END-PERIODIC TRAIN TRACK MAPS AND DYNAMICS ON FREE-BY-CYCLIC GROUPS

A Dissertation
Submitted to
the Temple University Graduate Board

In Partial Fulfillment
of the Requirements for the Degree of
DOCTOR OF PHILOSOPHY

---

by
Ruth Meadow-MacLeod
Aug 2024

Examining Committee Members:

Samuel Taylor, Advisory Chair, Mathematics
David Futer, Mathematics
Matthew Stover, Mathematics
Radhika Gupta, School of Mathematics, Tata Institute of Fundamental Research

ABSTRACT

An overarching question in the study of 3-manifolds is that of how a given 3- manifold fibers over the circle. Thurston and Fried investigated this question for a closed hyperbolic 3-manifold M. They found an open fibered cone in $H^1(M, \mathbb{R})$ for which each integral point corresponds to a fibration of M over S^1. Landry, Minsky, and Taylor extended the investigation of the fibered cone to its boundary. They found infinite-type surfaces with end-periodic maps naturally occurring in the boundary.

Dowdall, Kapovich, and Leininger investigated fibrations over the circle for the mapping torus of an irreducible train track map on a finite graph with no valence-1 vertices. They found an open fibered cone in the first cohomology of the mapping torus for which each integral point corresponds to a fibration of the mapping torus over S^1.

The research conducted in this thesis lives in this setting. We will address the natural question of the significance of the boundary of the fibered cone for this mapping torus. We construct finite graphs which naturally arise in the boundary of the fibered cone. We then use these finite graphs to construct infinite graphs in the boundary with natural maps with properties analogous to those of end-periodic surface homeomorphisms.

ACKNOWLEDGEMENTS

First, thank you to my advisor, Samuel Taylor, for the many rounds of editing and his patience with my writing. I also want to thank my dissertation committee for helping me grow over the years and collectively teaching me about all the concepts I used in this thesis.

Thanks to my cat Mochi, whose companionship kept me sane during many late nights writing, and to my partners for being patient and encouraging during this long journey.

To my parents, I would not have made it here without your unending support; thank you. To my fellow graduate students, I'm so grateful for the community we have built together, and I'm glad not to have gone through it alone.

TABLE OF CONTENTS

Page

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Throughout this thesis, we will work with graphs (Definition 1.5), graph maps (Section 1.2.1), and their mapping tori (Definition 1.7). Our work aims to construct graphs and graph maps that naturally arise in semiflows on free-by-cyclic groups. Before we begin constructing graphs, we will discuss some of our sources of motivation. The motivation starts with the work of Thurston and Fried in the setting of 3-manifolds, which we will outline in the following section.

1.1 Motivation from 3-Manifolds

For a closed hyperbolic 3-manifold M fibering over S^1, the Thurston norm on $H^1(M, \mathbb{R})$ determines a unit ball, which is a convex polyhedron with finitely many rationally defined sides.

Definition 1.1. A **fibered class** of a closed hyperbolic 3-manifold M is a class $\alpha \in H^1(M, \mathbb{R})$ such that there is a fibration $h : M \to S^1$ of M over S^1 for which α is the map induced by h on $H_1(M, \mathbb{R})$. A **fiber surface** S is a surface in M that is a fiber of such an h, and we say that S is **dual** to the cohomology class induced by h.

Thurston proved that there is a set of open cones on top dimensional faces of the unit ball of the Thurston norm with the property that a primitive integral point is in the set of cones if and only if it is a fibered class [11].

Definition 1.2. For a 3-manifold M with a flow (Definition 1.8) φ, a surface S in M is a **cross section** of φ if S is transverse to φ and S intersects all flowlines of φ infinitely often in both the forward and backward flow directions.

Each of these fibered cones has a canonical pseudo-Anosov suspension flow φ, which is a flow on M with a cross section S such that the first return map of S under φ (the map sending each point in S to the first point at which it is flowed back into the S by φ) is pseudo-Anosov (a surface homeomorphism with no periodic isotopy classes of curves [4]). Fried expanded on Thurston's work by proving that every fiber surface dual to a cohomology class in the cone is a cross section of φ. Moreover, because the fiber surface is a cross section of φ, it has a well-defined first return map under φ, and that first return map is necessarily pseudo-Anosov [5].

1.1.1 At the Boundary

The discussion in Section 1.1 explains the structure in the interiors of fibered cones. However, until recently, the boundary of a fibered cone was mysterious. We now turn our attention to the boundary of a fibered cone. Building on previous work of Mosher [10], Landry, Minsky, and Taylor prove the existence of transverse surfaces in M dual to classes in the boundary of a fibered cone [7]. As an application of the existence of these surfaces, they then use a procedure called spinning (which is directly analogous to a procedure we will define in Chapter 6) which takes in a cross section and such a transverse surface dual to a class in the boundary of the fibered cone and produces an infinite type surface in M. This infinite type surface in M is also transverse to φ and meets every flowline infinitely often

in both the forward and backward direction. They call the first return maps of infinite type graphs produced in this way **spun pseudo-Anosov** [7]. Spun pseudo-Anosov maps are a class of maps on infinite-type surfaces with many structural similarities to pseudo-Anosov maps.

Definition 1.3. For a surface homeomorphism $h : S \to S$, we say that h is **end-periodic** if there is some integer $m > 0$ such that for each end E of S, there is a neighborhood U_E of E so that either (i) or (ii) holds:

1. $h^m(U_E) \subsetneq U_E$ and $\{h^{nm}(U_E)\}$ forms a neighborhood basis of E for $n \in \mathbb{Z}_{\geq 0}$. In this case, E is called an **attracting** end.

2. $h^{-m}(U_E) \subsetneq U_E$ and $\{h^{-nm}(U_E)\}$ forms a neighborhood basis of E for $n \in \mathbb{Z}_{\geq 0}$. In this case, E is called **repelling** end.

Landry, Minsky, and Taylor prove that an end-periodic surface homeomorphism $j : L \to L$ is isotopic to a spun pseudo-Anosov representative $h : L \to L$ if and only if j is atoroidal, i.e. j has no invariant essential finite multicurve up to isotopy [7, Theorem A]. They also prove that when j is atoroidal, the spun pseudo-Anosov representative $h : L \to L$ is transverse to the pseudo-Anosov suspension flow [7, Theorem B].

Definition 1.4. For a topological space X and a map $j : X \to X$, the **growth rate of fixed points of** j is

$$\mathrm{gr}(j) \coloneqq \limsup_{N \to \infty} \left(\#\{\mathrm{Fix}(j^N)\}\right)^{1/N}.$$

Expanding on Thurston and Fried's work in the fibered cone, Landry, Minsky, and Taylor prove that all limit points of growth rates in the fibered cone are themselves growth rates of spun pseudo-Anosov maps [7, Theorem 7.4]. They also prove in [7, Theorem C] that for a spun pseudo-Anosov map $h : L \to L$,

1. $\text{gr}(h) = \inf_{h \simeq j}(\text{gr}(j))$ over all end-periodic j.

2. $\text{gr}(h) = \max_{\alpha,\beta} \limsup_{n\to\infty} \sqrt[n]{i(\alpha, h^n(\beta))}$, where α and β are essential simple closed curves and i takes their intersection number.

3. The $\log$ of $\text{gr}(h)$ is the topological entropy of the restriction of h to the unique largest invariant compact set.

Although we draw our motivation from this work, our work will take place in a different setting. Rather than 3-manifolds, we will work with a 2-dimensional cell complex, and rather than fiber surfaces, our fibers will be graphs. In the next section, we set up the terminology we need to define our 2-dimensional cell complex and other objects of interest.

1.2 Terminology and Objects

We begin with some basic definitions that we will use throughout this thesis.

Definition 1.5. A **graph** is 1-dimensional cell complex. The 0-cells are called **vertices**, and the 1-cells are called **edges**. A **directed graph** is a graph such that every 1-cell has an orientation.

Definition 1.6. A topological graph map is a continuous map between graphs, sending vertices to vertices and edges to edge paths.

Definition 1.7. Given a topological graph map $g : G \to G$, the **mapping torus** of g is $Z_g \coloneqq G \times [0,1]/(x,1) \sim (g(x),0)$, which is a cell complex with the natural cell structure described in Section 2.1.

Definition 1.8. A **semiflow** φ on a topological space X, is a continuous map $\varphi : X \times \mathbb{R}_{\geq 0} \to X$ such that for all $x \in X$ and $s,t \in \mathbb{R}_{\geq 0}$, $\varphi(x,0) = x$ and $\varphi(\varphi(x,t)) = \varphi(x,s+t)$. We will also write $\varphi(x,t)$ as $\varphi^t(x)$.

A **flow** on X is defined similarly, except we replace $X \times \mathbb{R}_{\geq 0}$ with $X \times \mathbb{R}$ in the domain of φ.

We will consider the natural upward semiflows on mapping tori of topological graph maps (Section 2.2).

Definition 1.9. For a semiflow φ on a topological space X, a **closed orbit** of φ is a loop $\ell : S^1 \to X$ in X such that when we pick any base point on ℓ and view ℓ as a path $\ell : [0,1] \to X$, there is some $s \in \mathbb{R}_{\geq 0}$ such that $\varphi(t, \ell(0)) : [0,s] \to X$ is a reparametrization of ℓ. We will consider closed orbits of φ equivalent up to the choice of base point and reparametrization. We will use O_φ to denote the set of closed orbits of φ.

The mapping torus of a topological graph map has a natural semiflow, which we will define in Section 2.2.

1.2.1 Graph Maps

In this section, we will fix a topological graph map f that we will use throughout this thesis. We start by defining some of its properties.

Definition 1.10. A **train track map** g, is a topological graph map $g : G \to G$ such that for all $n \geq 0$, g^n is an immersion on the interior of each edge of G.

Definition 1.11. A topological graph map $g : G \to G$ for a finite graph G, is **irreducible** if the only proper invariant subgraphs of G are subsets of the vertices of G.

Definition 1.12. A topological graph map, $g : G \to G$, is **expanding** if for every edge $e \in G$, $|g^n(e)| \to \infty$ as $n \to \infty$.

Definition 1.13. A topological graph map, $g : G \to G$ is **weakly expanding** if no interval in G is sent homeomorphically to itself under any power of g.

Definition 1.14. A topological graph map, $g : G \to G$, is **expanding on all scales** if for every $x \in G$, every neighborhood U of x, and every edge e of G, there exists a positive integer m and an open interval $W \subset U$ so that g^m maps W homeomorphically onto e.

We now fix $f : \Gamma \to \Gamma$ as an irreducible train track topological graph map such that Γ is finite with no valence-1 vertices. Then we fix Z_f as the corresponding mapping torus and ψ the natural semiflow on Z_f. Note that f is expanding on all scales by [2, Lemma A.2], and that expanding on all scales implies weakly expanding. Therefore, f is expanding on all scales and weakly expanding.

1.2.2 Graphs in the Mapping Torus

Our primary objects of interest will be graphs inside of Z_f that are smoothly embedded (see Definition 2.3, this essentially means the embedding is piece-wise smooth) in Z_f and are transverse to ψ in an appropriate sense which we will now define.

Definition 1.15. A smoothly embedded graph H in the mappingtorus Z_g of $f : \Gamma \to \Gamma$ is **transverse to** the natural semiflow ψ_g if it has an oriented I-bundle neighborhood $G \times [-1, 1]$ such that any segment of a flowline of ψ intersecting $G \times [-1, 1]$ flows from $G \times \{-1\}$ to $G \times \{1\}$, intersecting G at exactly one point in between.

Remark 1.1. Dowdall, Kapovich, and Leininger, whose results we will discuss in Section 1.3, use a different definition of transverse defined in [2, Section 5.1], which is more concerned with the semiflow parametrization than our definition which is more topological. However, the transverse graphs they construct are also transverse by our definition, and the transverse graphs we will construct can be easily modified to be transverse by their definition.

Definition 1.16. A **cross section** of the natural semiflow ψ_g on the mapping torus Z_g of $g : G \to G$ is a finite smoothly embedded graph Θ in Z_g that is transverse to ψ_g such that $\{s \in \mathbb{R}_{\geq 0} : {\psi_g}^s(x) \in \Theta\}$ is unbounded for every $x \in Z_g$.

Definition 1.17. For a graph H that is smoothly embedded in Z_g and is transverse to the natural semiflow ψ_g, if every $x \in H$ is such that ψ_g eventually flows x back into G ($\exists t \in \mathbb{R}_{>0}$ such that ${\psi_g}^t(x) \in G$), we can define the **first return map** $f_H : H \to H$ of H by $f_H(x) = {\psi_g}^t(x)$, where t is the minimal $t \in (0, \infty)$ such that ${\psi_g}^t(x) \in H$. Due to the fact that H is smoothly embedded and transverse to ψ_g, a minimal such t exists if any such t does.

Remark 1.2. The first return map of a smoothly embedded transverse graph must be continuous.

Note that all cross sections in a mapping torus Z_g of $g : G \to G$ have first return maps by definition. Other graphs in Z_g may also sometimes have first return maps. We will later construct infinite smoothly embedded graphs with first return maps by spinning (Chapter 6) cross sections around smoothly embedded finite graphs that are transverse to ψ_g.

Claim 1.1. If a finite smoothly embedded transverse graph in Z_f has a first return map under ψ, then that graph is a cross section.

We can justify the above claim by observing that f is expanding on all scales, which implies that the graph is eventually flowed over all of Z_f by ψ. Therefore, the fact that the graph has a first return map under ψ implies that every flowline intersects the graph infinitely often, and consequently the graph is a cross section.

Definition 1.18. For a smoothly embedded graph H in Z_g with an oriented I-bundle neighborhood (such as a finite graph transverse to ψ_g or a finite co-oriented track (Definition 8.1)), we say that H is **dual** to a cohomology class $\mu \in H^1(Z_g, \mathbb{R})$ if for any homology class $\alpha \in H_1(Z_g, \mathbb{R})$ and for any loop a in Z_g that is representative of α, counting the oriented intersections of a with H yields $\mu(\alpha)$.

1.3 Motivation from Free-by-Cyclic Groups

We have now set up enough definitions and objects to discuss some of Dowdall, Kapovich, and Leininger's results about Z_f which are analogous to the results of Thurston and Fried discussed in Section 1.1.

Dowdall, Kapovich, and Leininger construct a $K(G,1)$ space X_f called the folded mapping torus in [2, Definition 4.4]. The folded mapping torus is a quotient of the mapping torus Z_f such that the quotient map is also a homotopy equivalence. The space X_f has an associated semiflow ([2, Section 4.2]) such that the quotient map sending Z_f to X_f respects the semiflows. We will work in this same setting of $f : \Gamma \to \Gamma$ and with similar objects as Dowdall, Kapovich, and Leininger do, however we will consider Z_f rather than X_f. Due to the fact that X_f and Z_f are homotopy equivalent, our results and those of Dowdall, Kapovich, and Leininger will pass back and forth between Z_f and X_f, with our result Theorem A implying their result [3, Proposition 5.10].

Dowdall, Kapovich, and Leininger prove that there is an open convex cone A_{X_f} in $H^1(X_f, \mathbb{R})$ such that A_{X_f} contains the map on $H_1(X_f, \mathbb{R})$ induced by $u_0 : \pi_1(X_f) \to \pi_1(X_f)/\pi_1(\Gamma) = \mathbb{Z}$ ([2, Proposition 1.2]) and every primitive integral point in A_{X_f} is dual to a π_1-injective cross section whose first return map is an expanding irreducible train track map ([2, Theorem B]). They go on to show ([3, Meta-Theorem I]) that there is an open, convex, polyhedral cone C_{X_f} in $H^1(X_f, \mathbb{R})$ with finitely many rationally defined sides which contains A_{X_f} and which can be equivalently defined as the following:

1. The open convex hull of rays through integral classes of $H^1(X_f, \mathbb{R})$ that are dual to cross sections,

2. $\{u \in H^1(X_f, \mathbb{R}) : u$ is positive on closed orbits of the semiflow $\}$,

Not only can C_{X_f} be defined as the open convex hull of rays through integral classes of $H^1(X_f, \mathbb{R})$ that are dual to cross sections, but Dowdall, Kapovich, and Leininger prove

that every primitive integral class in the interior of C_{X_f} has a dual cross section in X_f ([3, Proposition 5.10]) such that its first return map under the semiflow is an expanding irreducible train track map ([3, Theorem B]).

Due to the fact that cross sections give us fibrations of the mapping torus over S^1, we will refer C_{X_f} as the **fibered cone**.

For a primitive integral class μ in the interior of the fibered cone that is dual to the cross section Θ_μ in X_f, let f_μ denote the first return map of Θ_μ. Let Q_μ be the stable quotient of $(f_\mu)_*$ and let ϕ_μ be the induced endomorphism on Q_μ. Then Dowdall, Kapovich, and Leininger prove that Q_μ is a finitely generated free group, f_μ is injective, and $\pi_1(X_f) \simeq Q_\mu *_{\phi_\mu}$ [3, Theorem B]. For this reason, we say that the semiflow on X_f is a semiflow on a free-by-cyclic group.

The concept of a stretch factor is also significant in this setting. Dowdall, Kapovich, and Leininger prove [3, Theorem F] which says that there is a real-analytic, degree -1 homogeneous function $h : C_{X_f} \to \mathbb{R}$ which is positive and convex such that

1. For any primitive integral class μ in C_{X_f} dual to the cross section Θ_μ, let f_μ be the first return map of Θ_μ, Q_μ be the stable quotient of $(f_\mu)_*$, and ϕ_μ be the induced endomorphism on Q_μ. Then $h(\mu)$ is the log of the stretch factor $\lambda(f_\mu)$ (Definition 4.5) as well as the log of the stretch factor $\lambda(\phi_\mu)$ defined in [3, Section 1.1],

2. h tends to infinity at the boundary of C_{X_f}.

We will show in Proposition 4.1 that the stretch factor of a train track map on a finite graph is equal to the growth rate of fixed points as long as some power of the map has a fixed point.

1.4 Results

In Section 1.1 and Section 1.3, we found many analogies between the setting of a closed hyperbolic 3-manifold M which fibers over the circle and the setting of Z_f with $f : \Gamma \to \Gamma$ being an irreducible train track topological graph map such that Γ is finite with no valence-1 vertices. Both settings have open, convex, polyhedral, fibered cones in the first cohomology of M or Z_f that have finitely many rationally defined sides. In both cones, the integral points correspond to cross sections of their respective flow or semiflow. In both settings, the first return maps of cross sections have analogous properties, either being pseudo-Anosov or expanding irreducible train track maps. Finally, both settings have a concept of growth rate or stretch factor defined in the fibered cone such that the limit points of these values in the fibered cone are of interest.

In Section 1.1, we discussed how Landry, Minsky, and Taylor extended our understanding of the fibered cone to include some knowledge about its boundary. Inspired by this work, we examine the boundary of the fibered cone for Z_f, which was previously mysterious. We prove the following theorem about the existence of transverse graphs dual to the boundary of the fibered cone.

Theorem A. *Given the irreducible train track topological graph map $f : \Gamma \to \Gamma$ on the finite graph Γ, which has no valence-1 vertices, suppose η is an integral cohomology class*

in the closure of the fibered cone of the mapping torus Z_f of $f : \Gamma \to \Gamma$. If η is positive on all closed orbits of vertices of Γ under the natural semiflow on Z_f, then η is dual to a finite transverse graph Σ that is smoothly embedded in Z_f.

In Section 8.2 and Section 8.3, we construct a smoothly embedded finite graph dual to η. Then, in Chapter 8, we prove Theorem A by altering the smooth embedding to make the graph transverse. Altering the embedding to be transverse is the bulk of the proof of Theorem A. Similarly, in the setting of 3-manifolds, transversality is the difficult part of the proof of the existence of transverse surfaces dual to the boundary of the fibered cone.

The hypothesis that the cohomology class η is positive on all closed orbits of vertices of Γ is automatically satisfied on the interior of the fibered cone. As such, Theorem A implies Proposition 5.10 of [3].

We will define a procedure called spinning in Chapter 6, which takes in two finite graphs and outputs an infinite graph. In Chapter 7, we will apply this procedure to a cross section Θ along with a finite smoothly embedded transverse graph Σ that is dual to $\eta \in H^1(Z_f, \mathbb{R})$ in the boundary of the fibered cone to get an infinite graph of finite valence with finitely many ends. In Section 7.1, we define a procedure on Z_f called blowing air in. We will use this procedure in Chapter 7 to turn this infinite graph of finite valence with finitely many ends into another infinite graph L_a of finite valence with finitely many ends, which has a first return map f_{L_a} that is topological graph map.

End-periodic homeomorphisms on infinite-type surfaces are an area of active research, but this concept has yet to be explored for infinite graphs. The observation that f_{L_a}

is infinite with finitely many ends leads us to define what it means for a topological graph map on an infinite graph to be end-periodic in Chapter 5. We then show in Chapter 7 that f_{L_a} is a weakly expanding, end-periodic, train track topological graph map.

In Chapter 4, we discuss and define various stretch factors and growth rates for topological graph maps, cohomology classes, and transverse graphs in a mapping torus Z_g. Some of these concepts can be applied to infinite graphs and topological graph maps on infinite graphs. At the end of Chapter 5, we prove the following theorem about growth rates and stretch factors of end-periodic topological graph maps.

Theorem B. *Suppose G is an infinite graph of finite valence with no valence-0 vertices and with finitely many ends, and $g : G \to G$ is an end-periodic (Definition 5.2) train track topological graph map. If some power of g fixes a point of G, then the stretch factor $\lambda(g)$ is equal to the growth rate of fixed points $gr(g)$.*

At the end of Section 7.3, we will apply Theorem B to f_{L_a} to prove Theorem C, which discusses limit points of stretch factors in the fibered cone.

Theorem C. *Given the irreducible train track topological graph map $f : \Gamma \to \Gamma$ on the finite graph Γ, which has no valence-1 vertices, and the mapping torus Z_f of f with the natural semiflow ψ, let Θ be a cross section of ψ which is dual to μ in the interior of the fibered cone of Z_f. Also, let Σ be a finite smoothly embedded transverse graph in Z_f, which is dual to η in the boundary of the fibered cone, and let L_a be a graph made by blowing air into Z_f (Section 7.2) with respect to a graph resulting from spinning (Section 6.2) Θ around Σ. Suppose Ξ_i is the cross section dual to $\xi_i = \mu + i\eta$ that has a first return map*

f_{Ξ_i} that is a train track topological graph map (existence due to [2, Theorem B]). Then the stretch factors of the first return maps f_{Ξ_i} of Ξ_i limit to the stretch factor of the first return map f_{L_a} of L_a as i goes to infinity, or

$$\lim_{i\to\infty} \lambda(f_{\Xi_i}) = \lambda(f_{L_a}).$$

Now, in Chapter 2 and Chapter 3, we will dig deeper into the structure of Z_f.

CHAPTER 2

BACKGROUND

As we work with various graphs, we will use G to denote a general graph (possibly infinite) with $g : G \to G$ a topological map, Z_g the corresponding mapping torus, and ψ_g the natural semiflow on Z_g. In this chapter, we will describe the cell structure on Z_g and define the natural semiflow ψ_g.

2.1 Cell Structure of Z_g

Recall Definition 1.7, which defines the mapping torus Z_g of $g : G \to G$. We will put a natural cell structure on Z_g as described below.

First, we put the natural cell complex structure on $[0, 1]$ where the 0-cells are $\{0\}$ and $\{1\}$ and the 1-cell is glued to $\{0\}$ on one end and $\{1\}$ on the other. Then, the space $G \times [0, 1]$ has the canonical cell complex structure determined by the cell complexes G and $[0, 1]$. In particular, the 0-cells are the points of the form $v \times \{0\}$ and $v \times \{1\}$ where v is a 0-cell of G. The 1-cells fall into two categories, which we will call vertical and horizontal. The **vertical** 1-cells are the sets of the form $v \times [0, 1]$ where v is a 0-cell of G. The **horizontal** 1-cells are the sets of the form $e \times \{0\}$ or $e \times \{1\}$ where e is a 1-cell of G. For each 1-cell e of G, there is a 2-cell whose interior has the form $\mathring{(}e) \times (0, 1)$. If e ends on 0-cells v_1 and v_2 such that e is oriented from v_1 to v_2, this 2-cell is glued in along $(e \times \{0\}) \cdot (v_2 \times [0, 1]) \cdot (\overline{e} \times \{1\}) \cdot (v_1 \times [1, 0])$.

The quotient of $G \times [0,1]$ identifying $(x,1) \sim (g(x),0)$ glues $e \times \{1\}$ to $g(e) \times \{0\}$. We then alter the cell structure on $G \times [0,1]$ by adding 0-cells to $e \times \{1\}$ for each 1-cell e of G at the points on $e \times \{1\}$ which will be identified to points of the form $v \times \{0\}$.

Now, because $g : G \to G$ is a topological graph map, the quotient of $G \times [0,1]$ identifying $(x,1) \sim (g(x),0)$ respects this new cell structure on $G \times [0,1]$, gluing 0-cells to 0-cells and 1-cells to 1-cells. Therefore, the quotient space $Z_f = G \times [0,1]/(x,1) \sim (g(x),0)$ has the quotient cell structure.

Due to the fact that all the 0-cells of $G \times [0,1]$ on $G \times \{1\}$ were glued to 0-cells on $G \times \{0\}$, the 0-cells of Z_f are the points of the form $v \times \{0\}$ where v is a 0-cell of G. The 1-cells of Z_f fall into two categories, which we call vertical and horizontal. The **vertical** 1-cells of Z_f are still sets of the form $v \times [0,1]$ where v is a 1-cell of G. Due to the fact that every horizontal 1-cell of $G \times [0,1]$ of the form $e \times \{1\}$ was glued to $g(e) \times \{0\}$, the **horizontal** 1-cells are sets of the form $e \times \{0\}$ where e is a 1-cell of G.

