STABLE COMMUTATOR LENGTH IN BAUMSLAG–SOLITAR GROUPS AND QUASIMORPHISMS FOR TREE ACTIONS

MATT CLAY, MAX FORESTER, AND JOEL LOUWSMA

Abstract. This paper has two parts, on Baumslag–Solitar groups and on general G–trees.

In the first part we establish bounds for stable commutator length (scl) in Baumslag–Solitar groups. For a certain class of elements, we further show that scl is computable and takes rational values. We also determine exactly which of these elements admit extremal surfaces.

In the second part we establish a universal lower bound of $1/12$ for scl of suitable elements of any group acting on a tree. This is achieved by constructing efficient quasimorphisms. Calculations in the group $BS(2,3)$ show that this is the best possible universal bound, thus answering a question of Calegari and Fujiwara. We also establish scl bounds for acylindrical tree actions.

Returning to Baumslag–Solitar groups, we show that their scl spectra have a uniform gap: no element has scl in the interval $(0, 1/12)$.

1. Introduction

Stable commutator length has been the subject of a significant amount of recent work, especially by Danny Calegari and his collaborators. See [6] for an introduction to stable commutator length and a desciption of much of this work. A major breakthrough in this area was Calegari's algorithm [7] for computing stable commutator length in free groups. This algorithm can also be used to compute stable commutator length in certain classes of groups that are built from free groups in simple ways. However, there are few other instances in which stable commutator length can be computed explicitly, with the exception of certain elements and classes of groups for which it is known to vanish.

Other work involves studying the *spectrum* of values taken by stable commutator length on a given group. In certain cases, this spectrum has been shown to have a *gap*, i.e. there is a range of values that are the stable commutator length of no element of the group. For example, results of this type have been shown for free groups [11], for word-hyperbolic groups [9], and recently for mapping class groups [2]. Such results have often involved constructing quasimorphisms with certain properties, thus relying on a dual interpretation of stable commutator length in terms of quasimorphisms.

The primary goal of this paper is to understand stable commutator length in Baumslag–Solitar groups. We obtain both quantitative and qualitative results. On the way to establishing the *gap theorem* below, we digress in Section 6 to construct efficient quasimorphisms in the completely general setting of groups acting on trees, and derive some consequences. These results may be of independent interest to some readers.

Stable commutator length in Baumslag–Solitar groups. We use the presentation $\langle\, a, t \mid ta^m t^{-1} = a^\ell \,\rangle$ for the Baumslag–Solitar group $BS(m, \ell)$, and we generally assume that $m \neq \ell$. Then, stable commutator length is defined exactly on the elements of t–exponent zero. We build on the approach taken in [4] and attempt to encode the computation of stable commutator length as the output of a linear programming problem. This approach used the notions of the *turn graph* and *turn circuits* to encode the geometric data of an admissible surface.

In the present setting, encoding this geometric data requires the use of a *weighted* turn graph instead, to account for winding numbers not present in the case of free groups. Even so, there is further winding data, and the natural encoding leads to an infinite-dimensional linear programming problem. By restricting to words of alternating t–shape, we are able to reduce to a finite-dimensional problem.

Theorem 1.1 (Theorem 5.2 and Corollary 5.3)**.** *Suppose $g \in BS(m, \ell)$, $m \neq \ell$, has alternating t–shape. Then there is a finite-dimensional, rational linear programming problem whose solution yields the stable commutator length of g. In particular,* $\mathrm{scl}(g)$ *is computable and is a rational number.*

More generally, the linear programming problem constructed in the proof of Theorem 1.1 is defined for any element g of t–exponent zero, and its solution provides a *lower bound* for $\mathrm{scl}(g)$ (see Theorem 4.3). What is difficult is to convert the solution into an admissible surface to obtain a matching upper bound; the encoding procedure from surfaces to vectors *loses* information, and not every vector can be realized by a surface.

In some cases the solution to the linear programming problem in Theorem 1.1 can be expressed in a closed formula. We show in Proposition 5.5 that if $m \nmid i$ and $\ell \nmid j$ then

$$\mathrm{scl}\big(ta^i t^{-1} a^j\big) \;=\; \frac{1}{2}\left(1 - \frac{\gcd(i,m)}{|m|} - \frac{\gcd(j,\ell)}{|\ell|}\right). \tag{1}$$

Next we characterize the elements of alternating t–shape for which there is a surface, known as an *extremal surface*, that realizes the infimum in the definition of stable commutator length. Such surfaces are important in applications of stable commutator length to problems in topology. It turns out that many elements have extremal surfaces, and many do not.

Theorem 5.7. *Let $g = \prod_{k=1}^{r} ta^{i_k} t^{-1} a^{j_k} \in BS(m, \ell)$, $m \neq \ell$. There is an extremal surface for g if and only if*

$$\ell \sum_{k=1}^{r} i_k = -m \sum_{k=1}^{r} j_k.$$

This allows us to find many examples of elements with rational stable commutator length for which no extremal surface exists. Previous examples of this phenomenon were found in free products of abelian groups of higher rank (see [8]).

Our last main result for Baumslag–Solitar groups is more qualitative in nature and concerns the scl spectrum.

Theorem 7.8 (Gap theorem)**.** *For every element $g \in BS(m,\ell)$, either* $\mathrm{scl}(g) = 0$ *or* $\mathrm{scl}(g) \geqslant 1/12$*.*

Thus, similar to hyperbolic groups, the spectrum has a gap above zero. This theorem is proved in Section 7, and it depends heavily on results in Section 6 to be discussed shortly. Nevertheless, these latter results do not apply to every element of $BS(m,\ell)$ (namely, those that are not *well-aligned*). To study stable commutator length of these left-over elements, we take advantage of special properties of the Bass–Serre trees for these groups. It is interesting to note that, in contrast with Theorem 6.11 below, it is the *failure* of acylindricity of these trees that is used in establishing the scl gap.

Stable commutator length in groups acting on trees. In order to prove the gap theorem we turn to the dual viewpoint of quasimorphisms on groups. According to Bavard Duality [1], a lower bound for $\mathrm{scl}(g)$ can be obtained by finding a homogeneous quasimorphism f on G with $f(g) = 1$ and of small defect. Indeed, if the defect of f is D then $\mathrm{scl}(g) \geqslant 1/2D$.

Many authors have constructed quasimorphisms on groups in settings involving negative curvature. For the most part these constructions are variants and generalizations of the Brooks counting quasimorphisms on free groups [18, 5, 15]. These settings include hyperbolic groups [12], groups acting on Gromov-hyperbolic spaces [13, 9], amalgamated free products and HNN extensions [14], and mapping class groups [3, 2].

One such result is Theorem D of [9], due to Calegari and Fujiwara. They showed that for any amalgamated product $G = A *_C B$ and any appropriately chosen hyperbolic element $g \in G$, there is a homogeneous quasimorphism f on G with $f(g) = 1$ and of defect at most 312. This bound is of interest since it is universal, independent of the group.

In Theorem 6.6 we construct efficient quasimorphisms, of defect at most 6, for any group acting on a tree. These are similar to the "small" counting quasimorphisms introduced by Epstein–Fujiwara [12], except that they are specifically tailored to the geometry of tree actions; the counting takes place in the tree rather than a Cayley graph. Moreover, by working directly with the homogenization of the counting quasimorphism, we obtain a further improvement in the defect.

Using the calculation (1) in the group $BS(2,3)$ (or alternatively, a different calculation in $\mathrm{PSL}(2,\mathbb{Z})$) we determine that 6 is the smallest possible defect that can be achieved in this generality, thus answering Question 8.4 of [9]. Expressed in terms of stable commutator length, the result can be stated as follows.

Theorem 6.9. *Suppose G acts on a simplicial tree T. If $g \in G$ is well-aligned then* $\mathrm{scl}(g) \geqslant 1/12$*.*

The same result holds for groups acting on $\mathbb{R}$–trees as well (Remark 6.7). Again, the bound of $1/12$ is the best possible. The condition of being *well-aligned* is necessary, and agrees with the double coset condition in [9] in the case of the Bass–Serre tree of an amalgam.

Not every hyperbolic element is well-aligned. Indeed, there are examples of 3–manifold groups that split as amalgams containing hyperbolic elements with very small stable commutator length; see [9]. If we consider trees that are *acylindrical* (see Section 6) then we can obtain an additional lower bound that applies to all hyperbolic elements. This bound is almost universal, depending only on the acylindricity constant. Alternatively, there is a genuinely uniform bound if one considers only elements of translation length greater than or equal to the acylindricity constant.

Theorem 6.11. *Suppose G acts K–acylindrically on a tree T and let N be the smallest integer greater than or equal to $\frac{K}{2}+1$.*

(i) *If $g \in G$ is hyperbolic then either* $\operatorname{scl}(g) = 0$ *or* $\operatorname{scl}(g) \geqslant 1/12N$.

(ii) *If $g \in G$ is hyperbolic and $|g| \geqslant K$ then either* $\operatorname{scl}(g) = 0$ *or* $\operatorname{scl}(g) \geqslant 1/24$.

In both cases, $\operatorname{scl}(g) = 0$ *if and only if g is conjugate to g^{-1}.*

Acknowledgments. Matt Clay is partially supported by NSF grant DMS-1006898. Max Forester is partially supported by NSF grant DMS-1105765.

2. Preliminaries

Stable commutator length. Stable commutator length may be defined as follows, according to Proposition 2.10 of [6].

Definition 2.1. Let $G = \pi_1(X)$ and suppose $\gamma\colon\ S^1 \to X$ represents the conjugacy class of $g \in G$. The *stable commutator length* of g is given by

$$\operatorname{scl}(g) \;=\; \inf_S \frac{-\chi(S)}{2n(S)}, \tag{2}$$

where S ranges over all singular surfaces $S \to X$ such that

- S is oriented and compact with $\partial S \neq \emptyset$
- S has no S^2 or D^2 components
- the restriction $\partial S \to X$ factors through γ; that is, there is a commutative diagram:

$$\begin{array}{ccc} \partial S & \longrightarrow & S \\ \big\downarrow & & \big\downarrow \\ S^1 & \xrightarrow{\ \gamma\ } & X \end{array}$$

- the total degree, $n(S)$, of the map $\partial S \to S^1$ (considered as a map of oriented 1–manifolds) is non-zero.

A surface S satisfying the conditions above is called an *admissible surface*. If, in addition, each component of ∂S maps to S^1 with *positive* degree, we call S a *positive admissible surface*. It is shown in Proposition 2.13 of [6] that the infimum in the definition of scl may be taken over positive admissible surfaces. Such surfaces (admissible or positive admissible) exist if and only if $g^k \in [G,G]$ for some nonzero integer k. If this does not occur then by convention $\mathrm{scl}(g) = \infty$ (the infimum of the empty set).

A surface $S \to X$ is said to be *extremal* if it realizes the infimum in (2). Notice that if this occurs, then $\mathrm{scl}(g)$ is a rational number.

In order to bound scl from above, one needs to construct an admissible surface realizing a given value of $\frac{-\chi(S)}{2n(S)}$. Sometimes a procedure for building a surface cannot be completed, leaving a surface with portions missing. The following result can be used in this situation.

Lemma 2.2. *Let S be a compact oriented surface with no S^2 or D^2 components, and whose boundary is expressed as two non-empty families of curves $\partial_1 S$ and $\partial_2 S$. Suppose $S \to X$ is a map taking the components of $\partial_1 S$ to group elements $a_1, \ldots, a_k \in \pi_1(X)$ and all components of $\partial_2 S$ to powers of the single element $g \in \pi_1(X)$, with total degree $n \neq 0$. Then there is an inequality*

$$\mathrm{scl}(g) \leqslant \frac{-\chi(S)}{2n} + \frac{1}{n}\Big(\sum_i \mathrm{scl}(a_i)\Big).$$

More generally, if one has defined scl for chains, the sum on the right hand side may be replaced by $\mathrm{scl}(\sum_i a_i)$, which may be finite even when the original sum was not.

Proof. We first show how to construct a cover of S that unwraps the curves in $\partial_1 S$ to give a collection of curves each of which is trivial in $H_1(X)$. Let b be the number of boundary components of S. Let c_i be the order of the conjugacy class of a_i in the abelianization of $\pi_1(X)$. If the conjugacy class of some a_i has infinite order in the abelianization of $\pi_1(X)$, then $\mathrm{scl}(a_i) = \infty$ and the lemma is tautological. Therefore we assume each c_i is finite. Let $M = \mathrm{lcm}(c_1, \ldots, c_k)$, and consider the prime factorization $M = p_1^{d_1} \cdots p_q^{d_q}$. We construct a tower of covers $S_q \to S_{q-1} \to \cdots \to S_1 \to S_0 = S$ as follows. For all i, the boundary ∂S_i will be partitioned into two families of curves $\partial_1 S_i$ and $\partial_2 S_i$, where the induced map $S_i \to X$ takes the curves in $\partial_1 S_i$ to powers of the elements $a_1, \ldots, a_k$ and the curves in $\partial_2 S_i$ to powers of the element g. For all i, $\partial_1 S_i$ will consist of exactly k curves and $\partial_2 S_i$ will consist of at least $b - k$ curves.

Suppose S_{i-1} has been constructed. Since $b - k \geqslant 1$, there is some integer e_i satisfying $k \leqslant e_i \leqslant b$ such that $e_i - 1$ is relatively prime to p_i, and hence to $p_i^{q_i}$. Therefore Lemma 1.12 of [6] shows that, for any e_i boundary components of S_{i-1}, there is a $p_i^{d_i}$–sheeted covering $S_i \to S_{i-1}$ that unwraps these e_i boundary components. We choose these e_i boundary components to be the k curves in $\partial_1 S_{i-1}$ and any $e_i - k$ curves in $\partial_2 S_{i-1}$. Then ∂S_i is also partitioned into two collections of curves: those in the preimage of $\partial_1 S_{i-1}$ are said to be in $\partial_1 S_i$, and those in the preimage of $\partial_2 S_{i-1}$ are said to be in $\partial_2 S_i$. By construction, $\partial_1 S_i$ consists of exactly k curves and $\partial_2 S_i$ consists of at least $b - k$ curves.

Iterating this procedure, we obtain a surface S_q that is a degree M cover of S. The induced map $S_q \to X$ takes the curves in $\partial_1 S_q$ to $a_1^M, \ldots, a_k^M$ and the curves in $\partial_2 S_q$ to powers of g with total degree nM. Note that, for each i, a_i^M is trivial in the abelianization of $\pi_1(X)$.

