Surface maps into free groups

Thesis by

Alden Walker

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Department of Mathematics
Caltech
Pasadena CA, 91125

2012
(Defended May 18, 2012)

Acknowledgements

I thank my advisor, Danny Calegari, for his guidance and many explanations. I thank Josh Sabloff, Lynne Butler, and Miller Maley for the math and opportunities. I thank my parents for their love and money. I thank Alexis for being awesome.

Abstract

We exploit the combinatorial properties of surface maps into free groups to prove several new results in the field of stable commutator length and bounded cohomology. We show that random homomorphisms between free groups are isometries of scl; we prove interesting properties of the scl unit ball; we describe a transfer construction for quasimorphisms and give an infinite family of chains whose scl it certifies; we linearize the dynamics of endomorphisms on free groups and use this to prove that random endomorphisms can be realized by surface immersions, which provides many examples of surface subgroups of HNN extensions of free groups; and finally, we give an algorithm to compute scl in free products of finite or infinite cyclic groups that generalizes and improves previous work.

Contents

List of Figures

Chapter 1

Introduction

1.1 Introduction

In this chapter, we give a very brief introduction to the subjects of this thesis and state our main results. In Chapter 2, we give more detailed background material, and the subsequent chapters form the main body.

1.2 Brief background

1.2.1 Stable commutator length

Let G be a group and let $X = K(G,1)$. There are homology chain groups $C_*(G,\mathbb{Z}) = C_*(X,\mathbb{Z})$, and we denote 1-boundaries by $B_1(G,\mathbb{Z})$. The space $B_1(G,\mathbb{Z})$ is the free $\mathbb{Z}$ module spanned by formal sums $\sum_{i=1}^n g_i$ of group elements with the property that there is a surface map $f : S \to X$ where S has n boundary components $\{(\partial S)_i\}$ and $f((\partial S)_i) = \gamma_i$ as loops in X, and γ_i is a loop representing the conjugacy class of g_i.

The study of stable commutator length is motivated by the question: given a collection of loops which bound a surface in X, what is the simplest surface that they bound? It is natural to stabilize the definition; that is, to allow sufaces to have multiple boundary components mapping to the same loop, and perhaps with degree greater than one. A surface map $f : S \to X$ is admissible for $\sum_i g_i$ if the following commutative diagram holds

$$\begin{array}{ccc} S & \xrightarrow{f} & X \\ \uparrow & & \uparrow \scriptstyle\coprod_{i=1}^n \gamma \\ \partial S & \xrightarrow{\partial f} & \coprod_{i=1}^n S^1_i \end{array},$$

where each S^1_i is a copy of S^1 and $\gamma_i : S^1_i \to X$ represents (the conjugacy class of) g_i. Furthermore, we define $n(S,f)$ to be the integer such that $\partial f_*[\partial S] = n(S,f)\left[\coprod_i S^1_i\right]$. The number $n(S,f)$ records how many times the surface wraps around each loop.

We define

$$\mathrm{scl}(\sum_i g_i) = \inf_{(S,f)} \frac{-\chi^-(S)}{2n(S,f)},$$

where $\chi^-(S)$ is the Euler charateristic disregarding disk and sphere components. It extends to $B_1(G;\mathbb{R})$, and we define $B_1^H(G) = B_1(G;\mathbb{R})/\langle xyx^{-1} = y, y^n = ny\rangle$. This will be our main space of study. We call an admissible surface map *extremal* for a chain if it realizes the infimum.

Stable commutator length is mysterious and complicated, even in groups as simple as free groups.

1.2.2 Quasimorphisms

A function $\phi : G \to \mathbb{R}$ is a *homogeneous quasimorphism* if there exists a finite real number $D(\phi) < \infty$ such that $|\phi(x) + \phi(y) - \phi(xy)| \leq D(\phi)$ for all $x, y \in G$, and $\phi(x^n) = n\phi(x)$ for all $x \in G$ and $n \in \mathbb{Z}$. We denote the space of homogeneous quasimorphisms by $Q(G)$. Note that $Q(G)$ contains the space $H^1(G)$ of real-valued homomorphisms of G.

Homogeneous quasimorphisms extend by linearity to give well-defined functions on $B_1^H(G)$. These functions are intimately connected with scl.

Theorem (Bavard duality, Theorem 2.79 in [8], generalizing [2])**.** *For $C \in B_1^H(G)$,*

$$\mathrm{scl}(C) = \sup_{\phi \in Q(G)/H^1(G)} \frac{\phi(C)}{2D(\phi)}.$$

This theorem says that the normed vector space $(Q(G)/H^1(G), D(\cdot))$ is the functional dual space of the (pseudo-) normed vector space $(B_1^H(G), \mathrm{scl})$. We call a quasimorphism *extremal* for a chain in $B_1^H(G)$ if it realizes the supremum.

When G is a free group, scl is a genuine norm on $B_1^H(G)$.

1.3 Outline and results

The general theme of this thesis is that we can leverage the combinatorial structure of surface maps into free groups in combination with an understanding of the geometry of surfaces to prove results involving stable commutator length and quasimorphisms. In addition to theorems, we also give several interesting constructions, and the final chapter gives an algorithm. These tools allow us to gain a better understanding of the situation, and they promote further study.

Let G and H be groups. We say that a homomorphism $\phi : G \to H$ is an *isometry* for scl if $\mathrm{scl}_H(\phi(C)) = \mathrm{scl}_G(C)$ for all $C \in B_1^H(G)$. As we will see, this is a complicated condition, and it is stronger than ϕ being injective.

Let F_k and F_l be free groups of rank k and l. We define a *random homomorphism* of length n $\phi : F_k \to F_\ell$ to be a homomorphism obtained by defining ϕ on the generators of F_k to be words selected uniformly at random from the ball of radius n in F_ℓ. We show that

Theorem (3.3.11; Random isometry theorem)**.** *A random homomorphism $\varphi : F_k \to F_l$ of length n between free groups of ranks k, l is an isometry of* scl *with probability $1 - O(C(k,l)^{-n})$ for some constant $C(k,l) > 1$.*

The random isometry theorem shows that an algebraically random map preserves scl. In free groups, it is natural to define a geometrically random surface map, and we show that these maps are extremal, which is the natural analog of isometric. For the purposes of this introduction, it suffices to know that a *fatgraph* is essentially a combinatorial description of a surface map into a free group (Figure 2.6, for the impatient).

Theorem (3.4.11; Random fatgraph theorem)**.** *For any combinatorial fatgraph $\hat{Y}$, if Y is a random fatgraph over F_k obtained by labeling the edges of $\hat{Y}$ by words of length n, then $S(Y)$ is extremal for $\partial S(Y)$ and is certified by the extremal quasimorphism H_Y, with probability $1 - O(C(\hat{Y},k)^{-n})$ for some constant $C(\hat{Y},k) > 1$.*

Here H_Y is a quasimorphism constructed in a combinatorial way from the labeling on Y.

The previous theorems show that, in a certain sense, the generic picture of scl in free groups is well understood (in addition, see [14]). However, a detailed picture is very complicated. A realization of a free group as the fundamental group of a hyperbolic surface induces a natural *rotation quasimorphism* via the circle action at infinity. Associated to such a rotation quasimorphism is a *geometric* face of the scl norm ball in $B_1^H(F)$. These faces are important and interesting. In Chapter 4, we explore the space of rotation quasimorphisms and geometric faces of the scl norm ball. We show some interesting structural properties of the set of geometric faces.

Theorem (4.3.8; Extremal rotations quasimorphisms)**.** *Let F_2 be a free group of rank 2. There exist infinitely many commutators for which the rotation quasimorphism from any surface realization is extremal. There exist infinitely many commutators for which there is no surface realization giving an extremal rotation quasimorphism.*

A slightly less general statement is true for free groups of arbitrary rank.

Rotation quasimorphisms induced by surface realizations are very natural and have many nice properties. However, it is trivial to see that we must search for larger classes of quasimorphisms, for a rotation quasimorphism from a surface realization takes integer values. There are many chains $C \in B_1^H(F_k)$ for which $\mathrm{scl}(C)$ is not in $\frac{1}{2}\mathbb{Z}$. Therefore, we cannot hope to find an extremal quasimorphism for most chains within the class of rotation quasimorphisms. We continue Chapter 4 with the construction of *tranfers* of quasimorphisms, and in particular, transfers of rotation quasimorphisms.

In the transfer construction, we lift a quasimorphism from a finite index subgroup to the whole group. This provides a new class of quasimorphisms, and we show that this is nontrivial by exhibiting an infinite family of examples for which they are extremal.

Proposition (4.4.8)**.** *Let* $w(n) = x^2 + y^n + yx^{-1}y^{n+1}x^{-2} \in F_2 = \langle x, y \rangle$*. Then* $\mathrm{scl}(w(n)) = \frac{2n+1}{2n+2}$*, and there is a family* ϕ_n *of transfers of rotation quasimorphisms from finite index subgroups of* F_2 *such that* ϕ_n *is extremal for* $w(n)$*.*

In addition, we show that in certain cases, we can take limits of the transfer procedure to obtain limiting quasimorphisms, and we give examples of chains for which these limit transfers are extremal.

Proposition (4.7.8)**.** *Let* F_2 *be a free group of rank* 2*. There exist sequences* ϕ_n *of transfers of rotation quasimorphisms on finite index subgroups of* F_2 *with index* n *such that* ϕ_n *converges to a weak-* limit* ϕ_∞*, and there are chains for which* ϕ_∞ *is extremal.*

In Chapter 5, we study the dynamics of endomorphisms on free groups. We call an endomorphism $\phi : F_k \to F_k$ *geometric* if there is a realization of F_k as the fundamental group of a smooth surface Σ and an immersion $\psi : \Sigma \to \Sigma$ such that $\psi_* = \phi$. A geometric automorphism must be in the mapping class group of a surface Σ with $\pi_1(\Sigma) \cong F_k$, and these elements are rather rare in $\mathrm{Out}(F_k)$. In contrast, we show:

Theorem (5.3.6; Geometric endomorphisms)**.** *A random endomorphism* $\phi : F_k \to F_k$ *of length* n *is geometric with probability* $1 - C(k)^{-n}$*.*

The theorem is proved by giving a simple combinatorial condition guaranteeing that an endomorphism is geometric.

In combination with results about scl and quasimorphisms, this produces many π_1-injective maps of surface groups into HNN extensions of free groups and gives simple conditions certifying their injectivity.

In Chapter 6, we give an algorithm, `scylla`, to compute scl in free products of cyclic groups. Here the cyclic groups can be finite or infinite, so this generalizes the `scallop` algorithm (on which `scylla` is based) given in Chapter 4 of [8]. Furthermore, `scylla` is polynomial time in the length of the input, the number of free factors, and the orders of the factors. This is in contrast to `scallop`, which is exponential in the rank. The algorithm `scylla` is therefore especially useful for computing scl when dealing with finite covers of free groups, which is of interest in Chapter 5.

Chapter 2

Background

The content of this thesis comprises several results involving stable commutator length, quasimorphisms, and surface maps. There is a significant amount of overlap in the background material required for the various chapters, so it is useful to collect it in one place. Therefore, in this chapter, we present a review of stable commutator length, quasimorphisms, and surface realizations of free groups.

2.1 Stable commutator length

2.1.1 Definition

Let G be a group and $[G,G]$ the commutator subgroup of G. If $g \in [G,G]$, then g can be written as a product of commutators, and we define the *commutator length* $\mathrm{cl}(g)$ to be the least number of commutators whose product is g.

Commutator length is complicated and difficult to understand, mostly because the set of commutators in a general group can contain elements which do not "look" like commutators. In fact, in a finite nonabelian simple group, every element is a commutator [27]! Changing the context from algebra to topology gives a more intuitive picture. The commutator length of g is the least genus of a surface with one boundary component which maps into a $K(G,1)$ in such a way that the image of its boundary loop maps to a loop representing the conjugacy class of g.

Definition 2.1.1. The *stable commutator length* $\mathrm{scl}(g)$ of g is defined

$$\mathrm{scl}(g) = \lim_{n\to\infty} \frac{\mathrm{cl}(g^n)}{n}.$$

The limit exists because the map $n \mapsto \mathrm{cl}(g^n)$ is clearly subadditive (and positive). The topological definition of cl motivates us to find a topological picture for scl. It will actually be more natural to define scl in a more general setting.

Let $C = \sum_{i=1}^{k} g_i$ be a formal sum of elements $g_i \in G$. If $[C] = 0$ in $H_1(G;\mathbb{Z})$, in which case $C \in B_1(G;\mathbb{Z})$, the space of boundaries, then there exists a surface map $S \to K(G,1)$ with n boundary components such that the images of the boundary components in $\pi_1(K(G,1)) = G$ map to the conjugacy classes of the g_i. We say that a surface map $f : S \to K(G,1)$ is *admissible* for the formal sum $\sum_{i=1}^{n} g_i$ if the following commutative diagram holds:

$$\begin{array}{ccc} S & \xrightarrow{f} & K(G,1) \\ \uparrow & & \uparrow{\scriptstyle \coprod_{i=1}^{n} \gamma} \\ \partial S & \xrightarrow{\partial f} & \coprod_{i=1}^{n} S_i^1 \end{array} \quad ,$$

where each S_i^1 is a copy of S^1 and $\gamma_i : S_i^1 \to K(G,1)$ represents (the conjugacy class of) g_i. Furthermore, we define $n(S,f)$ to be the integer such that $\partial f_*[\partial S] = n(S,f)\left[\coprod_i S_i^1\right]$. Informally, a surface map is admissible for a collection of loops if its boundary maps to some multiple of the loops. Note that an admissible surface map may have multiple boundary components mapping to the same copy of S^1, and in fact with different degrees.

Definition 2.1.2. Let the function $\chi^-(\cdot)$ be the Euler characteristic disregarding disk and sphere components. Then define

$$\operatorname{scl}\left(\sum_{i=1}^{k} g_i\right) = \inf_{S,f} \frac{-\chi^-(S)}{2n(S,f)},$$

where the infimum is taken over all surface maps (S,f) which are admissible for $\sum_{i=1}^{k} g_i$.

Remark 2.1.3. If $g \in [G,G]$, then $g \in B_1(G;\mathbb{Z})$. It is a proposition that Definitions 2.1.1 and 2.1.2 agree for such elements of $B_1(G;\mathbb{Z})$. See Propositions 2.10 and 2.74 in [8]. We have omitted a more algebraic generalization of Definition 2.1.1 because we will not need it.

Because scl has nice linearity properties (Lemmas 2.75 and 2.76 in [8]), it extends by linearity to $B_1(G;\mathbb{Q})$. Furthermore, it is continuous, so it extends by continuity to $B_1(G;\mathbb{R})$. Henceforth, we will supress the coefficient group in the notation and always take it to be $\mathbb{R}$, so $B_1(G) = B_1(G;\mathbb{R})$.

Definition 2.1.4. Let $H(G)$ be the subspace of $B_1(G)$ spanned by all elements of the form $hgh^{-1} - g$ and $g^n - ng$ for $g, h \in G$ and $n \in \mathbb{Z}$. We define $B_1^H(G) = B_1(G)/H(G)$.

Lemma 2.1.5 (§2.6.1 in [8])**.** *For any group G,* scl *gives a well-defined function on $B_1^H(G)$; that is,* scl *vanishes on $H(G)$.*

It is more natural to work in $B_1^H(G)$. We have the following properties of scl, from §2.6.1 in [8]:

Homogeneity. For $g \in G$ and $n > 0$, $\operatorname{scl}(g^n) = n\operatorname{scl}(g)$.

Conjugation invariance. For $g, h \in G$, $\operatorname{scl}(hgh^{-1}) = \operatorname{scl}(g)$.

Sublinearity. For $x, y \in B_1^H((G)$, $\mathrm{scl}(x + y) \leq \mathrm{scl}(x) + \mathrm{scl}(y)$.

In other words, scl is always a pseudo-norm on $B_1^H(G)$. It is not always a norm because it might vanish on a nonzero element. However, if G is word hyperbolic, or in particular free, then scl is a genuine norm [10].

2.1.2 Extremal surfaces and examples

Intuitively, $\mathrm{scl}(C)$ gives the most "efficient" surface which bounds loops representing C in $K(G, 1)$. It is not true in general that the infimum is actually realized by a surface map. For example, there are groups for which scl takes transcendental values on certain elements (see [33]). In particular, these values are not rational, so they cannot be realized by finite surface maps.

However, sometimes there is a surface map realizing the infimum. That is, for a chain $C \in B_1^H(G)$, a surface map $f : S \to K(G, 1)$ such that $\mathrm{scl}(C) = -\chi^-(S)/2n(S, f)$. We call such surfaces *extremal* for C.

Example 2.1.6. Let G be any group, and let $[g, h]$ be a commutator. Then $\mathrm{scl}([g, h]) \leq 1/2$. To see this, observe, that $[g, h]$ is the boundary of a tautological once-punctured torus in a $K(G, 1)$. Therefore, $\mathrm{scl}([g, h]) \leq -(-1)/2 = 1/2$.

Example 2.1.7. Let $G = \langle a, b\rangle$. Then $a + b + a^{-1}b^{-1} \in B_1^H(G)$, and $\mathrm{scl}(a + b + a^{-1}b^{-1}) \leq 1/2$. To see this, observe that there is a simple thrice-punctured sphere with three boundary components mapping to a, b, and $a^{-1}b^{-1}$. Since a thrice-punctured sphere has Euler characteristic -1, we have $\mathrm{scl}(a + b + a^{-1}b^{-1}) \leq -(-1)/2$. In fact, this surface is extremal.

Theorem 2.1.8 (§4.1 in [8])**.** *For any chain $C \in B_1^H(F)$ for F a free group, an extremal surface for C exists.*

Figures 2.3, 2.4a, and 2.4b give some pictures of scl in free groups.

2.2 Quasimorphisms

2.2.1 Definition

Let G be a group, and let $\phi : G \to \mathbb{R}$ be a map. If $\phi(g) + \phi(h) - \phi(gh) = 0$ for all $g, h \in G$, then h is a homomorphism. We generalize this:

Definition 2.2.1. Let $\phi : G \to \mathbb{R}$. If there exists a finite D such that $|\phi(g) + \phi(h) - \phi(gh)| \leq D$, then we call ϕ a *quasimorphism*. The least such number D is called the *defect* of ϕ and is denoted by $D(\phi)$.

Let $\hat{Q}(G)$ denote the real vector space of quasimorphisms on G. If $\phi \in \hat{Q}(G)$ satisfies $\phi(g^{-1}) = -\phi(g)$, then we call ϕ *antisymmetric*. The *antisymmetrization* ϕ' is defined $\phi'(x) = \frac{1}{2}(\phi(x) - \phi(x^{-1}))$. Note that antisymmetrization does not increase defect. If $\phi(g^n) = n\phi(g)$ for all $n \in \mathbb{Z}$, then we call ϕ *homogeneous*. Let $Q(G)$ denote the space of homogeneous quasimorphisms. If $\phi \in \hat{Q}(G)$, then we define the homogenization $\bar{\phi}$:

$$\bar{\phi}(g) = \lim_{n\to\infty} \frac{\phi(g^n)}{n}.$$

This limit exists, and the homogenization satisfies $|\phi(g) - \bar{\phi}(g)| \leq D(\phi)$, by Lemma 2.21 in [8]. Furthermore, $D(\phi) \leq D(\bar{\phi}) \leq 2D(\phi)$ by Corollary 2.59 in [8]. Notice that the space $H^1(G)$ of homomorphisms on G sits inside the space of homogeneous quasimorphisms.

2.2.2 Duality with scl

Homogeneous quasimorphisms extend by linearity to give well-defined functions on $B_1^H(G)$. These functions are intimately connected with scl.

Theorem 2.2.2 (Bavard duality, Theorem 2.79 in [8], generalizing [2])**.** *For $C \in B_1^H(G)$,*

$$\mathrm{scl}(C) = \sup_{\phi \in Q(G)/H^1(G)} \frac{\phi(C)}{2D(\phi)}.$$

This theorem says that the normed vector space $(Q(G)/H^1(G), 2D(\cdot))$ is the functional dual space of the (pseudo-) normed vector space $(B_1^H(G), \mathrm{scl})$. If a quasimorphism realizes the supremum, then we call it *extremal* for C. By the Hahn-Banach theorem, an extremal quasimorphism for C always exists. This is in contrast to extremal surfaces for C, which might not exist.

Observe that if $\phi \in Q(G)$ is any quasimorphism, and $f : S \to K(G,1)$ is any surface map admissible for C, then

$$\frac{\phi(C)}{2D(\phi)} \leq \mathrm{scl}(C) \leq \frac{-\chi^-(S)}{2n(S,f)}.$$

The most satisfying way to compute scl, then, is to exhibit a surface map and a homogeneous quasimorphism which gave matching upper and lower bounds.

2.2.3 Counting quasimorphisms

In free groups, there is a very natural class of quasimorphisms: those that count subwords.

Definition 2.2.3. Let σ and w be words in a free group F. Define $C'_\sigma(w)$ to be the number of (possibly overlapping) occurrences of σ as a subword of w. Define the *counting function* C_σ to be

$$C_\sigma(w) = \lim_{n\to\infty} \frac{C'_\sigma(w^n)}{n}.$$

The *big counting quasimorphism* H_σ is then

$$H_\sigma(w) = C_\sigma(w) - C_{\sigma^{-1}}(w).$$

Big counting quasimorphisms are easy to compute. Let $\overline{w}$ denote the cyclic, and cyclically reduced, word obtained from w by writing w around a circle and deleting any cancelling subwords at the beginning and end of w.

Lemma 2.2.4. $C_\sigma(w) = C'_\sigma(\overline{w})$.

Proof. Any subword of w^n is a cyclic subword of $\overline{w}$. $\square$

We can also count disjoint subwords:

Definition 2.2.5. Let $c'_\sigma(w)$ count the maximum number of *disjoint* copies of σ as a subword of w. As above, we define $c'_\sigma(w) = \lim_{n\to\infty} c'_\sigma(w^n)/n$. The *small counting quasimorphism* h_σ is

$$h_\sigma(w) = c_\sigma(w) - c_{\sigma^{-1}}(w).$$

Both the big and small counting quasimorphisms are, as advertised, quasimorphisms, and they are homogeneous. They were introduced by Brooks in [5] and generalized to hyperbolic groups by Epstein-Fujiwara [19] and Fujiwara [20]. Also see [10].

Big counting quasimorphisms are easier to compute, but small counting quasimorphisms have a uniformly bounded defect.

Lemma 2.2.6 (Corollary of Proposition 2.30, [8])**.** *For any small counting quasimorphism h_σ, we have $D(h_\sigma) \le 6$.*

Note that there is a difference in notation in [8]; our counting quasimorphisms are already homogenized.

Brooks counting quasimorphisms are simple linear combinations of counting functions. One might wonder whether any other linear combinations yield quasimorphisms. The answer is no.

Lemma 2.2.7. *If S is a finite set of words, and $\sum_{\sigma\in S} C_\sigma$ is a homogeneous quasimorphism, then $\sum_{\sigma\in S} C_\sigma = \frac{1}{2}\sum_{\sigma\in S} H_\sigma$. The analogous statement is true for small counting quasimorphisms.*

Proof. Since $\sum_{\sigma\in S} C_\sigma$ is a homogeneous quasimorphism, it is its own antisymmetrization. But its antisymmetrization is exactly the linear combination of Brooks counting quasimorphisms above. $\square$

2.2.4 Rotation quasimorphisms

In this section, we describe a class of quasimorphisms arising naturally from circle actions.

There is a universal central extension

$$0 \to \mathbb{Z} \to \mathrm{Homeo}^+(S^1)^\sim \to \mathrm{Homeo}^+(S^1) \to 0.$$

The group $\mathrm{Homeo}^+(S^1)^\sim$ can be thought of as the group of orientation preserving homeomorphisms of $\mathbb{R}$ which commute with integer translation, or equivalently, the group of homeomorphisms of $\mathbb{R}$ which cover homeomorphisms of S^1.

There are a few different ways to define the rotation number on $\mathrm{Homeo}(S^1)$ and $\mathrm{Homeo}^+(S^1)^\sim$. We define $\mathrm{rot}^\sim : \mathrm{Homeo}^+(S^1)^\sim \to \mathbb{R}$ by

$$\mathrm{rot}^\sim(f) = \lim_{n\to\infty} \frac{f^{\circ n}(0)}{n}.$$

This limit is well defined, and any other point can take the place of 0; the resulting limit is the same. This map descends to $\mathrm{rot} : \mathrm{Homeo}^+(S^1) \to \mathbb{R}/\mathbb{Z}$, which is Poincaré's rotation number. If we let $t \in \mathrm{Homeo}^+(S^1)^\sim$ be the generator of the center, so t is integer translation by 1, then we observe that $\mathrm{rot}(t^{\circ n} \circ f) = n + \mathrm{rot}(f)$. The function $\mathrm{rot}^\sim : \mathrm{Homeo}^+(S^1)^\sim \to \mathbb{R}$ is a homogeneous quasimorphism with defect 1.

For any action of a group on a circle $\rho : G \to \mathrm{Homeo}^+(S^1)$, there is an associated bounded Euler class $e_b(\rho,\mathbb{Z}) \in H_b^2(G;\mathbb{Z})$. The image $e(\rho;\mathbb{Z})$ in $H^2(G;\mathbb{Z})$ is the ordinary Euler class. The representation lifts to $\rho^\sim : G \to \mathrm{Homeo}^+(S^1)^\sim$ if and only if the ordinary Euler class $e(\rho,\mathbb{Z})$ vanishes. If the representation does lift, then the different lifts are parameterized by $H^1(G,\mathbb{Z})$ (essentially, by choosing a particular lift for each generator).

The Euler class determines the representation up to semiconjugacy:

Theorem 2.2.8 (Ghys [21])**.** *Two actions $\rho_1, \rho_2 : G \to \mathrm{Homeo}(S^1)$ give rise to the same $e_b(\rho;\mathbb{Z}) \in H_b^2(G;\mathbb{Z})$ if and only if ρ_1 and ρ_2 are* semiconjugate*; i.e., if and only if there is some third action ρ and monotone maps $\pi_1, \pi_2 : S^1 \to S^1$ such that $\rho_i\pi_i = \pi_i\rho$ for $i = 1, 2$.*

For background on the Euler class and rotation number, see [6], [31], and [22].

We will be interested in the pullback of the rotation quasimorphism via representations, and we therefore need to know how the defect behaves under pullback.

Proposition 2.2.9 (Thurston [31], Proposition 3.3)**.** *Let $\rho : G \to \mathrm{Homeo}^+(S^1)$ be a representation with associated* $\mathrm{rot}^\sim$*. Then either ρ is semiconjugate to a group of rigid rotations, in which case $D(\mathrm{rot}^\sim) = 0$, or else ρ is semiconjugate to a minimal action with centralizer of finite order n, and $D(\mathrm{rot}^\sim) = 1/n$.*

2.3 Free groups

2.3.1 Notation

We let F_k be a free group of rank k. Usually, our free groups will be generated by a, b, $\ldots$, and sometimes by $a_1, \ldots, a_k$. Where convenient, we will denote the inverses of generators by capital letters, so $A = a^{-1}$. We let X_k be a $K(F_k, 1)$. For simplicity, we will always take X_k to be a rose. That is, X_k is a simplicial complex with a single 0-cell and k 1-cells, each thought of as a directed edge labeled by a generator of F_k.

2.3.2 Stallings graphs

It is very convenient to think of free groups and subgroups of free groups as the fundamental groups of graphs.

Definition 2.3.1. A *Stallings graph* [30] G over a free group F_k is a directed graph with edges labeled by generators $a_1, \ldots, a_k$ of F_k.

A directed loop in G gives a word in F_k by reading edge labels in order as appropriate (going across an edge with the opposite orientation reads off the inverse of the generator). Therefore, a Stallings graph over F_k has a natural identification with a subgroup of F_k. Every subgroup H of F_k has an associated Stallings graph: for example, take a set of generators for H, and take the Stallings graph which is the rose on a set of loops, each labeled with a generator of H.

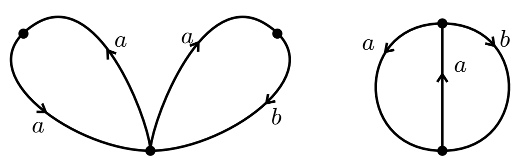

Figure 2.1: A Stallings graph over $\langle a, b\rangle$ with fundamental group the subgroup $\langle aa, ab\rangle$. Folding the two middle edges together produces a folded Stallings graph for the same subgroup (right).

A vertex v of a Stallings graph is *folded* if there is at most one incident incoming edge with any label and at most one incident outgoing edge with any label. A Stallings graph is *folded* if each of its vertices is folded. If a Stallings graph contains a vertex which is not folded, or *unfolded*, then we can perform a local transformation of the Stallings graph which folds the offending edges together. If there are two edges with the same label and same origin and destination, then we simply remove one of them — we think of it as folding the edges together. If there are two incoming (resp outgoing) edges e_1 and e_2 with the same direction and label, then we identify together the origin (resp destination) vertices of e_1 and e_2 and remove one of e_1 and e_2. Repeatedly performing the folding operation on a finite graph will eventually produce a folded graph. Observe that the folding operation does not change the subgroup of F_k associated with the Stallings graph. Therefore, every subgroup of F_k is associated with a folded Stallings graph. See Figure 2.1.

2.3.3 Subgroups and covers

One of the strengths of using Stallings graphs to understand subgroups of F_k is that the finite index subgroups of F_k are exactly the labeled covers of the Stallings graph X_k. For example, this provides a simple proof of the following.

Lemma 2.3.2. *Given a finitely generated proper subgroup $H \subsetneq F_k$, there is a finite index proper subgroup $N \subsetneq F_k$ so that $H \subseteq N$.*

Proof. Take a Stallings graph for H and fold it. If H is a proper subgroup, the resulting graph will not be a $K(F_k, 1)$. The only reason it is not a cover is that it is missing some edges. Add edges (and a vertex, if necessary) so that (1) there is more than one vertex and (2) every vertex has an incoming and outgoing edge labeled with each generator. The result is a finite cover of X_k whose fundamental group is therefore a finite index subgroup N containing H. $\square$

Of course, this is also consequence of the more powerful (and harder to prove) fact that finitely generated subgroups of a free group are separable (free groups are LERF) (see [23]).

2.3.4 The scl norm ball

In a free group F_k, scl is a norm on $B_1^H(F_k)$ (see [10]). It is natural to wonder what the scl norm ball looks like, since this is a generalization of wondering what the scl spectrum looks like. In fact, the scl norm ball in a free group is an infinite dimensional polyhedron, in the sense that its intersection with any finite-dimensional subspace of $B_1^H(F_k)$ is a polyhedron. This is a consequence of Calegari's `scallop` algorithm to compute scl in free groups; see [7], §4.1.

While they are polyhedra, these finite-dimensional slices of the norm ball can be quite complicated. In the large scale, however, scl looks like an L^1 norm; see [14].

2.3.5 Pictures of scl in free groups

It is motivating to observe the complexity of scl through experiments. Pictures of scl tend to be simple at a large scale and complex at a small scale, and this is formalized in [11] and [14].

Figure 2.2a show a histogram of the scl values of many random words of length 40 in F_2. Observe that the histogram looks moderately Gaussian, except for the large spike at 1.5. Comparing a cumulative histogram to the appropriate Gaussian cumulative distribution function in Figure 2.2b shows the marked similarity, again except for the jump at 1.5. We remark that while certain convergence results are known ([11] and [14] again), it is not known whether this apparent convergence is true.

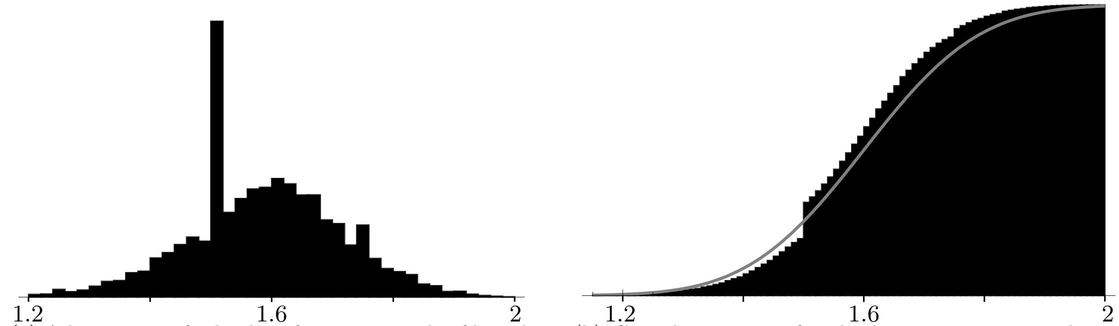

(a) A histogram of scl values for many words of length 40.

(b) Cumulative sums for the histogram compared to a Gaussian CDF.

Figure 2.2: Comparing a histogram of scl values to its cumulative distribution function.

We also remark that experimental data can be misleading: [14] proves a convergence result which occurs so slowly that we cannot hope to observe it with our current experimental capabilities.

Compare Figures 2.2a and 2.2b to a histogram of the same data with very small bins, as shown in Figure 2.3. Note the obvious complexity and self-similarily. This is discussed in more detail in [9]. See also [13].

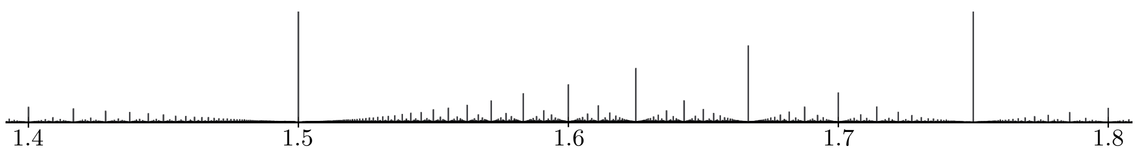

Figure 2.3: A histogram of scl values for many words of length 40 in F_2. Some of the vertical bars are not to scale.

Figure 2.4a shows the scl unit ball in the 2-dimensional subspace of $B_1^H(F_2)$ spanned by the chains $ABabbAbaBB$ and $AABBabbb + Ba$. For example, the starred vertex of the polygon, which is located at $(6/5, 4/5)$, indicates that

$$\mathrm{scl}(6ABabbAbaBB + 4AABBabbb + 4Ba) = 5.$$

Even with these relatively short chains, there is evident complexity in the unit ball. Figure 2.4b shows the unit ball in the 2-dimensional subspace spanned by $baaaaaBBABBBAAbAbbAb$ and $bbAbaaBAAABabaaBBBAb$. The similarity to an L^1 norm is starting to appear in the overall diamond shape, but the small scale complexity is enormous.

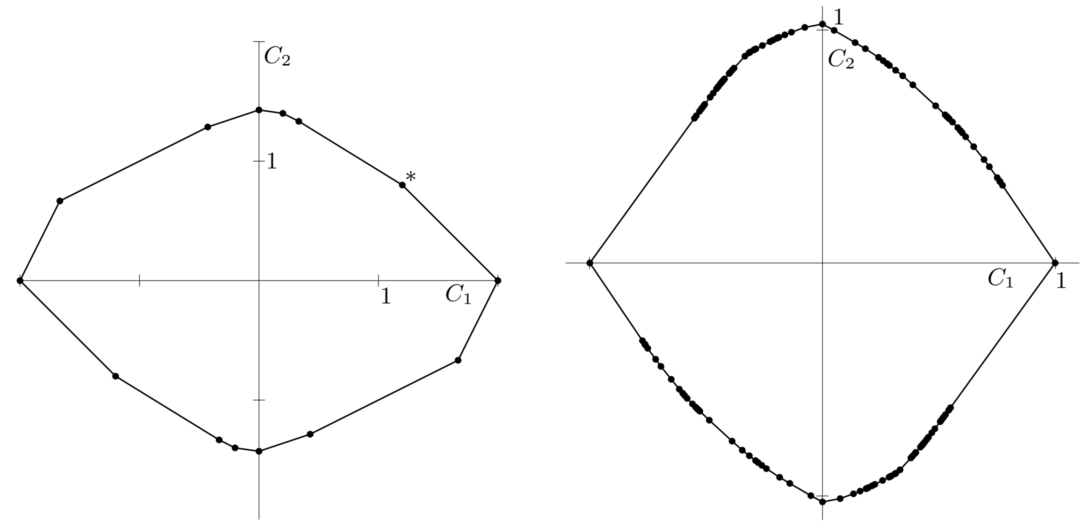

(a) The unit ball in the subspace spanned by the chains $C_1 = ABabbAbaBB$ and $C_2 = AABBabbb+Ba$.

