Topology and its Applications 180 (2015) 44–63

Contents lists available at ScienceDirect

Topology and its Applications

www.elsevier.com/locate/topol

Mapping tori of small dilatation expanding train-track maps

Yael Algom-Kfir *, Kasra Rafi

ARTICLE INFO

Article history:
Received 23 December 2013
Received in revised form 6 October 2014
Accepted 6 November 2014
Available online xxxx

Keywords:
$\mathrm{Out}(F_n)$
Mapping tori
Geometric group theory

ABSTRACT

An expanding train-track map on a graph of rank n is P-small if its dilatation is bounded above by $\sqrt[n]{\mathsf{P}}$. We prove that for every P there is a finite list of mapping tori $X_1, \dots, X_{\mathsf{A}}$, with A depending only on P and not n, so that the mapping torus associated with every P-small expanding train-track map can be obtained by surgery on some X_i. We also show that, given an integer $\mathsf{P} > 0$, there is a bound M depending only on P and not n, so that the fundamental group of the mapping torus of any P-small expanding train-track map has a presentation with less than M generators and M relations. We also provide some bounds for the smallest possible dilatation.

1. Introduction

Let G be a simplicial graph. A map $f\colon G \to G$ is a train-track map if vertices are mapped to vertices and, for every edge e and positive integer i, the i-th iterate $f^i(e)$ is an immersed edge path. The notion of a train-track map was first introduced by Bestvina and Handel [4] as a normal form for certain outer automorphisms of a free group. This is analogous to the train-track representation of a mapping class as developed by Thurston.

To a graph self-map, one can associate a *transition matrix* T_f (with size equal to the number of edges). We say f is an expanding train-track (ett) map if its transition matrix T_f is expanding (see Definition 2.3). Define the dilatation of f, λ_f, to be the largest modulus of an eigenvalue of T_f and define the rank of G, $\mathrm{rank}(G)$, to be the rank of fundamental group of G. Note that, $\mathrm{rank}(G)$ is not equal to the size of the matrix T_f.

Guided by analogies with mapping class group, we are interested in studying expanding train-track maps with small dilatation. As in the mapping class group setting, one expects the minimum possible dilation to converge to 1 as the rank of G goes to infinity. For a pseudo-Anosov mapping class $\phi\colon S \to S$, define

$$\underline{\lambda}^{pA}(n) = \min\{\lambda_\phi \mid \phi \text{ is a pseudo-Anosov map on a closed surface } S \text{ with } |\chi(S)| = n\}$$

* Corresponding author.
E-mail addresses: algomy@gmail.com (Y. Algom-Kfir), rafi@math.toronto.edu (K. Rafi).

http://dx.doi.org/10.1016/j.topol.2014.11.003
0166-8641/

[9,1,6] gave estimates for $\underline{\lambda}^{pA}(n)$

$$\frac{\log 2}{6} \le \liminf_n \bigl(n \log \underline{\lambda}^{pA}(n)\bigr) \le 2\log\left(\frac{3+\sqrt{5}}{2}\right).$$

We similarly define

$$\underline{\lambda}(n) = \min\bigl\{\lambda_f \bigm| f\colon\ G \to G \text{ is an ett and } \operatorname{rank}(G) = n\bigr\},$$

and, in Sections 6 and 7, we show that

$$\frac{\log 3}{3n-3} \le \log \underline{\lambda}(n) < \frac{\log 2}{n}. \tag{1}$$

The right-hand inequality is strict for every n (see examples in Section 6), however, we believe that it is asymptotically strict.

Conjecture 1.1. *For every n,*

$$\lim_{n\to\infty} n \log \underline{\lambda}(n) = \log 2.$$

In general, one would like to understand the structure of all ett maps with small dilatation. For $\mathsf{P} > 1$, we say an ett map $f\colon\ G \to G$ is P-small if

$$\lambda_f \le \sqrt[n]{\mathsf{P}}, \quad \text{where } n = \operatorname{rank}(G).$$

Similarly, a pseudo-Anosov map $\phi\colon\ S \to S$ on a closed surface S is P-small if

$$\lambda_\phi \le \sqrt[n]{\mathsf{P}}, \quad \text{where } n = \bigl|\chi(S)\bigr|.$$

We now consider the family of P-small maps as P is fixed and n goes to infinity. Our main theorem is an analogue of the following theorem of Farb–Leininger–Margalit.

Recall that a mapping torus associated with a map $f\colon\ X \to X$ is the space

$$M_f = X \times [0,1]/\sim \quad \text{with } (x,0) \sim \bigl(f(x),1\bigr).$$

By Van Kampen's theorem, this topological construction corresponds to an HNN extension of the fundamental group.

Theorem 1.2. *([5,2]) For each $\mathsf{P} > 1$ there exist finitely many complete, non-compact, hyperbolic 3-manifolds $M_1, \dots, M_r$ that fiber over $\mathbb{S}^1$, with the property that for any P-small pseudo-Anosov homeomorphism ϕ of any surface S, there exists a Dehn filling M_i' of M_i for some i, and a fiberation of $\Psi\colon\ M_i' \to \mathbb{S}^1$ such that ϕ is the monodromy of Ψ.*

We define a notion *surgery* for mapping tori of graphs and prove the following:

Theorem A. *For every $\mathsf{P} > 1$, there is a finite set of 2-complexes, which are mapping tori of self maps of graphs, $X_1, \dots, X_{\mathsf{A}}$, so that: If $f\colon\ G \to G$ is a P-small expanding train-track map on a graph G, then M_f is homeomorphic to a 2-complex that is obtained by surgery on some X_i.*

In particular, surgery does not alter the number of *essential 2-cells* in the mapping torus and only changes the structure of the edges. As a result we are able to prove a universal boundedness result on presentations of the fundamental groups of mapping tori associated with ett maps.

Theorem B. *There is a number* M *depending only on* P *(and not on* n*) so that if* $\phi \in \mathrm{Out}(F_n)$ *is a* P*-small ett automorphism then* Γ_ϕ *has a presentation with at most* M *generators and* M *relations.*

The paper is organized as follows. Section 2 provides background and sets notation for train-track maps of $\mathrm{Out}(F_n)$ elements, 2-complexes, and mapping tori. In Section 3 we describe how given a mapping torus we may replace its vertical graph with another to obtain a new mapping torus. Section 4 is dedicated to the proof of Theorem A. Section 5 is where Theorem B is proved. In Section 6 we prove the upper bound for Eq. (1) and in Section 7 we prove the lower bound.

2. Background

We review in Section 2.1 properties of graph maps and the outer automorphisms which they represent. We recall Definition 2.10 of a graph Δ_f called the *derivative of* f which will be useful for several purposes including the definition of surgery of mapping tori.

In Section 2.2 we set the notation for 2-complexes, and define the operations of subdivision and its inverse. We then define the notion of removing a sub-1-complex from the 1-skeleton of a 2-complex (Section 2.3).

Section 2.4 is devoted to describing a 2-complex structure of a mapping torus. We give a set of necessary and sufficient conditions for a 2-complex to be isomorphic to a mapping torus of a graph map.

2.1. Train-track maps and dilatation

Once and for all, fix a basis of the free group $\{x_1, \dots, x_n\}$. Let R be a graph with one vertex $*$, and n edges attached to it at both ends forming n loops. Identify the edges of R with the free basis.

Definition 2.1. A marked graph is a finite 1-complex G together with a homotopy class $[\tau]$ of homotopy equivalences $\tau\colon R \to G$.

Let Φ be an automorphism of the free group F_n. The automorphism Φ induces a map $f_\Phi\colon R \to R$ in the obvious way.

Definition 2.2. The map $f\colon G \to G$ is a *topological representative* of $\phi \in \mathrm{Out}(F_n)$ if: the f-image of any edge is a vertex or an immersed path beginning and ending at vertices, and for any homotopy inverse $\eta\colon G \to R$ of τ, the map $\mu \circ f \circ \tau$ is freely homotopic to f_Φ for $\Phi \in \phi$.

We are interested in M_f the mapping torus of f and its fundamental group. By the Van Kampen theorem, the fundamental group of the mapping torus M_f is an HNN extension of F_n, which can be presented as

$$\Gamma_f = F_n *_\phi = \langle x_1, \dots, x_n, t \mid t x_i t^{-1} = \Phi(x_i) \rangle$$

Observe that the universal cover $\widetilde{M_f}$ is contractible, since it is a graph of contractible spaces.

Let m be the number of edges in G. *The transition matrix* associated with f is an $m \times m$ matrix $M = (a_{ij})$ so that[1]

[1] Sometimes the transition matrix is defined as the transpose of this matrix.

$$a_{ij} \text{ is the number of times } f(e_i) \text{ crosses } e_j \text{ in either direction.}$$

Observe that M is a non-negative matrix. For $k \geq 1$, let a_{ij}^k be the (i, j) entry in M^k.