Fix $\epsilon > 0$. For all N relatively prime to $k-1$, we can construct a further cover $S_{q,N} \to S_q$ such that the curves in $\partial S_{q,N}$ are again partitioned into classes $\partial_1 S_{q,N}$ and $\partial_2 S_{q,N}$, where the curves in $\partial_1 S_{q,N}$ map to $a_1^{MN}, \ldots, a_k^{MN}$ in X. Choose N sufficiently large that, for all i, the element a_i^{MN} bounds an admissible surface S_i' that approximates $\operatorname{scl}\bigl(a_i^{MN}\bigr)$ to within ϵ/k. Since $\operatorname{scl}\bigl(a_i^{MN}\bigr) = MN \operatorname{scl}(a_i)$, we can also regard S_i' as an admissible surface for a_i that approximates $\operatorname{scl}(a_i)$ within ϵ/kMN. More precisely,

$$\frac{-\chi(S_i')}{2MN} \;\leqslant\; \operatorname{scl}(a_i) + \frac{\epsilon}{kMN}$$

for each i. Now join the surfaces S_i' along their boundaries to the corresponding curves in $\partial_1 S_{q,N}$. We thus obtain an admissible surface S'' for g, with $n(S'') = nMN$. We have

$$\begin{aligned} \frac{-\chi(S'')}{2n(S'')} &= \frac{-\chi(S_{q,N}) + \sum_i -\chi(S_i')}{2nMN} \\ &= \frac{-MN\chi(S) + \sum_i -\chi(S_i')}{2nMN} \\ &\leqslant \frac{-\chi(S)}{2n} + \frac{1}{n}\sum_i \Bigl(\operatorname{scl}(a_i) + \frac{\epsilon}{kNM}\Bigr) \\ &= \frac{-\chi(S)}{2n} + \frac{1}{n}\Bigl(\sum_i \operatorname{scl}(a_i)\Bigr) + \frac{\epsilon}{nMN}. \end{aligned}$$

Hence $\operatorname{scl}(g) \;\leqslant\; \frac{-\chi(S)}{2n} + \frac{1}{n}\bigl(\sum_i \operatorname{scl}(a_i)\bigr)$. □

Baumslag–Solitar groups. Before discussing Baumslag–Solitar groups per se, we make a general observation:

Lemma 2.3. *In any group G, if t and a are elements satisfying the Baumslag–Solitar relation $ta^m t^{-1} = a^\ell$ with $m \neq \ell$ then* $\operatorname{scl}(a) = 0$.

Proof. For any space X with fundamental group G there is a singular annulus $S \to X$, whose oriented boundary components represent a^m and $a^{-\ell}$ respectively (since a^m and a^ℓ are conjugate in G). This surface can be made admissible with $\chi(S) = 0$ and $n(S) = m - \ell \neq 0$, so $\operatorname{scl}(a) = 0$. □

The *Baumslag–Solitar group* $BS(m, \ell)$ is defined by the presentation

$$\langle\, a, t \mid ta^m t^{-1} = a^\ell \,\rangle. \tag{3}$$

The corresponding presentation 2–complex will be denoted $X_{m,\ell}$, or simply X, in this paper. One thinks of X as being constructed by attaching both ends of an annulus to a circle, by covering maps of degrees m and ℓ respectively; see Section 3.

Clearly, $BS(1,1)$ is $\mathbb{Z}\times\mathbb{Z}$ and $BS(1,-1)$ is the Klein bottle group. The cases $BS(m,\pm m)$ are also of special interest. By constructing a suitable covering space of X, one finds that this group contains a subgroup of index $2m$ isomorphic to $F_{2m-1}\times\mathbb{Z}$. In particular, stable commutator length can be computed in $BS(m,\pm m)$ and is always rational, using the rationality theorem for free groups [7] and results from [6] (such as Proposition 2.80) on subgroups of finite index.

In this paper we will study stable commutator length in $BS(m,\ell)$ under the standing assumption that $m\neq\ell$.

Remark 2.4. The abelianization of $BS(m,\ell)$ is $\mathbb{Z}\times\mathbb{Z}_{|m-\ell|}$ with generators t and a respectively. Since we are assuming that $m\neq\ell$, an element of $BS(m,\ell)$ has finite order in the abelianization if and only if it has t–exponent zero. Thus scl is finite on exactly these elements.

Definition 2.5. Given a word w in the letters $a^{\pm1}$ and $t^{\pm1}$ we denote by $|w|_t$ the *t–length* of w. That is, $|w|_t$ is the number of occurrences of t and t^{-1} in w.

Given an element $g\in BS(m,\ell)$ we denote by $|g|_t$ the *t–length* of the conjugacy class of g. That is, $|g|_t$ is the minimum value of $|w|_t$ over all words w that represent a conjugate of g.

Remark 2.6. Any element $g\in BS(m,\ell)$ has a conjugate that can be expressed as

$$w=t^{\epsilon_1}a^{k_1}t^{\epsilon_2}\cdots t^{\epsilon_n}a^{k_n}, \tag{4}$$

where:

- $\epsilon_i\in\{1,-1\}$ for $i=1,\dots,n$,
- $m\nmid k_i$ if $\epsilon_i=1$ and $\epsilon_{i+1}=-1$,
- $\ell\nmid k_i$ if $\epsilon_i=-1$ and $\epsilon_{i+1}=1$, and
- $|g|_t=|w|_t=n$.

The subscripts in the second and third bullet are read modulo n. We refer to such a representative word of the conjugacy class of g as *cyclically reduced.*

Up to cyclic permutation, the cyclically reduced word representing a conjugacy class is not unique. Two other modifications to the word (4) can be made, resulting in cyclically reduced words representing the same element:

$$a^ita^j \;\leftrightarrow\; a^{i-\ell}ta^{j+m} \quad\text{and}\quad a^it^{-1}a^j \;\leftrightarrow\; a^{i+m}t^{-1}a^{j-\ell}.$$

Collins' Lemma [10, 17] characterizes precisely when two cyclically reduced words represent the same conjugacy class. It implies easily that modulo the two moves above and cyclic permutation, the expression (4) is unique.

3. Surfaces in $X_{m,\ell}$

Transversality. Transversality will be used to convert a singular admissible surface $S\to X$ into a more combinatorial object. We will follow the approach from [4], which treated the case of surfaces mapping into graphs.

Recall that $X=X_{m,\ell}$ is the presentation 2–complex for the presentation (3). We can build X in the following way. Let A be the annulus $S^1\times[-1,1]$, and let C be a space homeomorphic to the circle. Fix orientations of S^1 and C and attach the boundary circles $S^1 \times \{\pm 1\}$ to C via covering maps of degrees m and ℓ respectively, to form X. Note that the natural map $\phi\colon\ A \to X$ is surjective, and maps the interior of A homeomorphically onto $X - C$. Thus we have an identification of $X - C$ with $S^1 \times (-1, 1)$.

The space X is also a cell complex with C as a subcomplex. The 1–skeleton of X may be taken to be C (having one 0–cell and one 1–cell, labeled a) along with an additional 1–cell labeled t, which is a fiber in A whose endpoints are attached to the 0–cell of C.

Let $C' = S^1 \times \{0\} \subset X - C$. This is a codimension-one submanifold. For any compact surface S and continuous map $f\colon\ S \to X$, we may perturb f by a small homotopy to make it *transverse* to C'. Then, $f^{-1}(C')$ is a properly embedded codimension-one submanifold $N \subset S$. By a further homotopy, we can arrange that N has an embedded I–bundle neighborhood $N \times [-1, 1] \subset S$ (with $N = N \times \{0\}$) such that $f^{-1}(X - C) = N \times (-1, 1)$ and

$$f|_{N\times(-1,1)}\colon\ N \times (-1, 1) \to S^1 \times (-1, 1)$$

is a map of the form $f_0 \times \mathrm{id}$.

Let $N_b \subset N$ be the union of the components that are intervals (rather than circles). Let $S_b \subset S$ be the subset $N_b \times [-1, 1]$, each component of which is a *band* $I \times [-1, 1]$ with $(I \times [-1, 1]) \cap \partial S = \partial I \times [-1, 1]$.

By a further homotopy of f in a neighborhood of ∂S, and using transversality for the map $S - (N \times (-1, 1)) \to C$, we can arrange that in addition to the structure given so far, there is a collar neighborhood $S_\partial \subset S$ on which f has a simple description. This map takes S_∂ into the 1–skeleton of X by a retraction onto ∂S followed by the restriction $\partial S \to X$. Each annulus component of S_∂ decomposes into squares that retract into ∂S and then map to X by the characteristic maps of 1–cells. These squares are labeled a– or t–squares depending on the 1–cell. The t–squares are exactly the components of $S_\partial \cap S_b$. In particular, each band ends in two t-squares, representing one instance each of t and t^{-1} along the boundary. See Figure 1.

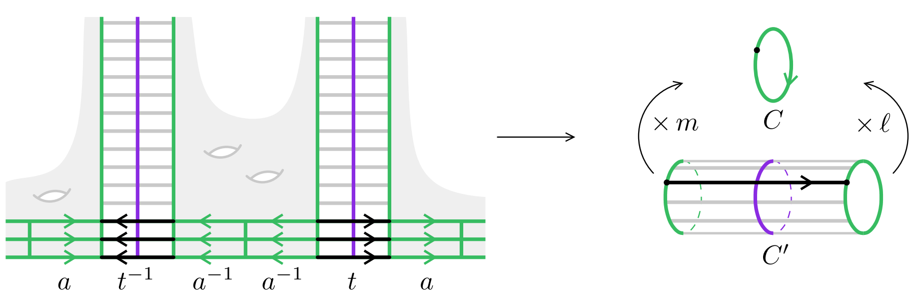

FIGURE 1. An admissible surface after the transversality procedure. The gray regions map into C.

Finally, we define $S_0 = S_\partial \cup S_b$ and $S_1 = S - \mathrm{int}(S_0)$. Observe that f maps ∂S_1 into C.

The boundary ∂S_0 decomposes into two subsets: ∂S, called the *outer boundary*, and components in the interior of S, called the *inner boundary*, denoted $\partial^- S_0$. Note that $\partial^- S_0 = S_0 \cap S_1 = \partial S_1$. In particular, components of the inner boundary map by f to loops in X representing conjugacy classes of powers of a.

Remark 3.1. Call a loop $f\colon S^1 \to X$ *regular* if S^1 can be decomposed into vertices and edges such that the restriction of f to each edge factors through the characteristic map of a 1–cell of X. Note that a regular map is completely described (up to reparametrization) by a cyclic word in the generators $a^{\pm 1}$, $t^{\pm 1}$ representing the conjugacy class of f in $\pi_1(X)$.

If a singular surface $S \to X$ has the property that its restriction to each boundary component is regular, then the transversality procedure described above can be performed rel boundary, so that the cyclic orderings of oriented a– and t–squares in S_∂ agree with the cyclic boundary words one started with.

Recall that $\mathrm{scl}(g)$ is the infimum of $\frac{-\chi(S)}{2n(S)}$ over all positive admissible surfaces. We will show how to compute $\mathrm{scl}(g)$ using the decomposition described above.

Choose a cyclically reduced word w representing the conjugacy class of g. For any positive admissible surface S, each boundary component maps by a loop representing a positive power of g in $\pi_1(X)$. Modify f by a homotopy to arrange that its boundary maps are regular, with corresponding cyclic words equal to positive powers of w. Then perform the transversality procedure given above, keeping the boundary map fixed (cf. Remark 3.1). At this point, the subsurfaces S_0, S_1 are defined. Each boundary component is labeled by a positive power of w and these powers add to $n(S)$.

Note that $\chi(S) = \chi(S_0) + \chi(S_1)$ since S_0 and S_1 meet along circles. Also,

$$\chi(S_0) \;=\; \frac{-n(S)\,|g|_t}{2},$$

as this is exactly the number of bands in S_b, each band connecting two instances of $t^{\pm 1}$ in $w^{n(S)}$ and contributing -1 to $\chi(S_0)$. (Note that $\chi(S_\partial) = 0$.) Let $\chi^+(S_1)$ denote the number of disk components in S_1. We have

$$\chi(S) \;=\; \frac{-n(S)\,|g|_t}{2} + \chi(S_1) \;\leqslant\; \frac{-n(S)\,|g|_t}{2} + \chi^+(S_1),$$

and therefore

$$\frac{-\chi(S)}{2n(S)} \;\geqslant\; \frac{|g|_t}{4} \;-\; \frac{\chi^+(S_1)}{2n(S)}.$$

From this, we conclude that

$$\mathrm{scl}(g) \;\geqslant\; \frac{|g|_t}{4} \;+\; \inf_S \frac{-\chi^+(S_1)}{2n(S)}, \tag{5}$$

where the infimum is taken over all positive admissible surfaces. In fact, the reverse of inequality (5) holds as well:

Lemma 3.2. *There is an equality*

$$\mathrm{scl}(g) \;=\; \frac{|g|_t}{4} \;+\; \inf_S \frac{-\chi^+(S_1)}{2n(S)}.$$

Proof. Given an admissible surface $S \to X$ decomposed as above, let S' be the union of S_0 and the disk components of S_1. Recall that the components of $\partial S'$ in $\partial^- S_0$ map to loops in X representing conjugacy classes of powers of a. Thus Lemma 2.2 and Lemma 2.3 imply

$$\mathrm{scl}(g) \;\leqslant\; \frac{-\chi(S')}{2n(S)} + \frac{1}{n(S)}\sum \mathrm{scl}(a^{p_i}) \;=\; \frac{-\chi(S')}{2n(S)}.$$

Since

$$\frac{-\chi(S')}{2n(S)} \;=\; \frac{-\chi(S_0)}{2n(S)} - \frac{\chi^+(S_1)}{2n(S)} \;=\; \frac{|g|_t}{4} - \frac{\chi^+(S_1)}{2n(S)}$$

and S was arbitrary, the reverse of inequality (5) holds, as desired. □

Lemma 3.3. *If S is an extremal surface for g, then S_1 consists only of disks and annuli.*

Proof. Let S_2 be the union of the components of S_1 that have nonnegative Euler characteristic, and let S_3 be the union of the components of S_1 that have negative Euler characteristic. Then S_2 consists only of disks and annuli and $\chi(S_2) = \chi^+(S_1)$. If S is extremal, we must have

$$\mathrm{scl}(g) \;=\; \frac{-\chi(S)}{2n(S)} \;=\; \frac{|g|_t}{4} - \frac{\chi(S_2)}{2n(S)} - \frac{\chi(S_3)}{2n(S)} \;=\; \frac{|g|_t}{4} - \frac{\chi^+(S_1)}{2n(S)} - \frac{\chi(S_3)}{2n(S)}.$$

Comparing with Lemma 3.2, this means $\chi(S_3) \geqslant 0$, meaning that S_3 must be empty. Thus S_1 consists only of disks and annuli. □

The weighted turn graph. As in [4], we use a graph to keep track of the combinatorics of the inner boundary $\partial^- S_0$.

Consider a cyclically reduced word w as in (4). A *turn* in w is a subword of the form a^k between two occurrences of $t^{\pm 1}$ considered as a cyclic word. The turns are indexed by the numbers $i = 1, \ldots, n$; the i^{th} turn is labeled by the subword $t^{\epsilon_i} a^{k_i} t^{\epsilon_{i+1}}$. A turn labeled $ta^k t^{-1}$ is of *type m*; a turn labeled $t^{-1}a^k t$ is of *type ℓ*; all other turns are of *mixed* type.

The *weighted turn graph* $\Gamma(w)$ is a directed graph with integer weights assigned to each vertex. The vertices correspond to the turns of w and the weight associated to the i^{th} turn is k_i. There is a directed edge from turn i to turn j whenever $-\epsilon_i = \epsilon_{j+1}$. In other words, if the label of a turn begins with $t^{\pm 1}$, then there is a directed edge from this turn to every other turn whose label ends with $t^{\mp 1}$. The vertices of $\Gamma(w)$ are partitioned into four subsets where the presence of a directed edge between two vertices depends only on which subsets the vertices lie in. Figure 2 shows the turn graph in schematic form.