(b) The unit ball in a subspace spanned by two random word of length 20.

Figure 2.4: Unit balls in the scl norm.

2.4 Fatgraphs

2.4.1 Fatgraphs

Definition 2.4.1. A *fatgraph*, also known as a ribbon graph, is a graph with a specified cyclic order on the edges at each vertex.

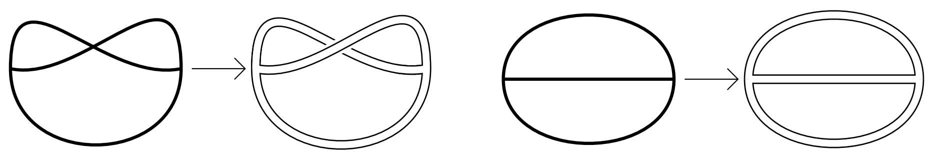

Figure 2.5: Two fatgraphs differing only in their cyclic orders, and their fattenings to a once-punctured torus and a trice-punctured sphere.

They are so called because a fatgraph Y can be "fattened" to a surface, which we will denote by $S(Y)$; the cyclic order at each vertex ensures that the surface obtained is well-defined. See Figure 2.5.

Definition 2.4.2. A fatgraph *over* F_k, or a *labeled* fatgraph, is a fatgraph with its edges labeled on both sides by generators of F_k, subject to the condition that if an edge has the label w on one side, then the other side of the edge is labeled by w^{-1}.

A labeled fatgraph Y determines in a surface map from $S(Y)$ to X_k: map each edge according to the labels. The labeling also shows the (conjugacy class of) the images of the boundary components of $S(Y)$: read off the labels around each boundary component. See Figure 2.6.

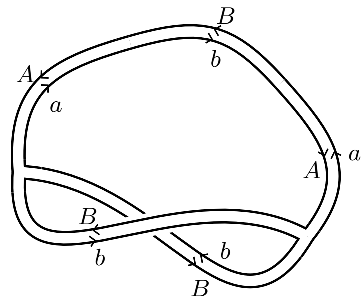

Figure 2.6: A labeled fatgraph; the induced surface map takes the boundary to the conjugacy class of the commutator $[abAB, BaBA] = abABBaBAbb$.

Conversely,

Lemma 2.4.3 ([18], Theorem 1.4)**.** *Let S be a surface with boundary, and let $g : S \to X_k$ be an incompressible map. Then g is carried by a labeled fatgraph map. That is, there exists a labeled fatgraph Y over F_k so that $g = i \circ h$, where $h : S \to S(Y)$ is a homeomorphism, and $i : S(Y) \to X_k$ is the map induced by the labeling of Y.*

Since it is the heart of many of our computational arguments, we record the following corollary of Lemma 2.4.3 as another lemma.

Lemma 2.4.4. *Let $C \in B_1^H(F_k)$ be a chain, and let $f : S \to X_k$ be an admissible map. Then there is a labeled fatgraph Y with induced surface map $i : S(Y) \to X_k$ such that $-\chi(S(Y)) \leq S$ and $i_*(\partial S(Y)) = f_*(\partial S)$.*

In other words, after possible compressions and homotopy, any admissible map is carried by a labeled fatgraph.

2.4.2 Stallings fatgraphs

Recall that a Stallings graph is folded if there is at most one outgoing edge and one incoming edge labeled by a generator at each vertex. We define a labeled fatgraph to be folded if its boundary words are reduced. Note that this is equivalent to saying that for every vertex, there are no cyclically *adjacent* edges with the same direction and label. Here adjacent is well-defined because of the cyclic order.

It is not the case that we may obtain a folded labeled fatgraph from an unfolded one by folding. The process of doing local folds can result in losing the fatgraph structure.

It is also not the case that a folded fatgraph is Stallings folded as an abstract graph, since non-ajacent edges at the same vertex might have the same label.

2.5 Surface realizations

2.5.1 Surface realizations

Definition 2.5.1. A *surface realization* of a free group F_k is a pair (Σ, f), where Σ is a smooth surface with boundary, and f is a homotopy equivalence $f : X_k \to \Sigma$.

We consider the set of surface realization modulo the following equivalence relation. Two surface realizations (Σ_1, f_1) and (Σ_2, f_2) are equivalent if there is a diffeomorphism $\phi : \Sigma_1 \to \Sigma_2$ such that $\phi \circ f_1$ is homotopic to f_2.

Because we have identified $F_k \cong \pi_1(X_k)$, a surface realization induces an isomorphism $f_* : F_k \to \pi_1(\Sigma)$ up to inner automorphisms of F_k.

Definition 2.5.2. A *hyperbolic surface realization* of F_k is a pair (Σ, f), where Σ is a hyperbolic surface with geodesic boundary, and $f : X_k \to \Sigma$ is a homotopy equivalence. Two hyperbolic surface realizations (Σ_1, f_1) and (Σ_2, f_2) are equivalent if there is an isometry $\phi : \Sigma_1 \to \Sigma_2$ such that $\phi \circ f_1$ is homotopic to f_2.

Definition 2.5.3 (Alternate definition)**.** A hyperbolic surface realization of F_k is a discrete faithful representation $\rho : F_k \to \mathrm{PSL}(2, \mathbb{R})$. Two hyperbolic surface realizations are equivalent if the representations are conjugate. In the above notation, we let (Σ_ρ, f_ρ) denote the hyperbolic surface realization given by ρ.

Given a hyperbolic surface realization, we can forget the hyperbolic structure to obtain a surface realization. We can therefore think of the set of surface realizations as π_0 of the space of hyperbolic surface realizations. Conversely, a surface realization can be stiffened to a hyperbolic surface realization by choosing a hyperbolic structure on Σ.

Given a free group F_k, there is a finite set of topological types of surfaces Σ with boundary such that $\pi_1(\Sigma) \cong F_k$. Therefore, the space of hyperbolic surface realizations decomposes as a union of Teichmüller spaces, with one component for each topological type of surface Σ and (because we mod out by the equivalence relation above) coset of the mapping class group of Σ in the outer automorphisms $\mathrm{Out}(F_k)$.

2.5.2 Rotation quasimorphisms from realizations

The action of $\mathrm{PSL}(2, \mathbb{R})$ on $\mathbb{H}^2$ induces an action on the ideal boundary $\partial \mathbb{H}^2 = S^1$ by homeomorphisms, giving rise to an inclusion $\mathrm{PSL}(2, \mathbb{R}) \to \mathrm{Homeo}^+(S^1)$.

Any representation $\rho : F_k \to \mathrm{PSL}(2, \mathbb{R}) \to \mathrm{Homeo}^+(S^1)$ lifts to $\rho^\sim : F_k \to \mathrm{Homeo}^+(S^1)^\sim$ because $H^2(F_k; \mathbb{Z}) = 0$. We denote by rot_ρ the associated class of $\mathrm{rot}^\sim$ in $Q(F_k)/H^1(F_k)$. Note that the rotation quasimorphism $\mathrm{rot}^\sim$ depends on the particular lift of the representation, but the class $\mathrm{rot}_\rho = [\mathrm{rot}^\sim] \in Q(F_k)/H^1(F_k)$ is well defined. Also note $[\mathrm{rot}^\sim] \in Q(F_k)/H^1(F_k; \mathbb{Z})$ is well defined.

If ρ is a discrete, faithful representation associated with a hyperbolic surface realization, the quasimorphism rot_ρ takes integer values. Because rot is a continuous function (on the space $\mathrm{Homeo}^+(S^1)$) and the Teichmüller component containing ρ is connected, it follows that rot_ρ is constant over the entire component. In other words,

Lemma 2.5.4. *Let Σ_{ρ_1} and Σ_{ρ_2} be two hyperbolic surface realizations in the same component. Then* $\mathrm{rot}_{\rho_1} = \mathrm{rot}_{\rho_2}$.

If (Σ, f) is a surface realization of F, we will denote the associated rotation quasimorphism by $\mathrm{rot}_{(\Sigma,f)}$, or just rot_Σ if f is understood.

If (Σ, f) is a hyperbolic surface realization, each Γ in B_1 has a canonical representative $\Gamma : \coprod_i S^1_i \to \Sigma$ (at least up to reparameterization), namely the geodesic representative. We say that Γ *virtually bounds an immersed surface* if there is some immersion $g : S \to \Sigma$ which is admissible for a geodesic representative Γ. This is equivalent to saying that we can find a hyperbolic structure on S with geodesic boundary for which $g : S \to \Sigma$ is a local isometry.

The following theorem indicates a deep relationship between the geometry and combinatorics of $B_1^H(F_k)$ and the set of surface realizations of F_k.

Theorem 2.5.5 (Calegari [7], Theorems A and B)**.** *Let (Σ, f) be a surface realization of F_k, and let $\partial\Sigma \in B_1^H(F_k)$ be the associated chain.*

1. *There is a codimension 1 face π_Σ of the unit ball of the* scl *norm, and the chain $\partial\Sigma$ projectively intersects its interior.*

2. *A chain Γ projectively intersects π_Σ if and only if, for any hyperbolic surface realization stiffening (Σ, f), the geodesic representative of Γ virtually bounds an immersed surface. Additionally, a surface admissible for Γ is extremal for Γ if and only if it is homotopic to an immersion with geodesic boundary.*

3. rot_Σ *is the unique (class of) quasimorphism in $Q(F_k)/H^1(F_k)$ with defect 1 which is extremal for elements projectively intersecting the interior of π_Σ.*

Theorem 2.5.5 implies the following ridigity result for the interaction of extremal surfaces and extremal quasimorphisms, valid for any group:

Corollary 2.5.6. *Let G be any group. Let $C \in B_1^H(G)$, and let $f : S \to K(G,1)$ be an extremal surface for C. Then any quasimorphism $\phi \in Q(G)$ with defect 1 which is extremal for C restricts to the tautological representation* rot_S *on $\pi_1(S)$, up to an element of $H^1(S;\mathbb{R})$.*

Proof. We have $f : S \to X$ admissible for C After replacing (S, f) by a finite cover if necessary, we may assume that $n(S) = 1$. Hence f takes ∂S to C, and therefore $f^*\phi$ is an element of $Q(\pi_1(S))$

extremal for ∂S. Note that $D(f^*\phi) = 1$, certified by its value on ∂S. By Theorem 2.5.5, this implies that $f^*\phi = \mathrm{rot}_S$ up to an element of $H^1(S;\mathbb{R})$. $\square$

In the preceeding corollary, it is important to note that the pullback of any extremal quasimorphism to any extremal surface gives the rotation number up to an element of $H^1(S;\mathbb{R})$, *not* $H^1(S;\mathbb{Z})$ or $H^1(G;\mathbb{R})$. This is a source of potential confusion.

A surface realization induces the rotation quasimorphism rot_Σ, and Theorem 2.5.5 shows that there is a strong relationship between geometry and the scl norm ball. The theorem associates a codimension one face of the scl norm ball with a surface realization.

Definition 2.5.7. We will call these codimension one faces *geometric faces* of the scl norm ball.

2.5.3 Cyclic orders and basic surface realizations

Definition 2.5.8. A *cyclic order* on a set S is a choice of total ordering $<_p$ on $S \setminus p$ for each $p \in S$, subject to the condition that if $p, q \in S$, then the total orders $<_p$ and $<_q$ differ by a cut. That is, if $x <_p y$, then $x <_q y$ unless $x <_p q <_p y$, in which case $y <_q x$.

The intuition for a cyclic order is, naturally, that a cyclic order on a set S arranges the elements of S cyclically, so while we cannot say if $x < y$, we can say whether any triple (x, y, z) is clockwise or counterclockwise. We will call a triple which is a counterclockwise *positive*. It is sometimes useful to think of a cyclic order as a map $O : S^3 \to \{-1, 0, +1\}$ which assigns ± 1 or 0 to a triple depending on whether it is counterclockwise, clockwise, or degenerate.

We will denote explicit cyclic orders by square brackets, so $[x, y, z, w]$ gives a cyclic order on the set $\{x, y, z, w\}$. It is the same cyclic order as, for example, $[y, z, w, x]$, but not the same as $[x, y, w, z]$.

Cyclic orders and fatgraphs give a simple combinatorial way to describe surface realizations. Let $S = \{a_1, a_1^{-1}, \ldots, a_k, a_k^{-1}\}$ be a set of semigroup generators for F_k. Given a cyclic order O on S, build a labeled fatgraph over F_k as follows. First create a single vertex v. Then attach $2k$ half-edges to v in such a way that the outgoing labels on each half-edge are arranged according to the cyclic order O on S. Then glue the unattached ends of the half edges together, producing a fatgraph with a single vertex and k edges. Call this labeled fatgraph Σ_O. There is a tautological map $f_O : X_k \to \Sigma_O$ obtained from following the labels.

Definition 2.5.9. The surface realization (f_O, Σ_O) as constructed above is the *basic fatgraph realization* induced by O.

See Figure 2.7 for an example.

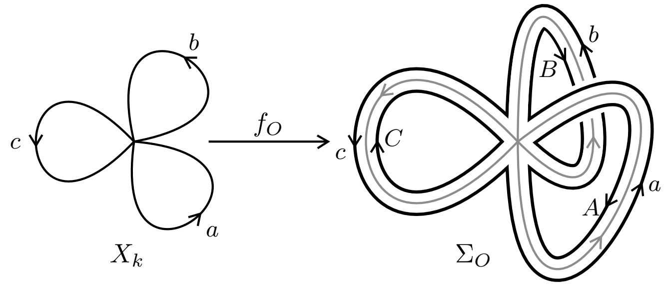

Figure 2.7: The basic fatgraph realization of $F_3 = \langle a, b, c\rangle$ induced by the cyclic order $[a, b, A, B, c, C]$. The surface Σ_O has two boundary components, given by the conjugacy classes of the words C and $aBAbc$.

2.5.4 A formula for rotation number

For a basic fatgraph realization, there is a simple combinatorial formula for the associated rotation number rot_{Σ_O}, which for simplicity we will denote by rot_O. The formula is given by Theorem 4.76 in [8].

If S is a set with a cyclic order O, and $x, y, z \in S$, we define the *distance* between x and y avoiding z to be the number $d_z(x, y)$ such that x and y have $d_z(x, y) - 1$ elements between them when arranged according to the total order $<_z$ on S. Then the *translation* from x to y avoiding z is

$$t_z(x,y) = \begin{cases} d_z(x,y) & \text{if } x <_z y \\ -d_z(x,y) & \text{if } y <_z x \end{cases} .$$

Theorem 2.5.10 (Theorem 4.76 in [8])**.**

$$\mathrm{rot}_O = \frac{1}{|S|} \sum_{x \neq y^{\pm 1}} t_{y^{-1}}(x,y) C_{yx}$$

While the formula may seem slightly arbitrary, this theorem is completely intuitive given the discussion in [8], §4.2.

Remark 2.5.11 (Important technical note). Theorem 2.5.10 gives the rotation quasimorphism induced by a cyclic order as a sum of counting functions. There is a difference in convention here which results in possible confusion: counting functions read subwords of the input word from left to right, but the circle action at infinity by F_k coming from the $\mathrm{PSL}(2, \mathbb{R})$ representation is a *left* action. In order to correctly compute the rotation quasimorphism, therefore, we find the counting function C_{yx}, rather than C_{xy}. The other option is to read the input word backwards.

Chapter 3

Isometric endomorphisms of free groups

3.1 Introduction

Definition 3.1.1. Given a homomorphism $\phi : G \to H$ between groups, we say that ϕ is an *isometry* for scl if $\mathrm{scl}_H(\phi(C)) = \mathrm{scl}_G(C)$ for all $C \in B_1^H(G)$.

It is natural to ask what conditions are sufficient to show that a homomorphism is an isometry. In this chapter, we show that a random homomorphism between free groups is an isometry, and we show that a random map of a fatgraph into a free group is extremal for its boundary. This chapter contains essentially the same material as [17].

More specifically, we prove

Theorem (3.3.11; Random isometry theorem)**.** *A random homomorphism $\varphi : F_k \to F_l$ of length n between free groups of ranks k, l is an isometry of* scl *with probability $1 - O(C(k,l)^{-n})$ for some constant $C(k,l) > 1$.*

By this theorem, the scl unit ball of a given free group is the same as infinitely many different subballs of itself.

We also prove

Theorem (3.4.11; Random fatgraph theorem)**.** *For any combinatorial fatgraph $\hat{Y}$, if Y is a random fatgraph over F_k obtained by labeling the edges of $\hat{Y}$ by words of length n, then $S(Y)$ is extremal for $\partial S(Y)$ and is certified by the extremal quasimorphism H_Y, with probability $1 - O(C(\hat{Y}, k)^{-n})$ for some constant $C(\hat{Y}, k) > 1$.*

Here the random labeling of a combinatorial fatgraph gives a "random" surface map, so this theorem says that a random surface is extremal.

3.2 Isometries and non-isometries of scl

Let $\phi : G \to H$ be a homomorphism between groups. Then since the image of a commutator is a commutator, $\mathrm{scl}_H(\phi(C)) \leq \mathrm{scl}_G(C)$ for any $C \in B_1^H(G)$. This inequality also holds for the commutator length of any element $g \in [G, G]$: we have $\mathrm{cl}(\phi(g)) \leq \mathrm{cl}(g)$. A natural question is whether a homomorphism is an *isometry* for scl; that is, whether $\mathrm{scl}_G(\phi(C)) = \mathrm{scl}_F(C)$ for all $C \in B_1^H(G)$. Analogously, we might ask if a homomorphism preserves commutator length. The latter implies the former, but not the other way around.

Example 3.2.1. An automorphism is an isometry.

Let us now specialize to free groups. A necessary condition for a homomorphism to be an isometry is that it is injective. However, this is not sufficient, as the following examples show.

Example 3.2.2. Let H be a genus 3 handlebody with boundary $S = \partial S$. Decompose S into two surfaces S_1 and S_2 of genus 1 and 2, respectively, separated by a curve $\gamma = \partial S_1 = \partial S_2$. Denote the inclusion $i : S_2 \to S$.

Let T be the subcomplex of the complex of curves in S containing all the curves which bound disks in H. By [24], Theorem 2.7, there is a pseudo-Anosov homeomorphism ϕ of the boundary ∂S such that $d(\phi(\gamma), T) \geq 2$. Since $\phi(\gamma)$ is disjoint (i.e., distance 1 in the curve complex) from all the curves in $\phi(S_2)$ representing elements of $\pi_1(S_2)$, this implies that $(\phi \circ i)_* : \pi_1(S_2) \to \pi_1(H)$ is injective. Now $\mathrm{scl}_{\pi_1(S_2)}(\partial S_2) = 1$, since a surface is extremal in its own surface group for its boundary component. However, by construction, $\partial(\phi(S_2)) = \phi(\gamma) = \partial(\phi(S_1))$, so $\phi(\gamma)$ bounds a torus, so $\mathrm{scl}_{\pi_1(H)}(\partial\phi(S_2)) \leq 1/2$. Therefore, while $(\phi \circ i)_*$ is injective, it is not an scl-isometry.

In fact, this construction gives an example of an injective map of free groups that does not preserve commutator length. This gives a negative answer to a question of Bardakov [1], who asked whether an injection of free groups must preserve commutator length.

Example 3.2.3 (Nongeometric covers). Let $H \subseteq F_k$ be a subgroup of finite index. If (Σ, f) is a surface realization of F_k, then there is a cover $p : (\Sigma', f') \to (\Sigma, f')$ so that (Σ', f') is a realization of H. We call this a *geometric* cover. However, it is not the case that every surface realization of H arises in this way. If (Σ, f) is a realization of H which does not cover a realization of F_k, then we say that this is a *nongeometric* cover. The inclusion map $p : H \to F_k$ is injective, but we claim that $\mathrm{scl}_{F_k}(\partial\Sigma) < \mathrm{scl}_H(\partial\Sigma)$ for any nongeometric cover Σ realizing H. Thus the inclusion map of any finite index subgroup is not an scl-isometry.

The proof of the claim relies on background developed more fully in [8], §4.2, but the essential ingredient is Theorem 2.5.5, and we give a sketch. If Σ is a nongeometric cover, then the conjugation action of F_k/H on H contains elements of $\mathrm{Out}(H)$ which are not in the mapping class group of Σ. Let ϕ be the action of such an element of F_k/H. Applying Theorem 2.5.5, we see that $\partial\Sigma$ lies in the interior of a codimension one face of the scl norma ball of H, and $\phi\partial\Sigma$ cannot lie in the same face (projectively). Therefore,

$$\mathrm{scl}_H(\partial\Sigma + \phi\partial\Sigma) < \mathrm{scl}_H(\partial\Sigma) + \mathrm{scl}_F(\phi\partial\Sigma).$$

Now we can compute $\mathrm{scl}_{F_k}(\partial\Sigma)$ using the formula from Corollary 2.81 in [8]. This formula, together with the above strict inequality, implies that $\mathrm{scl}_{F_k}(\partial\Sigma) < \mathrm{scl}_H(\partial\Sigma)$.

This holds for any nongeometric cover, so for any finite index subgroup, there are infinitely many elements of $B_1^H(H)$ whose scl strictly decreases under the (injective) inclusion map.

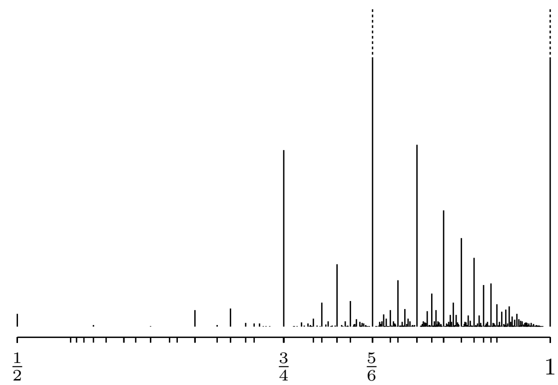

Figure 3.1: A histogram of $\mathrm{scl}_{F_2}(i_*\partial\Sigma)$ for 7500 random surface realizations of an index 2 subgroup.

Figure 3.1 shows a histogram of $\mathrm{scl}(i_*\partial\Sigma)$ for many realizations Σ of an index 2 subgroup $G \subseteq F_2$, where Σ is a four-punctured sphere. For all these realizations, we have $\mathrm{scl}_G(\partial\Sigma) = 1$, but under the inclusion map $i : G \to F_2$ the scl must strictly decrease for nongeometric covers. The large spike at 1 is all the geometric covers.

3.3 Random homomorphisms are usually isometries

3.3.1 Bounds on scl

Let F_k be freely generated by $x_1, \ldots, x_k$; if $\Gamma \in B_1^H(F_k)$, we denote by $|\Gamma|_i$ the number of times that x_i and x_i^{-1} appear in Γ.

Proposition 3.3.1. *With notation as above, we have an inequality*

$$\mathrm{scl}(\Gamma) \leq \frac{|\Gamma| - \max_i |\Gamma|_i}{4}$$

for any $\Gamma \in B_1^H(F_k)$.

Proof. Without loss of generality, we may assume that $\max_i |\Gamma|_i = |\Gamma|_n$. Given an extremal surface for Γ, we may assume it has a fatgraph structure by Lemma 2.4.4. Now cut out all rectangles corresponding to matched pairs of x_1 and x_1^{-1}. Note that this actually means cutting out at most $|\Gamma|_1/2$ thrice-punctured spheres, What is left is an immersed subsurface S' of S, and because we cut out thrice-punctured spheres, $-\chi(S) \leq -\chi(S') + |\Gamma|_i/2$. An essential immersed subsurface of an extremal surface is also extremal, by [7] (essentially, Theorem 2.5.5). Consequently S' is extremal for its boundary Γ', which lies in $B_1^H(\langle x_2, \ldots, x_n\rangle)$. We therefore have the inequality $\mathrm{scl}(\Gamma) \leq \mathrm{scl}(\Gamma') + |\Gamma|_1/4$. Repeating this argument $n-1$ times yields

$$\mathrm{scl}(\Gamma) \leq \mathrm{scl}(\Gamma'') + |\Gamma|_1/4 + \cdots + |\Gamma|_{n-1}/4$$

where $\mathrm{scl}(\Gamma'') = 0$, since $\Gamma'' \in B_1^H(\langle x_n\rangle)$. The proof follows. □

Example 3.3.2. The bound in Proposition 3.3.1 is sharp, which we show by a family of examples. We first recall the free product formula ([8], §2.7), which says that if G_1 and G_2 are arbitrary groups, and $g_i \in G_i'$ have infinite order, then $\mathrm{scl}_{G_1 * G_2}(g_1 g_2) = \mathrm{scl}_{G_1}(g_1) + \mathrm{scl}_{G_2}(g_2) + 1/2$.

Now, let F_k be freely generated by $x_1, \ldots, x_k$ as above, and define

$$w_k = [x_1, x_2][x_3, x_4] \cdots [x_{k-1}, x_k]$$

if k is even, and

$$w_k = [x_1, x_2][x_3, x_4] \cdots [x_{k-4}, x_{k-3}] x_{k-2} x_{k-1} x_k x_{k-1}^{-1} x_k x_{k-2}^{-1} x_k^{-2}$$

if k is odd.

For each i, we have $\mathrm{scl}([x_i, x_{i+1}]) = 1/2$. Moreover, using `scallop` ([15]) one can check that

$$\mathrm{scl}(x_{k-2} x_{k-1} x_k x_{k-1}^{-1} x_k x_{k-2}^{-1} x_k^{-2}) = 1.$$

The free product formula then shows that $\mathrm{scl}(w_k) = (k-1)/2$, so Proposition 3.3.1 is sharp for all k.

3.3.2 A small cancellation condition

Definition 3.3.3. Let A be a set, and let $F(A)$ be the free group on A. Let U be a subset of $F(A)$ with $U \cap U^{-1} = \emptyset$, and let S denote the set $U \cup U^{-1}$. We say that U satisfies condition (SA) if the following is true:

(SA1) if $x, y \in S$ and y is not equal to x^{-1}, then xy is reduced; and

(SA2) if $x, y \in S$ and y is not equal to x or x^{-1}, then any common subword s of x and y has length strictly less than $|x|/12$; and

(SA3) if $x \in S$ and a subword s appears in at least two different positions in x (possibly overlapping) then the length of s is strictly less than $|x|/12$.

Let B be a set, and $\varphi : B \to U$ a bijection. Extend φ to a homomorphism $\varphi : F(B) \to F(A)$. We say φ satisfies condition (SA) if U satisfies condition (SA).

Note that except for condition (SA1), this is the small cancellation condition $C'(1/12)$. We will show the following:

Proposition 3.3.4. *Let $\varphi : F(B) \to F(A)$ be a homomorphism satisfying condition (SA). Then φ is an isometry of* scl*.*

Condition (SA1) for φ means that if g is a cyclically reduced word in $F(B)$, then the word in $F(A)$ obtained by replacing each letter of g by its image under φ is also cyclically reduced. This condition is quite restrictive — in particular it implies that $|A| \geq |B|$, and even under these conditions it is not "generic" — but we will show how to dispense with it in §3.3.3. However, its inclusion simplifies the arguments in this section.

Example 3.3.5. The set $\{aa, bb\}$ satisfies (SA1). The set $\{ab, ba\}$ satisfies (SA1).

Suppose $\varphi : F(B) \to F(A)$ satisfies condition (SA), and let Y be a fatgraph with $\partial S(Y)$ in the image of φ, i.e., such that $\partial S(Y)$ is a collection of cyclically reduced words of the form $\varphi(g)$. By condition (SA1), each $\varphi(g)$ is obtained by concatenating words of the form $\varphi(x^{\pm})$ for $x \in B$. We call these subwords *segments* of $\partial S(Y)$, as distinct from the decomposition into *arcs* associated with the fatgraph structure. An arc is a pair of subwords of $\partial S(Y)$ which are glued together in the fatgraph structure.

Definition 3.3.6. A *perfect match* in Y is a pair of segments $\varphi(x), \varphi(x^{-1})$ contained in a pair of arcs of $\partial S(Y)$ that are matched by the pairing. A *partial match* in Y is a pair of segments $\varphi(x), \varphi(x^{-1})$ containing subsegments s, s^{-1} in "corresponding" locations in $\varphi(x)$ and $\varphi(x^{-1})$ that are matched by the pairing.

The existence of a perfect match will let us replace Y with a "simpler" fatgraph. This is the key to an inductive proof of Proposition 3.3.4. The next lemma shows how to modify a fatgraph Y to promote a partial match to a perfect match.

Lemma 3.3.7. *Suppose Y contains a partial match. Then there is Y' containing a perfect match with $S(Y')$ homotopic to $S(Y)$ and $\partial S(Y) = \partial S(Y')$.*

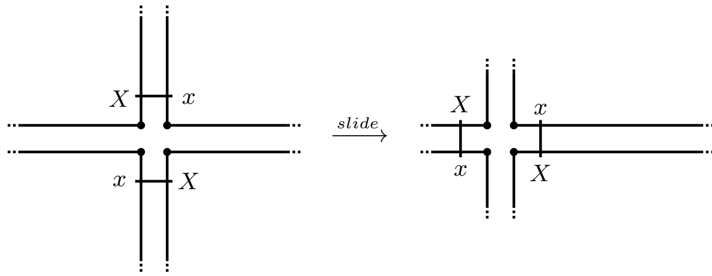

Figure 3.2: A local move to replace a partial match with a perfect match.

Proof. The fatgraph Y can be modified by a certain local move, illustrated in Figure 3.2.

This move increases the length of the paired subsegments by 1. Perform the move repeatedly to obtain a perfect match. $\square$

Each vertex v of Y of valence $|v|$ contributes $(|v|-2)/2$ to $-\chi(Y)$, in the sense that $-\chi(Y) = \sum_v (|v|-2)/2$. Since each vertex v of Y is in the image of $|v|$ vertices in ∂S, we assign a weight of $(|v|-2)/2|v|$ to each vertex of ∂S.

Lemma 3.3.8. *Let Y be a fatgraph with $\partial S(Y) = \varphi(\Gamma)$ and suppose that φ satisfies (SA). Then either Y contains a partial match, or $-\chi(Y) > |\Gamma|$.*

Proof. Observe that $\partial S(Y)$ decomposes into $|\Gamma|$ segments, corresponding to the letters of Γ. Suppose Y does not contain a partial match. Then since each vertex contributes $(|v|-2)/2|v|$ to $-\chi(Y)$, it suffices to show that each segment of ∂Y contains at least six vertices in its interior.

Suppose not. Then some segment $\varphi(x)$ of ∂Y contains a subsegment s of length at least $|\varphi(x)|/6$ that does not contain a vertex in its interior. Either s contains a possibly smaller subsegment s' which is paired with some entire segment $\varphi(y)$, or at least half of s is paired with some s^{-1} in some $\varphi(y)$. In either case, since s is not a partial match by hypothesis, we contradict either (SA2) or (SA3).

Thus each segment contributes at least $7 \times ((3-2)/2 \cdot 3) = 7/6$ to $-\chi(Y)$, and the lemma is proved. $\square$

We now give the proof of Proposition 3.3.4.

Proof of Proposition 3.3.4. Suppose $\varphi : F(B) \to F(A)$ satisfies (SA) but is not isometric.

Let Y be a fatgraph with $\partial S(Y) = \varphi(\Gamma)$ so that $\mathrm{scl}(\varphi(\Gamma)) \le -\chi(S(Y))/2 < \mathrm{scl}(\Gamma)$ (the existence of such a Y follows from Lemma 2.4.3; for instance, we could take Y to be extremal). We will construct a new Y' with $\partial S(Y') = \varphi(\Gamma')$ satisfying $\mathrm{scl}(\varphi(\Gamma')) \le -\chi(S(Y'))/2 < \mathrm{scl}(\Gamma')$, and such that Y' is shorter than Y. By induction on the size of Y we will obtain a contradiction.

By Proposition 3.3.1 and Lemma 3.3.8, Y contains a partial match, and by Lemma 3.3.7 we can modify Y without affecting $\partial S(Y)$ or $\chi(Y)$ so that it contains a perfect match. A perfect match cobounds a rectangle in $S = S(Y)$ that can be cut out, replacing S with a “simpler” surface S' for which $\partial S'$ is also in the image of φ. By Lemma 2.4.3, there is some surface S'' with $-\chi(S'') \leq -\chi(S')$ and $\partial S'' = \partial S'$, anda fatgraph Y' with $S(Y') = S''$.

In the degenerate case that S'' is a disk, necessarily S is an annulus, and both boundary components of S consist entirely of perfect matches; hence $\Gamma = g + g^{-1}$ and $\mathrm{scl}(\Gamma) = \mathrm{scl}(\varphi(\Gamma)) = 0$ in this case, contrary to hypothesis. Otherwise $\partial S'' = \partial S' = \varphi(\Gamma')$ for some Γ', and satisfies $-\chi(S(Y')) \leq -\chi(S') = -\chi(S(Y)) - 1$.

On the other hand, Γ can be obtained from Γ' by gluing on a pair of pants; hence $\mathrm{scl}(\Gamma) \leq \mathrm{scl}(\Gamma') + 1/2$. We have the “diagram of inequalities” shown in Figure 3.3, from which we deduce

$$\begin{array}{ccccc}
\mathrm{scl}(\varphi(\Gamma)) & \leq & -\chi(S(Y))/2 & < & \mathrm{scl}(\Gamma) \\
 & & \rotatebox{-90}{$\leq$} & & \rotatebox{90}{$\leq$} \\
\mathrm{scl}(\varphi(\Gamma')) + 1/2 & \leq & -\chi(S(Y'))/2 + 1/2 & & \mathrm{scl}(\Gamma') + 1/2
\end{array}$$

Figure 3.3: Diagram of inequalities for the proof of Proposition 3.3.4

$\mathrm{scl}(\varphi(\Gamma')) \leq -\chi(S(Y'))/2 < \mathrm{scl}(\Gamma')$ as claimed. Since each reduction step reduces the length of $\partial S(Y)$, we obtain a contradiction. $\square$

3.3.3 A generic small cancellation condition

In this section, we weaken condition (SA) by allowing partial cancellation of adjacent words $\varphi(x)$ and $\varphi(y)$. Providing we quantify and control the amount of this cancellation, we obtain a new condition (A) (defined below) which holds with high probability, and which implies isometry.

If two successive letters x, y in a fatgraph do not cancel, but some suffix of $\varphi(x)$ cancels some prefix of $\varphi(y)$, we encode this pictorially by adding a *tag* to our fatgraph. A tag is an edge, one vertex of which is 1-valent. The two sides of the tag are then labeled by the maximal canceling segments in $\varphi(x)$ and $\varphi(y)$. If Γ is a chain, and Y is a fatgraph with $\partial S(Y)$ equal to the cyclically reduced representative of $\varphi(\Gamma)$, then we can add tags to Y to produce a fatgraph Y' so that $\partial S(Y')$ is equal to the(possibly unreduced) chain $\varphi(\Gamma)$.

Definition 3.3.9. Let A be a set, and let $F(A)$ be the free group on A. Let U be a subset of $F(A)$ with $U \cap U^{-1} = \emptyset$, and let Σ denote the set $U \cup U^{-1}$. We say that U satisfies condition (A) if there is some non-negative real number T such that the following is true:

(A1) the maximal length of a tag is T; and

(A2) if $x, y \in \Sigma$ and y is not equal to x or x^{-1}, then any common subword s of x and y has length strictly less than $(|x| - 2T)/12$; and

(A3) if $x \in \Sigma$ and a subword s appears in at least two different positions in x (possibly overlapping) then the length of s is strictly less than $(|x| - 2T)/12$.