Definition 2.3. We say M is *expanding* if for all $1 \leq i, j \leq m$

$$\limsup_k a_{ij}^k = \infty.$$

Let λ be the largest-modulus of an eigenvalue of M. Perron–Frobenius theory states that λ is real and ≥ 1. Recall further that if $\lambda = 1$ then M is a permutation matrix and f is a homeomorphism. The number λ is called the *Perron–Frobenius eigenvalue* of f.

The following definition is due to Bestvina–Handel [4].

Definition 2.4. A topological representative f of ϕ is a *train-track map* if $f^k(e)$ is immersed for all edges $e \in G$ and all powers $k > 0$. The map f is an *expanding* train-track map (ett) if it is a train-track map, and its transition matrix is expanding.

We give an equivalent description of a train-track map. Endow the graph G with an orientation once and for all. For a directed edge e, we denote its initial vertex by $i(e)$ and its terminal vertex by $ter(e)$.

Definition 2.5. A pair of directed edges $\{e, e'\}$ is called a *turn* if $i(e) = i(e')$.

Definition 2.6. Let $Df(e)$ denote the first edge in the path $f(e)$. A turn $\{e, e'\}$ is *pre-degenerate* (PD) if $Df(e) = Df(e')$.

It follows from the work of Nielsen [8] and Stallings [10] that if f is not a homeomorphism then it has a PD turn. The map Df sends a turn to a turn.

Definition 2.7. A turn $\{e, e'\}$ is *illegal* if it is PD or if it is mapped to a PD turn under a positive iterate of Df, otherwise, it is *legal*. An edge path is legal if it crosses no illegal turns.

The proof of the following proposition is clear.

Proposition 2.8. *A topological representative* f: $G \to G$ *is a train-track map if and only if* $f(e)$ *is legal for every edge* e.

Definition 2.9. *A trap* is an oriented connected graph with the property that every vertex has a unique edge initiating from it. Topologically, a trap is a union of directed trees and a directed circle so that each tree is glued to the circle at a distinct vertex. The edges of the trees are directed towards the circle.

Given a graph map, its derivative graph (definition below) is an example of a union of traps.

Definition 2.10. Let f: $G \to G$ be a graph map that is a topological representative of an automorphism ϕ, and we assume that no edge is taken by f to a vertex. The derivative graph of f, Δ_f is constructed as follows: There is a vertex in Δ_f for each directed edge of G and an edge from e to e' if $Df(e) = e'$.

Proposition 2.11. *A topological representative* f: $G \to G$ *is a train-track map iff* $f^i(e)$ *does not contain a backtracking segment (an edge followed by its inverse) for* $1 \leq i \leq m$ *where* m *is the number of edges in* G.

Proof. We must show that if a turn does not become PD after $m-1$ iterations of f then it is legal. The graph Δ_f is a union of disjoint traps. In order to check if a turn $\{e, e'\}$ is illegal we start with $v_e, v_{e'} \in \Delta_f$ corresponding to the directed edges e, e' in G. We form the sequences a_i and b_i of vertices in Δ_f starting with $a_1 = v_e$, $b_1 = v_{e'}$ and a_i, b_i are the terminal vertices of the directed edges initiating at a_{i-1}, b_{i-1}. The sequences a_i and b_i are getting trapped by the oriented circles (the traps) of Δ_f. The sequences a_i and b_i reach a cycle (or two disjoint cycles) after no more than $m-1$ steps. If both a_i and b_i are not on the same directed circle or if they are on a directed circle and $a_i \neq b_i$ then they will never coincide. □

We deduce that deciding if a map has the train-track property is a finite check.

Definition 2.12. Let $\vec{v}$ be the unit length positive right eigenvector of M corresponding to the Perron–Frobenius eigenvalue λ_f i.e. $Mv = \lambda v$. *The natural metric* on G induced by a train-track map f is given by setting the length of the edge e_i to be v_i the i-th coordinate of $\vec{v}$.

When G is endowed with the natural metric $len(f(e)) = \lambda len(e)$ for every edge e in G. The next proposition shows that the dilatation can be associated with the automorphism represented by f.

Proposition 2.13. *([3]) If $\phi \in \operatorname{Out}(F_n)$ is an ett automorphism represented by a PF train-track map $f\colon G \to G$ then for any non-periodic conjugacy class w in F_n, and for any basis X of F_n:*

$$\log \lambda_f = \lim_{k\to\infty} \frac{\log |\phi^k(w)|_X}{k} \tag{2}$$

Where $|w|_X$ denotes the length in the basis X of the cyclically reduced word equivalent to w. In particular, λ_f is the same for all train-track representatives f of ϕ.

Sketch of the proof. This proof is essentially written down in [3], but we include it for completeness. Let w be a conjugacy class of a word, it is represented in G by an immersed loop α. For a closed path γ in G we denote by $f_\#(\gamma)$ the immersed loop homotopic to the loop $f(\gamma)$. Consider the infinite sequence $\alpha_k = f_\#^k(\alpha)$. Endow G with the natural metric, and consider the length of α_k. If it is bounded, then α is preperiodic (eventually periodic) because there are only finitely many loops whose length is smaller than any given length. Thus w is ϕ pre-periodic. Since ϕ is an automorphism, w preperiodic implies that it is periodic. We conclude that if w is not periodic, the length of α_k is unbounded.

Let $\mathrm{BCC}(f)$ the bounded cancellation constant of f (see [3]), i.e. let $\tilde{f}\colon \tilde{G} \to \tilde{G}$ be a lift of f to the universal cover of G, then for any geodesic $[x, y]$ whose length is greater than $\mathrm{BCC}(f)$, its image $f(x) \neq f(y)$. There is a threshold

$$\mathsf{T} = \frac{4\mathrm{BCC}(f)}{\lambda - 1}$$

such that if

(1) δ, β, γ are immersed segments in G,
(2) β is legal, and
(3) $len(\beta) > \mathsf{T}$

then

$$len\big(f_\#^k(\delta \cdot \beta \cdot \gamma)\big) \geq c\lambda^k$$

The reason is that the middle segment $f_\#^k(\beta)$ grows like λ^k, and the threshold T is large enough so that (3) guarantees that when we reduce $f_\#^k(\delta)f_\#^k(\beta)f_\#^k(\gamma)$ the cancelled segments do not affect the exponential growth. See [3] for more details.

If w is not periodic, then the length of α_k grows to infinity. Since f is a train-track map, the number of illegal turns in α_k is non-increasing. Therefore, there is some j so that $f_\#^j(\alpha)$ contains a legal segment of length greater than T. Thus, $len(f_\#^k(\alpha)) \geq c\lambda^k$ for $k \geq j$ (we absorbed $\lambda^{-j} len(f^j(w))$ in c). The word length with respect to the basis X of F_n, is quasi-isometric to lengths of immersed paths in G. Thus, up to a multiplicative error $|\phi^k(w)|_X \asymp \lambda^k$ hence the limit on the right of Eq. (2) is $\log \lambda_f$. $\square$

2.2. Subdivision of a 2-complex and its inverse operation

We give the usual definition of a 2-complex.

Definition 2.14. A *2-complex* X is a tuple $(\mathcal{V}, \mathcal{E}, \mathcal{C}, \mu)$. The set $\mathcal{V}$ is the set of vertices, $\mathcal{E}$ is the set of edges, $\mathcal{C}$ is the set of 2-cells each of which is a polygonal subset of $\mathbb{R}^2$. The set μ is a collection of maps, called gluing maps. It is the union of two sets: gluing maps of edges $\mu_\mathcal{E} = \{\mu_e \mid e \in \mathcal{E}\}$, and gluing maps of cells $\mu_\mathcal{C} = \{\mu_c \mid c \in \mathcal{C}\}$. For every $e \in \mathcal{E}$, $\mu_e\colon \partial e \to \mathcal{V}$ is called a gluing map of e. The 1-skeleton of X is the graph

$$X^{(1)} = \bigcup_{e\in\mathcal{E}} e \cup_{\mu_\mathcal{E}} \mathcal{V}$$

For $c \in \mathcal{C}$, $\mu_c\colon \partial c \to X^{(1)}$ is called the gluing map of $c \in \mathcal{C}$. We require that when $\mu_{c_1}(e) \cap \mu_{c_2}(e) \neq \emptyset$ then $\mu_{c_1}^{-1} \circ \mu_{c_2}$ is piecewise affine. These maps can be defined by specifying a directed path in $X^{(1)}$ for every edge of ∂c. The total space is

$$X = \bigcup_{c\in\mathcal{C}} c \cup_{\mu_\mathcal{C}} X^{(1)}$$

We will preform subdivision of cells as follows:

Definition 2.15 *(Subdivision).* Let c be a 2-cell of X so that ∂c contains more than 3 vertices. One may subdivide c into two cells without changing the number of vertices by adding a diagonal e between two non-adjacent vertices v, w to the set $\mathcal{E}$, and replacing $c \in \mathcal{C}$ with two cells c_1, c_2 formed by subdividing c along the diagonal from v to w. We say that this complex is obtained from X by subdividing c along (v, w).

