The Structure of Automorphic Conjugacy in the Free Group of Rank Two

Bilal Khan

Abstract. In the study of the automorphism group of a free group $F = F(X)$ on a set X, J. H. C. Whitehead introduced a graph whose vertices are elements of F, where two vertices are connected if and only if the corresponding elements of F are related by one of a specially chosen set of generators of $Aut(F)$. Here we give a precise structural description of Whitehead's graph for the case where $F = F_2$ is the free group of rank two. This description allows us to quantify relationships between the natural length function $|\ |$ of F_2, and the action of $Aut(F_2)$ on F_2. As an application, we show that Whitehead's algorithm for testing automorphic conjugacy in F_2 runs in time that is at most quadratic in the length of the elements.

1. The Automorphism Graph of F_2

To start, let X denote a **base set** of elements $\{a, b, \ldots\}$, and let $X^{-1} = \{A, B, \ldots\}$ be the set consisting of the corresponding formal inverses of elements from X. We call the elements of $X \cup X^{-1}$ **letters**, and denote the **free group** on the set X as $F(X)$.

The elements of the free group can be taken as the set of freely reduced words of finite length over the alphabet $X \cup X^{-1}$, where by **freely reduced** we mean words which contain no subword of the form xx^{-1} or xx^{-1} for any $x \in X$. Multiplication of elements of $F(X)$ is simply concatenation of words, followed by **free reduction**, which is to say repeated cancellation of all subwords of the form xx^{-1} or xx^{-1} for $x \in X$. The unique empty word of length 0 plays the role of the identity element. It is well-known that given two sets X and Y the free group $F(X) \cong F(Y)$ if and only if $|X| = |Y|$. This justifies denoting such a free group as $F_{|X|}$, since upto isomorphism the group depends only on the cardinality of the base set. The group $F_{|X|}$ is called *the* free group of **rank** $|X|$. This work considers F_2, the free group of rank two on the set $X = \{a, b\}$.

Recall that for any group G, the set of automorphisms of G again forms a group, denoted $Aut(G)$, in which composition of automorphisms plays the role of multiplication. Given a group G, two elements $g, h \in G$ are said to be **automorphic conjugates** if there exists an automorphism $\phi \in Aut(G)$ for which $\phi(g) = h$.

This work concerns the properties of $Aut(F_2)$, the group of automorphisms of F_2. In general, a structural description of the orbits of F_n under the action of

$Aut(F_n)$ has several algorithmic applications. The principal one we will consider here pertains to the problem of testing automorphic conjugacy:

DEFINITION 1.0.1. *For each $n \in \mathbb{N}$, let* $\mathbf{AUT\text{-}CONJ}_n(u, v)$ *be the following decision problem:*

INPUT*:* $u, v \in F_n$

OUTPUT*:* $\begin{cases} \textit{1} & \textit{if } u \textit{ and } v \textit{ are automorphic conjugates in } F_n. \\ \textit{0} & \textit{otherwise.} \end{cases}$

Note that **AUT-CONJ** is different from the much simpler problem of testing (ordinary) conjugacy:

DEFINITION 1.0.2. *For each $n \in \mathbb{N}$, let* $\mathbf{CONJ}_n(u, v)$ *be the following decision problem:*

INPUT*:* $u, v \in F_n$

OUTPUT*:* $\begin{cases} \textit{1} & \textit{if } \exists w \in F_n \textit{ s.t. } u^w = v. \\ \textit{0} & \textit{otherwise.} \end{cases}$

To hilight the difference between $\mathbf{CONJ}_n$ and $\mathbf{AUT\text{-}CONJ}_n$, note that if u, v are conjugates in the ordinary sense, then certainly they are automorphic conjugates—since conjugation is an inner automorphism. On the other hand not every automorphism of F_n is inner, so it is possible for u and v to be automorphic conjugates but not be conjugates in the ordinary sense.