The edges of the turn graph come in *dual pairs*: if $e \in \Gamma(w)$ is an edge from turn i to turn j, then one verifies easily that there is also an edge $\bar{e}$ from turn $j+1$ to turn $i-1$, and moreover $\bar{\bar{e}} = e$. See Figure 3.

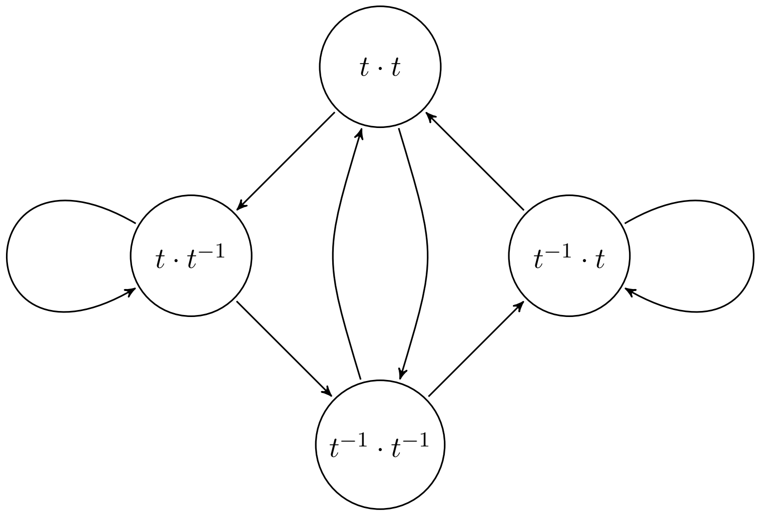

FIGURE 2. A schematic picture of the turn graph $\Gamma(w)$.

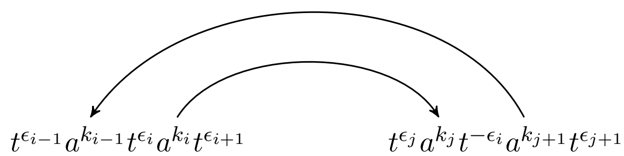

FIGURE 3. Dual edge pairs in w. Whenever there is an edge from turn i to turn j, there must also be an edge from turn $j+1$ to turn $i-1$. The labels of these turns share occurrences of $t^{\pm 1}$ with the labels of the original two turns.

A directed circuit in $\Gamma(w)$ is of *type* m or *type* ℓ if every vertex it visits corresponds to a turn of type m or of type ℓ, respectively. Otherwise, the circuit is of *mixed* type. The *weight* $\omega(\gamma)$ of a directed circuit γ is the sum of the weights of the vertices it visits (counted with multiplicity). Given a directed circuit γ, define

$$\mu(\gamma) = \begin{cases} m & \text{if } \gamma \text{ is of type } m \\ \ell & \text{if } \gamma \text{ is of type } \ell \\ \gcd(m,\ell) & \text{otherwise.} \end{cases}$$

A directed circuit γ is a *potential disk* if $\omega(\gamma) \equiv 0 \mod \mu(\gamma)$.

Turn circuits. Let $S \to X$ be a positive admissible surface whose boundary map is regular and labeled by $w^{n(S)}$. Decomposing S as $S_0 \cup S_1$, each inner boundary component of S_0 can be described as follows. Traversing the curve in the positively oriented direction, one alternately follows the boundary arcs (or *sides*) of bands in S_b and visits turns of w along S_∂; such a visit consists in traversing the inner edges of some a–squares before proceeding up along another side of a band (cf. Figure 1).

If the side of the band leads from turn i to turn j, then $(t^{\epsilon_i})^{-1} = t^{\epsilon_j}$ and therefore there is an edge in $\Gamma(w)$ from turn i to turn j. In this way, $\partial^- S_0$ gives rise to a finite collection (possibly with repetitions) of directed circuits in $\Gamma(w)$, called the *turn circuits* for S_0.

Since ∂S is labeled by $w^{n(S)}$, there are $n(S)$ occurrences of each turn on ∂S. The turn circuits do not contain the information of which particular instances of turns are joined bands, nor do they record how many times the band corresponding to a given edge in the circuit wraps around the annulus $X - C$.

Remark 3.4. Given two cyclically reduced words w, w' representing the same conjugacy class in $BS(m, \ell)$, there is an isomorphism $\Gamma(w) \to \Gamma(w')$ of the underlying directed graph structure that respects vertex type and edge duality but not necessarily the vertex weights. However, a directed circuit is a potential disk with one sets of weights if and only if it is a potential disk with the other set. The difference in weights of a type m vertex is a multiple of m, the difference in weights of a type ℓ vertex is a multiple of ℓ, and the difference in weights of mixed type vertex is a multiple of $\gcd(m, \ell)$. See Remark 2.6.

In what follows, only the property of being a potential disk is used and therefore this ambiguity in the weighed turn graph associated to a conjugacy class is not an issue.

Lemma 3.5. *Suppose γ is a turn circuit for S_0 that corresponds to an inner boundary component in $\partial^- S_0$ that bounds a disk in S_1. Then γ is a potential disk.*

Proof. For any band in S_b, the core arc (a component of N_b) maps to C' as a loop of some degree d. The two sides then map to C as loops of degrees dm and $d\ell$ respectively.

If ι is the side of a band that leads from a turn labeled $ta^k t^*$ to a turn labeled $t^* a^{k'} t^{-1}$ then the map $\iota \to C$ has degree a multiple of m. Likewise if ι leads from a turn labeled $t^{-1} a^k t^*$ to a turn labeled $t^* a^{k'} t$ then ι maps to C with degree a multiple of ℓ.

Therefore the total degree of an inner boundary component corresponding to a turn circuit γ is $\omega(\gamma) + dm + d'\ell$ for some integers d, d'. If γ is of type m then $d' = 0$. Likewise, if γ is of type ℓ then $d = 0$.

If the boundary component actually bounds a disk in S_1 then this total degree is 0. Hence $\omega(\gamma) \equiv 0 \mod \mu(\gamma)$ and therefore γ is a potential disk. ☐

4. Linear optimization

We would like to convert the optimization problem in Lemma 3.2 to a problem of optimizing a certain linear functional on a vector space whose coordinates correspond to possible potential disks, subject to certain linear constraints. Here the functional would count the number of potential disks, and the constraints would arise from the pairing of edges in the turn graph. The objective would then be to compute stable commutator length using classical linear programming.

The main difficulty in such an approach is arranging that the optimization takes place over a *finite dimensional* object. In this section, we show how to convert an admissible surface to a vector in a finite dimensional vector space in such a way that the number of disk components of S_1 is less than the value of an appropriate linear functional. We thus obtain computable, rational lower bounds for the stable commutator length of elements of Baumslag–Solitar groups (Theorem 4.3). In Section 5, we will show that these bounds are sharp for a certain class of elements.

We construct the finite dimensional vector space as follows. Let w be a conjugate of g of the form given in Remark 2.6. Let $M = \max\{|m|, |\ell|\}$. We consider two sets of directed circuits in $\Gamma(w)$:

- $\mathfrak{X}$: the set of potential disks that are a sum of not more than M embedded circuits, and
- $\mathfrak{Y}$: the set of all embedded circuits.

Note that both $\mathfrak{X}$ and $\mathfrak{Y}$ are finite sets and that they may have some circuits in common. Enumerate these sets as $\mathfrak{X} = \{\alpha_1, \ldots, \alpha_p\}$ and $\mathfrak{Y} = \{\beta_1, \ldots, \beta_q\}$. Let $\mathbb{X}$ be a p–dimensional real vector space with basis $\{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$, and let $\mathbb{Y}$ be a q–dimensional real vector space with basis $\{\mathbf{y}_1, \ldots, \mathbf{y}_q\}$. Equip both $\mathbb{X}$ and $\mathbb{Y}$ with an inner product that makes the respective bases orthonormal. By Remark 3.4, the vector spaces $\mathbb{X}$ and $\mathbb{Y}$ depend only on the conjugacy class in $BS(m, \ell)$ represented by w. Abusing notation, we let $\{\mathbf{x}_1, \ldots, \mathbf{x}_p, \mathbf{y}_1, \ldots, \mathbf{y}_q\}$ denote the corresponding orthonormal basis of $\mathbb{X} \oplus \mathbb{Y}$. This is the vector space with which we will work.

The linear functional on the vector space $\mathbb{X} \oplus \mathbb{Y}$ whose values will be compared with the number of disk components of S_1 is the functional that is the sum of the coordinates corresponding to $\mathbb{X}$, i.e. the functional that takes in $\mathbf{u} \in \mathbb{X} \oplus \mathbb{Y}$ and gives out $|\mathbf{u}|_{\mathbb{X}} := \sum_{i=1}^{p} \mathbf{u} \cdot \mathbf{x}_i$. One thinks of this functional as counting the number of potential disks.

There are additional linear functionals on $\mathbb{X} \oplus \mathbb{Y}$ that count the number of times turn circuits in a given collection visit a specific vertex or edge. For each vertex $v \in \Gamma(w)$, define $F_v \colon \mathbb{X} \oplus \mathbb{Y} \to \mathbb{R}$ by letting $F_v(\mathbf{x}_i)$ be the number of times α_i visits v, letting $F_v(\mathbf{y}_i)$ be the number of times β_i visits v, and extending by linearity. For each edge $e \subset \Gamma(w)$, define $F_e \colon \mathbb{X} \oplus \mathbb{Y} \to \mathbb{R}$ by letting $F_e(\mathbf{x}_i)$ be the number of times α_i traverses e, letting $F_v(\mathbf{y}_i)$ be the number of times β_i traverses e, and extending by linearity.

One thinks of the next lemma as saying that, if a collection of turn circuits traverses each edge the same number of times as its dual edge, then this collection of turn circuits visits each vertex the same number of times.

Lemma 4.1. *If $F_e(\mathbf{u}) = F_{\bar{e}}(\mathbf{u})$ for each dual edge pair $e, \bar{e}$ of $\Gamma(w)$, then $F_v(\mathbf{u}) = F_{v'}(\mathbf{u})$ for any vertices $v, v' \in \Gamma(w)$.*

Proof. For a vertex $v \in \Gamma(w)$, let $E^+(v)$ be the set of directed edges that are outgoing from v and let $E^-(v)$ be the set directed edges that are incoming to v. Then

$$F_v(\mathbf{u}) = \sum_{e \in E^+(v)} F_e(\mathbf{u}) = \sum_{e \in E^-(v)} F_e(\mathbf{u}).$$

First suppose that v corresponds to turn i and v' corresponds to turn $i-1$. Then edge duality gives a pairing between edges in $E^+(v)$ and edges in $E^-(v')$. Since

$F_e(\mathbf{u}) = F_{\bar{e}}(\mathbf{u})$ for every dual edge pair $e, \bar{e}$, we have that

$$F_v(\mathbf{u}) = \sum_{e \in E^+(v)} F_e(\mathbf{u}) = \sum_{e \in E^+(v)} F_{\bar{e}}(\mathbf{u}) = \sum_{e \in E^-(v')} F_e(\mathbf{u}) = F_{v'}(\mathbf{u}).$$

Letting i vary, we obtain a similar statement for all pairs of vertices corresponding to adjacent turns. It follows that $F_v(\mathbf{u}) = F_{v'}(\mathbf{u})$ for any vertices $v, v' \in \Gamma(w)$. □

Let $C \subset \mathbb{X} \oplus \mathbb{Y}$ be the cone of non-negative vectors $\mathbf{u}$ such that $F_e(\mathbf{u}) = F_{\bar{e}}(\mathbf{u})$ for every dual edge pair $e, \bar{e}$ of $\Gamma(w)$. In light of the lemma, we denote by $F\colon C \to \mathbb{R}$ the function $F_v\big|_C$ for any vertex $v \in \Gamma(w)$.

The following proposition shows how to convert an admissible surface into a vector $\mathbf{u} \in C$ in such a way that $|\mathbf{u}|_{\mathbb{X}}$ is at least the number of disk components of S_1.

Proposition 4.2. *Given an admissible surface $S \to X$, there is a vector $\mathbf{u} \in C$ such that*

$$\frac{|\mathbf{u}|_{\mathbb{X}}}{F(\mathbf{u})} \geqslant \frac{\chi^+(S_1)}{n(S)}. \tag{6}$$

Proof. Suppose the surface S has been decomposed as $S_0 \cup S_1$ as described in Section 3, and consider the collection of turn circuits for S_0. Let γ be one of the turn circuits in this collection. As a cycle we can decompose γ as a sum of embedded circuits, i.e. $\gamma = \beta_{i_1} + \cdots + \beta_{i_k}$. This decomposition may not be unique, but we only need its existence.

For each turn circuit γ, we will construct a corresponding vector $\mathbf{u}(\gamma)$, depending on this decomposition and on whether the corresponding boundary component of $\partial^- S_0$ bounds a disk in S_1. If the corresponding inner boundary component of $\partial^- S_0$ does not bound a disk in S_1, we define

$$\mathbf{u}(\gamma) = \sum_{j=1}^{k} \mathbf{y}_{i_j}.$$

Otherwise, the corresponding inner boundary component of $\partial^- S_0$ does bound a disk in S_1, in which case Lemma 3.5 implies that γ is a potential disk. If $k \leqslant M$, then $\gamma \in \mathfrak{X}$; say $\gamma = \alpha_i$. In this case we define $\mathbf{u}(\gamma) = \mathbf{x}_i$. Otherwise, if $k > M$, we proceed as follows. For each β_{i_j}, let $\mu(\beta_{i_j})\beta_{i_j}$ denote the sum of $\mu(\beta_{i_j})$ copies of β_{i_j}. Notice that $\mu(\beta_{i_j})\beta_{i_j}$ is a potential disk that is not the sum of more than M embedded circuits. Hence $\mu(\beta_{i_j})\beta_{i_j} \in \mathfrak{X}$, so $\mu(\beta_{i_j})\beta_{i_j} = \alpha_{i'_j}$ for some $i'_j \in \{1, \ldots, p\}$. In this case we define

$$\mathbf{u}(\gamma) = \sum_{j=1}^{k} \frac{1}{\mu(\beta_{i_j})} \mathbf{x}_{i'_j}.$$

The vector we will consider is $\mathbf{u} = \sum_\gamma \mathbf{u}(\gamma)$, where this sum is taken over all γ in the collection of turn circuits for S_0 (with multiplicity). Establishing the following three claims will complete the proof of the proposition.

(i) $\mathbf{u} \in C$,
(ii) $F(\mathbf{u}) = n(S)$, and
(iii) $|\mathbf{u}|_{\mathbb{X}} \geqslant \chi^+(S_1)$.