Definition 2.16. A *side* of a 2-cell in X is a pair $s = (c, \alpha)$ where α is an edge of ∂c in $\mathbb{R}^2$.

Definition 2.17. Let v be a vertex in X, it is inessential if $\mathrm{val}(v) = 2$ and there is no edge of X that is adjacent to v at both ends then v is inessential.

Definition 2.18. Suppose e is an edge in $X^{(1)}$ so that $\mu_\mathcal{C}^{-1}(e)$ is the union of exactly two subintervals J_1, J_2 with $J_1 \subseteq (c_1, \alpha_1)$ and $J_2 \subseteq (c_2, \alpha_2)$ and $c_1 \neq c_2$. We say that e is an *inessential* edge.

Definition 2.19 *(Removing an edge).* If X is a 2-complex and $e \in \mathcal{E}$ is an inessential edge we can glue the cells whose boundary contains J_1, J_2 and remove e from the set of edges. We will say that the new complex is obtained from X by removing the removable edge e.

2.3. Cutting a 2-complex along a graph

We define the operation of cutting a 2-complex X along a 1-sub-complex U. We sketch an example in Fig. 1.

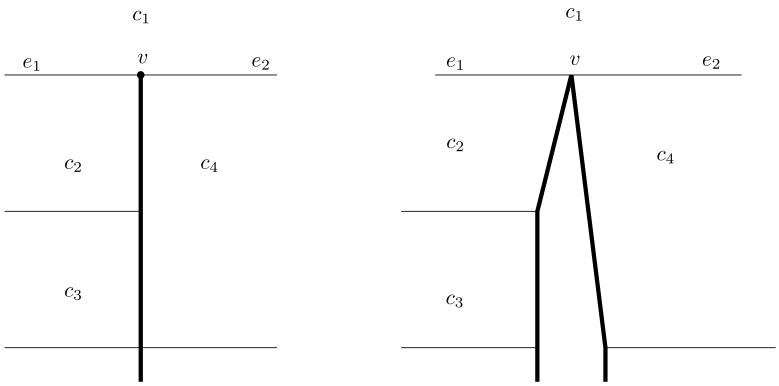

Fig. 1. Cutting a 2-complex along a graph. The heavier edges are the edges of U. When we cut along U we duplicate the edges of U so they correspond to the sides mapping to them. Note that the left side of c_4 is mapped to a single edge in X_U whereas it mapped to a multiple edge path in X. The vertices of U correspond to pre-vertices after gluing two if they are adjacent to a common non-U-edge. Notice that the pre-vertices corresponding to v in c_2 and v in c_4 are equivalent since they are adjacent in the cell c_1.

Definition 2.20. Let X be a 2-complex and U a 1-sub-complex, U is *admissible* if for each side (c, α) and any parameterization $\rho\colon I \to \alpha$,

$$\mu_c \circ \rho(t_0) \in U \quad \text{for some } t_0 \in I \quad \Longrightarrow \quad \mu_c \circ \rho(t) \in U \quad \text{for all } t \in U$$

By subdividing the sides of the polygon, e.g. turning a square to a pentagon, we can turn any 1-sub-complex into an admissible one. We require this condition for ease of description of the complex cut along U.

Definition 2.21. A pre-vertex of a 2-complex is a pair (c, t) so that $c \in \mathcal{C}$ and $t \in \partial c$ and $\mu_c(t)$ is a vertex. A pre-edge is a pair (c, σ) with σ a segment of ∂c so that $\mu_c(\sigma)$ is a single edge of $X^{(1)}$.

Definition 2.22. We define an equivalence relation on the pre-vertices by the following defining relations: The pre-vertices (c_1, t), (c_2, s) are equivalent $(c_1, t) \sim (c_2, s)$ if:

(1) $\mu_{c_1}(t) = \mu_{c_2}(s) = v$ and,
(2) if $v \in U$ we further require that there exist pre-edges (c_1, α_1), (c_2, α_2) and parameterizations $\rho_i\colon I \to \alpha_i$ so that $\rho_1(0) = t$, $\rho_2(0) = s$ and for all $r \in I$,

$$\mu_{c_1} \circ \rho_1(r) = \mu_{c_2} \circ \rho_2(r) \notin U.$$

We give another description of the equivalence classes:

Proposition 2.23. *There is a bijective correspondence between equivalence classes of pre-vertices and the set:*

$$\mathcal{V}_U = \left(X^{(0)} \cap U^C\right) \cup \left\{(u, e) \mid u \in U \cap X^{(0)}, e \in X^{(1)} \cap U^C, u \in \partial e\right\} / \sim$$

where U^C denotes the complement of U and $\sim$ is generated by the relation: $(u, e) \sim (u', e')$ if $u = u'$ and there exists a cell c and an interval $J \in \partial c$ so that $\mu_c(J) = ee'$.

Proof. We define the correspondence

$$\mathcal{F}\colon \left\{(c, t) \mid c \in X^{(2)}, t \in \partial c, \mu_c(t) \in X^{(0)}\right\} \to \left\{(u, e) \mid u \in U \cap X^{(0)}, e \in X^{(1)} \cap U^C, u \in \partial e\right\}$$

as follows. If $\mu_c(t) \in U^C$ then $\mathcal{F}(c,t) = \mu_c(t)$. Otherwise, let (c,α) be a non-U-side so that $t \in \partial\alpha$. Define $\mathcal{F}(c,t) = (\mu_c(t), \mu_c[\alpha])$. This map is surjective and descends to a map from equivalence classes of pre-vertices to $\mathcal{V}_U$. Since generating relations map to generating relations this map is injective. □

Definition 2.24 *(Cutting along a graph).* Let X be a 2-complex, U a sub-1-complex, the complex X_U obtained from X by cutting along U is the 2-complex

$$(\mathcal{V}_U, \mathcal{E}_U, \mathcal{C}_U, \{\mu^U\}).$$

The vertex set $\mathcal{V}_U$ is the set of equivalence classes of vertices by the relation defined in Definition 2.22. The edge set is defined by

$$\mathcal{E}_U = (X^{(1)} \cap U^C) \cup \{(c,\alpha) \mid \mu_c[\alpha] \in X^{(1)} \cap U\}.$$

The gluing maps μ_e^U, μ_c^U are the obvious ones. See Fig. 1 for an example.

There is a cellular quotient map $\chi : X_U \to X$, that restricts to the identity on the interiors of the 2-cells, and on the edges in $X^{(1)}$. For $[c,\alpha] \in \mathcal{E}_U$ so that $\mu_c[\alpha] \subset U$ we have $\chi[c,\alpha] = \mu_c[\alpha]$. Similarly for $[c,t] \in \mathcal{V}_U$ with $\mu_c(t) \in U$, $\chi[c,t] = \mu_c(t)$.

Definition 2.25. The pre-U-graph in X_U is $K := \chi^{-1}(U)$. It consists of vertices $[c,t]$ with $c \in \mathcal{C}(X)$, $t \in \partial X$ and $\mu_c(t) \in U$; and edges $[c,\alpha]$ with $c \in \mathcal{C}(X)$, and $\mu_c(\alpha) \subset U$.

2.4. Mapping tori

Let G be a graph. The mapping torus of a map $f\colon G \to G$ is

$$M_f = G \times [0,1]/\sim$$

where the equivalence is generated by the relations $(x,0) \sim (f(x),1)$. We describe the 2-complex structure on M_f.

(1) The set of vertices $\mathcal{V}$ equals the set of vertices of G.
(2) The set of edges $\mathcal{E}$ is the disjoint union of: the set of edges of G which we call horizontal edges, and denote it by HE; and its complement that we call the set of vertical edges and denote by VE. Each vertical edge has the form $(v, f(v))$ for $v \in \mathcal{V}$. We orient a vertical edge from v to $f(v)$. We orient the edges of G arbitrarily.
(3) There is one square c for every $e \in \mathrm{HE}$. We describe the gluing map of c. The top edge of c is mapped to e, the left side l of c is mapped to the vertical edge $(i(e), f(i(e)))$, and the right side r of c is glued to the vertical edge $(ter(e), f(ter(e)))$. The bottom edge is glued to the edge path $f(e)$.

It is straight-forward to check if a 2-complex is a mapping torus, and we give the relevant conditions in the next proposition.