Due to the fact that the quotient map did not identify the interiors of 2-cells with anything, for each 1-cell e of G, there is still a 2-cell whose interior has the form $\mathring{(}e) \times (0,1)$. The quotient map did, however, affect the gluing of the 2-cell. If e has endpoints v_1 and v_2 such that e is oriented from v_1 to v_2, the boundary of the 2-cell $e \times [0,1]$ is glued in along $(e \times \{0\}) \cdot (v_2 \times [0,1]) \cdot (\overline{ge} \times \{0\}) \cdot (v_1 \times [1,0])$. We will call the 2-cells of this cell structure **rectangles** or **rectangular 2-cells** because we will depict them as rectangles such as Figure 2.1. We will call the 2-skeleton $Z_g^{(2)}$ and the 1-skeleton $Z_g^{(1)}$.

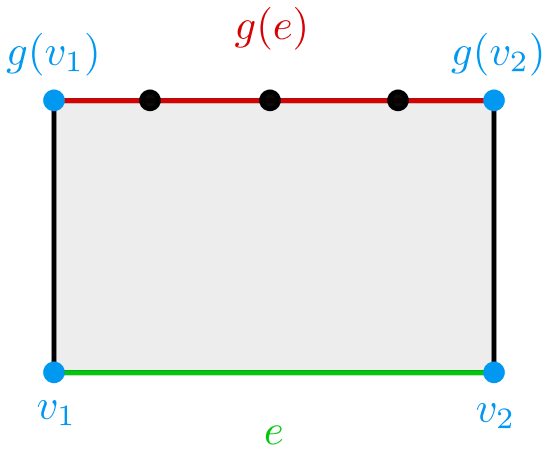

Figure 2.1: A 2-cell of $Z_f^{(2)}$.

Claim 2.1. For any smoothly embedded graph H in Z_g with an oriented I-bundle neighborhood. If H intersects each 2-cell of $Z_g^{(2)}$ in a finite disjoint union of intervals, then H is dual to a unique cohomology class in $H^1(Z_g, \mathbb{R})$.

The above claim is true because H intersects each 1-cell of $Z_g^{(1)}$ finitely many times, which implies that counting oriented intersections of loops in $Z_g^{(1)}$ with H is a well-defined map $\pi_1(Z_g^{(1)}) \to \mathbb{Z}$. Any loop ℓ in Z_g can be homotoped within 2-cells of $Z_g^{(2)}$ to be in $Z_g^{(1)}$, and if ℓ intersects H finitely many times, then homotoping it within 2-cells to be in $Z_g^{(1)}$ does not change its oriented intersection number with H. Therefore, counting oriented intersection numbers with H determines a well-defined map $\mu_H : \pi_1(Z_g) \to \mathbb{Z}$. For two loops in Z_g that each intersect H finitely many times, the oriented intersection number of their composition with H is the sum of both of their oriented intersection numbers with H. The oriented intersection number of the trivial loop with H is zero, and therefore, the map μ_H determines a cohomology class to which H is dual by construction.

2.2 The Semiflow ψ_g

The mapping torus $Z_g = G \times [0,1]/(x,1) \sim (g(x),0)$ has a natural upward semiflow determined by flowing points in the positive direction along $[0,1]$ at unit speed. We will call this semiflow ψ_g.

For the sake of visualizing ψ and graphs in the interiors of rectangular 2-cells, as we construct the cell structure of Z_g in Section 2.1 it is useful to identify the interior of each edge of G with $(0,1)$ by homeomorphism at the outset. As a result, the interiors of 2-cells of $G \times [0,1]$ are naturally homeomorphic to $e \times (0,1)$ or $(0,1) \times (0,1)$ such that flowlines of ψ_g are vertical lines in $(0,1) \times (0,1)$. These homeomorphisms allow us to draw accurate pictures of graphs in rectangles such that the semiflow direction is vertical.

Remark 2.1. By its construction, the semiflow ψ_g has the property that $\psi_g(n,(x,0)) = (g^n(x),0)$ for any integer $n \geq 0$, and $\psi_g(n,(x,0))$ is the nth time that ψ_g flows $(x,0) \in Z_g$ back into $G \times \{0\}$.

Some other definitions we will use that are related to the natural semiflow are as follows:

Definition 2.1. A **vertex leaf** of ψ_g is a connected component of the subset of Z_g that is eventually flowed into a vertex of G under ψ_g.

Definition 2.2. Fixing a small $\epsilon > 0$, for each rectangular 2-cell, r, with bottom 1-cell e, the set $\{{\psi_g}^t(x) : t \in (0,\epsilon), x \in \mathring{e}\} \cup \{{\psi_g}^t(y) : t \in (0,\epsilon),\ y \in Z_g \text{ such that } {\psi_g}^\epsilon(y) \in \mathring{e}\}$ has a finite number (which is constant as $\epsilon \to 0$) of connected components which we will call **sheets**.

Any horizontal 1-cell has one sheet above it but can have multiple sheets below it, as in Figure 2.2.

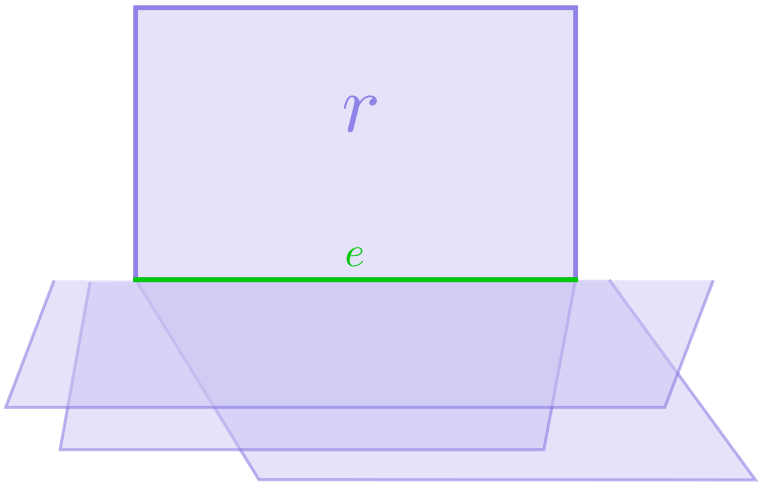

Figure 2.2: The edge e has three sheets below it. Note that the bottom side of r is a potentially proper subset of the top side of the rectangles below r.

2.3 Smoothly Embedded Graphs in the Mapping Torus

In this section, we will define what it means for a graph to be smoothly embedded in Z_g. We will also define other properties and objects that will later be applied to smoothly embedded graphs.

Definition 2.3. We will say that a graph H is **smoothly embedded** in Z_g, if it is embedded in Z_g such that any interval of H inside a rectangular 2-cell of $Z_g^{(2)}$ is smoothly embedded into that 2-cell, and any interval of H inside a 1-cell of $Z_g^{(1)}$ is smoothly embedded into that 1-cell.

We will later discuss smoothly embedded graphs in a finite subdivision of the natural cell structure of Z_g (Section 8.1). In that case, the definition is the same except that intervals of the graph in subdivided 2-cells and 1-cells will be smoothly embedded into the

subdivided 2-cells and 1-cells. A graph that is smoothly embedded in the subdivided cell structure can be easily upgraded to being smoothly embedded in the original cell structure by smoothing at the points where it intersects the added 1-cells or 0-cells of the subdivision.

The following definitions introduce a concept that we will apply to smoothly embedded graphs in Chapter 8 to prove Theorem A.

Definition 2.4. For a closed interval I that is properly embedded in a disk D such that the endpoints of I map into ∂D, and $\mathring{I}$ maps into $\mathring{D}$, a **co-orientation** of I is a choice of one of the connected components of $D \setminus I$.

Definition 2.5. For two intervals I_1 and I_2 properly embedded in the same disk D such that their images are disjoint, we will say that their co-orientations are **compatible** if the choices of connected components of $D \setminus I_1$ and $D \setminus I_2$ have non-empty intersection.

In this thesis, the intervals we will be co-orienting will be edges of smoothly embedded graphs in Z_g, and the disks will be 2-cells. For any interval I in a smoothly embedded transverse graph such that the interior of I is mapped into a 2-cell and the endpoints of I are mapped into the 1-skeleton, I has a well-defined upward co-orientation in the 2-cell by letting the co-orientation be in the direction of ψ_g.

In Chapter 3, we will discuss the flow graph, which is a smooth embedding of the transition graph of $g : G \to G$ into Z_g. The transition graph of $g : G \to G$ is defined as follows.

Definition 2.6. For a topological graph map, $g : G \to G$, the **transition graph**, D_g, has a vertex, v_a, for each edge, a of Γ. Then D_g has one oriented edge $e_{a,b}$, from v_a to v_b for each occurrence of b or b^{-1} in $g(a)$.

CHAPTER 3

THE FLOW GRAPH

The flow graph is a valuable construction for encoding the semiflow on the mapping torus Z_g of g. In this section, we will define the flow graph and discuss some of its properties, which we will use to prove various statements about growth rates and stretch factors in Chapter 4 and Chapter 5 including Theorem B.

The **flow graph** D_g is a smooth embedding of the transition graph of g into Z_g. The flow graph has exactly one vertex on the interior of each horizontal 1-cell of $Z_g^{(1)}$. As depicted in Figure 3.1, in the interior of each rectangular 2-cell r of $Z_g^{(2)}$, D_g has exactly one edge for each vertex in the top arc of r, such that the edge is directed from the one vertex on the bottom 1-cell of r to that vertex in the top arc of r.

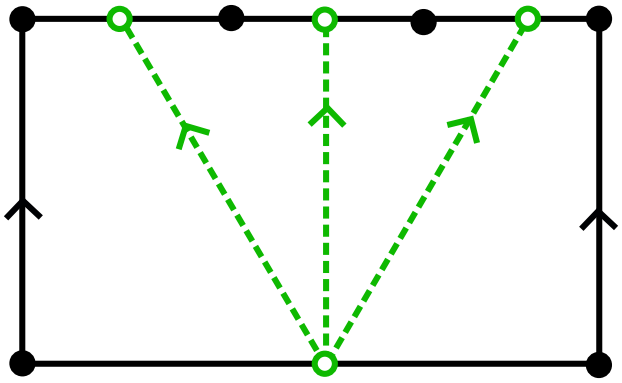

Figure 3.1: The green dotted edges and hollow green vertices show the flow graph, which is the smooth embedding of the transition graph into Z_g.

Definition 3.1. A directed graph D is **strongly connected** if for any pair of vertices v_1 and v_2 in D, there is a directed edge path in D starting at v_1 and ending at v_2.

Lemma 3.1. *If g is irreducible, then the flow graph D_g is strongly connected.*

Proof. For any pair of edges a and b in G, let v_a and v_b be the corresponding vertices of D_g. Due to the fact that g is irreducible, some power of g must map a over b, otherwise $\bigcup_{n\in\mathbb{Z}_{\geq 0}} g^n(a)$ would be a proper invariant subgraph of G that is not a set of vertices of G. This implies the existence of a directed path from v_a to v_b in D_g, which implies that D_g is strongly connected. □

Remark 3.1. The above lemma shows that D_f is strongly connected because f is irreducible.

The following definitions will allow us to categorize directed loops in the flow graph.

Definition 3.2. We will say that a set of rectangular 2-cells of the universal cover $\tilde{Z}_g$ is **aligned** if there exists a continuous path composed of exactly one vertical 1-cell from each rectangle. We will call such a path an **aligned side** of those 2-cells.

For example, the rectangular 2-cells in Figure 3.2 are aligned.

Definition 3.3. We will call a directed cycle, γ, of D_g a **vertex loop** when we can pick a base point b of γ on an oriented 1-cell e such that a path lift $\tilde{\gamma}$ of γ to the universal cover of Z_g satisfies the following properties:

- $\tilde{\gamma}$ passes through a sequence of aligned rectangular 2-cells,
- $\tilde{\gamma}$ starts and ends on lifts of e that are both adjacent to the same aligned side, and are either both oriented towards or both oriented away from that aligned side of that sequence of aligned rectangular 2-cells.

Vertex loops will be considered equivalent up to choice of base point. See Figure 3.2 for an example of a vertex loop.

By regularity of the universal cover, if one lift of γ for some choice of base point has this property, then all lifts of γ for all choices of base point do.

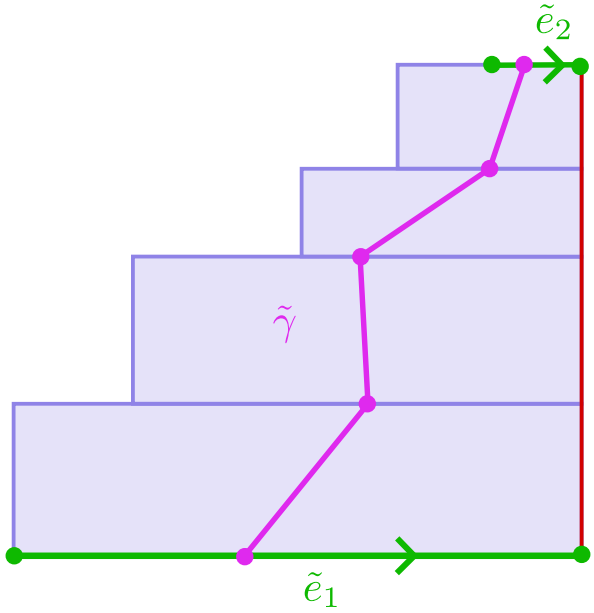

Figure 3.2: $\tilde{\gamma}$ is a lift of a vertex loop γ, and $\tilde{e}_1$ and $\tilde{e}_2$ are both lifts of an edge e. The rectangles $\tilde{\gamma}$ passes through are aligned, with the red path on the right as their aligned side.

Remark 3.2. For a topological graph map $g : G \to G$, every vertex loop γ of D_g is homotopic to some closed vertex orbit of G. In particular, it is homotopic to a closed vertex orbit that has a path lift that is an aligned side of the rectangles that $\tilde{\gamma}$ passes through. Thus, any closed vertex orbit in Z_g of a vertex v corresponds in this way to at most as many vertex loops as the valence of v, and if v has finite valence greater than zero, any closed vertex orbit of v has some power that is homotopic to a vertex loop.

Now, we will introduce a type of equivalence that applies to paths and loops in the mapping torus. We will apply this equivalence to closed orbits of ψ and loops in smoothly embedded graphs. In this chapter, the loops in smoothly embedded graphs which we will apply the following definition to are directed cycles in the flow graph.

Definition 3.4. We will say that two paths $p_{1,2} : I \to Z_g$ (or mapped into any cover of Z_g) such that $p_i^{-1}(Z_g^{(1)})$ is a discrete set of points for $i = 1, 2$ are **in-step** if they have lifts to the universal cover that pass through the exact same sequence of 1-cells and rectangular 2-cells of the universal cover. See Figure 3.3 for an example of lifts of paths that are in-step. Similarly, a loop will be considered in-step with a path (or another loop) if it has a choice of base point so that as a path starting and ending at that base point, it is in-step with that path (or loop).

Note that loops that are in-step with one another are homotopic, and being in-step is a relation on the loops it applies to.

Remark 3.3. Suppose two vertex loops of D_g have path lifts (after choosing base points), that pass through the same sequence of 1-cells and 2-cells of $\tilde{Z}_g^{(1)}$ and $\tilde{Z}_g^{(2)}$. Then the two

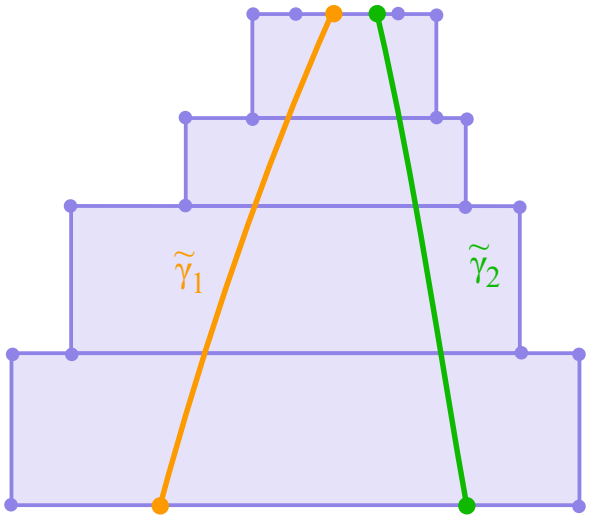

Figure 3.3: $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ are lifts of two paths that are in-step with one another.

vertex loops must be the same because each horizontal 1-cell of $Z_g^{(1)}$ has exactly one vertex of D_g on it, and each pair of vertices on the boundary of a rectangular 2-cell is connected by at most one edge of D_g.

Claim 3.1. Suppose ℓ is a lopp in Z_g that only passes through interiors of horizontal 1-cells of $Z_g^{(1)}$ and 2-cells of $Z_g^{(2)}$ such that any time ℓ passes through a 2-cell it does so from the bottom horizontal 1-cell (relative to ψ_g) to the top. Then ℓ in-step with a unique directed cycle of D_g.

The above claim can be seen by taking a path lift (after choosing a base point) of ℓ to the universal cover $\tilde{Z}_g$. The construction of D_g is such that each horizontal 1-cell of $Z_g^{(1)}$ has exactly one vertex of D_g on it. Each pair of vertices on the boundary of a rectangular 2-cell such that one vertex is on the bottom horizontal 1-cell (relative to ψ_g) and the other is on the top is connected by exactly one edge of D_g directed from the bottom to the top. Therefore, there is a unique directed path $\tilde{\ell}'$ in $\tilde{D}_g$ that passes through the same sequence

of 1-cells and 2-cells in the universal cover of Z_g as $\tilde{\ell}$ does, and this is a path lift of the unique directed cycle of D_g that is in-step with ℓ. By the regularity of the universal cover, this directed cycle ℓ' does not depend on the choice of base point of ℓ.

The following lemma and proposition will be used throughout this thesis alongside Remark 3.2 to relate directed cycles of the flow graph D_g with closed orbits of ψ_g

Lemma 3.2. *For a weakly expanding topological graph map $g : G \to G$, the following two scenarios cannot happen:*

1. *γ_1 and γ_2 are distinct closed orbits of ψ_g that are not closed vertex orbits, and γ_1 and γ_2 are in-step with one another as in Figure 3.4 (the definition Definition 3.4 of in-step does not apply to vertex orbits),*

2. *γ_1 is a closed orbits of ψ_g that is not a closed vertex orbit, and it is in-step with a vertex loop γ of D_g as in Figure 3.5.*

Proof. In scenario (2), let $\tilde{\gamma}_1$ and $\tilde{\gamma}$ be path lifts of γ_1 and γ (after choosing base points) to the universal cover of Z_g that pass through the same sequence of rectangles and 1-cells. As a result of γ being a vertex loop, the sequence of rectangles $\tilde{\gamma}$ passes through has an aligned side, which is a path lift $\tilde{\gamma}_2$ of some vertex loop γ (after choosing a base point). Suppose e is the oriented edge of G that the chosen base point of γ_1 is on, and suppose without loss of generality that the lifts of e that $\tilde{\gamma}_1$ start and end on are both oriented towards $\tilde{\gamma}_2$.

In scenario (1), let $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ be path lifts of γ_1 and γ_2 respectively (after choosing base points) to the universal cover of Z_g that pass through the same sequence of rectangles and 1-cells.

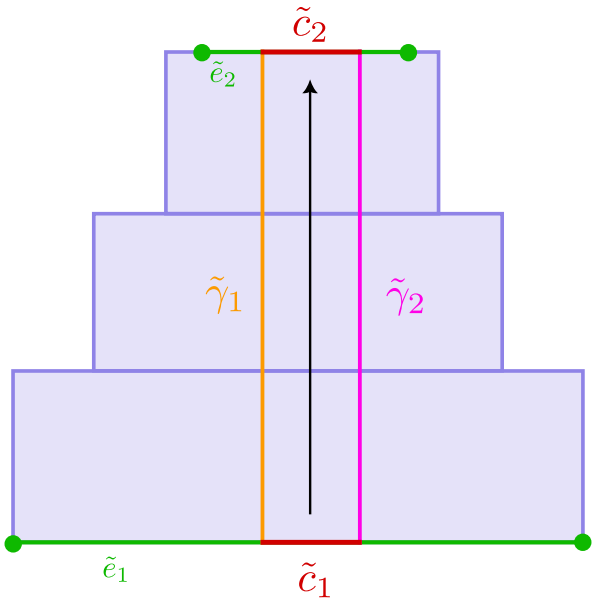

Figure 3.4: $\tilde{\gamma}$ and $\tilde{\gamma}'$ are lifts of two closed orbits of ψ_g that are not vertex orbits.

In both scenarios, let $\tilde{c}_1$ be the interval between the base points of $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ and $\tilde{c}_2$ be the interval between their endpoints. Due to the fact that $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ are lifts of distinct closed orbits, they start and end on different lifts of their chosen distinct base points.

Therefore, in scenario (1), if e is the edge of G on which γ_1 and γ_2 have their chosen base points, $\tilde{c}_1$ and $\tilde{c}_2$ must both be lifts of the unique sub-interval c of e between those two points. In scenario (2), $\tilde{c}_1$ and $\tilde{c}_2$ must both be lifts of the unique sub-interval c of e between the chosen base points of γ_1 and γ_2 that is oriented towards the base point of γ_2.

Thus, in both scenarios, the flow $\tilde{\psi}_g$ sends $\tilde{c}_1$ homeomorphically to $\tilde{c}_2$, implying that a power of g maps c homeomorphically to itself (if $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ pass through m rectangles, then g^m maps c homeomorphically to itself). This observation contradicts the fact that g is weakly expanding, so these scenarios are impossible. □

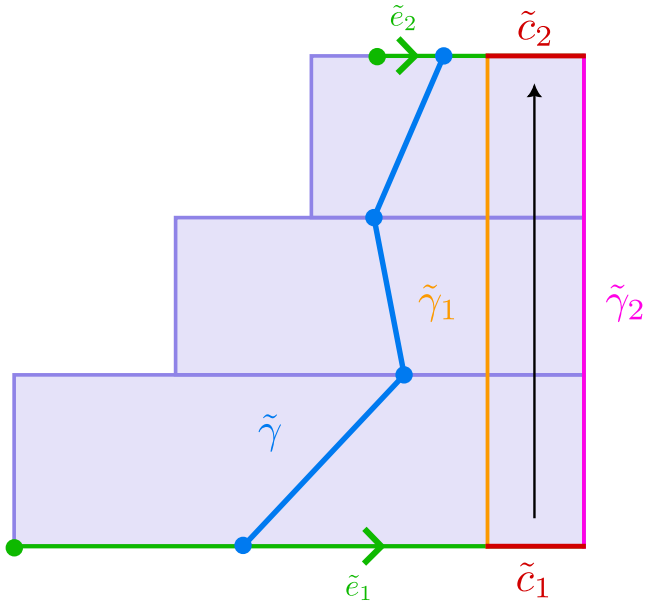

Figure 3.5: $\tilde{\gamma}$ is a lift of a vertex loop of D_g while $\tilde{\gamma}_1$ is a lift of a closed orbit of ψ_g that is not a vertex orbit.

Proposition 3.1. *For a weakly expanding train track map, $g : G \to G$, there is a bijection between the set of directed cycles of D_g that are not vertex loops and the set closed orbits of ψ_g that are not closed vertex orbits, such that the bijection preserves primitivity and powers.*

Moreover, the directed cycle γ' is in-step with the corresponding closed orbit γ.

Proof. First, we will define a map sending closed orbits γ of ψ_g that are not closed vertex orbits to directed cycles γ' of D_g that are not vertex loops. For any closed orbit γ that is not a closed vertex orbit, by Claim 3.1, γ is naturally in-step with a directed cycle γ' of D_g. By scenario (2) of Lemma 3.2, γ' must not be a vertex loop.

By Remark 3.3, distinct directed cycles of D_g cannot be in-step with one another. Therefore, more than one directed cycle of D_g cannot be in-step with γ, and therefore,

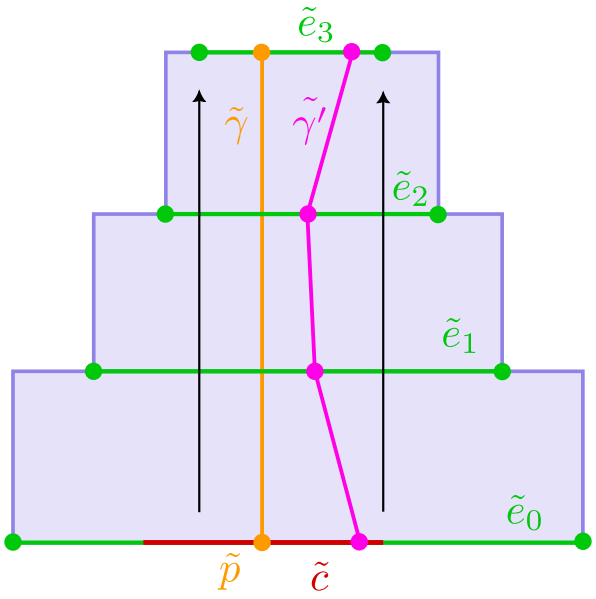

Figure 3.6: $\tilde{\gamma}$ is a lift of a closed orbit, while $\tilde{\gamma}'$ is a lift of a directed cycle in D_g.

this choice of γ' is well-defined. Similarly, the map is injective because scenario (1) of Lemma 3.2 tells us that distinct closed orbits that are not closed vertex orbits cannot be in-step with one another. Therefore, more than one cannot be in-step with γ'.

To show this map is onto, let γ' be a primitive directed cycle of the flow graph D_g that is not a vertex loop. Consider a path lift $\tilde{\gamma}'$ of γ' (after choosing a base point) to the universal cover, $\tilde{Z}_g$ of Z_g. Let $\tilde{e}_0, \tilde{e}_1, \ldots, \tilde{e}_{m-1}, \tilde{e}_m$ be the sequence of horizontal edges of $\tilde{Z}_g^{(1)}$, that $\tilde{\gamma}'$ passes through, starting on $\tilde{e}_0$ and ending on $\tilde{e}_m$ as in Figure 3.6 ($e_0 = e_m$ is the horizontal 1-cell on which γ' has it's chosen base point). Let $\tilde{c}$ be the segment of $\tilde{e}_0$ that $\tilde{\psi}_g$ flows onto $\tilde{e}_m$. Thus, g^m maps c onto e_0. Due to the fact that γ' is not a vertex loop, $\tilde{c}$ is contained in the interior of $\tilde{e}_0$. Due to the fact that g is a train track map, g^m maps c homeomorphically onto e_0, and hence over itself. Therefore, g^m fixes some point, p, on c. Due to the fact that $\tilde{\psi}_g$ flows $\tilde{p}$ (the lift of p to $\tilde{e}_0$) through the same sequence of rectangular

2-cells as $\tilde{\gamma}'$ passes through, the closed orbit of p is in-step with γ'. Due to the fact that c is contained in the interior of e_0, p is not a vertex of G, so the closed orbit of p is not a closed vertex orbit. Therefore, this map sends the closed orbit of p to γ', and so it is onto.