(i): The vector $\mathbf{u}(\gamma)$ was constructed so that $F_e(\mathbf{u}(\gamma))$ counts the number of times the turn circuit γ traverses the edge e. Thus $F_e(\mathbf{u})$ records the number of times turn circuits for S_0 traverse e. Every time an edge e is traversed by a turn circuit for S_0, there is a band in S_b one side of which represents e. The other side of this band represents $\bar{e}$, so therefore we have that $F_e(\mathbf{u}) = F_{\bar{e}}(\mathbf{u})$ for all edges e. Thus $\mathbf{u} \in C$.

(ii): The vector $\mathbf{u}(\gamma)$ was also constructed so that $F_v(\mathbf{u}(\gamma))$ counts the number of times the turn circuit γ visits the vertex v. Therefore $F(\mathbf{u})$ records the number of times turn circuits for S_0 visit any given vertex. As each turn occurs once in w, each vertex must be visited exactly $n(S)$ times. Thus $F(\mathbf{u}) = n(S)$.

(iii): Let γ be a turn circuit for S_0, and suppose the corresponding inner boundary component of S_0 bounds a disk in S_1. Decompose γ as a sum $\beta_{i_1} + \cdots + \beta_{i_k}$ of embedded circuits as above. If $k \leqslant M$, then $\mathbf{u}(\gamma) = \mathbf{x}_i$ for some $i \in \{1, \ldots, p\}$, and thus $|\mathbf{u}(\gamma)|_{\mathbb{X}} = 1$. Otherwise, we have

$$|\mathbf{u}(\gamma)|_{\mathbb{X}} = \sum_{j=1}^{k} \frac{1}{\mu(\beta_{i_j})} \geqslant \sum_{j=1}^{k} \frac{1}{M} = \frac{k}{M} \geqslant 1.$$

In either case, $|\mathbf{u}(\gamma)|_{\mathbb{X}} \geqslant 1$. As $|\cdot|_{\mathbb{X}}$ is a linear functional, we thus have that $|\mathbf{u}|_{\mathbb{X}} \geqslant \chi^+(S_1)$. □

Theorem 4.3. *Let $g \in BS(m, \ell)$, $m \neq \ell$, be of t–exponent zero. Then there is a computable, finite sided, rational polyhedron $P \subset \mathbb{X} \oplus \mathbb{Y}$ such that*

$$\operatorname{scl}(g) \geqslant \frac{|g|_t}{4} - \frac{1}{2} \max\left\{ |\mathbf{u}|_{\mathbb{X}} \mid \mathbf{u} \text{ is a vertex of } P \right\}. \tag{7}$$

Proof. Let $P = F^{-1}(1)$. If V is the number of vertices of $\Gamma(w)$, we can extend F to a linear functional $\widetilde{F} \colon \mathbb{X} \oplus \mathbb{Y} \to \mathbb{R}$ by setting $\widetilde{F}(\mathbf{u}) = \frac{1}{V} \sum_v F_v(\mathbf{u})$, where the sum is taken over all vertices of $\Gamma(w)$. The linear functional $\widetilde{F}$ is positive on all basis vectors of $\mathbb{X} \oplus \mathbb{Y}$, and hence a level set of $\widetilde{F}$ intersects the positive cone in a compact set. Clearly $P = C \cap \widetilde{F}^{-1}(1)$. Thus P is a finite sided, rational, compact polyhedron.

By Lemma 3.2, Proposition 4.2, and the linearity of $|\cdot|_{\mathbb{X}}$ and F, we have that

$$\begin{aligned}
\operatorname{scl}(g) &= \frac{|g|_t}{4} + \frac{1}{2} \inf_S \frac{-\chi^+(S_1)}{n(S)} \\
&= \frac{|g|_t}{4} - \frac{1}{2} \sup_S \frac{\chi^+(S_1)}{n(S)} \\
&\geqslant \frac{|g|_t}{4} - \frac{1}{2} \sup_{\mathbf{u} \in C} \frac{|\mathbf{u}|_{\mathbb{X}}}{F(\mathbf{u})} \\
&\geqslant \frac{|g|_t}{4} - \frac{1}{2} \sup_{\mathbf{u} \in P} |\mathbf{u}|_{\mathbb{X}} .
\end{aligned}$$

As P is a finite sided, compact polyhedron and $|\cdot|_{\mathbb{X}}$ is a linear functional, the supremum is realized at one of the vertices of P. This gives (7). □

Remark 4.4. If $\mathbf{u}$ is a vertex of P that maximizes $|\mathbf{u}|_{\mathbb{X}}$ in P, then $\mathbf{u} \cdot \mathbf{y}_i = 0$ for all $i \in \{1, \ldots, q\}$. Indeed, suppose not and let $\beta_i \in \mathfrak{Y}$ be such that $\mathbf{u} \cdot \mathbf{y}_i = c > 0$. Then there is some $\alpha_{i'} \in \mathfrak{X}$ such that $\mu(\beta_i)\beta_i = \alpha_{i'}$. One then observes that $\mathbf{u}' = \mathbf{u} - c\mathbf{y}_i + \frac{c}{\mu(\beta_i)}\mathbf{x}_{i'} \in P$ and $|\mathbf{u}'|_{\mathbb{X}} > |\mathbf{u}|_{\mathbb{X}}$.

The linear programming problem described in this section has been implemented using `Sage` [19] and is available from the first author's webpage. The number of embedded circuits in $\Gamma(w)$ is on the order of $|w|_t!$ and so the algorithm is only practical for elements with small t–length.

5. Elements of alternating t–shape

The bounds given in Theorem 4.3 are not always sharp, as we will point out in Remark 5.4. However, we show in Theorem 5.2 that these bounds are sharp for a class of elements that have what we call *alternating t–shape*. We thus show that stable commutator length is computable and rational for such elements. We also characterize which elements of alternating t–shape admit extremal surfaces (Theorem 5.7).

Definition 5.1. We say that an element $g \in BS(m, \ell)$ has *alternating t–shape* if it has a conjugate of the form given in Remark 2.6 where n is even and $\epsilon_i = (-1)^{i-1}$.

In this section, we restrict attention to elements of alternating t–shape and express this conjugate as

$$w = \prod_{k=1}^{r} t a^{i_k} t^{-1} a^{j_k}. \tag{8}$$

Note that if g has alternating t–shape then it has t–exponent zero. Hence stable commutator length is finite for elements of alternating t–shape.

Constructing surfaces. Let P be as in the proof of Theorem 4.3, and let

$$L(g) = \frac{|g|_t}{4} - \frac{1}{2} \max \left\{ |\mathbf{u}|_{\mathbb{X}} \mid \mathbf{u} \text{ is a vertex of } P \right\}.$$

To show that the lower bound $L(g)$ on stable commutator length is sharp, we would like to find a surface S that gives the same upper bound on stable commutator length. Specifically, given a vertex $\mathbf{u} \in P$, we want to construct a corresponding surface $S = S_0 \cup S_1$ of the type discussed in Section 3, where ∂S maps to loops representing conjugacy classes of powers of g and ∂S_1 maps to loops representing conjugacy classes of powers of a. Such a surface S_0 can be built (in fact, many such surfaces can be built); the construction is given in the proof of Theorem 5.2. The difficulty is arranging S_0 so that its inner boundary components can be efficiently capped off by S_1.

If the degree of each inner boundary component of S_0 were zero, we could take each component of S_1 to be a disk. In this case, we would have $|\mathbf{u}|_{\mathbb{X}} = \chi^+(S_1) = \chi(S_1)$ and

$$\operatorname{scl}(g) \leqslant \frac{-\chi(S)}{2n(S)} = \frac{|g|_t}{4} - \frac{\chi(S_1)}{2n(S)} = L(g) \leqslant \operatorname{scl}(g).$$

This would mean the bound in Theorem 4.3 is sharp and the surface S is extremal.

It may not be the case that all inner boundary components of S_0 can be made to have degree zero. Nevertheless, when g has alternating t–shape, we can control the number of inner boundary components of S_0 that have nonzero degree in such a way as to show that there are surfaces S for which $\frac{-\chi(S)}{2n(S)}$ is arbitrarily close to $L(g)$. The details are given in the proof of Theorem 5.2. In this way, we establish that the lower bound given in Theorem 4.3 is sharp for elements of alternating t–shape.

Theorem 5.2. *Let $g \in BS(m, \ell)$, $m \neq \ell$, have alternating t–shape. Then*

$$\mathrm{scl}(g) = L(g).$$

Proof. We will show that $\mathrm{scl}(g) < L(g) + \epsilon$ for all $\epsilon > 0$. Note that, since g is of alternating t–shape, all circuits in the turn graph are either of type m or of type ℓ, not of mixed type. Let $\mathbf{u}$ be a vertex of P on which $|\cdot|_{\mathbb{X}}$ is maximal. Since P is a rational polyhedron on which all coordinates are nonnegative, all coordinates of $\mathbf{u}$ are nonnegative rational numbers. By Remark 4.4, $\mathbf{u}$ has nonzero entries only in coordinates corresponding to $\mathbb{X}$. Let K denote the number of edges in the turn graph. Let N be an integer such that each coordinate of $N\mathbf{u}$ is a nonnegative integer and such that $N > K/2\epsilon$ (so that $K/2N < \epsilon$).

Each coordinate $\mathbf{x}_i$ of $\mathbb{X}$ represents a directed circuit γ in the turn graph. For each such directed circuit γ of length n, we consider a $2n$–gon with alternate sides labeled by the powers of a corresponding to the vertices of the turn graph through which γ passes and alternate sides labeled by the intervening edges of the turn graph traversed by γ. See Figure 4.

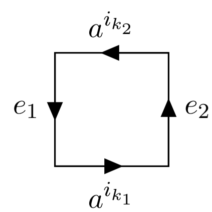

FIGURE 4. A sample polygon corresponding to a circuit γ of length 2.

For each i we take $N\mathbf{u}\cdot\mathbf{x}_i$ copies of the polygon corresponding to $\mathbf{x}_i$, thus obtaining a collection $Q_1, \ldots, Q_s$ of polygons. Since $F_e(N\mathbf{u}) = F_{\bar{e}}(N\mathbf{u})$, there exists a pairing of the edges of these polygons corresponding to edges of the turn graph such that each edge labeled by e on a polygon Q_i is paired with an edge labeled by $\bar{e}$ on a polygon Q_j. Let Δ be the graph dual to this pairing, i.e. the graph with a vertex for each polygon Q_i and an edge between the vertex corresponding to Q_i and the vertex corresponding to Q_j for each edge of Q_i that is paired with an edge from Q_j. The graph Δ may have many components. However, we can adjust the pairings of edges of polygons to obtain some control over the number of components of Δ. Suppose Q_{i_1} and Q_{i_2} are polygons where an edge labeled e of Q_{i_1} has been paired with an edge labeled $\bar{e}$ of Q_{i_2}, and suppose Q_{j_1} and Q_{j_2} are polygons in another component of Δ where an edge labeled e of Q_{j_1} has been paired with an edge labeled $\bar{e}$ of Q_{j_2}. Then we can modify the pairing of edges to instead pair the edge labeled e of Q_{i_1} with the edge labeled $\bar{e}$ of Q_{j_2} and the edge labeled e of Q_{j_1} with the edge labeled $\bar{e}$ of Q_{i_2}. The graph Δ corresponding to this pairing will have one fewer component than the graph corresponding to the original pairing. See Figure 5. Such a modification can be done any time there are two components of Δ on which edges with the same labels have been paired. Therefore, we can arrange that the number of components of Δ is no more than K, the number of edges in the turn graph. Note that Δ is naturally a bipartite graph, with vertices partitioned into those corresponding to turn circuits of type m ("type m vertices") and those corresponding to turn circuits of type ℓ ("type ℓ vertices").

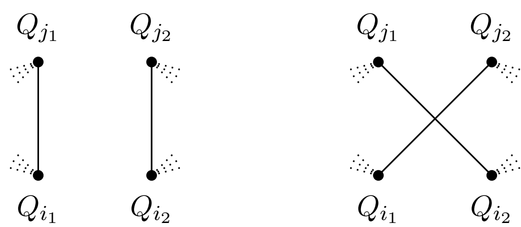

FIGURE 5. Modifying the edges in Δ to reduce the number of components.

If a vertex $v \in \Delta$ corresponds to a turn circuit γ, we define the weight of v to be $\omega(v) := \omega(\gamma)$. If v is a type m vertex we have that $m \mid \omega(v)$, and if v is a type ℓ vertex we have that $\ell \mid \omega(v)$. We wish to assign an integer $\omega(e)$ to each edge $e \in \Delta$ such that, whenever v is of type m, we have

$$\omega(v) - m \sum_{e \ni v} \omega(e) = 0, \tag{9}$$

and, whenever v is of type ℓ, we have

$$\omega(v) + \ell \sum_{e \ni v} \omega(e) = 0. \tag{10}$$

On each connected component of Δ, proceed as follows. For each type m vertex v, choose a preferred edge e_v emanating from v. For each v, set $\omega(e_v) = \omega(v)/m$, and let $\omega(e) = 0$ for all other edges e. This makes (9) hold for all vertices of type m. Now choose a preferred type ℓ vertex v_0. For another type ℓ vertex v_1, let

$$\Omega = \frac{\omega(v_1)}{\ell} + \sum_{e \ni v_1} \omega(e).$$

Choose a path $e_1, \ldots, e_k$ connecting v_1 to v_0, and modify the weights $\omega(e_i)$ by decreasing $\omega(e_i)$ by Ω whenever i is odd and increasing $\omega(e_i)$ by Ω whenever i is even. For all vertices other than v_1 and v_0, this does not change the quantities in (9) and (10). Moreover, this causes (10) to now be true for v_1. Fixing v_0 and letting v_1 vary over all type ℓ vertices other than v_0, we obtain edge weights $\omega(e)$ such that (9) and (10) are true for all vertices on this component of Δ except for v_0. Thus we obtain edge weights $\omega(e)$ such that (9) and (10) are true for all vertices except for one vertex in each component of Δ.

We now proceed to build a surface. Rather than building a surface from the polygons Q_i, we use them to build a band surface S_0, then attempt to fill various components of ∂S_0 with disks. For each pairing of an edge of Q_i with an edge of Q_j, insert a rectangle with sides labeled by t, $a^{m\omega(e)}$, t^{-1}, and $a^{-\ell\omega(e)}$, where $\omega(e)$ is the weight assigned to the corresponding edge of Δ. See Figure 6.

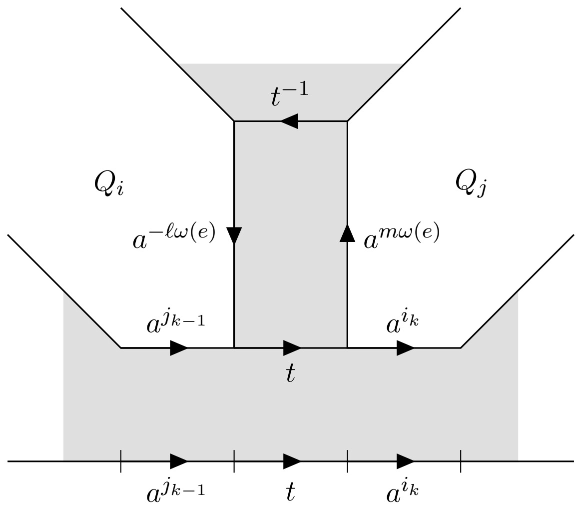

FIGURE 6. Using the polygons Q_i to build a band surface S_0. The surface S_0 is shaded above.