Proposition 2.26. *The 2-complex X is a mapping torus of a map $f\colon G \to G$ iff:*

(1) *The set of edges $\mathcal{E}$ of X may be partitioned into two sets: the set of vertical edges* VE *and the set of horizontal edges* HE*.*

(2) *We define the graph on the vertical edges:*

$$U = \bigcup_{e \in \mathrm{VE}} e \cup_\mu \mathcal{V}.$$

This graph may be endowed with an orientation to make it a union of traps. Equivalently, there is a bijection $a\colon \mathcal{V} \to \mathrm{VE}$.

(3) *There is a bijection* $\mathrm{top}\colon \mathcal{C} \to \mathrm{HE}$.

(4) *We define the horizontal graph by*

$$G = \bigcup_{e \in \mathrm{HE}} e \cup_\mu \mathcal{V}.$$

Choose an orientation on G. For $c \in \mathcal{C}$, the map μ_c is given by the path $\mathrm{top}(c)rw^{-1}l^{-1}$ *where r, l are positively oriented edges in U and w is an edge path in G. We denote the edge path w by* $\mathrm{bot}(c)$.

Proof. If X satisfies items (1)–(4) we define the map $f\colon G \to G$ that takes each vertex $v \in G$ to the terminal endpoint of $a(v)$ defined in item (2), and for e an edge in G we let $f(e) = \mathrm{bot}(c)$ where c is the cell so that $\mathrm{top}(c) = e$. We check that f is well defined: since $erw^{-1}l^{-1}$ is a connected path, $ter(r) = ter(w)$ and $ter(l) = i(w)$, hence $f(ter(e)) = ter(f(e))$ and $f(i(e)) = i(f(e))$. The maps given in (2), (4) are exactly those induced by f. □

3. Surgery of mapping tori

In this section we describe an operation where we excise the vertical graph from a mapping torus and glue in a new graph in such a way that yields a new mapping torus. We shall call this operation surgery.

Definition 3.1. Let $f\colon G \to G$ be a map, we call and edge e *mixed* if its image is a concatenation of more than one edge. We call an edge e *dynamic* if it is contained in the image of a mixed edge, or if there is more than one edge that maps onto it.

Definition 3.2. We define an equivalence relation on the vertices of G, generated by: $v \sim w$ if $f(v) = w$. We define an equivalence relation on the edges of G generated by $e \sim f(e)$ if $f(e)$ is a single edge that is non-dynamic. An equivalence class of edges will be called a *stack*.

Each stack has the form $\varepsilon = \{e, f(e), f^2(e), \dots, f^s(e)\}$ for $s \geq 0$ where e is a dynamic edge and $f^s(e)$ is possibly a mixed edge. We call $f^s(e)$ *the bottom edge* of the stack ε. Note that e is the only dynamic edge in ε and $f^i(e)$ for $i = 0, \dots, s-1$ are not mixed. The quotient of G under these equivalences is a graph denoted by Q, and the quotient map will be denoted by $p\colon G \to Q$.

Definition 3.3. *The archetype* of f is a map $f_Q : Q \to Q$ defined as follows: f_Q fixes every vertex of Q, and $f_Q(\varepsilon) = p(f(e))$ where e is the bottom edge of ε.

Remark 3.4. The archetype of an expanding train-track map need not be a train-track map, need not be expanding and might not even be a homotopy equivalence. The most that can be said is that $f_Q : Q \to Q$ is a composition of a sequence of Stallings folds, i.e. homotopy equivalences that identify a pair of adjacent edges, and a pinch map, i.e. a map that is a homeomorphism on the graph minus its vertices.

We define *a surgery of a mapping torus*. Let X be a mapping torus of the map $f\colon G \to G$. Let U be the vertical subgraph, as in Proposition 2.26(2). Every connected component of U is an oriented graph, in fact a trap, whose vertices are the elements of an equivalence class described in Definition 3.2.

We want to remove U from X, glue in a different graph S, and redivide the 2-cells so that the output is again a mapping torus. We first remove redundant edges of X. For each non-dynamic edge, e of G there are exactly two 2-cells c_e, $c_{f^{-1}(e)}$ that are attached to e. Therefore e is an inessential edge and we may remove e (Definition 2.19). Additionally if removing an inessential edge creates an inessential vertex, then we remove the vertex as well.

By repeating this procedure for all the non-dynamic edges in a single stack ε, one obtains a single 2-cell R_ε for every stack ε, we call it the rectangle corresponding to ε. The attaching map μ_{R_ε} sends the top edge of R_ε to the dynamic edge e in ε, the bottom edge of ε to $f(e')$ for $e' = f^s(e)$ the bottom edge of ε. The right side of R_ε is sent to the edge path from $ter(e)$ to $f(ter(e'))$ and the left side of R_ε is mapped from $i(e)$ to $f(i(e'))$.

Definition 3.5. The 2-complex obtained from X by removing all the non-dynamic edges is called the *the floor-plan* of X and denoted $\mathring{X}$.

Proposition 3.6. *Let X^f be the mapping torus of f and X^{f_Q} the mapping torus of f_Q. Let U be vertical graph in X^f, and W the vertical graph in X^{f_Q} then*

$$\mathring{X}^f_U = X^{f_Q}_W$$

Proof. Let $\varepsilon = \{e, \dots, e' = f^s(e)\}$ be a stack with e the dynamic edge and e' the bottom edge of ε. There is a bijective correspondence between the stack ε and its dynamic edge: $D(\varepsilon) = e$.

Let us consider the cell structure of X^{f_Q}: There is one 2-cell c_ε for each stack ε, $\mathrm{top}(c_\varepsilon) = \varepsilon$ and $\mathrm{bot}(c_\varepsilon) = [f(e')]$. In $\mathring{X}^f$ there is one 2-cell $R_{[e]}$ for each dynamic edge e, and $\mathrm{top}(R_{[e]}) = e$, $\mathrm{bot}(R_{[e]}) = f(e')$. Define the map

$$D\colon \{\text{2-cells of } X^{f_Q}\} \to \{\text{2-cells of } \mathring{X}^f\} \tag{3}$$

by setting $D(c_\varepsilon) = R_\varepsilon$. There is also a correspondence between the horizontal edges: $D(\varepsilon) = e$. We now consider $X^{f_Q}_W$ and $\mathring{X}^f_U$. Eq. (3) defines a correspondence between the sets of 2-cells. There is also a bijective correspondence between the edges coming from horizontal edges in the two complexes. Let l_ε denote the left side of the rectangle c_ε, let $l_{\bar{\varepsilon}}$ be the right side. Let l_e be the right side of R_ε and $l_{\bar{e}}$ be the right side. We set $D(l_\varepsilon) = l_e$ and $D(l_{\bar{\varepsilon}}) = l_{\bar{e}}$. This shows that D defines a bijective correspondence between the vertical edges of $X^{f_Q}_W$ and $\mathring{X}^f_U$. Notice that:

$$D(\mathrm{top}(c_\varepsilon)) = \mathrm{top}(D(c_\varepsilon)) \qquad D(\mathrm{bot}(c_\varepsilon)) = \mathrm{bot}(D(c_\varepsilon)) \qquad D(l_\varepsilon) = l_e \quad D(l_{\bar{\varepsilon}}) = l_{\bar{e}}.$$

Moreover, note that the equivalence of the vertices in X^{f_Q}, $\mathring{X}^f$ (Proposition 2.23) defining the vertices of $X^{f_Q}_W$, $\mathring{X}^f_U$ is completely determined by the top and bottom gluing maps of the cells c_ε, R_ε. Thus, D descends to a map from $X^{f_Q}_W \to X^f_U$. □

Let $\mathring{X}_U$ be $\mathring{X}$ cut along U, see Definition 2.24. Recall that K denotes the pre-image of U under the quotient map $\chi\colon \mathring{X}_U \to \mathring{X}$, see Definition 2.25. The next lemma will help us to determine K in our running example. Recall the derivative graph in Definition 2.10.

Proposition 3.7. *The graph K is isomorphic to the quotient of Δ_{f_Q} defined by the equivalence relation generated by the following condition: The vertices v_ε, $v_{\varepsilon'}$ corresponding to the edges ε, ε', are equivalent if there exists an edge ε'' so that*

$$f_Q(\varepsilon'') = \dots \bar{\varepsilon}\varepsilon' \dots$$

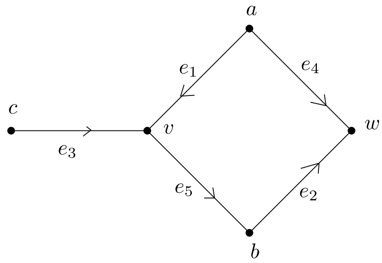

Fig. 2. Consider the graph above with the graph map f defined by $e_i \mapsto e_{i+1}$ for $i = 1, \dots, 4$ and $e_5 \mapsto e_3\bar{e_1}e_4$. The stacks corresponding to it are $x = \{e_1, e_2\}$, $y = \{e_3\}$, $z = \{e_4, e_5\}$.