We have now shown that there is a bijection between closed orbits γ of ψ_g that are not closed vertex orbits to directed cycles γ' of D_g that are not vertex loops, such that γ and γ' are in-step.

Finally, to show this bijection preserves primitivity and powers, suppose a closed orbit γ that is not a closed vertex orbit is a power of α. Then α also must not be a closed vertex orbit, and therefore the bijection sends α to some corresponding directed cycle α' that is in-step with α. Therefore, a power of α' is in-step with γ', and so is equal to γ' as distinct directed cycles of D_g cannot be in-step with one another. This observation implies that γ' is also a power of a different directed cycle. Similarly, if γ' is a power, so much γ be. Therefore, this bijection sends primitive closed orbits to primitive directed cycles, and its inverse sends primitive directed cycles to primitive closed. □

CHAPTER 4

GROWTH RATES AND STRETCH FACTORS

In this chapter, we will define several different concepts of growth rates and discuss how they relate to one another and when they are equivalent. We will use the results in this chapter to prove Theorem B in Chapter 5 and Theorem C in Section 7.3.2. Recall that we have already defined one type of growth (Definition 1.4), and recall the definition of O_{ψ_g} from Definition 1.9.

4.1 Growth of Closed Orbits

We first define some growth rates on cohomology classes and transverse graphs in Z_g that measure the growth of sets of closed orbits of ψ_g.

Definition 4.1. For a cohomology class, $\xi \in H^1(Z_g, \mathbb{R})$, let the **growth rate of closed orbits of ψ_g on ξ** be

$$\mathrm{gr}_{\psi_g}(\xi) \coloneqq \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : \xi(\gamma) \leq N\}\right)^{1/N}.$$

Remark 4.1. For Z_f, when $\xi \in H^1(Z_f, \mathbb{R})$ is not in the fibered cone, $\mathrm{gr}_{\psi}(\xi)$ is infinite because there is some closed orbit of ψ on which ξ is not positive.

Definition 4.2. For a transverse graph H smoothly embedded in Z_g, let the **growth rate of closed orbits of ψ_g on H** be

$$\mathrm{gr}_{\psi_g}(H) \coloneqq \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : \gamma \text{ intersects } H \text{ at } N \text{ times }\}\right)^{1/N}.$$

This lemma is one which we will use throughout this chapter as well as in Chapter 5. Our first application of the lemma will be to show that when the transverse graph H has

a first return map f_H, the newly defined growth rate of closed orbits on H is equal to the previously defined growth rate of fixed points of f_H.

Lemma 4.1. *Let A_N and B_N be sets for $N \in \mathbb{Z}_{>0}$, and define sequences $a_N = |A_N|$ and $b_N = |B_N|$. Suppose that b_N does not converge to 0 as N goes to infinity, and suppose for all $N \in \mathbb{Z}_{>0}$, $a_N \leq b_N \leq Na_N$. Then*

$$\limsup_{N\to\infty}(a_N)^{1/N} = \limsup_{N\to\infty}(b_N)^{1/N} = \limsup_{N\to\infty}\left(\sum_{i=1}^{N} a_i\right)^{1/N}.$$

The equality also holds with the sum starting at $i = 0$ rather than $i = 1$.

Note that if a_N does not converge to 0 as N goes to infinity, letting $B_N = A_N$ implies that

$$\limsup_{N\to\infty}(a_N)^{1/N} = \limsup_{N\to\infty}\left(\sum_{i=1}^{N} a_i\right)^{1/N}.$$

Proof. The first equality is easier to show. Due to the fact that $a_N \leq b_N$ we have

$$\limsup_{N\to\infty}(a_N)^{1/N} \leq \limsup_{N\to\infty}(b_N)^{1/N},$$

and because $a_N \geq \frac{b_N}{N}$ we have

$$\limsup_{N\to\infty}(a_N)^{1/N} \geq \limsup_{N\to\infty}(\frac{1}{N}\cdot b_N)^{1/N} = \lim_{N\to\infty}(\frac{1}{N})^{1/N}\cdot\limsup_{N\to\infty}(b_N)^{1/N}$$

$$= 1\cdot\limsup_{n\to\infty}(b_N)^{1/N} = \limsup_{N\to\infty}(b_N)^{1/N}.$$

Using the first equality, one direction of the second equality is also simple

$$\limsup_{N\to\infty}\left(\sum_{i=1}^{N} a_i\right)^{1/N} \geq \limsup_{N\to\infty}(a_N)^{1/N} = \limsup_{N\to\infty}(b_N)^{1/N}.$$

For the other direction of the second equality, define the sequence $h_N = \max\{b_i : i \in \mathbb{Z}, 1 \leq i \leq N\}$ (if we want the sum to start at $i = 0$, then take the maximum over $0 \leq i \leq N$), and let $1 \leq R_N \leq N$ be the maximum index such that $h_N = b_N$. Therefore, R_N is a positive, non-decreasing integer sequence. Either b_N is bounded by B, or R_N has a monotone increasing subsequence R_{N_k}.

If b_N is bounded by B, then we have

$$\limsup_{N\to\infty} \left(\sum_{i=1}^{N} a_i\right)^{1/N} \leq \limsup_{N\to\infty} \left(\sum_{i=1}^{N} b_i\right)^{1/N} \leq \limsup_{N\to\infty} (N \cdot B)^{1/N} = 1.$$

Due to the fact that b_N is bounded and does not converge to 0,

$$\limsup_{N\to\infty} (b_N)^{1/N} = 1,$$

which gives us the second equality in the case where b_N is bounded.

If b_N is not bounded, then

$$\limsup_{N\to\infty} \left(\sum_{i=1}^{N} a_i\right)^{1/N} \leq \limsup_{N\to\infty} \left(\sum_{i=1}^{N} b_i\right)^{1/N} \leq \limsup_{N\to\infty} \left(\sum_{i=1}^{N} b_{R_N}\right)^{1/N}$$

$$= \limsup_{N\to\infty} (N \cdot b_{R_N})^{1/N} = \lim_{N\to\infty} (N)^{1/N} \cdot \limsup_{N\to\infty} (b_{R_N})^{1/N} = \limsup_{N\to\infty} (b_{R_N})^{1/N}.$$

Due to the fact that $R_N \leq N$, we have

$$\limsup_{N\to\infty} (b_{R_N})^{1/N} \leq \limsup_{N\to\infty} (b_{R_N})^{1/R_N}.$$

Let $b_{R_{N_j}}$ be the subsequence such that

$$\lim_{j\to\infty} (b_{R_{N_j}})^{1/R_{N_J}} = \limsup_{N\to\infty} (b_{R_N})^{1/R_N}.$$

Then we can pass to a subsequence $b_{R_{N_{j_k}}}$ so that $R_{N_{j_k}}$ is monotone increasing, meaning that $(b_{R_{N_{j_k}}})^{1/R_{N_{j_k}}}$ is a subsequence of $(b_N)^{1/N}$, and

$$\lim_{k\to\infty}(b_{R_{N_{j_k}}})^{1/N_{j_k}} = \limsup_{N\to\infty}(b_{R_N})^{1/N},$$

implying that

$$\limsup_{N\to\infty}(b_{R_N})^{1/R_N} \leq \limsup_{N\to\infty}(b_N)^{1/N}. \qquad \square$$

Corollary 4.1. *For a smoothly embedded transverse graph H in Z_g with first return map $f_H : H \to H$, if some power of f_H has a fixed point, then*

$$\limsup_{N\to\infty}\left(\#\{\gamma \in O_{\psi_g} : \gamma \textit{ intersects } H \textit{ at } N \textit{ times }\}\right)^{1/N}$$

$$= \limsup_{N\to\infty}\left(\#\{\textit{Fix}(f_H^N)\}\right)^{1/N}$$

$$= \limsup_{N\to\infty}\left(\#\{\gamma \in O_{\psi_g} : \gamma \textit{ intersects } H \textit{ at most } N \textit{ times }\}\right)^{1/N}.$$

So, the growth rate of fixed points $gr(f_H)$ of f_H is equal to the growth rate of closed orbits $gr_{\psi_g}(H)$ of ψ_g on H.

Proof. If γ intersects H at N times, there are between 1 and N distinct points in the intersection of γ and H. Each intersection point is a fixed point of f_H^N. If x is a fixed point of f_H^N, then by the definition of the first return map f_H, x must be the nth point on the forward flow of x under ψ_g that intersects H. Therefore, the forward flow of x under H starting at x and until that nth intersection with H must be a closed orbit that intersects H at N times. Then let A_N be the set of closed orbits γ that intersect H at N times, and B_N be the number of fixed points of f_H^N. Letting $a_N = |A_N|$ and $b_N = |B_N|$, we have that $a_N \leq b_N \leq Na_N$.

The sequence b_n does not converge to 0 as N goes to infinity because some power f_H^k has a fixed point. So for every multiple mk of k, f_H^{mk} has some fixed point. Therefore, we can apply Lemma 4.1 to get the desired equality:

$$\limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : \gamma \text{ intersects } H \text{ at } N \text{ times }\}\right)^{1/N} = \limsup_{N\to\infty} \left(\#\{\text{Fix}(f_H^N)\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : \gamma \text{ intersects } H \text{ at most } N \text{ times }\}\right)^{1/N}. \qquad \square$$

The above corollary implies that when the transverse graph H is dual to the cohomology class ξ, the two concepts of growth on ξ and H defined at the start of this section are equal.

4.2 Growth on Directed Graphs

We now define some growth rates of cohomology classes of directed graphs.

Definition 4.3. For a directed graph, D, and a cohomology class, $a \in H^1(D, \mathbb{R})$ that is positive on directed cycles, the **growth rate of** a **on** D is

$$\text{gr}(a; D) \coloneqq \limsup_{N\to\infty} \left(\#\{\gamma : \gamma \text{ is a directed cycle of } D, \text{ and } a(\gamma) = N\}\right)^{1/N}.$$

Definition 4.4. For a directed graph D, and cohomology classes, $a, b \in H^1(D, \mathbb{R})$. If a is positive on directed cycles and b is non-negative on directed cycles, then the **growth rate of** a **on** D **relative to** b is

$$\text{gr}(a; b; D)$$

$$\coloneqq \limsup_{N\to\infty} \left(\#\{\gamma : \gamma \text{ is a directed cycle of } D, a(\gamma) = N, \text{ and } b(\gamma) = 0\}\right)^{1/N}.$$

The following lemma will later be used to prove Theorem C. Its first application will be to relate the growth rate of closed orbits on a transverse graph in Z_g to the growth rate of the corresponding cohomology class on D_g.

Lemma 4.2. *Let G be a graph with map $g : G \to G$ such that G has finitely many periodic vertices, and the periodic vertices of G are all finite valence. Suppose that $\alpha \in H^1(Z_g, \mathbb{R})$ is positive on closed orbits of O_{ψ_g} and $\beta \in H^1(Z_g, \mathbb{R})$ is non-negative on closed orbits of O_{ψ_g} such that β is zero on some closed orbit of ψ_g. Letting a be the pullback of α to D_g and b be the pullback of β to D_g, then*

$$gr(a; b; D_g) = \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : \beta(\gamma) = 0, \alpha(\gamma) = N\}\right)^{1/N}.$$

Proof. First note that by Remark 3.2 and Proposition 3.1, all directed cycles of D_g are homotopic to closed orbits of ψ, and therefore a is positive on all directed cycles of D_g and b is non-negative on directed cycles of D_g. Also, due to the fact that β is zero on some closed orbit of ψ_g, and α must be postive on that orbit, $\#\{\gamma : \gamma \in O_{\psi_g}, \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}$ does not limit to zero as N goes to infinity. Similarly, by Remark 3.2 and Proposition 3.1, $\#\{\gamma : \gamma \text{ is a directed cycle of } D_g, a(\gamma) = N, \text{ and } b(\gamma) = 0\}$ does not limit to zero as N goes to infinity

By Remark 3.2, and because G has finitely many periodic vertices, which are each finite valence, the number of primitive vertex loops of D_g is bounded above by some integer V_1.

Then we have

$$|\{\gamma : \gamma \text{ is a non-vertex loop directed cycle of } D_g, a(\gamma) = N, \text{ and } b(\gamma) = 0\}|$$

$$\leq |\{\gamma : \gamma \text{ is a directed cycle of } D_g, a(\gamma) = N, \text{ and } b(\gamma) = 0\}|$$

$$\leq |\{\gamma : \gamma \text{ is a non-vertex loop directed cycle of } D_g, a(\gamma) = N, \text{ and } b(\gamma) = 0\}| + V_1.$$

Therefore, as N goes to infinity these sets satisfy the inequality in Lemma 4.1. Then because $\#\{\gamma : \gamma \text{ is a directed cycle of } D_g, a(\gamma) = N, \text{ and } b(\gamma) = 0\}$ does not limit to zero as N goes to infinity, Lemma 4.1 applies and gives us

$$\limsup_{N\to\infty} \left(\#\{\gamma : \gamma \text{ is a directed cycle of } D_g, a(\gamma) = N, \text{ and } b(\gamma) = 0\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \Big(\#\{\gamma : \gamma \text{ is a non-vertex loop directed cycle of } D_g,$$

$$a(\gamma) = N, \text{ and } b(\gamma) = 0\}\Big)^{1/N}.$$

Similarly, because G has finitely many periodic vertices, the number of primitive vertex orbits of ψ_g is bounded above by some integer V_2. Then we have

$$|\{\gamma : \gamma \in O_{\psi_g}, \gamma \text{ is not a vertex orbit, } \alpha(\gamma) = N, \text{ and } b(\gamma) = 0\}|$$

$$\leq |\{\gamma : \gamma \in O_{\psi_g}, \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}|$$

$$\leq |\{\gamma : \gamma \in O_{\psi_g}, \gamma \text{ is not a vertex orbit, } \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}| + V_2.$$

Therefore, as N goes to infinity, these sets satisfy the inequality in Lemma 4.1. Then because $\#\{\gamma : \gamma \in O_{\psi_g}, \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}$ does not limit to zero as N goes to infinity, Lemma 4.1 applies and gives us

$$\limsup_{N\to\infty} \left(\#\{\gamma : \gamma \in O_{\psi_g}, \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}\right)^{1/N}$$
$$= \limsup_{N\to\infty} \left(\#\{\gamma : \gamma \in O_{\psi_g}, \gamma \text{ is not a vertex orbit, } \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}\right)^{1/N}.$$

The above observations in combination with Proposition 3.1 gives us the following sequence of equalities, with the middle equality due to Proposition 3.1.

$$\limsup_{N\to\infty} \left(\#\{\gamma : \gamma \text{ is a directed cycle of } D_g, a(\gamma) = N, \text{ and } b(\gamma) = 0\}\right)^{1/N}$$
$$= \limsup_{N\to\infty} \left(\#\{\gamma : \gamma \text{ is a non-vertex loop directed cycle of } D_g,\right.$$
$$\left. a(\gamma) = N, \text{ and } b(\gamma) = 0\}\right)^{1/N}.$$
$$= \limsup_{N\to\infty} \left(\#\{\gamma : \gamma \in O_{\psi_g}, \gamma \text{ is not a vertex orbit, } \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}\right)^{1/N}.$$
$$\limsup_{N\to\infty} \left(\#\{\gamma : \gamma \in O_{\psi_g}, \alpha(\gamma) = N, \text{ and } \beta(\gamma) = 0\}\right)^{1/N}$$

The first line of the string of equalities above is the definition of $\text{gr}(a; b; D_g)$; thus, we have completed the proof. □

Corollary 4.2. *Let G be a topological graph with map $g : G \to G$ such that Z_g has finitely many periodic vertices, and the periodic vertices of G are all finite valence. If A is a*

smoothly embedded transverse graph in Z_g dual to $\alpha \in H^1(Z_g, \mathbb{R})$ such that α is positive on closed orbits of ψ_g and a is the pullback of α to D_g, then

$$gr_{\psi_g}(A) = gr(a; D_g).$$

Proof. Letting 0_c represent the trivial cohomology class in $H^1(Z_g, \mathbb{R})$ and $H^1(D_g, \mathbb{R})$, we have

$$\text{gr}(a; D_g) \coloneqq \limsup_{N\to\infty} \left(\#\{\gamma : \gamma \text{ is a directed cycle of } D_g \text{ and } a(\gamma) = N\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma : \gamma \text{ is a directed cycle of } D_g, a(\gamma) = N, \text{ and } 0_c(\gamma) = 0\}\right)^{1/N}$$

$$\coloneqq \text{gr}(a; 0_c; D_g).$$

Then by Lemma 4.2,

$$\text{gr}(a; 0_c; D_g) = \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : 0_c(\gamma) = 0, \alpha(\gamma) = N\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : \alpha(\gamma) = N\}\right)^{1/N}.$$

Due to the fact that α is dual to A, the above is equal to

$$= \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi_g} : \gamma \text{ intersects } A \text{ at } N \text{ times }\}\right)^{1/N} \coloneqq \text{gr}_{\psi_g}(A). \qquad \square$$

4.3 Growth of Edge Paths

Finally, for a topological graph map $g : G \to G$, we define a growth rate of combinatorial lengths of edge paths under iterates of g, and we call this concept of growth the stretch factor of g.

Definition 4.5. For a topological graph map $g : G \to G$, we define

$$\lambda(g) \coloneqq \sup_{e \in G} \Big\{ \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e))} \Big\},$$

where $\ell(g(e))$ is the combinatorial length of the edge path $g(e)$, and we call $\lambda(g)$ the **stretch factor** of $g : G \to G$.

The following lemma will later be used to prove Theorem B. However, its first application will be in the proof of Proposition 4.1, which describes the conditions under which the stretch factor of a topological graph map is equal to its growth rate of fixed points.

Lemma 4.3. *For a finite subgraph $G' \subseteq G$ of a graph G with topological graph map $g : G \to G$,*

$$\max_{e \in G'} \Big\{ \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e))} \Big\} = \limsup_{N \to \infty} \sqrt[N]{\sum_{e \in G'} \ell(g^N(e))} \Big\}$$

Proof. Let e' be the edge of G' realizing the maximum

$$\max_{e \in G'} \Big\{ \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e))} \Big\}.$$

Due to the fact that $\ell(g^N(e')) \leq \sum_{e \in G'} \ell(g^N(e))$, we have one direction of the proof,

$$\max_{e \in G'} \Big\{ \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e))} \Big\} = \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e'))} \leq \limsup_{N \to \infty} \sqrt[N]{\sum_{e \in G'} \ell(g^N(e))}.$$

For the other direction of the proof,

$$\limsup_{N\to\infty} \sqrt[N]{\sum_{e\in G'} \ell(g^N(e))} \leq \limsup_{N\to\infty} \sqrt[N]{\left|G'\right| \cdot \max_{e\in G'} \left\{\ell(g^N(e))\right\}}$$

$$= \limsup_{N\to\infty} \left(\left|G'\right|\right)^{1/N} \cdot \left(\max_{e\in G'} \left\{\ell(g^N(e))\right\}\right)^{1/N}$$

$$= \lim_{N\to\infty} \left(\left|G'\right|\right)^{1/N} \cdot \limsup_{N\to\infty} \left(\max_{e\in G'} \left\{\ell(g^N(e))\right\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\max_{e\in G'} \left\{\ell(g^N(e))\right\}\right)^{1/N}$$

Let k index the subsequence of $\left(\max_{e\in G'} \left\{\ell(g^N(e))\right\}\right)^{1/N}$ whose limit realizes the above $\limsup$. Then because G' is finite we can pass to a further subsequence indexed by j so that $\max_{e\in G'} \left\{\ell(g^{N_{k_j}}(e))\right\}$ is realized by the same edge e_1 for all j. Then we can complete the second direction of the proof,

$$\limsup_{N\to\infty} \left(\max_{e\in G'} \left\{\ell(g^N(e))\right\}\right)^{1/N} = \limsup_{j\to\infty} \left(\max_{e\in G'} \left\{\ell(g^{N_{k_j}}(e))\right\}\right)^{1/N}$$

$$\leq \limsup_{j\to\infty} \left(\left\{\ell(g^{N_{k_j}}(e_1))\right\}\right)^{1/N} \leq \max_{e\in G'} \left\{\limsup_{N\to\infty} \left(\ell(g^N(e))\right)^{1/N}\right\}. \qquad \square$$

The above lemma implies that for a finite graph G with topological graph map $g : G \to G$, we can alternately define

$$\lambda(g) = \limsup_{N\to\infty} \sqrt[N]{\sum_{e\in G} \ell(g^N(e))}.$$

The following proposition will later be used in the proof of Theorem C.

Proposition 4.1. *For a finite graph G with a train track topological graph map $g : G \to G$, if some power of $g : G \to G$ has a fixed point, then the stretch factor $\lambda(g)$ is equal to the growth rate of fixed points $gr(g)$.*

Proof.

$$\lambda(g) \coloneqq \sup_{e \in G} \left\{ \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e))} \right\} = \max_{e \in G} \left\{ \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e))} \right\}.$$

By Lemma 4.3,

$$\max_{e \in G} \left\{ \limsup_{N \to \infty} \sqrt[N]{\ell(g^N(e))} \right\} = \limsup_{N \to \infty} \sqrt[N]{\sum_{e \in G} \ell(g^N(e))}.$$

Due to the fact that g is a train track topological graph map, the combinatorial length $\ell(g^N(e))$ of the edge path $g^N(e)$ is the number of directed paths of length N based at e in the transition graph D_g, so the above is equal to

$$= \limsup_{N \to \infty} \sqrt[N]{\#\{\text{directed paths in } D_g \text{ of length } N\}}.$$

The number of directed paths of length N in D_g is non-decreasing, implying that the limit as N goes to infinity of the number of directed paths of length N in D_g is not zero. Therefore, we can apply Lemma 4.1 to get

$$\limsup_{N \to \infty} \sqrt[N]{\#\{\text{directed paths in } D_g \text{ of length } N\}}$$

$$= \limsup_{N \to \infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\text{directed paths in of length } i \text{ in } D_g\}}.$$

If a fixed point of some power of g is in the interior of an edge, then that power of $g : G \to G$ must map that edge over itself, and so D_g has some directed cycle. If instead

a fixed point of some power of g is a vertex of G, then some power of g must map an edge adjacent to that vertex over itself. In either case, some power of g maps an edge over itself, and therefore D_g has a directed cycle, and therefore we can apply [9, Lemma 3.1] to get

$$\limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\text{directed paths in of length } i \text{ in } D_g\}}$$
$$= \limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\text{directed cycles of length } i \text{ in } D_g\}}.$$

Finally, because D_g has a directed cycle, the limit as N goes to infinity of the number of directed cycles of length N in D_g is not zero, and so we can apply Lemma 4.1 again to get

$$\limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\text{ directed cycles of length } i \text{ in } D_g\}}$$
$$= \limsup_{N\to\infty} \sqrt[N]{\#\{\text{ directed cycles in of length } N \text{ in } D_g\}}.$$

We now consider the transverse smooth embedding $G \times \{0.5\}$ of G into $Z_g := G \times [0,1]/(x,1) \sim (g(x),0)$. Let μ_0 be the cohomology class dual to $G \times \{0.5\}$ which exists by Claim 2.1. Then, letting $\overline{\mu}_0$ be the pullback of μ_0 to $H^1(D_g, \mathbb{R})$, the length of a directed cycle of D_g is the number of times that the directed cycle intersects $G \times \{0.5\}$, which is also its combinatorial length as an edge path in D_g, so

$$\limsup_{N\to\infty} \sqrt[N]{\#\{\text{ directed cycles in of length } N \text{ in } D_g\}} = \text{gr}(\overline{\mu}_0; D_g).$$

Then, using the fact that $G \times \{0.5\}$ is a smoothly embedded transverse graph and μ_0 is positive on closed orbits, Corollary 4.2 gives us

$$\text{gr}(\overline{\mu}_0; D_g) = \text{gr}_{\psi_g}(G \times \{0.5\}).$$

Due to the fact that any closed orbit of ψ_g intersects $G \times \{0.5\}$ exactly once for each time it intersects $G \times \{0\}$ and vice versa,

$$\text{gr}_{\psi_g}(G \times \{0.5\}) = \text{gr}_{\psi_g}(G \times \{0\}).$$

Corollary 4.1 tells us that $\text{gr}_{\psi_g}(G \times \{0\})$ is equal to the growth rate of fixed points of the first return map of G determined by the smooth embedding $G \times \{0\}$, and Remark 2.1 tells us that the first return map of G determined by the smooth embedding $G \times \{0\}$ is $g : G \to G$. Therefore, the growth rate of fixed points of the first return map of $G \times \{0\}$ is equal to $\text{gr}(g)$, and we have completed the proof. □

CHAPTER 5

END-PERIODIC TRAIN TRACK MAPS

The topic of end-periodic homeomorphisms on infinite-type surfaces is an area of active research. However, this topic has not yet been explored for infinite graphs. In this section, we will define and discuss what it means for a topological graph map on an infinite graph to be end-periodic. At the end of this chapter, we will prove Theorem B, which we will later use in Section 7.3.2 to prove Theorem C.

Definition 5.1. An **end** of G is an equivalence class of sequences $\cdots U_3 \subseteq U_2 \subseteq U_1$ where U_i is a connected component of $G \setminus K_i$ for some exhaustion of G by finite subgraphs $K_1 \subseteq K_2 \subseteq K_3 \subseteq \cdots$. Two such sequences of nested connected components are considered equivalent if each connected component of one sequence contains all but finitely many connected components of the other sequence.