Note that the edges of these rectangles labeled t and t^{-1}, together with the edges of the polygons labeled by powers of a, form paths that would map to the 1–skeleton of X. By construction, these paths correspond exactly to powers of w. To each of these paths, attach an annulus $S^1 \times [0, 1]$ labeled on both sides by this power of w. The rectangles and annuli together form S_0, shown in Figure 6. Note that each rectangle maps naturally to the 2–cell of X with degree $\omega(e)$. The annuli map to the 1–skeleton of X, as indicated by the labels, with the map factoring through the projection $S^1 \times [0, 1] \to S^1$.

As in Section 3, we refer to the boundary components of S_0 that map to a power of w as outer boundary components and to those corresponding to a polygon Q_i as inner boundary components. Each of the inner boundary components of S_0 maps to a power of a; let d be the number of components of the inner boundary for which this power is zero. The powers of a on the inner boundary components are exactly the quantities on the left-hand sides of (9) and (10). The weights $\omega(e)$ have been chosen so that these quantities are zero for all but K components of the inner boundary, so therefore $d \geqslant s - K$. Fill these d components of the inner boundary with disks, and call the resulting surface S.

Each inner boundary component of S maps to a power of a, and Lemma 2.3 says that $\mathrm{scl}(a) = 0$. Therefore, applying Lemma 2.2, we have that

$$\begin{aligned}
\mathrm{scl}(g) &\leqslant \frac{-\chi(S)}{2n(S)} \\
&= \frac{|g|_t}{4} - \frac{d}{2N} \\
&\leqslant \frac{|g|_t}{4} - \frac{s-K}{2N} \\
&= \frac{|g|_t}{4} - \frac{s}{2N} + \frac{K}{2N} \\
&< \frac{|g|_t}{4} - \frac{1}{2}\max\{|\mathbf{u}|_{\mathbb{X}} \mid \mathbf{u} \text{ is a vertex of } P\} + \epsilon \\
&= L(g) + \epsilon.
\end{aligned}$$

Thus $\mathrm{scl}(g) = L(g)$, as desired. □

Corollary 5.3. *If $g \in BS(m, \ell)$, $m \neq \ell$, has alternating t–shape, then* $\mathrm{scl}(g)$ *is rational.*

Proof. Since each vertex $\mathbf{u}$ of P has rational coordinates and $|\mathbf{u}|_{\mathbb{X}}$ is the sum of certain of these coordinates, we know that $L(g)$ is rational. Therefore it follows from Theorem 5.2 that $\mathrm{scl}(g)$ is rational. □

Remark 5.4. In general one suspects the inequality in Theorem 4.3 is strict. For example, the function $|\mathbf{u}|_{\mathbb{X}}$ has a unique maximum on the polyhedron P in Theorem 4.3 for the element $a^2t^2at^{-1}at^{-1} \in BS(2,3)$. When attempting to build a surface from this unique optimal vertex $\mathbf{u}$, it turns out that every component of the dual graph Δ has a constant proportion of vertices that cannot be filled, regardless of the edge weights.

Therefore, in contrast to Theorem 5.2, where all but at most a constant number of vertices can be filled, there is no sequence of surfaces associated to $\mathbf{u}$ to for which $\frac{\chi^+(S)}{n(S)}$ approaches $|\mathbf{u}|_{\mathbb{X}}$.

An explicit formula. When $|g|_t = 2$ the turn graph consists of two vertices, each adjacent to a one-edge loop. In this case the vector space $\mathbb{X}$ is essentially two dimensional and the linear optimization problem can be solved easily by hand, resulting in a formula for stable commutator length for such elements.

This calculation is interesting for two reasons. First, it is rare that one can derive a formula for scl in non-trivial cases. Second, the minimal value for scl among all "well-aligned" elements (see Definition 6.8 and Theorem 6.9) is realized by an element of this type.

Proposition 5.5. *In the group $BS(m, \ell)$ with $m \neq \ell$, if $m \nmid i$ and $\ell \nmid j$ then*

$$\mathrm{scl}(ta^it^{-1}a^j) \;=\; \frac{1}{2}\left(1 - \frac{\gcd(i,m)}{|m|} - \frac{\gcd(j,\ell)}{|\ell|}\right).$$

The divisibility hypotheses simply mean that the word $ta^it^{-1}a^j$ is cyclically reduced (cf. Remark 2.6).

Proof. The turn graph for the word $ta^it^{-1}a^j$ is as shown in Figure 7.

FIGURE 7. The turn graph for $ta^it^{-1}a^j$.

There are two types of potential disks:

(i) Circuits of type m that traverse the left loop of the turn graph p times, where $m \mid pi$.
(ii) Circuits of type ℓ that traverse the right loop of the turn graph q times, where $\ell \mid qj$.

Note that the condition $m \mid pi$ is equivalent to $\frac{|m|}{\gcd(i,m)} \mid p$, and the condition $\ell \mid qj$ is equivalent to $\frac{|\ell|}{\gcd(j,\ell)} \mid q$. Suppose $p = \frac{k|m|}{\gcd(i,m)}$, where $p \leqslant \max\{|m|, |\ell|\}$, for some positive integer k, and let $\mathbf{x}_{i_k}$ be the corresponding basis vector of $\mathbb{X}$. We claim that, if $k > 1$ and $\mathbf{u}$ is a vertex of P that maximizes $|\mathbf{u}|_{\mathbb{X}}$, then $\mathbf{u} \cdot \mathbf{x}_{i_k} = 0$. Indeed, suppose not, and consider the vector $\mathbf{u}' = \mathbf{u} - (\mathbf{u} \cdot \mathbf{x}_{i_k})\mathbf{x}_{i_k} + k(\mathbf{u} \cdot \mathbf{x}_{i_k})\mathbf{x}_{i_1}$. Then $F_e(\mathbf{u}') = F_e(\mathbf{u})$ for all e and $F(\mathbf{u}') = F(\mathbf{u}) = 1$, but

$$\left|\mathbf{u}'\right|_{\mathbb{X}} = |\mathbf{u}|_{\mathbb{X}} - \mathbf{u} \cdot \mathbf{x}_{i_k} + k(\mathbf{u} \cdot \mathbf{x}_{i_k}) = |\mathbf{u}|_{\mathbb{X}} + (k-1)(\mathbf{u} \cdot \mathbf{x}_{i_k}) > |\mathbf{u}|_{\mathbb{X}}.$$

A similar argument applies to coordinates of $\mathbb{X}$ corresponding to potential disks of type ℓ. Thus, if $\mathbf{u}$ is a vertex of P that maximizes $|\mathbf{u}|_{\mathbb{X}}$, only two coordinates of $\mathbf{u}$ are nonzero, one corresponding to a potential disk of type m where $p = \frac{|m|}{\gcd(i,m)}$ and the other corresponding to a potential disk of type ℓ where $q = \frac{|\ell|}{\gcd(j,\ell)}$.

Let c be the value of the coordinate corresponding to this potential disk of type m, and let d be the value of the coordinate corresponding to this potential disk of type ℓ. Then the conditions $F_e(\mathbf{u}) = F_{\bar{e}}(\mathbf{u})$ and $F(\mathbf{u}) = 1$ become

$$\frac{|m|}{\gcd(i,m)} c = \frac{|\ell|}{\gcd(j,\ell)} d = 1.$$

Therefore we have that $c = \frac{\gcd(i,m)}{|m|}$ and $d = \frac{\gcd(j,\ell)}{|\ell|}$. This means that

$$\begin{aligned} \operatorname{scl}(ta^it^{-1}a^j) &= \frac{|ta^it^{-1}a^j|_t}{4} - \frac{1}{2}\max\{|\mathbf{u}|_{\mathbb{X}} \mid \mathbf{u} \text{ is a vertex of } P\} \\ &= \frac{1}{2} - \frac{1}{2}\left(\frac{\gcd(i,m)}{|m|} + \frac{\gcd(j,\ell)}{|\ell|}\right), \end{aligned}$$

as desired. □

Extremal surfaces. We now characterize the elements $g \in BS(m,\ell)$ of alternating t–shape for which an extremal surface exists.

Lemma 5.6. *Suppose S is an admissible surface for some $g \in BS(m, \ell)$, $m \neq \ell$, of alternating t–shape that has been decomposed as described in Section 3. If S is extremal, then S_1 consists only of disks.*

Proof. If S is extremal, we know by Lemma 3.3 that S_1 consists of only disks and annuli. Suppose that some component of S_1 is an annulus. This means that some component of the inner boundary of S_0 does not bound a disk in S. Using the construction from the proof of Proposition 4.2, there is a $\mathbf{u} \in P$ such that $|\mathbf{u}|_{\mathbb{X}} \geqslant \chi^+(S_1)$ and $\mathbf{u} \cdot \mathbf{y}_i > 0$ for some i. Remark 4.4 shows how to find $\mathbf{u'} \in P$ such that $|\mathbf{u'}|_{\mathbb{X}} > |\mathbf{u}|_{\mathbb{X}}$, so we have $|\mathbf{u'}|_{\mathbb{X}} > \chi^+(S_1)$. But then Theorem 5.2 shows that S is not extremal. Thus S_1 cannot have an annular component, meaning it consists only of disks. □

Theorem 5.7. *Let $g = \prod_{k=1}^{r} ta^{i_k}t^{-1}a^{j_k} \in BS(m, \ell)$, $m \neq \ell$. There is an extremal surface for g if and only if*

$$\ell \sum_{k=1}^{r} i_k = -m \sum_{k=1}^{r} j_k. \tag{11}$$

Proof. The status of equation (11) does not change under cancellation of $t^\epsilon t^{-\epsilon}$ pairs in g, nor under applications of the defining relator in $BS(m, \ell)$; hence we may assume without loss of generality that g is cyclically reduced.

First, suppose g has an extremal surface S. Decompose S as described in Section 3. By Lemma 5.6, S_1 consists only of disks. Let Δ be the graph that has a vertex for each component of S_1 and an edge for each band of S_b that connects the vertices corresponding to the two disks it adjoins. There is a weight function on the vertices of Δ, where $\omega(v)$ is the total degree of a at all vertices of the circuit in the turn graph corresponding to v. There is also a natural weight function w on the edges of Δ, where $\omega(e)$ is the signed degree of the map from the band corresponding to e to the 2–cell of X. Since all vertices of Δ bound disks, we know that whenever v corresponds to a circuit of type m, we have

$$\omega(v) - m \sum_{e \ni v} \omega(e) = 0,$$

and, whenever v corresponds to a circuit of type ℓ, we have

$$\omega(v) + \ell \sum_{e \ni v} \omega(e) = 0.$$

Summing over all vertices of type m, we obtain

$$m \sum_{e \in \Delta} \omega(e) = \sum_{\substack{v \in \Delta \\ \text{of type } m}} \omega(v) = n(S) \sum_{k=1}^{r} i_k. \tag{12}$$

Summing over all vertices of type ℓ, we obtain

$$-\ell \sum_{e \in \Delta} \omega(e) = \sum_{\substack{v \in \Delta \\ \text{of type } \ell}} \omega(v) = n(S) \sum_{k=1}^{r} j_k. \tag{13}$$

Multiplying (12) by ℓ and (13) by $-m$ and combining gives (11).

Conversely, suppose the element g satisfies (11). Let S_0 and Δ be as in the proof of Theorem 5.2. Restrict to one connected component of S_0, and let Δ_0 be the corresponding connected component of Δ. Then (9) holds for all vertices of type m in Δ_0. Let N_0 be the power of w corresponding to the image of the map on the outer boundary of this component of S_0. Summing over all vertices of type m in Δ_0, we have that

$$m \sum_{e \in \Delta_0} \omega(e) = \sum_{\substack{v \in \Delta_0 \\ \text{of type } m}} \omega(v) = N_0 \sum_{k=1}^{r} i_k \tag{14}$$

also holds. Multiplying (14) by ℓ and combining with (11) shows that

$$-\ell \sum_{e \in \Delta_0} \omega(e) = \sum_{\substack{v \in \Delta_0 \\ \text{of type } \ell}} \omega(v) = N_0 \sum_{k=1}^{r} j_k. \tag{15}$$

Since the procedure in the proof of Theorem 5.2 ensures that (10) holds for all but one $v \in \Delta_0$ of type ℓ, (15) implies that (10) in fact holds for all $v \in \Delta_0$. The same argument applies to each component of Δ, so hence (9) and (10) hold for all $v \in \Delta$. Thus all inner boundary components of S_0 can be filled with disks, meaning the resulting surface achieves the lower bound on $\mathrm{scl}(g)$ given by linear programming. Hence this surface is extremal. □

Remark 5.8. Corollary 5.3 and Theorem 5.7 combine to give many examples of elements for which stable commutator length is rational but for which no extremal surface exists. Previous examples of this phenomenon were found in free products of abelian groups of higher rank. See [8].

6. Quasimorphisms on groups acting on trees

We now turn our attention from analyzing scl for a single element in $BS(m, \ell)$ to analyzing properties of the scl spectrum $\mathrm{scl}(BS(m, \ell)) \subset \mathbb{R}$. Our main theorem about the spectrum (Theorem 7.8) shows that either $\mathrm{scl}(g) = 0$ or $\mathrm{scl}(g) \geqslant 1/12$. In other words, there is a *gap* in the spectrum. The proof has two parts. In this section we will provide a general condition (*well-aligned*) for an element g in a group G acting on a tree that implies $\mathrm{scl}(g) \geqslant 1/12$ (Theorem 6.9). This is not quite enough for the Gap Theorem for Baumslag–Solitar groups. In Section 7 we use the specific structure of $BS(m, \ell)$ as an HNN extension and show that if a stronger form of the well-aligned property does not hold, then $\mathrm{scl}(g) = 0$ (Theorem 7.6 and Proposition 7.7).

The key to the argument in Theorem 6.9 is the construction of a certain function $f\colon BS(m, \ell) \to \mathbb{R}$ for each hyperbolic element $g \in BS(m, \ell)$, satisfying certain properties described below, that provides a lower bound on $\mathrm{scl}(g)$.

The material in this section applies to any group G.

Quasimorphisms and stable commutator length. The functions we will construct are homogeneous quasimorphisms.

Definition 6.1. A function $f\colon G \to \mathbb{R}$ is called a *quasimorphism* if there is a number D such that

$$|f(gh) - f(g) - f(h)| \ \leqslant\ D \tag{16}$$

for all $g, h \in G$. The smallest such D is called the *defect* of f. A quasimorphism f is *homogeneous* if $f(g^n) = nf(g)$ for all $g \in G$, $n \in \mathbb{Z}$.

Bavard Duality [1] provides the link between homogeneous quasimorphisms and stable commutator length. We only need one direction of this link.