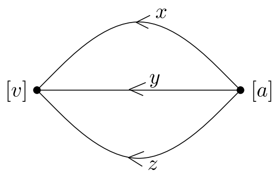

Fig. 3. The quotient graph Q of the map described in Fig. 2. The archetype is the map $f_Q\colon Q \to Q$ so that $x \mapsto x$, $y \mapsto y$ and $z \mapsto yXz$.

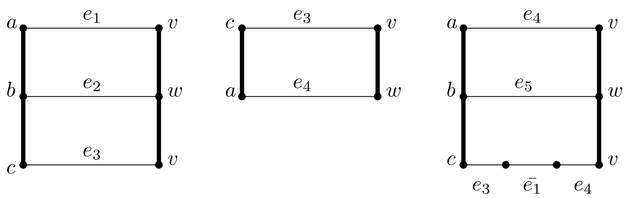

Fig. 4. The mapping torus of the map of Fig. 2. The heavy edges mark the vertical subgraph. The floor plan of X is obtained by removing all of the middle edges from the rectangles above.

Proof. By Proposition 3.6, the 2-complexes $\mathring{X}_U^f$ and $X_W^{f_Q}$ are isomorphic, and the pre-images of the vertical graphs in U, W are isomorphic. So we will denote by K the pre-image of the vertical graph in $X_W^{f_Q}$. The edges of K are the left and right sides of the cells c_ε. We denote them by $l_\varepsilon, r_\varepsilon$ accordingly. Define a map

$$\begin{aligned} \widetilde{P}\colon \quad & V(\Delta_{f_Q}) \to \{\text{pre-vertices of } X^{f_Q}\} \\ & v_\varepsilon \mapsto (c_\varepsilon, i(\varepsilon)). \end{aligned}$$

We post-compose with the quotient map of equivalence relation of Definition 2.22 to get a map

$$P\colon V(\Delta_{f_Q}) \to \mathcal{V}\big(X_W^{f_Q}\big)$$

This map induces a map on the edges because if $(v_\varepsilon, v_{\varepsilon'})$ is an edge in Δ_{f_Q} then $f_Q(\varepsilon) = \varepsilon'$ and thus the left side of c_ε is an edge from $i(\varepsilon)$ to $i(\varepsilon')$ in $X_W^{f_Q}$. Furthermore, P is injective on the edges.

Since $\widetilde{P}$ is injective, if $P(v_\varepsilon) = P(v_{\varepsilon'})$ then $(c_\varepsilon, i(\varepsilon)) \sim (c_{\varepsilon'}, i(\varepsilon'))$. This equivalence is generated by relations of the type $\mu_{\varepsilon''}(\text{bot}(c_{\varepsilon''})) \supset \bar{\varepsilon}\varepsilon'$. The statement follows. □

For the example of Fig. 2, whose archetype is described in Fig. 3, the mapping torus is the space described in Fig. 4. The graph Δ_{f_Q} has six vertices: x, y, z, $\bar{x}$, $\bar{y}$, $\bar{z}$, see Fig. 5. Proposition 3.7 implies that in K, $\bar{y} \sim \bar{x}$ and $x \sim z$ since $f(z) = y\bar{x}z$. We denote the edges of $\mathring{X}_U^f$ corresponding to the left side of R_ε by l_ε and the right side by $l_{\bar{\varepsilon}}$. In the example, K has two connected components depicted in Fig. 6.

We now wish to glue $\mathring{X}_U$ to a different graph S via a map $\psi\colon K \to S$.

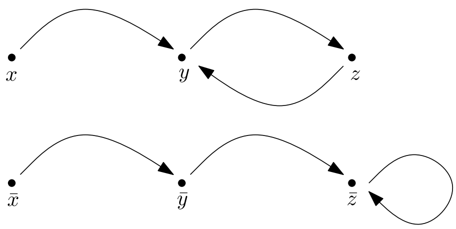

Fig. 5. The derivative graph Δ_{f_Q} of f_Q. The graph K is a quotient of this by $x \sim z$ and $\bar{x} \sim \bar{y}$.

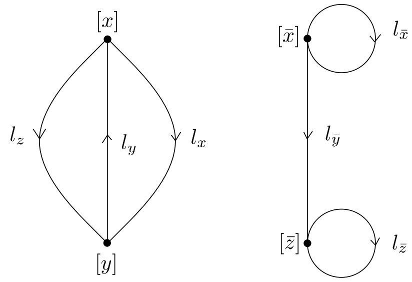

Fig. 6. The K-graph, i.e. pre-image of the vertical graph.

Definition 3.8. An oriented graph S is called *admissible* if it is a union of traps.

We shall denote by $M(S)$ the set of midpoints of edges of S.

Definition 3.9. Let X be a mapping torus of f, U its vertical subgraph and let S be an admissible graph. A map $\psi : K \to S$ is a *filling* of (X, U) if:

(1) ψ takes vertices to vertices, and maps edges to edge paths,
(2) if l_ε, $l_{\bar{\varepsilon}}$ are the vertical edges in the same rectangle R_ε in X_U, then

$$\#\psi|_{l_\varepsilon}^{-1}\big(M(S)\big) = \#\psi|_{l_{\bar{\varepsilon}}}^{-1}\big(M(S)\big)$$

We define the *hight* of R_ε to be $ht_\psi(R_\varepsilon) := \#\psi|_{l_\varepsilon}^{-1}(M(S))$. We now define the surgery operation using the data (X, U, ψ).

Definition 3.10. Let X be a mapping torus of a map f: $G \to G$, and let U be the vertical graph. Let S be an oriented admissible graph and ψ a filling of (X, U). The complex obtained from X by ψ-filling is denoted $X(U, S, \psi)$ and it is constructed from

$$\mathring{X}_U \cup_\psi S$$

by adding $ht_\psi(R_\varepsilon) - 1$ parallel edges between l_ε and $l_{\bar{\varepsilon}}$ in R_ε connecting corresponding ψ pre-images of $V(S)$.

Example 3.11. We continue with the running example. We let S be the graph with two connected components S_1, S_2, S_1 is a 3-edge circle, S_2 is a 2-edge circle. We describe a map ψ in Fig. 7.

The 2-cells of the 2-complex $X(U, S, \psi)$ are depicted in Fig. 8. They are glued to S via the map ψ in Fig. 7.

The 2-complex is a mapping torus of the map in Fig. 9.

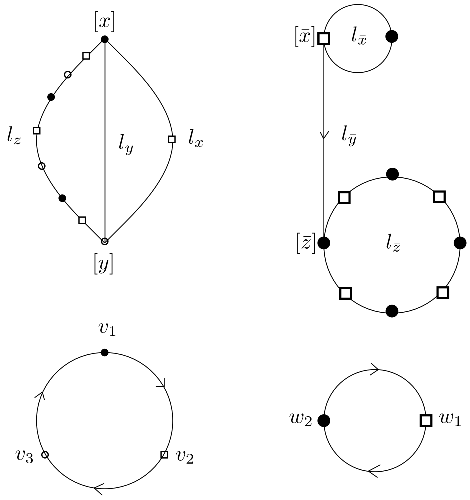

Fig. 7. The graph S is the union of the two circles on the bottom of the figure. The filling map $\psi\colon\ K \to S$ is the map sending the marked points in K_i to the corresponding vertices of S_i, and extends linearly on the edges between the marked points. Notice that for $a = x, y, z$ the height of l_a is equal to that of $l_{\bar{a}}$. Thus ψ satisfies the conditions for a filling.

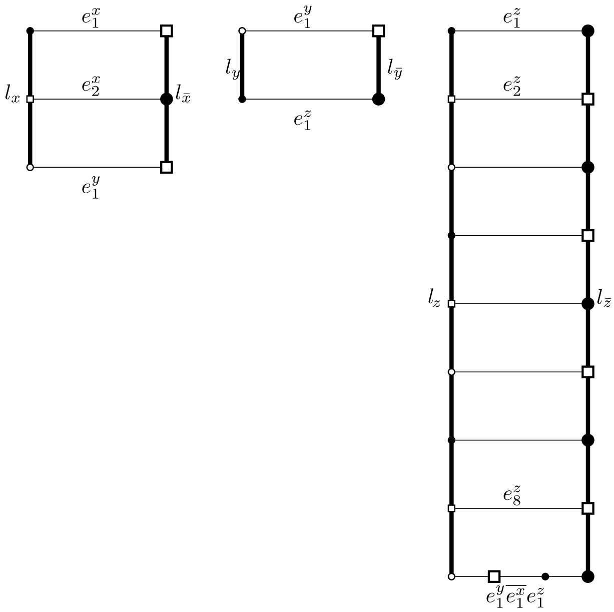

Fig. 8. The 2-complex $X(U, S, \psi)$.