Such a sequence of nested connected components representing and end E is called a **neighborhood basis** of E. A **neighborhood** of E is an open set U such that every neighborhood basis of E is eventually contained in U, and a neighborhood basis for any other end has only finitely many terms intersecting U.

Definition 5.2. We will say that a topological graph map $g : G \to G$ on an infinite graph G is **end-periodic** if there is some integer $m > 0$ such that for each end E of G, there is a neighborhood U_E of E so g is a graph isometric embedding (sends edges to edges) on the largest subgraph of G contained in U_E, and either (i) or (ii) holds:

1. $g^m(U_E) \subsetneq U_E$ and $\{g^{nm}(U_E)\}$ forms a neighborhood basis of E for $n \in \mathbb{Z}_{\geq 0}$. In this case, E is called an **attracting** end.

2. $g^{-m}(U_E) \subsetneq U_E$ and $\{g^{-nm}(U_E)\}$ forms a neighborhood basis of E for $n \in \mathbb{Z}_{\geq 0}$. In this case, E is called **repelling** end.

We will now prove a lemma that relates the property of being end-periodic to the structure of the transition graph. Then, we will use the lemmas to prove Theorem B.

Lemma 5.1. *An end-periodic topological graph map $g : G \to G$ on a graph G with finitely many ends has a transition graph D_g with the following properties. There is some finite subgraph $D'_g \subset D_g$ such that the connected components of $D_g \setminus D'_g$ are either one of finitely many infinite directed rays that originate on D'_g, or one of finitely many directed trees that terminate on D'_g (any maximal directed edge path in the tree terminates on the same vertex in D'_g).*

Proof. For each end e of G, pick a neighborhood satisfying Definition 5.2, and let $\{U_e\}_{e \in E}$ be the set of those neighborhoods, and let m be the integer from Definition 5.2. Due to the fact that G has finitely many ends, we can choose these end neighborhoods to be pairwise disjoint, and there is some finite subgraph G_1 of G such that each connected component of $G \setminus G_1$ is contained in a unique U_e.

Recall from the definition of end periodic, that for each attracting end e, $\{g^{mn}(U_e)\}_{n \in \mathbb{Z}_{\geq 0}}$ is a neighborhood basis for e, whereas if e is repelling, $\{g^{-mn}(U_e)\}_{n \in \mathbb{Z}_{\geq 0}}$ is a neighborhood basis for e. Then, since G_1 is finite, there is some

integer $k > 1$ such that for each attracting end e, $g^{mn}(U_e)$ and $g(g^{mn}(U_e))$ are each disjoint from G_1 for $n \in \mathbb{Z} \geq k$, and for each repelling end e, $g^{-mn}(U_e)$ and $g(g^{-mn}(U_e))$ are each disjoint from G_1 for $n \in \mathbb{Z} \geq k$. We can also choose k to be large enough that $g(G_1)$ does not intersect $g^{mk}(U_e)$ for any attracting end U_e, or $g^{-mn}(U_e)$ for repelling end U_e.

Then we can take $\{g^{mk}(U_e)\}_{e \in E+}$ to be neighborhoods for the attracting ends, and we can take $\{g^{-mk}(U_e)\}_{e \in E-}$ to be neighborhoods for the repelling ends. These end neighborhoods still satisfy the definition of end-periodic and are pairwise disjoint.

There is again some finite subgraph G_2 of G such that each connected component of $G \setminus G_2$ is contained in a unique $\{g^{mk}(U_e)\}$ for attracting end e or $\{g^{-mk}(U_e)\}$ for repelling end e. Therefore, G_1 is necessarily contained in G_2.

Let D_{g_1} be the subgraph of D_g induced by the vertices of D_g that correspond to edges of G_1, and let D_{g_2} be the subgraph of D_g induced by the vertices of D_g corresponding to edges of G_2. Then $D_{g_1} \subseteq D_{g_2}$, and D_{g_1} and D_{g_2} contain finitely many vertices because G_1 and G_2 are finite. Therefore, D_{g_1} and D_{g_2} are finite directed graphs because $g : G \to G$ is a topological graph map, meaning each vertex of D_g has finitely many outward directed edges originating on it.

Observe that D_g is finite valence. Every edge of D_g not contained in D_{g_2} either originates on D_{g_2} or on a vertex outside D_{g_2}. Based on the same reasoning above, there are finitely many of the first type of edges. Every vertex of D_g outside of D_{g_1} (and therefore every vertex outside of D_{g_2}) corresponds to an edge of G contained in one of the chosen end neighborhoods U_e, meaning that g is a graph isometric embedding on that edge. Therefore,

each vertex of D_g outside of D_{g_1} has exactly one directed edge originating on it. The fact that g is a graph isometric embedding on the end neighborhoods also implies that at most finitely many of these directed edges originating at vertices outside D_{g_2} can terminate on a single vertex. Specifically, no more than the number of ends because g does not send two edges of G within the same end neighborhood U_e over the same edge. Therefore, D_g is finite valence.

Due to the fact that $\{g^{mn}(U_e)\}_{n\in\mathbb{Z}\geq 0}$ or $\{g^{-mn}(U_e)\}_{n\in\mathbb{Z}\geq 0}$ is a neighborhood basis for the end e, there can be no directed cycles of D_g containing vertices outside of D_{g_1}. This observation implies that the connected components of $D_g \setminus D_{g_1}$ are directed trees, and the same is true of $D_g \setminus D_{g_2}$.

For a neighborhood U_e of an attracting end e, there are finitely many edges of G intersecting U_e that are not contained in $g^m(U_e)$, and the images of these edges under powers of g^m must cover U_e. Similarly, for a neighborhood U_e of a repelling end e, there are finitely many edges of G intersecting U_e that are not contained in $g^{-m}(U_e)$. The pre-images of these edges under powers of g^m must cover U_e. Therefore, every vertex in D_g outside of D_{g_2} must be connected to D_{g_2} by some directed path. This observation implies that each connected component of $D_g \setminus D_{g_2}$ is a directed tree that is either directed inward towards D_{g_2}, meaning there is a single vertex in D_{g_2} such that every maximal directed edge path in the tree terminates at that vertex, or the tree is directed outwards away from D_{g_2}, meaning there is a subset of vertices in D_{g_2} such that every maximal directed edge path in the tree is based at one of those vertices.

Let D_g^n be the union of D_{g_2} with all the vertices and edges of D_g contained within distance n of D_g. Due to the fact that the outward directed trees originate on the finite graph of finite valence D_{g_2} because every vertex outside of D_{g_2} has exactly one directed edge originating on it, there is some integer j such that each connected component of $D_g \setminus D_g^j$ is either one of finitely many disjoint outward directed rays or an inward directed tree. Letting $D_g^j =: D'_g$, The graph D'_g is finite because D_{g_2} is finite, and D_g is finite valence. This observation also implies that there are finitely many connected components of $D_g \setminus D'_g$. □

Lemma 5.2. *Suppose G is an infinite graph G of finite valence with finitely many ends and $g : G \to G$ is an end-periodic topological graph map such that some power of g fixes a point of G. Letting D_g be the flow graph of $g : G \to G$ and D'_g be the finite subgraph as in Lemma 5.1, we have*

$$\limsup_{N\to\infty} \sqrt[N]{\#\{\textit{ directed paths in } D_g \textit{ of length } N \textit{ that are based in } D'_g\}}$$

$$= \limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\textit{ directed paths in of length } i \textit{ in } D'_g\}}.$$

Proof. We will prove this equality by showing that each side is less than or equal to the other. Let $V(D_g)$ be the vertex set of D_g.

If a fixed point of some power of g is in the interior of an edge, then that power of $g : G \to G$ must map that edge over itself, and so D_g has some directed cycle. If instead a fixed point of some power of g is a vertex of G, then some power of g must map an edge adjacent to that vertex over itself. In either case, some power of g maps an edge over itself;

therefore, D_g has a directed cycle. Due to the fact that some power of $g : G \to G$ does have a fixed point, D_g has a directed cycle, which is necessarily contained in D'_g.

One direction is simple; due to the fact that D'_g contains a directed cycle, the limit as N goes to infinity of the number of directed paths of length N in D'_g is not zero, meaning we can apply Lemma 4.1 to get

$$\limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\text{ directed paths in of length } i \text{ in } D'_g\}}$$

$$= \limsup_{N\to\infty} \sqrt[N]{\#\{\text{ directed paths in } D'_g \text{ of length } N\}}.$$

Then, because

$$\#\{\text{ directed paths of length } N \text{ in } D'_g\}$$

$$\leq \#\{\text{ directed paths in } D_g \text{ of length } N \text{ that are based in } D'_g\},$$

we have

$$\limsup_{N\to\infty} \sqrt[N]{\#\{\text{ directed paths of length } N \text{ in } D'_g\}}$$

$$\leq \limsup_{N\to\infty} \sqrt[N]{\#\{\text{ directed paths in } D_g \text{ of length } N \text{ that are based in } D'_g\}}.$$

For the other direction, note that given any directed path of length N based in D'_g, for some $0 \leq i \leq N$ that directed path is composed of a directed path of length i contained in D'_g concatenated with a directed path of length $N - i$ contained in one of the finitely many directed rays of D_g originating on D'_g described in Lemma 5.1. Letting L be the number of those directed rays of D_g originating on D'_g, we have,

$$\limsup_{N\to\infty} \sqrt[N]{\#\{\text{ directed paths in } D_g \text{ of length } N \text{ that are based in } D'_g\}}$$

$$\leq \limsup_{N\to\infty} \sqrt[N]{L \cdot \sum_{i=0}^{N} \#\{\text{ directed paths of length } i \text{ in } D'_g\}}$$

$$= \lim_{N\to\infty} \sqrt[N]{L} \cdot \limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\text{ directed paths of length } i \text{ in } D'_g\}}$$

$$= \limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{\text{ directed paths of length } i \text{ in } D'_g\}}. \qquad \square$$

Now, we will use the previous two lemmas to prove Theorem B.

Theorem B. *Suppose G is an infinite graph of finite valence with no valence-0 vertices and with finitely many ends, and $g : G \to G$ is an end-periodic train track topological graph map. If some power of g fixes a point of G, then the stretch factor $\lambda(g)$ is equal to the growth rate of fixed points $gr(g)$.*

Proof. Let D_g be the transition graph of G, and let $V(D_g)$ be its vertex set. Let D'_g be the finite subgraph of D_g as in Lemma 5.1.

The fact that a power of g fixes a point of G implies that Z_g has a closed orbit under ψ_g. By Remark 3.2 and Proposition 3.1, the fact that G has no valence-0 vertices implies that D_g has a directed cycle, which is necessarily contained in D'_g.

For any edge e_1 of G that is contained in one of the finitely many directed rays of D_g that originate on D'_g as described in Lemma 5.1, because $g^N(e_1)$ is still in that same directed ray, $\ell(g^N(e_1)) = 1$. For any edge e' in G such that $e' \in V(D'_g)$, $\ell(g^N(e')) > 1$ for all $N \in Z_{\geq 0}$. So $1 \leq \limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e'))}$, and therefore we have

$$\limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e_1))} = \limsup_{N\to\infty} \sqrt[N]{1} = 1 \leq \limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e'))}.$$

For any edge e_2 of G that is contained in one of the directed trees of D_g that terminates on D'_g as described in Lemma 5.1, there is some minimal integer $m > 0$ so that $g^m(e_2) \in V(D'_g)$. Because m is minimal, $g^m(e_2)$ must be a single edge $e' \in V(D'_g)$ of G. Therefore, for $0 \leq N \leq m$, $\ell(g^N(e_2)) = 1$, so

$$\limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e_2))} = \limsup_{N\to\infty} \sqrt[m+N]{\ell(g^{m+N}(e_2))}$$

$$= \limsup_{N\to\infty} \sqrt[m+N]{\ell(g^N(e'))} \leq \limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e'))}$$

By the above remarks, we can calculate the stretch factor by taking the supremum only over the finite set of edges of G that are in the vertex set of D'_g, or

$$\lambda(g) = \sup_{e\in V(D_g)} \left\{ \limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e))} \right\} = \max_{e\in V(D'_g)} \left\{ \limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e))} \right\}.$$

By Lemma 4.3,

$$\max_{e\in V(D'_g)} \left\{ \limsup_{N\to\infty} \sqrt[N]{\ell(g^N(e))} \right\} = \limsup_{N\to\infty} \sqrt[N]{\sum_{e\in V(D'_g)} \ell(g^N(e))}.$$

Due to the fact that g is a train track topological graph map, the combinatorial length $\ell(g^N(e))$ of the edge path $g^N(e)$ is the number of directed paths in the transition graph D_g of length N based at e, so the above is equal to

$$= \limsup_{N\to\infty} \sqrt[N]{\#\{ \text{ directed paths in } D_g \text{ of length } N \text{ that are based in } D'_g \}}.$$

By Lemma 5.2, the above is equal to

$$= \limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{ \text{ directed paths in of length } i \text{ in } D'_g\}}.$$

The fact that D'_g contains a directed cycle allows us to apply [9, Lemma 3.1] to get

$$\limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{ \text{ directed paths in of length } i \text{ in } D'_g\}}$$

$$= \limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{ \text{ directed cycles of length } i \text{ in } D'_g\}}.$$

We now apply the fact that D'_g contains a directed cycle again. The existence of this directed cycle implies that the limit as N goes to infinity of the number of directed cycles of length N in D'_g is not zero. Thus we can apply Lemma 4.1 again to get

$$\limsup_{N\to\infty} \sqrt[N]{\sum_{i=0}^{N} \#\{ \text{ directed cycles of length } i \text{ in } D'_g\}}$$

$$= \limsup_{N\to\infty} \sqrt[N]{\#\{ \text{ directed cycles in of length } N \text{ in } D'_g\}}.$$

Due to the fact that all directed cycles of D_g are contained in D'_g, we have

$$\limsup_{N\to\infty} \sqrt[N]{\#\{ \text{ directed cycles in of length } N \text{ in } D'_g\}}$$

$$= \limsup_{N\to\infty} \sqrt[N]{\#\{ \text{ directed cycles in of length } N \text{ in } D_g\}}$$

We now consider the transverse smooth embedding $G \times \{0.5\}$ of G into $Z_g := G \times [0,1]/(x,1) \sim (g(x),0)$. Let μ_0 be the cohomology class dual to $G \times \{0.5\}$ which exists by Claim 2.1. Then, letting $\overline{\mu}_0$ be the pullback of μ_0 to $H^1(D_g, \mathbb{R})$, the length of a directed cycle of D_g is the number of times that the directed cycle intersects $G \times \{0.5\}$ which is also its combinatorial length as an edge path in D_g, so

$$\limsup_{N\to\infty} \sqrt[N]{\#\{ \text{ directed cycles in of length } N \text{ in } D_g\}} = \text{gr}(\overline{\mu}_0; D_g).$$

Then, using the fact that $G \times \{0.5\}$ is a smoothly embedded transverse graph and μ_0 is positive on closed orbits, Corollary 4.2 gives us

$$\text{gr}(\overline{\mu}_0; D_g) = \text{gr}_{\psi_g}(G \times \{0.5\}).$$

Due to the fact that any closed orbit of ψ_g intersects $G \times \{0.5\}$ exactly once for each time it intersects $G \times \{0\}$ and vice versa,

$$\text{gr}_{\psi_g}(G \times \{0.5\}) = \text{gr}_{\psi_g}(G \times \{0\}).$$

Corollary 4.1 tells us that $\text{gr}_{\psi_g}(G \times \{0\})$ is equal to the growth rate of fixed points of the first return map of G determined by the smooth embedding $G \times \{0\}$, and Remark 2.1 tells us that the first return map of G determined by the smooth embedding $G \times \{0\}$ is $g : G \to G$. Therefore, the growth rate of fixed points of the first return map of $G \times \{0\}$ is equal to $\text{gr}(g)$, and we have completed the proof. $\square$

CHAPTER 6

THE SPINNING CONSTRUCTION

This chapter will define what taking a co-oriented cut and paste of two graphs means. Then we will use this to define a procedure called spinning G_1 around G_2, which takes in two graphs G_1 and G_2 and outputs a non-unique, infinite, smoothly embedded transverse graph $G_1 +_{sp} G_2$ as depicted in Figure 6.2. In Chapter 7 will use this spinning procedure in combination with the operation defined in Section 7.1 to construct an infinite graph L_a, which we will use in Theorem C.

6.1 Co-Oriented Cut and Paste

Let G_1 and G_2 be two smoothly embedded transverse graphs in Z_g that are co-oriented in the positive semiflow direction inside 2-cells of Z_g. Also, suppose that $G_1 \cap G_2$ is a discrete set of points that are contained in the interiors of 2-cells of Z_g and do not intersect the vertices of G_1 or G_2. Then let W_1 and W_2 be their respective oriented I-bundle neighborhoods satisfying Definition 1.15. We will define an operation on G_1 and G_2 called a **co-oriented cut and paste**, which alters G_1 and G_2 within small neighborhoods of their intersection points using their co-orientations and yields a graph, $G_1 + G_2$.

To take a co-oriented cut and paste of G_1 and G_2, we perform the operation depicted in Figure 6.1 within $W_1 \cap W_2$, which resolves the intersections between G_1 and G_2 so that the new graph has a well-defined co-orientation.

Remark 6.1. Each vertex of G_1+G_2 has a neighborhood in Z_g such that the neighborhood's intersection with G_1+G_2 is either equal to its intersection with G_1 or with G_2 but not both.

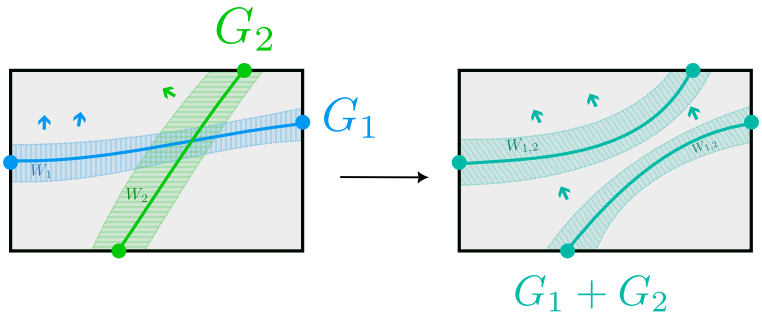

Figure 6.1: Co-oriented cut and paste sum of G_1 and G_2.

Remark 6.2. A co-oriented cut and paste $G_1 + G_2$ is smoothly embedded and transverse to ψ_g with neighborhood $W_{1,2}$ depicted in Figure 6.1.

Claim 6.1. Suppose G_1 and G_2 are two smoothly embedded transverse graphs in Z_g such that we can take their co-oriented cut and paste. If G_1 and G_2 are dual to cohomology classes ξ_1 and ξ_2 respectively, the co-oriented cut and paste $G_1 + G_2$ is dual to $\xi_1 + \xi_2$.

We can see that the above is true because Claim 2.1 tells us the $G_1 + G_2$ is dual to some class. The set of intersection points of $G_1 + G_2$ with $Z_g^{(1)}$ is the union of intersection points of G_1 with $Z_g^{(1)}$ and G_2 with $Z_g^{(1)}$. Therefore, the fact that any loop in Z_g can be homotoped into $Z_g^{(1)}$ implies the claim.

Claim 6.2. For a cross section G_1 of ψ, and a possibly infinite smoothly embedded transverse graph G_2, which may not have a first return map, a co-oriented cut and paste $G_1 + G_2$ does have a well-defined first return map. If G_2 is also finite, their co-oriented cut and paste is a cross section.

The above claim is valid because ψ eventually flows any point on $G_1 + G_2$ into G_1, and therefore back into itself. If G_2 is also finite, then the co-oriented cut and paste of G_1 and G_2 is a finite graph which is smoothly embedded and transverse by Remark 6.2. Then,

because the co-oriented cut and paste has a first return map, Claim 1.1 implies that it is a cross section.

6.2 Spinning G_1 around G_2

Suppose G_1 and G_2 are two smoothly embedded transverse finite graphs in Z_f that do not intersect 0-cells of $Z_f^{(0)}$. Let W_1 and W_2 be the neighborhoods of G_1 and G_2 that satisfy the definition of being transverse to ψ, and let h_1 and h_2 be the homeomorphisms $h_1 : G_1 \times (-1, 1) \to W_1$ and $h_2 : G_1 \times (-1, 1) \to W_2$ such that $h_1(G_1 \times \{0\})$ is G_1, and $h_2(G_2 \times \{0\})$ is G_2.

To **spin** G_1 **around** G_2 is to perform an operation on G_1 and G_2 which will yield an infinite smoothly embedded transverse graph $G_1 +_{sp} G_2$ in the manner outlined below and as depicted in Figure 6.2. While constructed from G_1 and G_2, this graph will not be unique.

First, isotope h_2 while keeping the restriction of h_2 to $G_2 \times \{0\}$ fixed, so that for every $x \in (-1, 1)$, $h_2(G_2 \times \{x\})$ is smoothly embedded, and any segment of ψ that passes from the bottom to the top of $h_2(G_2 \times (-1, 1))$ intersects $h_2(G_2 \times \{x\})$ exactly once along the way.

Then, shrink the domain of h_2 to $G_2 \times (-\epsilon_2, \epsilon_2)$ so that the closures of the connected components of the intersection of $h_2(G_2 \times (-\epsilon_2, \epsilon_2))$ with $Z_g^{(1)}$ have exactly one point of intersection with G_2, are pairwise disjoint, and are disjoint from the 0-cells of $Z_g^{(0)}$. Also let ϵ_2 be small enough that for each rectangle r, each connected component of the intersection

of $h_2(G_2 \times (-\epsilon_2, \epsilon_2))$ with the interior of r has some flowline of ψ contained in r that enters and exits the connected component while inside of r.

We define $G_{2,n} \coloneqq h_2(G_2 \times \{\frac{\epsilon}{n}\})$, which is smoothly embedded and transverse with the transverse neighborhood $h_2(G \times (-\frac{\epsilon_2}{n+\frac{1}{4n^2}}, \frac{\epsilon_2}{n+\frac{1}{4n_2}}))$.

Then define

$$\overline{G}_2 = \bigcup_{|n| \in \mathbb{Z}_{\geq 2}} G_{2,n}.$$

Isotope h_1 so that the images of h_1 and of h_2 do not intersect on $Z_g^{(1)}$ while maintaining that G_1 is smoothly embedded and transverse. Also, let this isotopy of h_1 be such that the intersection points between G_1 and $\overline{G}_2$ are contained in the interiors of 2-cells and are not on either of their vertices. This isotopy is possible because the vertices of G_1 and $\overline{G}_2$ are locally finite in Z_g. Finally, let this isotopy of h_1 be such that the vertices of $G_1 = h_1(G_1 \times \{0\})$ are on vertex leaves w. This is possible because vertex leaves are dense in Z_f [3, Section 4].

Let the output of spinning G_1 around G_2 be the co-oriented cut and paste of $\overline{G}_2$ and $G_1 \setminus \{G_1 \cap G_2\}$, and call this output $G_1 +_{sp} G_2$. Note that the alterations we have made to h_1 and h_2 ensure that $\overline{G}_2$ and $G_1 \setminus \{G_1 \cap G_2\}$ have the properties required to be able to take their co-oriented cut and paste.

Remark 6.3. The graph $\overline{G}_2$ is a transverse smoothly embedded graph because the transverse neighborhoods of the $G_{2,n}$ mentioned above are disjoint from one another. Thus, by Remark 6.2, $G_1 +_{sp} G_2$ is a smoothly embedded transverse graph as well.

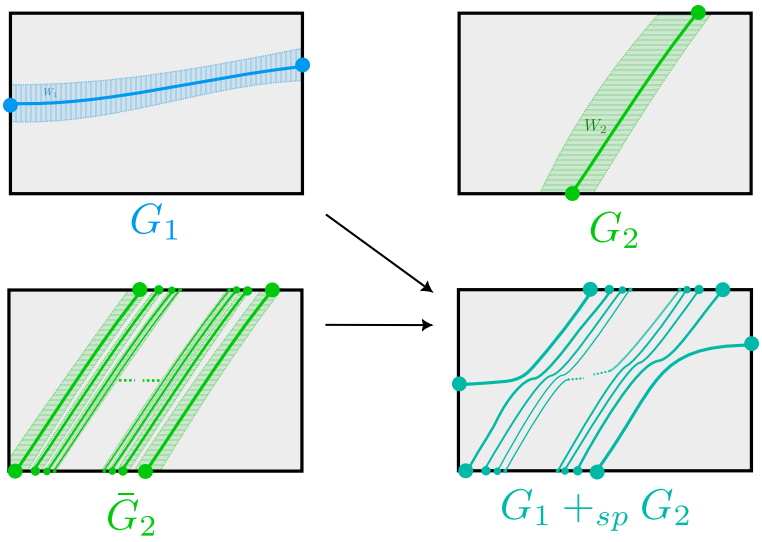

Figure 6.2: Spinning G_1 around G_2.

Observe that $G_1 +_{sp} G_2$ is infinite because $\overline{G}_2$ is infinite, but $G_1 +_{sp} G_2$ has bounded valence by Remark 6.1.

We will also define some finite graphs that we will use as tools in the proof of Proposition 7.1. These graphs are constructed similarly to spinning G_1 around G_2. As with spinning, suppose G_1 and G_2 are two smoothly embedded transverse finite graphs in Z_f that do not intersect 0-cells of $Z_f^{(0)}$, and let h_1 and h_2 be the same homeomorphisms described above. For $i \in \mathbb{Z}_{>0}$, define $G_1 +_{(2)} iG_2$ to be the co-oriented cut and paste of G_1 with

$$2 * iG_2 := \bigcup_{n \in \mathbb{Z}, 2 \leq |n| \leq 1+i} h_2(G_2 \times \{\frac{\epsilon_2}{n}\}).$$

Remark 6.4. The finite graph $2 * iG_2$ is smoothly embedded and transverse with analogous I-bundle neighborhood to the one in Remark 6.3. Thus, by Remark 6.2, the finite graph $G_1 +_{(2)} iG_2$ is smoothly embedded and transverse to ψ_g.