The result that $\mathbf{CONJ}_n$ is decidable is folklore. The algorithm attributed to Greendlinger is as follows: Take two cycle graphs of lengths $|u|$ and $|v|$ respectively. Write u clockwise on the edges of the first, and v along the second—these labelled graphs are called "circular words". Now perform **cyclic free reduction** on these circular words, which is to say repeatedly contract all pairs of consecutive edges with labels x, x^{-1} or x, x^{-1} (for $x \in X$). This reduction process terminates since the original words are of finite length, and their length strictly decreases at each reduction step. Upon termination of cyclic free reduction, check to see if the two circles are equal graphs, as drawn. If so, output 1. Otherwise output 0. It is not difficult to show that this procedure is correct, and can be implemented in at $O(|u|^2 \log n + |v|^2 \log n)$ time. This will be revisited in Section 4, where algorithmic issues are addressed.

In 1936, J. H. C. Whitehead proved [27, 28] that $\mathbf{AUT\text{-}CONJ}_n$ is also decidable. His argument, which will be outlined in Section 4, provided a bound of $O(2^n 2^{|u|+|v|})$. Until the recent work of A. Miasnikov and V. Shpilrain [17], this was the best known analysis of Whitehead's Algorithm. In their paper (to appear), using techniques are quite different from what is carried out here, Miasnikov and Shpilrain obtained the first polynomial-time analysis for $\mathbf{CONJ}_2$. Their analysis showed that in F_2 Whitehead's algorithm always terminates in time $O(\min(|u|, |v|)^4)$. In this work, we will provide a structural description of the orbits of F_2, making it possible for sharper analysis of Whitehead's algorithm in the case of $\mathbf{AUT\text{-}CONJ}_2$. We will show that Whitehead's algorithm always terminates in time $O(\min(|u|, |v|)^2)$. *This will bring the upper-bound complexity of best known algorithm for* $\mathbf{AUT\text{-}CONJ}_2$ *in line with the upper-bound complexity of the best known algorithm for* $\mathbf{CONJ}_2$.

1.1. Whitehead's Graph. It is well-known [13, pp. 31] that the automorphism group of F_n is generated by the set of elementary Whitehead automorphisms, which we will denote as W_n. At the end of this section, we will present the Whitehead automorphisms W_2 which generate $Aut(F_2)$. First, let us see how they will be used.

DEFINITION 1.1.1. *The* **Whitehead Graph** *of F_n is the labelled directed graph*

$$\Gamma_n = (F_n, E_n)$$

where $(u, v) \in E_n$ if there is some $\phi \in W_n$ satisfying $\phi(u) = v$. The directed edges of Γ_n are equipped with a labelling $L_n : E_n \to 2^{W_n}$ which satisfies $\phi \in L_n(u, v) \Leftrightarrow \phi u = v$.

REMARK 1.1.2. *Since the automorphisms W_n generate all of $Aut(F_n)$ [27, 28], it follows that two vertices u_1, u_2 are connected by a path in Γ_n if and only if they are conjugate via an automorphism of F_n.*

The previous remark suggests that a good understanding of the structure of Γ_n might be used to devise efficient algorithms for testing automorphic conjugacy in F_n. This is precisely what we aim to accomplish in the subsequent sections, for the special case when $n = 2$.

NOTATION 1.1.3. *Since hereafter we shall be considering F_2 and $Aut(F_2)$ almost to the exclusion of free groups of rank $n > 2$, we adopt the following simplifying notation: $\Gamma = \Gamma_2$, $W = W_2$.*

We will now describe the automorphisms W precisely. In this exposition, we partition W into three subsets: the 8 length-preserving automorphisms Π, the 8 basic shifts Ψ, and the 4 conjugations Θ.