Let B be a set, and $\varphi : B \to U$ a bijection. Extend φ to a homomorphism $\varphi : F(B) \to F(A)$. We say φ satisfies condition (A) if U satisfies condition (A).

Notice that condition (SA) is a special case of condition (A) when $T = 0$.

Proposition 3.3.10. *Let $\varphi : F(B) \to F(A)$ be an homomorphism between free groups satisfying condition (A). Then φ is an isometry of* scl*. That is,* $\mathrm{scl}(\Gamma) = \mathrm{scl}(\varphi(\Gamma))$ *for all chains $\Gamma \in B_1^H(F(B))$. In particular,* $\mathrm{scl}(g) = \mathrm{scl}(\varphi(g))$ *for all $g \in [F(B), F(B)]$.*

Proof. The proof is essentially the same as that of Proposition 3.3.4, except that we need to be slightly more careful computing $\chi(Y)$. We call the edges in a tag *ghost edges*, and define the valence of a vertex v to be the number of non-ghost edges incident to it. Then $-\chi(Y) = \sum_v (|v| - 2)/2$ where the sum is taken over all "interior" vertices v — i.e., those which are not the endpoint of a tag.

The proof of Lemma 3.3.8 goes through exactly as before, showing that either Y contains a partial match, or $-\chi(Y) > |\Gamma|$. To see this, simply repeat the proof of Lemma 3.3.8 applied to Y with the tags "cut off". Partial matches can be improved to perfect matches as in Lemma 3.3.7. Note that this move might unfold a tag.

If Y is a fatgraph with $\partial S(Y) = \varphi(\Gamma)$ and $\mathrm{scl}(\varphi(\Gamma)) \leq -\chi(S(Y))/2 < \mathrm{scl}(\Gamma)$, we can find a perfect match and cut out a rectangle, and the induction argument proceeds exactly as in the proof of Proposition 3.3.4. ∎

3.3.4 Random homomorphisms are usually isometries

Fix k, l integers ≥ 2. We now explain the sense in which a random homomorphism from F_k to F_ℓ will satisfy condition (A). Fix an integer n, and let $F_\ell(\leq n)$ denote the set of reduced words in F_ℓ (in a fixed free generating set) of length at most n. Define a *random homomorphism of length $\leq n$* to be the homomorphism $\varphi : F_k \to F_l$ sending a (fixed) free generating set for F_k to k randomly chosen elements of $F_\ell(\leq n)$ (with the uniform distribution).

Theorem 3.3.11 (Random Isometry Theorem)**.** *A random homomorphism $\varphi : F_k \to F_l$ of length n between free groups of ranks k, l is an isometry of* scl *with probability $1 - O(C(k,l)^{-n})$ for some constant $C(k,l) > 1$.*

Proof. By Proposition 3.3.10 it suffices to show that a random homomorphism satisfies condition (A) with sufficiently high probability.

Let $u_1, \ldots, u_k$ be the images of a fixed free generating set for F_k, thought of as random reduced words of length $\leq n$ in a fixed free generating set and their inverses for F_l. First of all, for any $\epsilon > 0$, we can assume with probability at least $1 - O(C^{-n})$ for some C that the length of every u_i is between n and $(1-\epsilon)n$. Secondly, the number of reduced words of length ϵn is (approximately)

$(2l-1)^{\epsilon n}$, so the chance that the maximal length of a tag is more than ϵn is at least $1 - O(C^{-n})$. So we restrict attention to the φ for which both of these condition hold.

If (A2) fails, there are indices i and j and a subword s of u_i of length at least $n(1-3\epsilon)/12 \geq n/13$ (for large n) so that either s or s^{-1} is a subword of u_j. The copies of $s^{\pm}$ are located at one of at most n different places in u_i and in u_j; the chance of such a match at one specific location is approximately $(2l-1)^{-n/13}$, so the chance that (A2) fails is at most $k^2n^2(2l-1)^{-n/13} = O(C^{-n})$ for suitable C.

Finally, if (A3) fails, there is an index i and a subword s of u_i of length at least $n/13$ that appears in at least two different locations. It is possible that s overlaps itself, but in any case there is a subword of length at least $|s|/3$ that is disjoint from some translate. If we examine two specific disjoint subsegments of length $n/39$, the chance that they match is approximately $(2l-1)^{-n/39}$. Hence the chance that (A3) fails is at most $kn^2(2l-1)^{-n/39} = O(C^{-n})$ for suitable C. Evidently C depends only on k and l. The lemma follows. □

Corollary 3.3.12. *Let $k, l \geq 2$ be integers. There are (many) isometric homomorphisms $\varphi : F_k \to F_l$.*

Lemma 3.3.13. *Let F be a finitely generated free group. The following holds:*

1. *if there are integral chains Γ_1, Γ_2 in $B_1^H(F)$ such that $\mathrm{scl}(\Gamma_i) = t_i$, then there is an integral chain Γ in $B_1^H(F)$ with $\mathrm{scl}(\Gamma) = t_1 + t_2$; and*

2. *if there are elements g_1, g_2 in F' such that $\mathrm{scl}(g_i) = t_i$, then there is an element $g \in F'$ with $\mathrm{scl}(g) = t_1 + t_2 + 1/2$.*

Proof. Let F_1, F_2 be copies of F, and let $\sigma_i : F \to F_i$ be an isomorphism. Then in case (1) the chain $\sigma_1(\Gamma_1) + \sigma_2(\Gamma_2)$ in $F_1 * F_2$ has scl equal to $t_1 + t_2$, and in case (2) the element $\sigma_1(g_1)\sigma_2(g_2)$ has scl equal to $t_1 + t_2 + 1/2$; see [8], §2.7. Now choose an isometric homomorphism from $F_1 * F_2$ to F, which exists by Corollary 3.3.12. □

Corollary 3.3.14. *Let F be a countable nonabelian free group. The image of F' under* scl *contains elements congruent to every element of $\mathbb{Q}$ mod $\mathbb{Z}$. Moreover, the image of F' under* scl *contains a well-ordered sequence of values with ordinal type ω^ω.*

Proof. These facts follow from Lemma 3.3.13 plus the Denominator Theorem and Limit Theorem from [9]. □

3.3.5 Isometry conjecture

Conjecture 3.3.15 (Isometry conjecture)**.** *Let $\varphi : F_2 \to F_k$ be any injective homomorphism from a free group of rank 2 to a free group F_k. Then φ is isometric.*

Remark 3.3.16. Since free groups are Hopfian by Malcev [28], any homomorphism from F_2 to a free group F_k is either injective, or factors through a cyclic group. Furthermore, since F_2 is not proper of finite index in any other free group, no counterexample to the conjecture can be constructed by the method of Example 3.2.3.

Since any free group admits an injective homomorphism into F_2, and since scl is monotone nonincreasing under any homomorphism between groups, to prove Conjecture 3.3.15 it suffices to prove it for endomorphisms $\varphi : F_2 \to F_2$.

Remark 3.3.17. In view of Example 3.2.3, rank 2 cannot be replaced with rank 3 in Conjecture 3.3.15.

Conjecture 3.3.15 has been tested experimentally on all cyclically reduced homologically trivial words of length 11 in F_2, and all endomorphisms $F_2 \to F_2$ sending $a \to a$ and b to a word of length 4 or 5. It has also been tested on thousands of "random" longer words and homomorphisms. The experiments were carried out with the program `scallop` ([15]).

3.4 Random fatgraph labelings are usually extremal

3.4.1 Labeling fatgraphs

We have been (and are) typically interested in *labeled* fatgraphs. However, in this section, we need the distinction between a fatgraph and a labeling of it to be very explicit. We use the notation $\hat{Y}$ for an abstract (unlabeled) fatgraph, and Y for a labeling of $\hat{Y}$ by words in F; i.e., a fatgraph over F. We recall here Definitions 2.4.1 and 2.4.2. A labeling *of length* n is a reduced labeling for which every edge of Y is a word of length n.

By our convention, boundary words in $\partial S(Y)$ must be cyclically reduced. For a labeling in which boundary words are not reduced, one can perform Stallings fatgraph folding as in §2.4.2.

3.4.2 The vertex quasimorphism construction

In this section, we construct a (counting) quasimorphism on F from a fatgraph Y over F. We will call this the vertex quasimorphism of Y. We will see that this vertex quasimorphism is typically extremal for $\partial S(Y)$.

Remark 3.4.1. This remark is actually a warning. In our notation, we take big and small counting quasimorphisms to be already homogenized. This differs from [17], on which this chapter is based.

Define a set σ_Y on a labeled fatgraph Y over F as follows: every boundary component of $S(Y)$ decomposes into a union of arcs, and each arc is labeled by an element of F. Between each pair of arcs is a vertex of $\partial S(Y)$ (associated to a vertex of Y). For each vertex of Y and each pair of incident arcs with labels u and v (u comes into the vertex; v leaves it), decompose u and v into $u = u_1u_2$, $v = v_1v_2$, where usually we expect u_1 and u_2 to each be approximately half the length of

u, and similarly for v_1, v_2, v, and add the word u_2v_1 to the set σ_Y. There is some flexibility here in the phrase "about half the length" which will not affect our later arguments; in fact this flexibility indicates possible other constructions, in which the pieces have different sizes, bounded length, etc.

Definition 3.4.2. Recall the definition of small counting quasimorphisms (Defintion 2.2.5). A *vertex quasimorphism* for Y is a small counting quasimorphism of the form h_{σ_Y}. See Figure 3.4 for an example. In this figure, σ_Y is the set

$$\sigma_Y = \{bbAb,\ aBAA,\ aaaa,\ AbAA,\ AbaB,\ BBaB\}.$$

Note that we have not broken the edges exactly in half, or even in the same place on either side.

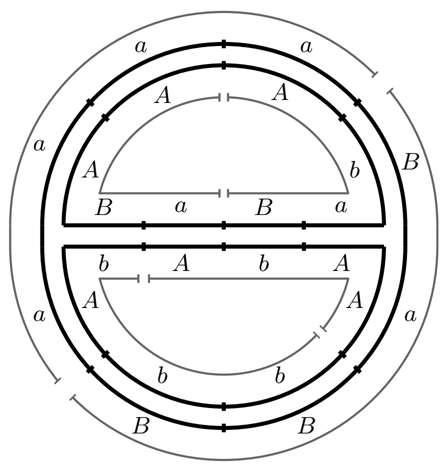

Figure 3.4: The vertex quasimorphism construction on a thrice-punctured sphere.

Lemma 3.4.3. *If no element of σ_Y^{-1} appears in the boundary $\partial S(Y)$, then there is an inequality $h_{\sigma_Y}(\partial S(Y)) \geq \sum_v |v|$, where the sum is taken over all vertices v, and $|v|$ is the valence of the vertex v.*

Proof. Since no element of σ_Y^{-1} appears in $\partial S(Y)$, we have $c_{\sigma_Y^{-1}}(\partial S(Y)) = 0$, so $h_{\sigma_Y}(\partial S(Y)) = c_{\sigma_Y}(\partial S(Y))$. For every vertex of Y and for each incident edge, we have a word in σ_Y. By construction, these words do not overlap in the boundary chain $\partial S(Y)$, so the value of $c_{\sigma_Y}(\partial S(Y))$ is at least as big as $\sum_v |v|$. □

Remark 3.4.4. Note that it is possible for a strict inequality in Lemma 3.4.3, since there may be many different ways to put disjoint copies of elements of σ_Y in $\partial S(Y)$. However, if Y is trivalent and σ_Y satisfies the hypotheses of the lemma, then there is an *equality* $h_{\sigma_Y}(\partial S(Y))$ is equal to three times the number of vertices of Y.

3.4.3 Trivalent fatgraphs are usually extremal

We say that a fatgraph Y over F satisfies condition (SB) if there is a choice of σ_Y as above so that no element of σ_Y^{-1} appears in $\partial S(Y)$.

Lemma 3.4.5. *If a trivalent labeled fatgraph Y satisfies condition (SB), then both $S(Y)$ and h_{σ_Y} are extremal for the boundary $\partial S(Y)$ and certify each other.*

Proof. For a trivalent graph, $h_{\sigma_Y}(\partial S(Y)) \geq 3V$, where V is the number of vertices, by Lemma 3.4.3. By Bavard duality, and Lemma 2.2.6 there is an inequality

$$\mathrm{scl}(\partial S(Y)) \geq \frac{3V}{2D(h_{\sigma_Y})} \geq \frac{V}{4}.$$

On the other hand, since Y is trivalent, the number of edges is $3V/2$, so $\chi(S(Y)) = -V/2$. Hence we get a chain of inequalities

$$\mathrm{scl}(\partial S(Y)) \geq \frac{3V}{2D(h_{\sigma_Y})} \geq \frac{V}{4} = \frac{-\chi(S(Y))}{2} \geq \mathrm{scl}(\partial S(Y)).$$

Hence each of these inequalities is actually an equality, and the lemma follows. □

We now show that condition (SB) is generic in a strong sense. Given $\hat{Y}$, we are interested in the set of Y with $\partial S(Y)$ reduced obtained by labeling the edges of $\hat{Y}$ by words of length at most n. For each n, this is a finite set, and we give it the uniform distribution.

Proposition 3.4.6. *For any combinatorial trivalent fatgraph $\hat{Y}$, if Y is a random fatgraph over F_k obtained by labeling the edges of $\hat{Y}$ by words of length n, then $S(Y)$ is extremal for $\partial S(Y)$ and is certified by some extremal quasimorphism h_{σ_Y}, with probability $1 - O(C(\hat{Y}, k)^{-n})$ for some constant $C(\hat{Y}, k) > 1$.*

Proof. The constant $C(\hat{Y}, k)$ depends only on the number of vertices of $\hat{Y}$. We make use of some elementary facts about random reduced strings in free groups.

If we label the edges of $\hat{Y}$ with random reduced words of length n, it is true that there may be some small amount of folding necessary in order to obtain a fatgraph with $\partial S(Y)$ cyclically reduced. However, the expected amount of letters to be folded is a constant independent of n, which is asymptotically insignificant, and may be safely disregarded here and elsewhere for simplicity.

Now consider some element w of σ_Y under some random labeling. The fatgraph Y over F_k will satisfy condition (SB) with the desired probability if the probability that w^{-1} appears (as a subword) in $\partial S(Y)$ is C^{-n}, because the number of elements of σ_Y is fixed (note that we are using the elementary but useful fact in probability theory that extremely rare events are almost independent).

If w^{-1} appears in $\partial S(Y)$, then at least half of it must appear as a subword of one of the edges of Y, so the probability that w^{-1} appears in $\partial S(Y)$ is certainly smaller than the probability that the prefix or suffix of w of length $n/2$ appears as a subword of an edge of Y. Let E denote the number of edges of $\hat{Y}$. The probability that a subword of length $n/2$ appears in a word of length n is approximately $(n/2)(2k)^{-n/2}$, so, as we must consider each edge and its inverse, the probability that w^{-1} appears is smaller than $2E(n/2)(2k)^{-n/2}$. By replacing k by a slightly smaller constant, we may disregard the $(n/2)$ multiplier, and the lemma is proved. □

3.4.4 Higher valence fatgraphs

For fatgraphs with higher valence vertices, the construction of a candidate extremal quasimorphism is significantly more delicate. We still take a linear combination of counting functions, but we require big counting quasimorphisms, and for a higher-valence vertex, we need to involve the counting functions which count words which cross the vertex but are not necessarily along the boundary. The same warning about the homogenization notation differing from [17] still applies.

For $m \geq 3$ let K_m be the complete graph on m vertices. Label the vertices $0, 1, 2, \ldots, m-1$. Define a weight w_m on directed edges (i, j) of K_m by the formula $w_m(i, i+\ell) = 3 - (6\ell/m)$ where indices are taken mod m.

Lemma 3.4.7. *The function $w_m(i, i+\ell) := 3 - (6\ell/m)$ is the unique function on directed edges of K_m with the following properties:*

1. *It is antisymmetric: $w_m(i, j) = -w_m(j, i)$.*
2. *It satisfies the inequality $|w_m(i, j)| \leq 3 - 6/m$ for all distinct i, j.*
3. *For every distinct triple i, j, ℓ, there is an equality $w_m(i, j) + w_m(j, \ell) + w_m(\ell, i) = \pm 3$ where the sign is positive if the natural cyclic order on i, j, ℓ is positive, and negative otherwise.*
4. *It satisfies $w_m(i, i+1) = 3 - 6/m$ for all i.*

Proof. Only uniqueness is not obvious. If we think of w_m as a simplicial 1-cochain on the underlying simplicial structure on the regular $m-1$ simplex then condition (3) determines the coboundary δw_m, so w_m is unique up to the coboundary of a function on vertices. But condition (4) says this coboundary is zero. □

For x a reduced word in F, parameterize x proportional to arclength as the interval $[-1, 1]$, and let $x[-t, t]$ denote the smallest subword containing the interval from $x(-t)$ to $x(t)$. Fix some small $\epsilon > 0$ and define the *stack function* S_x to be the following integral of big counting functions:

$$S_x = \frac{1}{1-\epsilon} \int_\epsilon^1 C_{x[-t,t]} dt.$$

The ϵ correction term ensures that the length of the shortest word in the support of S is at least $\epsilon|x|$. If x is quite long, this word will also be quite long, and ensure that there are no “accidents” in what follows. The constant ϵ we need is of order $1/(\max_v |v|)$; we leave it implicit in what follows, and in practice ignore it.

Remark 3.4.8. The function S_x is actually a finite rational sum of ordinary big counting functions, since $x[-t,t]$ takes on only finitely many values. We can make it into a genuine integral by first applying the (isometric) endomorphism φ_m to F which takes every generator to its mth power, and then taking $\lim_{m\to\infty} \varphi_m^* S_{\varphi_m(x)}$ in place of S_x. However, this is superfluous for our purposes here.

We are now in a position to define the quasimorphism H_Y.

Definition 3.4.9. Let Y be a fatgraph over F_k, and suppose that every edge has length $\geq 2n$. For each vertex v, denote the set of oriented subarcs in ∂Y of length n ending at v by $x_i(v)$, where the index i runs from 0 to $|v|-1$ and the cyclic order of indices agrees with the cyclic order of edges at v. Denote the inverse of $x_i(v)$ by $X_i(v)$.

Then define

$$H_Y = \sum_v \sum_{i,i+\ell\%|v|} (3-(6\ell/|v|))(S_{x_i(v)X_{i+\ell}(v)} - S_{x_{i+\ell}(v)X_i(v)})$$

(note that the factor $3-(6\ell/|v|)$ is $w_{|v|}(i,i+\ell)$ from Lemma 3.4.7).

Let σ denote a word of the form $x_i(v)X_j(v)$ or its inverse. In other words, the σ are the words in the support of H_Y. Now say that Y satisfies condition (B) if, whenever some $\sigma[a,b]$ appears as a subword of some other σ', or some $\sigma[a,b]$ or its inverse appears twice in σ, then $(b-a)$ is not too big — explicitly, $(b-a) < 6/4(\max_v |v|)$. Hereafter we denote $\delta := 6/4(\max_v |v|)$.

Lemma 3.4.10. *Suppose Y satisfies condition (B). Then $D(H_Y) \leq 6$.*

Proof. Recall that H_Y is a homogeneous quasimorphism. In order to prove the lemma, it suffices to show that the defect of the un-homogenized quasimorphism is at most 3. This reduces to showing that the value of this quasimorphism on a tripod is at most 3. Condition (B) says that if two distinct σ,σ' overlap a junction on one side of a tripod, then $S_\sigma, S_{\sigma'}$ each contributes at most δ to the defect. So we can assume that on at least one side, there is a unique $\sigma = x_i(v)X_j(v)$ with a subword of definite size that overlaps a junction. Again, without loss of generality, we can assume that the junction is at $\sigma(t)$ where $t \in [-1+\delta, 1-\delta]$. By condition (B), if σ' on another side has a subword of definite size that overlaps the junction, it either contributes at most δ, or else we must have $\sigma' = x_\ell(v)X_i(v)$ or $\sigma' = x_j(v)X_\ell(v)$. So the only case to consider is when the three incoming directed edges at the junction are suffixes of $x_i(v), x_j(v), x_\ell(v)$ of length $1 \geq s \geq t \geq u \geq 0$ respectively. But in this case the total contribution to the defect is $u(w_{|v|}(i,j)+w_{|v|}(j,\ell)+w_{|v|}(\ell,i))+(t-u)w_{|v|}(i,j)$. Since $|w_{|v|}(i,j)+w_{|v|}(j,\ell)+w_{|v|}(\ell,i)| = 3$ and $|w_{|v|}(i,j)| < 3$, this defect is ≤ 3, as claimed. $\square$

Theorem 3.4.11 (Random fatgraph theorem)**.** *For any combinatorial fatgraph $\hat{Y}$, if Y is a random fatgraph over F_k obtained by labeling the edges of $\hat{Y}$ by words of length n, then $S(Y)$ is extremal for $\partial S(Y)$ and is certified by the extremal quasimorphism H_Y, with probability $1 - O(C(\hat{Y}, k)^{-n})$ for some constant $C(\hat{Y}, k) > 1$.*

Proof. The argument is a minor variation on the arguments above, so we just give a sketch of the idea.

It suffices to show that a random Y satisfies condition (B) with probability $1-O(C^{-n})$ for some C. But this is obvious, since the $x_i(v)$ are independent, and for any constant $\kappa > 0$, two random words in F of length n do not have overlapping segments of length bigger than κn, and a random word of length n does not have a segment of length bigger than κn that appears twice, in either case with probability $1 - O(C^{-n})$. □

Remark 3.4.12. Since $\chi(S(Y)) \in \mathbb{Z}$, the boundary $\partial S(Y)$ satisfies $\mathrm{scl}(\partial S(Y)) \in \frac{1}{2}\mathbb{Z}$. On the other hand, Theorem 3.4.11 does *not* imply anything about the structure of scl for generic *chains* of a particular length. A random homologically trivial word (or chain) in a hyperbolic group of length n has scl of size $O(n/\log n)$ (see [11]), so a random homologically trivial word of length n *conditioned to have genus bounded by some constant*, will be very unusual.

In fact, computer experiments suggest that the expected denominator of $\mathrm{scl}(w)$ is a proper function of the length of a (random) word w.

There are only finitely many distinct combinatorial fatgraphs with a given Euler characteristic, so if we specialize $\hat{Y}$ to have a single boundary component (recall this depends only on the combinatorics of $\hat{Y}$ and not on any particular immersion), then then we see that for any integer m there is a constant C depending on m so that a random word of length n conditioned to have commutator length at most m has commutator length exactly m and $\mathrm{scl} = m - 1/2$, with probability $1 - O(C^{-n})$.

3.4.5 Experimental data

Our main purpose in this section is to give experimental confirmation of our results and to estimate the constants $C(\hat{Y}, k)$. However, it is worth mentioning that vertex quasimorphisms provide quickly verifiable rigorous (lower) bounds on scl.

3.4.5.1 Fast rigorous lower bounds on scl

Although not every chain admits an extremal surface which is *certified* by a vertex quasimorphism, it happens much more frequently that a vertex quasimorphism certifies good lower bounds on scl. For example, if Y is not trivalent, a quasimorphism of the form h_{σ_Y} will never be extremal; but if the average valence of Y is close to 3, the lower bound one obtains might be quite good.

Because verifying condition (B) requires only checking the (non)-existence of certain words as subwords of the boundary $\partial S(Y)$, plus a small cancellation condition, it is possible to verify the defect of a vertex quasimorphism in polynomial time. This compares favorably to the problem of computing the defect of an arbitrary linear combination of big counting quasimorphisms (or even a single big counting quasimorphism) for which the best known algorithms are exponential.

Example 3.4.13. It is rare for (short) words or chains to admit extremal trivalent fatgraphs. A cyclic word is *alternating* if it contains no $a^{\pm 2}$ or $b^{\pm 2}$ substring; for example, $baBABAbaBabA$ is alternating, with $\mathrm{scl} = 5/6$. An extremal fatgraph for an alternating word necessarily has all vertices of even valence, since the edge labels at each vertex must alternate between one of $a^{\pm}$ and one of $b^{\pm}$.

3.4.5.2 Experimental calculation of constants $C(\hat{Y}, k)$

While *certifying* a vertex quasimorphism is easy, *finding* one is much harder. To verify our asymptotic results, we can be content with breaking the edges of the fatgraph into uniform pieces and checking whether condition (B) is satisfied. However, for a given fatgraph, it might be the case (and usually is) that while a naive assignment of words for H_Y fails, a more careful choice succeeds. To check whether there is *any* vertex quasimorphism is (naively) exponential, and this makes large experiments difficult.

However for *trivalent* fatgraphs, condition (SB) on σ_Y is much simpler. In particular, whether or not a collection of words satisfies (SB) depends only on the (local) no-overlap condition, plus the "constant" condition of certain words not appearing in $\partial S(Y)$. This makes this a priori infeasible problem of checking whether there is any vertex quasimorphism for a particular fatgraph possible with the use of a "meet-in-the-middle" time-space trade-off.

Using this method, we can experimentally estimate the best possible constants $C(\hat{Y}, k)$, at least in the case of trivalent $\hat{Y}$. Figure 3.5 shows some data on the likelihood that a random labeling of a trivalent fatgraph with four vertices admits a vertex quasimorphism. The linear dependence of $-\log(P(\text{fail}))$ on label length is evident. We can calculate a best fit slope and y-intercept for these lines, which gives a best fit line of $1.47336n - 1.42772$, or equivalently, $P(\text{success}) = 1 - 4.16918(4.36387)^{-n}$. Note that the lower right graph is the least likely to admit a vertex quasimorphism; this is heuristically reasonable, since self-loops at vertices handicap the graph by forcing a shorter length on some words in σ_Y. A best fit for this line yields $P(\text{success}) \geq 1 - 82.3971(3.19827)^{-n}$.

3.4.5.3 Using homomorphisms to improve success rate

When a particular labeling Y does not admit a vertex quasimorphism, it might still be possible to find an extremal quasimorphism by applying a homomorphism ϕ to Y. If (the folded fatgraph)

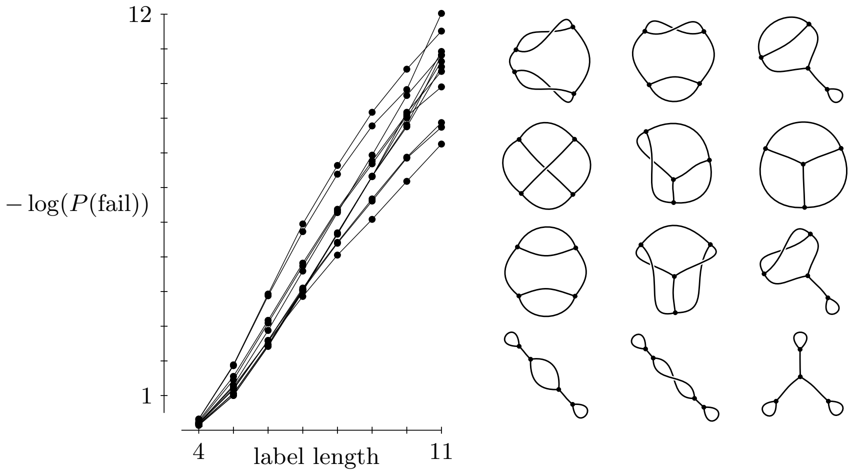

Figure 3.5: A plot of $-\log$ of the failure rate for random labelings of lengths between 4 and 11, plotted for each trivalent fatgraph with four vertices. Each dot represents 500,000 trials. The fatgraphs themselves are arranged left to right, top to bottom in decreasing order of $-\log$ of failure at length 11, so the tripod in the lower right is the "hardest" to find vertex quasimorphisms for. The pictures were created using `wallop` [32].

$\phi(Y)$ admits an extremal vertex quasimorphism $H_{\phi(Y)}$, and folding does not change the Euler characteristic of the fatgraph, then the quasimorphism $\phi^* H_{\phi(Y)}$ is extremal for $\partial S(Y)$.

Because the edges of $\phi(Y)$ are no longer random (and in particular, distinct edge labels will necessarily share long common subwords), it is not clear that applying a homomorphism will affect our success rate. In fact, it turns out to help significantly, especially for shorter labelings. Figure 3.6 shows $-\log$ of the failure rate for a particular fatgraph compared with $-\log$ of the failure rate after applying many random homomorphisms. We decrease the probability of failure by a factor of about 5. Interestingly, changing the length of the homomorphism or the number of homomorphisms that we try does not seem to significantly alter our success with this procedure.

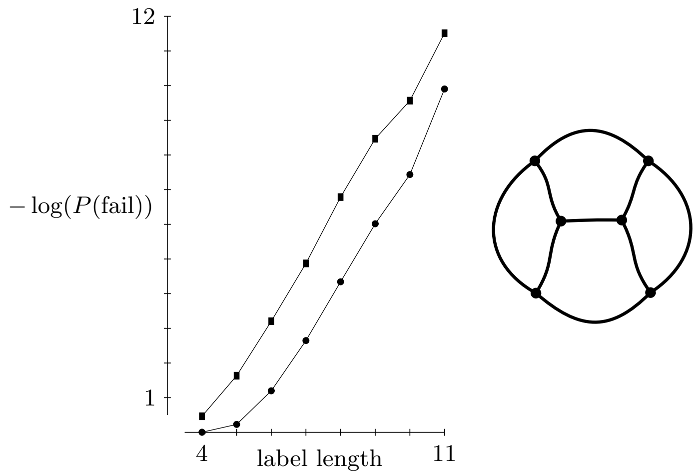

Figure 3.6: A plot of $-\log$ of the failure rate for labelings of the fatgraph (circles), and $-\log$ of the failure rate after acting by many random homomorphisms (squares).

Chapter 4

Cyclic orders and quasimorphisms

In this chapter, we look more closely at the set of surface realizations of a free group. Using fatgraphs and cyclic orders, we give simple combinatorial conditions which certify when rotation quasimorphisms are extremal. This produces many interesting examples. We also give a construction of a new kind of quasimorphism, the *transfer*, and we show that in certain cases, we can take limits to produce yet more quasimorphisms.

The material in this chapter is a subset of the material from the upcoming [12].

4.1 Cyclic orders and compatibility

Recall the definition of a cyclic order (Definition 2.5.8). Also recall that from a cyclic order O on a set of semigroup generators of F_k, we get the basic fatgraph realization (Σ_O, f_O). An example is shown in Figure 2.7. We will be looking more closely at cyclic orders, and we need some new terminology.

Suppose A is a set with a cyclic order O, and $B \subseteq A$ is a subset of A with its own cyclic order O'. Then O is *positively compatible* with O' if O' is the restriction of O to B, and O is *negatively compatible* with O' if O' is the negative of the restriction of O to B. It is *compatible* if it is either negatively or positively compatible. If $\{b_i\}_{i=1}^n \subseteq A$ is a finite subset of A with its own cyclic order O' given by the cyclic tuple $[b_1, \dots, b_k]$, where we have relabeled the b_i without loss of generality, then it is simple to check whether O' is compatible with O. It is sufficient to check that $O(b_i, b_j, b_k)$ is either always 1 or always -1 for all (i, j, k) which are positively ordered in the natural cyclic order on $[1, \cdots, k]$.

Example 4.1.1. Take the natural cyclic order O given by $[1, \dots, 5]$. Then $O(1, 2, 3) = 1$, $O(5, 2, 3) = 1$, and $O(2, 5, 4) = -1$. Also, the cyclic order $[1, 2, 3, 4]$ is positively compatible with O and $[4, 2, 1, 5]$ is negatively compatible. The cyclic order $[1, 2, 5, 4]$ is not compatible with O.

A cyclic order on the semigroup generators of F_k induces a cyclic order on all of F_k in the following way: given three distinct words g_1, g_2, g_3 in F_k, remove any common initial word, and embed the remaining ends in the Cayley graph starting from a given point. There will be a unique tripod contained in the union of the embeddings, and if the tuple (g_1, g_2, g_3) is cyclically ordered in the same way as the generators labeling the tripod, then the tuple is positively ordered. See Example 4.1.2.

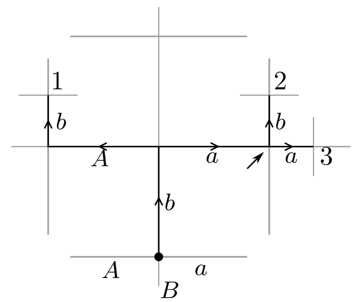

Figure 4.1: Determining if a tuple of words is cyclically in order by using a cyclic order on the generators. The bold circle indicates the basepoint, and the arrow indicates the tripod which determines the cyclic order. See Example 4.1.2.

Example 4.1.2. Give the standard generating set of F_2 the cyclic order $[a, b, A, B]$, and consider the tuple (bAB, bab, baa). Figure 4.1 shows the embeddings and the tripod which remains after removing the common initial segment b. The counterclockwise order of the words around this tripod is $(1, 3, 2)$. Thus, we say that the tuple (bAB, bab, baa) is negatively cyclically ordered in the cyclic order on F_2 determined by $[a, b, A, B]$.

4.2 Realizations and immersions

4.2.1 Fatgraph realizations

For a free group F_k of rank k, there are $(2k-1)!$ basic fatgraph realizations, because there are that many distinct cyclic orders on the semigroup generators.

By precomposing with an automorphism of F_k, we can give a fatgraph structure to any surface realization of F_k: suppose (Σ_O, f_O) is a basic fatgraph realization. Then acting by the automorphism $\phi \in \mathrm{Out}(F_2)$ gives the surface realization $(\Sigma_O, f_O \circ \phi^{-1})$, which we will call a *fatgraph realization.* Again, every surface realization can be represented as a fatgraph realization by choosing a fatgraph structure.

Unfortunately, we need a bit of notation. We will denote by (ϕ, O) the fatgraph realization of F_k obtained by acting by ϕ on the basic fatgraph realization induced by the order O on the semigroup generators of F_k. We denote the rotation quasimorphism obtained on F_k by this surface realization by $\mathrm{rot}_{(\phi,O)}$.

4.2.2 Surface maps into fatgraph realizations

Recall X_k is our standard $K(F_k, 1)$. Suppose that $g : S \to X_k$ is an incompressible surface map, and (ϕ, O) is a fatgraph realization (ϕ might be the identity). Then $\phi^{-1} \circ g : S \to X_k$ has a fatgraph representative, which by abuse of notation we also denote by S. Note there are two fatgraphs in the picture: the target fatgraph surface realization of F_k, and the fatgraph representative of the surface map $\phi^{-1} \circ g$. Now S is a labeled fatgraph over F_k, and as such it has a cyclic order on the edges at each vertex. In addition, the order O induces *another* cyclic order on the edges at each vertex of S by simply applying O to the labels on the edges. We will call this cyclic order on each vertex the pullback of O. Note this cyclic order might not be well-defined if the labels are not all distinct around a vertex.

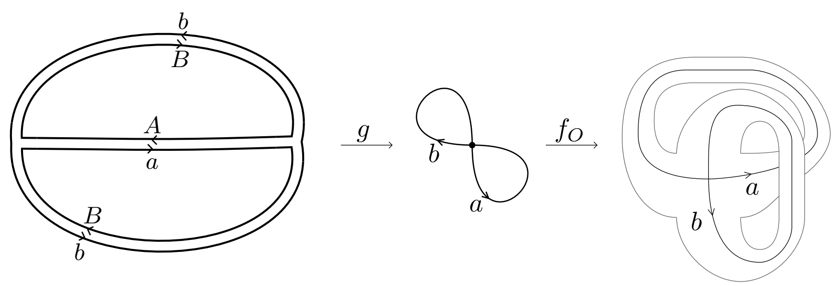

Figure 4.2: Comparing the cyclic orders at each vertex on the edges of the fatgraph coming from the fatgraph structure and from the ordering on labels from the pullback of the order O on semigroup generators of F_2. See Example 4.2.1.