CHAPTER 7

THE FIRST RETURN MAP OF THE BOUNDARY GRAPH

For the duration of this chapter, we will fix Θ as a cross section of ψ dual to the cohomology class μ in the interior of the fibered cone and Σ as a finite smoothly embedded graph in Z_f that is transverse to ψ and dual to the cohomology class η in the boundary of the fibered cone. Unlike cross sections of ψ, Σ does not have a well-defined first return map under ψ because, unlike a cross section, some points on Σ may never flow back into Σ under ψ. We can use spinning to create an infinite graph from Σ and Θ that does have a first return map.

Fix $L = \Theta +_{sp} \Sigma$ as a result of spinning Θ around Σ as in Section 6.2. By Remark 6.3, L is a smoothly embedded graph. By Claim 6.2, L does have a well-defined first return map under ψ, which we will call f_L.

However, to get the first return map f_L to be a topological graph map (i.e., to get the first return map to send vertices to vertices), we need to make some modifications to Z_f and simultaneously to L. We will do this by blowing air into Z_f with respect to L as described in Section 7.2.

Fix Z_f^a as a cell complex resulting from blowing air into Z_f with respect to L. Let ψ_a be the semiflow on Z_f^a as in Section 7.1, and let π_a be the quotient map sending Z_f^a back to Z_f by a homotopy equivalence as in Section 7.1.

In Section 7.3, we will argue that ${\pi_a}^{-1}(L)$ has a well-defined first return map under ψ_a. In Section 7.2.2, we will describe how to put a cell complex structure on ${\pi_a}^{-1}(L)$ so

that it is an infinite graph L_a of finite valence. We will do this in such a way that the first return map f_{L_a} must be a topological graph map.

In Section 7.3, we will discuss the properties of L_a and f_{L_a}, and at the end of this chapter we will use these properties along with results from Chapter 3, Chapter 4, and Chapter 5 to prove Theorem C.

7.1 Blowing Air into a Triangle

In this section, we will define an operation on rectangular 2-cells of Z_g called blowing air into a triangle, which will simultaneously modify a smoothly embedded graph in Z_g and Z_g. In Section 7.2, we will use this operation to construct the infinite graph L_a, which we will use in Theorem C.

Let r be a rectangular 2-cell of $Z_g^{(2)}$ with bottom 1-cell e such that r has more than one sheet below it. Let T be an embedded triangle in r such that one side of T is a segment e' of e, another side of T is a segment in the interior of r that is a forward flow of an endpoint of e' under ψ_g, and the final side of T is smoothly embedded in the interior of r.

To **blow air into** T is to perform the following operation on Z_g, depicted in Figure 7.1. Make a copy, T_i, of T for each sheet below r. Let $Z_g|_{e'}$ be the result of cutting Z_g along e', and define space resulting from this operation to be $Z_g^T \coloneqq \left(Z_g|_{e'} \setminus \text{int}(T) \right) \cup_{h_i} \{T_i\}$, where h_i glues the sides of T_i in the interior of r to the same place those sides of T were, and h_i glues the bottom side of T_i to the copy of e' on the ith sheet below r.

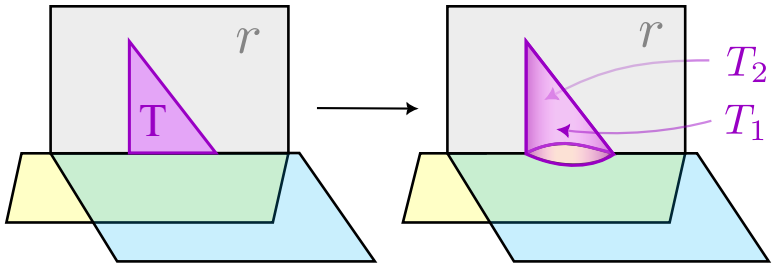

Figure 7.1: Blowing air into T.

Note that there is a quotient map $\pi_T : Z_g^T \to Z_g$ defined by quotienting $T_i = T_j$ for all i, j, and this projection map is a homotopy equivalence. This map is also a homeomorphism and an immersion in the interiors of its 2-cells.

There is also a semiflow ψ_T defined on ${Z_g}^T$ by letting $\psi_T^t(x)$ be $\pi_T^{-1}(\psi_g^t(\pi_T(x)))$ for $t \in \mathbb{R}_{\geq 0}$ and $x \in {Z_g}^T$ such that $\psi_g^t(\pi_T(x)) \in Z_g \setminus T$. For $t \in \mathbb{R}_{\geq 0}$ and $x \in {Z_g}^T$ such that $\psi_g^t(\pi_T(x)) \in T$, let $\psi_T^t(x)$ be the unique point in the set $\pi_T^{-1}(\psi_g^t(\pi_T(x)))$ such that ψ_T is continuous.

Remark 7.1. For any $t \in \mathbb{R}_{\geq 0}$ and $x \in Z_f^a$ we have $\pi_T(\psi_T^t(x)) = \psi^t(\pi_T(x))$.

Remark 7.2. This procedure of blowing air in can be applied to finitely many disjoint triangles simultaneously.

7.2 Blowing Air into Z_f with Respect to L

We will now use the procedure defined in Section 7.1 to describe a process called **blowing air into Z_f with respect to L** which will modify Z_f and L simultaneously.

To blow air into Z_f with respect to L, we will choose a non-unique finite collection of disjoint triangles in Z_f as described in Section 7.2.1, then we will blow air into all of the triangles simultaneously, yielding the space Z_f^a with the quotient map π_a sending Z_f^a back

to Z_f as in Section 7.1. In Section 7.2.2, we describe how to get the graph L_a by putting a cell complex structure on $\pi_a^{-1}(L)$.

7.2.1 Choosing Triangles

Let $h_2 : \Sigma \times (-\epsilon_2, \epsilon_2) \to W_2$ be the homeomorphism determining an I-bundle neighborhood of Σ as described in Section 6.2. Similarly, let $h_1 : \Theta \times (-\epsilon_1, \epsilon_1) \to W_1$ be the homeomorphism determining the I-bundle neighborhood of Θ as described in Section 6.2.

Given a rectangular 2-cell r with bottom 1-cell e in $Z_f^{(1)}$ such that e has multiple sheets below it, consider the intersection of $h_2(\Sigma \times (-\epsilon_2, \epsilon_2))$ with e. In the process of creating L as in Section 6.2, we chose ϵ_2 so that the closures of the connected components of $h_2(\Sigma \times (-\epsilon_2, \epsilon_2))$ did not intersect 0-cells of $Z_f^{(0)}$. Therefore, this intersection between $h_2(\Sigma \times (-\epsilon_2, \epsilon_2))$ and e is a set of open intervals whose closures are contained in the interior of e.

Consider one such open interval I in e. Let the closure of I be the bottom of a triangle. See Figure 7.2 for examples of how this triangle will look depending on how Σ sits in r above that interval.

By the choice of ϵ_2 as in Section 6.2, the forward flowline of one of the endpoints of I inside r flows through and exits $h_2(\Sigma \times (-\epsilon_2, \epsilon_2))$ while still inside r. Let the vertical side of T be a forward flowline of that endpoint of I inside of r that ends in the interior of r. Specifically, let that flowline pass through every connected component of the intersection of L with the interior of r that is adjacent to I. Let the flowline terminate before passing through any connected components that are not adjacent to I.

In Section 6.2, we isotoped h_1 so that Θ does not intersect I. By this observation and the choice of the vertical side of T described above, it is possible to choose the final side of T to be some interval that is smoothly embedded and transverse to ψ in the interior of r and does not intersect L.

Remark 7.3. This choice of triangle is such that L only intersects the bottom and vertical sides of the triangle.

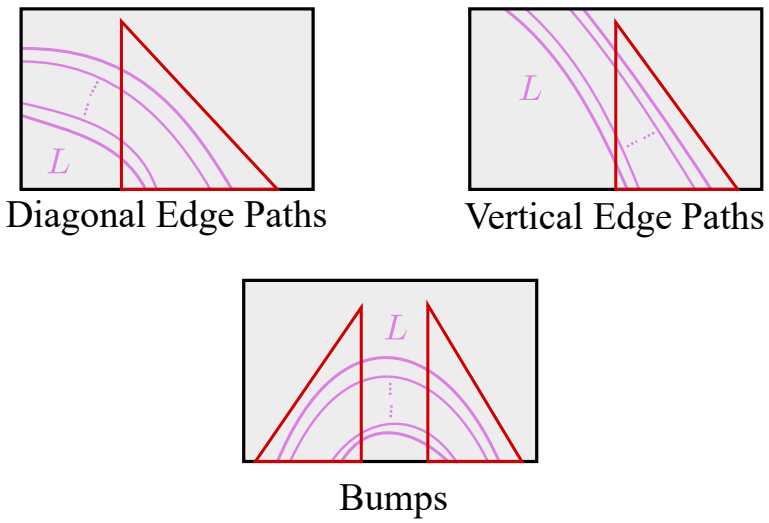

Figure 7.2: We choose a collection of triangles to blow air into based on L.

To blow air into Z_f with respect to L, we pick such a triangle for each interval of $h_2(\Sigma \times (-\epsilon_2, \epsilon_2)) \cap Z_f^{(1)}$ on the bottom of a rectangle with multiple sheets below it, and we blow air into all of these triangles simultaneously. Note that these triangles are disjoint because, in Section 6.2, we chose ϵ_2 so that the closures of the intervals are disjoint, meaning that the bottom sides of the triangles are disjoint. Also, there are finitely many of these intervals because Σ is finite, and so there are finitely many of these disjoint triangles. Therefore, it is possible to blow air into all of them simultaneously by Remark 7.2.

7.2.2 Cell Complex on L_a

Let L_a be the graph whose smoothly embedded image in Z_f^a is the set $\pi_a^{-1}(L)$, and has its vertices embedded as follows. The graph L_a has a vertex at every point which π_a sends to a vertical 1-cell of $Z_f^{(1)}$. We will call the collection of such vertices V_0. The graph L_a also has a vertex at every point x in $\pi_a^{-1}(L)$ such that x is not in V_0 and x has valence more than 2 as in Figure 7.3. We will call that collection of vertices V_1. All of the vertices in V_1 were vertices of Θ on a horizontal 1-cell of $Z_f^{(1)}$, or they are on the vertical sides of one of the finitely many triangles used in blowing air into Z_f with respect to L to get L_a. Therefore, because in the construction of L, we homotoped Θ so that its vertices were on vertex leaves (Section 6.2), and because there were finitely many triangles chosen which all have their vertical sides on vertex leaves, the vertices in V_1 all lie on a finite collection of flowlines on vertex leaves. Finally, the graph L_a will also have a vertex at every point of L that intersects a forward flowline under ψ_a of any vertex in V_0 or V_1.

Remark 7.4. Due to the fact that the vertices of V_0 and V_1 lie on a finite collection of flowlines on vertex leaves, the set of vertices of L_a will be locally finite in $\pi_a^{-1}(L)$, and therefore this vertex placement determines a cell complex structure on L_a.

Remark 7.5. Note that while L_a is an infinite graph, its vertices are finite valence because Z_f^a is a finite cell complex.

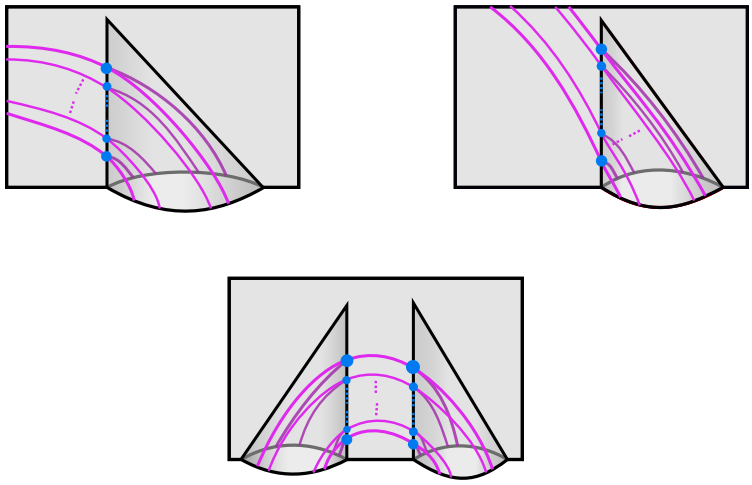

Figure 7.3: We put a set V_1 of vertices in the natural locations highlighted in blue.

7.3 L_a and its First Return Map

We begin by remarking on the existence of the first return map if L_a.

Remark 7.6. Due to the fact that L has a first return map under ψ, Remark 7.1 and Remark 7.2 imply that L_a also has a first return map under ψ_a which we will call f_{L_a}.

This section will discuss some of the properties of L_a and f_{L_a}. The following theorem summarizes the most relevant properties L and L_a that we will cover in this section before we begin discussing the stretch factor of f_{L_a}.

Theorem 7.1. *The graph L_a has a well-defined first return map f_{L_a} that is an end-periodic, weakly expanding, train track, topological graph map. The infinite graphs L and L_a are homotopy equivalent, and they both have finitely many connected components. Also, the first return maps f_{L_a} and f_L commute with the quotient map π_a.*

The map f_{L_a} must be a topological graph map for the stretch factor of f_{L_a} to be well-defined, and so the following lemma is necessary to the proof of Theorem C.

Lemma 7.1. *The first return map $f_{L_a} : L_a \to L_a$ is a topological graph map.*

Proof. By construction, every forward flow of a vertex of L_a under ψ_a is also a vertex of L_a. Therefore, f_{L_a} sends vertices to vertices. Due to the fact that f_{L_a} is a train track map that sends vertices to vertices, it also sends edges to edge paths, and so it is a topological graph map. □

The following observation about the first return maps f_L and f_{L_a} will be useful throughout this section. We will first use the observation to prove that the growth rates of fixed points of f_L and f_{L_a} are equal, which will contribute to the proof of Theorem C.

Remark 7.7. By Remark 7.1 and Remark 7.2, the first return maps f_L and f_{L_a} commute with π_a.

Lemma 7.2. *For any $n \in \mathbb{Z}_{>0}$, ${f_{L_a}}^n$ and ${f_L}^n$ have the same number of fixed points, and therefore the growth rate of fixed points of f_{L_a} and the growth rate of fixed points of f_L are equal.*

Proof. By Remark 7.7, if x is a fixed point of ${f_L}^n$ and x_1 and x_2 are two lifts of x to Z_f^a, then x_1 and x_2 are mapped to lifts of x under the map ${f_{L_a}}^n$.

By the construction of L_a and Z_f^a, if x is not inside some triangle T chosen to blow air into Z_f with respect to L to create Z_f^a, then x has only one lift in Z_f^a, so the lift $x_1 = x_2$ is fixed by ${f_{L_a}}^n$. If, instead, x was inside some triangle T, then the forward flows of any two lifts of x under ψ_a agree by the time they have exited their respective lifts of T. Therefore, ${f_{L_a}}^n(x_1) = {f_{L_a}}^n(x_2) = x_3$ and x_3 is the only lift of x that is fixed by ${f_{L_a}}^n$.

Conversely, for every point y of L_a that is fixed by ${f_{L_a}}^n$, $\pi_a(y)$ is fixed by ${f_L}^n$. Therefore, ${f_{L_a}}^n$ and ${f_L}^n$ have the same number of fixed points, and so the growth rate of fixed points of f_{L_a} is equal to the growth rate of fixed points of f_L. □

Now, we will prove that both L and L_a have finitely many connected components. This fact about L_a will allow us to prove Proposition 7.2, which we will use to prove Theorem C at the end of this chapter.

Proposition 7.1. *The graph, L, has finitely many connected components.*

Proof. Let W_2 be the I-bundle neighborhood of Σ with $h_2 : \Sigma \times (-\epsilon_2, \epsilon_2) \to W_2$ the homeomorphism described in Section 6.2. Similarly, let W_1 be the I-bundle neighborhood of Θ with $h_1 : \Theta \times (-\epsilon_1, \epsilon_1) \to W_1$ the homeomorphism described in Section 6.2.

Let M_n to be the graph $\Theta +_{(2)} n\Sigma$ defined in Section 6.2. By Claim 6.2, M_n is a cross section.

Due to the fact that η is in the boundary of the fibered cone, it is zero on some closed orbit γ of ψ. Therefore, γ does not intersect W_2, and so γ intersects M_n exactly where it intersected Θ. Due to the fact that Θ is a cross section of ψ, it intersects γ in a finite positive number of points. Let k be that number. Thus M_n intersects γ at k points. Due to the fact that M_n is a cross section, every connected component of M_n must intersect γ. Therefore M_n has at most k connected components for any n.

Suppose that L has infinitely many connected components. Define the following open neighborhood of Σ in Z_f, $U_n \coloneqq h_2(\Sigma \times (-\frac{\epsilon_2}{n+\frac{1}{2}}, \frac{\epsilon_2}{n+\frac{1}{2}}))$. Due to the fact that

$\{L \backslash U_n\}_{n \in \mathbb{Z}_{\geq 2}}$ is a nested open cover for L with $L \backslash U_n \subsetneq L \backslash U_{n+1}$, the number of connected components of $L \setminus U_n$ must tend to infinity.

By their constructions, $L \setminus U_n$ is equal to $M_n \setminus U_n$. Therefore, the number of connected components of $M_n \setminus U_n$ must also tend to infinity.

Also, observe that $M_n \cap U_n$ is homeomorphic for all integers $n \geq 2$. Thus, as n goes to infinity, the number of connected components of $M_n \cap U_n$ and the number of points in $M_n \cap \partial U_n$ are both constant finite values.

By Remark 6.1, the valence of M_n is bounded as n goes to infinity.

Putting all these observations together, because the number of points in $M_n \cap \partial U_n$ is constant, and the valence of those points is bounded as n goes to infinity, the fact that the number of connected components of $M_n \setminus U_n$ tends to infinity, implies that the number of connected components of M_n must tend to infinity. This implication contradicts our earlier observation that M_n has at most k connected components. Therefore, L must have finitely many connected components. □

Corollary 7.1. *The graph L_a has finitely many connected components.*

Proof. By Remark 7.8, L_a has the same number of connected components as L, so by Proposition 7.1 we are done. □

Proposition 7.2. *The graph L_a has finitely many ends, and the first return map f_{L_a} is end-periodic.*

Proof. Let W_2 be an I-bundle neighborhood of Σ with homeomorphism $h_2 : \Sigma \times (-\epsilon_2, \epsilon_2) \to W_2$ as described in Section 6.2.

First, we will show that L_a has finitely many ends. Define the open neighborhood of Σ in Z_f, $U_n \coloneqq h_2(\Sigma \times (-\frac{\epsilon_2}{n+\frac{1}{2}}, \frac{\epsilon_2}{n+\frac{1}{2}}))$, and let $\overline{U}_n$ be $\pi_a^{-1}(U_n)$. Letting K_n be the largest subgraph of L_a that does not intersect $\overline{U}_n$, we have that $K_1 \subsetneq K_2 \subsetneq \cdots$ gives an exhaustion of L_a by finite subgraphs. Note that every vertex in K_{n+1} that was not in K_n must be contained in

$$\overline{U}_n \setminus \overline{U}_{n+1}$$

$$= \pi_a^{-1}\big(h_2(\Sigma \times (\frac{\epsilon_2}{n+1+\frac{1}{2}}, \frac{\epsilon_2}{n+\frac{1}{2}}))\big) \bigcup \pi_a^{-1}\big(h_2(\Sigma \times (-\frac{\epsilon_2}{n+\frac{1}{2}}, -\frac{\epsilon_2}{n+1+\frac{1}{2}}))\big).$$

The construction of L_a by blowing air into Z_f with respect to L as described in Section 7.2 is such that L_a has the same number of connected components in $\overline{U}_n \setminus \overline{U}_{n+1}$ as L does in $U_n \setminus U_{n+1}$. The construction of L by spinning Θ around Σ as described in Section 6.2 is such that L has a constant finite number of connected components in $U_n \setminus U_{n+1}$. Therefore, L_a has a constant finite number of connected components in $\overline{U}_n \setminus \overline{U}_{n+1}$. By the construction of the cell complex structure on L_a described in Section 7.2.2, L_a has a constant number of vertices in $\overline{U}_n \setminus \overline{U}_{n+1}$, and because L_a has finitely many connected components in $\overline{U}_n \setminus \overline{U}_{n+1}$, Remark 7.4 implies that that constant number of vertices is finite. Also, by Remark 7.5, the vertices of L_a have bounded valence. These observations combined with by Corollary 7.1 imply that there can be only a bounded number of connected components in $L_a \setminus K_n$ as n goes to infinity. Due to the fact that K_n gives an exhaustion of L_a by finite subgraphs, this implies that L_a has finitely many ends.

Now that we know L_a has finitely many ends, we can choose an integer $j > 1$ so that each connected component of $L_a \setminus K_j$ that is infinite is a neighborhood of a unique

end. By the construction of the cell complex on L_a in Section 7.2.2, we can choose j to be large enough that both f_{L_a} and ${f_{L_a}}^{-1}$ are graph isometric embeddings on $L_a \setminus K_j$.

By the construction of L in Section 6.2, The connected components of $L_a \setminus K_j$ are partitioned into the set of connected components contained in

$\overline{U}_- := \pi_a^{-1}\big(h_2(\Sigma \times (0, \epsilon_2))\big)$ and those contained in $\overline{U}_+ := \pi_a^{-1}\big(h_2(-\epsilon_2, 0))\big)$.

Let W_+ be the union of the connected components of $L_a \setminus K_j$ in $\overline{U}_+$, and let W_- be the union of the connected components of $L_a \setminus K_j$ in $\overline{U}_-$. Then f_{L_a} is such that $\bigcap_{n \in \mathbb{Z}_{>0}} {f_{L_a}}^n(W_+) = \emptyset$ and $\bigcap_{n \in \mathbb{Z}_{>0}} {f_{L_a}}^{-n}(W_-) = \emptyset$. Thus, the iterated images of the components exhaust the ends.

Then because there are finitely many infinite connected components in W_+ and W_-, there must be some interger m such that for each infinite connected component $w_+ \subset W_+$ and $w_- \subset W_-$, ${f_{L_a}}^m(w_+) \subseteq w_+$ and $w_- \subseteq {f_{L_a}}^m(w_-)$. The fact that $W_+ \setminus L_a(W_+)$ and $L_a(W_-) \setminus W_1$ are both finite completes the proof that L_a is end-periodic. □

7.3.1 Other Observations about L_a and its First Return Map

Here, we will discuss some properties of L_a and f_{L_a} that will not contribute to our main theorems.

Proposition 7.3. *The first return map f_{L_a} is weakly expanding.*

Proof. Suppose f_{L_a} is not weakly expanding. Then there must be some interval $\hat{I}$ of L_a and some $n > 0$ such that ${f_{L_a}}^n$ sends $\hat{I}$ homeomorphically to itself. Combined with the fact that f_{L_a} is a topological graph map, this means that if $\hat{I}$ contains k vertices of L_a, ${f_{L_a}}^{nk}$ must send those vertices back to themselves. Therefore, we can shrink $\hat{I}$ to be contained

in the interior of an edge of L_a. This observation implies that either $\hat{I}$ is contained inside of a triangle that had air blown into it to create Z_f^a, or $\hat{I}$ is disjoint from such triangles. Therefore, $\pi_a(\hat{I})$ is a homeomorphism onto its image.

Let I be $\pi_a(\hat{I})$. By Remark 7.7, $({f_L}^n \circ \pi_a)|_{\hat{I}} = (\pi_a \circ {f_{L_a}}^n)|_{\hat{I}}$. The right side is a composition of homeomorphisms, and so ${f_L}^n$ must send I homeomorphically to itself.

Due to the fact that $\hat{I}$ is contained in an edge of L_a, I does not intersect any vertical 1-cells of $Z_f^{(1)}$. If $\tilde{I}_1$ is a lift of I to the universal cover of Z_f, let $\tilde{r}_1$ be the rectangular 2-cell in $\tilde{Z_f}^{(2)}$ that is the lowest (relative to $\tilde{\psi}$) 2-cell containing $\tilde{I}_1$. Then because I does not intersect any vertical 1-cells and L is transverse to ψ, there is an interval $\tilde{I}_1'$ on the bottom horizontal 1-cell of $\tilde{r}_1$ such that $\tilde{I}_1'$ projects homeomorphically onto $\tilde{I}_1$ by flowing along $\tilde{\psi}$ as in Figure 7.4. Call this projection map p_1.

Due to the fact that ${f_L}^n$ sends I homeomorphically to itself, ${\tilde{f}_L}^n$ sends $\tilde{I}_1$ homeomorphically to another lift $\tilde{I}_2$ of I. As with $\tilde{I}_1$, let $\tilde{r}_2$ be the corresponding lowest rectangle containing $\tilde{I}_2$ and let $\tilde{I}_2'$ be the interval on the bottom of $\tilde{r}_2$ that $\tilde{\psi}$ projects homeomorphically onto $\tilde{I}_2$ by projection map p_2 as in Figure 7.4.