An automorphism ϕ in $Aut(F_2)$ is called a **length-preserving** or **permutation** automorphism if $|u| = |\phi(u)|$ for all u in F_2. It is easy to see that there are 8 length-preserving automorphisms, which together form an 8 element subgroup of $Aut(F_2)$ of isomorphism type D_4. For convenience, we shall take the following 3 automorphisms as a (non-minimal) set of generators for this subgroup:

$$\pi_a : \begin{cases} a & \mapsto & A \\ b & \mapsto & b \end{cases} \qquad \pi_b : \begin{cases} a & \mapsto & a \\ b & \mapsto & B \end{cases} \qquad \hat{\pi} : \begin{cases} a & \mapsto & b \\ b & \mapsto & a \end{cases}$$

For $x \in X \cup X^{-1}$, we shall hereafter abbreviate $\hat{\pi}(x)$ as $\hat{x}$.

We denote the 8 element group generated by $\{\pi_a, \pi_b, \hat{\pi}\}$ as

$$\Pi = \{1,\ \pi_a,\ \pi_b,\ \hat{\pi},\ \pi_a\pi_b,\ \pi_a\hat{\pi},\ \pi_b\hat{\pi},\ \pi_a\pi_b\hat{\pi}\}.$$

The 8 **basic shift automorphisms** in Ψ are obtained by taking the following two automorphisms for each $x \in X^\delta$, where $\delta = \pm 1$:

$$\psi_x \quad : \quad \begin{cases} x & \mapsto & x\hat{x}^\delta \\ \hat{x} & \mapsto & \hat{x} \end{cases} \qquad {}_x\psi \quad : \quad \begin{cases} x & \mapsto & \hat{x}^\delta x \\ \hat{x} & \mapsto & \hat{x} \end{cases}$$

REMARK 1.1.4. *It is easy to check from these definitions that*

$$(\psi_x)^{-1} \quad \equiv_{F_2} \quad {}_{x^{-1}}\psi$$

for each $x \in X^{\pm 1}$.

The **conjugation automorphisms** Θ are obtained by considering the following automorphism for each $x \in X^\delta$, where $\delta = \pm 1$:

$$\theta_x \quad : \quad \begin{cases} x & \mapsto & x \\ \hat{x} & \mapsto & x^{-1}\hat{x}x \end{cases}$$

The automorphisms Θ can be obtained by composing appropriate elements from Ψ. Thus, we restrict ourselves to the smaller generating set $(\Pi \cup \Psi) \subsetneq W$. Clearly, this will not change the fundamental equivalence of the problems of automorphic conjugacy and determining whether two vertices lie in the same connected component of Γ (see Remark 1.1.2, pp. 3).

partitioning of Γ into connected components.

1.2. Symmetries in Whitehead's Graph. In the two subsections that follow, we will consider two symmetries, which give rise to **conjugacy** and **permutation** equivalence relations, respectively.

1.2.1. *Conjugacy Equivalence.* As noted in the paragraph following Definition 1.0.2 (pp. 2), it is easy to test whether two elements of F_2 are conjugates (in the ordinary sense). This suggests that we somehow "eliminate" ordinary conjugacy from Γ. Let us do this now, formally.

DEFINITION 1.2.1. *The* **conjugacy equivalence relation** $\mathcal{J}$ *on* F_2 *is defined by making* $g \sim_{\mathcal{J}} g'$ *iff* $g^w = g'$ *for some* w *in* F_2. *Clearly* $\mathcal{J}$ *is an equivalence relation since (i)* $g^1 = g$, *and (ii)* $g^w = g'$ *implies* $(g')^{w^{-1}} = g$, *and (iii)* $g^{w_1} = g'$, $(g')^{w_2} = g''$ *implies* $g^{w_1 w_2} = g''$. *We denote the equivalence classes in* F_2 *modulo* $\mathcal{J}$ *as* $\tilde{F}_2 = F_2/\mathcal{J}$, *and the elements of* $\tilde{F}_2$ *are referred to as conjugacy classes of* F_2.