Proposition 6.2 (Bavard Duality, easy direction)**.** *Given $g \in G$, suppose there is a homogeneous quasimorphism f with defect at most D such that $f(g) = 1$. Then* $\operatorname{scl}(g) \geqslant 1/2D$*.*

Proof. One checks easily using (16) that if g^n is a product of m commutators then $|f(g^n)| \leqslant 2mD$. Hence $1 = |f(g)| \leqslant 2\operatorname{cl}(g^n)D/n$, and taking the limit as $n \to \infty$ gives the desired result. □

Hence to derive a large lower bound for $\operatorname{scl}(g)$ one tries to construct a homogenous quasimorphism $f\colon G \to \mathbb{R}$ such that $f(g) = 1$ and with defect as small as possible.

***G*–trees.** We consider simplicial trees not just as combinatorial objects, but also as metric spaces with each edge being isometric to an interval of length one. A *segment* in a tree T is a subset $\alpha \subset T$ that is isometric to a closed segment in $\mathbb{R}$.

Suppose we are given an action of G on a simplicial tree T, always assumed to be without inversions. Every element $g \in G$ has a *characteristic subtree* T_g, consisting of those points $x \in T$ where the displacement function $x \mapsto d(x, gx)$ achieves its minimum. This minimum is denoted $|g|$, and called the *translation length* of g, or simply the *length* of g. If the length is zero then T_g is the set of fixed points of g and we call g *elliptic*. Otherwise, T_g is a linear subtree on which g acts by a shift of amplitude $|g|$. In this case T_g is called the *axis* of g and g is *hyperbolic*. Note that T_g has a natural orientation, given by the direction of the shift by g. A *fundamental domain* for g is a segment (of length $|g|$) contained in the axis, of the form $[x, gx]$. We specifically allow x to be a point in the interior of an edge.

If $k \neq 0$ then g^k has the same type (elliptic or hyperbolic) as g. If g is hyperbolic then $T_{g^k} = T_g$ and $\left|g^k\right| = |k|\,|g|$. Also, $\left|hgh^{-1}\right| = |g|$ for all g, h.

Remark 6.3. There is an easy way to identify the axis of a hyperbolic element $g \in G$. Namely, if α is an oriented segment or edge in T, then α is on the axis if and only if α and $g\alpha$ are *coherently oriented* in T, i.e., there is an oriented segment that contains both α and $g\alpha$ as oriented subsegments. When this occurs, if x is any point in α, then the segment $[x, gx]$ is a fundamental domain for g.

Definition 6.4. Let γ be an oriented segment in T. The *reverse* of γ is the same segment with the opposite orientation, denoted $\overline{\gamma}$. A *copy of* γ is a segment of the form $g\gamma$ for some $g \in G$.

If g is hyperbolic then the quotient of T_g by the action of $\langle g \rangle$ is a circuit of length $|g|$. A *copy of* γ in $T_g/\langle g \rangle$ is the image of a copy of γ in T_g, provided that $|\gamma| \leqslant |g|$. (If $|\gamma| > |g|$ then there are no copies of γ in $T_g/\langle g \rangle$.) We say that two segments *overlap* if their intersection is a non-trivial segment.

Let γ be an oriented segment in T. For an oriented segment α, let $c_\gamma(\alpha)$ be the maximal number of non-overlapping positively oriented copies of γ in α. Note that $c_\gamma(\overline{\alpha}) = c_{\overline{\gamma}}(\alpha)$. Also define

$$f_\gamma(\alpha) \ = \ c_\gamma(\alpha) - c_{\overline{\gamma}}(\alpha).$$

If $g \in G$ is hyperbolic, let $c_\gamma(g)$ be the maximal number of non-overlapping positively oriented copies of γ in $T_g/\langle g \rangle$. If g is elliptic, let $c_\gamma(g) = 0$. In either case, define

$$f_\gamma(g) \ = \ c_\gamma(g) - c_{\overline{\gamma}}(g) \tag{17}$$

and

$$h_\gamma(g) \ = \ \lim_{n\to\infty} \frac{f_\gamma(g^n)}{n}. \tag{18}$$

We will see shortly that f_γ is a quasimorphism. Therefore, by [6, Lemma 2.21], the limit defining h_γ exists and h_γ is a homogeneous quasimorphism.

Lemma 6.5. *Let $g \in G$ be hyperbolic and suppose that a fundamental domain for g is expressed as a concatenation of non-overlapping segments $\alpha_1, \ldots, \alpha_k$, each given the same orientation as T_g. Then for any γ there is an estimate*

$$\sum_i c_\gamma(\alpha_i) \ \leqslant \ c_\gamma(g) \ \leqslant \ k + \sum_i c_\gamma(\alpha_i).$$

In the situation of the lemma, we will refer to the images in $T_g/\langle g \rangle$ of the endpoints of the segments α_i as *junctures*. There are k junctures in $T_g/\langle g \rangle$.

Proof. Start with maximal collections of non-overlapping copies of γ in the segments α_i. The union of these sets of copies projects to a non-overlapping collection in $T_g/\langle g \rangle$, yielding the first inequality. For the second inequality, start with a maximal collection of non-overlapping copies of γ in $T_g/\langle g \rangle$. At most k of these copies contain junctures in their interiors. Each remaining copy lifts to a copy of γ in one of the segments α_i, and no two of these lifts overlap. Hence $\sum_i c_\gamma(\alpha_i) \geqslant c_\gamma(g) - k$. ☐

The main technical result of this section is the following theorem.

Theorem 6.6. *Suppose G acts on a simplicial tree T. Let γ be an oriented segment in T (with endpoints possibly not at vertices). Then the functions f_γ and h_γ defined in* (17) *and* (18) *are quasimorphisms on G with defect at most* 6.

Proof. We will prove the result for h_γ directly. Replacing "n" throughout by "1" yields a proof of the result for f_γ.

Fix elements $g, h \in G$. We wish to show that $|h_\gamma(gh) - h_\gamma(g) - h_\gamma(h)| \leqslant 6$. There are several cases, corresponding to different configurations of the characteristic subtrees T_g, T_h, and T_{gh}.

Case I: T_g and T_h are disjoint. Let ρ be the segment joining T_h to T_g, oriented from T_h and towards T_g. Let $\rho' = g\overline{\rho}$ (which is a copy of $\overline{\rho}$).

If g and h are both hyperbolic, let α and β be fundamental domains for h and g respectively, as indicated in Figure 8. Note that gh has a fundamental domain

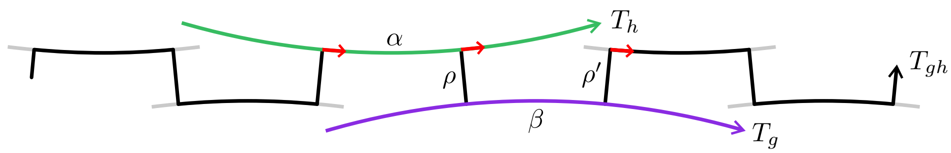

FIGURE 8. Case I, g and h hyperbolic. T_h is green, T_g is purple, T_{gh} is black. The red edges are of the form e, he, and ghe.

given by the concatenation $\alpha \cdot \rho \cdot \beta \cdot \rho'$, by Remark 6.3. More generally, $(gh)^n$ has a fundamental domain made of n copies each of α, ρ, β, and $\overline{\rho}$. Lemma 6.5 yields

$$\begin{aligned} n(c_\gamma(\alpha) + c_\gamma(\rho) + c_\gamma(\beta) + c_\gamma(\overline{\rho})) \ &\leqslant\ c_\gamma((gh)^n) \\ &\leqslant\ n(c_\gamma(\alpha) + c_\gamma(\rho) + c_\gamma(\beta) + c_\gamma(\overline{\rho})) + 4n \end{aligned} \tag{19}$$

and

$$\begin{aligned} n(c_{\overline{\gamma}}(\alpha) + c_{\overline{\gamma}}(\rho) + c_{\overline{\gamma}}(\beta) + c_{\overline{\gamma}}(\overline{\rho})) \ &\leqslant\ c_{\overline{\gamma}}((gh)^n) \\ &\leqslant\ n(c_{\overline{\gamma}}(\alpha) + c_{\overline{\gamma}}(\rho) + c_{\overline{\gamma}}(\beta) + c_{\overline{\gamma}}(\overline{\rho})) + 4n. \end{aligned} \tag{20}$$

Subtracting (20) from (19) yields

$$n(f_\gamma(\alpha) + f_\gamma(\beta)) - 4n \ \leqslant\ f_\gamma((gh)^n) \ \leqslant\ n(f_\gamma(\alpha) + f_\gamma(\beta)) + 4n. \tag{21}$$

Since h^n and g^n have fundamental domains made of n copies of α and β respectively, Lemma 6.5 also yields, in a similar way,

$$nf_\gamma(\alpha) - n \ \leqslant\ f_\gamma(h^n) \ \leqslant\ nf_\gamma(\alpha) + n \tag{22}$$

and

$$nf_\gamma(\beta) - n \ \leqslant\ f_\gamma(g^n) \ \leqslant\ nf_\gamma(\beta) + n. \tag{23}$$

Subtracting (22) and (23) from (21) yields

$$-6n \ \leqslant\ f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n) \ \leqslant\ 6n.$$

Dividing by n and taking a limit, we obtain $|h_\gamma(gh) - h_\gamma(g) - h_\gamma(h)| \ \leqslant\ 6$, as desired.

Remark. In the argument just given, the fundamental domains for h, g, and gh respectively were decomposed into one, one, and four segments; hence $T_h/\langle h\rangle$, $T_g/\langle g\rangle$, and $T_{gh}/\langle gh\rangle$ contained a total of six junctures. These junctures were the only source of defect, since the individual segments always contributed zero to $|f_\gamma(gh) - f_\gamma(g) - f_\gamma(h)|$. Every case below follows the same pattern: the defect will be bounded above by the total number of junctures appearing in the quotient circuits. In what follows, we will describe the structure of $T_h/\langle h\rangle$, $T_g/\langle g\rangle$, and $T_{gh}/\langle gh\rangle$ in each case and leave some of the details of the estimates to the reader.

Returning to Case I, suppose g and h are both elliptic. Then g^n and h^n are also elliptic, and $(gh)^n$ has a fundamental domain made of n copies of ρ and n copies of $\overline{\rho}$. See Figure 9. Using Lemma 6.5 one obtains

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \; = \; |f_\gamma((gh)^n)| \; \leqslant \; 2n,$$

for a defect of at most 2.

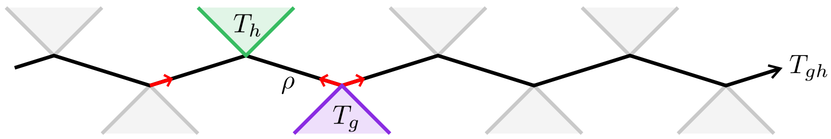

FIGURE 9. Case I, g and h elliptic. T_h is green, T_g is purple, T_{gh} is black. The red edges are of the form e, he, and ghe.

If one of g and h is elliptic, say g, then h^n has a fundamental domain given by n copies of α, and $(gh)^n$ has a fundamental domain given by n copies each of α, ρ, and $\overline{\rho}$. The estimate given by Lemma 6.5 becomes

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \; = \; |f_\gamma((gh)^n) - f_\gamma(h^n)| \; \leqslant \; 4n,$$

for a defect of at most 4.

Case II: g and h are hyperbolic, with positive overlap. That is, T_g and T_h intersect in a segment, on which T_g and T_h induce the same orientation. Let e be an edge in $T_g \cap T_h$, oriented coherently with T_g and T_h. Let v be the terminal vertex of e. Since $e \in T_h$, the edges $h^{-1}e$ and e are coherently oriented and $\alpha = [h^{-1}v, v]$ is a fundamental domain for h. Similarly, e and ge are coherently oriented and $\beta = [v, gv]$ is a fundamental domain for g.

Let all the edges of α and β be given orientations from T_h and T_g respectively. Since g and h both move e in the same direction (that is, into the same component of $T - \{e\}$), e separates $h^{-1}e$ from ge. It follows that the edges of α and of β are all coherently oriented in T. Hence α and β do not overlap, and $\alpha \cdot \beta = [h^{-1}v, gv]$ is a fundamental domain for gh.

With a total of four junctures (one for h, one for g, two for gh), the estimate given by Lemma 6.5 becomes

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \; \leqslant \; 4n,$$

for a defect of at most 4.

Case III: g and h are hyperbolic, with negative overlap. That is, T_g and T_h intersect in a segment, on which T_g and T_h induce opposite orientations. Let Δ be the length (possibly infinite) of $T_g \cap T_h$. There are several sub-cases, according to the relative sizes of $|g|$, $|h|$, and Δ.

Sub-case III-A: $\Delta \leqslant |g|, |h|$, not all three numbers equal. Let ρ be the segment $T_h \cap T_g$, oriented coherently with T_h. There is a fundamental domain for h of the form $\alpha \cdot \rho$, and similarly, a fundamental domain for g of the form $\overline{\rho} \cdot \beta$; see Figure

10. Then $\alpha \cdot \beta$ is a fundamental domain for gh. (By assumption, at least one of α, β is a non-trivial segment, and gh is hyperbolic.)

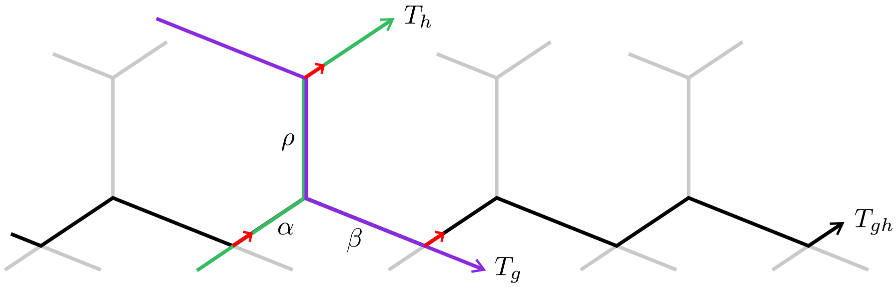

FIGURE 10. Case III-A. T_h is green, T_g is purple, T_{gh} is black. The red edges are of the form e, he, and ghe.

The quotient circuits have at most six junctures: two for g, two for h, and two for gh. Lemma 6.5 leads to an estimate

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \;\leqslant\; 6n,$$

for a defect of at most 6.

Sub-case III-B: $|g| = |h| \leqslant \Delta$. In this case, we show that gh is elliptic. Let $\alpha \subset T_h \cap T_g$ be a fundamental domain for h. Then $\overline{\alpha}$ is a fundamental domain for g, and gh fixes the initial endpoint of α. With two junctures in total, we obtain the estimate

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \;=\; |-f_\gamma(g^n) - f_\gamma(h^n)| \;\leqslant\; 2n,$$

for a defect of at most 2.

Sub-case III-C: $|h| < |g|, \Delta$. There is a simplicial fundamental domain α for h such that if e is the initial edge of α, then $\alpha \cdot he$ is contained in $T_h \cap T_g$. Then, there is a fundamental domain for g of the form $\overline{\alpha} \cdot \beta$; see Figure 11. By considering the location of ghe, one finds that β is a fundamental domain for gh. The three circuits have a total of four junctures, and we obtain

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \;\leqslant\; 4n,$$

for a defect of at most 4.