Example 4.2.1. Figure 4.2 illustrates the different cyclic orders on vertices coming from both the intrinsic fatgraph order and the order induced by the labels. Here the fatgraph realization of F_2 is $(\text{id}, [b, a, B, A])$, as shown by the map f_O. The surface map $g : S \to X_k$ induces a fatgraph structure on the surface S (in the picture, the surface map is implicit from the labeling), and this fatgraph structure includes the cyclic orders on the vertices, which are indicated here by the planar embedding.

The fatgraph realization gives the cyclic order $[b, A, B, A]$ on semigroup generators for F_2, and this induces cyclic orders on the vertices of S. In this case, the vertex on the left has outgoing edges (b, a, B), which is positive in both the S order and the pullback of O. However, the vertex on the right has outgoing edges (B, b, A), which is positive in the fatgraph order but negative in the pullback of O.

Theorem 4.2.2. *Let $g : S \to X_k$ be a surface map. Let (ϕ, O) be a fatgraph realization of F_k. Give S a labeled fatgraph structure induced by the map $\phi^{-1} \circ g : S \to X_k$. Then the following are equivalent.*

1. *The intrinsic fatgraph orders on all vertices in S are compatible (all positively or all negatively)*

with the orders given by the pullback of O.

2. *The induced quasimorphism* $\mathrm{rot}_{(\phi,O)}$ *is extremal for* ∂S*.*

3. *The map* $f_O \circ \phi^{-1} \circ g : S \to \Sigma_O$ *is homotopic to an immersion with geodesic boundary.*

Proof. That (2) and (3) are equivalent is the content of Theorem 2.5.5. The equivalence of (1) can be seen intuitively in Figure 4.3. However, the quickest way to see this is to appeal to the fact that the rotation quasimorphism induced by a realization has the area form (divided by -2π) on the hyperbolic surface Σ_O as its coboundary (see [7]). Decompose each vertex of the fatgraph into tripods. There is a pleated representative of S which has one ideal triangle for each tripod. The rotation quasimorphism applied to ∂S will therefore compute the sum of $\pm 1/2$ over these triangles, depending on whether they are positive or negative. Since $-\chi(S)$ is $1/2$ times the number of triangles, we have that $\mathrm{rot}_{(\phi,O)}(\partial S) = -\chi(S)$ if and only if all tripods are correctly cyclically ordered around all vertices. □

Condition (1) in Theorem 4.2.2 is a simple combinatorial check, provided that we have a fatgraph representative of $\phi^{-1} \circ g : S \to X_k$. However, in practice, the map $g : S \to X_k$ is given as a labeled fatgraph, and if ϕ is not the identity, it appears tricky to obtain a fatgraph representative of $\phi^{-1} \circ g : S \to X_k$, which is the fatgraph of interest in Theorem 4.2.2. In fact, it is possible to push a fatgraph map forward under ϕ^{-1} with a minimum of effort: simply relabel the edges by applying ϕ^{-1} and fold (in the sense of Stallings fatgraph folding) the resulting fatgraph. See Figure 4.3, (A)–(C).

In addition, if $g : S \to X_k$ is a surface map and (ϕ, O) a fatgraph realization of F, it is possible to check condition (1) of Theorem 4.2.2 "all at once," rather than on each vertex seperately. Consider S as a labelled fatgraph obtained from $\phi^{-1} \circ g : S \to X_k$. Take a maximal subgraph T of S — the remaining edges of S determine a set of generators $\{g_i\}_{i=1}^k$ for $\pi_1(S)$. Since $\pi_1(S) \subseteq F_k$, the cyclic order O determines a cyclic order on the g_i. Furthermore, the cyclic order on each vertex of S determines a planar embedding of T and thus a cyclic order O' on the g_i. Then clearly

Corollary 4.2.3. *Condition (1) of Theorem 4.2.2 is is satisfied if and only if O' is compatible with O, in the above notation.*

This idea will be pursued in detail in Section 4.6.

4.3 Examples and consequences

Theorem 4.2.2 gives an explicit condition by which to check the extremality of a particular fatgraph realization on a chain C, given an extremal surface for C. This condition of compatible cyclic orders is simple to check for any particular example, but we are particularly interested in conditions which say something about the extremality of families of realizations.

The complication is that it is difficult to know when there is (or is not) an automorphism which acts on a fatgraph in such a way that after folding, we are left with a new fatgraph whose vertices are all compatible with a cyclic order, as in the following example.

Example 4.3.1. See Figure 4.3. The tautological torus (A) bounding $[BaBA, bb]$ does not bound an immersed surface with geodesic boundary in any basic fatgraph realization (for example, it is evident in the realization $[a, A, B, b]$, shown in (D), that we cannot remove the twist without introducing others). However, if we act by the automorphism $a \mapsto A$, $b \mapsto ba$ (so we push forward by its

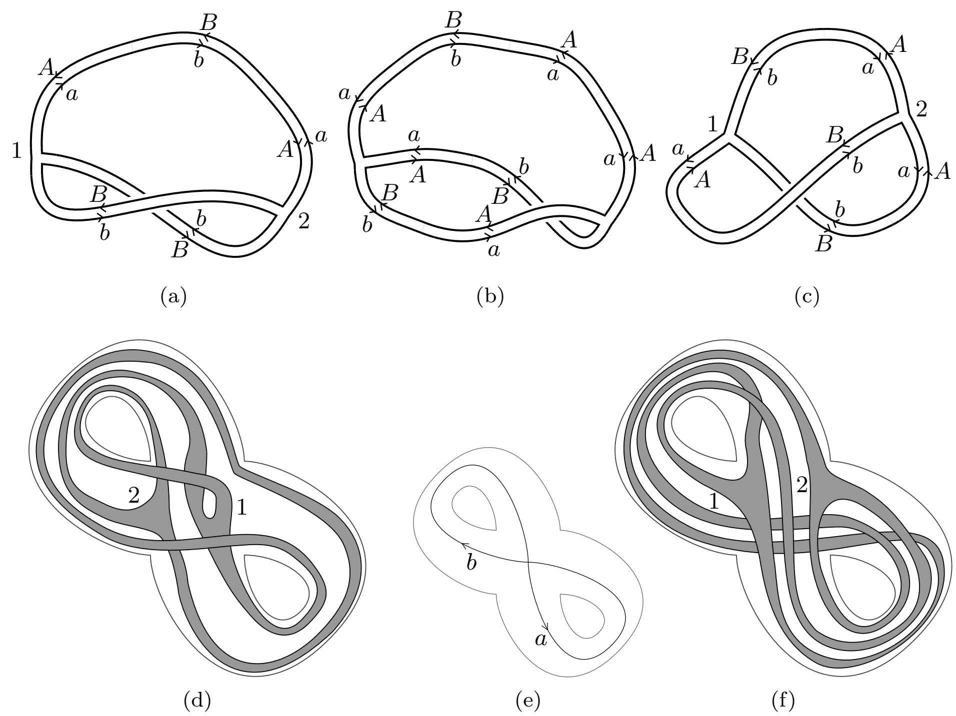

Figure 4.3: See Example 4.3.1. Corresponding vertices in the fatgraphs and immersions are indicated with numbers.

inverse, but this automorphism has order two, so it is its own inverse), then we get the fatgraph (B), and after folding to (C), the resulting fatgraph does bound an immersed surface (F) in the basic realization corresponding to the order $[a, A, B, b]$ (which is shown in (E)). Therefore, the commutator $[BaBA, bb]$ bounds an immersed surface in the fatgraph realization $(\{a \mapsto A,\, b \mapsto ba\},\, [a, A, B, b])$, and the associated rotation quasimorphism is extremal.

This example illustrates how acting by automorphisms can cause certain surfaces to become immersed, but it should also be clear that complicated combinatorial processes are at work. Consequently, while Theorem 4.2.2 gives a sufficient condition for a fatgraph realization to be extremal for a chain, it gives no indication as to how to produce such a realization or how to show that none exist.

We can, however, make some general statements, which rely on combinatorial facts about automorphisms of free groups. Sources for these types of techniques include [3], [26] and [25]. We will be concerned in this section exclusively with *cyclic* words in a free group, which means that we will consider only conjugacy classes of words, or equivalently baseless loops in $X_k = K(F_k, 1)$. A subword of a cyclic word, however, can be non-cyclically reduced.

Because we only will look at cyclic words, we will only be interested in $\mathrm{Out}(F_k)$, so we will remove any conjugation from automorphisms. Let $\phi \in Out(F_k)$ be such an automorphism, and suppose that w is a cyclic word which we write as a product of subwords $w = gh$, where there is no cancellation in the product gh. To compute $\phi(w)$, we can compute the (cyclic) product $\phi(w) = \phi(g)\phi(h)$. There might be some cancellation in the product $\phi(g)\phi(h)$. In fact, $\phi(g)$ might completely disappear. For example, set $\phi(a) = a$, $\phi(b) = Ab$, $g = a$, and $h = b$. Then $w = ab$, and $\phi(w) = \phi(a)\phi(b) = aAb = b$. Note that the remaining letter comes from $\phi(h)$. However, we generally expect some of $\phi(g)$ to remain.

Definition 4.3.2. We call a (non-cyclic) word g *blocking for* ϕ if for all h so that the cyclic product gh is reduced, we have that some of $\phi(g)$ remains in the product $\phi(g)\phi(h)$.

Lemma 4.3.3. *Let F_k be a free group and $\phi \in \mathrm{Out}(F_k)$ be irreducible. Then there exists a word $g \in F_k$ which is blocking for ϕ^n for all $n \geq 0$.*

Proof. The best way to see this is through train track maps; we need to recall the standard terminology which can be found in any of the above named sources. The key is that ϕ can be represented as a train track map with uniformly expanding dynamics: say that the length of the image of any edge is multiplied by $\lambda > 1$. In addition, for any irreducible automorphism ϕ, there is a bounded cancellation constant C_ϕ; see [25] §3. This means that if g and h are such that the product is reduced, then the length of the segment which cancels (on either side in the cyclic multiplication) in $\phi(g)\phi(h)$ is less than C_ϕ. Choose a loop g with length $|g|$ such that $\lambda|g| - 2C_\phi > |g|$. After applying ϕ to gh, we must therefore find a segment from $\phi(g)$ of length at least $|g|$, so g is blocking for ϕ, and the length of what remains of ϕ is at least $|g|$; we may repeat this argument as many times as we like, so g is blocking for ϕ^n for all n. $\square$

Lemma 4.3.4. *The word $[a, b]^3$ is blocking for all $\phi \in \mathrm{Out}(F_2)$.*

Proof. The case of F_2 is special because $[a, b]$ is invariant (up to conjugation) under $\mathrm{Out}(F_2)$. This is a consequence of the fact that the mapping class group of the once-punctured torus is equal to the outer automorphism group of its fundamental group. Note that higher rank free groups do not have $\mathrm{Out}(F_2)$- invariant words.

Therefore, let $g = [a, b]^3$, and suppose that we have the cyclic word gh and $\phi \in \mathrm{Out}(F_2)$ such that $\phi(g)$ is completely cancelled by $\phi(h)$. Then $\phi(h)$ must begin or end with $[a, b]^{-1}$, so h must begin or end with $[a, b]^{-1}$, which contradicts gh being reduced. $\square$

Attempting to produce words which are blocking for all $\phi \in \mathrm{Out}(F_k)$ for higher rank free groups F_k is surprisingly slippery. However,

Conjecture 4.3.5. *There exists $g \in F$, where F is any rank free group, such that g is blocking for all $\phi \in \mathrm{Out}(F)$.*

A reasonable approach here is to generalize the results in [4] to (un)stable laminations for automorphisms of free groups, then to pick a word which does not align with the lamination for any automorphism. The case of finite-order (and non-irreducible) automorphisms would have to be considered separately.

Construction 4.3.6. The point of these blocking words is the following construction. Let g be a blocking word for $\phi \in \mathrm{Out}(F_k)$. Consider the fatgraphs shown in Figure 4.4. Everything is assumed to be reduced. These fatgraphs have four key features:

1. The surface map induced by the labels is extremal (because it bounds a commutator).
2. The outgoing edge labels at the vertices are the same.
3. The fatgraph (A) is positively or negatively immersed in any basic fatgraph realization. The fatgraph (B) is not immersed in any basic fatgraph realization.
4. The same is true for the fatgraph realization (ϕ, O), for any order O.

The last two require a little explanation. For (3), if the cyclic orders on the vertices coming from the fatgraph structure are the same (as in (A)), then no matter what cyclic order induces the basic fatgraph realization of F_k, it will be positively or negatively compatible with the order in (A). However, if the cyclic orders are different (as in (B)), then no matter what the cyclic order associated with the fatgraph realization is, it cannot be compatible with both vertices.

The most important is (4): since g is blocking for ϕ, we can apply ϕ to the fatgraphs (A) and (B), and we know that the folding process cannot cross any of the blocking words g (and g^{-1}), so the fatgraph must fold identically around each vertex. Therefore, whatever the folded fatgraphs are, they still have two vertices with the same outgoing labels and the same (A) or different (B) cyclic orders induced by the fatgraph structure.

Lemma 4.3.7. *Let F_k be a free group. For every basic fatgraph realization O and irreducible automorphism $\phi \in \mathrm{Out}(F_k)$, there exist infinitely many commutators C such that $\mathrm{rot}_{(\phi^n,O)}$ is extremal for C for all n and infinitely many commutators C' such that there is no n so that $\mathrm{rot}_{(\phi^n,O)}$ is extremal for C'.*

Proof. This is simply Figure 4.4 (A) (for the first case) and (B) (for the second case), using the blocking word g provided by Lemma 4.3.3. Note that it might be the case that $\mathrm{rot}_{(\phi^n,O)}(C)/2 = -\mathrm{scl}(C)$;

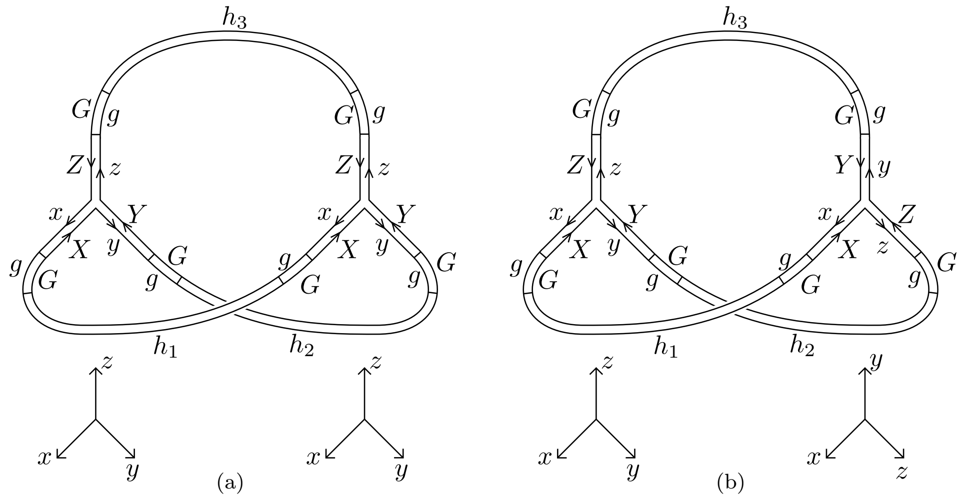

Figure 4.4: Blocked fatgraphs bounding commutators; in (A), the vertices have the same cyclic orders, while in (B), they differ. The outgoing edge labels are indicated below each vertex. For clarity, capital letters denote inverses of words, and the inverse labels on the middle segments h_i are not shown.

recall that our definition of extremal includes this. We get infinitely many of these commutators because it does not matter what the words h_i in the middle are, as long as everything is reduced. □

Using Lemma 4.3.4, we can do better in a free group of rank 2. This is a fairly simple consequence of our context, but we elevate it to the status of a theorem because it is interesting and it is one of our main results.

Theorem 4.3.8. *Let F_2 be a free group of rank 2. There exist infinitely many commutators C for which* $\mathrm{rot}_{(\phi,O)}$ *is extremal for C for all fatgraph realizations (ϕ, O) (and thus all surface realizations). There exist infinitely many commutators C' for which there is no extremal surface realization.*

Proof. Again, this is simply Figure 4.4 (A) and (B), but this time we may use the blocking word provided by Lemma 4.3.4 to obtain the stronger conclusion. □

Recall now the definition of a geometric face from Definition 2.5.7. These lemmas allow us to draw some interesting conclusions about geometric faces of the scl norm ball. One technicality is that while we say that a realization is extremal for C if $|\mathrm{rot}(C)|/2 = \mathrm{scl}(C)$, in terms of the scl norm ball, this means that C is contained in the closure of *either* the face dual to rot or the face dual to $-$rot. Let T denote (the closure of) a geometric face. The scl ball is invariant under reflection through the origin because $\mathrm{scl}(C) = \mathrm{scl}(C^{-1})$; denote the matching inverse face of T by $-T$.

Corollary 4.3.9. *Let T be a geometric face of the* scl *norm ball in $B_1^H(F_k)$ for a free group F_k, and let $\phi \in \mathrm{Out}(F_k)$ be irreducible. Then either T or $-T$ has a common intersection point with infinitely many geometric faces obtained by acting by ϕ^n.*

Proof. A chain is in the face $\phi \cdot T$ if and only if $\phi(C)$ is contained in T. Thus, repeat Figure 4.4 (A) with a blocking word for ϕ to obtain a commutator C. After acting by any power of ϕ, either C or C^{-1} is contained in T, and the lemma follows. □

Remark 4.3.10. By using different commutators, we can obtain different patterns of intersection.

Corollary 4.3.11. *There are infinitely many commutators C in $B_1^H(F_2)$ so that each one partitions the set $\mathcal{T}$ of geometric faces into two sets; $\mathcal{T}_0$, those that contain C, and $\mathcal{T}_0$, those that contain $C^{-1} = -C$.*

Proof. This is the same argument, now using Lemma 4.3.4. □

4.4 Transfer

Discrete, faithful representations $F_k \to \mathrm{PSL}(2, \mathbb{R})$ consist only of hyperbolic and parabolic elements, and since everything fixes a point, the induced rotation quasimorphism can take only integer values. Therefore, rotation quasimorphisms induced by discrete, faithful representations can be extremal only for chains with scl in $\frac{1}{2}\mathbb{Z}$.

Since many chains have scl with larger denominators (in fact, the scl spectrum is dense in $\mathbb{Q}/\mathbb{Z}$, by Corollary 3.3.14), we need to consider a wider class to have a hope of finding extremal quasimorphisms for general chains. The idea of transfer provides very natural candidate quasimorphisms.

4.4.1 Transfer of quasimorphisms

In this section, we will supress the subscript in the notation F_k, because it simplifies the notation.

Let G be a finite index subgroup of a free group F of rank k. There is a finite cover $K(G, 1) \to X_k$ which is just a graph covering of the rose. Let $\gamma \in X_k$ be a loop. We may lift γ to its total preimage in $K(G, 1)$, which is typically a collection of loops. This map extends by linearity to $B_1^H(F)$, and we denote it by T_F^G, so $T_F^G : B_1^H(F) \to B_1^H(G)$. Technically, we must prove that it is well-defined and has the image claimed.

Lemma 4.4.1. *The image of T_F^G is contained in $B_1^H(G)$, and T_F^G is a (well-defined) homomorphism.*

Proof. That the image of T_F^G is contained in $B_1^H(G)$ is a consequence of [8], Proposition 2.80. Briefly, the problem is to lift a surface S bounding $\Gamma \in B_1^H(F)$ to one bounding $T_F^G(\Gamma)$. Since there is a finite index normal subgroup K contained in G, we may take the (finite) cover of S corresponding to the kernel of $\pi_1(S) \to F/K$; this surface then lifts to $K(G, 1)$

To show T_F^G is well-defined, we merely observe that the total lifting procedure commutes with taking powers and conjugation. □

Remark 4.4.2. There is an inclusion $I : B_1^H(G) \to B_1^H(F)$; note that $I \circ T_F^G$ is multiplication by $[F : G]$.

The stable commutator length of the transfer of a chain is determined by the stable commutator length of the chain:

Lemma 4.4.3 ([8], Proposition 2.80)**.** *For all* $\Gamma \in B_1^H(F)$*,*

$$\mathrm{scl}_F(\Gamma) = \frac{1}{[F : G]}\mathrm{scl}_G(T_F^G(C)).$$

There is a natural map induced on quasimorphisms by T_F^G: define ${}^*T_F^G : Q(G) \to Q(F)$ by

$${}^*T_F^G(\phi)(\Gamma) = \frac{1}{[F : G]}\phi(T_F^G(\Gamma)).$$

For convenience, we denote ${}^*T_F^G(\phi)$ by ϕ_F^G.

Lemma 4.4.4.

$$D(\phi_F^G) \leq D(\phi),$$

where the defect of ϕ *is over* G *and the defect of* ϕ_F^G *is over* F*.*

Proof. First, ϕ_F^G is homogeneous because T_F^G and ϕ are. Therefore, its defect is the supremum of its value over all thrice-punctured spheres. Let $C \in B_1^H(F)$ be a thrice-punctured sphere, and let $n = [F : G]$. By Lemma 4.4.3,

$$\frac{1}{2} \geq \mathrm{scl}_F(C) = \frac{1}{n}\mathrm{scl}_G(T_F^G(C)).$$

Therefore,

$$\frac{1}{2} \geq \frac{1}{n}\frac{\phi(T_F^G(C))}{2D(\phi)} = \frac{\phi_F^G(C)}{2D(\phi)},$$

so $D(\phi) \geq \phi_F^G(C)$. Since this holds for all thrice-punctured spheres C, $D(\phi) \geq D(\phi_F^G)$. $\square$

4.4.2 Transfer of rotation quasimorphisms

Let G be a finite index subgroup of F, and let $\rho : G \to \mathrm{PSL}(2, \mathbb{R})$ be a discrete, faithful representation. By the above, $(\mathrm{rot}_\rho)_F^G$ is a homogeneous quasimorphism on F of defect at most 1, and it takes values in $\frac{1}{[F:G]}\mathbb{Z}$. We record the following observation as a lemma.

Lemma 4.4.5. *Let* C *be a chain in* $B_1^H(F)$ *represented by the collection of loops* Γ*, where* F*,* G*, and* ρ *as as above. Then* $(\mathrm{rot}_\rho)_F^G$ *is extremal for* C *if and only if the total preimage of* Γ *bounds an immersed surface in the realization of* G *corresponding to* ρ*.*

Naturally, we are most interested in the case where Γ does not bound an immersed surface in any realization of F.

In anticipation of the fact that there exist nontrivial examples of it, we define a *vimersion* to be a chain $\Gamma \in B_1^H(F)$, together with a finite index subgroup $G \subseteq F$ and a discrete, faithful representation $\rho : F \to \mathrm{PSL}(2,\mathbb{R})$ such that $(\mathrm{rot}_\rho)_F^G$ is extremal for Γ, or equivalently, that $T_F^G(\Gamma)$ bounds an immersed surface in the realization of G corresponding to ρ.

4.4.3 Example

In this section we describe an explicit infinite family of vimersions.

Definition 4.4.6. For each positive integer n, let w_n denote the chain in $F = \langle x, y\rangle$ defined by the formula

$$w_n = x^2 + y^n + yXY^{n+1}X.$$

We will see that $\mathrm{scl}(w_n) = (2n+1)/(2n+2)$.

For each n, let G_n denote the normal subgroup of F of index $(n+1)$ which is the kernel of the homomorphism $F \to \mathbb{Z}/(n+1)\mathbb{Z}$ sending $x \to 0$ and $y \to 1$. In terms of generators, we have

$$G_n = \langle x, yxY, y^2xY^2, \dots, y^nxY^n, y^{n+1}\rangle.$$

Denote y^ixY^i by a_i for $i = 0, \cdots, n$, and denote y^{n+1} by b.

In terms of these generators, the transfer w_n' of w_n has the form

$$w_n' = a_0^2 + a_1^2 + \cdots + a_n^2 + b^n + A_1BA_1A_2BA_2 \cdots A_nBA_nbA_0BA_0.$$

A realization of G_n, and therefore also a rotation quasimorphism on G_n, is given implicitly by a cyclic ordering on the generators of G_n and their inverses, as in [8], Theorem 4.76. This particular realization has underlying surface $\Sigma_{0,n+3}$ a sphere with $n+3$ holes, and the associated cyclic ordering is

$$B, b, A_n, a_n, \dots, A_1, a_1, A_0, a_0.$$

We denote the associated rotation quasimorphism by $\mathrm{rot}_{0,n+3}$.

By the rotation quasimorphism formula, Theorem 2.5.10, we compute

$$(\mathrm{rot}_{0,n+3})_F^G(w_n) = \frac{1}{n+1}\mathrm{rot}_{0,n+3}(w_n') = \frac{(2n+4)(2n+1)}{(n+1)(2n+4)} = \frac{2n+1}{n+1}.$$

The defect of rot_F^G is (at most) 1, so we obtain a lower bound $\mathrm{scl}(w_n) \geq (2n+1)/(2n+2)$. To see that this lower bound is an equality, we will show that the geodesic representatives of the chains w_n'

virtually bound immersed surfaces in the realizations $\Sigma_{0,n+3}$. We show this by induction.

Note that it suffices to show that representatives of the homotopy class of the w'_n which are minimal with respect to Reidemeister moves (*efficient*) virtually bound immersed surfaces. Consider the case $n = 1$, and for simplicity denote a_0, a_1, b by a, b, c. In this notation, $w'_1 = a^2 + b^2 + c + ACABCBc$. The surface $\Sigma_{0,4}$ is a disk with three holes, corresponding to the conjugacy classes of a, b and c. Figure 4.5 indicates an efficient representative of the chain $ACABCBc$. From the figure, it is clear that w'_1 bounds an immersed 5-holed sphere.

Remark 4.4.7. The immersed sphere is more easily seen if one takes a double cover of D^2 "branched" over the puncture corresponding to the loop a. The loop $ACABCBc$ has two preimages in this cover, each of which bounds a (visually obvious) 5-holed sphere.

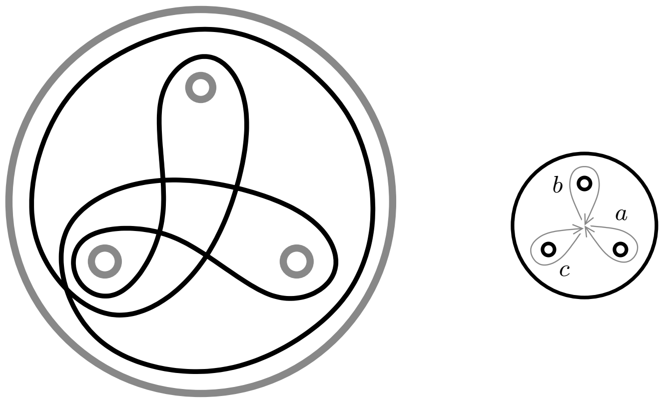

Figure 4.5: The loop in $\Sigma_{0,4}$ associated with the conjugacy class $ACABCBc$ (generators for π_1 are indicated to the right).

The immersed surfaces bounding the chains w'_n may all be constructed by induction. We illustrate the case $n = 2$; the other cases are all essentially the same. For simplicity, if we denote a_0, a_1, a_2, b by a, b, c, d the "interesting" (non-boundary parallel) loop in w'_2 corresponds to the conjugacy class $ADABDBCDCd$. As in Figure 4.6, take a 2-fold branched cover over the puncture corresponding to the loop a. In this cover, $ADABDBCDCd$ lifts to two loops as indicated in the figure, and after filling in four punctures, and pushing across a bigon, one gets a curve isotopic to the "interesting" component of w'_1. Hence w'_2 bounds an immersed 7-holed sphere.

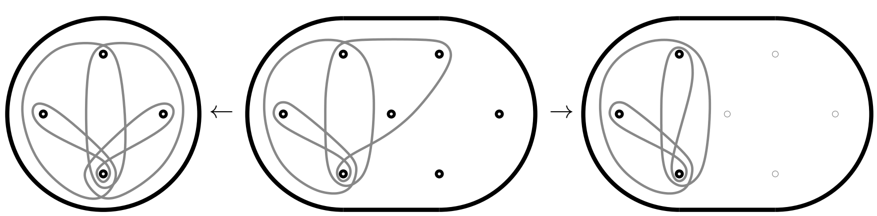

Figure 4.6: A preimage of $ADABDBCDCd$ in a double cover can be pushed over two punctures to become a copy of $ACABCBc$.

This picture carries over to higher n: the "interesting" component of w'_n lifts to two loops in a 2-fold cover of $\Sigma_{0,n+3}$ "branched" over the puncture corresponding to a_0. After filling in half the punctures in this cover together with the puncture corresponding to a_0, one copy can be pushed across a bigon to get a curve isotopic to the "interesting" component of w'_{n-1}, which, together with suitable boundary parallel loops, bounds an immersed surface by the induction hypothesis.

We therefore obtain a proof of the following:

Proposition 4.4.8. *For each $n \geq 1$, the chain $w_n \in F_2$ as above satisfies* $\mathrm{scl}(w_n) = (2n+1)/(2n+2)$. *Moreover, the transfer of w_n bounds an immersed surface in (gives a vimersion with) $G_n \to \pi_1(\Sigma_{0,n+3})$ where G_n is the subgroup of index $n+1$ in F_2, also as above.*

4.5 Symplectic rotation number

Regular $\mathrm{SL}(2,\mathbb{R})$ representations and transfer are special cases of pulling back *symplectic* rotation number from symplectic representations. The symplectic rotation number is a class function $\mathrm{rot} : \mathrm{Sp}(2n,\mathbb{R}) \to S^1$. Restricted to $U(n)$, it is simply det; more generally, to compute $\mathrm{rot}(A)$, we find the 2-dimensional subspaces of $\mathbb{R}^{2n}$ which are fixed and rotated by A (called *eigenplanes*, even though they are not just dialated), and rot is the sum of the angles of these rotations. This computation can also be expressed in terms of the Jordan blocks of A over $\mathbb{R}$.

The function rot lifts to the universal central extension $\widetilde{\mathrm{Sp}}(2n,\mathbb{R})$, and it is a quasimorphism with defect (at most) n. A representation $F_k \to \mathrm{Sp}(2n,\mathbb{R})$ lifts, and we pull back to obtain a quasimorphism on F_k. That is,

$$\begin{array}{ccccc} F & \xrightarrow{\tilde{\rho}} & \widetilde{\mathrm{Sp}}(2n,\mathbb{R}) & \xrightarrow{\widetilde{\mathrm{rot}}} & \mathbb{R} \\ & & \downarrow & & \downarrow \\ & & \mathrm{Sp}(2n,\mathbb{R}) & \xrightarrow{\mathrm{rot}} & S^1 \end{array}.$$

Clearly, $\mathrm{SL}(2,\mathbb{R})$ representations are the special case $n = 1$. Transfer is the special case in which we produce an $\mathrm{Sp}(2n,\mathbb{R})$ representation by taking the direct sum of the n different $\mathrm{SL}(2,\mathbb{R})$ representations on the subgroup obtained by acting by conjugation in the larger group. See [8] §5.2.3.

Unfortunately, representations of free groups into symplectic groups are far less amenable to computation, and geometric intuition is more difficult to come by.

4.6 Ends

In this section, we combine the ideas of the preceeding sections to provide another perspective and to define a new class of quasimorphisms called *limit transfers*.

4.6.1 Cyclic orders and the extremality of transfer

4.6.1.1 Notation

For the following material, it is useful to compare a $K(F,1)$ for a free group to a $K(G,1)$ for a subgroup $G \subseteq F$. Therefore, we add some redundant notation: $K_F = X_k = K(F_k, 1)$.

4.6.1.2 Transfer of fatgraphs

Recall the transfer function from Section 4.4, that is, $T_F^G : B_1^H(F) \to B_1^H(G)$, where $G \subseteq F$ has finite index, and $T_F^G(C)$ is the total preimage of (loops representing) C in the covering space (graph) of $K_F = K(F,1)$ corresponding to G; note this covering space K_G is a $K(G,1)$. Suppose that $f : S \to K_F$ is an admissible surface for C. Now f has a fatgraph representative, and that is how we will think of S. From this perspective, it is obvious that we obtain a new surface $f' : S' \to K_G$ by taking the total preimage of the fatgraph S, and the image of S' bounds $T_G^F(C)$. We have the following commutative diagram, in which the vertical maps are covers of degree $[F : G]$.

$$\begin{array}{ccc} S' & \xrightarrow{f'} & K_G \\ \downarrow & & \downarrow \\ S & \xrightarrow{f} & K_F \end{array}$$

Because S' covers S with degree $[F : G]$ and bounds $T_F^G(C)$, Lemma 4.4.3 implies that S' is extremal for $T_F^G(C)$ if and only if S is extremal for C.

Now suppose that we have a surface realization of G, which we realize as the fatgraph realization (ϕ, O), so $f_O \circ \phi^{-1} : K_G \to \Sigma_O$ is the surface realization. As discussed in Section 4.4, we are interested in the case that the surface realization (f_O, Σ_O) does not cover a surface realization of F. That is, we have the following commutative diagram, which Lemma 4.6.1 references.

$$\begin{array}{ccccc} S' & \xrightarrow{f'} & K_G & \xrightarrow{f_O \circ \phi^{-1}} & \Sigma_O \\ \downarrow & & \downarrow & & \\ S & \xrightarrow{f} & K_F & & \end{array}$$

Lemma 4.6.1. *The hypothesis of Theorem 4.2.2 applies to the fatgraph S' and surface realization (ϕ, O) of G (in the above commutative diagram) if and only if $(\mathrm{rot}_{(\phi,O)})_F^G$ is extremal for C.*

Proof. This is the combination of Theorem 4.2.2 and the above observations, including Lemma 4.4.3. $\square$

Lemma 4.6.1 says that we can detect whether the transfer of rot for a realization of G is extremal for for a chain $C \in B_1^H(F)$ by checking a combinatorial condition on the transfer of an extremal fatgraph for C; that is, that the cyclic order on each vertex of the fatgraph representative of $\phi^{-1} \circ f' : S' \to K_G$ is compatible with the cyclic order given by applying the order O to the labels at each vertex. Unfortunately, while it is simple to take the total preimage of the fatgraph S, it is not so easy to combinatorially produce the set of vertices and labels after applying the automorphism ϕ^{-1}. The purpose of this section is to describe another way to check this combinatorial condition which is more feasible.

4.6.1.3 Ends

If $K_1 \subseteq K_2 \subseteq \cdots$ are compact sets whose union covers a space X, then an *end* of X is a sequence of sets $U_1 \supseteq U_2 \supseteq \cdots$, where U_i is a connected component of $X \setminus K_i$. Essentially, the ends are the connected components of X "at infinity." The ends of a group are the ends of its Cayley graph, and the ends of a free group can be identified, after picking a basepoint, with right-infinite words in the generating set; this is a Cantor set, and we give it that topology, so automorphisms of the free group extend to continuous maps on the ends. In addition, a cyclic order (ϕ, O) on the free group F extends to a cyclic order on the ends: apply the automorphism ϕ^{-1} and then the basic cyclic order O.

If $S \to K_F$ is a fatgraph map, then we lift to a map on the universal covers $\widetilde{S} \to \widetilde{K}_F$, which identifies the ends of $\widetilde{S}$ with a subset of the ends of F. Choose a basepoint and a fundamental domain T (note T is a tree) in $\widetilde{S}$ such that T contains a neighborhood of every vertex in it, so the leaves of T are all partial edges. For each leaf of this tree, pick some end of $\widetilde{S}$ from the basepoint which exits T through the leaf. We call a collection E of ends of $\widetilde{S}$ *exhausting* if it has this property; that is, if for each leaf of T, there is exactly one end in E which exits T at that leaf. Note that an exhausting set of ends gets an induced cyclic order from S.

Suppose that we have a basic fatgraph realization (id, O) of F. Checking the hypothesis of Theorem 4.2.2 for S and O requires checking each vertex of S to see whether its cyclic order is compatible with the one induced by O. However, this is equivalent to checking whether the intrinsic cyclic order on an exhausting set of ends E is compatible with the order that O gives E. This is essentially the same approach as Corollary 4.2.3.

