Since the fibers of these two maps are the same (Lemma 9.3), it follows that the quotients are homeomorphic. Let $\eta : V \rightarrow \widetilde{N}^\circ$ be the homeomorphism for which $\eta \circ f = \tilde{p}$.

We use the homeomorphism η and the homeomorphism $V \cong S^\circ \times \mathbb{R}$ from Proposition 9.4 to identify \widetilde{N}° with $S^\circ \times \mathbb{R}$. Conjugating by this homeomorphism we obtain an action of $\langle \hat{\delta} \rangle$ on \widetilde{N}° . We abuse notation and simply refer to the conjugate homeomorphism by the same name:

$$\hat{\delta} : \widetilde{N}^\circ \rightarrow \widetilde{N}^\circ.$$

Now, N° is the image of M° via the quotient map $p : M^\circ \rightarrow N^\circ$, and since the covering map $q : \widetilde{M}^\circ \rightarrow M^\circ$ is also a quotient map, the composition $p \circ q : \widetilde{M}^\circ \rightarrow N^\circ$ is a quotient map. We also have the covering map

$$r : \widetilde{N}^\circ \rightarrow \widetilde{N}^\circ / \langle \hat{\delta} \rangle$$

which is a quotient map, and hence so is $r \circ \tilde{p}$. Since the fibers of $p \circ q$ are the same as those of $r \circ \tilde{p}$, it follows that

$$N^\circ \cong \widetilde{N}^\circ / \langle \hat{\delta} \rangle.$$

General covering space theory implies that $\pi_1(N^\circ)$ fits into a short exact sequence:

$$1 \rightarrow \pi_1(\widetilde{N}^\circ) \rightarrow \pi_1(N^\circ) \rightarrow \mathbb{Z} \rightarrow 1.$$

The monodromy $\mathbb{Z} \rightarrow \text{Out}(\pi_1(\widetilde{N}^\circ))$ is given by $1 \mapsto \hat{\delta}_*$. On the other hand, the isomorphism $\pi_1(\widetilde{N}^\circ) \rightarrow \pi_1(S^\circ)$ induced by projecting onto the first factor, conjugates $\hat{\delta}_*$ to ϕ_*^{-1} by Proposition 9.4. Therefore, $\pi_1(N^\circ)$ is an extension of \mathbb{Z} by $\pi_1(S^\circ)$ with monodromy $\mathbb{Z} \rightarrow$

$\text{Out}(\pi_1(S^\circ))$ given by $1 \mapsto \phi_*$. The fundamental group $\pi_1(M^\circ)$ also has such a description. Since \mathbb{Z} is free, each sequence splits, and it follows that $\pi_1(M^\circ)$ is isomorphic to $\pi_1(N^\circ)$. Moreover, the map $p : M^\circ \rightarrow N^\circ$ induces this isomorphism

$$p_* : \pi_1(M^\circ) \rightarrow \pi_1(N^\circ)$$

as can be seen in the lift $\tilde{p} : \tilde{M}^\circ \rightarrow \tilde{N}^\circ$.

Since N° is covered by $S^\circ \times \mathbb{R}$ and S° is a hyperbolic surface, it follows that the universal cover of N° is \mathbb{R}^3 . Therefore N° is irreducible. Furthermore, N° is the complement of a 1-subcomplex of the cell complex N , which has K -bounded complexity by Proposition 8.4. Since N can be subdivided into a simplicial complex (see the proof of Proposition 2.1), it follows that N° is tame, that is, it is homeomorphic to the interior of a compact 3-manifold with boundary.

Because N° is tame, we can remove a product neighborhood of the ends to produce a *compact core* \hat{N}° for N° , which is a compact submanifold for which the inclusion is a homotopy equivalence. Since p is a proper map, $p^{-1}(\hat{N}^\circ)$ is a compact subset of M° . On the other hand, M° is also tame, and so we can also remove product neighborhoods of the ends of M° to arrive at a compact core \hat{M}° for M° that contains $p^{-1}(\hat{N}^\circ)$.

Using the product structure on the complement of \hat{M}° in M° and of \hat{N}° in N° , there are strong deformation retractions $h_M : M^\circ \rightarrow \hat{M}^\circ$ and $h_N : N^\circ \rightarrow \hat{N}^\circ$. The map $F = h_N \circ p \circ h_M$ is therefore homotopic to p , and moreover the restriction $F|_{\hat{M}^\circ} : \hat{M}^\circ \rightarrow \hat{N}^\circ$, satisfies the hypotheses of Waldhausen's Theorem (Theorem 9.2): the only thing to verify is that the boundary subgroups are mapped injectively, but that is clear from the construction of p . Therefore, $F|_{\hat{M}^\circ}$ is homotopic to a homeomorphism. Using this homotopy and the product neighborhoods of the ends one can construct a homotopy from F , and hence from p , to a homeomorphism. \square

References

- [FLP] *Travaux de Thurston sur les surfaces*, volume 66 of *Astérisque*. Société Mathématique de France, Paris, 1979. Séminaire Orsay, With an English summary.
- [Ag] Ian Agol. Small 3-manifolds of large genus. *Geom. Dedicata*, 102:53–64, 2003.
- [AY] Pierre Arnoux and Jean-Christophe Yoccoz. Construction de difféomorphismes pseudo-Anosov. *C. R. Acad. Sci. Paris Sér. I Math.*, 292(1):75–78, 1981.
- [Ba] Max Bauer. An upper bound for the least dilatation. *Trans. Amer. Math. Soc.*, 330(1):361–370, 1992.
- [BH] M. Bestvina and M. Handel. Train-tracks for surface homeomorphisms. *Topology*, 34(1):109–140, 1995.
- [BrH] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [Br] Jeffrey F. Brock. Weil-Petersson translation distance and volumes of mapping tori. *Comm. Anal. Geom.*, 11(5):987–999, 2003.