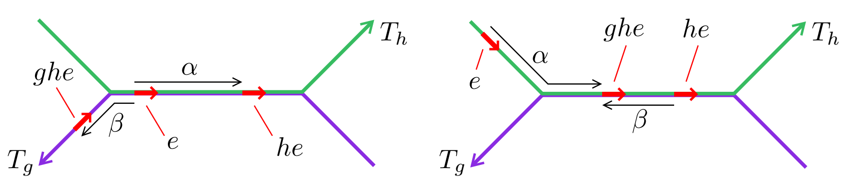

FIGURE 11. Cases III-C (left) and III-D (right). T_h is green, T_g is purple.

Sub-case III-D: $|g| < |h|, \Delta$. Fundamental domains β, α, and $\alpha \cdot \overline{\beta}$ for g, gh, and h respectively can be constructed in a similar fashion as in Case III-C; see Figure 11.

Alternatively, this case reduces to Case III-C, replacing g and h by h^{-1} and g^{-1} respectively.

Case IV: g is elliptic, h is hyperbolic, $T_g \cap T_h \neq \emptyset$. If $T_g \cap T_h$ contains an edge e, let $\alpha \subset T_h$ be the fundamental domain starting with $h^{-1}e$. Then α is also a fundamental domain for gh. This leads to an estimate

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \;=\; |f_\gamma((gh)^n) - f_\gamma(h^n)| \;\leqslant\; 2n,$$

and a defect of at most 2.

If $T_g \cap T_h$ is a single vertex v, let $e \in T_h$ be the coherently oriented edge with initial vertex v, and let α be the fundamental domain $[h^{-1}v, v]$. If $ge \notin \alpha$ then α is a fundamental domain for gh also, and we obtain a defect of at most 2 as above.

So now assume that $ge \in \alpha$, i.e. that ge separates $h^{-1}v$ from v. Note that $h^{-1}e$ and ge are not coherently oriented, so the characteristic subtree T_{gh} will not contain these edges.

We have that $gh(\alpha) \cap \alpha$ contains the edge ge. Consider the length of $gh(\alpha) \cap \alpha$. If this length is $|\alpha|/2$ or greater, then gh fixes the midpoint of α. Then

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \;=\; |-f_\gamma(h^n)| \;\leqslant\; n,$$

giving a defect of at most 1.

Otherwise, there is a subsegment $\beta \subset \alpha$, centered on the midpoint of α, of maximal size so that β does not overlap $gh\beta$. We can write α as a concatenation $\alpha_1 \cdot \beta \cdot \alpha_2$, where $\alpha_2 = gh\overline{\alpha}_1$. See Figure 12. Now β is a fundamental domain for gh,

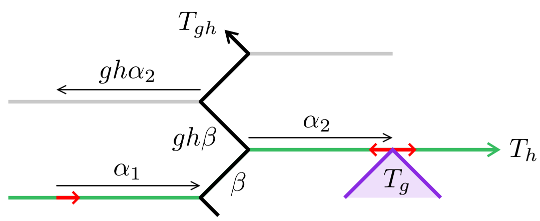

FIGURE 12. Case IV. The element gh takes α_1 to $\overline{\alpha}_2$. Left to right, the red edges are $h^{-1}e$, ge, and e.

and we have a total of four junctures (three in $T_h/\langle h\rangle$ and one in $T_{gh}/\langle gh\rangle$). Thus we have

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \;=\; |f_\gamma((gh)^n) - f_\gamma(h^n)| \;\leqslant\; 4n,$$

and a defect of at most 4.

Case V: h is elliptic, g is hyperbolic, $T_g \cap T_h \neq \emptyset$. This case is covered by Case IV, replacing g and h by h^{-1} and g^{-1} respectively.

Case VI: the remaining cases. If g and h are hyperbolic and T_g and T_h intersect in one point, then the configuration closely resembles the first one discussed in Case I, except that the copies of ρ have been shrunk to have length zero. That is, there are fundamental domains α and β for h and g respectively, such that $\alpha \cdot \beta$ is a fundamental domain for gh. With four junctures, we obtain

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| \leqslant 4n,$$

for a defect of at most 4.

Lastly, if g and h have a common fixed point, then

$$|f_\gamma((gh)^n) - f_\gamma(g^n) - f_\gamma(h^n)| = 0$$

for all n. □

Remark 6.7. The functions f_γ and h_γ can be defined in the more general setting of a group acting on an $\mathbb{R}$–tree. The proof of Theorem 6.6 goes through in this setting, with only superficial modifications (essentially, removing any mention of *edges*, and using small segments instead).

Well-aligned elements. We now consider elements $g \in G$ for which we can find a segment $\gamma \subset T$ such that $h_\gamma(g) = 1$.

Definition 6.8. Given a G–tree T, a hyperbolic element $g \in G$ is *well-aligned* if there does not exist an element $h \in G$ such that $ghgh^{-1}$ fixes an edge of T_g. This property is the G–tree analogue of the double coset condition from [9, Theorem D].

Theorem 6.9. *Suppose G acts on a simplicial tree T. If $g \in G$ is well-aligned then* $\mathrm{scl}(g) \geqslant 1/12$.

Proof. Let $\gamma = [x, gx] \subset T_g$ be a fundamental domain for g where x is *not* a vertex of T. We know that $c_\gamma(g^n) = n$ for all n. If $c_{\overline{\gamma}}(g^n) > 0$ for some n, then there is a copy of $\overline{\gamma}$ in T_g. That is, there is an element h such that $h\gamma$ lies in T_g with the opposite orientation. So $h(T_g) = T_{hgh^{-1}}$ has negative overlap with T_g along a segment containing $h\gamma$. The element $ghgh^{-1}$ fixes one of the endpoints of $h\gamma$, since g and hgh^{-1} shift it in opposite directions inside $T_g \cap T_{hgh^{-1}}$. This endpoint is in the interior of an edge $e \subset T_g$, and so $ghgh^{-1}$ fixes e. Hence, if g is well-aligned, we must have $h_\gamma(g) = 1$. By Theorem 6.6, h_γ is a homogeneous quasimorphism with defect at most 6, and so Proposition 6.2 implies that $\mathrm{scl}(g) \geqslant 1/12$. □

This bound is in fact *optimal*. Both in HNN extensions and in amalgamated free products, there are examples of elements g with $\mathrm{scl}(g) = 1/12$ that are well-aligned with respect to the action on the associated Bass–Serre tree, as we now explain. This answers Question 8.4 from [9].

Theorem 6.10. *Let $g = tat^{-1}a \in BS(2,3)$ and let T be the Bass–Serre tree associated to the splitting of $BS(2,3)$ as an HNN extension $\langle a \rangle *_{\langle ta^2t^{-1} = a^3 \rangle}$. Then g is well-aligned and* $\mathrm{scl}(g) = 1/12$. *In particular, the bound in Theorem 6.9 is optimal.*

Proof. Denote the vertex of T stabilized by $\langle a \rangle$ by v_0 and let $v_1 = tv_0$. The vertices along the axis of g are: $\{g^n v_0, g^n v_1\}_{n \in \mathbb{Z}}$.

If $ghgh^{-1}$ fixes an edge $e \subset T_g$, then we also see that $hgh^{-1}g$ fixes $g^{-1}e \subset T_g$. Replacing h by hg^k for some k (which does not affect hgh^{-1}), we can arrange that h fixes a vertex of T_g. By further replacing h by a conjugate $g^k hg^{-k}$, we can arrange that the vertex fixed by h is either v_0 or v_1; the elements $ghgh^{-1}$ and $hgh^{-1}g$ still fix edges of T_g.

First assume that h fixes v_0, and so $h = a^r$ for some $r \in \mathbb{Z}$. In this case

$$ghgh^{-1} = tat^{-1}a^{1+r}tat^{-1}a^{1-r}.$$

If $ghgh^{-1}$ is elliptic (which it necessarily is if it fixes an edge), then this expression cannot be cyclically reduced (Remark 2.6). Hence we find that $r \equiv \pm 1 \mod 3$. If $r \equiv 1 \mod 3$, then $hgh^{-1}g = a^5$; if $r \equiv -1 \mod 3$ then $ghgh^{-1} = a^5$. In either case, the element does not fix an edge in T, giving a contradiction.

Similarly, if h fixes v_1, then we have $h = ta^rt^{-1}$ for some $r \in \mathbb{Z}$, and so

$$hgh^{-1}g = ta^{1+r}t^{-1}ata^{1-r}t^{-1}a.$$

Again, this expression cannot be cyclically reduced if $hgh^{-1}g$ is elliptic, and so $r \equiv 1 \mod 2$. Again, we find that $hgh^{-1}g = a^5$, giving a contradiction for the same reason as above. Therefore g is well-aligned as claimed.

Finally, $\mathrm{scl}(tat^{-1}a) = 1/12$ by Proposition 5.5. □

The bound in Theorem 6.9 is still optimal if one restricts to amalgamated free products. In the free product $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/3\mathbb{Z} \cong \mathrm{PSL}(2, \mathbb{Z})$, no nontrivial element fixes an edge of the associated Bass–Serre tree, so every hyperbolic element that is not conjugate to its inverse is well-aligned. The group $\mathrm{PSL}(2, \mathbb{Z})$ has a finite index free subgroup, and therefore stable commutator length can be computed in this group by using a relationship between stable commutator length in a group and a finite index subgroup from [6] together with Calegari's algorithm for computing stable commutator length in free groups [7]. This is described explicitly in [16]. The element $\left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$ is an example of an element that has stable commutator length 1/12.

Acylindrical trees. We conclude this section by adding a moderate restriction, *acylindricity*, to the tree action. We can then say something about hyperbolic elements that are not necessarily well-aligned.

Acylindricity has been used previously in the context of counting quasimorphisms on Gromov-hyperbolic spaces, cf. [9]. For a tree, the definition is particularly simple to state. A group acts *K–acylindrically* on a tree if the stabilizer of any segment of length K is trivial.

Theorem 6.11. *Suppose G acts K–acylindrically on a tree T and let N be the smallest integer greater than or equal to $\frac{K}{2} + 1$.*

(i) *If $g \in G$ is hyperbolic then either* $\mathrm{scl}(g) = 0$ *or* $\mathrm{scl}(g) \geqslant 1/12N$.

(ii) *If $g \in G$ is hyperbolic and $|g| \geqslant K$ then either* $\mathrm{scl}(g) = 0$ *or* $\mathrm{scl}(g) \geqslant 1/24$.

In both cases, $\mathrm{scl}(g) = 0$ *if and only if g is conjugate to g^{-1}.*

Proof. First note that if $|g| = 1$ then a fundamental domain for g maps to a single loop in the quotient graph of T, which implies that g has infinite order in the abelianization of G, and $\mathrm{scl}(g) = \infty$. (In fact, the same conclusion holds whenever $|g|$ is odd.) Thus we may assume that $|g| \geqslant 2$.

Observe that if for some $h \in G$, we have $\left|T_g \cap T_{hgh^{-1}}\right| \geqslant K + |g|$ where g and hgh^{-1} shift in opposite directions, then $ghgh^{-1}$ fixes a segment of length K and hence $g = hg^{-1}h^{-1}$. In particular, $\mathrm{scl}(g) = 0$.

For (i) note that $\left|g^N\right| = N\,|g| \geqslant \frac{K}{2}\,|g| + |g| \geqslant K + |g|$. Taking γ to be a fundamental domain for g^N, if $h_\gamma(g^N) < 1$ then there is an h as above and $g = hg^{-1}h^{-1}$, $\mathrm{scl}(g) = 0$. Otherwise, $h_\gamma(g^N) = 1$ and $\mathrm{scl}(g) = \frac{1}{N}\,\mathrm{scl}(g^N) \geqslant 1/12N$ by Theorem 6.6 and Proposition 6.2.

For (ii) let $N = 2$ and apply the same reasoning: $\left|g^2\right| = 2\,|g| \geqslant K + |g|$, and either $\mathrm{scl}(g) = 0$ or $\mathrm{scl}(g) = \frac{1}{2}\,\mathrm{scl}(g^2) \geqslant 1/24$. □

7. The gap theorem

In this section we consider G–trees of a particular form, for which we can improve upon the "well-aligned" condition in Theorem 6.9 without any trade-off in the lower bound of 1/12.

Let G be an HNN extension $A *_C$, with stable letter t, such that the edge groups C and C^t are central in A. Let T be the Bass–Serre tree associated to this HNN extension. Such a tree has special properties, given below in Lemma 7.1 and Proposition 7.2.

This class of HNN extensions obviously includes the Baumslag–Solitar groups. It is still true (cf. Remark 2.4) that an element cannot have finite stable commutator length if its t–exponent is non-zero.

We use the following notation for a G–tree T: if $X \subset T$ is any subset, then G_X denotes the *stabilizer* of X, which is the subgroup of G consisting of those elements that fix X pointwise.

Lemma 7.1. *Suppose $S \subset T$ is a subtree and $h \in G$ fixes a vertex $v \in S$. Then $G_S = G_{hS}$.*

Proof. Let e be an edge in S with endpoint v. Then $G_S \subset G_e$. Since G_e is central in G_v, h commutes with G_S, and therefore $G_{hS} = hG_Sh^{-1} = G_S$. □

Consider the t–exponent homomorphism $\phi\colon G \to \mathbb{Z}$ (sending t to 1 and A to 0). There is an action of $\mathbb{Z}$ on $\mathbb{R}$ by integer translations. Letting G act on $\mathbb{R}$ via ϕ, there is also a G–equivariant map $F\colon T \to \mathbb{R}$. This map is just the natural map from T to the universal cover of the quotient graph T/G. The action of an element g on T projects by F to a translation by $\phi(g)$ on $\mathbb{R}$. We think of F as a *height function* on T. Then, the elements of t–exponent zero act on T by height-preserving automorphisms.

Proposition 7.2. *If $S \subset T$ is a subtree and $\sigma \subset S$ is a finite subtree such that $F(\sigma) = F(S)$ then $G_\sigma = G_S$.*

Proof. Case I: σ is a segment mapped by F injectively to $\mathbb{R}$.

First suppose that S is a finite subtree. Let $\{S_i\}$ be the subtrees obtained as the closures of the components of $S - \sigma$. Then $G_S = G_\sigma \cap \bigcap_i G_{S_i}$. Fixing i, we will prove by induction on the number of edges of S_i that $G_\sigma \subset G_{S_i}$. It then follows that $G_\sigma = G_S$.

The base case is that S_i is a single edge e with one vertex v on σ. Since $F(S_i) \subset F(\sigma)$, there is an edge e' on σ with endpoint v and an element $h \in G_v$ taking e to e'. Indeed, edges incident to v are in the same G_v–orbit if and only if they have the same image under F. By Lemma 7.1 we have that $G_{S_i} = G_{e'} \supset G_\sigma$.

For larger S_i, let $e \in S_i$ be the edge with endpoint $v \in \sigma$. Again there is an edge e' on σ with endpoint v and an element $h \in G_v$ taking e to e'. Again, $G_{S_i} = G_{hS_i}$ by Lemma 7.1. But now $hS_i = e' \cup S_i'$ where S_i' has fewer edges than S_i. By induction, $G_\sigma \subset G_{S_i'}$. Since $G_\sigma \subset G_{e'}$, we now have $G_\sigma \subset (G_{e'} \cap G_{S_i'}) = G_{S_i}$.