The point of this approach is that while vertices and labels behave badly under automorphisms, the condition of being exhausting is invariant:

Lemma 4.6.2. *If E is an exhausting collection of ends for a fatgraph map $f : S \to K_F$, then $\phi \cdot E$ is an exhausting collection of of ends for $\phi \circ f : S \to K_F$.*

Proof. The automorphism ϕ acts on the Cayley graph of F, and this action restricts to its action on $\widetilde{S}$, which covers its action on S. Thus, ϕ takes a fundamental domain of $\widetilde{S}$ to one of $\phi \cdot \widetilde{S}$, so for each leaf of $\phi \cdot \widetilde{S}$, there is exactly one end in $\phi \cdot E$ which exits that leaf. $\square$

Remark 4.6.3. Given a fatgraph, we can produce an exhausting collection of ends, and we know where those ends exit the fundamental domain for the universal cover of the fatgraph. Even though we know that acting by an automorphism preserves the property of being exhausting, we do not a priori know where exactly the acted-on ends exit the fundamental domain for the acted-on fatgraph. This is an unimportant inconvenience which allows Lemma 4.6.2 to work.

Lemma 4.6.4. *Let $f : S \to K_F$ be a fatgraph map and E an exhausting collection of ends for S. Then Theorem 4.2.2 applies to S and the realization (ϕ, O) if and only if the intrinsic cyclic order on $\phi^{-1} \cdot E$ is compatible with the order given by O.*

Proof. To check whether S satisfies Theorem 4.2.2 for (id, O), we need only check whether the intrinsic cyclic order on E is compatible with the one given by O. Thus, to check for S after acting by ϕ, it suffices to check this condition on an exhausting collection of ends of $\phi^{-1} \circ f : S \to K_F$, and $\phi^{-1} \cdot E$ is such a collection, by Lemma 4.6.2. □

Remark 4.6.5. The group $\mathrm{Out}(F)$ acts on the left on fatgraph realizations by precomposition by inverses. However, $\mathrm{Out}(F)$ acts on the left on tuples of ends by applying the automorphism itself. Therefore, it is *not* the case that $\mathrm{Out}(F)$ acts on pairs of a surface realization together with a fatgraph map into it. There is no contradiction here; only possible confusion.

4.6.1.4 Ends of subgroups

If G is finite index in F, then elements of G form a net in the Cayley graph of F. Equivalently, there are only finitely many (left) cosets of G in F, so for any word $w \in F$, there is a word v of (uniformly) bounded length so that $wv \in G$. Therefore, any right-infinite word in F may be written as a right-infinite word in G. In other words, while a priori the ends of G are a subset of the ends of F, actually the ends of G *are* the ends of F.

Therefore, an automorphism of G induces a continuous map on the ends of F, and a cyclic order on G induces a cyclic order on the ends of F.

4.6.1.5 Transfer of ends

We now return to transfer of fatgraphs, and we show how to determine the extremality of a particular transfer rotation quasimorphism on a chain. Let G be a finite index subgroup of F; by passing to a subgroup, we may assume that G is normal. Let $n = [F : G]$, and let $C \in B_1^H(F)$ be a chain with extremal fatgraph surface map $f : S \to K_F$. Let E be an exhausting set of ends for S, and recall that E is also a set of ends of G. Let (ϕ, O) be a surface realization of G.

Proposition 4.6.6. *In the above notation, $(\mathrm{rot}_{(\phi,O)})_F^G$ is extremal for C if and only if the cyclic order determined by (ϕ, O) on the ends of G is compatible with the intrinsic cyclic order on $\psi_i \cdot E$, for all $\psi_1, \ldots, \psi_n \in F/G$, where $\cdot$ is the conjugation action.*

Proof. First, let us consider another way to obtain the fatgraph representative of the total preimage $f' : S' \to K_G$. Let $C(F)$ be the Cayley graph of F; this is the universal cover of K_F. The total preimage $\widetilde{S}$ of the fatgraph S in K_F is a collection of infinite trees. Let T be a fundamental domain for $\widetilde{S}$ under the action of F, and let $T_1, \ldots, T_n$ be the translates of T under the conjugation action of F/G. Then, as $\widetilde{S}$ also covers S' under the action of G, the union of the T_i is a fundamental domain for $\widetilde{S}$ under G. By Lemma 4.6.2, each collection $\psi_i \cdot E$ is exhausting for T_i, and by Theorem 4.2.2 and Lemma 4.6.4, the compatibility of the order (ϕ, O) on each $\psi_i \cdot E$ with the intrinsic order is equivalent to each T_i (and thus the union S') immersing in K_G. Lemma 4.6.1 gives the equivalence of this with $(\mathrm{rot}_{(\phi,O)})_F^G$ being extremal for C. □

4.6.2 Examples

In this section, we consider two examples, one which gives a comprehensible picture of Proposition 4.6.6 in action, and another which reviews the examples of Section 4.4.3 from this new perspective. We also show some pictures describing the set of three-tuples of ends which are invariant under some finite-order automorphisms induced by conjugation in a larger group.

Example 4.6.7. Let $F = \langle x, y\rangle$ and let $C = xyy+YYY+XY$. Let $G = \langle xx, xy, yy\rangle$, where we denote the generators by a, b, and c, respectively. Then $F/G = \langle [x]\rangle$, and F/G acts on G by conjugation by x, which, in terms of generators for G, gives

$$[x] : \begin{array}{c} a \\ b \\ c \end{array} \mapsto \begin{array}{c} a \\ acB \\ bcB \end{array}.$$

Figure 4.7 shows an extremal fatgraph S for C and how to take the total preimage S' of this fatgraph in the cover corresponding to G. The total preimage is shown twice, on the left for clarity, and on the right to indicate how to lift from the $K(F, 1)$ to the $K(G, 1)$ (the bold section indicates one of the preimages, which is half of S').

Next, we rewrite the fatgraph S' with labels in G, as shown in Figure 4.8. This is most easily accomplished in the following way: first, decide which vertex to map to the basepoint. Here, we choose the lower right; note this means the upper right will also map to the basepoint, since they differ by the path yy, which is in G. Next, drag the vertices of the fatgraph by pulling apart the edges, so all vertices now map to the basepoint. Here we have pulled apart the top and bottom edges. All the loops in what is left over must be in G, so we label them appropriately, and fold, obtaining the total preimage, labeled in G, on the right.

Now let us see how this corresponds to the ends of S. Figure 4.9 shows a maximal tree in S, at bottom. The square vertex gives the basepoint, so this tree indicates the following exhausting set of

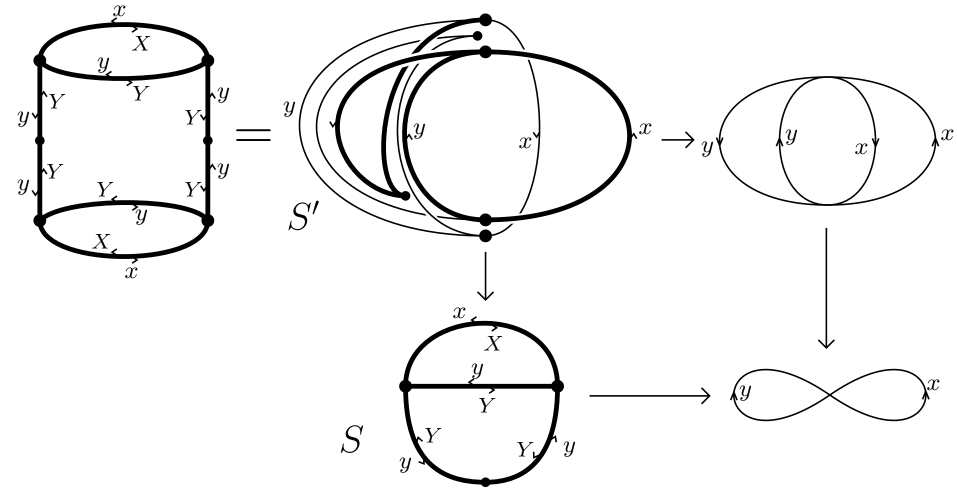

Figure 4.7: Taking the total preimage of a fatgraph S in a cover corresponding to the subgroup $\langle xx, xy, yy\rangle$. The vertical maps are covers, and most of the labels on S' in the middle are omitted for clarity. See Example 4.6.7.

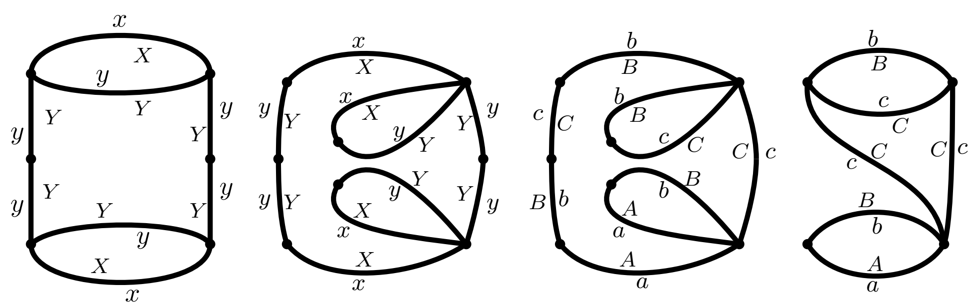

Figure 4.8: Rewriting labels on a fatgraph as generators in G, and folding. For example, for the outside boundary, $yyxyyx = (yy)(xy)(yy)(YX)(xx) = cbcBa$, and on the inside, $xY = (xy)(YY) = bC$.

ends for S (starting going up and proceeding counterclockwise):

$$[(xY),\, (yX),\, (y),\, (Y)],$$

where parentheses denote infinite power, so $x(y) = x(y)^\infty$. If we rewrite these ends in G, we get

$$E = [(bC),\, (cB),\, (c),\, (C)].$$

And acting by $[x] \in F/G$,

$$[x] \cdot E = [(acBbCB),\, (bcBbCA),\, (bcB),\, (bCB)] = [(aB),\, (bA),\, b(c),\, b(C)].$$

The ends E and $[x] \cdot E$ together contain vertices with the same labelings as all the vertices of S', so checking the compatibility of the intrinsic orders with an order on G will determine whether S' immerses in the surface realization. It is not quite true that we can lift the maximal tree in S to

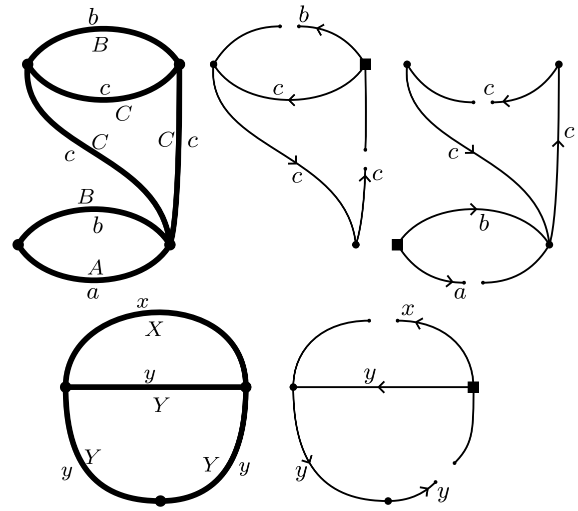

Figure 4.9: A maximal tree, which gives ends of S (bottom), and two trees containing all the vertices in S', as obtained from E and $[x] \cdot E$ as ends in G. These trees are almost obtained by lifting the tree in S and acting by $[x]$.

one in S' and "act by $[x]$" to get two trees; this is not exactly well-defined, since it is unclear what to do with the half-edges, and folding can rearrange things. However, after rewriting the ends of S to obtain E and $[x] \cdot E$, we can use these ends to give two trees in S' which together contain all the vertices, as shown. As is evident, we are "almost" just lifting the tree and acting by $[x]$.

In this case, the order $[c, C, b, B, a, A]$ is compatible with both E and $[x] \cdot E$, so S' immerses in this surface realization, and the transfer of the rotation quasimorphism will be extremal for C. For this example, though, transfer is more pedagogical than necessary, since by checking the orders on the vertices of S, we see that it immerses in the basic fatgraph realization given by the order $[X, x, y, Y]$!

Example 4.6.8. Here we consider one of the infinite family of vimersions in Section 4.4.3 from the perspective of Proposition 4.6.6.

Let $F = \langle x, y \rangle$, and $G = \langle x, yxY, y^2 \rangle$, where we denote $a = x$, $b = yxY$, and $c = y^2$. Note that $F/G = \mathbb{Z}/2\mathbb{Z} \cong \langle [y] \rangle$, and F/G acts on G by conjugation by y. In generators,

$$[y] : \begin{array}{ccc} a & & b \\ b & \mapsto & caC \\ c & & c \end{array} .$$

Let $C = x^2 + y + yXY^2X$. An extremal fatgraph for C is shown in Figure 4.10.

Figure 4.10 shows a maximal tree in an extremal fatgraph for C, which lifts to a fundamental domain in the universal cover for the action of F. Let the square indicate the basepoint in the universal cover; then a set of ends of the fatgraph which exits the leaves gives an exhausting set of ends for the fatgraph. Such a set of ends is given by generators of the fundamental group of the

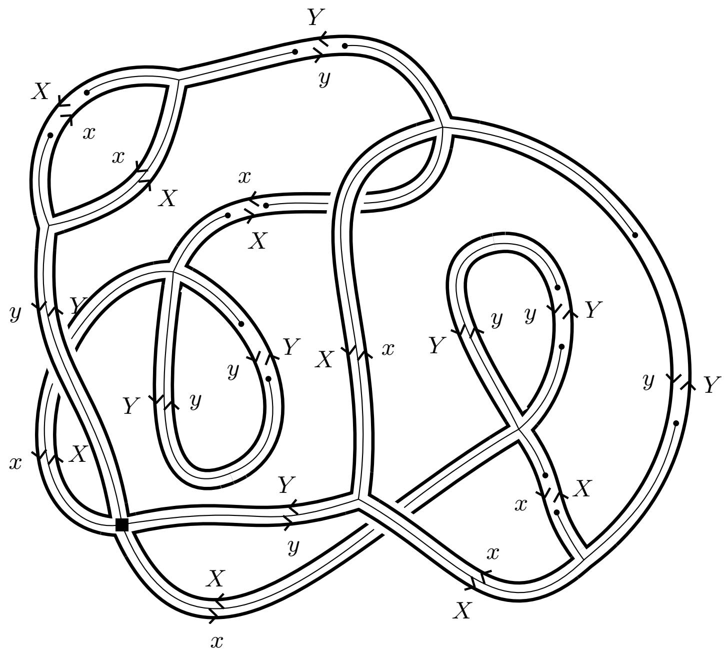

Figure 4.10: An extremal fatgraph for $C = x^2 + y + yXY^2X$ and a maximal tree.

fatgraph (starting going down from the basepoint, and proceeding counterclockwise):

$$E = [(xxxY),\, (xYYX),\, (xyyX),\, (yXYXY),\, (yXXX),\, (yxxx),\, (yxyxY),$$
$$(yxYxy),\, (YXyXY),\, (YXXy),\, (YxxY),\, (XYYx),\, (Xyyx),\, (XXXY)],$$

where the parentheses, as before, denote infinite powers, so $x(Y) = x(Y)^\infty$.

As ends of G, we have

$$E = [(aaaCbbb),\, a(C),\, a(c),\, B(AACBB),\, (BBBcAAA),$$
$$(bbbcaaa),\, bc(aabbc),\, (bacaCbc),\, (CBcACAB),\, C(B),$$
$$(Cbb),\, A(C),\, A(c),\, (AAACBBB)],$$

and acting by F/G gives

$$[y] \cdot E = [(bbbaaaC),\, b(C),\, b(c),\, cAC(BBAAC),\, (cAAABBB),$$
$$(caaabbb),\, ca(bbcaa),\, (caCbcba),\, (ABCBcAC),\, (A),$$
$$(aaC),\, B(C),\, B(c),\, (BBBAAAC)].$$

Note that both of these are positively cyclically ordered in the order on G given by $[C, c, A, a, B, b]$. Observant readers will note that this order is the inverse of the reverse of the cyclic order from Section 4.4.3; this is due to an inversion of conventions.

Example 4.6.9. Lemma 4.6.1 shows that in order to determine if a transfer of F to a finite index subgroup G with surface realization (ϕ, O) is extremal for a chain C with extremal fatgraph S, it suffices to find an exhausting set of ends E for S, write E as ends in G, and compare the cyclic order given by (ϕ, O) with the intrinsic cyclic orders on $\psi_i \cdot E$ for all $\psi_i \in F/G$. If they all agree, then $(\mathrm{rot}_{(\phi,O)})_F^G$ is extremal for C.

If L is any cyclically ordered set $[x_1, \ldots, x_k]$, then the entire cyclic list of elements is not necessary in order to specify the cyclic order: it suffices to give enough ordered triples to reconstruct the order. For example, the ordered triples

$$\{\, [x_1, x_2, x_3],\ [x_1, x_3, x_4],\ \ldots,\ [x_1, x_{k-1}, x_k] \,\}$$

determine the order on L.

Therefore, breaking E into a similar set of ordered triples, we observe that it suffices for all the triples to be invariantly ordered after acting by all $\psi_i \in F/G$, so one way to find an invariantly ordered set of ends is to build it out of invariantly ordered triples. From this perspective, it is clearly useful to understand the set of all invariantly ordered triples of ends (for a given finite index subgroup and cyclic order), and in fact we can produce pictures of these sets.

A basic cyclic order on a free group F gives a planar embedding of the Cayley graph in a disk, say by making all the angles equal and all edges the same length in the Poincaré disk model of $\mathbb{H}^2$. Thus, a right-infinite word gives a sequence of points in the interior of the disk which converge to a point on the boundary S^1. This maps the Cantor set of ends of F continuously into S^1; denote this map $I : \partial F \to S^1$. The map I is not injective; in fact it is two-to-one on a countable number of points, but it certainly suffices to draw pictures. After acting on the basic cyclic order by an automorphism of F, this gives a way to draw ∂F in S^1 using any cyclic order. Figure 4.6.9 (A) illustrates the map I with the order $[a, b, A, B]$. Plotted are $I((ab)^\infty)$, a point, and $I(bA*)$, where $bA*$ denotes the set of all ends which start with bA; this is a cylinder set of the Cantor set, and it maps to an interval in S^1, as shown.

Now let G be a finite index subgroup of F, and suppose we have a cyclic order (ϕ, O) on G. Then the map

$$I \times I \times I : \partial G \times \partial G \times \partial G \to S^1 \times S^1 \times S^1$$

sends triples of ends of G into the 3-torus $\mathbb{T}^3$. In practice, of course, we cannot draw $\mathbb{T}^3$, but we can identify S^1 with the unit interval $[0, 1]$ and draw $[0, 1]^3$, which gives $\mathbb{T}^3$ when opposite faces are identified.

Let $\mathcal{E}^G_{(\phi,O)}$ be the set of triples of ends of G which are invariantly ordered by (ϕ, O) under all of F/G. The group G acts on $\mathcal{E}^G_{(\phi,O)}$ on the left by conjugation (really, by pre-pending), so in order to find a set of coset representatives of $\mathcal{E}^G_{(\phi,O)}$ under this action, we may assume that the first letters of a given triple are distinct.

Because F/G acts continuously on ∂G, a prefix of an end determines a prefix of its image. This means we can determine the cyclic order on the image of a triple using only a triple of finite prefixes of the ends. The lengths of the prefixes necessary to determine the cyclic order depends on the triple and on the action of F/G, but this can all be determined.

Therefore, to obtain a set of coset representatives for $\mathcal{E}^G_{(\phi,O)}$ under the action of F/G, it suffices to do the following:

1. List all triples of distinct generators of G, each of which gives a product of three cylinder sets of ends. From now on, we will write "cylinder set" to mean a product of three cylinder sets, i.e., a basic open set in the product of three Cantor sets.

2. For each triple, subdivide the cylinder set into smaller cylinder sets (append all possible combinations of letters to the generators to obtain many more triples) until the cyclic order on each one is determined, even after acting by F/G.

3. For each triple, check whether its order in (ϕ, O) in invariant under the action of F/G.

4. Each invariant triple gives a cylinder set, which after acting by I gives a solid box in $\mathbb{T}^3$ of invariantly ordered triples.

Figure 4.11 shows examples of these methods. In all cases, "the invariant triples" means coset representatives for the set $\mathcal{E}^G_{(\phi,O)}$ of all invariant triples, as described above. In (A), we see how to map ∂F to S^1 using the cyclic order $[x, y, X, Y]$. As illustrated, the end $(ab)^\infty$ maps to the point 0.44, and the cylinder set given by the prefix maps to the interval $[0, 3]$. In (B), we show the triples of ends of $G = \langle xx, xy, yy\rangle = \langle a, b, c\rangle$ invariant in the order $[c, C, b, B, a, A]$ under the conjugation action of $F/G = \langle [x]\rangle$, as in Example 4.6.7. Axes are shown in the background for reference. Recall that this is a picture of $\mathbb{T}^3$, so all opposite faces of the cube $[0, 1]^3$ are identified.

Recall the definition of $G_3 = \langle x, yxY, y^2xY^2, y^3xY^3, y^4\rangle$, where we denote the generators by a, b, c, d, e, respectively, and we give it the cyclic order determined by the order

$$O = [E, e, D, d, C, c, B, b, A, a].$$

This subgroup and order makes an appearance in the vimersion example in Section 4.4.3 and (at a lower degree) Example 4.6.7, though note we have taken the inverse reverse of it again because it makes the picture slightly nicer. Figure 4.11 (C) shows the triples in G_3 invariant in the order under the conjugation action of $[y]$. In this case, we have removed one-third of the triples so that it is possible to see inside.

We can also reinterpret this as a picture in F: each invariant cylinder set in G determines a cylinder set in F by simply applying the inclusion map. Using the order $[Y, y, X, x]$ on F, we plot

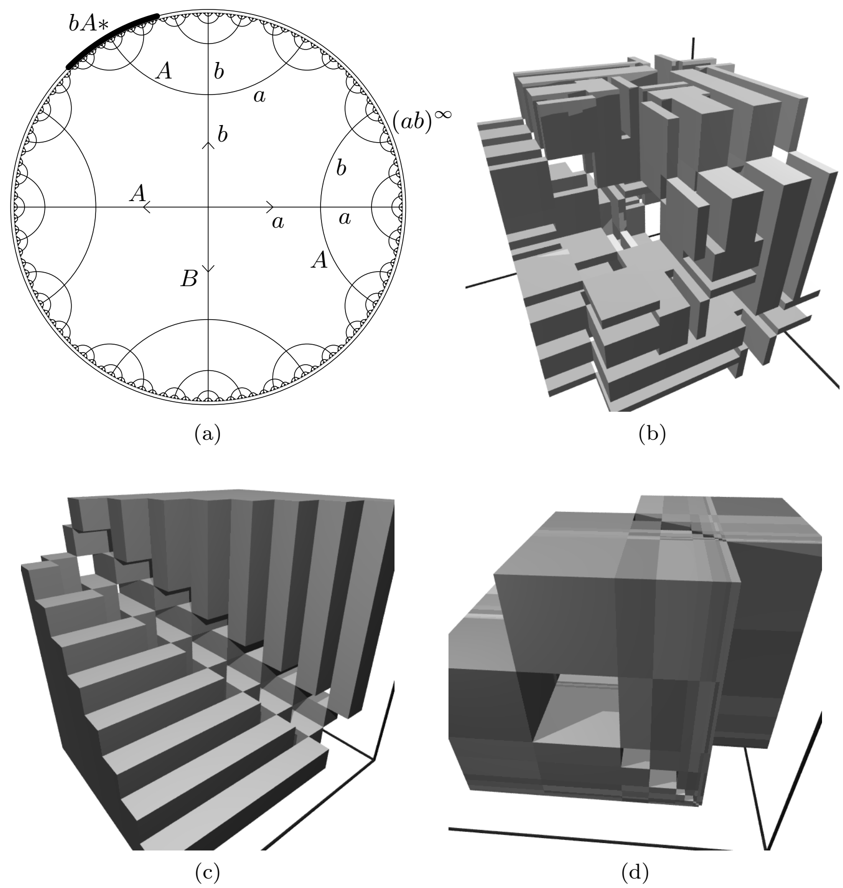

Figure 4.11: Mapping the Cayley graph of F into S^1, and pictures of invariant triples of ends, as described in Example 4.6.9.

these cylinder sets in (D) as sets in $(\partial F)^3$ (mapped to $\mathbb{T}^3$). This gives a picture of the triples of ends of F which, as ends in G, are invariantly ordered under F/G. For this picture, we have not removed the one-third of the triples that we did for (C). In addition, each cylinder set has been given a slightly different random grayscale value, so it it possible to see the entire set of invariant triples as a union of cylinder sets.

Example 4.6.10. Example 4.6.9 shows that we can create very explicit pictures showing which 3-tuples of ends are invariant under the action of F/G for finite-index G. However, a more interesting question is that of invariant 4-tuples. For, suppose that we have a tuple T of four cylinder sets of ends for which there is a finite index $G \subseteq F$ and a fatgraph realization (ϕ, O) of G such that for any four tuple of ends $t \in T$, the cyclic order on $\psi \cdot t$ is compatible with that given by (ϕ, O). Now pick some $t \in T$ so that t is made of infinite powers. Then t gives an exhausting set of ends for a commutator or a thrice-punctured sphere, and Proposition 4.6.6 shows that $(\mathrm{rot}_{(\phi,O)})_F^G$ is extremal for this commutator or thrice-punctured sphere.

In a sense, a picture of 4-tuples invariant as above gives a picture of all the commutators and thrice-punctured spheres for which the transfer of rot is extremal(!). Of course, we cannot make

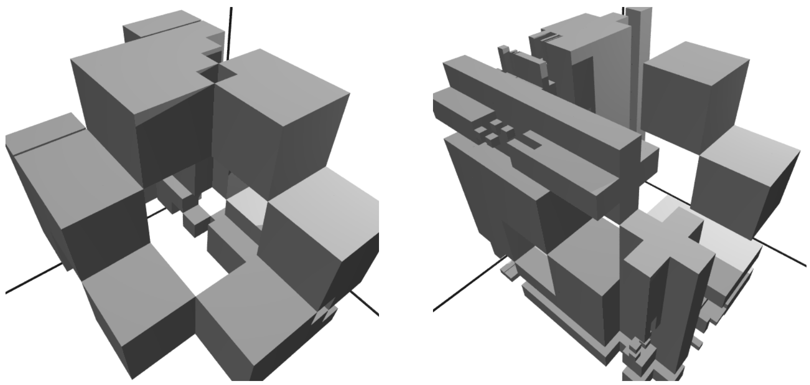

(a) (b)

Figure 4.12: Invariant 4-tuples of ends, as described in Example 4.6.10.

4-dimensional pictures. It will have to suffice to take 3-dimensional slices of the set of all invariant 4-tuples. Figure 4.12 shows two examples of this. Let $G_1 = \langle x,\, yxY,\, yy\rangle$ and $G_2 = \langle xx, yy, xy\rangle$. The picture on the left shows all 3-tuples (t_0, t_1, t_2) such that there is some basic fatgraph realization (id, O) of G_1 such that $(x^\infty, t_0, t_1, t_2)$ and its translate under the action of F/G_1 are correctly ordered. Thus, this picture shows all commutators and thrice-punctured spheres (with one end close to x^∞) for which there is some basic fatgraph realization of G_1 which induces an extremal transfer rotation quasimorphism. The picture on the right is the same, except that the subgroup G_1 is replaced with G_2.

4.7 Limit transfers

Section 4.6.1 shows that if we have a chain C and an extremal fatgraph S for C with exhausting set of ends E, then if we can find a finite index subgroup G and a cyclic order on G so that the action of F/G on E leaves the cyclic order on E invariant, then we obtain an extremal quasimorphism for C. Though we have given some examples which show this process in action, it is a priori rather difficult to come up with the subgroup G and cyclic order on G. In this section, we show that if we are given C, where C is in a large class of commutators and thrice-punctured spheres, we can find a family of finite index subgroups G_i and basic cyclic orders O_i on the G_i so that $(\mathrm{rot}_{(\mathrm{id},O_i)})_F^{G_i}$ is almost extremal for C, and while there is no well-defined limit of the G_i, there *is* a well defined limit of the $(\mathrm{rot}_{(\mathrm{id},O_i)})_F^{G_i}$, and this quasimorphism is extremal for C.

4.7.1 λ-orders

Let $S = \{s_0, s_1, \ldots, s_k\}$ be a finte set, and let λ be a real number which is rationally independent from π. The λ-*order* on S is the cyclic order obtained by the injection $S \to S^1$ defined by $s_i \mapsto \lambda i (\mathrm{mod} 2\pi)$.

Lemma 4.7.1. *Let $S = \{s_0, s_1, \ldots, s_k\}$ be a finte set, and let $n < i < j < m$ be four indices such that $i - n \neq m - j$. Then for any cyclic order O on $S' = \{s_n, s_i, s_j, s_m\}$, there is a λ such that the λ-order on S restricts to O.*

Proof. Note that the cyclic order given by a λ-order to a tuple is independent under shifting all the indices by the same amount. Thus we may assume without loss of generality that $n = 0$. We need to show that there is a λ which realizes each of the 6 cyclic orders on S':

$$\begin{array}{cc} [s_0, s_i, s_j, s_m] & [s_0, s_m, s_j, s_i] \\ [s_0, s_m, s_i, s_j] & [s_0, s_j, s_i, s_m] \\ [s_0, s_i, s_m, s_j] & [s_0, s_j, s_m, s_i] \end{array} .$$

Each row gives a cyclic order and its reverse. If a λ-order restricts to one of the cyclic orders above on S', then the order for $-\lambda$ restricts to the reverse, so it suffices for us to exhibit a λ which gives one of the orders in each row. It is best to imagine the four points $0, i, j, m$ in S^1 being dialated by changing λ; each of m, j, and i wraps around as they pass 2π. Starting with λ small, we get the first row; as m passes 2π the first time, we get the second row; finally, either m passes i, or j wraps around 2π, and we obtain the third order. We now give the precise values of λ.

Choose any λ such that $0 < \lambda < 2\pi/m$. Then note $\lambda m < 2\pi$, so this λ gives the first row.

Next, choose any λ so that $2\pi/m < \lambda < \min(2\pi/j, 2\pi/(m-i))$. Then we identify S' with the four (cyclically ordered) points in $[0, \lambda m - 2\pi, \lambda i, \lambda j]$, and this λ gives the second row.

Finally, the third row. There are two possibilities. If $m - i > j$, then $\frac{2\pi}{j} > \frac{2\pi}{m-i}$, so it is possible to select $\lambda \in \left(\frac{2\pi}{m-i}, \min(\frac{2\pi}{j}, \frac{2\pi}{m-j})\right)$, which gives the cyclic order $[s_0, s_i, s_m, s_n]$ and the third row. If $m - i < j$, then $\frac{2\pi}{j} < \frac{2\pi}{m-i}$, and we can select $\lambda \in \left(\frac{2\pi}{j}, \min(\frac{2\pi}{i}, \frac{2\pi}{m-i})\right)$, giving the cyclic order $[s_0, s_j, s_m, s_i]$ and the third row. By the assumption in the lemma that $i - n \neq m - j$ (recall $n = 0$), we needn't consider the case that $m - i = j$. $\square$

Remark 4.7.2. The assumption that $i - n \neq m - j$ is necessary: if we have equality, there is no λ-order which restricts to either of the orders in the third row.

4.7.2 Limit transfers

Let $G_n = \langle x,\, yxY,\, y^2xY^2,\, \ldots,\, y^nxY^n,\, y^{n+1}\rangle$, as in Section 4.4.3, and as before, denote y^ixY^i by a_i and y^{n+1} by b. Then $F/G_n \cong \mathbb{Z}/(n+1)\mathbb{Z}$ acts on G_n by conjugation by powers of y, and F/G is generated by ϕ, where $\phi(g) = ygY$. In the new generators, we have $\phi(a_i) = a_{i+1}$ if $i < n$, $\phi(a_n) = ba_0B$, and $\phi(b) = b$.

If $\lambda \in \mathbb{R}$ is rationally independent from 2π, and $\epsilon = \pm 1$, then a (λ, ϵ)-*order*, a variation of a λ-order, on G_n is given as a basic cyclic order on semigroup generators by sending $a_i \mapsto \lambda i (\bmod 2\pi) \in S^1$ (and $b \mapsto \lambda(n+1)(\bmod 2\pi)$). The inverse A_i is placed in the cyclic order immediately after a_0 if $\epsilon = 1$, and immediately preceding if $\epsilon = -1$, and similarly for B.

Proposition 4.7.3. *Let $\{G_n\}_{n>0}$ be as above, and let λ rationally independent from 2π and $\epsilon \in \{\pm 1\}$ be fixed. Give G_n the (λ, ϵ)-order, and let rot_n be the sequence of rotation quasimorphisms on the G_n given by these orders. The sequence of transfers $(\mathrm{rot}_n)_F^G$ then converges in the weak-* topology on $Q(F)$ to a quasimorphism of defect 1.*

The proof consists of some technical lemmas. In the following, let $C \in B_1^H(F)$, and let S be an extremal surface for C, thought of as a fatgraph. Let $\mathcal{T}$ be a set of triples of ends which determine the cyclic order on a maximal tree in S (see Example 4.6.9). Let O_n be the (λ, ϵ) order on G_n as described in Proposition 4.7.3.

Lemma 4.7.4.

$$(\mathrm{rot}_n)_F^G(C) = \frac{1}{2(n+1)} \sum_{t \in \mathcal{T}} \sum_{\phi \in F/G_n} O_n(\phi \cdot t)$$

Proof. Give the surface realization of G_n corresponding to O_n a hyperbolic structure. The area form here is then the coboundary of the rotation quasimorphism rot_n (see [8],§4.2 or [7]). Let S' be the total preimage of an extremal surface for C, and give S' the structure of a pleated surface. Then

$$(\mathrm{rot}_n)_F^G(C) = \frac{1}{n+1} \frac{\mathrm{area}(S')}{2\pi},$$

where the area of S' is the signed area computed as a pleated surface. Therefore, it suffices to show that

$$\frac{\mathrm{area}(S')}{2\pi} = \frac{1}{2} \sum_{t \in \mathcal{T}} \sum_{\phi \in F/G} O_n(\phi \cdot t).$$

Consider the decomposition of S' as a pleated surface. Build the graph G dual to the triangles; we observe this graph is a deformation retraction of S', and it has the same topological type and ends as the (fat)graph which is the total preimage of the extremal surface fatgraph for C. There is one vertex of valence three in G for each triangle, and whether the triangle is positive or negative in the pleated surface depends on the cyclic order given by O_n to this vertex. Note that the cyclic orders on these vertices are exactly given by the cyclic orders on the conjugates of the triples in $\mathcal{T}$, so the sum on the right gives a signed count of the triangles in the pleated surface S'. The triangles are ideal, and ideal triangles have area π, so we have proved the claim and the lemma. $\square$

Lemma 4.7.5. *Let $T = (t_1, t_2, t_3)$ be a triple of ends in F. Let G be any finite index normal subgroup of F, and O a cyclic order on G determining a surface realization. Then for any $w \in F$,*

$$\frac{1}{[F:G]} \sum_{\phi \in F/G} O(\phi \cdot (t_1, t_2, t_3)) = \frac{1}{[F:G]} \sum_{\phi \in F/G} O(\phi \cdot (wt_1, wt_2, wt_3)).$$

Proof. Since F acts by conjugation on F/G, the two sums are the same except re-indexed by multiplying everything in F/G by the element $wG \in F/G$. $\square$

Lemma 4.7.6. *For all $(t_1, t_2, t_3) \in \mathcal{T}$, the limit*

$$\lim_{n\to\infty} \frac{1}{n+1} \sum_{\phi \in F/G_n} O_n(\phi \cdot (t_1, t_2, t_3))$$

exists.