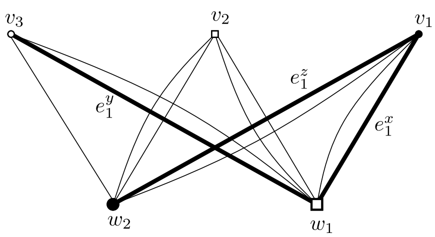

Fig. 9. This is the horizontal graph of the mapping torus of Fig. 8. The map is $e_1^x \to e_2^x \to e_1^y$, $e_1^y \to e_1^z$, and $e_1^z \to e_2^z \to e_3^z \to e_4^z \to e_5^z \to e_6^z \to e_7^z \to e_8^z \to e_1^y\overline{e_1^x}e_z^1$. The heavy edges denote the edges e_1^x, e_1^y and e_1^z that are the new dynamic edges.

Recall the map $\chi\colon \mathring{X}_U \to \mathring{X}$ in Definition 2.24 and notice that it restricts to a homeomorphism on $\mathring{X}_U - K$. Denote by $\tau\colon K \to U$ the restriction of χ to K.

Proposition 3.12. *For any mapping torus X,*

(1) $X(U, U, \tau) = X$
(2) *If $Z = X(U, S, \psi)$ then $X = Z(S, U, \tau)$.*

Proof. To prove (1), notice that $\mathring{X} = \mathring{X}_U \cup_\tau U$. The original complex X is obtained from $\mathring{X}$ by subdividing each c_e into $ht(c_e) = |\varepsilon|$ cells. Therefore, $X = X(U, U, \tau)$. To see (2), note that $\mathring{X}_U = \mathring{Z}_S$, therefore $X = X(U, U, \tau) = Z(S, U, \tau)$. □

Proposition 3.13. *If X, U, S, ψ are as above then $X(U, S, \psi)$ is a mapping torus.*

Proof. We need to check properties 1–4 in Proposition 2.26. The vertical graph of $X(U, S, \psi)$ is S, the other edges are the horizontal edges. By assumption, S is a union of traps so it satisfies (2) in Proposition 2.26. To check (3): notice that there is a bijection from the set of cells of X to the horizontal edges. After removing the non-dynamic edges, there is a bijection from the set of rectangles $\mathcal{R}$ to the edges of $\mathring{X}$ not in U. This induces a bijection from the rectangles in X_U to the edges outside K, which survives after gluing K along S. So there is a bijection between the 2-cells of $\mathring{X}_U \cup_\psi S$ and the edges outside of S. Subdivision adds to each new cell and a new top edge for it. So at the end of the subdivision process, there is a bijection top: $\mathcal{C} \to \mathrm{HE}$. Property (4) of Proposition 2.26 requires that if $\mathrm{top}(c) r w^{-1} l^{-1}$ is the boundary map of ∂c then r, l are edges and w is an edge path. After adding the parallel horizontal edges each vertical side is mapped to a single edge of S. Moreover, the bottom paths of the cells are either a single new edge or one of the "old" bottom paths of the complex X. Hence w is an edge path and the proof is complete. □

Proposition 3.14. *Given X a mapping torus of $f\colon G \to G$ and U the vertical graph, let S be a union of $I = |\mathcal{V}_Q|$ disjoint 1-edge circles. Let $\psi\colon K \to S$ be the map sending each edge of K_i to the single edge of S_i. Then $X(U, S, \psi)$ is the mapping torus of f_Q.*

Proof. This follows from Proposition 3.6. □

4. Finiteness

We start by proving that if f is P-small then the quotient graph Q and the archetype of f are uniformly finite.

Definition 4.1. Let $\mathsf{B} > 0$ a map is B-bounded if the image of every edge is a concatenation of at most B edges.

Proposition 4.2. *For every P there are positive integers E, B so that if f is P-small then*

(1) *Q has at most E edges and vertices.*
(2) *f_Q is B-bounded.*

The following proof is a slightly modified version of the proof appearing in [5]. We repeat it in order to note that the proofs hold when the transition matrix T_f is expanding and not Perron–Frobenius as in [7] and [5].

Let A be an $m \times m$ expanding matrix. Define a graph Γ with m vertices and a_{ij} edges from v_i to v_j. Let $d_{\text{out}}(v_i)$ as the number of edges coming out of v_i, i.e. the row sum of the i-th row. Let $d_{\text{in}}(v)$ be the number of edges coming into v_i, i.e. the column sum of the i-th column.

We denote by a^k_{ij} the i, j entry of A^k. We note that a^k_{ij} is equal to the number of directed paths from v_i to v_j of length equal to k.

Proposition 4.3. *Let A be an expanding matrix with PF eigenvalue λ then,*

$$\min_i \sum_j a_{ij} \leq \lambda.$$

Additionally, for each integer $k > 1$,

$$\min_i \sum_j a^k_{ij} \leq \lambda^k$$

Proof. Assume for contradiction that $\min_i \sum_j a_{ij} = \mu > \lambda$. Let $\mathbf{x} = [x_1, \ldots, x_m]$ be vector with positive entries and let $\mathbf{y} = [y_1, \ldots, y_m] = \mathbf{x}A$. Then $\mathbf{y}$ is positive and $\sum_i y_i \geq \mu \sum_i x_i$. Which means, for $k \geq 1$, $|\mathbf{x}A^k|_{L^1} > \mu^k |\mathbf{x}|_{L^1}$. This is a contradiction.

The second assertion follows from the fact that a power of an expanding matrix is itself expanding. □

Lemma 4.4. *([7,5]) Let A be a non-negative expanding $m \times m$ matrix with PF eigenvalue λ then*

$$1 + \sum_{v \in \Gamma} \bigl(d_{\text{out}}(v) - 1\bigr) = 1 + \sum_{v \in \Gamma} \bigl(d_{\text{in}}(v) - 1\bigr) \leq \lambda^m \tag{4}$$

Proof. Let $E(\Gamma)$ be the set of edges, then

$$\sum_{v \in \Gamma} d_{\text{in}}(v) = \bigl|E(\Gamma)\bigr| = \sum_{v \in \Gamma} d_{\text{out}}(v)$$

this establishes the first equality in (4).

Let i be the index of the smallest row sum of A^m. Since A is expanding, for each j there is a path from v_i to v_j. Therefore, there is a directed tree in Γ rooted at v_i, with edges directed away from v_i, that contains all of the vertices of Γ. We denote this tree by $T(v_i)$. The tree $T(v_i)$ contains exactly $m-1$ edges. Therefore,

$$1 + \sum_{v \in \Gamma} \bigl(d_{\text{out}}(v) - 1\bigr) = 1 + \sum_{v \in \Gamma} d_{\text{out}}(v) - m = \bigl|E(\Gamma)\bigr| - \bigl|E\bigl(T(v_i)\bigr)\bigr| \tag{5}$$

Consider the map

$$F\colon \{\text{paths of length } m \text{ initiating at } v_i\} \to E(G) - E\bigl(T(v_i)\bigr)$$

defined by letting $F(\alpha)$ be the first edge in the edge path α that is not in $T(v_i)$. Note that since the length of α is m, and $T(v_i)$ is a tree, there must be an edge in α that is not contained in $T(v_i)$. Moreover, since A is expanding, F is onto. This implies that

$$\bigl|E(\Gamma)\bigr| - \bigl|E\bigl(T(v_i)\bigr)\bigr| \leq \bigl|\{\text{paths of length } m \text{ initiating at } v_i\}\bigr| = \sum_j a^m_{ij} \tag{6}$$

By our choice of i and by Proposition 4.3, $\sum_j a^m_{ij} \leq \lambda^m$. Together with Eqs. (5) and (6) this implies the lemma. □

Proof of Proposition 4.2. If f is P-small then $\lambda_f^n \le \mathsf{P}$. Since $m \le 3n-3$ then, $\lambda^m \le \mathsf{P}^3$. Let T_f be the transition matrix of f, and Γ_f the adjacency graph constructed above. By Lemma 4.4 we have,

$$\sum_{v\in\Gamma_f}\big(d_{\text{out}}(v)-1\big) \le \mathsf{P}^3-1.$$

Thus for each v, $d_{\text{out}}(v) \le \mathsf{P}$ and letting $\mathsf{B} := \mathsf{P}^3$ we have that for each edge e, $|f(e)| \le \mathsf{B}$. A vertex $v \in \Gamma_f$ has $d_{\text{out}}(v) > 1$ iff v corresponds to a mixed edge of f. Thus the number of mixed edges is $\le \mathsf{P}^3$. Thus, number of edges that are images of mixed edges is bounded by P^6.