Now consider an arbitrary subtree S. We need to show that G_σ fixes S pointwise. But every point x in S is in a finite subtree S' containing σ, and G_σ fixes S' pointwise; hence G_σ fixes x.

Case II: σ is an arbitrary finite subtree of S. Fixing the image $F(\sigma)$, we proceed by induction on the number of edges of σ. The base case is when this number is smallest, namely the length of $F(\sigma)$. Then Case I applies. If there are more edges than this, there must be a vertex $v \in \sigma$ and a pair of edges $e_0, e_1 \in \sigma$ incident to v, with $F(e_0) = F(e_1)$.

Decompose S into two subtrees $S = S_0 \cup S_1$ with $S_0 \cap S_1 = \{v\}$ and $e_0 \in S_0$, $e_1 \in S_1$. Let $\sigma_i = S_i \cap \sigma$. There is an element $h \in G_v$ such that $he_0 = e_1$. Let $\sigma' = h\sigma_0 \cup \sigma_1$ and $S' = hS_0 \cup S_1$. Note that σ' has fewer edges than σ. Also, $F(\sigma) = F(\sigma')$ and $F(S) = F(S')$, and $G_\sigma = G_{\sigma_0} \cap G_{\sigma_1} = G_{h\sigma_0} \cap G_{\sigma_1} = G_{\sigma'}$ by Lemma 7.1. Similarly, $G_S = G_{S_0} \cap G_{S_1} = G_{hS_0} \cap G_{S_1} = G_{S'}$. By the induction hypothesis, $G_{\sigma'} = G_{S'}$, and therefore $G_\sigma = G_S$. □

Now consider a hyperbolic element g with t–exponent zero. The axis T_g has the property that $F(T_g)$ is a finite interval. To see this, let γ be a fundamental domain, and note that $T_g = \bigcup_n g^n\gamma$. The t–exponent condition implies that $F(g^n\gamma) = F(\gamma)$ for all n, and hence $F(T_g) = F(\gamma)$.

Definitions 7.3. We call a vertex v on T_g *extremal* if $F(v)$ is an endpoint of $F(T_g)$. A segment $\sigma \subset T_g$ is *stable* if $F(\sigma) = F(T_g)$ and σ contains no extremal vertex in its interior (equivalently, no proper subsegment σ' satisfies $F(\sigma') = F(T_g)$). See Figure 13. Note that if σ and τ are stable segments, then they do not overlap, unless they are equal.

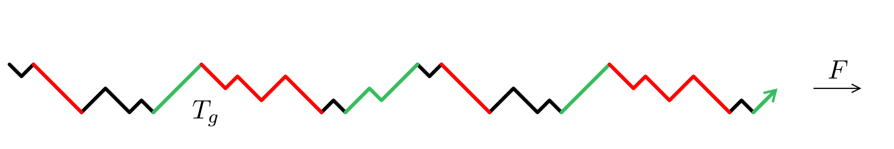

FIGURE 13. Stable segments along T_g, in green and red.

The natural orientation of T_g defines a linear ordering $<_g$ on the stable segments of T_g. The "larger" end is the attracting end of T_g; that is, $\sigma <_g g\sigma$ always holds. We say that $\sigma \leqslant_g \tau$ if $\sigma <_g \tau$ or $\sigma = \tau$.

Remark 7.4. If γ is a fundamental domain for g whose endpoints are extremal, then every stable segment either does not overlap with γ or is contained in γ. Moreover, γ contains a copy of every stable segment. (Being a fundamental domain, it overlaps with a copy of every non-trivial segment in T_g.) If γ is a fundamental domain that starts with a stable segment then its endpoints are extremal, as the endpoints of a stable segment are extremal.

Proposition 7.2 immediately implies:

Corollary 7.5. *If $\sigma \subset T_g$ is stable then $G_\sigma = G_{T_g}$.*

The main technical result of this section is:

Theorem 7.6. *Let $G = A*_C$ with stable letter t, and C, C^t central in A. Let $g \in G$ be a hyperbolic element with t–exponent zero. Then either:*

(i) *there is a fundamental domain γ for g such that $h_\gamma(g) = 1$, or*
(ii) *there is an element h such that $h(T_g) = \overline{T}_g$.*

Conclusion (i) implies by Proposition 6.2 and Theorem 6.6 that $\operatorname{scl}(g) \geqslant 1/12$.

Proof. Let $\alpha \subset T_g$ be a stable segment and let γ be the fundamental domain for g that starts with α. Note that γ has extremal endpoints. If $h_\gamma(g) < 1$ then there is an element h such that $h\gamma$ lies in T_g with the opposite orientation and overlaps with α. Note that h fixes a point in γ. In particular h is elliptic, and hence acts as a height-preserving automorphism of T. Now $h\gamma$ is a fundamental domain for g^{-1} with extremal endpoints, and so $h\gamma$ contains α (which is stable for $T_{g^{-1}}$ as well as for T_g).

The segment $\beta = h^{-1}\alpha$ is a stable segment for T_g contained in γ. Clearly $\alpha \leqslant_g \beta$, and as the endpoints of α have different heights, $\alpha \neq \beta$; therefore $\alpha <_g \beta$. Note that $h\alpha = \beta$, since h acts as a reflection on the segment $\gamma \cap h\gamma$. Hence the element h^2 fixes the stable segment α. Therefore, by Proposition 7.2, h^2 fixes $T_g \cup h(T_g)$. That is, h acts as an involution on this entire subtree of T.

Claim. If there is a stable segment $\rho \subset T_g \cap h(T_g)$ such that either $\rho <_g \alpha$ or $\beta <_g \rho$, then conclusion (ii) holds.

Proof of Claim. Since h acts as a reflection on the segment $T_g \cap h(T_g)$, if $\rho \subset T_g \cap h(T_g)$ and $\rho <_g \alpha$, then $h\rho \subset T_g \cap h(T_g)$ and $\beta <_g h\rho$. Thus we only need to verify the claim in the $\beta <_g \rho$ case.

Let σ be the $<_g$–smallest stable segment in $h\gamma$. Observe that $\sigma \subset T_g \cap h(T_g)$. The translate $g\sigma$ is the $<_g$–smallest stable segment in $gh\gamma$, which has a common endpoint with β. Hence $g\sigma \leqslant_g \tau$ for any τ satisfying $\beta <_g \tau$. In particular, $g\sigma \leqslant_g \rho$. Since $\beta, \rho \subset T_g \cap h(T_g)$ and $\beta <_g g\sigma \leqslant_g \rho$, it follows that $g\sigma \subset T_g \cap h(T_g)$.

Note that $h(T_g) = T_{hgh^{-1}}$, and hgh^{-1} takes $g\sigma$ to σ. Thus $ghgh^{-1}$ fixes $g\sigma$. Since this is a stable segment, $ghgh^{-1}$ must fix all of T_g by Corollary 7.5. This implies that hgh^{-1} acts on T_g as a translation, of the same amplitude but opposite direction as g. Hence $T_{hgh^{-1}} = \overline{T}_g$. □

Returning to the proof of Theorem 7.6, assume that the Claim does not apply. Then α and β are the $<_g$–smallest and $<_g$–largest stable segments in $T_g \cap h(T_g)$ respectively. It follows that β is also the $<_g$–largest stable segment in γ; otherwise, if $\beta <_g \rho$ and $\rho \subset \gamma$, then $h\rho <_g \alpha$ and $h\rho \subset T_g \cap h(T_g)$, contradicting that α is smallest.

Note that h takes stable segments to stable segments, and does not take any stable segment to itself (since the endpoints have different heights). Hence the stable segments of γ may be enumerated in order as $\alpha = \alpha_1, \dots, \alpha_n, \beta_n, \dots, \beta_1 = \beta$ where h interchanges α_i and β_i.

Now let γ' be the fundamental domain for g starting with β_n. Assuming that conclusion (i) does not hold, we have $h_{\gamma'}(g) < 1$, and so there is an elliptic element k such that $k\gamma'$ lies in T_g with the opposite orientation and contains β_n.

The configuration of T_g, $k(T_g)$, β_n, and $k\beta_n$ is exactly analogous to that of T_g, $h(T_g)$, α, and β. In particular, the Claim is applicable to this situation. If the Claim does not apply, then we conclude as above that β_n and $k\beta_n$ are the $<_g$–smallest and $<_g$–largest stable segments in $T_g \cap k(T_g)$ and that $k\beta_n$ is the $<_g$–largest stable segment in γ'.

The stable segments in γ' are, in order: $\beta_n, \dots, \beta_1, g\alpha_1, \dots, g\alpha_n$, and the element k interchanges β_i and $g\alpha_i$. Thus

$$kh\alpha = kh\alpha_1 = k\beta_1 = g\alpha_1 = g\alpha.$$

Since kh and g agree on the stable segment α, they agree on all of $T_g \cup h(T_g) \cup k(T_g)$, by Proposition 7.2. Similarly, h and k both act as involutions on $T_g \cup h(T_g) \cup k(T_g)$. Now

$$hg^{-1}\beta = h(kh)^{-1}\beta = hhk\beta = k\beta = g\alpha,$$

which implies that $g\alpha \subset T_g \cap h(T_g)$. However, $\beta <_g g\alpha$ and β is the $<_g$–largest stable segment in $T_g \cap h(T_g)$. This contradiction establishes the theorem. ☐

The next proposition concerns conclusion (ii) in Theorem 7.6. It is a variant of the observation that if an element is conjugate to its inverse, then it has scl zero.

Proposition 7.7. *Suppose G acts on a tree T and* scl *vanishes on the elliptic elements of G. If g is hyperbolic and there is an element h such that $h(T_g) = \overline{T}_g$, then* $\mathrm{scl}(g) = 0$.

Proof. Since $h(T_g) = T_{hgh^{-1}}$, the element $ghgh^{-1}$ fixes T_g pointwise. Similarly, $g^n hg^n h^{-1}$ fixes T_g for every n. Thus there are elliptic elements a_n such that $g^n hg^n h^{-1} = a_n$. This equation can be realized by a surface of genus zero and three boundary components, labeled by g^n, g^n, and a_n^{-1} respectively. Lemma 2.2 now implies that

$$\mathrm{scl}(g) \;\leqslant\; \frac{1}{4n} + \frac{\mathrm{scl}(a_n^{-1})}{2n}.$$

Hence $\mathrm{scl}(g) \leqslant 1/4n$ for all $n > 0$. ☐

Theorem 7.8 (Gap theorem)**.** *For every element $g \in BS(m, \ell)$, either* $\mathrm{scl}(g) = 0$ *or* $\mathrm{scl}(g) \geqslant 1/12$.

Proof. Every elliptic element g is conjugate to a power of a, and therefore $\mathrm{scl}(g) = 0$ for elliptic elements by Lemma 2.3. If g is hyperbolic then Theorem 7.6 and Proposition 7.7 imply that $\mathrm{scl}(g) = 0$ or $\mathrm{scl}(g) \geqslant 1/12$. □

Remark 7.9. If m and ℓ are both odd, then conclusion (ii) of Theorem 7.6 can never occur, since the element h would fix a vertex of T_g and exchange two adjacent edges, yielding an element of order two in $\mathbb{Z}/m\mathbb{Z}$ or $\mathbb{Z}/\ell\mathbb{Z}$. Therefore, $\mathrm{scl}(g) \geqslant 1/12$ for *every* hyperbolic element g in $BS(m, \ell)$.

If either m or ℓ is even, say $m = 2k$, then for $g = ta^kt^{-1}a^rta^kt^{-1}a^s \in BS(m, \ell)$ where $r + s = \ell$ we have $\mathrm{scl}(g) = 0$. Indeed, taking $h = ta^{-k}t^{-1}$ one checks that $ghgh^{-1} = a^{4\ell} \in \langle a^\ell \rangle = G_{T_g}$. Thus $h(T_g) = \overline{T}_g$ and so by Proposition 7.7 we have $\mathrm{scl}(g) = 0$.

References

1. Christophe Bavard, *Longueur stable des commutateurs*, Enseign. Math. (2) **37** (1991), no. 1-2, 109–150. MR 1115747 (92g:20051)
2. Mladen Bestvina, Ken Bromberg, and Koji Fujiwara, *Stable commutator length on mapping class groups*, preprint, arxiv:1306.2394v1.
3. Mladen Bestvina and Koji Fujiwara, *Bounded cohomology of subgroups of mapping class groups*, Geom. Topol. **6** (2002), 69–89 (electronic). MR 1914565 (2003f:57003)
4. Noel Brady, Matt Clay, and Max Forester, *Turn graphs and extremal surfaces in free groups*, Topology and geometry in dimension three, Contemp. Math., vol. 560, Amer. Math. Soc., Providence, RI, 2011, pp. 171–178. MR 2866930 (2012h:57001)
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, *scl*, MSJ Memoirs, vol. 20, Mathematical Society of Japan, Tokyo, 2009. MR 2527432
7. ______, *Stable commutator length is rational in free groups*, J. Amer. Math. Soc. **22** (2009), no. 4, 941–961. MR 2525776
8. ______, *scl, sails, and surgery*, J. Topol. **4** (2011), no. 2, 305–326. MR 2805993 (2012g:57004)
9. 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)
10. Donald J. Collins, *On embedding groups and the conjugacy problem*, J. London Math. Soc. (2) **1** (1969), 674–682. MR 0252489 (40 #5709)
11. Andrew J. Duncan and James Howie, *The genus problem for one-relator products of locally indicable groups*, Math. Z. **208** (1991), no. 2, 225–237. MR 1128707 (92i:20040)
12. 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)
13. 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)
14. ______, *The second bounded cohomology of an amalgamated free product of groups*, Trans. Amer. Math. Soc. **352** (2000), no. 3, 1113–1129. MR 1491864 (2000j:20098)
15. R. I. Grigorchuk, *Some results on bounded cohomology*, Combinatorial and geometric group theory (Edinburgh, 1993), London Math. Soc. Lecture Note Ser., vol. 204, Cambridge Univ. Press, Cambridge, 1995, pp. 111–163. MR 1320279 (96j:20073)
16. Joel Louwsma, *Extremality of the rotation quasimorphism on the modular group*, Ph.D. thesis, California Institute of Technology, CaltechTHESIS:05132010-155930760, 2011.

17. Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*, Classics in Mathematics, Springer-Verlag, Berlin, 2001, Reprint of the 1977 edition. MR 1812024 (2001i:20064)
18. A. H. Rhemtulla, *A problem of bounded expressibility in free products*, Proc. Cambridge Philos. Soc. **64** (1968), 573–584. MR 0225889 (37 #1480)
19. W. A. Stein et al., *Sage Mathematics Software (Version 5.11)*, The Sage Development Team, 2013, `http://www.sagemath.org`.

Department of Mathematics, University of Arkansas, Fayetteville, AR 72701, USA
E-mail address: `mattclay@uark.edu`

Department of Mathematics, University of Oklahoma, Norman, OK 73019, USA
E-mail address: `forester@math.ou.edu`

Department of Mathematics, Niagara University, Niagara University, NY 14109, USA
E-mail address: `jlouwsma@niagara.edu`