Proof. There are a few cases. In the first, case, assume that none of t_1, t_2, t_3 are $y^{\pm\infty}$. By Lemma 4.7.5, we may assume that the ends begin with distinct generators in G_n — say these generators are a_i, a_j, and a_k. Let $K = \max(i,j,k) - \min(i,j,k)$. In this case, acting by F/G_n is just a rigid rotation by λ, which clearly preserves order, except for possibly K elements of F/G_n. Let $e \in \{\pm 1\}$ denote this dominant order. The limit is then $e\left[1 - \frac{K}{n+1}\right] \to e$.

Now suppose that one of t_1, t_2, t_3 is $y^{\pm\infty}$, say $t_1 = y^{\infty}$. Again we may assume that the other ends begin with distinct generators a_i and a_j, and without loss of generality, $j > i$. Then acting by F/G_n shifts t_1 relative to a_i and a_j by a shift of λ. Since λ is rationally independent from 2π, this action is ergodic, and the limit is then $\pm\lambda(j-i)(\mathrm{mod} 2\pi)$, with the $\pm$ depending on the intrinsic order (t_1, t_2, t_3).

If two of the ends are $\pm y^{\pm\infty}$, then the shift action of F/G_n keeps the other end always to the same side, and the limit is ± 1. $\square$

Lemma 4.7.7. *For all $n > 0$, $(\mathrm{rot}_n)_F^G(x) = (\mathrm{rot}_n)_F^G(y) = 0$.*

Proof. The total preimage of x is $\sum_{i=0}^n a_n$, and the total preimage of y is $(n+1)b$. In both cases, all the words in the total preimage are simply powers of generators, on which rot_n is zero. For example, we may apply the formula in [8] §4.2. $\square$

Proof of Proposition 4.7.3. First observe that combining Lemma 4.7.4 and Lemma 4.7.6 implies that $\lim_{n\to\infty}(\mathrm{rot}_n)_F^G(C)$ converges for all $C \in B_1^H(F)$.

Now, given $w \in F$, there are $k, l \in \mathbb{Z}$ so that $w + x^k + y^l \in B_1^H(F)$. Then by Lemma 4.7.7,

$$\lim_{n\to\infty} (\mathrm{rot}_n)_F^G(w) = \lim_{n\to\infty} (\mathrm{rot}_n)_F^G(w + x^k + x^l),$$

and the limit on the right exists. The weak-* limit quasimorphism must have defect 1 because $(\mathrm{rot}_n)_F^G$ does, for all n. □

We denote by $\mathrm{rot}_\infty^{(\lambda,\epsilon)}$ the weak-* limit of $(\mathrm{rot}_n)_F^G$ when G_n is given the (λ, ϵ) order.

Proposition 4.7.8. *Let C be a chain of the form:*

$$[y^i x^{\pm 1} w_1 x^{\pm 1} y^j,\ y^k x^{\pm 1} w_2 x^{\pm 1} y^l],$$

or

$$g_1 + g_2 + g_1^{-1} g_2^{-1},$$

where

$$g_1 = y^i x^{\pm 1} w_1 x^{\pm 1} y^k, \quad \textit{and} \quad g_2 = y^l x^{\pm 1} w_2 x^{\pm 1} y^j,$$

and w_1 and w_2 are any words such that the above are locally reduced around the w_i, and i, j, k, and l are such that any cancellation above does not remove any $x^{\pm 1}$. Furthermore, subject to the condition in both cases that if we order i, j, k, and l and denote them by $m_1 < \cdots < m_4$, then $m_4 - m_3 \neq m_2 - m_1$.

Then there exists λ and ϵ such that $\mathrm{rot}_\infty^{(\lambda,\epsilon)}$ is extremal for C.

Proof. The rather contrived hypotheses are best explained by a schematic of a fatgraph which bounds C, which is shown in Figure 4.13, on the left. To get a commutator, we glue 1 to 3 and 2 to 4, and $w_1 = h_1 h_3^{-1}$, and to get a chain of the second type, we glue 1 to 2 and 3 to 4 and $w_1 = h_2 h_1^{-1}$. The

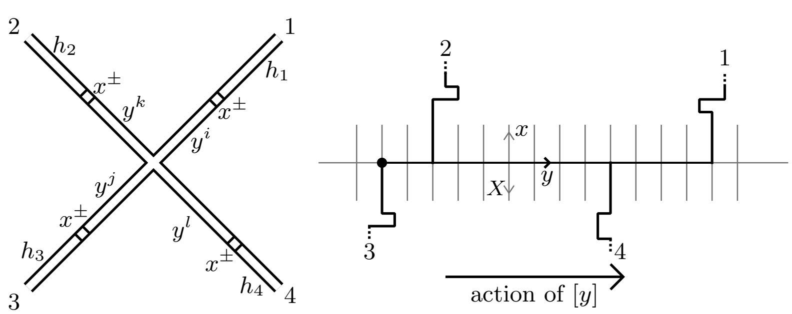

Figure 4.13: Illustrating a fatgraph bounding C and the action of F/G_n for the proof of Proposition 4.7.8.

hypotheses of the lemma serve to guarantee that a set of four exhausting ends for the fatgraph has the form shown on the right — when drawn in the Cayley graph of F, the core of the four ends (the common middle segment) is contained in the axis corresponding to translation by y. The action of

$[y] \in F/G_n$ serves to shift everything to the right. In the picture, the ends in the order $[3, 4, 1, 2]$ have first letters $[a_0^{-1}, a_9^{-1}, a_{13}, a_2]$ In general, the assumed structure of the chains guarantees that a set of exhausting ends (computed in G_n for n sufficiently large) is $[a_i \ldots, a_j \ldots, a_k \ldots, a_l \ldots]$, perhaps after a permutation. By Lemma 4.7.1, there exists λ such that the four ends are correctly oriented. Furthermore, λ does not depend on n for n sufficiently large, and the action of F/G_n is simply a rigid rotation and does not change the given order, except for a fixed finite number of elements. As there are 2 triples in $\mathcal{T}$ (as defined above) for C,

$$\lim_{n\to\infty} \frac{1}{2(n+1)} \sum_{t\in\mathcal{T}} \sum_{\phi\in F/G_n} O_n(\phi \cdot t) = 1,$$

So we compute $(\mathrm{rot}_n)_F^G(C)$ using Lemma 4.7.4, then, to see that

$$\mathrm{rot}_\infty^{(\lambda,\epsilon)}(C) = \lim_{n\to\infty} (\mathrm{rot}_n)_F^G(C) = 1.$$

Since the thrice-punctured sphere or once-punctured torus which bounds C has Euler characteristic -1, which gives $\mathrm{scl}(C) \leq 1/2$, the surface and the quasimorphism $\mathrm{rot}_\infty^{(\lambda,\epsilon)}$ certify each other on C, and both are extremal. $\square$

The hypothesis of Proposition 4.7.8 can be slightly weakened, at the cost of the cleanliness of λ-orders.

Lemma 4.7.9. *In Proposition 4.7.8, the hypothesis that $m_4 - m_3 \neq m_2 - m_1$ is not necessary, provided that the ends of the fatgraph bounding C are in configuration (A) or (B), as shown below (not (C)).*

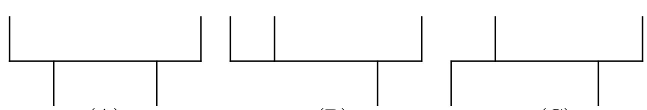

(A) (B) (C)

Proof. The idea is to replace the (λ, ϵ) order with two different λ orders, one for the generators a_i, and another for their inverses a_i^{-1}. The two orders use the same λ value, so the action of F/G_n still acts as a shift, and we still get the weak-* convergence of Proposition 4.7.3; however, we shift the λ order for the inverses. In this way, it is trivial to obtain any ordering for the configuration (B). The configuration (A) is slightly harder, but it suffices to observe that we can shift the two downward pointing ends freely such that they overlap in three different ways: if the ends are labeled 1, 2, 3, 4 clockwise from top left, then these shifts obtain orders $[1, 3, 4, 2]$, $[1, 3, 2, 4]$, and $[1, 2, 3, 4]$. Together with their negatives, this gives all 6 orders.

This technique fails for configuration (C) because the intervals we shift are the same size; here only 4 orders are possible with shifted λ orders. It is still possible in this case to select orders on G_n so that $(\mathrm{rot}_n)_F^G(C) \to 1$. Create an identification of the generators and inverses with points in S^1 as we shift: observe that as we shift the four ends, there are always two new generators, whose position in S^1 we may choose freely. However, the weak-* convergence of $(\mathrm{rot}_n)_F^G$ in this case is much more difficult. $\square$

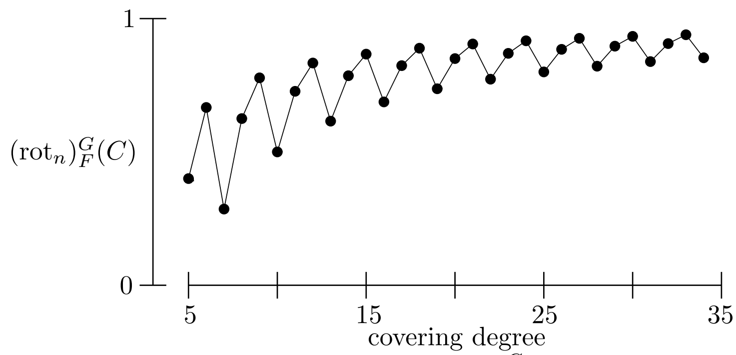

Figure 4.14: Exhibiting that the weak-* limit of $(\mathrm{rot}_n)_F^G$ is extremal for C. (See Example 4.7.10).

Example 4.7.10. Let $C = [yxxyXYxY^3,\ y^4xyXyxY^5]$. Then Proposition 4.7.8 applies to C with $i = 1$, $j = 3$, $k = 4$, $l = 5$, $w_1 = [x, y]$, and $w_2 = yXy$. We need to find a (λ, ϵ) order on G_n (for $n > 5$) such that $[a_1, a_5, a_3, a_4]$ is correctly ordered. We consult Lemma 4.7.1 to see that we may choose any λ with $\pi/2 < \lambda < 2\pi/3$ and λ rationally independent from π to achieve this. For example, $\lambda = 2$. The ϵ parameter does not matter because we do not have both a generator a_i and its inverse, so pick $\epsilon = 1$.

Figure 4.14 shows $(\mathrm{rot}_n)_F^G(C)$ for the range of covering degrees $(n + 1) = 5, \ldots, 35$ for $(\lambda, \epsilon) = (2, 1)$. Evidently, the weak-* limit $\mathrm{rot}_\infty^{(2,1)}$ is extremal for the given C, as guaranteed by Proposition 4.7.8.

Chapter 5

Dynamics and endomorphisms

The topologically minimal surfaces produced by scl hold the promise of producing surface groups. It is straightforward, given the fact that scl-extremal surfaces exist in free groups, that if $G = F_k *_{<w>} F_\ell$, and $< w >$ is cyclic, and $H_2(G, \mathbb{Z})$ is nontrivial, then G has a surface subgroup. We would like to extend this result to graphs of free groups. In particular, we would like to produce examples of surface subgroups of HNN extensions of free groups.

In this chapter, we study the dynamics of endomorphisms of free groups, and we use this to give conditions guaranteeing the existence of surface subgroups of HNN extensions by these endomorphisms.

Throughout this chapter, the ranks of our free groups are not particularly important, so we supress the subscript notation F_k in favor of the generic F. We will specify the rank if necessary.

This chapter contains material from the upcoming [16].

5.1 Traintracks

In this section, we provide some background on traintracks in free groups. This brief sketch is extracted from [16].

One of the problems when studying the dynamics of endomorphisms is that chains tend to get extremely large upon repeated application of an endomorphism. Therefore, it is completely intractible to experimentally compute scl for $\mathrm{scl}(\phi^k(C))$ for a chain C and endomorphism ϕ when k is larger than, say, 4. The issue is that the dimension of the vector space parameterizing surfaces bounding C grows. Another, related, issue is studying the value of a counting quasimorphism f on the images of a chain; that is, $f(\phi^k(C))$. In both of these situations, we can make the problem finite and bounded if we focus on subwords of a bounded length in C. This proves very useful.

Let F be a free group generated by the symmetric generating set S. Let A_ℓ be the set of reduced words in the generators of length ℓ. Let $\mathbb{R}[A_\ell]$ be the $\mathbb{R}$-vector space spanned by the set A_ℓ. We define $h : \mathbb{R}[A_\ell] \to H_1(F)$ on generators by taking $h(w)$ to be the (homology class of) the first letter of w. We define $s : \mathbb{R}[A_\ell] \to \mathbb{R}[A_{\ell-1}]$ and $t : \mathbb{R}[A_\ell] \to \mathbb{R}[A_{\ell-1}]$ (the *source* and *target* maps) on generators by setting $s(w)$ to be the subword of w which is the first $\ell - 1$ letters and setting $t(w)$ to be the subword of w which is the last $\ell - 1$ letters. Then we define $\partial : \mathbb{R}[A_\ell] \to \mathbb{R}[A_{\ell-1}]$ by $\partial(x) = s(x) - t(x)$.

Definition 5.1.1. The *verbal traintrack* T_ℓ is the directed graph with vertices $A_{\ell-1}$ and edges A_ℓ, where the source and target maps (note they are well-defined on A_ℓ) give the source and target of every edge.

Definition 5.1.2. Now define the *weight space* W_ℓ to be the kernel $W_\ell = \ker[\partial \oplus h]$. This is the space of homologically trivial weights on the verbal traintrack for F.

Recall $B_1(F)$ is the space of boundaries in the group homology of F, or more explicitly, the space of formal sums of words which are homology trivially. There is a natural map $\Psi_\ell : B_1(F) \to W_\ell$ defined by taking a chain to the sum of its subwords of length ℓ. It is necessary to cyclically reduce the words first.

Example 5.1.3. Let $C = abABAbaB$. Then

$$\Psi_2(C) = 1(ab) + 1(bA) + 1(AB) + 1(Ab) + 1(ba) + 1(aB) + 1(Ba),$$

where we have written the weight which is 1 on w and 0 elsewhere as (w). Now set $C' = abAB + AbaB$. Notice that $\Psi_2(C) = \Psi_2(C')$; that is, they have the same subwords of length 2. However, $\Psi_3(C) \neq \Psi_3(C')$.

The space W_ℓ is the natural space in which to work if we want to study the dynamics of an endomorphism on a certain length scale. See Lemma 5.2.2, which says that the action of an endomorphism can be linearized to act on W_ℓ under the hypothesis that the endomorphism is *expanding*.

It is also useful to work in the space W_ℓ when we want to compute scl, but we only care about the local properties of the chain, not the chain itself.

In all of these cases, we are interested in the space of *positive weights*, those weights which have nonnegative coefficients on every word. We will denote the space of positive weights by W_ℓ^+. One of the main results of [16], §3 is the following.

Definition 5.1.4. Define scl on W_ℓ^+ by

$$\mathrm{scl}(x) = \inf\{\mathrm{scl}(C) \mid C \in B_1(F),\ \Psi_\ell(C) = x\}.$$

Theorem 5.1.5 (Theorem 3.23 in [16] (paraphrased))**.** *For $x \in W_\ell^+$, the computation of* $\mathrm{scl}(x)$ *reduces to a linear programming problem. Furthermore, the infimum is achieved, and there is an extremal surface map $S \to K(F, 1)$ with $\Psi_\ell(\partial S) = x$ and* $\mathrm{scl}(\partial S)$ *realizing the infimum.*

5.2 Endomorphisms and HNN extensions

In this section, we study the dynamics of certain endomorphisms through their action on traintracks. This discussion is analogous to, and motivated by, the similar technology for free and surface group automorphisms. The general idea is that an endomorphism which is *expanding* (defined precisely below) gives an action on W_ℓ; this action reveals properties of the endomorphism.

5.2.1 Endomorphism actions on traintracks

Let F is a free group with symmetric generating set S. We denote the word length of $w \in F$ with respect to S by $|w|_S$. If x, y, z are words in F such that xyz is reduced, and $\phi : F \to F$ is an endomorphism, we will denote the letters in $\phi(xyz)$ which originate in $\phi(y)$ by $\phi(y; xyz)$.

Definition 5.2.1. An endomorphism $\phi : F \to F$ is *expanding* (with respect to S) at length scale $K > 0$ with dilation $C > 1$ if for all words s, w, t with $|w|_S \geq K$ and swt reduced, we have $|\phi(w; swt)|_S \geq C|w|_S$.

An endomorphism $\phi : F \to F$ is *virtually expanding* if there is some power n such that ϕ^n is expanding.

We remark that while being expanding depends on the generating set, being virtually expanding does not. Recall that if S' is another generating set, there exists λ so that if swt is reduced, then $(1/\lambda)|w|_{S'} \leq |w|_S \leq \lambda|w|_{S'}$, where $|w|_{S'}$ refers to the length of the section of swt written in the basis S' corresponding to w. Suppose that ϕ^n is expanding with respect to the generating set S at scale $K > 0$ with dilation $C > 1$. Then if swt is reduced and $|w|_{S'} \geq K\lambda$, we have $|w|_S \geq K$, so if

$$|\phi^{nm}(w; swt)|_{S'} \geq \frac{1}{\lambda}|\phi^{nm}(w; swt)|_S \geq \frac{C^m}{\lambda}|w|_S \geq \frac{C^m}{\lambda^2}|w|_{S'}$$

Choosing m large enough makes $C^m/\lambda^2 > 1$.

Lemma 5.2.2. *If ϕ is expanding with respect to the generating set S with factor $C > 2$ at scale $K > 0$, then for ℓ sufficiently large, that is, $\ell > C(K-2)/(C-2)$, ϕ gives a well-defined action ϕ_* on W_ℓ which commutes with any forgetful map $W_\ell \to W_{\ell-k}$ (for $\ell - k$ large enough).*

Proof. Let w be a weight in W_ℓ which is 1 on a single edge and 0 elsewhere. We will think of w as the word of length ℓ labeling the edge, and we define the image $\phi_*(w)$ as follows. Write $w = w_0w_1w_2$, where $|w_1|_S = K$ and $|w_0|_S, |w_1|_S \geq (\ell - K - 1)/2$; that is, w_1 is the middle K letters. Define $\phi_*(w)$ to be (the weight representing) the collection of words of length ℓ which begin in the image $\phi(w_1; w_0w_1w_2)$. We need to show that this is well-defined; in other words, we must show that regardless of s and t, if swt is reduced, then the set of subwords of $\phi(w; swt)$ originating in the middle K letters of w remains the same.

Since $\ell > C(K-2)/(C-2)$, we have $2\ell/(\ell-k-2) < C$. Now define m by setting $\ell = mK$, so $2mK/(k(m-1)-2) < C$. Because ϕ is expanding, the length of $\phi(w_2; w_0w_1w_2)$ which cannot be cancelled on the right is at least

$$\left\lfloor \frac{m-1}{2}K \right\rfloor C > \left(\frac{m-1}{2}K - 1\right) C > \left(\frac{m-1}{2}K - 1\right) \frac{2mK}{K(m-1)-2} = mK = \ell.$$

□

Remark 5.2.3. Recall that the verbal traintrack T_ℓ of length ℓ can be thought of as a simplicial complex. A reduced word in F gives a simplicial path in T_ℓ, and ϕ acts on such paths, although the action is not simplicial, because we must tighten the result of applying ϕ before it gives a path in T_ℓ. The condition of being expanding essentially guarantees that the average derivative of ϕ on any given path is bounded below (and is larger than 1). Choosing a large enough scale means that the action of ϕ is locally well-defined.

Remark 5.2.4. We could also ask that an endomorphism be expanding for certain words. We remark that an irreducible automorphism is expanding at a sufficiently large scale for subwords of its invariant lamination.

5.2.2 Norm-realizing surfaces in HNN extensions

If $\phi : F \to F$ is an endomorphism, the HNN extension of F by ϕ is the group $\langle F, t \,|\, txt^{-1} = \phi(x)\rangle$. We will denote this group by F_ϕ. Note that we obtain a $K(F_\phi, 1)$ by the quotient

$$X_\phi = K(F_\phi, 1) = X \times [0,1]/\{(x,0) = (\phi(x), 1)\},$$

Where $X = K(F,1)$ (we have supressed the subscript in X_k), and where here ϕ denotes the self-map of X induced by the endomorphism ϕ.

The homology of an HNN extension of F is given by the homology of F and the action of ϕ on it. In particular,

Lemma 5.2.5.

$$H_2(F_\phi) = H_2(X_\phi, \mathbb{R}) \cong \{[C] \in H_1(F) \,|\, [C - \phi(C)] = 0\},$$

further, the isomorphism is induced by the inclusion map on $X \times \{0\}$.

Proof. By Mayer-Vietoris. □

There is a norm (The Gromov-Thurston norm) on $H_2(F_\phi)$, given by

$$\|\alpha\| = \inf_{[\Sigma] = n\alpha} -\chi(\Sigma)/n(\Sigma),$$

where the infimum is taken over all surfaces representing $n\alpha$. We will call a surface Σ with $[\Sigma] = \alpha$ *norm-realizing* if $-\chi(\Sigma)/n(\Sigma) = \|\alpha\|$. Note that there may not be a norm-realizing surface for a given homology class.

Lemma 5.2.6. *Let $\alpha \in H_2(F_\phi)$ be represented by the chain $C - \phi(C) \in B_1(F)$. Then*

$$\|\alpha\| = \lim_{k\to\infty} \frac{2}{k} \mathrm{scl}\left(C - \phi^k(C)\right).$$

Proof. First observe that the limit exists because the expression is (non-strictly) monotone decreasing and bounded below. We claim that if C' is another chain with $[C] = [C']$, then

$$\lim_{k\to\infty} \frac{2}{k} \mathrm{scl}\left(C - \phi^k(C)\right) = \lim_{k\to\infty} \frac{2}{k} \mathrm{scl}\left(C' - \phi^k(C')\right).$$

Since $[C - C'] = 0$, we can compute $s := \mathrm{scl}(C - C')$. Note that $\mathrm{scl}(\phi^k(C - C')) \leq s$ for all $k \geq 0$ because we can simply push forward any surface bounding $C - C'$. Therefore,

$$\begin{aligned}\frac{2}{k} \mathrm{scl}\left(C - \phi^k(C)\right) &= \frac{2}{k} \mathrm{scl}\left(C - C' + C' - \phi^k(C) + \phi^k(C') - \phi^k(C')\right) \\ &\leq \frac{2}{k}\left[s + \mathrm{scl}\left(C' - \phi^k(C')\right) + s\right].\end{aligned}$$

Taking $k \to \infty$ gives one inequality, and the other is obtained by exchanging the roles of C and C'.

Now we prove the claim in the lemma. By the equality above, we may assume that C consists of a single word. Take some surface Σ admissible for $C - \phi^k(C)$ with $\left|-\chi(\Sigma)/2n(\Sigma) - \mathrm{scl}(C - \phi^k(C))\right| < \epsilon$. We may assume for simplicity that Σ has exactly two boundary components, one mapping to nC and the other to $n\phi^k(C)$ (see [8] Chapter 2). By gluing these two boundaries together, which is possible in F_ϕ, we obtain a closed surface $\bar{\Sigma}$, and $\chi(\bar{\Sigma}) = \chi(\Sigma)$. Note that the cylinder which we glue in to attach the boundaries of Σ will wrap k times around the mapping torus of ϕ; that is, it crosses $X \times \{0\}$ a total of k times. By Lemma 5.2.5, then, we see that $[\bar{\Sigma}] = n(\Sigma)k\alpha \in H_2(F_\phi)$, so

$$\|\alpha\| \leq \frac{-\chi(\bar{\Sigma})}{n(\bar{\Sigma})k} \leq \frac{2}{k}\mathrm{scl}\left(C - \phi^k(C)\right) + \frac{2\epsilon}{k}.$$

Since ϵ is arbitrary for each k, taking $k \to \infty$ gives one inequality.

To prove the other direction, observe that the Mayer-Vietoris sequence from Lemma 5.2.5 says that every representative of α in $H_2(F_\phi)$ arises as the gluing of some surface bounding $C' - \phi(C')$ with $[C] = [C']$. Let us be given some surface Σ with $|-\chi(\Sigma)/n(\Sigma) - \|\alpha\|| < \epsilon$, and take the associated chain C'; note $n(\Sigma) = 1$. Then

$$\frac{2}{k}\mathrm{scl}\left(C' - \phi^k(C')\right) \leq 2\mathrm{scl}\left(C' - \phi(C')\right) \leq -\chi(\Sigma) \leq \|\alpha\| + \epsilon.$$

The expression on the left is monotone decreasing, and ϵ is arbitrary; this proves the inequality in the other direction and concludes the proof of the lemma. □

Lemma 5.2.7. *Let $\alpha \in H_2(F_\phi)$ be fixed by the action of the endomorphism $\phi : F \to F$. Then*

$$\|\alpha\| = \inf_{[C]=\alpha} 2\mathrm{scl}(C - \phi(C)),$$

where the infimum is taken over all $C \in B_1^H(F)$.

Proof. The proof is very similar to Lemma 5.2.6. To prove the inequality in one direction (that the norm is larger), we take some surface Σ so that $|-\chi(\Sigma)/n(\Sigma) - \|\alpha\|| < \epsilon$. Then by Mayer-Vietoris, Σ is obtained from a surface bounding $C - \phi(C)$ with $[C] = n(\Sigma)\alpha$. Therefore,

$$\mathrm{scl}\left(\frac{1}{n(\Sigma)}(C - \phi(C))\right) \leq \frac{-\chi(\Sigma)}{2n(\Sigma)}.$$

This holds for all ϵ, so the infimum on the right must be at most the norm $\frac{1}{2}\|\alpha\|$.

The other direction is essentially the reverse procedure: take some surface Σ which is very close to realizing the infimum of scl. Gluing Σ along its boundary $C - \phi(C)$ with $[C] = \alpha$ shows that the norm $\|\alpha\|$ must be less than $-\chi(\Sigma)$, and thus is at most the infimum. □

Lemma 5.2.8. *Let $C \in B^+$ with $[C] = \alpha \in H_2(F_\phi)$, such that $\phi_*(\alpha) = \alpha$. Then*

$$\|\alpha\| \geq 2 \inf_{[C]=\alpha} \mathrm{scl}(\Phi_\ell(C) - \Phi_\ell(\phi(C))),$$

where the scl *is the traintrack* scl.

Proof. This is essentially a corollary of Lemma 5.2.7; we merely recall that $\mathrm{scl}(\Phi_\ell(C) - \Phi_\ell(\phi(C))) \leq \mathrm{scl}(C - \phi(C))$. □

Lemma 5.2.9. *If $f : \Sigma \to K(F_\phi, 1)$ is a norm-realizing surface map, then $f_* : \pi_1(\Sigma) \to F_\phi$ is injective.*

Proof. Suppose that f_* is not injective, so there is an immersed loop γ in Σ whose image is nulhomotopic. Because surface groups are LERF ([29]), there is a finite cover $\widetilde{\Sigma}$ of Σ so that a lift $\widetilde{\gamma}$ of γ is embedded. Compress $\widetilde{\Sigma}$ along $\widetilde{\gamma}$ to obtain a surface with smaller genus representing a multiple of α. But this is a contradiction, since a finite cover of a norm-realizing surface is norm-realizing. □

We can use Lemma 5.2.8 to certify a norm-realizing surface, for if ϕ gives a well-defined action on W_ℓ, then the infimum on the right is a finite-dimensional problem, and in fact can be computed using a variation on the algorithm from Theorem 5.1.5. It is not necessarily the case that an extremal surface Σ for $\mathrm{scl}(\Phi_\ell(C) - \Phi_\ell(\phi(C)))$ will have boundary which is of the form $C - \phi(C)$. However, it might, and if it satisfies $-\chi(\Sigma)/2n(\Sigma) = \inf_{[C]=\alpha} \mathrm{scl}(\Phi_\ell(C) - \Phi_\ell(\phi(C)))$, then we know that Σ is norm-realizing (and injective).

5.3 Geometric endomorphisms

In this section, we describe *geometric* endomorphisms, which can be realized as surface map immersions. This property, which turns out to be generic, introduces structure we can exploit in the search for norm-realizing and injective surfaces. Much of the discussion benefits from the context of traintracks.

5.3.1 Tripods and cyclic orders

A *tripod* is a cyclic triple of generators, written with square brackets to emphasize the invariance under cyclic permutation, e.g., $[a, b, B]$. The *reverse* or *negative* of a tripod is the negative of the cyclic order; reverses will be denoted by minus signs, e.g., $-[a, b, B] = [a, B, b]$.

Lemma 5.3.1. *Let ϕ be an expanding endomorphism of F at scale $K = 1$. Then ϕ induces a well-defined map on tripods which commutes with negation.*

Proof. Given a tripod $T = [x, y, z]$, apply ϕ to T by applying ϕ to each of x, y, z and Stallings folding the resulting labeled tripod. The generator labels at the junction of the three legs then give the tripod $\phi(T)$. The expanding condition guarantees that this junction exists because the folding cannot pass the midpoint of each leg.

That this map commutes with negation is obvious from the definition. □

Tripods are the building blocks of cyclic orders. Recall that if S is a set and $A, B \subseteq S$, then cyclic orders on A and B are said to be *compatible* if they are simultaneously induced by a cyclic order on S.

Lemma 5.3.2 ([6], Lemma 2.37)**.** *Let S be a set. A cyclic order on all triples of distinct elements on S is compatible if and only if for every subset $Q \subseteq S$ with four elements, the cyclic order on triples of distinct elements of Q is compatible. In this case, these circular orderings are uniquely compatible, and determine a circular ordering on S.*

As a consequence, in order to define a cyclic order on a set S, it suffices to define a cyclic order on all tripods in S which is compatible on all sets of four elements. Such a cyclic order can be encoded by a map O from the set of tripods of S to $\{\pm 1\}$ which records the orientation sign of each tripod.

5.3.2 Geometric endomorphisms

If $\phi : F \to F$ is an endomorphism and (Σ, f) is a hyperbolic surface realization of F, then ϕ induces a homotopy class of map $\phi_* = f \circ \phi \circ f^{-1} : \Sigma \to \Sigma$, where f^{-1} denotes the homotopy inverse of f.

Definition 5.3.3. An endomorphism $\phi : F \to F$ is geometric with respect to the hyperbolic surface realization (Σ, f) if ϕ_* is homotopic to an immersion with geodesic boundary.

An endomorphism is *virtually geometric* if some power of it is geometric.

Recall the definition of the basic surface realization (Σ_O, f_O) induced by the cyclic order O on semigroup generators of F (Definition 2.5.9).

Proposition 5.3.4. *Let F have rank r, and let $\phi : F \to F$ be expanding at scale $K = 1$ with dilation $C > 1$. Then there is some $n \in \mathbb{Z}$ and a cyclic order O on the elements of S so that $1 \le n \le 4\binom{2r}{3}$ and ϕ^n is geometric with respect to the basic realization (Σ_O, f_O).*

In preparation for the proof, we need some discussion. Consider an ordered tuple (x_1, x_2, x_3, x_4) of four distinct generators of F. We will call such a thing a *quadpod*[1]. Though we will be interested in unordered sets of four generators, it is necessary to introduce an order for the purposes of exposition.

A quadpod can be visualized as a graph vertex with four outgoing incident labeled edges. The first element of the tuple is defined to be the label for the edge outgoing to the right. As long as ϕ is expanding at scale $K = 1$, We can apply ϕ to a quadpod by applying it to each edge and Stallings folding. Of course ϕ need not preserve the order. In fact, the image of a quadpod under ϕ need not be a quadpod. See Figure 5.1.

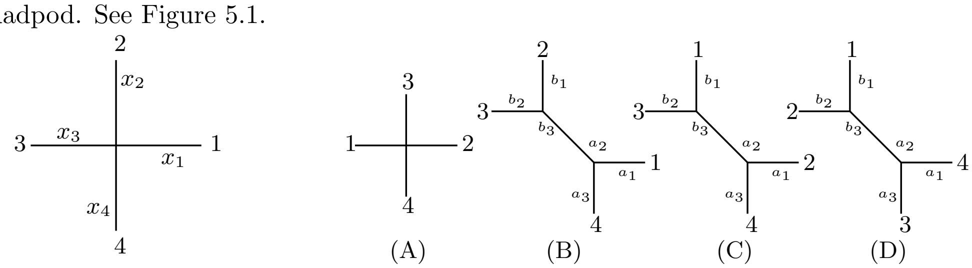

Figure 5.1: The possible results (right) of applying the endomorphism ϕ expanding at scale $K = 1$ to a quadpod (left). The numbers indicate how the edges fold. The letters x_i, a_i, and b_i indicate generators of F labeling the outgoing edges incident the vertices. The cyclic orders induced on the right by drawing the figures in the plane are meaningless.

We need some functions to record pieces of quadpods. Denote the set of all quadpods by Q, and let $q = (x_1, x_2, x_3, x_4) \in Q$. Recall that x_1 is outgoing to the right. We define R, U, L, D all functions from Q to T, the set of tripods, by

$$R(q) = [x_4, x_1, x_2],\ U(q) = [x_1, x_2, x_3],\ L(q) = [x_2, x_3, x_4],\ D(q) = [x_3, x_4, x_1].$$

[1]The better-sounding *tetrapod* unfortunately shares its first letter with tripod.

Note these are the right, top (up), left, and bottom (down) tripods in the quadpod q.

Lemma 5.3.5. *Let $q \in Q$ be a quadpod and let the endomorphism $\phi : F \to F$ be expanding at scale $K = 1$. If $\phi(q)$ is* not *a quadpod (as in Figure 5.1, (B), (C), and (D)), then one of the following holds:*

1. $\phi(R(q)) = \phi(D(q))$ *and* $\phi(L(q)) = \phi(U(q))$
2. $\phi(R(q)) = \phi(U(q))$ *and* $\phi(L(q)) = \phi(D(q))$
3. $\phi(R(q)) = -\phi(L(q))$ *and* $\phi(U(q)) = -\phi(D(q))$

Proof. The proof is basically by Figure 5.1. Because $\phi(q)$ is not a quadpod, there must be folding, so the result must look like (B), (C), or (D) in Figure 5.1. It is important to note that there is no "correct" or "positive" orientation given on the triples $\{a_1, a_2, a_3\}$ and $\{b_1, b_2, b_3\}$, even though it is necessary to draw them in the plane.

Consider case (B) from Figure 5.1, in which edge 1 folds with edge 4. In this case, we have

$$\begin{aligned} \phi(R(q)) &= \epsilon[a_1, a_2, a_3], & \phi(U(q)) &= \delta[b_1, b_2, b_3], \\ \phi(L(q)) &= \delta[b_1, b_2, b_3], & \phi(D(q)) &= \epsilon[a_1, a_2, a_3] \end{aligned}$$

where $\epsilon, \delta = \pm 1$. No matter what the orientation on the triples is (we do not know it), we do know that $\phi(R(q))$ and $\phi(D(q))$ are equal and given the same orientation, and similarly with $\phi(U(q))$ and $\phi(L(q))$. This gives case (1) of the lemma.

Now case (C). Here, we have

$$\begin{aligned} \phi(R(q)) &= \epsilon[a_3, a_2, a_1], & \phi(U(q)) &= \delta[b_1, b_3, b_2], \\ \phi(L(q)) &= \epsilon[a_1, a_2, a_3], & \phi(D(q)) &= \delta[b_2, b_3, b_1]. \end{aligned}$$

Being careful to remember the signs on the tripods, we see this is case (3) of the lemma.

And case (D). Here,

$$\begin{aligned} \phi(R(q)) &= \epsilon[b_3, b_1, b_2], & \phi(U(q)) &= \epsilon[b_1, b_2, b_3], \\ \phi(L(q)) &= \delta[a_2, a_3, a_1], & \phi(D(q)) &= \delta[a_3, a_1, a_2], \end{aligned}$$

which gives case (2) of the lemma. □

Proof of Proposition 5.3.4. Let S be the set of symmetric generators of F. Let T be the set of all tripods for rank r (i.e., tripods of elements of S) and Q the set of all quadpods. The endomorphism ϕ acts on T by Lemma 5.3.1. We let $\mathcal{L}_\phi \subseteq T$ be the limit set of ϕ, which is the set of tripods in the image of ϕ^n for all n. Note that ϕ acts as a permutation on $\mathcal{L}_\phi$. Similarly, we define $\mathbb{Q}_\phi$ to be the set of quadpods in the image of ϕ^n for all n.