Then due to the fact that p_1, p_2, $\tilde{f}_L$ and $\tilde{f}$ are all determined by ψ, there is some integer $m > 0$ so that $(p_2 \circ \tilde{f}^m)|_{\tilde{I}_1'} = ({\tilde{f}_L}^n \circ p_1)|_{\tilde{I}_1'}$. Recalling that ${f_L}^n$ is a homeomorphism on I, ${\tilde{f}_L}^n$ must also be a homeomorphism on $\tilde{I}_1$, and so ${\tilde{f}_L}^n \circ p_1$ is a composition of homeomorphisms. Therefore, $\tilde{f}^m$ must also be a homeomorphism on $\tilde{I}_1'$ that sends $\tilde{I}_1'$ to $\tilde{I}_2'$. Recalling that $\tilde{I}_1'$ and $\tilde{I}_2'$ are both lifts of I, this implies that f^m sends I homeomorphically

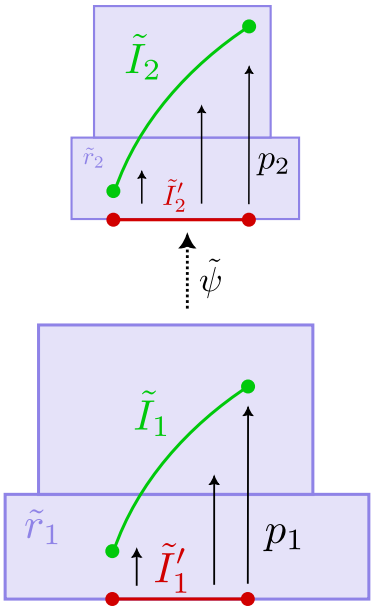

Figure 7.4: Projection maps p_1 and p_2 send the intervals $\tilde{I}'_1$ and $\tilde{I}'_2$ to $\tilde{I}_1$ and $\tilde{I}_2$ respectively by homeomorphisms.

to itself. This contradicts the fact that f is weakly expanding, so f_{L_a} must also be weakly expanding. □

Proposition 7.4. *The first return map f_{L_a} is a train track map.*

Proof. Suppose f_{L_a} is not a train track map. Then there must be some point $\hat{x}$ in the interior of some edge $\hat{e}$ of L_a, and some $n > 0$ such that ${f_{L_a}}^n$ is not an immersion at $\hat{x}$. Consider a sub-interval $\hat{I}$ of $\hat{e}$ such that $\hat{I}$ contains $\hat{x}$. Let I be $\pi_a(\hat{I})$ and x be $\pi_a(x)$. By the construction of L_a, and because x is in the interior of an edge of L_a, x cannot be on a 0-cell of $Z_f^{(0)}$ or a vertical 1-cell of $Z_f^{(1)}$. If x is in the interior of a rectangular 2-cell of $Z_f^{(2)}$, then we can shrink $\hat{I}$ so that I is also contained in the interior of that 2-cell. Otherwise, if x is on a horizontal 1-cell of $Z_f^{(1)}$, we can shrink $\hat{I}$ so that $I \setminus x$ is contained in the interiors of two rectangular 2-cells of $Z_f^{(2)}$ (the one above x, one of the ones below x).

The fact that $\hat{I}$ is in the interior of an edge of L_a implies that either $\hat{I}$ is contained inside of a triangle which had air blown into it to create Z_f^a, or $\hat{I}$ is disjoint from such triangles. Therefore, $\pi_a(\hat{I})$ is an immersion on $\hat{I}$.

By Remark 7.7, $({f_L}^n \circ \pi_a)|_{\hat{I}} = (\pi_a \circ {f_{L_a}}^n)|_{\hat{I}}$. The right side is not an immersion at $\hat{x}$, while π_a is an immersion on $\hat{I}$, so ${f_L}^n$ must not be an immersion at x.

Due to the fact that $\hat{I}$ is contained in an edge of L_a, I does not intersect any vertical 1-cells of $Z_f^{(1)}$. If $\tilde{I}_1$ is a lift of I to the universal cover of Z_f, let $\tilde{r}_1$ be the rectangular 2-cell in $\tilde{Z_f}^{(2)}$ that is the lowest (relative to $\tilde{\psi}$) 2-cell containing $\tilde{I}_1$. Then because I does not intersect any vertical 1-cells and L is transverse to ψ, there is an interval $\tilde{I}_1'$ on the bottom horizontal 1-cell of $\tilde{r}_1$ such that $\tilde{\psi}$ projects $\tilde{I}_1'$ by an immersion onto $\tilde{I}_1$ as in Figure 7.5. Call this projection map p_1.

The map ${\tilde{f}_L}^n$ sends $\tilde{I}_1$ to another interval $\tilde{I}_2$ which is a lift of ${f_L}^n(I)$. We can project $\tilde{I}_2$ back into $\tilde{\Gamma}$ by flowing it forward along $\tilde{\psi}$ as in Figure 7.5. This projection map is an immersion, which we will call p_2.

Due to the fact that p_1, p_2, $\tilde{f}$, and $\tilde{f}_L$ are all determined by $\tilde{\psi}$, there is some integer $m > 0$ such that $\tilde{f}^m = p_2 \circ {\tilde{f}_L}^n \circ p_1$. Recalling that ${f_L}^n$ is not an immersion on I, ${\tilde{f}_L}^n$ must not be an immersion on $\tilde{I}_1$. Therefore, $p_2 \circ {\tilde{f}_L}^n \circ p_1$ is not an immersion on $\tilde{I}_1'$, so $\tilde{f}^m$ is also not an immersion on $\tilde{I}_1'$, and hence f^m is not an immersion on I_1' the projection of $\tilde{I}_1'$ to Z_f. Recalling that I_1' is contained in an edge of Γ, this contradicts the fact that f is a train track map, and therefore f_{L_a} must also be a train track map. $\square$

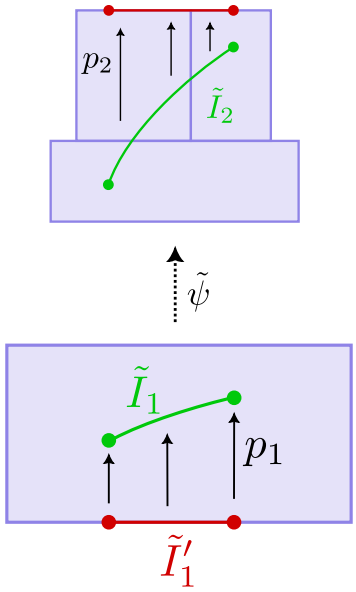

Figure 7.5: The maps p_1 projects $\tilde{I}'_1$ onto $\tilde{I}_1$ by an immersion, while the map p_2 projects $\tilde{I}_2$ into $\tilde{\Gamma}$ by an immersion.

Remark 7.8. By Remark 7.3, the restriction of π_a to L_a is a homotopy equivalence between L_a and L.

Claim 7.1. The space $Z_f \setminus \Sigma$ is homotopy equivalent to the mapping tori Z_{f_L} (which is defined in the same way as Definition 1.7 without the requirement that L be a graph or f_L a topological graph map) and $Z_{f_{L_a}}$. Letting m be the number of connected components of L and of L_a, and letting f_{L_0} be ${f_L}^m$ restricted to a connected component L_0 of L, and $f_{L_{a,0}}$ be ${f_L}^m$ restricted to a connected component $L_{a,0}$ of L_a, $Z_f \setminus \Sigma$ is also homotopy equivalent to $Z_{f_{L_0}}$ and to $Z_{f_{L_{a,0}}}$.

To justify the above claim, observe that for any $\epsilon > 0$, and any $x \in \Sigma$, the segment $\{\psi^t(x) : t \in (0, \epsilon)\}$ intersects L. Also, for any $x' \in Z_f$ such that $\psi^\epsilon(x') = x$ the segment $\{\psi^t(x') : t \in (0, \epsilon)\}$ intersects L. Therefore, we can reparametrize ψ to get a semiflow on $Z_f \setminus \Sigma$, and the reparametrized semiflow determines the same first return map on L. Also, L

is still transverse to this reparametrized semiflow. The fact that f is expanding on all scales means that ψ flows L over all of Z_f. Combined with the previous observations, this implies that the reparametrization of ψ on $Z_f \setminus \Sigma$ flows L over all of $Z_f \setminus \Sigma$. Then because L is transverse to the reparametrized semiflow and has a first return map, $Z_f \setminus \Sigma$ is homotopy equivalent to the mapping torus Z_{f_L}. By Remark 7.7 and Remark 7.8, it is also homotopy equivalent to $Z_{f_{L_a}}$. Finally, because f is expanding on all scales, Z_f must be connected, so $Z_f \setminus \Sigma$ must also be connected. Therefore, Z_{f_L} and $Z_{f_{L_a}}$ must also be connected. This observation implies that f_L and f_{L_a} must act transitively on their connected components, which gives us the last sentence of Claim 7.1.

7.3.2 The Stretch Factor of f_{L_a}

Recall that $\lambda(f_{L_a})$ is the stretch factor of f_{L_a} defined by Definition 4.5, and recall that $\mathrm{gr}(f_{L_a})$ is the growth rate of fixed points of f_{L_a} defined by Definition 1.4. We will define another concept of growth for f_{L_a}, then show that all three of them are equal, and use their equivalence to prove Theorem C.

We define the new concept of growth,

$$\mathrm{gr}_\psi(\Theta;\Sigma) \coloneqq \mathrm{gr}_\psi(\mu;\eta)$$

$$\coloneqq \limsup_{N\to\infty} \left(\#\{\gamma \in O_\psi : \gamma \text{ does not intersect } \Sigma \text{ but } \gamma \text{ intersects } \Theta \text{ at } N \text{ times }\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma \in O_\psi : \eta(\gamma) = 0, \mu(\gamma) = N\}\right)^{1/N}.$$

The above expressions are equal because Σ is dual to η and Θ is dual to μ.

Lemma 7.3. *The growth rate of fixed points $gr(f_{L_a})$ of f_{L_a} is equal to $gr_\psi(\Theta;\Sigma)$.*

Proof. By the construction of L, a closed orbit $\gamma \in O_\psi$ intersects L infinitely many times if and only if γ intersects Σ, and if γ does not intersect Σ, then γ intersects L_a at exactly the times it intersects Θ. Therefore, we have

$$\mathrm{gr}_\psi(\Theta;\Sigma)$$

$$:= \limsup_{N\to\infty} \left(\#\{\gamma \in O_\psi : \gamma \text{ does not intersect } \Sigma \text{ but } \gamma \text{ intersects } \Theta \text{ at } N \text{ times }\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma \in O_\psi : \gamma \text{ intersects } L \text{ at } N \text{ times }\}\right)^{1/N}.$$

Due to the fact that Σ is dual to η in the boundary of the fibered cone, there is some closed orbit $\gamma \in O_\psi$ such that $\eta(\gamma) = 0$. So, γ does not intersect Σ. However, γ does intersect Θ because Θ is a cross section. Then γ must intersect L a positive finite number of times. Supposing m is that positive finite number, ${f_L}^m$ fixes those points of intersection of γ with L. Therefore, we can apply Corollary 4.1 to get

$$= \limsup_{N\to\infty} \left(\#\{\gamma \in O_\psi : \gamma \text{ intersects } L \text{ at } N \text{ times }\}\right)^{1/N} = \mathrm{gr}(f_L).$$

Finally, by Lemma 7.2, the above is equal to $\mathrm{gr}(f_{L_a})$, which completes the proof.

□

Corollary 7.2. *The stretch factor $\lambda(f_{L_a})$ of f_{L_a}, is equal to $gr_\psi(\Theta;\Sigma)$.*

Proof. By Lemma 7.3 $\mathrm{gr}_\psi(\Theta;\Sigma)$ is equal to the growth rate of fixed points of f_{L_a}.

Note that because η is in the boundary of the fibered cone, there is some closed orbit of ψ that misses Σ; however, Θ is a cross section of ψ, and therefore intersects this closed orbit a finite positive number of times. The construction of L described in Section 6.2 is such that L intersects this closed orbit the same finite positive number of times as Θ does. Those intersection points between the closed orbit and L are fixed points of ${f_L}^n$ where n is the number of intersection points between the closed orbit of ψ and L. Therefore, by Lemma 7.2, ${f_{L_a}}^n$ also has some fixed point.

By the construction of L_a described in Section 6.2 and Section 7.2, L_a must be an infinite graph. Also, observe that by the construction of the cell complex structure on L_a described in Section 7.2.2, L_a is a graph of finite valence with no valence-0 vertices.

Finally, note that by Proposition 7.2, f_{L_a} is end-periodic.

These observations mean that we can apply Theorem B to f_{L_a} to get that the growth rate of fixed points of f_{L_a} is equal to the stretch factor of f_{L_a}. □

Now, we use the previous lemma and corollary to prove Theorem C.

Theorem C. *Given the irreducible train track topological graph map $f : \Gamma \to \Gamma$ on the finite graph Γ, which has no valence-1 vertices, and the mapping torus Z_f of f with the natural semiflow ψ, let Θ be a cross section of ψ that is dual to μ in the interior of the fibered cone of Z_f. Also, let Σ be a finite smoothly embedded transverse graph in Z_f that is dual to η in the boundary of the fibered cone, and let L_a be a graph made by blowing*

air into Z_f with respect to a graph resulting from spinning Θ around Σ. Suppose Ξ_i is the cross section dual to $\xi_i = \mu + i\eta$ that has a first return map f_{Ξ_i} that is a train track topological graph map. Then the stretch factors of the first return maps f_{Ξ_i} of Ξ_i limit to the stretch factor of the first return map f_{L_a} of L_a as i goes to infinity, or

$$\lim_{i\to\infty} \lambda(f_{\Xi_i}) = \lambda(f_{L_a}). \tag{C}$$

Proof. Due to the fact that Ξ_i is a cross section of ψ, some power of f_{Ξ_i} has a fixed point, and therefore we can apply Corollary 4.1 and Proposition 4.1 to replace $\lambda(f_{\Xi_i})$ with $\text{gr}_\psi(\Xi)$. By Corollary 7.2, we can replace $\lambda(f_{L_a})$ with $\text{gr}_\psi(\Theta;\Sigma)$. So we will prove the theorem by proving that

$$\lim_{i\to\infty} \text{gr}_\psi(\Xi_i) = \text{gr}_\psi(\Theta;\Sigma).$$

To prove the above equality, consider the flow graph, D_f smoothly embedded in Z_f, and let $\alpha_i \in H^1(D_f)$ be the pullback of $\xi_i = \mu + i\eta$. Let α be the pullback of μ to $H^1(D_f)$ and β be the pullback of η to $H^1(D_f)$. Define D'_f to be the subgraph of D_f that is covered by directed cycles γ such that $\beta(\gamma) = 0$, implying that the α_i all pull back to α on D'_f.

For any directed cycle $\gamma \in D_f$ not contained in D'_f, Remark 3.2 and Proposition 3.1 tell us that γ is homotopic to a closed orbit of ψ in Z_f. We will call this closed orbit γ'. Due to the fact that η is in the closure of the fibered cone, $\eta(\gamma') \geq 0$, implying that $\beta(\gamma) \geq 0$. Due to the fact that γ is not contained in D'_f, $\beta(\gamma) > 0$. Therefore, $\alpha_i(\gamma) = \eta(\gamma) + i\beta(\gamma)$

goes to infinity as i goes to infinity. This means α_i satisfies condition (1) of [8, Lemma 9.12].

As noted in [8], condition (2) of [8, Lemma 9.12] can be replaced by the condition that D'_f contains a directed cycle. Due to the fact that η is in the boundary of the fibered cone, there is some closed orbit γ_0 of ψ such that $\eta(\gamma_0) = 0$. If γ_0 is not a vertex orbit, then by Proposition 3.1, γ_0 is in-step with and therefore also homotopic to a unique directed cycle ℓ_0 of D_f. If γ_0 is a vertex orbit, then by Remark 3.2, there is some integer $m \geq 1$ so that γ_0^m is homotopic to some directed cycle ℓ_0 of D_f. Now, because $\eta(\gamma_0) = 0$, and either ℓ_0 is homotopic to γ_0 or to γ_0^m in Z_f, we have that $\beta(\ell_0) = 0$. Therefore, D'_f contains ℓ_0. Thus, D'_f does contain a directed cycle, and we can apply [8, Lemma 9.12] which gives us that,

$$\lim_{i\to\infty} \text{gr}(\alpha_i; D_f) = \text{gr}(\alpha; D'_f).$$

We will show that $\text{gr}_\psi(\Xi_i) = \text{gr}(\alpha_i; D_f)$, which combined with the above equality implies that $\lim_{i\to\infty} \text{gr}_\psi(\Xi_i) = \text{gr}(\alpha; D'_f)$. By Corollary 4.2,

$$\text{gr}_\psi(\Xi_i) \coloneqq \limsup_{N\to\infty} \left(\#\{\gamma \in O_\psi : (\mu + i\eta)(\gamma) = N\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma \text{ a directed cycle in } D_f : \alpha_i(\gamma) = N\}\right)^{1/N} =: \text{gr}(\alpha_i; D_f),$$

Finally, we will complete the proof by showing that $\text{gr}_\psi(\Theta; \Sigma) = \text{gr}(\alpha; D'_f)$. Due to the fact that η is zero on some closed orbit of ψ, Lemma 4.2 applies and gives us

$$\mathrm{gr}_{\psi}(\Theta;\Sigma) \coloneqq \limsup_{N\to\infty} \left(\#\{\gamma \in O_{\psi} : \eta(\gamma) = 0, \mu(\gamma) = N\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma \text{ a directed cycle in } D_f : \beta(\gamma) = 0, \alpha(\gamma) = N\}\right)^{1/N}$$

$$= \limsup_{N\to\infty} \left(\#\{\gamma \text{ a directed cycle in } D'_f : \alpha(\gamma) = N\}\right)^{1/N} =: \mathrm{gr}(\alpha; D'_f). \qquad \square$$

In Chapter 8, we will discuss when there exist finite transverse smoothly embedded graphs dual to the boundary of the fibered cone, such as the graph Σ, which we used in this chapter to create L_a. Due to [2, Theorem B], the sequence of cross sections Ξ_i as in Theorem C always exists, so whenever we have a graph such as Σ, the equality in Theorem C is realized.

CHAPTER 8

EXISTENCE OF A BOUNDARY GRAPH

In this section, we will show that under the right conditions, there are finite transverse graphs in Z_f that are dual to integral classes in the closure of the fibered cone. We will prove the following theorem at the end of this chapter in Section 8.5.

Theorem A. *Given the irreducible train track topological graph map $f : \Gamma \to \Gamma$ on the finite graph Γ, which has no valence-1 vertices, suppose η is an integral cohomology class in the closure of the fibered cone of the mapping torus Z_f of $f : \Gamma \to \Gamma$. If η is positive on all closed orbits of vertices of Γ under the natural semiflow on Z_f, then η is dual to a finite transverse graph Σ that is smoothly embedded in Z_f.*

Graphs such as Σ in the statement of Theorem A (finite graphs that are smoothly embedded in Z_f, transverse to ψ, and dual to classes in the boundary of the fibered cone) were used in Chapter 7 along with tools from Chapter 6 and Section 7.1 to prove Theorem C.

We will prove Theorem A in this chapter by subdividing the 1-skeleton of Z_f in Section 8.1, then discussing the existence of non-negative integral cocycle representatives on this subdivided 1-skeleton in Section 8.2. Then, in Section 8.3, we create smoothly embedded graphs with oriented I-bundle neighborhoods that are dual to integral classes in the closure of the fibered cone. We discuss the structure of such graphs in Section 8.4 and Section 8.4.2. Finally, in Section 8.5, we explain how to alter the smooth embeddings

of these graphs to be transverse to ψ, and we put all these concepts together to prove theorem A.

8.1 Setup

First, we subdivide the cell structure of Z_f, as pictured in Figure 8.1. Subdivide along the image of D_f by adding a new 0-cell to the interior of each horizontal 1-cell from Γ, then add new 1-cells (in bijection with directed edges of D_f) within each rectangular 2-cell that connect the new 0-cell on the bottom of edge the 2-cell to the new 0-cells on the top arc of the 2-cell by straight lines. The graph made of only these new 1-cells and 0-cells is the image of D_f. Then, we subdivide further by adding new 1-cells within each rectangular 2-cell that connect the new 0-cell on the bottom edge of the 2-cell to the original 0-cells on the top arc of the 2-cell by straight lines. Call the 0-skeleton of this new subdivided cell structure $Z_{f_{sub}}^{(0)}$, the 1-skeleton $Z_{f_{sub}}^{(1)}$, and the 2-skeleton $Z_{f_{sub}}^{(2)}$. We will call the 2-cells of this cell structure **triangular 2-cells**. As the vertical 1-cells of $Z_{f_{sub}}^{(1)}$ are flowlines of ψ, we will orient them in the direction of ψ. Each of the other new 1-cells is contained in the interior of some rectangular 2-cell, with a top and bottom arc determined by ψ. We will orient each of the new 1-cells towards the top arc of the rectangular 2-cell containing them, and as a result, all of the 1-cells except for the horizontal ones will have orientations.

8.2 Existence of a Non-Negative Integral Cocycle Representative

The following lemma gives us the existence of non-negative cocycle representatives on $Z_{f_{sub}}^{(1)}$ of integral classes in the fibered cone.

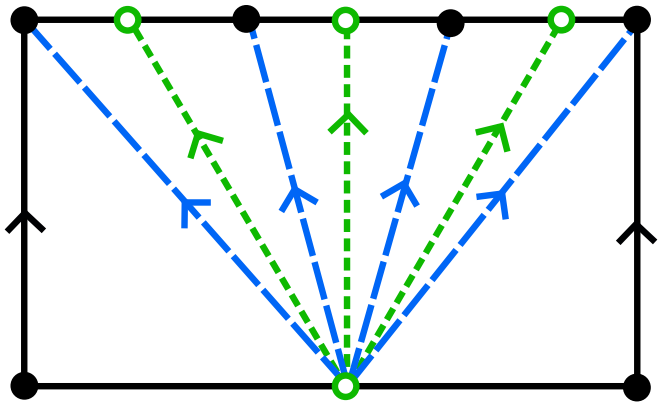

Figure 8.1: The subdivision of a rectangular 2-cell of Z_f. The green open circles are the new 0-cells, and the green dotted lines are some of the new 1-cells. The image of D_f in Z_f icomprisesthose green 0-cells and 1-cells.

Lemma 8.1. *For an integral cohomology class $\eta \in H^1(Z_f, \mathbb{R})$ which is non-negative on closed orbits of ψ, we can choose orientations on the horizontal 1-cells of ${Z_f}_{sub}^{(1)}$ so that η has a non-negative integral cocycle representative $\bar{\eta}$ on that oriented 1-skeleton $\overset{o}{{Z_f}}_{sub}^{(1)}$.*

Proof. The 1-cells of ${Z_f}_{sub}^{(1)}$ that are not horizontal all have natural orientation in the positive semiflow direction.

Consider the subset of all the new 1-cells in the subdivided cell complex, all the 0-cells in the subdivided cell complex, and all the vertical 1-cells, and call the union of these cells K. Then K is ${Z_f}_{sub}^{(1)}$ minus the interiors of horizontal 1-cells. Pull back η to K. We can easily construct some integral cocycle representative, $\bar{\eta}'_K$, of the pullback of η to K. We do this by picking a maximal tree in K and defining $\bar{\eta}'_K$ to be zero on the edges of that tree. This definition of $\bar{\eta}$ on the tree determines the value that $\bar{\eta}'_K$ must assign to the other edges of K in order for it to represent the pullback of η. Due to the fact that each edge of K which is not in the maximal tree is part of some loop in K such that all other edges of the loop are in the maximal tree, and because η is an integral cohomology class,

$\bar{\eta}'_K$ assigns integer values to these edges outside the maximal tree. So $\bar{\eta}'_K$ is an integral cocycle representative.

Note that all the directed cycles of K are either cycles in D_f or closed orbits of vertices under ψ. By Proposition 3.1 and Remark 3.2, every directed cycle of D_f is homotopic to some closed orbit of ψ. Due to the fact that η is non-negative on closed orbits of ψ, this implies that the pullback of η to K is non-negative on all directed cycles of K

Therefore, by [6, Proposition 1.13 and Remark 1.12] there is a non-negative integral cocycle, $\bar{\eta}_K$, on K that represents the pullback of η to K.

We can then extend $\bar{\eta}_K$ to a cochain, $\bar{\eta}$, on all of $Z_{f_{sub}}^{(1)}$ by orienting e and defining $\bar{\eta}(e)$ for each horizontal 1-cell e of $Z_{f_{sub}}^{(1)}$. Each horizontal 1-cell of $Z_{f_{sub}}^{(1)}$ is adjacent to one triangular 2-cell of $Z_{f_{sub}}^{(2)}$, that is directly above it relative to ψ, and a collection of triangular 2-cells below it. Consider the triangular 2-cell T, that is above e as depicted in Figure 8.2. Let the other two 1-cells on the boundary of T be e_a and e_b with $\bar{\eta}_K(e_a) = a$ and $\bar{\eta}_K(e_b) = b$. Pick an orientation of e, and define $\bar{\eta}(e)$ so that $\bar{\eta}(\partial S) = 0$. If $\bar{\eta}(e)$ is negative, switch the orientation of e so that $\bar{\eta}(e)$ is positive. Now we have that $\bar{\eta}(e) = |a - b|$.

For any triangular 2-cell T' below e, let the other 1-cells of T' be e_c and e_d. Then because $\partial T + \partial T'$ is a boundary in Z_f and $\partial T + \partial T' = e_a - e_b - e_d + e_c$ is a chain in K, we have that $\bar{\eta}_K(\partial T + \partial T') = \eta(\partial T + \partial T') = 0$, implying that $\bar{\eta}(\partial T + \partial T') = 0$. Combined with the fact that $\bar{\eta}(\partial T) = 0$, this implies that $\bar{\eta}(\partial T') = 0$. Thus, we have that $\bar{\eta}$ is a cocycle on $Z_{f_{sub}}^{(1)}$, which we have constructed to be non-negative and integer-valued

on 1-cells. Also, $\bar{\eta}$ must represent η on Z_f because any loop in Z_f can be homotoped into K. $\square$

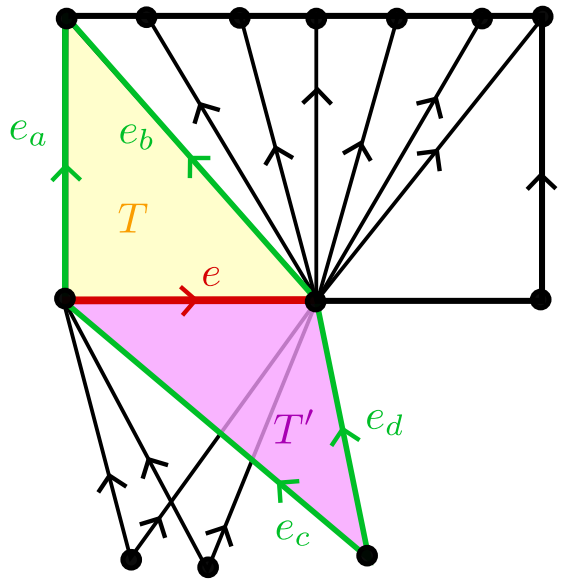

Figure 8.2: The 1-cell e is not part of K, but $\bar{\eta}_K$ can be extended to $\bar{\eta}$, which is defined on e.