If e is an edge in G such that there are distinct edges e', e'' in G so that $f(e') = f(e'') = e$ then the vertex corresponding to e in Γ_f has $d_{\text{in}}(v_e) > 1$ so the number of such vertices is $\le \mathsf{P}^3$.

The set of dynamic edges is a union of images of mixed edges and edges with a preimage containing more than one edge. The number of dynamic edges is thus bounded by $\mathsf{E} := \mathsf{P}^6 + \mathsf{P}^3$. There is a bijection between edges in Q and dynamic edges of G, thus the number of edges in G is bounded above by E. □

Proposition 4.5. *For every* $\mathsf{P} > 1$ *there is a constant* A *and maps* $g_1, \dots, g_{\mathsf{A}}$ *such that for every* P*-small map* f *its archetype is one of the maps* $g_1, \dots, g_{\mathsf{A}}$*.*

Proof. When f is P-small, f_Q is B-bounded by Proposition 4.2. There are only finitely many graphs with less than E edges, and there are only finitely many self maps of those graphs that are B-bounded. We list these maps as $g_1, \dots, g_{\mathsf{A}}$. Hence $f_Q = g_i$ for some i. □

Theorem A. *For every* $\mathsf{P} > 1$*, there is a finite set of* 2*-complexes, which are mapping tori of self maps of graphs,* $X_1, \dots, X_{\mathsf{A}}$*, so that: If* $f\colon G \to G$ *is a* P*-small expanding train-track map on a graph* G*, then* M_f *is homeomorphic to a* 2*-complex that is obtained by surgery on some* X_i*.*

Proof. Let $g_1, \dots, g_{\mathsf{A}}$ be the maps from Proposition 4.5 and let X_i be the mapping torus of g_i for $i = 1, \dots, \mathsf{A}$. There is some i such that $f_Q = g_i$. By Proposition 3.14 X_i can be obtained from M_f by surgery. Since this operation is invertible, by Proposition 3.12, M_f can be constructed from X_i by surgery. □

5. Bounded presentations

Proposition 5.1. *If* $f\colon G \to G$ *is a* P*-small automorphism then* M_f *has a finite CW structure where the number of cells only depends on* P*.*

Proof. Let U be the vertical graph in $X := M_f$, and let $\mathring{X}$ be the floor plan of X. Let S be the vertical subgraph of M_{f_Q}, then by Proposition 3.14 we have $\mathring{X}_U = (M_{f_Q})_S$. By Proposition 4.2, the number of i-cells in M_{f_Q} is bounded by E for all $i = 0, 1, 2$. Therefore, E bounds the number of 2-cells and the number of horizontal edges in $\mathring{X}$. Even though the number of vertices of $\mathring{X}$ may be large, we shall argue that most of them are removable. Recall the map $\tau\colon K \to U$ from Definition 2.25. A vertex of U is a *natural vertex* if its valence is $\neq 2$ or if it is an image of a vertex of K. Thus, for an unnatural vertex $u \in U$, its valence in $\mathring{X}^{(1)}$ is also 2. Therefore, all unnatural vertices are removable. The complex X' that we obtain after removing all of the unnatural vertices is the complex that satisfies the claim. It is left to argue that the number of natural vertices of U may be bounded by a function that only depends on P.

The set of natural vertices is the union

$$\tau(\text{vertices of } K) \cup \big\{v \bigm| \operatorname{val}(v) = 1\big\} \cup \big\{v \bigm| \operatorname{val}(v) \ge 3\big\}$$

The first set in this union is bounded by E. Next we show that the cardinality of the set of valence 3 vertices is bounded by the cardinality of the set of valence 1 vertices.

Each component of U is a trap, a union of disjoint trees glued to a circle at one (valence 1) point of the tree. In a tree it is easy to verify (for example by induction on the edges) that

$$\left|\left\{v \mid \mathrm{val}(v)=1\right\}\right|-1 \geq \left|\left\{v \mid \mathrm{val}(v)\geq 3\right\}\right|$$

A vertex in U with valence ≥ 3 has to lie on one of the directed trees. Therefore the number of all vertices in U with valence ≥ 3 is smaller than the number of valence 1 vertices.

Finally, we claim

$$\tau(\text{vertices of } K) \supseteq \left\{v \mid \mathrm{val}(v)=1\right\}$$

Consider $\tau\colon\ K \to U$, if τ is not onto the vertices then there is a vertex in M_f with no horizontal edge adjacent to it. Thus G has an isolated vertex, but we assumed that G is connected. Therefore, τ must be onto. Since τ is locally injective on edges, if x is in the interior of an edge in U then $\mathrm{val}(\tau(v)) \geq 2$. Thus every valence 1 vertex in U is an image of a valence 1 vertex in K. This finishes the proof. □

Theorem B. *There is a number* M *depending only on* P *(and not on n) so that if $\phi \in \mathrm{Out}(F_n)$ is a* P*-small ett automorphism then Γ_ϕ has a presentation with at most* M *generators and* M *relations.*

Proof. Consider M_f with the CW structure obtained in Proposition 5.1. The complex $\widetilde{M_f}$ is contractible, therefore, the presentation of the fundamental group may be read from the CW structure of M_f whose number of cells is bounded by a function of P. □

6. Examples

In this section we give examples that have a small dilatation relative to their rank. These examples verify the upper bound for $\underline{\lambda}(n)$.

Example 6.1. Consider the rose with n leaves $x_1, \ldots, x_n$ and the map f defined by

$$x_1 \to x_2 \to \ldots \to x_n \to x_1 x_2.$$

This is a train-track map since it is positive. It is easy to verify that it is expanding. The dilatation of f is computed by declaring one of the edges, x_1, to have unit length and computing the other edge lengths by requiring that f stretches each edge by the factor λ. Thus the lengths of $x_2, \ldots, x_n$ are $\lambda, \ldots, \lambda^{n-1}$ respectively and from $f(x_n) = x_1 x_2$ we get the equation:

$$\lambda^n = 1 + \lambda \tag{7}$$

Let t_n be the root of this equation. Clearly, $\lim_{n\to\infty} t_n = 1$. By taking log in Eq. (7) we see that $\lim_{n\to\infty} n\log(t_n) = \log(2)$.

Define

$$\underline{\lambda}(n) = \inf\left\{\lambda_f \mid f\colon\ G \to G \text{ with } rk(\pi_1 G) = n\right\}$$

The example above implies

$$\lim_{n\to\infty} n\log\left(\underline{\lambda}(n)\right) \leq \log(2).$$

In Section 7 we provide a lower bound for the value of $\underline{\lambda}(n)$.

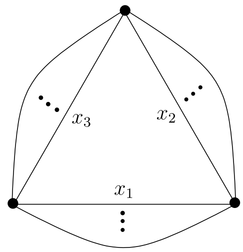

Fig. 10. The rank of this graph is the number of edges subtracted by 2. Edges map homeomorphically $x_1 \to x_2 \to x_3 \to x_4 \to \ldots \to x_{n+2}$.

Question 6.2. *What are the asymptotic of $\underline{\lambda}(n)$? In other words, what is the limit $\lim_{n\to\infty} n \log \underline{\lambda}(n)$?*

Example 6.3. We give a slightly better example than the rose example. Consider the graph in Fig. 10. The rank of this graph is the number of edges subtracted by 2. The edges map homeomorphically

$$x_1 \to x_2 \to x_3 \to x_4 \to \ldots \to x_n \to x_{n+1} \to x_{n+2}.$$

When $n \bmod 3 = 0$ then x_{n+2} is parallel to x_3 and we set $f(x_{n+2}) = X_3X_2$. If $n \bmod 3 = 1$ then x_{n+2} is parallel to x_1 and we set $f(x_{n+2}) = X_1X_3$. If $n \bmod 3 = 2$ then x_{n+2} is parallel to x_2 and we set $f(x_{n+2}) = X_2X_1$. One may verify using Proposition 2.11 that f is an expanding train-track map in all of these cases. The equation for the dilatation of f is one of the three equations:

$$\lambda^{n+2} = \lambda + \lambda^2$$
$$\lambda^{n+2} = 1 + \lambda^2$$
$$\lambda^{n+2} = 1 + \lambda$$

The roots of these are slightly smaller than of Eq. (7). However, they still satisfy $\lim_{n\to\infty} n \log \underline{\lambda}(n) = \log(2)$.

7. A lower-bound for dilatation

In this section we provide a lower-bound for $\underline{\lambda}(n)$. This does not match the upper-bound provided by examples in the previous section. However, we arrive naturally at Example 6.3 as, possibly, the ett with the smallest dilation (see Remark 7.4).

Let $f\colon G \to G$ be an expanding train-track map. Endow G with the natural metric as in Definition 2.12 so that f is λ_f-Lipschitz.