Because ϕ acts as a permutation on $\mathbb{Q}_\phi$ and $\mathcal{L}_\phi$, it is possible to choose $n \leq 4\binom{2r}{3}$ such that ϕ^n acts as the identity on $\mathcal{L}_\phi$ and $\mathbb{Q}_\phi$ and $\phi^n(t) \in \mathcal{L}_\phi$ for all t. This will be the n such that ϕ^n is geometric.

Now choose any cyclic order on S by defining $O : T \to \{\pm 1\}$ such that O is compatible on all sets of four elements of S. The order O is arbitrary, and ϕ^n will almost certainly not preserve it.

Now define $\widehat{O} : T \to \{\pm 1\}$ by

$$\widehat{O}(t) = O(\phi^n(t)).$$

The function $\widehat{O}$ is preserved by ϕ^n, since by the definition of n,

$$\widehat{O}(\phi^n(t)) = O(\phi^n(\phi^n(t))) = O(\phi^n(t)) = \widehat{O}(t).$$

In addition, $\widehat{O}$ commutes with tripod negation because ϕ^n does. However, it is not a priori clear that $\widehat{O}$ defines a cyclic order on S. By Lemma 5.3.2, it suffices to check that $\widehat{O}$ is compatible on all sets of four elements of S.

To this end, let us be given some set $Q \subseteq S$ with four elements. Give these elements an arbitrary order and put them in the quadpod $q = (x_1, x_2, x_3, x_4)$. There are two general cases:

If $\phi^n(q)$ is a quadpod, then by construction, ϕ^n fixes q. In particular, it fixes all of the constituent tripods in q, so all of these tripods are in $\mathcal{L}_\phi$. But the order O, and thus $\widehat{O}$, is compatible on these tripods, so there is an order on Q which induces all the orders of the tripods.

If $\phi^n(q)$ is *not* a quadpod, then we apply Lemma 5.3.5 with the endomorphism ϕ^n, as follows. Now we case out what happens. For example, suppose that condition (1) holds with $\widehat{O}(R(q)) = +1$ and $\widehat{O}(L(q)) = -1$. Then $\widehat{O}(D(q)) = +1$ and $\widehat{O}(U(q)) = -1$. It is simple to check that the order $[x_1, x_3, x_2, x_4]$ restricts to the orders on all these tripods. Now do this for all cases. In the interest of completeness, we list them all, including the ones with duplicate signs, in Figure 5.2. We omit exact negatives for conciseness (and sanity), so it suffices to show only those cases for which $\widehat{O}(R(q)) = +1$.

Case	$\widehat{O}(R(q))$	$\widehat{O}(L(q))$	$\widehat{O}(D(q))$	$\widehat{O}(U(q))$	compatible order
(1)	1	1	1	1	$[x_1, x_2, x_3, x_4]$
(1)	1	-1	1	-1	$[x_1, x_3, x_2, x_4]$
(2)	1	1	1	1	$[x_1, x_2, x_3, x_4]$
(2)	1	-1	-1	1	$[x_1, x_2, x_4, x_3]$
(3)	1	-1	1	-1	$[x_1, x_3, x_2, x_4]$
(3)	1	-1	-1	1	$[x_1, x_2, x_4, x_3]$

Figure 5.2: All possibilities for the values of $\widehat{O}$ on the constituent tripods of q, and the orders on Q which induce these values of $\widehat{O}$.

In all cases, there is a cyclic order on Q which induces the required signs on all the tripods. By Lemma 5.3.2, then, $\widehat{O}$ is compatible on all tripods and gives a cyclic order on S. Recall that by construction, ϕ^n preserves $\widehat{O}$. Therefore, ϕ^n is geometric with respect to the the basic hyerbolic surface realization $(\Sigma_{\widehat{O}}, f_{\widehat{O}})$. □

Theorem 5.3.6. *A random endomorphism $\phi : F_k \to F_k$ of length n is virtually geometric with probability $1 - C(k)^{-n}$.*

Proof. By Proposition 5.3.4, it suffices to check that a random endomorphism of length n is expanding at scale $K = 1$. Notice, however, that expanding at scale $K = 1$ immediately follows if the targets of the generators have bounded cancellation. In particular, condition (A) from §3.3.3 on the set of words obtained by applying ϕ to the generators guarantees that ϕ is expanding. By the proof of Theorem 3.3.11, we see that condition (A) will be satisfied with probability $1 - C(k)^{-n}$, so the theorem follows. □

5.3.3 Equivariant rotation quasimorphisms

We recall the rigidity theorem, Theorem 2.5.5, relating rotation quasimorphisms and surface realizations. If (Σ, f) is a hyperbolic surface realization, recall we denote the associated rotation quasimorphism by rot_Σ.

Lemma 5.3.7. *Let F be a free group and $\phi : F \to F$ geometric with respect to (Σ, f). Then $\mathrm{rot}_\Sigma = \mathrm{rot}_\Sigma \circ \phi$ as elements of Q/H^1. In particular, $\mathrm{rot}_\Sigma(C) = \mathrm{rot}_\Sigma(\phi(C))$ for $C \in B_1^H(F)$.*

Proof (smooth). The area form on Σ is the coboundary of $[\mathrm{rot}_\Sigma] \in Q/H^1$. Let $g : \Sigma' \to \Sigma$ be a pleated surface. Because ϕ is homotopic to an immersion, $\phi \circ g : \Sigma' \to \Sigma$ is homotopic to a pleated surface with the same signs on each triangle. Therefore, $\delta[\mathrm{rot}_\Sigma] = \delta[\mathrm{rot}_\Sigma \circ \phi]$ in $H_b^2(F)$, so rot_Σ and $\mathrm{rot}_\Sigma \circ \phi$ differ by a homomorphism; the claim follows. □

Proof (combinatorial). Let $C \in B_1^H$, and take any fatgraph S' bounding C. The value of rot_Σ is obtained by counting the signs on the tripods on Σ'. By definition, ϕ preserves all these signs, so $\mathrm{rot}_\Sigma(C) = \mathrm{rot}_\Sigma(\phi(C))$. □

Lemma 5.3.8. *Let F be a free group and ϕ geometric with respect to (Σ, f). If $C \in C_1(F)$ and $C - \phi(C) \in B_1(F)$, then if $\mathrm{rot}_\Sigma(C - \phi(C)) = 2\mathrm{scl}(C - \phi(C))$, then $\mathrm{rot}_\Sigma(C - \phi^n(C)) = 2\mathrm{scl}(C - \phi^n(C))$ for all $n \geq 1$.*

Proof. If $\mathrm{rot}_\Sigma(C - \phi(C)) = 2\mathrm{scl}(C - \phi(C))$, then $C - \phi(C)$ bounds an immersed surface Σ' in Σ. Since ϕ is homotopic to an immersion, we can glue Σ' to $\phi(\Sigma')$ along the boundary $\phi(C)$ to obtain an immersed surface bounding $C - \phi(C) + \phi(C - \phi(C)) = C - \phi^2(C)$. Therefore, $\mathrm{rot}_\Sigma(C - \phi^2(C)) = 2\mathrm{scl}(C - \phi^2(C))$. Iterating this construction proves the lemma. □

Lemma 5.3.9. *Let $C, C' \in C_1(F)$ be chains representing the same homology class, and let $\phi : F \to F$ fix this homology class. Let ϕ be geometric with respect to (Σ, f). Then $\mathrm{rot}_\Sigma(C - \phi(C)) = \mathrm{rot}_\Sigma(C' - \phi(C'))$.*

Proof. Since $[C] = [C']$, we have $C - C' \in B_1(F)$. Using the invariance of rot_Σ under ϕ, we see

$$\mathrm{rot}_\Sigma(C - \phi(C) - C' + \phi(C')) = \mathrm{rot}_\Sigma(C - C') - \mathrm{rot}_\Sigma(\phi(C - C')) = 0.$$

□

5.4 Norm-realizing surfaces in geometric HNN extensions

5.4.1 Geometric HNN extensions

Recall that if $\phi : F \to F$ is an endomorphism, the HNN extensions of F by ϕ is the group $\langle F, t \,|\, txt^{-1} = \phi(x)\rangle$. We will denote this group by F_ϕ. Recall that we obtain a $K(F_\phi, 1)$ by the quotient

$$X_\phi = K(F_\phi, 1) = X \times [0,1]/\{(x,0) = (\phi(x), 1)\},$$

where here ϕ denotes the self-map of X induced by the endomorphism ϕ.

Definition 5.4.1. An HNN extension F_ϕ is *geometric* if ϕ is geometric.

Corollary 5.4.2 (Corollary of Proposition 5.3.4)**.** *Let F have rank r, and let $\phi : F \to F$ be expanding at scale $K = 1$ with dilation $C > 1$. Then there is some $n \in \mathbb{Z}$ so that $1 \le n \le 4\binom{2r}{3}$ and F_{ϕ^n} is geometric.*

5.4.2 Norm-realizing surfaces in geometric HNN extensions

Proposition 5.4.3. *Let $\phi : F \to F$ be geometric with respect to the surface realization (Σ, f). Suppose that ϕ fixes the homology class of the chain $C \in C_1(F)$. If $\mathrm{rot}_\Sigma(C - \phi(C)) = 2\mathrm{scl}(C - \phi(C))$, i.e. if rot_Σ is extremal for $C - \phi(C)$, then the surface map obtained by gluing the boundaries of any extremal surface for $C - \phi(C)$ is norm-realizing.*

Proof. Let $\alpha \in H_2(F_\phi)$ denote the homology class in F_ϕ associated with $C - \phi(C)$. By assumption, $\mathrm{rot}_\Sigma(C - \phi(C)) = 2\mathrm{scl}(C - \phi(C))$. By Lemmas 5.3.8 and 5.3.7, we have

$$k\mathrm{rot}_\Sigma(C - \phi(C)) = \mathrm{rot}_\Sigma(C - \phi^k(C)) = 2\mathrm{scl}(C - \phi^k(C)).$$

Therefore, by Lemma 5.2.6

$$\|\alpha\| = \lim_{k\to\infty} \frac{2}{k}\mathrm{scl}\left(C - \phi^k(C)\right) = \mathrm{rot}_\Sigma(C - \phi(C)).$$

Now take any surface map $g : S \to X_\phi$ which is extremal for $C - \phi(C)$ in X. By taking covers, we may assume that the boundaries of S map with the same degrees to C and $-\phi(C)$, so that we may glue it together to form a closed surface map $\bar{g} : \bar{S} \to X_\phi$ representing a multiple of the homology class α. Since S is extremal, we have

$$\|\alpha\| = \mathrm{rot}_\Sigma(C - \phi(C)) = \frac{-\chi(S)}{n(S)} = \frac{-\chi(\bar{S})}{n(\bar{S})}.$$

That is, $\bar{S}$ is norm-realizing. $\square$

Lemma 5.2.8 can certify norm-realizing surfaces, but the combination of Propositions 5.3.4 and 5.4.3 provides a nice geometric interpretation, for it implies that most endomorphisms (and certainly almost all random endomorphisms) are geometric, and finding a norm-realizing surface reduces to checking a single equality (the extremality of rot_Σ). This produces a large number of positive examples of norm-realizing, and thus injective, surfaces in HNN extensions, and we provide a few below.

Example 5.4.4. Let $F = \langle a, b\rangle$ and $\phi(a) = aBaB$, $\phi(b) = bbABa$. Then $\phi([b]) = [b]$. Let (Σ, f) be the basic surface realization given by the cyclic order $abAB$. Checking the set of tripods shows that ϕ is geometric with respect to (Σ, f). Another way to show this is to simply check that $\mathrm{rot}_\Sigma(\phi([b, a])) = 1$, since $[b, a]$ is the boundary of this realization, and this shows the image surface is immersed. Finally, we check that $\mathrm{rot}_\Sigma(b - bbABa) = 1$ and $\mathrm{scl}(b - bbABa) = 1/2$. By Proposition 5.4.3, any extremal surface for $b - bbABa$ then gives a norm-realizing surface. Using `scallop` [15] (made easier by `wallop` [32]), we can find a particular example: a genus 1 surface with boundary $bb + AbaBBAbaBB$. A set of generators for this surface group is $\{bb, Abba, AAba, t\}$, where recall $F_\phi = \langle a, b, t \,|\, tat^{-1} = \phi(a), tbt^{-1} = \phi(b)\rangle$

Example 5.4.5. Let $F = \langle a, b, c\rangle$, and let $\phi : F \to F$ be defined

$$a \mapsto BCABc, \quad b \mapsto CCAca, \quad c \mapsto aBAcA.$$

Then ϕ is geometric with respect to the realization Σ given by the cyclic order $acCbAB$, and $\phi([a] + [CC]) = [a] + [CC]$. We also have $\mathrm{scl}(a + CC - BCABc + (aBAcA)^2) = 3/2$, and $\mathrm{rot}_\Sigma(a + CC - BCABc + (aBAcA)^2) = 3$. Therefore, any extremal surface for this chain gives a norm-realization surface. For example, there is a genus zero surface (punctured sphere) with 8 boundaries which bounds twice the chain. Gluing this surface produces a closed genus 4 norm-realizing surface in F_ϕ.

Chapter 6

Scylla

This chapter gives a generalization of Calegari's `scallop`. The `scallop` algorithm is described in detail in [8] §4, and an implementation is available [15]. The current distribution of `scallop` also includes an implementation of `scylla`.

6.1 Overview

The `scallop` algorithm computes the scl of elements of $B_1^H(F_k)$ by expressing it as the solution to a linear programming problem in the following way. Recall Definition 2.1.2, so for $\Gamma \in B_1^H(F_k)$, we have $\mathrm{scl}(\Gamma)$ as an infimum over all admissible surfaces maps. Calegari observes that it is enough to consider labeled fatgraph maps, e.g., by Lemma 2.4.4. Such structures are parameterized by a polyhedron in a linear space V, and the Euler characteristic expression to be minimized is a linear function on this space. Doing linear programming on this parameterizing space, then, computes $\mathrm{scl}(\Gamma)$.

The dimension of the parameterizing space V is $O(|\Gamma|^{2r})$, where r is the rank of F. Therefore, `scallop` is polynomial in the chain length $|\Gamma|$ and exponential in the rank. Note that the *space* complexity is also polynomial in $|\Gamma|$ and exponential in r. In practice, depending on the particular problem instance, either time or space might be a limiting factor.

The `scylla` algorithm comprises two improvements to `scallop`. First, we give a combinatorial structure on surface maps into a free group consisting of smaller pieces; this gives a linear programming problem with $O(|\Gamma|^3)$ columns which computes $\mathrm{scl}(\Gamma)$. Note this is independent of rank. Second, we extend the algorithm to compute scl in free products of infinite *or finite* cyclic groups. Let $G = \mathbb{Z}^{*k} * \mathbb{Z}/o_1\mathbb{Z} * \cdots * \mathbb{Z}/o_m\mathbb{Z}$, and let $K = \sum_{i=1}^m o_i$. Then for $\Gamma \in B_1^H(G)$, `scylla` computes $\mathrm{scl}(\Gamma)$ with a linear programming problem with $O((K+1)|\Gamma|^3)$ columns and $O((K+1)|\Gamma|^2)$ rows.

6.2 Improving the combinatorial structure in `scallop`

Recall that `scylla` contains two improvements to `scallop`: it is uniformly polynomial in the chain length, irrespective of rank, and it computes in free products of infinite and finite cyclic groups. Since these improvements are somewhat disjoint, we discuss them separately. First, we show how to remove the exponential dependence on rank.

Let F_k be a free group of rank k and $\Gamma \in B_1^H(F_k)$. We will write $\Gamma = \sum_i w_i$, expressing Γ as a formal sum of words. Note that while $B_1^H(F_k)$ is a vector space over $\mathbb{R}$, it suffices to compute scl over $\mathbb{Q}$ by continuity and over $\mathbb{Z}$ by clearing denominators and using homogeneity. We use the term *generator* to mean a generator of F_k. The formal sum of words $\sum_i w_i$ in these generators represents Γ, and we will call a particular letter in a particular location in one of these words a *letter*. It is important to distinguish between a generator and a particular occurrence of that generator in the chain Γ (i.e., a letter). Two letters are inverse if the generators they denote are inverse. We use $\Gamma_{i,j}$ to denote letter j of word i of Γ, with indices starting at 0. If a chain is written out, we will use similar subscripts to reference letters, so if $\Gamma = abAABB + ab$, then $a_{0,0}$ denotes the a at index 0 of word 0, and $B_{0,4}$ denotes the B at index 4 in the word 0.

It is important in what follows to be mindful of the fact that the tedious notation is important only for bookkeeping reasons — the idea of the algorithm is quite straightforward and is contained in the figures.

Let $f : S \to K(F_k, 1)$ be an admissible map for Γ. By compressing S (which can only reduce $-\chi(S)$), we may assume that f is given by a fatgraph structure on S. This is the content of Lemma 2.4.4. See Figure 6.1 for an example of an admissible fatgraph map. By abuse of notation, we refer to the fatgraph by S, and we think of it as already fattened. Each edge of S is a rectangle, and the two long edges of each rectangle are labeled by a letters in Γ which are inverses.

On each rectangle in S, we find two letters, say x and y, which must be inverses, and we denote the rectangle with these long edges labeled by x and y by $[x, y]$. The rectangle $[x, y]$ is the same as $[y, x]$. The short edges of rectangle $R = [x, y]$ are called *rectangle interface edges*, and they are classified by which letters in Γ label the adjacent long edges. If an interface edge lies between long edges labeled by x and y, we denote it by $ri(x, y)$. The interface edge $ri(x, y)$ is *not* the same as $ri(y, x)$. The ri stands for "rectangle interface".

Thus, in the rectangle $[x, y]$, there are four edges (in order, reading counterclockwise): the edge labeled by x, the interface edge $ri(x, y)$, the edge labeled by y, and the interface edge $ri(y, x)$. See Figure 6.2.

Each vertex in the fatgraph can be thought of as a polygon whose sides are *polygon interface edges*. Each polygon interface edge is associated with a unique rectangle interface edge, in the following way. At each side of the polygon, there is an incident rectangle whose long edges are

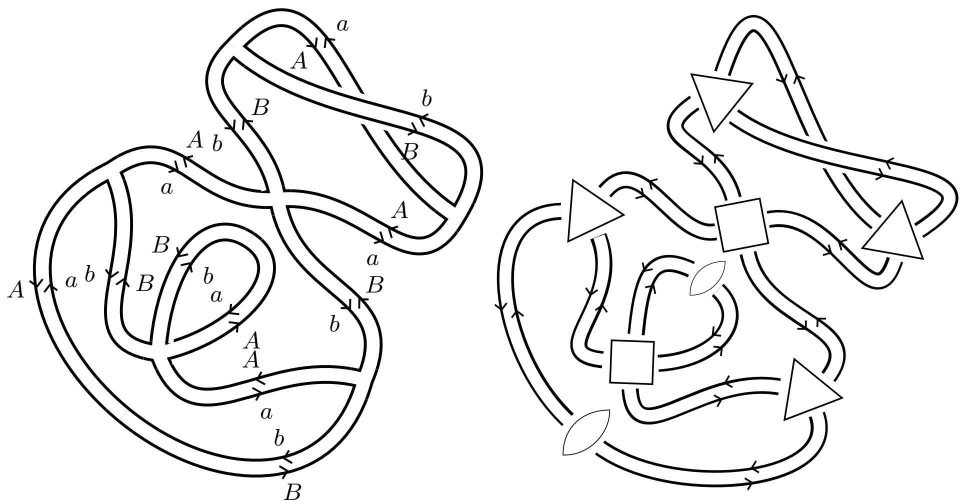

Figure 6.1: A fatgraph structure on an extremal (in particular, admissible) surface for the chain $abAABB + ab$, and the same fatgraph split into rectangles and polygons.

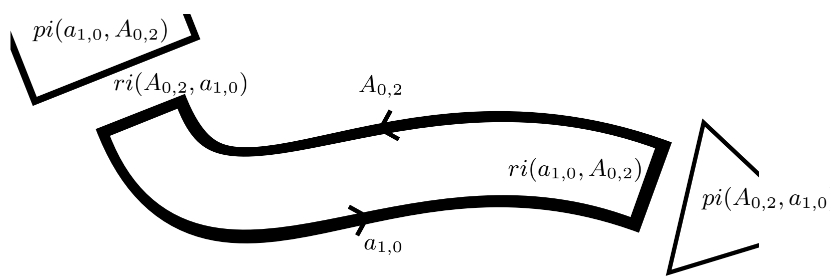

Figure 6.2: The rectangle $[a_{1,0}, A_{0,2}]$ for the chain $abAABB + ab$, pictured as an enlarged piece of Figure 6.1, as found towards the lower right of Figure 6.1, with sides and interface edges labeled.

labeled by an outgoing letter, say x, and an incoming letter, say y. Such a polygon side will be denoted by $pi(x, y)$, for "polygon interface". Notice that if the rectangle $[x, y]$ is attached to a polygon along the short edge $ri(x, y)$, the corresponding edge on the polygon is $pi(y, x)$. Also note that while there is only a single possible rectangle containing any given rectangle interface edge, the same polygon interface edge may be part of many different polygons.

Every polygon satisfies the constraint that $pi(x, \Gamma_{i,j})$ is followed counterclockwise by $pi(\Gamma_{i,j+1}, y)$, for some letters x and y, and where we take indices in word i modulo the length of word i. See Figure 6.3.

The set of all possible polygons is infinite because there is no a priori bound on the number of sides they may have. The `scallop` algorithm handles this by assuming that the polygons have at most $2r$ sides. This is not rigorous, but it is generically correct, and it can be made rigorous by applying it to the chain obtained from Γ by applying the map which takes each generator g to g^3. See [8], §4.1.

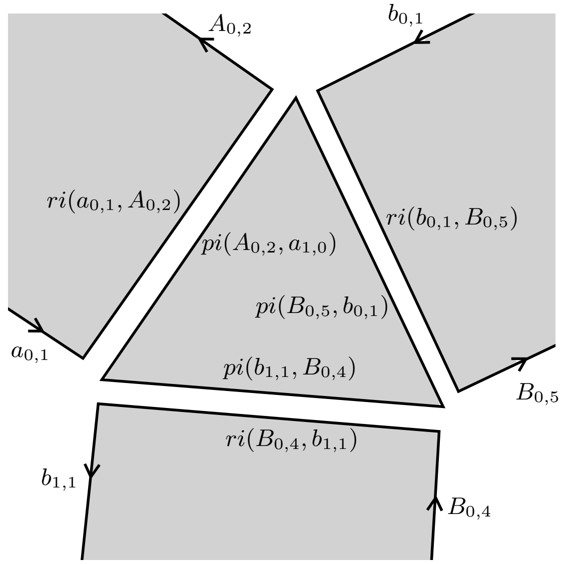

Figure 6.3: A polygon and incident rectangles, with rectangle sides and interface edges labeled. This example is the lower-right hand triangle in Figure 6.1.

The `scylla` algorithm breaks the polygons into smaller pieces of uniformly bounded size; we then observe that there can be only finitely many types of these pieces. Any polygon can be cut (possibly in multiple ways) into finitely many *polygon pieces* so that each piece has at most 3 sides. See Figure 6.4.

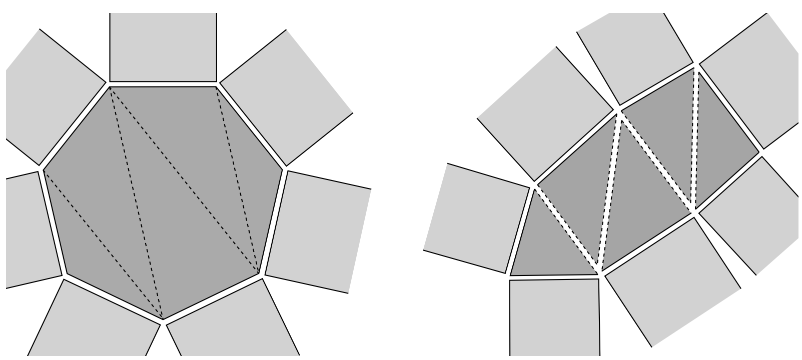

Figure 6.4: Cutting a polygon (left, shown with rectangles attached) into polygon pieces (right) with at most 4 sides.

The polygon pieces have two kinds of edges: polygon interface edges inherited from the polygons we cut and *central edges* which are newly created by the cutting procedure. Central edges are denoted by a pair $c(x, y)$, where x and y are letters in Γ, as follows. Each central edge starts at some vertex v_1 of a polygon and ends at some vertex v_2. Because each polygon piece is oriented, "start" and "end" are induced from the orientation. There are labeled rectangle sides incoming to and outgoing from both v_1 and v_2. The central edge in question is given the label $c(x, y)$, where x is the label *incoming* to vertex v_1 and y is the label *outgoing* from v_2.

Under this scheme, a central edge of a polygon piece which lies between interface edges $pi(x, y)$ and $pi(z, w)$ will be labeled $c(y, z)$. Further, if two central edges $c(x, y)$ and $c(z, w)$ are adjacent (with $c(z, w)$ following $c(x, y)$), then, somewhat counterintuitively, the letter y follows z cyclically in Γ.

Note that cutting a polygon into two polygon pieces will generate the central edge $c(\Gamma_{i,j}, \Gamma_{k,l})$ on one polygon piece and $c(\Gamma_{k,l-1}, \Gamma_{i,j+1})$ on the other. See Figure 6.5.

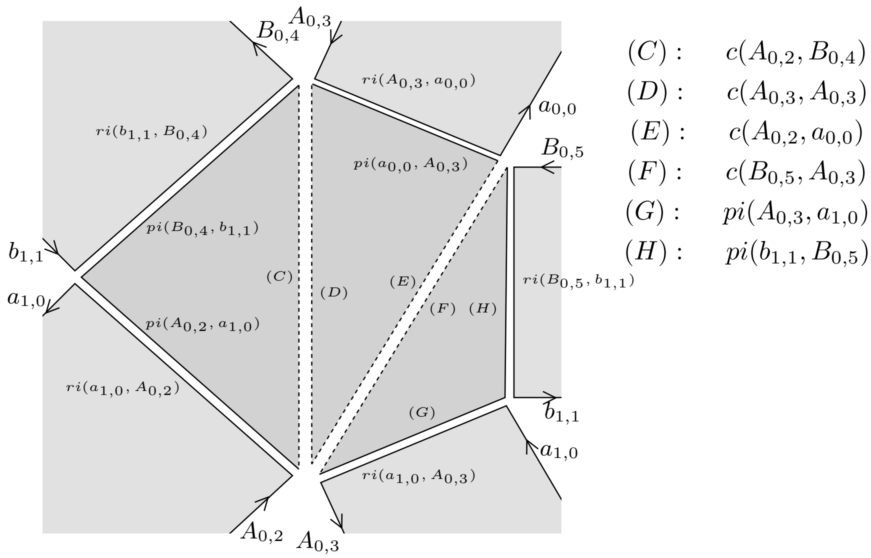

Figure 6.5: Labeling the central and interface edges of the polygon pieces for a hypothetical collection of polygon pieces and rectangles for the chain $abAABB + ab$. For clarity, some of the labels have been moved off to the right, as indicated.

Note that while if $pi(x, y)$ is an interface edge, then x and y must be inverses, there is no such requirement for central edges.

Thus, given a fatgraph map $f : S \to K(F_k, 1)$, we may cut it into rectangles and polygons, and we may cut the polygons further into polygon pieces. Each polygon piece has at most 3 sides, and each side is labeled by a pair of letters in Γ. Since the letters on either side of a corner of a polygon piece determine each other, we see that there are at most $|\Gamma|^3$ labeled polygon pieces.

Given $\Gamma \in B_1^H(F_k)$, let V_Γ be the vector space over $\mathbb{Q}$ spanned by all possible labeled polygon pieces and rectangles for Γ. If (S, f) is a fatgraph map for Γ, then the discussion above shows how to cut S into polygon pieces and rectangles. The *cutting vector* for this decomposition is the vector in V_Γ which records how many of each rectangle and polygon piece are present in S. Since there might be many ways to cut the polygons into polygon pieces, there might be many cutting vectors for a single fatgraph map. Similarly, if an admissible surface admits a fatgraph structure, it might admit many fatgraph structures, so for a single admissible surface, there might be many cutting vectors.

For each pair of letters (x, y) in Γ such that x and y are inverses, we define an *interface linear functional* $I_{(x,y)}$. It suffices to define $I_{(x,y)}$ on a basis for V_Γ, that is, on rectangles and polygon pieces. On the rectangle $[x, y]$, $I_{(x,y)}$ takes the value 1. It is zero for all other rectangles. On a polygon piece p, $I_{(x,y)}$ takes the value which is negative the number of times that $pi(y, x)$ appears as a side of p.

For each pair of letters of Γ, we define a *central linear functional* $C_{(x,y)}$, as follows. These functionals take the value 0 on all rectangles. On a polygon piece p, $C_{(\Gamma_{i,j},\Gamma_{k,l})}$ gives the number of times that $c(\Gamma_{i,j}, \Gamma_{k,l})$ appears as a side of p, minus the number of times that $c(\Gamma_{k,l-1}, \Gamma_{i,j+1})$ appears.

Note that $C_{(\Gamma_{i,j},\Gamma_{k,l})} = -C_{(\Gamma_{k,l-1},\Gamma_{i,j+1})}$. Therefore, it is only necessary to use half of these linear functions. The pairing is slightly awkward to express, so we will ignore this technicality in our discussion.

Given a fatgraph map (S, f) and a cutting vector $v_{(S,f)} \in V_\Gamma$, we observe that

1. $v_{(S,f)}$ is nonnegative.

2. $I_{(x,y)}(v_{(S,f)}) = 0$ for all $x = y^{-1}$ letters in Γ.

3. $C_{(x,y)}(v_{(S,f)}) = 0$ for all x, y letters in Γ.

We would like to use the converse; that any vector satisfying these constraints can be used to build a fatgraph. This is not quite true. Any suitable collection of poylgons pieces and rectangles can be used to build an admissible surface, but it might not be a fatgraph. We will see, however, that any *extremal* (as defined later) collection does yield a fatgraph.

Further, it is possible to determine the Euler characteristic of an admissible fatgraph surface from *any* cutting vector. We define a linear function χ on V_Γ by defining it on our basis of rectangles and polygon pieces. On rectangles, χ is zero. On a polygon piece p with n sides, we set $\chi(p) = (2-n)/2$.

Lemma 6.2.1. *Let $\Gamma \in B_1^H(F)$ and $v \in V_\Gamma$ be any cutting vector for a fatgraph map (S, f) for Γ. Then $\chi(S) = \chi(v)$.*

Proof. The fatgraph (S, f) is homotopy equivalent to a graph whose vertices are the polygons and whose edges are the rectangles. This graph, in turn, is homotopy equivalent to a graph G whose vertices are polygon pieces and rectangles, and two vertices are joined by an edge exactly when the corresponding polygon pieces or rectangles share an interface or central edge.

Let $V = \{v_i\}_i$ and E be the vertices and edges of G. We compute $\chi(S) = \chi(G)$ as $|V| - |E|$. Let $|v_i|$ denote the valence of vertex v_i. Then $|E| = \sum_i |v_i|/2$. Therefore,

$$\chi(S) = \sum_i 1 - \frac{|v_i|}{2} = \sum_i \frac{2 - |v_i|}{2},$$

and the expression on the right coincides with $\chi(v)$. $\square$

Finally, we define a collection of linear functions which compute the image in homology of a surface map. Recall we have written $\Gamma \in B_1^H(F_k)$ as the formal sum of words

$$\Gamma = w_0 + \cdots + w_{k-1}.$$

For each $0 \leq i < k$, we define a linear function N_i which is zero on polygon pieces. On the rectangle $[x, y]$, N_i is 1 if x or y is the first letter of w_i and 0 otherwise.

Lemma 6.2.2. *Let Γ be as above and γ_i be a loop in a $K(F_k, 1)$ representing the word w_i. Let $\langle[\gamma_i]\rangle$ denote the span of the vector $[\gamma_i]$ in $B_1^H(F_k)$. Let (S, f) be an admissible fatgraph, and let $v \in V_\Gamma$ be any cutting vector. There is an induced map $\partial f_{*,i} : H_1(\partial S) \to \langle[\gamma_i]\rangle \cong \mathbb{R}$, and $N_i(v)[\gamma_i] = \partial f_{*,i}([\partial S])$.*

Proof. By the construction of the vector space V_Γ, any surface for which v is a cutting vector will have boundary components whose images in $\pi_1(K(F_k, 1)) = F_k$ are conjugates of powers of the w_i, so in $B_1^H(F_k)$, the image of $[\partial S]$ is in the subspace spanned by the collection of all $[\gamma_j]$. Projecting to the ith component, which we denote by $\partial f_{*,i}$, ignores all loops except γ_i, so the image of $\partial f_{*,i}[\partial S]$ in $B_1^H(F_k)$ lies in the subspace spanned by $[\gamma_i]$. To find $\partial f_{*,i}[\partial S]$ as a multiple of $[\gamma_i]$, we just need to count the number of times that the word w_i appears in the labeled boundaries of ∂S. Therefore, we simply need to count the number of times that the first letter of w_i appears, which is $N_i(v)$. $\square$

We are now ready to construct the linear programming problem. Let $\Gamma = w_0 + \cdots + w_{k-1} \in B_1^H(F_k)$, and define $P_\Gamma \subseteq V_\Gamma$ to be the set of vectors v such that

1. v is nonnegative.

2. $I_{(x,y)}(v) = 0$ for all $x = y^{-1}$ letters in Γ.

3. $C_{(x,y)}(v) = 0$ for all x, y letters in Γ.

4. $N_i(v) = 1$ for $0 \leq i < k$.

We call P_Γ the *admissible polyhedron.* Note that P_Γ is, in fact, a polyhedron.

Lemma 6.2.3. *Let (S, f) be a fatgraph map for Γ. For any cutting vector v for (S, f), there is some scalar $k \in \mathbb{Q}$ so that $kv \in P_\Gamma$. Further,*

$$\frac{-\chi(kv)}{2} = \frac{-\chi(S)}{2n(S, f)}.$$

Proof. If v is a cutting vector for a fatgraph map (S, f), then v satisfies conditions (1) through (3) in the definition of the admissible polyhedron, and any scalar multiple of v does too. Write

$\Gamma = \sum_i w_i$ as above. Because (S, f) bounds Γ in $B_1^H(F_k)$, all of the $N_i(v)$ are equal to some integer K. Therefore, $(1/K)v$ lies in the admissible polyhedron, as desired.

Because all of the $N_i(v)$ are equal, $N_1(v) = n(S, f)$, so since χ and N_1 are linear,

$$\frac{-\chi(S)}{2n(S,f)} = \frac{-\chi(v)}{2N_1(v)} = \frac{-\chi(kv)}{2N_1(kv)} = \frac{-\chi(kv)}{2}.$$

□

Lemma 6.2.4. *Let $\Gamma \in B_1^H(F_k)$ and $v \in P_\Gamma$ (not necessarily integral) so that $-\chi$ is minimized over P_Γ at v. Then there exists a scalar $k \in \mathbb{Q}$ and (possibly many) admissible fatgraphs (S, f) for Γ such that kv is a cutting vector for (S, f). Furthermore,*

$$\frac{-\chi(v)}{2} = \frac{-\chi(S)}{2n(S,f)}.$$

Proof. Because v is a rational vector, there is some scalar $k \in \mathbb{Q}$ so that kv is integral.