8.3 Creating the Graph Σ

In this section, we will constructively prove that integral classes in the closure of the fibered cone are dual to finite smoothly embedded graphs in Z_f with specific properties outlined in Lemma 8.2. We will then define some terms (Definition 8.1 and Definition 8.2) for graphs with subsets of these properties and discuss such graphs more generally.

Lemma 8.2. *Given an integral cohomology class η in the closure of the fibered cone, if $\bar{\eta}$ is a non-negative integral cocycle representative on the oriented 1-skeleton $\overset{o}{Z}{}^{(1)}_{f_{sub}}$ coming from a choice of orientations on the horizontal 1-cells of $Z^{(1)}_{f_{sub}}$, there is a graph Σ mapped into Z_f satisfying the following properties:*

1. *For each* 1*-cell* e *of* $\overset{o}{Z}{}^{(1)}_{f\,sub}$, Σ *has* $\bar{\eta}(e)$ *vertices mapped to distinct points along the interior of* e,

2. *The interior of any edge of* Σ *is smoothly embedded in the interior of some triangular* 2*-cell of* $Z^{(2)}_{f\,sub}$,

3. *Interiors of distinct edges of* Σ *do not intersect,*

4. *The subset of* Σ *mapped into the closure of any triangular* 2*-cell is a disjoint collection of edges of* Σ,

5. *Every edge of* Σ *can be co-oriented in the* 2*-cell it is mapped into such that the edge's co-orientation agrees with the orientations of the* 1*-cells that its endpoints are mapped into. See Figure 8.3 for an example of this property.*

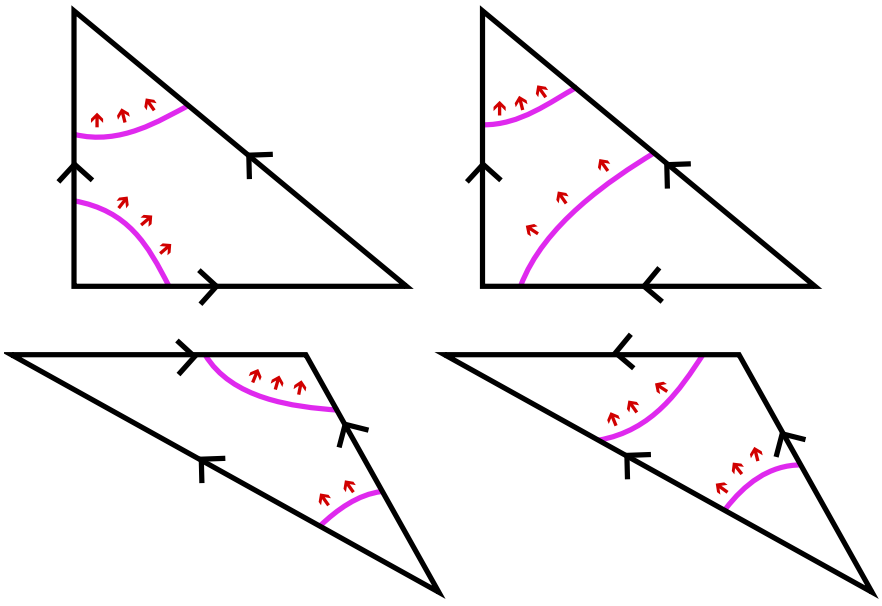

Figure 8.3: These triangles are examples of triangular 2-cells with purple edges depicted inside them. The small red arrows show the co-orientations on these edges.

Proof. To prove this, we will construct the continuous image of a graph of Σ in Z_f so that it has these properties. By abuse of notation, we will also call the image Σ. By Lemma 8.1, there is a non-negative integral representative $\bar{\eta}$, of η. Then for each 1-cell e, of $Z_{f_{sub}}^{(1)}$, let Σ have $\bar{\eta}(e)$ vertices mapped in the interior of e. Thus, Σ will satisfy property 1.

Given a triangular 2-cell D of $Z_{f_{sub}}^{(2)}$, let e be the horizontal 1-cell on its boundary and e_a and e_b the other two 1-cells of its boundary. In Lemma 8.1 we constructed the cocycle representative $\bar{\eta}$, such that $\bar{\eta}(e) = |\bar{\eta}(e_a) - \bar{\eta}(e_b)|$. Without loss of generality, say $\bar{\eta}(e_a) \geq \bar{\eta}(e_b)$. Then $\bar{\eta}(e) = \bar{\eta}(e_a) - \bar{\eta}(e_b)$ and e is oriented such that it either has the same initial 0-cell or the same terminal 0-cell as e_a depending on whether e is at the bottom of D or the top, respectively as depicted in Figure 8.4. Then, we pair each vertex on e with a vertex on e_a. In particular, if v is the ith vertex on e from the shared 0-cell of e and e_a, then v will be paired with the ith vertex on e_a from the shared 0-cell. Then, for each such pair, we smoothly embed an edge of Σ within the 2-cell so that its interior is in the interior of D, and its boundary is the pair of vertices, so these edges satisfy property 2. Then, because e and e_a share the same initial or terminal 0-cell, we co-orient each of these edges in the direction of the orientation of e_a, and it will agree with the direction of the orientation of e, therefore these edges satisfy property 5.

Each vertex along e is adjacent to exactly one edge of Σ in D. There may still be more vertices along e_a that were not paired with vertices of e, and there may be vertices on e_b. We will pair these vertices on e_a with the ones on e_b, and this will exhaust any remaining vertices because $\bar{\eta}(e_b) = \bar{\eta}(e_a) - \bar{\eta}(e)$. In particular, if v is the ith vertex on

e_a from the shared 0-cell of e_a and e_b, then v will be paired with the ith vertex on e_b from the shared 0-cell. Then, for each such pair, we smoothly embed an edge of Σ as above so that it satisfies property 2. Due to the fact that e_a and e_b were oriented according to ψ, they share the same initial or terminal 0-cell, and as above, we co-orient each of these edges in the direction of the orientation of e_a, and it will agree with the direction of the orientation of e_b. Therefore, these edges satisfy property 5.

Now, for each vertex on the boundary of D, we have smoothly embedded exactly one edge of Σ into D adjacent to that vertex, so property 4 is satisfied. Finally, because each pair of edges on the boundary of D was chosen so that the edge connecting them did not separate any other pair of vertices, we can homotope Σ within the interior of D to satisfy property 3, while still preserving the other properties. □

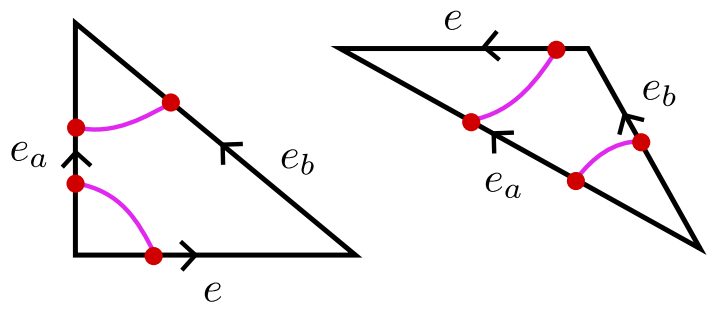

Figure 8.4: Two images of triangular 2-cells are shown with horizontal 1-cells labelled e, and the other 1-cells labelled e_a and e_b. The vertices of Σ are shown as red circles, and the edges of Σ are purple lines that connect the vertices as described above.

Remark 8.1. Due to the fact that $\bar{\eta}$ is finite on each 1-cell of $\overset{o}{Z}{}^{(1)}_{f_{sub}}$, Σ is finite.

Definition 8.1. We will call any finite graph G mapped into Z_f with its vertices mapped disjointly into $Z^{(1)}_{f_{sub}}$ that (replacing Σ with G in the statements) satisfies properties (2) through (5) of Lemma 8.2 a **co-oriented track** in Z_f.

Although a co-oriented track isn't automatically transverse to ψ, it does have an oriented I-bundle neighborhood in Z_f (such a graph in a 2-complex is sometimes called a track [1]).

Claim 8.1. A co-oriented track is smoothly embedded into Z_f.

In particular, a co-oriented track is smoothly embedded in Z_f because its vertices are mapped disjointly into $Z_{f_{sub}}^{(1)}$ combined with properties (2) and (3) of Lemma 8.2.

Definition 8.2. Let G be a co-oriented track in Z_f that is dual to some cohomology class α. If there is some oriented 1-skeleton $\overset{o}{Z}{}^{(1)}_{f_{sub}}$ coming from a choice of orientations on the horizontal 1-cells of $Z_{f_{sub}}^{(1)}$, and some non-negative cocycle representative $\overline{\alpha}$ of α on $\overset{o}{Z}{}^{(1)}_{f_{sub}}$ such that G satisfies the properties of Lemma 8.2 (replacing Σ with G and $\bar{\eta}$ with $\overline{\alpha}$ in the statements), then we will say G is a **co-oriented track relative to** $\overline{\alpha}$.

Claim 8.2. Given a non-negative integral cocycle representative $\bar{\eta}$ on $Z_{f_{sub}}^{(1)}$ for some cohomology class, a co-oriented track relative to $\bar{\eta}$ is dual to η.

By Claim 2.1, the co-oriented track is dual to some unique cohomology class in $H^1(Z_f, \mathbb{R})$, the fact that $\bar{\eta}$ represents η and property 1 of Lemma 8.2 implies that that cohomology class must be η. Therefore, the above claim holds.

For any non-negative cohomology class η in the closure of the fibered cone, Lemma 8.2 gives us the existence of a finite graph Σ that is smoothly embedded in Z_f and dual to η. What's left is to do to prove Theorem A is to smoothly embed Σ more carefully into Z_f so that it is transverse to ψ. We will do this by re-defining the embedding in a way that only moves edges within each triangular 2-cell, so the new embedding is homotopic

to the original and still corresponds to the same cocycle representative. However, before we do that, we will make some observations about the structure of co-oriented tracks in general and how they are mapped into Z_f.

8.4 The Structure of a Co-Oriented Track

In this section, we will uncover the structure of co-oriented tracks, with all discussion applying to general co-oriented tracks until Lemma 8.6 at the end of the section, where we will consider co-oriented tracks that are dual to cohomology classes in the closure of the fibered cone.

8.4.1 Categorizing Edges of Co-Oriented Tracks

We start by categorizing the edges of a co-oriented track. For a co-oriented track G, we will consider the minimal edge paths with both endpoints on $Z_f^{(1)}$. Each of these edge paths is contained in a rectangular 2-cell of $Z_f^{(2)}$ and has its endpoints on the boundary of that rectangular 2-cell. In considering these edge paths, we are returning to the original un-subdivided 1-skeleton of Z_f and viewing these edge paths as the natural edges of G in that cell structure. We will define and label some categories of these edge paths. Such an edge path will be called **vertical** if its endpoints are on the top and bottom arcs of the rectangular 2-cell, **horizontal** if its endpoints are on the left and right vertical sides of the rectangular 2-cell (after fixing a homomorphism between the rectangle and $(0,1) \times (0,1)$), **diagonal** if one of its endpoints is on the top or bottom arc of the rectangular 2-cell and the other endpoint is on the left or right arc of the rectangular 2-cell, or a **bump** if both of its

endpoints are either on the top or bottom arc of the rectangular 2-cell. We give examples of each type of edge path in Figure 8.5.

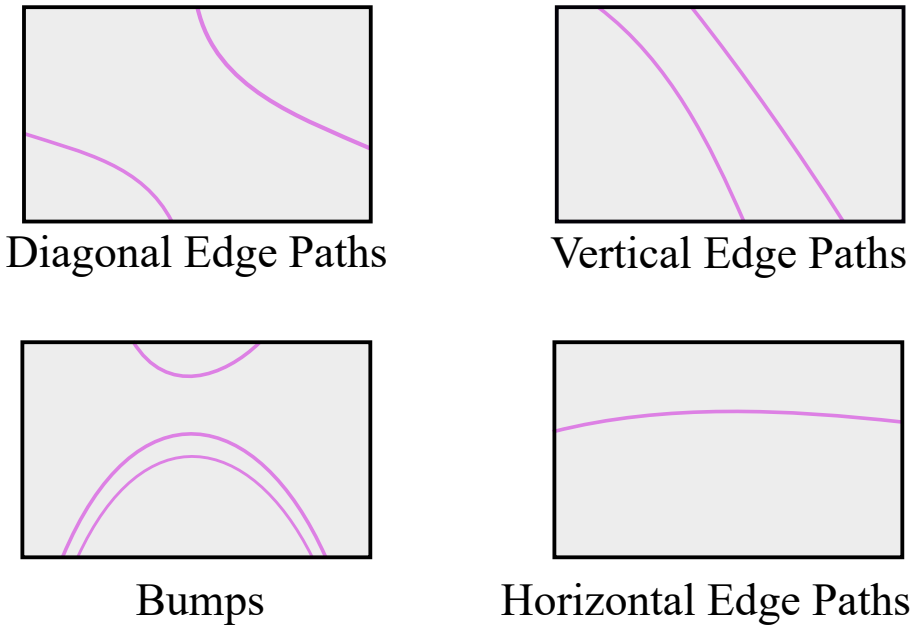

Figure 8.5: Categories of edge paths of a co-oriented track.

By the properties (2) through (5) of Lemma 8.2, it isn't possible for any of these edge paths to have both endpoints on the left or right arc of a rectangular 2-cell. Indeed, if such an edge path h, were to have both endpoints on the right side of the rectangle (after fixing a homomorphism between the rectangle and $(0,1) \times (0,1)$), then that would mean there was some triangular 2-cell in that rectangle where some arc of h contained in that triangular 2-cell had both endpoints on the right 1-cell of that triangular 2-cell. However, this would contradict property 5 of Lemma 8.2.

To further divide the vertical edge paths into subcategories, we will orient them towards their endpoint, which is at the top of the rectangle, and we will make the following observation.

Claim 8.3. If a collection of vertical edge paths of a co-oriented track meet at some mutual endpoint point on a horizontal 1-cell e, then at most one of these vertical edge paths is oriented away from x.

Claim 8.3 is true due to the fact that $f(e)$ is well-defined, meaning there is at most one rectangular 2-cell with e as its bottom arc. This claim then tells us that any two distinct oriented loops composed of vertical edge paths are disjoint from one another.

Definition 8.3. We will call the oriented loops of co-oriented tracks composed of vertical edge paths **vertical loops**.

8.4.2 Structure of Vertical Edge Paths in Co-Oriented Tracks

Now, we can further categorize the vertical edge paths of co-oriented tracks into those that are part of some vertical loop and those that are not.

The following lemma relating the co-orientation of a vertical edge path with its intersections with $Z_{f_{sub}}^{(1)}$ will be used throughout the proof of Theorem A.

Lemma 8.3. *For a co-oriented track in Z_f, suppose ℓ is a vertical edge path in some rectangle r and that v is the vertex where ℓ meets the bottom arc of r. Let e_1 and e_2 be the two horizontal 1-cells in $Z_{f_{sub}}^{(1)}$ which make up the bottom arc of r. Then, if ℓ is co-oriented towards the side of the bottom arc that e_2 is on, v must be on e_2 as in Figure 8.6.*

Proof. Fix a homeomorphism between the interior of r and $(0,1) \times (0,1)$ as in Section 2.1 so that we can refer to left and right. Without loss of generality, suppose that the co-orientation of ℓ is to the right in r, and suppose that e_2 is on the right side of the bottom

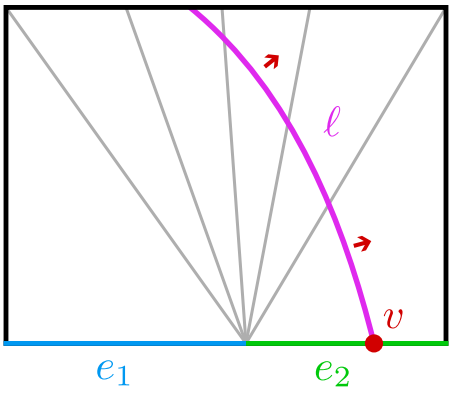

Figure 8.6: ℓ is co-oriented towards e_2, and so it has a vertex v on e_2.

arc of r as in Figure 8.7. Then the lemma states that v should be on e_2. Suppose instead that v is on e_1. Then e_1 must be oriented to the right due to the co-orientation of ℓ. Due to the fact that e_1 is on the left side of the bottom arc of r, there is a diagonal 1-cell of $Z^{(1)}_{f_{sub}}$ in r that ℓ must pass through. However, the orientation on that diagonal 1-cell induces the opposite co-orientation on ℓ, as pictured in Figure 8.7, which contradicts the well-defined co-orientation on Σ. Therefore, v must be on e_2. □

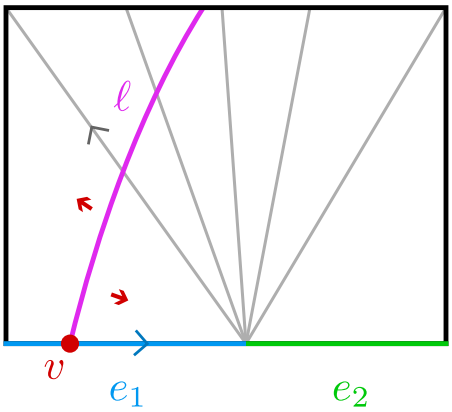

Figure 8.7: If ℓ has a vertex on e_1, then it does not have a well-defined co-orientation, which is a contradiction.

We now make the following observations about the structure of the vertical edge paths that are not part of some vertical loop.

Lemma 8.4. *For a co-oriented track in Z_f, the collection of vertical edge paths that are not part of any vertical loop are partitioned into disjoint oriented trees which each have a* root vertex *so that all maximal directed edge paths terminate at the root.*

Proof. By Claim 8.3, if multiple vertical edge paths meet at a mutual endpoint x, then at most, one is oriented away from x (oriented up). This observation implies that any loops in the union of vertical edge paths must be directed loops. Therefore, any connected component of the union of all vertical edge paths that are not part of a vertical loop must be a tree.

In such a tree, if two maximal directed edge paths terminate on distinct vertices, then because these vertices are in the same connected component, there is some path from one to the other. Either this path will contradict the maximality of one of the directed edges terminating on one of the vertices, or it will contradict Claim 8.3. □

This lemma combined again with Claim 8.3 tells us that the root vertex is the only part of a vertical tree that can also be part of a vertical loop. Every other vertex of a vertical loop or vertical tree has some edge path oriented away from it. Claim 8.3 tells us that vertical loops are disjoint, and Lemma 8.4 tells us that other vertical edge paths cannot connect vertical loops. Therefore, any connected component of the set of vertical edge paths contains at most one vertical loop.

Definition 8.4. Given a co-oriented track in Z_f, a **vertical tree** is a connected component of the union of vertical edge paths that are not part of any vertical loop. See Figure 8.8 for a picture of a vertical tree.

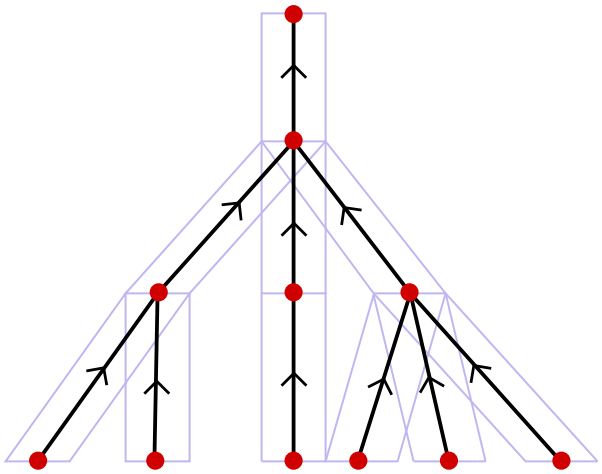

Figure 8.8: A vertical tree with the root vertex at the top. The vertices depicted in this tree are vertices of Σ that are on horizontal 1-cells of $Z_f^{(1)}$, and the edges of the tree are vertical edge paths of Σ. The outlines of rectangles around the tree show how a neighborhood of the vertical tree sits in the universal cover of Z_f.

Properties of Vertical Loops in Co-Oriented Tracks

The following lemma tells us that distinct vertical loops that are not homotopic are, in some sense, far away from one another.

Lemma 8.5. *Suppose ℓ_α and ℓ_β are two vertical edge paths of a co-oriented track in Z_f, such that ℓ_α and ℓ_β are within the same rectangular 2-cell. Suppose that ℓ_α and ℓ_β are subsets of some vertical loops, γ_α and γ_β respectively. If ℓ_α and ℓ_β have compatible co-orientations, then γ_α and γ_β are distinct vertical loops that are in-step.*

Proof. Suppose two vertical edge paths, ℓ_α, and ℓ_β, are within the same rectangular 2-cell with compatible co-orientations. Let ℓ_α be co-oriented towards ℓ_β. Also suppose ℓ_α and ℓ_β

are part of vertical loops, γ_α and γ_β respectively. Following the loops $\bar{\gamma}_\alpha$ and $\bar{\gamma}_\beta$ (γ_α and γ_β with the opposite orientations) starting at the bottom vertices of ℓ_α and ℓ_β, if they never split into different sheets at the bottom of a rectangle, then γ_α and γ_β are in-step.

Suppose instead that $\bar{\gamma}_\alpha$ and $\bar{\gamma}_\beta$ do split into different sheets at the bottom of some rectangle r. See this scenario depicted in Figure 8.9. The vertical edge paths in $\bar{\gamma}_\alpha$ and $\bar{\gamma}_\beta$ that are above that split within r, must still have compatible co-orientations, and that vertical edge path in $\bar{\gamma}_\alpha$ must still be co-oriented towards that vertical edge path of $\bar{\gamma}_\beta$. We will view $\bar{\gamma}_\alpha$ and $\bar{\gamma}_\beta$ as paths that start and end at their vertex at the bottom of r where this split occurs. For $\bar{\gamma}_\alpha$ call that vertex v_0, and for $\bar{\gamma}_\beta$ call that vertex w. Note that the 1-cells of the bottom arc of r that v_0 and w are on must be oriented in the same direction by property 5 of Lemma 8.2, and that by the choice that ℓ_α is the one co-oriented towards ℓ_β, the 1-cell that v_0 is on is oriented towards w.

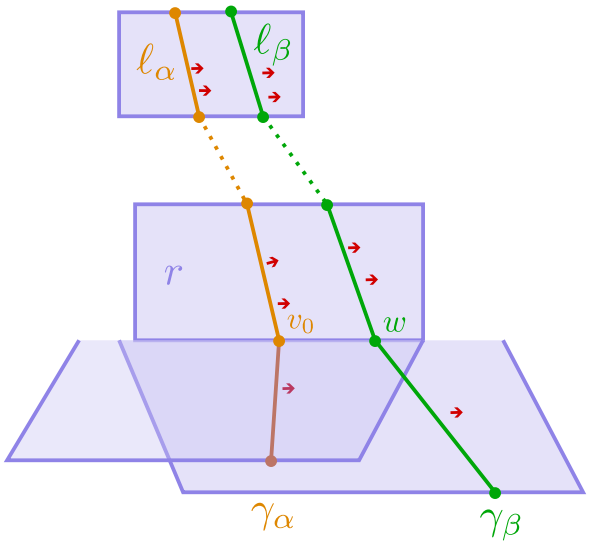

Figure 8.9: Vertical loops γ_α and γ_β split into different sheets at the bottom of the rectangle r.

Consider the lift of $\bar{\gamma}_\beta$ to the universal cover of Z_f. Call this lift $\tilde{\gamma}_\beta$ and let $\tilde{r}_n, \tilde{r}_{n-1}, \ldots, \tilde{r}_1$ be the sequence of rectangles that $\tilde{\gamma}_\beta$ passes through.

We fix a homeomorphism between the interior of each rectangle $\tilde{r}_n, \tilde{r}_{n-1}, \ldots, \tilde{r}_1$ and $(0,1) \times (0,1)$ as in Section 2.1 so that we can refer to left and right. We choose the homeomorphisms so that the co-orientation of $\tilde{\gamma}_\beta$ is to the right in each rectangle $\tilde{r}_n, \tilde{r}_{n-1}, \ldots, \tilde{r}_1$. By 5 of Lemma 8.2, orienting the bottom side of $\tilde{r}_i$ to the right using the interior of $\tilde{r}_i$ yields the same orientation as orienting the bottom side of $\tilde{r}_i$ to the right using the interior of $\tilde{r}_{i-1}$ does for $i = 2, \ldots n$. For any integer $i > 0$, we can similarly orient the rectangles for the lift $\tilde{\gamma}_\beta^i$ of the ith power of $\bar{\gamma}_\beta$ with the same base point so that $\tilde{\gamma}_\beta^i$ is co-oriented to the right in them.

Note that there is also a lift, $\tilde{\gamma}_\alpha$ of γ_α which is based on the top arc of $\tilde{r}_n$, and let $\tilde{v}_0$ be the vertex of $\tilde{\gamma}_\alpha$ that is on that arc. Also note that $\tilde{v}_0$ is to the left of $\tilde{\gamma}_\beta^i$ and is on some 1-cell of $\tilde{Z}_{f_{sub}}^{(1)}$ that is oriented to the right.

Property 4 of Lemma 8.2 says that $\tilde{v}_0$ has some unique lift of an edge path $\tilde{e}$ below it in $\tilde{r}_n$. Property 5 of Lemma 8.2 says that $\tilde{e}$ must be co-oriented to the right, implying that $\tilde{e}$ cannot end on the left side of $\tilde{r}_n$. Property 3 of Lemma 8.2 implies that $\tilde{e}$ also cannot end on the right side of $\tilde{r}_n$, because $\tilde{\gamma}_\beta^i$ is to the right of $\tilde{v}_0$. Finally, if $\tilde{e}$ ends on the top arc of $\tilde{r}_n$, then by property 5 of Lemma 8.2, $\tilde{e}$ must end on a 1-cell that is oriented to the left and is between the 1-cells that $\tilde{v}_0$ and $\tilde{w}$ are on. This observation implies that the bottom arc of r has at least three 1-cells (the one v_0 is on, the one w is on, and the one in between), which

contradicts the construction of $Z_{f_{sub}}^{(1)}$ having exactly two 1-cells on the bottom arc of each rectangle. Therefore, $\tilde{e}$ must end on the bottom arc of $\tilde{r}_n$.