Convention 7.1. We denote by e the smallest edge of G and we scale G so that the length of e is 1.

Let E be the number of edges of G and V be the number vertices of G.

Claim 7.2. *For every pair of edges e', e'' in G, there exists a $k < \mathsf{E}$ so that $f(e')$ contains e''.*

Proof. Consider the directed graph Γ corresponding to the transition matrix of f. Since f is expanding Γ is connected. Thus there is a path from e' to e'' and its length is smaller than $\mathsf{E} - 1$.

This provides a quick lower-bound for λ_f. Namely, let e_{mix} be a mixed edge. Then $f(e_{mix})$ contains at least two edges and hence has a length of at least 2. Assuming e_{mix} is contained in $f^k(e)$ we have

$$\lambda_f^{k+1} = \left|f^{k+1}(e)\right| \geq \left|f(e_{mix})\right| \geq 2 \tag{8}$$

Since $k+1 \leq \mathsf{E} \leq 3n-3$, we have

$$\log \lambda_f \geq \frac{\log 2}{3n-3}.$$

Using the same type of argument, with more care, we can replace $\log(2)$ with $\log(3)$ in the above estimate.

Theorem 7.3. *If f is an expanding train-track map with dilatation λ_f then*

$$\log \lambda_f \geq \frac{\log 3}{3n-3}.$$

Proof. We consider the following cases separately:

(1) There is more than one mixed edge, or
(2) there is one mixed edge and its image contains three edges or more, or
(3) there is exactly one mixed edge and its image contains two edges.

For each edge e' let $k(e')$ be the first natural number so that $f(e)$ contains e' where e is the smallest edge. Suppose we are in case (1). Let $k = \max\{k(e') \mid e' \text{ is a mixed edge }\}$. Then $|f^{k+1}(e)| \geq 3$ and replacing 2 with 3 in Eq. (8) gives us:

$$\log \lambda_f \geq \frac{\log 3}{3n-3}.$$

In case (2), let e' be a mixed edge with $|f(e')| > 3$ and let $k := k(e')$ then as before we get $|f^{k+1}(e)| \geq 3$ and the conclusion of the theorem follows.

Suppose we are in case (3). Let e' be the mixed edge and $k := k(e')$. We claim that e must be contained in $f(e')$. Indeed since f is onto, e must be contained in the image of some edge of G. If e is contained in the image of e'' and $e'' \neq e'$ then e'' is not mixed. Therefore, $f(e'') = e$ or $\bar{e}$ and $\ell(e) = \lambda\ell(e'')$. But we have assumed that e is the smallest edge so this is impossible. Thus e is not in the image of any edge other than e'. Thus we denote $e_1 = e$ and $e_{i+1} = f(e_i)$ for $i = 1, \dots, k$. If $k < \mathsf{E}$, let e_{k+1} be the other edge in the image of e_k and let $e_{i+1} = f(e_i)$ for all $i = k+1, \dots, \mathsf{E}$. If we let $e_s = f(e_{\mathsf{E}})$ then $s \leq k$. Indeed if $s \geq k+1$ then we would have $\lambda^{\mathsf{E}-s}\ell(e_s) = \ell(e_{\mathsf{E}})$ and $\lambda\ell(e_{\mathsf{E}}) = \ell(e_s)$. This is a contradiction so $s \leq k$

We divide the vertices into equivalence classes as in Definition 3.2. We call each equivalence class a vertex orbit. We claim that there has to be only one vertex orbit. Otherwise, we have at least two vertex orbits A and B. There are 3-types of edges, those connecting a vertex in A to A, A to B or B to B. The type of an edge is preserved unless the edge is a mixed edge. That is, there a mixed edge for each type. Since we have only one mixed edge, there should be only one type of edge, which is from A to B (G is connected). But this implies that G is a bi-partite graph and the image of the mixed edge has a length at least 3 which has been dealt with in case (2).

We now analyze the case where there is only one vertex orbit. We argue that the set of images of vertices is the entire vertex set. The only vertex that may violate this is an endpoint v of the edge e_1. The vertex v is also an endpoint of e_{k+1} (otherwise it would be an image of a vertex). Moreover, if v is not an image of a vertex, than v is not an endpoint of any edge e_i for $i \neq 1, e_{k+1}$. If e_1 or e_{k+1} were loops then one of the ends of e_k would map to v. Thus, if v is not an image of a vertex then v has valence 2, but this contradicts our hypothesis that all vertices of G have valence ≥ 3.

Label the vertices $0, 1, \dots, (\mathsf{V}-1)$ with $f(i) = i+1$ and let e be the edge $[0, a]$. Then for $i = 1, \dots, k$, e_i has the form $[i, i+a]$, but there may be more than one edge of type $[i, i+a]$. The edge e_k is of the form

$(\mathsf{V}-1-a, \mathsf{V}-1)$ or $(\mathsf{V}-1, a-1)$. If the latter, then e_1 or e_{k+1} are loops. If e_{k+1} is a loop than so is e_s and therefore e_1. Thus the graph is a rose with n petals and in this case

$$\lambda_f \sim \frac{\log 2}{n} \tag{9}$$

so the theorem holds. Therefore, we have the case that $e_k = [\mathsf{V}-1-a, \mathsf{V}-1]$. Since $e_1 = [0, a]$ we have that $e_{k+1} = [a, \mathsf{V}-a]$ or $[0, \mathsf{V}-a]$. If $\gcd(a, \mathsf{V}) \neq 1$ then in both cases we will get a disconnected graph. Therefore,

$$\gcd(a, \mathsf{V}) = 1. \tag{10}$$

We again deal with two cases: $k = \mathsf{E}$ and $k < \mathsf{E}$. If $k = \mathsf{E}$, the last edge $[E-1, E+a-1]$ is mapped to a path of length two where one edge is $e = [0, a]$ and the other one an edge adjacent to e (either $[a, 2a]$ or $[-a, 0]$). That is, the end points of $[\mathsf{E}-1, \mathsf{E}+a-1]$ are mapped to either $[-a, a]$ or $[0, 2a]$. Either way, we have

$$\mathsf{V} = 3a \tag{11}$$

If $k < \mathsf{E}$ then we have $e_{\mathsf{E}} = [i, i-2a]$ or $[i, i-a]$ and $e_s = [i, i+a]$ so we either have $-2a \equiv a$ or $-a \equiv a$ which gives

$$3a \equiv 0 \quad \text{or} \quad 2a \equiv 0 \tag{12}$$

Eqs. (11) and (10) implies $\mathsf{V} = 3$. Eqs. (12) and (10) implies $\mathsf{V} = 3$ or $\mathsf{V} = 2$. In each of these cases Eq. (9) holds and the theorem follows. □

Remark 7.4. In fact, it seems that having more than one mixed edge should increase λ_f even further and it is reasonable to conjecture that Example 6.3 is the ett with lowest dilatation number.

Acknowledgements

The authors would like to thank PCMI for their hospitality. We express our gratitude to Catherine Pfaff and Lee Mosher for inspiring conversations, and Martin Lustig for helpful corrections.

References

[1] John W. Aaber, Nathan Dunfield, Closed surface bundles of least volume, Algebr. Geom. Topol. 10 (4) (2010) 2315–2342.
[2] Ian Agol, Ideal triangulations of pseudo-Anosov mapping tori, in: Topology and Geometry in Dimension Three, in: Contemp. Math., vol. 560, Amer. Math. Soc., Providence, RI, 2011, pp. 1–17.
[3] M. Bestvina, M. Feighn, M. Handel, Laminations, trees, and irreducible automorphisms of free groups, Geom. Funct. Anal. 7 (2) (1997) 215–244.
[4] Mladen Bestvina, Michael Handel, Train tracks and automorphisms of free groups, Ann. Math. 135 (1) (1992) 1–51.
[5] Benson Farb, Christopher J. Leininger, Margalit Dan, Small dilatation pseudo-Anosov homeomorphisms and 3-manifolds, Adv. Math. 228 (3) (2011) 1466–1502.
[6] Eriko Hironaka, Small dilatation mapping classes coming from the simplest hyperbolic braid, Algebr. Geom. Topol. 10 (4) (2010) 2041–2060.
[7] Ji-Young Ham, Won Taek Song, The minimum dilatation of pseudo-Anosov 5-braids, Exp. Math. 16 (2) (2007) 167–179.
[8] Jakob Nielsen, Die Isomorphismengruppe der freien Gruppen, Math. Ann. 91 (3–4) (1924) 169–209.
[9] R.C. Penner, Bounds on least dilatations, Proc. Am. Math. Soc. 113 (2) (1991) 443–450.
[10] John R. Stallings, Topology of finite graphs, Invent. Math. 71 (3) (1983) 551–565.