The integer vector $kv \in V_\Gamma$ corresponds to a collection of rectangles and polygon pieces. The fact that the linear functionals $C_{(x,y)}$ vanish ensures that amongst all sides of all polygon pieces, there are exactly the same number of central edge $c(\Gamma_{i,j}, \Gamma_{k,l})$ as there are of the matching central edge $c(\Gamma_{k,l-1}, \Gamma_{i,j+1})$. For every pair of central edges, glue the collection of sides of polygon pieces with those edges arbitrarily. In a similar way, the vanishing of all the linear functionals $I_{(x,y)}$ ensures that amongst all rectangles and polygon pieces, there are exactly the same number of rectangle interface edge $ri(x, y)$ as there are of polygon interface edge $pi(y, x)$. For each interface edge pair, then, glue the collection arbitrarily. By sending all polygons to the basepoint of $K(F_k, 1)$ and all the rectangles appropriately around the 1-cells, we obtain a surface and map (S, f). By construction of the polygon pieces and rectangles, this surface bounds some multiple of Γ.

By the same argument as Lemma 6.2.1, $\chi(kv) = \chi(S)$, and by the same argument as Lemma 6.2.3, $-\chi(v)/2 = -\chi(S)/2n(S, f)$.

The result of this gluing a priori *might not be a fatgraph.* There are two bad things that could happen: (1) a sequence of polygon pieces with two central edges could be glued in a loop, yielding a polygon with a puncture, and (2) a sequence of poylgon pieces could be glued so that the resulting object is a loop, not a simply connected polygon.

In the case (1), we get a puncture in the surface which maps to the basepoint. We can therefore fill in this puncture, reducing $-\chi$, to get a new admissible surface map (S', f'). Apply Lemma 2.4.4 to get an admissible fatgraph and a cutting vector v'. Applying Lemma 6.2.3, find k' so that $k'v' \in P_\Gamma$. Then

$$-\chi(v)/2 = -\chi(S)/2n(S, f) > -\chi(S')/2n(S', f') = -\chi(k'v')/2.$$

Since v minimizes $-\chi$ over P_Γ, this is a contradiction. Thus, (1) cannot happen.

In the case (2), we get a compressible loop in the glued up surface. in this case, too, we compress along it to get a new surface with smaller $-\chi$, yielding another contradiction.

Therefore, the result of this gluing *is* a labeled fatgraph (together with its induced map). By construction, it can be cut into the polygon pieces and rectangles that we started with, so kv is a cutting vector for (S, f). □

6.3 The `scylla` algorithm for free groups

Proposition 6.3.1. *Let* $\Gamma \in B_1^H(F_k)$. *Then*

$$\mathrm{scl}(\Gamma) = \inf_{v \in P_\Gamma} \frac{-\chi(v)}{2}.$$

Proof. Given any admissible surface (S, f), we may apply Lemma 2.4.4 followed by Lemma 6.2.3 to obtain some vector $v \in P_\Gamma$ so that $-\chi(v)/2 \leq -\chi(S)/2n(S, f)$. Thus,

$$\inf_{v \in P_\Gamma} \frac{-\chi(v)}{2} \leq \inf_{(S,f)} \frac{-\chi(S)}{2n(S, f)}.$$

The reverse inequality is immediately implied by Lemma 6.2.4, so we have

$$\inf_{v \in P_\Gamma} \frac{-\chi(v)}{2} = \inf_{(S,f)} \frac{-\chi(S)}{2n(S, f)} = \mathrm{scl}(\Gamma),$$

□

Observe that this expresses the computation of $\mathrm{scl}(\Gamma)$ as a linear programming problem, since P_Γ is a polyhedron and $-\chi$ is a linear function. Furthermore, the complexity is independent of rank, because the rank does not affect the construction of the rectangles or polygon pieces.

6.4 Generalizing to finite cyclic free factors

We now describe the complete `scylla` algorithm, which computes scl in free products of cyclic groups with finite cyclic free factors allowed. The approach is very similar to the free group case, but the rectangles are replaced by more general objects.

It should be noted that free products of finite cyclic groups are virtually free. Using Corollary 2.81 in [8] and `scallop`, it is theoretically possible to compute scl in free products of finite cyclic groups. However, the rank of the free group in which the computation is done is essentially the product of the orders of the finite cyclic groups. Given `scallop`'s exponential dependence on rank, this method is not feasible. Even using the first improvement in `scylla` to remove the exponential dependence on rank does not fix the problem, for the size of the computation still grows too quickly. Therefore, it is necessary to compute directly in the finite cyclic groups.

Recall that a fatgraph map (necessarily into a free group) can be described by a collection of labeled rectangles and polygon pieces, appropriately glued together along interface and central edges. Let $G = \mathbb{Z}^{*k} * \mathbb{Z}/o_1\mathbb{Z} * \cdots \mathbb{Z}/o_m\mathbb{Z}$ and let $\Gamma \in B_1^H(G)$. A *cyclic fatgraph bounding* Γ is given by a collection of the following objects, glued along interface and central edges: central polygon pieces (previously called polygon pieces); rectangles, whose side labels are letters in Γ representing generators in the factor $\mathbb{Z}^{*k}$; and group polygons. A group polygon is associated with a finite free factor $\mathbb{Z}/o_i\mathbb{Z}$, and it has two kinds of sides, which alternate: *group polygon interface edges*, which are glued to polygon interface edges, and labeled sides, each labeled by a letter in Γ representing the generator of $\mathbb{Z}/o_i\mathbb{Z}$. We will denote by $gi(x, y)$ the group polygon interface edge between two sides labeled with letters x and y. If y immediately (cyclically) follows x in Γ, then the group polygon interface edge $gi(x, y)$ is a *dummy* edge, and it is not allowed to be glued to a polygon. A group polygon in factor $\mathbb{Z}/o_i\mathbb{Z}$ contains exactly o_i interface edges and o_i labeled sides. See Figure 6.6.

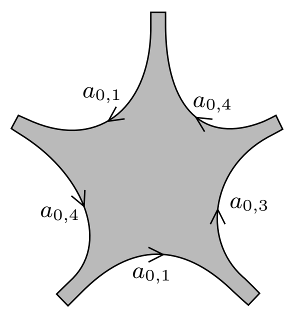

Figure 6.6: A group polygon in $\mathbb{Z}/5\mathbb{Z}$ generated by a. Here the group polygon is in a hypothetical cyclic fatgraph bounding the chain $aabaaaB$ in $\mathbb{Z}/5\mathbb{Z} * \mathbb{Z}$.

Note that for the definition of a cyclic fatgraph, we have done away with the polygons, and we are only using polygon pieces. This makes the technicalities of Lemma 6.2.4 easier. Note that we are able to give surfaces with punctures and compressible loops cyclic fatgraph structures. This is fine, because these surface maps will never be extremal.

We need to observe the following, but it does not quite deserve to be elevated to the status of a lemma.

Remark 6.4.1. A cyclic fatgraph bounding Γ, as defined above, induces a surface map $f : S \to K(G, 1)$ which bounds Γ in $B_1^H(G)$. All that we need to check is that the obvious map on polygon pieces (to the basepoint) and rectangles (around infinite free factors, as before) extends to group polygons. But the group polygons in factor $\mathbb{Z}/o_i\mathbb{Z}$ all have o_i sides, so their boundaries do bound disks in a $K(G, 1)$.

In light of Remark 6.4.1, we will conflate a cyclic fatgraph bounding Γ and the (admissible) surface map it induces.

Much more interesting is the converse. Note that this lemma is intuitively quite clear. A glance at Figure 6.9 is highly recommended before reading the proof.

Lemma 6.4.2. *Let $G = \mathbb{Z}^{*k} * \mathbb{Z}/o_1\mathbb{Z} * \cdots \mathbb{Z}/o_m\mathbb{Z}$ and $\Gamma \in B_1^H(G)$. Let (S, f) be an admissble surface for Γ. There is a cyclic fatgraph (S', f') with the same boundary image in $B_1^H(G)$ such that $-\chi(S') \leq -\chi(S)$.*

Proof. First, we take a form for Γ which includes only positive powers of the generators of the finite factors. To understand the set of possible surface maps into a $K(G, 1)$, it suffices to consider the 2-skeleton, which we construct now. Take a single vertex, and $k+m$ 1-cells, which forms a free group of rank $k + m$. Now glue on m 2-cells $\{T_i\}_{i=1}^m$ by gluing cell i to the $(k + i)$th 1-cell (which forms a copy of S^1) via a map on the boundary of degree o_i. Denote this CW-complex by X. Clearly $\pi_1(X) \cong G$, and since we may take any map between CW-complexes to be cellular, any surface map into a $K(G, 1)$ is homotopy equivalent to a map to X.

Now, $\Gamma \in B_1^H(G)$, and let (S, f) be an admissible map for Γ. After possibly homotoping f, we may assume that f is cellular, and we may assume that the boundary components of S are wrapped around the generators of the finite factors only in a positive direction. Suppose that a cell C in S maps via f to the cell T_i in X. By the homotopy extension property, C is a branched cover of T_i. Since f takes ∂C to the $(k+i)$th 1-cell while factoring through ∂T_i, the map of C into X must be a map whose degree on the boundary is a multiple of the order o_i. Further, after homotopy, we may assume that any branch point in C maps to a given point in T_i, say p_i.

Now excise a small neighborhood P_i of p_i in each cell T_i to obtain the space X', and remove the preimage of each P_i from S to obtain a surface map $f' : S' \to X'$. This gives

$$\begin{array}{ccc} S & \xrightarrow{f} & X \\ \uparrow & & \uparrow \\ S' & \xrightarrow{f'} & X' \end{array},$$

where the vertical arrows are inclusions, and S and X are obtained from S' and X' by gluing in some disks. Since $\pi_1(X')$ is a free group with the same generators as $\pi_1(X) \cong G$, we may find a fatgraph map (S'', f'') obtained from (S', f') through homotopy and compressions. Note the image of $\partial S''$ and ∂S are the same in $B_1^H(G)$.

The fatgraph S'' has some special boundary components which bound powers of the generators of the finite factors; these correspond to the components of S' to which we glue disks in order to obtain S. Note that any two of these boundary components cannot run across both sides of the same rectangle in S''. One would have to contain the generator of a finite factor and the other its inverse, but by construction these boundary components can contain only the inverse of the generator (because the boundary of S contains only positive powers of these generators).

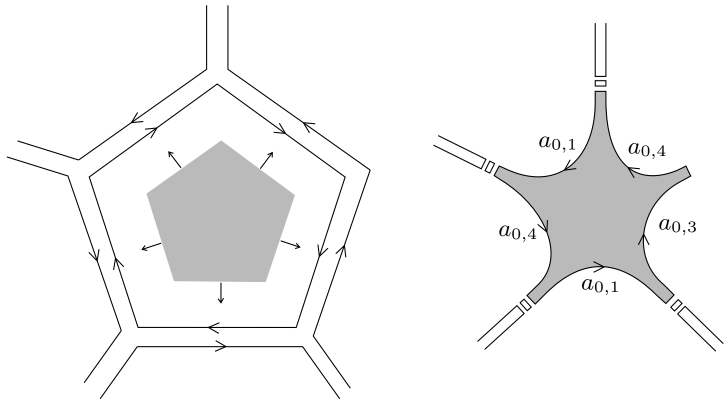

Figure 6.7: Creating a group polygon by gluing a disk to a boundary component of S'. The small squares on the right are central polygons. Note the dummy edge created when a central polygon with two sides (a vertex of the fatgraph S'' with only two incident rectangles) lies along a boundary component to which we glue a disk. As an example, we provide possible labels for a group polygon bounding $aabaaaB$ in $\mathbb{Z}/5\mathbb{Z} * \mathbb{Z}$.

We replace these boundary components with group polygons as in Figure 6.7, obtaining a cyclic fatgraph bounding Γ in $B_1^H(G)$ and its induced map (S''', f'''), where S''' is the cyclic fatgraph bounding Γ. Suppose that we remove K disks from S (and glue K disks to $-\chi(S'')$). Then

$$-\chi(S''') - K = -\chi(S'') \leq -\chi(S') = -\chi(S) - K$$

as desired.

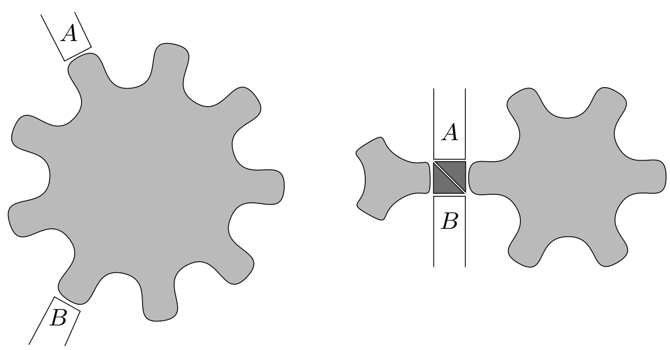

Figure 6.8: Pinching off o_i of the Mo_i sides of a group polygon with too many sides (left) to create a real group polygon with o_i sides, two central polygon pieces, and a smaller bad "group polygon". Here $o_i = 3$ and the new central polygon is darkened. All the interface edges might be dummy edges or be glued to polygons, but we emphasize what happens around the interface edges by labels A and B for clarity.

Observant readers will note that, in fact, the "group polygons" that we just created for each factor $\mathbb{Z}/o_i\mathbb{Z}$ might not have exactly o_i labeled sides — they might have Mo_i labeled sides, where

$M \geq 1$ is any integer. To resolve this and produce an honest cyclic fatgraph structure, we perform the move indicated in Figure 6.8, which replaces a "group polygon" with Mo_i sides with: an honest group polygon with o_i sides, two central polygon pieces, and a "group polygon" with $(M-1)o_i$ sides. Repeated applications of this move produce a cyclic fatgraph obviously homotopy equivalent to (S''', f'''). This is the cyclic fatgraph of the lemma, and this completes the proof. $\square$

Example 6.4.3. Let $G = \mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ generated by a and b, respectively. The chains ab and $aabab$ are both contained in $B_1^H(G)$, because $(ab)^6$ and $(aabab)^6$ are homologically trivial. We show cyclic fatgraphs bounding these chains in Figure 6.9. Group polygons are shown darkened, and the central polygons (bigons) lie between the group polygons. Note the dummy group polygon interface edges in the surface on the right. Also note that the objects which look like rectangles are actually group polygons with 2 interface sides.

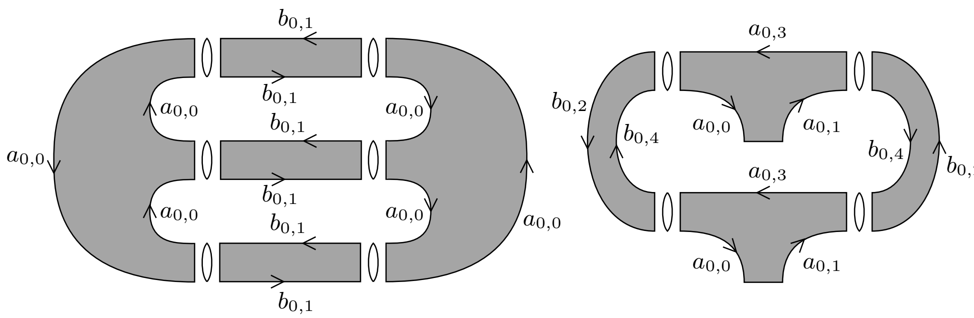

Figure 6.9: See Example 6.4.3. Cyclic fatgraphs bounding the chains ab and $aabab$ in the group $G = \mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$. These exhibit $\mathrm{scl}_G(ab) \leq 1/12$ and $\mathrm{scl}_G(aabab) = 0$. In fact, they are both extremal.

We review the current context. We have shown that every admissible surface can be given a cyclic fatgraph structure (after possibly compressing, etc.). In a similar way as the fatgraph structure does for free groups, this allows us to parameterize the space of all admissible surfaces and do linear programming to compute scl. We simply cut the cyclic fatgraphs into central polygon pieces, rectangles, and group polygons, show that we can compute Euler characteristic from these pieces, and then show how to build a polyhedron in an appropriate vector space.

However, the complexity of such an algorithm is prohibitive, since the number of possible labeled group polygons in a factor $\mathbb{Z}/o_i\mathbb{Z}$ can be $O(|\Gamma|^{o_i})$. Therefore, we must break our surfaces into even smaller pieces. We do this by building the group polygons themselves out of smaller *group teeth.*

Consider a group polygon g in free factor i with o_i sides. Pick any labeled side, which will be labeled with some letter x, and call g a group polygon *based at x*. For any given group polygon, there are multiple represention as a based group polygon, but certainly any group polygon can be represented in this way. Now set the chosen labeled side to have index 0, and index the remaining sides counterclockwise from 0. Also index the interface sides from 0. Thus, interface side i lies between labeled sides i and $i+1$, indices taken modulo o_i. A *group tooth* is a single labeled interface side at a particular index in a group polygon based at a particular letter x. We will denote it by $gt(y, z, i, x)$, where the interface edge $gi(y, z)$ is at index i in a group polygon based at x. Note that $y = x$ if $i = 0$, and $z = x$ if $i = o_i - 1$. See Figure 6.10. The nomenclature comes from the resemblance of the group teeth to the teeth in a bicycle sprocket.

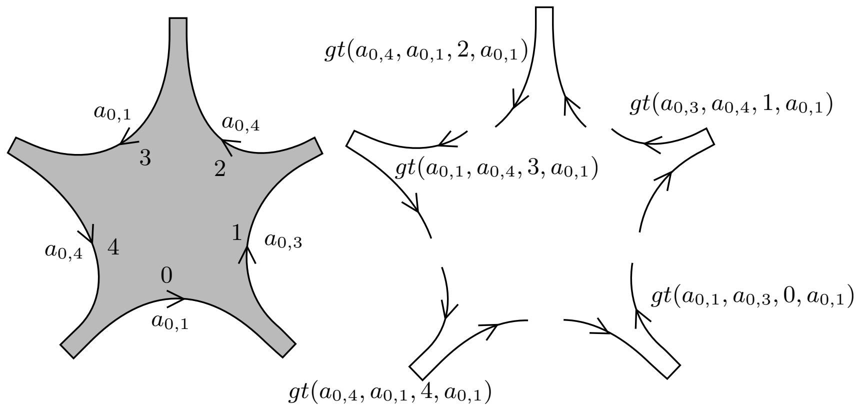

Figure 6.10: Cutting the group polygon from Figure 6.6 into group teeth. The bottom labeled side is arbitrarily chosen as the base.

We are finally ready to build the linear programming problem. Given $G = \mathbb{Z}^{*k} * \mathbb{Z}/o_1\mathbb{Z} * \cdots \mathbb{Z}/o_m\mathbb{Z}$ and $\Gamma \in B_1^H(G)$, let the V_Γ be the vector space over $\mathbb{Q}$ spanned by all labeled rectangles, polygon pieces, and group teeth. Define linear functionals $I_{(x,y)}$ and $C_{(x,y)}$ as before, where $I_{(x,y)}$ is 1 on a group tooth exactly when the group tooth is of the form $gt(x, y, \cdot, \cdot)$, and all $C_{(x,y)}$ vanish on all group teeth. Note that the $I_{(x,y)}$ are defined for $x \neq y^{-1}$ for the finite free factors (and for $x = y^{-1}$ for the infinite free factors, as before).

Define a new set of linear functionals $T_{i,y,x}$, where x and y are letters in Γ in free factor j, and $0 \leq i \leq o_j - 1$. The functional $T_{i,y,x}$ is 1 on a group tooth of the form $gt(\cdot, y, i-1, x)$ and -1 on a group tooth of the form $gt(y, \cdot, i, x)$ and 0 on all other group teeth, rectangles, and polygon pieces. The vanishing of all these linear functionals ensures that the group teeth can be glued together to obtain group polygons.

Observe that any cyclic fatgraph bounding Γ can be cut into central polygon pieces, rectangles, and group teeth. An integer vector v in V_Γ then records how many of each piece we have. We will call such v a *cutting vector* for the cyclic fatgraph. Note that there may be many different cyclic fatgraph structures on the same admissible surface. Also note that a cyclic fatgraph consists of central polygon pieces, rectangles, and group polygons, but the vector space V_Γ records central poylgon pieces, rectangles, and *group teeth.*

Lemma 6.4.4. *Let $\Gamma \in B_1^H(G)$ and $v \in V_\Gamma$ an integer vector such that*

1. *v is nonnegative.*

2. *$I_{(x,y)}(v) = 0$ for all x, y letters in Γ in the same free factor.*

3. *$C_{(x,y)}(v) = 0$ for all x, y letters in Γ.*

4. *$T_{i,y,x}(v) = 0$ for all x, y letters in the same free factor j, and $0 \le i \le o_j - 1$.*

Then there is a cyclic fatgraph map (S, f) admissible for Γ such that v is a cutting vector for (S, f).

Proof. Because v is an integer vector, it corresponds to a collection of central polygons, rectangles, and group teeth. As with the proof of Lemma 6.2.4, we simply need to show that we can glue everything up in some way to construct a surface. Unlike this lemma, though, we do not care that some boundary components might map with degree zero. First, consider the group teeth. Fix some letter x in Γ in factor j and index $i = 1$. The fact that $T_{1,y,x}(v) = 0$ for all y in Γ means that there are the same number of group teeth at position 0 based at x whose second letter is y as there are group teeth at position 1 based at x whose first letter is y. We choose some bijection between these teeth and glue them. Then proceed to the next letter y in Γ. After gluing all the group teeth at indices 1 and 2 based at x, similarly glue all the group teeth at indices i and $i + 1$, etc. Proceeding around the indices, we arrive at indices $o_j - 1$ and 0. Each group tooth at 0 is glued to one at 1, which is glued to one at 2, etc., all the way to $o_j - 1$. That is, each group tooth at 0 is glued, through a succesion of teeth counterclockwise, to a group tooth at $o_j - 1$. Because group teeth are "based", a group tooth at 0 always starts with letter x, and a group tooth at $o_j - 1$ always ends with letter x, so we may glue up these matched group teeth. This creates a collection of group polygons with the correct number (o_j) of sides.

By constructing these group polygons, we have not changed the numbers of any of the interface sides, so conditions (1) through (3) still hold. At this point, we simply proceed to glue up the interface and central edges: at each interface edge, we have the same number of central polygon interface edges as we do rectangle and group polygon interface edges — choose some bijection, and glue. Similarly glue the central polygon edges to obtain a cyclic fatgraph. By construction, this cyclic fatgraph bounds some multiple of Γ in $B_1^H(G)$, and it gives an induced surface map admissible for Γ. Since we built it out of the pieces specified by v, we can obviously cut it into pieces to give v, as desired.

Again, we reiterate that we might obtain a surface with punctures or compressible loops. This is fine. ∎

The definition of χ becomes slightly more complicated. On rectangles, χ is zero. On a polygon piece p with n sides, we set $\chi(p) = (2 - n)/2$. On a group tooth $gt(y, z, i, x)$ in free factor j,

$$\chi(gt(y,z,i,x)) = \begin{cases} \frac{1}{o_j} & \text{if } z \text{ cyclically follows } y \text{ in } \Gamma, \\ \frac{1}{o_j} - \frac{1}{2} & \text{otherwise.} \end{cases}$$

Lemma 6.4.5. *Let $\Gamma \in B_1^H(F)$ and (S, f) be an cyclic fatgraph map with a cutting vector $v \in V_\Gamma$. Then $\chi(S) = \chi(v)$.*

Proof. A cyclic fatgraph is homotopy equivalent to a graph G whose vertices are polygon pieces, rectangles, and group polygons, and where two vertices are joined by an edge exactly when the corresponding pieces are glued along an interface or central edge. We have

$$\chi(G) = \sum_w 1 - \frac{|w|}{2},$$

where the sum is taken over all vertices $w \in G$. Observe that for the rectangles and polygon pieces, as before, χ applied to that basis vector computes the appropriate expression $1 - |w|/2$. Note that just from the group teeth recorded in v, we *do not know* how many or what types of group polygons appear in the cyclic fatgraph S. However, we can still compute the expression $\sum_w (1 - |w|/2)$, where the sum is taken over all group polygons assembled *in any way* from the group teeth. For each group polygon, we know how many group teeth it contains, so we know how many group polygons there must be, and we know, by observing how many dummy interface edges there are, how many of the interface sides are actually connected. Therefore, we know the total valence of all the group polygons. This computation is exactly carried out by the linear functional χ defined above. $\square$

Next, the homology linear functionals. Write $\Gamma \in B_1^H(G)$ as the formal sum of words

$$\Gamma = w_0 + \cdots + w_{k-1}.$$

For each $0 \leq i < k$, we define a linear function N_i which is zero on polygon pieces. On the rectangle $[x, y]$, N_i is 1 if x or y is the first letter of w_i and 0 otherwise. On a group tooth $gt(y, z, i, x)$, N_i is 1 or 0 depending on whether or not y is the first letter of w_i. Lemma 6.2.2 holds with the same proof; we record it here. Recall that possible punctures (boundary components with degree zero) do not affect this computation.

Lemma 6.4.6. *Let Γ be as above and γ_i be a loop in a $K(G, 1)$ representing the word w_i. Let $\langle[\gamma_i]\rangle$ denote the span of the vector $[\gamma_i]$ in $B_1^H(G)$. Let (S, f) be an admissible cyclic fatgraph, and let $v \in V_\Gamma$ be any cutting vector. There is an induced map $\partial f_{*,i} : H_1(\partial S) \to \langle[\gamma_i]\rangle \cong \mathbb{R}$, and $N_i(v)[\gamma_i] = \partial f_{*,i}([\partial S])$.*

We are now ready to construct the linear programming problem. Let $\Gamma = w_0 + \cdots + w_{k-1} \in B_1^H(G)$, and define $P_\Gamma \subseteq V_\Gamma$ to be the set of vectors v such that

1. v is nonnegative.

2. $I_{(x,y)}(v) = 0$ for all x, y letters in Γ in the same free factor of G.

3. $C_{(x,y)}(v) = 0$ for all x, y letters in Γ.

4. $T_{i,y,x}(v) = 0$ for all x, y letters in the same free factor j, and $0 \leq i \leq o_j - 1$.

5. $N_i(v) = 1$ for $0 \leq i < k$.

We call P_Γ the *admissible polyhedron.* Note that P_Γ is, in fact, a polyhedron.

By combining Lemmas 6.4.4 and the same proofs as Lemmas 6.2.3 and 6.2.4, we get the following analogous statements.

Lemma 6.4.7. *Let (S, f) be a cyclic fatgraph bounding Γ. For any cutting vector v for (S, f), there is some scalar $k \in \mathbb{Q}$ so that $kv \in P_\Gamma$. Further,*

$$\frac{-\chi(kv)}{2} = \frac{-\chi(S)}{2n(S,f)}.$$

Lemma 6.4.8. *Let $\Gamma \in B_1^H(G)$ and $v \in P_\Gamma$ (not necessarily integral). Then there exists a scalar $k \in \mathbb{Q}$ and (possibly many) cyclic fatgraphs (S, f) bounding Γ such that kv is a cutting vector for (S, f). Furthermore,*

$$\frac{-\chi(v)}{2} = \frac{-\chi(S)}{2n(S,f)}.$$

6.5 The complete `scylla` algorithm

Proposition 6.5.1. *Let $\Gamma \in B_1^H(G)$. Then*

$$\mathrm{scl}(\Gamma) = \inf_{v \in P_\Gamma} \frac{-\chi(v)}{2}.$$

Proof. Given any admissible surface (S, f), we may apply Lemma 6.4.2 followed by Lemma 6.4.7 to obtain some vector $v \in P_\Gamma$ so that $-\chi(v)/2 \leq -\chi(S)/2n(S, f)$. Thus,

$$\inf_{v \in P_\Gamma} \frac{-\chi(v)}{2} \leq \inf_{(S,f)} \frac{-\chi(S)}{2n(S,f)}.$$

The reverse inequality is immediately implied by Lemma 6.4.8, so we have

$$\inf_{v \in P_\Gamma} \frac{-\chi(v)}{2} = \inf_{(S,f)} \frac{-\chi(S)}{2n(S,f)} = \mathrm{scl}(\Gamma).$$

□

6.6 Complexity and comparison with `scallop`

The complete `scylla` algorithm computes $\mathrm{scl}(\Gamma)$ in $\mathbb{Z}^{*k} * \mathbb{Z}/o_1\mathbb{Z} * \cdots \mathbb{Z}/o_m\mathbb{Z}$ with a linear programming problem with $O((K+1)|\Gamma|^3)$ columns and $O((K+1)|\Gamma|^2)$ rows, where $K = \sum_{i=1}^m o_i$. This follows from simply observing that there are fewer than $K|\Gamma|^3$ labeled group teeth, fewer than $|\Gamma|^3$ polygon pieces, and fewer than $(K+1)|\Gamma|^2$ constraints. The 1 is added so the statement remains true if $K = 0$.

The `scallop` algorithm computes scl (in free groups) with a linear program with $O(|\Gamma|^{2r})$ columns and $O(|\Gamma|^2)$ rows. This is non-rigorous (it is rigorous for alternating words), but quite accurate in general; on moderately sized chains of length up through 30, `scallop` was found to have a success rate of about 97%. A slight modification to the scallop algorithm in rank 2, which we will call m5, makes the number of columns $O(|\Gamma|^5)$ and makes the algorithm far more accurate. Completely rigorous `scallop` has $O((3|\Gamma|)^{2r})$ columns. For small rank, especially rank 2, the complexity for `scallop` is competitive with `scylla`. We provide the following chart, which indicates the algorithm which generally performs better with reasonably sized chains under reasonable conditions.

	`scallop`	`scylla`
rank 2 (nonrigorous)	*	
rank 2 (m5 nonrigorous)	*	
rank 2 rigorous		*
higher rank		*
finite free factors	basically N/A	*

It is interesting to note that if we fix word length and increase the rank, the running time of `scylla` actually *improves*. This is because it has far fewer rectangles to consider.

Let us now do a more detailed study. It is not difficult to see that in a random word v of length n in a rank r free group, we expect the following (we have disregarded some lower-order terms):

	`scallop`	`scylla`
Interface edges (and rectangles)	$\frac{n^2}{4r}$	$\frac{n^2}{2r}$
Central edges		$\frac{n^2}{2}$
Polygons	$\frac{n^2}{(2r)^2(2r-1)}\left(\frac{n}{2r}\right)^{2r-2}$	$\frac{n^3}{2r} + \frac{n^3}{(2r)^2} + \frac{n^3}{(2r)^2(2r-1)}$
Nonzeros per polygon	2r	3

Let us compare problem sizes by comparing the total number of matrix entries (other reasonable measures include number of nonzero entries and a weighted product of the rows and columns). For ranks 3 and higher, `scylla` quickly produces a smaller problem: in rank 3, the problem becomes smaller when $n > 72$. However, in rank 2, we must wait until $n > 1472$. Of course it is slightly unfair to compare `scallop` and `scylla` since `scylla` is rigorous, and we can mimic the behavior of `scallop` by considering only polygon pieces with a single central edge. This reduces the number of polygon pieces that `scylla` must consider, and causes `scylla` to generate a smaller (non-rigorous) linear program when $n > 320$.

Bibliography

[1] V. Bardakov, *Computation of commutator length in free groups*, Algebra and Logic **39** (2000), no. 4, 224–251.

[2] Christophe Bavard, *Longueur stable des commutateurs*, Enseign. Math. (2) **37** (1991), no. 1-2, 109–150. MR 1115747 (92g:20051)

[3] M. Bestvina, M. Feighn, and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, Geom. Funct. Anal. **7** (1997), no. 2, 215–244. MR 1445386 (98c:20045)

[4] Joan S. Birman and Caroline Series, *Geodesics with bounded intersection number on surfaces are sparsely distributed*, Topology **24** (1985), no. 2, 217–225. MR 793185 (87f:57012)

[5] Robert Brooks, *Some remarks on bounded cohomology*, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), Ann. of Math. Stud., vol. 97, Princeton Univ. Press, Princeton, N.J., 1981, pp. 53–63. MR 624804 (83a:57038)

[6] Danny Calegari, *Foliations and the geometry of 3-manifolds*, Oxford Mathematical Monographs, Oxford University Press, Oxford, 2007. MR 2327361 (2008k:57048)

[7] ———, *Faces of the scl norm ball*, Geom. Topol. **13** (2009), no. 3, 1313–1336. MR 2496047 (2010e:20062)

[8] ———, *scl*, MSJ Memoirs, vol. 20, Mathematical Society of Japan, Tokyo, 2009. MR 2527432 (2011b:57003)

[9] ———, *scl, sails, and surgery*, J. Topol. **4** (2011), no. 2, 305–326. MR 2805993

[10] Danny Calegari and Koji Fujiwara, *Stable commutator length in word-hyperbolic groups*, Groups Geom. Dyn. **4** (2010), no. 1, 59–90. MR 2566301 (2011a:20109)

[11] Danny Calegari and Joseph Maher, *Statistics and compression of scl*, preprint; arXiv:1008.4952.

[12] Danny Calegari and Alden Walker, *Circles and quasimorphisms*, in preparation.

[13] ______, *Integer hulls of linear polyhedra and scl in families*, Trans. Amer. Math. Soc., to appear.

[14] ______, *Random rigidity in the free group*, preprint; arXiv:1104.1768.

[15] ______, `scallop`, computer program available from the authors' webpages.

[16] ______, *Traintracks and endomorphisms*, in preparation.

[17] ______, *Isometric endomorphisms of free groups*, New York J. Math. **17** (2011), 713–743. MR 2851070

[18] Marc Culler, *Using surfaces to solve equations in free groups*, Topology **20** (1981), no. 2, 133–145. MR 605653 (82c:20052)

[19] David B. A. Epstein and Koji Fujiwara, *The second bounded cohomology of word-hyperbolic groups*, Topology **36** (1997), no. 6, 1275–1289. MR 1452851 (98k:20088)

[20] Koji Fujiwara, *The second bounded cohomology of a group acting on a Gromov-hyperbolic space*, Proc. London Math. Soc. (3) **76** (1998), no. 1, 70–94. MR 1476898 (99c:20072)

[21] Étienne Ghys, *Groupes d'homéomorphismes du cercle et cohomologie bornée*, The Lefschetz centennial conference, Part III (Mexico City, 1984), Contemp. Math., vol. 58, Amer. Math. Soc., Providence, RI, 1987, pp. 81–106. MR 893858 (88m:58024)

[22] ______, *Groups acting on the circle*, Enseign. Math. (2) **47** (2001), no. 3-4, 329–407. MR 1876932 (2003a:37032)

[23] Marshall Hall, Jr., *Subgroups of finite index in free groups*, Canadian J. Math. **1** (1949), 187–190. MR 0028836 (10,506a)

[24] John Hempel, *3-manifolds as viewed from the curve complex*, Topology **40** (1997), 631–657.

[25] Ilya Kapovich and Martin Lustig, *Invariant laminations for irreducible automorphisms of free groups*, preprint; arXiv:1104.1265.

[26] Gilbert Levitt and Martin Lustig, *Automorphisms of free groups have asymptotically periodic dynamics*, J. Reine Angew. Math. **619** (2008), 1–36. MR 2414945 (2009m:20034)

[27] Martin W. Liebeck, E. A. O'Brien, Aner Shalev, and Pham Huu Tiep, *The Ore conjecture*, J. Eur. Math. Soc. (JEMS) **12** (2010), no. 4, 939–1008. MR 2654085 (2011e:20016)

[28] A. Malcev, *On isomorphic matrix representations of infinite groups*, Rec. Math. [Mat. Sbornik] N.S. **8 (50)** (1940), 405–422. MR 0003420 (2,216d)

[29] Peter Scott, *Subgroups of surface groups are almost geometric*, J. London Math. Soc. (2) **17** (1978), no. 3, 555–565. MR 0494062 (58 #12996)

[30] John R. Stallings, *Topology of finite graphs*, Inventiones Mathematicae **71** (1983), 551–565, 10.1007/BF02095993.

[31] William P. Thurston, *Three-manifolds, foliations and circles, i*, preprint; arXiv math/9712268.

[32] Alden Walker, `wallop`, computer program available from the author's webpage.

[33] Dongping Zhuang, *Irrational stable commutator length in finitely presented groups*, J. Mod. Dyn. **2** (2008), no. 3, 499–507. MR 2417483 (2009k:20075)