Now, $\tilde{e}$ is a unique lift of a vertical edge path that is based at $\tilde{v}_0$ and in-step with $\tilde{\gamma}_\beta^i$. By this same reasoning, there must be another unique lift of a vertical edge path below $\tilde{e}$ that is in-step with $\tilde{\gamma}_\beta^i$, and so on. So, for all $i > 0$, there is some unique lift of a sequence of vertical edge paths, $\tilde{p}_i$ that is based at $\tilde{v}_0$ and in-step with $\tilde{\gamma}_\beta^i$. Figure 8.10 depicts this scenario when $i = 1$.

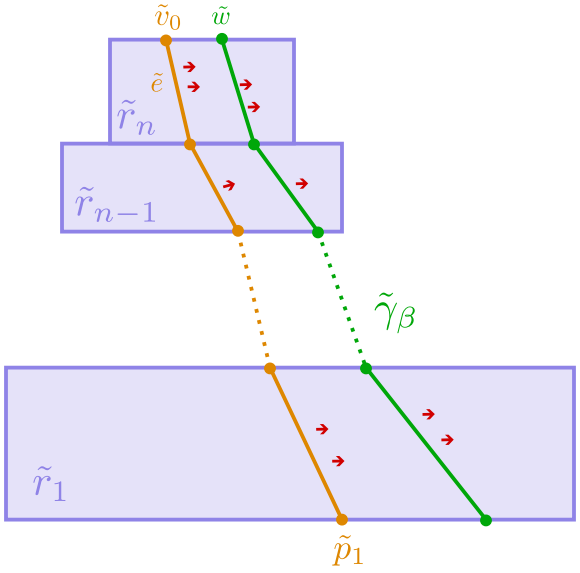

Figure 8.10: The green path on the right is the lift $\tilde{\gamma}_\beta$ in the universal cover. The orange path on the left is not the lift $\tilde{\gamma}_\alpha$, but a lift of a different path also based at v_0.

Therefore, corresponding to each $\tilde{\gamma}_\beta^i$, there is some lift of a vertex $\tilde{v}_i$ that is the bottom vertex of $\tilde{p}_i$. This results in an infinite sequence of vertex lifts $\{\tilde{v}_i\}$. If any pair of these are lifts of the same vertex, this creates a vertical loop in the projection p_i of $\tilde{p}_i$ to Z_f. By Claim 8.3, the interval of p_i starting at v_0 and ending at the repeated vertex must be part of some vertical loop. However, this vertical loop in p_i is distinct from γ_α because γ_α

enters a different sheet below v_0 than this loop. This vertical loop and γ_α both contain the vertex v_0, meaning they are not disjoint and thus contradicting Claim 8.3.

Therefore, any two distinct elements of the sequence $\{\tilde{v}_i\}$ must be lifts of distinct vertices. This observation implies that the co-oriented track has infinitely many vertices, which contradicts its finiteness. Hence, γ_α and γ_β are in-step. □

Lemma 8.5 implies that each vertical loop intersects any given rectangle at most twice, and those intersections with that rectangle must have non-compatible co-orientations.

The following claim will allow us to effectively forget about the subdivision of the cell structure of Z_f in Section 8.5.1.

Claim 8.4. Suppose G is a co-oriented track in Z_f. If γ_1 and γ_2 are two vertical loops that are in-step with compatible co-orientations, then not only do their lifts $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ pass through the same sequence of horizontal 1-cells of $\tilde{Z_f}^{(1)}$, they also pass through the same sequence of 1-cells of $\tilde{Z_f}^{(1)}_{sub}$.

As a result of being in-step with one another, $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ pass through the same sequence of horizontal 1-cells of $\tilde{Z_f}^{(1)}_{sub}$. To justify the above claim, note that by property 5 of Lemma 8.2, they must pass through the same sequence of diagonal 1-cells of $\tilde{Z_f}^{(1)}_{sub}$ in between the horizontal 1-cells.

The following lemma will associate vertical loops with closed orbits of ψ. This lemma is where we need the hypothesis in Theorem A that η is positive on all closed vertex orbits of Γ.

Lemma 8.6. *Suppose η is an integral cohomology class in the closure of the fibered cone and is positive on all closed vertex orbits in O_ψ, and $\bar{\eta}$ is a non-negative integral cocycle representative on the oriented 1-skeleton $\overset{o}{Z}{}^{(1)}_{f_{sub}}$ coming from a choice of orientations on the horizontal 1-cells of $Z^{(1)}_{f_{sub}}$. In a co-oriented track Σ relative to $\bar{\eta}$, any vertical loop is in-step with a unique closed orbit of ψ that is not a closed vertex orbit.*

Proof. By Claim 3.1, the vertical loop ℓ is naturally in-step with a unique directed cycle ℓ' of D_f. If ℓ' is not a vertex loop, then by Proposition 3.1, we are done. Suppose instead that ℓ' is a vertex loop of D_f. Then by Remark 3.2, ℓ' is homotopic a closed vertex orbit γ in Z_f. By the assumption that η is positive on all closed orbits of vertices of Γ, $\eta(\gamma) > 0$.

Due to the fact that ℓ is in-step with ℓ' while ℓ' is in-step with γ, ℓ must also be in-step with γ, implying that γ is homotopic into Σ. Due to the fact that Σ has an oriented I-bundle neighborhood in Z_f, any loop in Σ can be homotoped off of Σ by pushing it along the I-bundle neighborhood. Therefore, γ can be homotoped to have no intersection with Σ, so because Σ is dual to η, $\eta(\gamma) = 0$. This observation contradicts the fact that $\eta(\gamma) > 0$, so ℓ' must not be a vertex loop, and we are done. □

8.5 Proof of Theorem A

Let η be as in the statement of Theorem A. By Lemma 8.1, η has a non-negative integral cocycle representative, $\bar{\eta}$, on the oriented 1-skeleton $\overset{o}{Z}{}^{(1)}_{f_{sub}}$ coming from a choice of orientations on the horizontal 1-cells of $Z^{(1)}_{f_{sub}}$. By Lemma 8.2, there exists a co-oriented track Σ relative to $\bar{\eta}$ in Z_f, and by Claim 8.2 and Remark 8.1, Σ is dual to η and finite. By Claim 8.1, Σ is smoothly embedded in Z_f.

What's left to do is to redefine the mapping of Σ into Z_f so that Σ is transverse to ψ, while still maintaining all the properties of Lemma 8.2 so that Σ remains smoothly embedded and dual to η. We will also refer to redefining this mapping as **placing** Σ. This new smooth embedding of Σ will be similar to the original one. Precisely, each vertex of Σ will be mapped into the same 1-cell of $Z_{f_{sub}}^{(1)}$ as it originally was, so the new smooth embedding will be homotopic to the original. So, one could also describe this process as changing the mapping by a homotopy rather than by redefining it.

Claim 8.5. If a co-oriented track in Z_f also has the following two properties, then that graph is transverse to ψ:

1. Each edge of the graph is transverse to flowlines of ψ within each 2-cell,
2. The co-orientation of each edge within each 2-cell is towards the side of the 2-cell that is above the edge relative to ψ.

To justify the above claim, note that any co-oriented track has an oriented I-bundle neighborhood. Property 1 of Claim 8.5 will ensure that we can choose the I-bundle neighborhood such that segments of ψ passing through the I-bundle neighborhood inside of 2-cells pass from the negative to the positive side of the I-bundle neighborhood, intersecting the co-oriented track exactly once in between. Property 2 above ensures that we can also choose the I-bundle neighborhood such that segments of ψ passing through the I-bundle neighborhood as well as the 1-skeleton of Z_f pass from the negative to the positive side of the I-bundle neighborhood, intersecting the co-oriented track exactly once in between. See Figure 8.11 for an example image where property 2 is met and one where it is not.

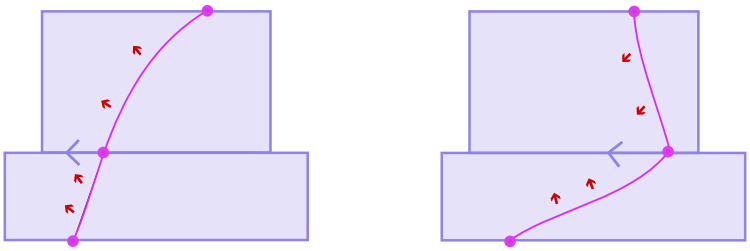

Figure 8.11: On the left is an example where condition (2) of Claim 8.5 is met, and Σ is transverse, and on the right condition (2) is not met, so Σ is not transverse.

As we place Σ, we will ensure that the mapping maintains the properties of Lemma 8.2 and satisfies the properties of Claim 8.5. By Claim 8.5, this will prove the theorem.

We will place the vertices that were on $Z_f^{(1)}$ (on the boundary of rectangular 2-cells) disjointly in the original mapping on the same 1-cells of $Z_{f_{sub}}^{(1)}$ as they were originally mapped into. Note that by property 5 of Lemma 8.2, the sequence of 1-cells that the edge paths must pass through are uniquely determined by where their endpoints were placed. Then within rectangular 2-cells, we will fill in the edge paths between pairs of those vertices, making sure the edge paths pass through the same sequence of 1-cells of $Z_{f_{sub}}^{(1)}$ as they did under the original mapping, ensuring that the new mapping satisfies properties (2) and (4) of Lemma 8.2. Due to the fact that all the vertices of Σ are mapped disjointly into the same 1-cells of $Z_{f_{sub}}^{(1)}$ as they were under the original mapping, and the original mapping satisfied properties (1) and (5) of Lemma 8.2, the redefined mapping must satisfy those properties as well.

8.5.1 Vertical Loops

We will start with the vertical loops of Σ. We take one representative of each equivalence class of vertical loops where the equivalence is being in-step and place these vertical loops. By Lemma 8.5, we can place these loops simultaneously and independently while maintaining property 3 of Lemma 8.2. We will take the vertical loop that is below all the others relative to ψ as the representative.

For each representative, ℓ, we place ℓ in Z_f by placing its path lift into the universal cover of Z_f. For the vertical loop ℓ, pick a vertex v_1 and view ℓ as a path starting and ending at v_1. Let $\tilde{\ell}$ be a path lift of ℓ to the universal cover $\tilde{Z}_f$ of Z_f, and let $\tilde{r}_1, \ldots, \tilde{r}_{n-1}, \tilde{r}_n$ be the sequence of rectangular 2-cells that $\tilde{\ell}$ passes through. Let $\tilde{e}_i \in \tilde{Z}_f^{(1)}$ be the horizontal 1-cell that is the bottom of $\tilde{r}_i$, and let γ be the unique closed orbit of ψ that is in-step with ℓ, which is not a closed vertex orbit by Lemma 8.6. Let $\tilde{\gamma}$ be the lift of γ to $\tilde{r}_1, \ldots, \tilde{r}_{n-1}, \tilde{r}_n$ and let $\tilde{p}$ be the point where $\tilde{\gamma}$ intersects $\tilde{e}_1$.

We fix a homeomorphism between the interior of each rectangle $\tilde{r}_1, \ldots, \tilde{r}_{n-1}, \tilde{r}_n$ and $(0,1) \times (0,1)$ as in Section 2.1 so that we can refer to left and right. We choose the homeomorphisms so that the co-orientation of $\tilde{\ell}$ is to the right in each rectangle $\tilde{r}_1, \ldots, \tilde{r}_{n-1}, \tilde{r}_n$. By 5 of Lemma 8.2, orienting $\tilde{e}_i$ to the right using the interior of $\tilde{r}_{i-1}$ yields the same orientation as orienting $\tilde{e}_i$ to the right using the interior of $\tilde{r}_i$ does for $i = 2, \ldots n$.

Let $\tilde{e}_{i,l}$ and $\tilde{e}_{i,r}$ be the horizontal 1-cells of $\overset{o}{Z}{}^{(1)}_{f\,sub}$ that make up $\tilde{e}_i$, such that $e_{i,r}$ is on the side of $\tilde{e}_i$. By Lemma 8.3, the original mapping of $\tilde{v}_i$ is on $\tilde{e}_{i,r}$. So, we start by placing

$\tilde{v}_1$ on $\tilde{e}_{1,r}$ to the right of $\tilde{p}$ as in Figure 8.12. There is another lift of v_1, $\tilde{v}_1'$ that is on the top arc of $\tilde{r}_n$, which must also be to the right of $\tilde{\gamma}$.

Consider the forward flow of $\tilde{v}_1$ under $\tilde{\psi}$ within $\tilde{r}_1$. Let $\tilde{q}$ be the point on $\tilde{e}_1$ such that $\psi^n(q) = v_1$. Also consider the flow of $\tilde{q}$ under $\tilde{\psi}$ within $\tilde{r}_1, \ldots, \tilde{r}_{n-1}, \tilde{r}_n$. Due to the fact that f is expanding on all scales and $\tilde{v}_1$ is to the right of $\tilde{p}$, the forward flow of $\tilde{v}_1$ will be to the right of the forward flow of $\tilde{q}$ within $\tilde{r}_1$. Therefore, we can place $\tilde{v}_2$ somewhere between these two flows, on $\tilde{e}_{2,r}$ as pictured in Figure 8.12 and satisfy condition (2) of Claim 8.5. Note that for this to be possible, when placing $\tilde{v}_1$, we should make sure it is far enough to the right such that its forward flow in $\tilde{r}_1$ ends on $\tilde{e}_{2,r}$. Then we can continue to place the rest of the vertices sequentially, placing $\tilde{v}_i$ on $\tilde{e}_{i,r}$ between the forward flow of $\tilde{v}_{i-1}$ and the forward flow of $\tilde{q}$.

We fill in the edges between vertices as straight lines within each rectangular 2-cell, ensuring that condition (1) of Claim 8.5 is satisfied. Due to the fact that each $\tilde{v}_i$ is placed between $\tilde{\gamma}$ and the forward flow of $\tilde{v}_{i-1}$ and is to the right of $\tilde{\gamma}$, this placement of ℓ satisfies condition (2) of Claim 8.5.

We have now placed one representative of each equivalence class of vertical loops. Then, for each equivalence class of vertical loops, we simultaneously place all other vertical loops in the equivalence class by taking the mapping of the representative and flowing it along ψ for some small distinct interval of time for each other element of the equivalence class, and letting the result be the image of the mapping of that other element. Due to the fact that the flow is non-degenerate, these elements of the equivalence class satisfy property

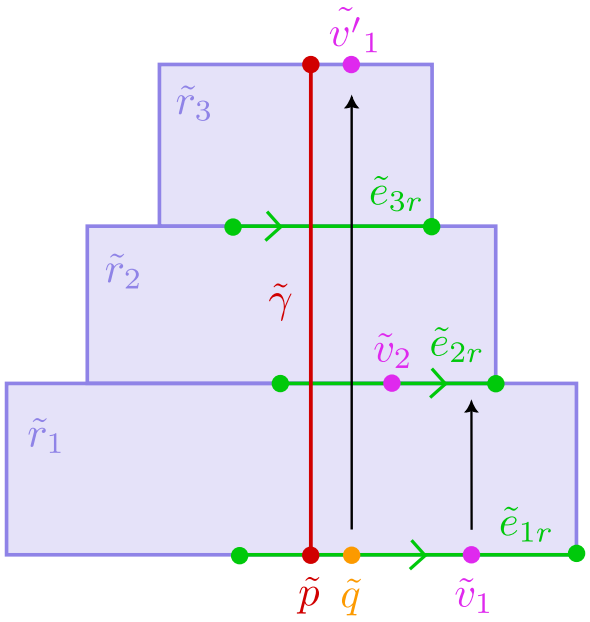

Figure 8.12: The vertex $\tilde{v}_2$ is placed between the two flows.

3 of Lemma 8.2. By Claim 8.4, the vertices of these vertical loops will still be on the same 1-cells as they were in the original mapping.

8.5.2 Vertical Trees

Now that all vertical loops have been placed, we will place the vertical trees. As we do this, we will continue to ensure that all vertices are mapped into the same 1-cell of $\overset{o}{Z}{}^{(1)}_{f_{sub}}$ as in the original mapping, and we will call this 1-cell the **pre-determined 1-cell**. We will also ensure that the ordering of the images of vertices within any 1-cell is the same as in the original mapping, and we will call this the **pre-determined order**. Doing so will allow the mapping to maintain property 3 of Lemma 8.2.

By Lemma 8.4 and Claim 8.3, each oriented tree is either disjoint from all vertical loops, or its root is on a vertical loop, and it is otherwise disjoint from all vertical loops.

We will work down from the root vertices to place the collection of vertical trees. First, we place any root vertices that have not yet been placed, map them into the pre-determined horizontal 1-cell and on the pre-determined side of any vertices that have already been placed on that 1-cell. Then we place all of the vertices on vertical trees that are one vertical edge path from a root vertex, then all that are two vertical edge paths away from a root vertex, and so on (filling in the vertical edge paths as straight lines as we go).

At any given step of this process, and in any given rectangular 2-cell r, some vertical edge paths and corresponding endpoint vertices may have already been placed in this rectangle. The only other vertices that may be on the bottom arc of r are root vertices, as all other vertices that have been placed so far have already had an edge path placed above them.

Claim 8.6. Suppose there is a root vertex w on the bottom arc of r, and suppose v is a vertex on the bottom arc of r that is the bottom of a vertical edge path in r. Fixing a homeomorphism between the interior of r and $(0,1) \times (0,1)$ as in Section 2.1 so that we can refer to left and right, suppose the vertical edge path above v in r is co-oriented to the right. Then, according to the pre-determined ordering of vertices, v must be to the right of w on the bottom of r.

The claim holds because, by Lemma 8.3, v must be on the right 1-cell of the bottom arc of r. So, if w is on the left 1-cell of the bottom arc of r, we are done. By property 5 of Lemma 8.2 that right 1-cell of the bottom arc of r must be oriented to the right and by Property 4 of Lemma 8.2 there must be an edge path e of Σ adjacent to w in r. Therefore,

if w is on the right 1-cell of the bottom arc of r, property 5 of Lemma 8.2 says that the other endpoint of e cannot be on the right vertical 1-cell of r nor the same 1-cell as w is on. The other endpoint of e cannot be on the top arc of r because w was a root vertex. Therefore, the other endpoint of e must be on the left 1-cell of r as depicted on the right side of Figure 8.13, or it must be the left 1-cell of the bottom arc of r as depicted on the left side of Figure 8.13. In either of these cases, property 3 of Lemma 8.2 implies that v must be to the right of w.

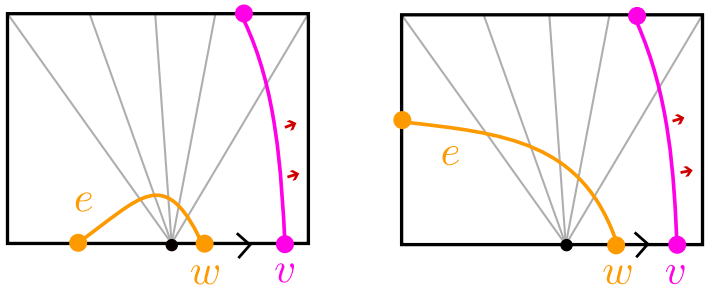

Figure 8.13: The vertex v must be to the right of root vertex w, as is the case in these two images.

For each rectangle, we fix a homeomorphism between the interior of the rectangle and $(0,1) \times (0,1)$ as in Section 2.1 so that we can refer to left and right. In any given rectangle r, for the first step of placing the collection of vertical trees, we pick out any root vertices on the top arc of r, and for all other steps, we pick out the vertices on the top arc of r that were just placed in the previous step. For each of those vertices on the top arc of r, we place the corresponding vertices (other endpoint of the adjacent vertical edge path) on the bottom arc of r on their pre-determined 1-cells, and in the pre-determined order with one another and with any previously placed vertices. We also make sure to place these corresponding vertices so that if the vertex on the top arc of r is on a 1-cell oriented to the

right (or left) of r, then the corresponding vertex on the bottom arc of r is placed such that ψ flows it into the top arc of r on the right (or left) side of the vertex it corresponds to. It is possible to place the vertices in this way by Lemma 8.3, and because any previously placed vertical edge paths will also have this property. Placing the vertices this way ensures that the mapping satisfies property 2 of Claim 8.5. We connect the corresponding vertices by straight edge paths in r, necessarily placing any vertices on that edge path on the correct 1-cell of $Z_{f_{sub}}^{(1)}$, and ensuring that the mapping satisfies property 1 of Claim 8.5. Due to the fact that all the trees are finite, this process will terminate.

Now, all the vertical trees and vertical loops have been placed.

8.5.3 Horizontal Edge Paths, Bumps, and Diagonal Edge Paths

We will now place the horizontal edge paths. If any vertices needed for these edge paths have not been placed yet, place them on the relevant vertical 1-cells anywhere that respects the pre-determined ordering of any previously placed vertices. The 2-cells that contain horizontal edge paths must still be empty at this stage, and so we can put in the horizontal edge paths straight across the 2-cells as in Figure 8.14, satisfying property 1 of Claim 8.5. The co-orientation of these edge paths will necessarily point to the side of the rectangle that is above them relative to ψ because the vertical 1-cells are oriented in that direction, so they satisfy property 2 of Claim 8.5. Property 3 of Lemma 8.2 will be satisfied because it was in the original mapping, and all vertices were placed in their pre-determined orders.

For the bumps and diagonal edge paths, place any other necessary vertices on the pre-determined 1-cell and in the pre-determined order with any previously placed vertices.

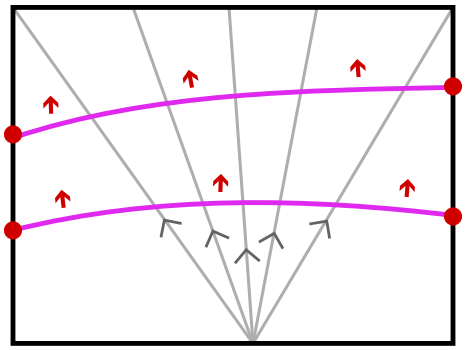

Figure 8.14: The red circles are vertices of Σ and the purple arcs are edge paths of Σ. The edge paths in this image are horizontal edge paths of Σ.

Then fill in the bump edge paths and diagonal edge paths, taking care to make them transverse to flowlines of ψ as in Figure 8.15 so that they satisfy property 1 of Claim 8.5. Once the bumps and diagonal edge paths are placed transversely to flowlines of ψ in each rectangle, the co-orientations on them must point above them relative to ψ, meaning they satisfy property 2 of Claim 8.5. For bumps, this is because any bump crosses some diagonal 1-cell of $\overset{o}{Z}{}^{(1)}_{f_{sub}}$, which is oriented in that direction. For diagonal edge paths, it is because they end on some vertical 1-cell, which is oriented in that direction. Again, property 3 of Claim 8.5 will be satisfied because it was in the original mapping, and all vertices were placed in their pre-determined orders.

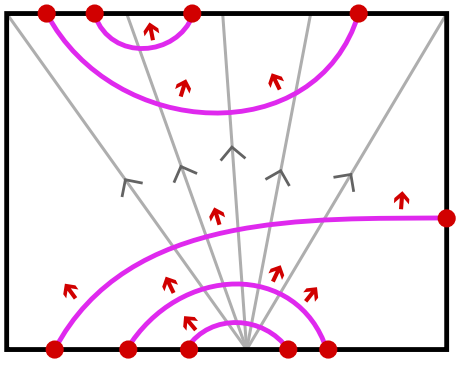

Figure 8.15: The red circles are vertices of Σ and the purple arcs are edge paths of Σ. The edge paths in this image are bumps and diagonal edge paths of Σ.

Now, all of Σ has been placed in Z_f so that the mapping satisfies the properties of Lemma 8.2 and Claim 8.5. Thus, Σ is a smoothly embedded co-oriented track in Z_f that is finite and dual to η and transverse to ψ. $\square$

REFERENCES

[1] Warren Dicks and Martin John Dunwoody. *Groups acting on graphs*, volume 17. Cambridge University Press, 1989.

[2] Spencer Dowdall, Ilya Kapovich, and Christopher J. Leininger. Dynamics on free-by-cyclic groups. *Geom. Topol.*, 19(5):2801–2899, 2015.

[3] Spencer Dowdall, Ilya Kapovich, and Christopher J. Leininger. McMullen polynomials and Lipschitz flows for free-by-cyclic groups. *J. Eur. Math. Soc. (JEMS)*, 19(11):3253–3353, 2017.

[4] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston's work on surfaces*, volume 48 of *Mathematical Notes*. Princeton University Press, Princeton, NJ, 2012. Translated from the 1979 French original by Djun M. Kim and Dan Margalit.

[5] David Fried. Fibrations over S^1 with pseudo-Anosov monodromy. *Travaux de Thurston sur les surfaces*, 66:251–266, 1979.

[6] François Gautero. Cross sections to semi-flows on 2-complexes. *Ergodic Theory Dynam. Systems*, 23(1):143–174, 2003.

[7] Michael P. Landry, Yair N. Minsky, and Samuel J. Taylor. Endperiodic maps via pseudo-anosov flows. *arXiv:2304.10620*, 2023.

[8] Michael P. Landry, Yair N. Minsky, and Samuel J. Taylor. Flows, growth rates, and the veering polynomial. *Ergodic Theory Dynam. Systems*, 43(9):3026–3107, 2023.

[9] Curtis T. McMullen. Entropy and the clique polynomial. *J. Topol.*, 8(1):184–212, 2015.

[10] Lee Mosher. Dynamical systems and the homology norm of a 3-manifold ii. *Inventiones mathematicae*, 107(2):243–282, 1992.

[11] W.P. Thurston. *A Norm for the Homology of 3-manifolds*. Memoirs of the American Mathematical Society. American Mathematical Society, 1986.