- [CB] Andrew J. Casson and Steven A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston*, volume 9 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1988.
- [CH] Jin-Hwan Cho and Ji-Young Ham. The minimal dilatation of a genus-two surface. *Experiment. Math.*, 17(3):257–267, 2008.
- [Fr] David Fried. Flow equivalence, hyperbolic systems and a new zeta function for flows. *Comment. Math. Helv.*, 57(2):237–259, 1982.
- [Ga] F. R. Gantmacher. *The theory of matrices. Vols. 1, 2*. Translated by K. A. Hirsch. Chelsea Publishing Co., New York, 1959.
- [HS] Ji-Young Ham and Won Taek Song. The minimum dilatation of pseudo-Anosov 5-braids. *Experiment. Math.*, 16(2):167–179, 2007.
- [Ha] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [He] John Hempel. *3-Manifolds*. Princeton University Press, Princeton, N. J., 1976. Ann. of Math. Studies, No. 86.
- [HK] Eriko Hironaka and Eiko Kin. A family of pseudo-Anosov braids with small dilatation. *Algebr. Geom. Topol.*, 6:699–738 (electronic), 2006.
- [Iv] N. V. Ivanov. Coefficients of expansion of pseudo-Anosov homeomorphisms. *Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI)*, 167(Issled. Topol. 6):111–116, 191, 1988.
- [KKT] Eiko Kin, Sadayoshi Kojima, and Mitsuhiko Takasawa. Entropy vs volume for pseudo-Anosov maps. Preprint, arXiv:0812.2941.
- [KT] Eiko Kin and Mitsuhiko Takasawa. Pseudo-Anosov braids with small entropy and the magic 3-manifold. Preprint, arXiv:0812.4589.
- [LT] Erwan Lanneau and Jean-Luc Thiffeault. On the minimum dilatation of pseudo-Anosov homeomorphisms on surfaces of small genera. Preprint.
- [Li] Michele Linch. A comparison of metrics on Teichmüller space. *Proc. Amer. Math. Soc.*, 43:349–352, 1974.
- [Mc1] Curtis T. McMullen. *Renormalization and 3-manifolds which fiber over the circle*, volume 142 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1996.
- [Mc2] Curtis T. McMullen. Polynomial invariants for fibered 3-manifolds and Teichmüller geodesics for foliations. *Ann. Sci. École Norm. Sup. (4)*, 33(4):519–560, 2000.
- [Mk] Hiroyuki Minakawa. Examples of pseudo-Anosov homeomorphisms with small dilatations. *J. Math. Sci. Univ. Tokyo*, 13(2):95–111, 2006.
- [NZ] Walter D. Neumann and Don Zagier. Volumes of hyperbolic three-manifolds. *Topology*, 24(3):307–332, 1985.

- [Oe] Ulrich Oertel. Homology branched surfaces: Thurston’s norm on $H_2(M^3)$. In *Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984)*, volume 112 of *London Math. Soc. Lecture Note Ser.*, pages 253–272. Cambridge Univ. Press, Cambridge, 1986.
- [Ot] Jean-Pierre Otal. Le théorème d’hyperbolisation pour les variétés fibrées de dimension 3. *Astérisque*, (235):x+159, 1996.
- [Pe] R. C. Penner. Bounds on least dilatations. *Proc. Amer. Math. Soc.*, 113(2):443–450, 1991.
- [So] Won Taek Song. Upper and lower bounds for the minimal positive entropy of pure braids. *Bull. London Math. Soc.*, 37(2):224–229, 2005.
- [SKL] Won Taek Song, Ki Hyoung Ko, and Jérôme E. Los. Entropies of braids. *J. Knot Theory Ramifications*, 11(4):647–666, 2002. Knots 2000 Korea, Vol. 2 (Yongpyong).
- [Th1] William P. Thurston. The geometry and topology of 3-manifolds, 1980. Princeton University Notes.
- [Th2] William P. Thurston. A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.*, 59(339):i–vi and 99–130, 1986.
- [Ti] D. Tischler. On fibering certain foliated manifolds over S^1 . *Topology*, 9:153–154, 1970.
- [Ts] Chia-yen Tsai. The asymptotic behavior of least pseudo-Anosov dilatations. Preprint, arXiv:0810.0261.
- [Ve] Rupert Venzke. Braid forcing, hyperbolic geometry, and pseudo-Anosov sequences of low entropy. California Institute of Technology, Ph. D. Dissertation, available at <http://etd.caltech.edu/etd/available/etd-05292008-085545/>
- [Zh] A. Yu. Zhironov. On the minimum dilation of pseudo-Anosov diffeomorphisms of a double torus. *Uspekhi Mat. Nauk*, 50(1(301)):197–198, 1995.

Benson Farb
 Dept. of Mathematics, University of Chicago
 5734 University Ave.
 Chicago, IL 60637
 farb@math.uchicago.edu

Christopher J. Leininger
 Dept. of Mathematics, University of Illinois at Urbana-Champaign
 273 Altgeld Hall, 1409 W. Green St.
 Urbana, IL 61802
 clein@math.uiuc.edu

Dan Margalit
 Dept. of Mathematics, Tufts University
 503 Boston Ave
 Medford, MA 02155
 dan.margalit@tufts.edu