NOTATION 1.2.2. *For any concrete element in* F_2 *(e.g.* $abbAB$*) we will denote its conjugacy class by enclosing the element in* $[,]$ *(e.g.* $[abbAB]$*). For a variable representing an element of* F_2 *(e.g.* g*) we denote its conjugacy class by a* ˜ *symbol (e.g.* $\tilde{g}$*). For a conjugacy class* $\tilde{g} \in \tilde{F}_2$ *we define* $|\tilde{g}|$ *to be the length of* g *after cyclic free reduction. We shall often consider an element* $\tilde{g} \in \tilde{F}$ *to be a cycle graph* $O_{\tilde{g}}$ *of length* $|\tilde{g}|$, *whose edges are labelled clockwise by successive letters in the cyclically reduced from of* g.

REMARK 1.2.3. *In this work, we adhere to several conventions regarding depictions of cyclic words as cycle graphs (see Figure 1). First, there is no distinguished vertex in the cycle graph–that is to say that one can read the cyclic word off from the cycle graph by beginning at any vertex. Second, the graphs have a distinguished orientation; we will assume that diagrams of cyclic words will always be read clockwise. Clearly, a directed edge labelled by* x *presents the same information as the reverse edge labelled by* x^{-1} *(where* x *is in* $\{a, b, A, B\}$*). Our final notational convention seeks to circumvent this ambiguity: edges are never labelled by negated generators* $\{A, B\}$*). When reading the cyclic word off from the cycle graph, it is implicitly assumed that the reader will invert the labels of all edges that are traversed in the direction opposite to their depicted orientation.*

NOTATION 1.2.4. *Given an element* w *in* $F = F(X)$ *and an automorphism* $\phi : F \to F$, *we define* $\phi(w)$ *to be the freely reduced form of the image of* w *under* ϕ. *We define* $\phi\lceil w \rceil$ *to be the unreduced word obtained by replacing every occurrence of* x *in* w *with* $\phi(x)$, *where* $x \in X^{\pm}$. *Figure 2 presents an illustration of this notation.*

OBSERVATION 1.2.5. *Given an automorphism* ϕ *of* $F = F(X)$ *and two elements* u, v *in* F *that are conjugate via* w. *Then* $u^w = v$ *implies that* $\phi(u^w) = \phi(v)$, *i.e. that* $\phi(u)^{\phi(w)} = \phi(v)$, *which is to say that* $\phi(u)$ *and* $\phi(v)$ *are conjugate via* $\phi(w)$.

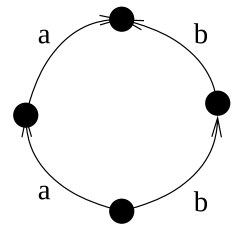

Figure 1: The cycle graph representing the conjugacy class of the words $aBBa$, $BaaB$ and $AbaBBaBa$.

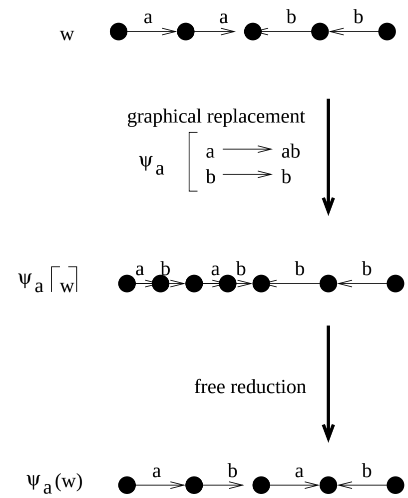

Figure 2: $\phi_a(aaBB)$ versus $\phi_a\lceil aaBB\rceil$.

In short, every automorphism ϕ induces well-defined map of conjugacy classes $\tilde{\phi}$: $\tilde{F} \to \tilde{F}$.

NOTATION 1.2.6. *Analogously, given a conjugacy class $\tilde{w}$ in $\tilde{F}$ and map of conjugacy classes $\tilde{\phi}$, we denote $\tilde{\phi}(\tilde{w})$ to be the cyclically reduced form of the image of $\tilde{w}$ under $\tilde{\phi}$. We denote $\tilde{\phi}\lceil\tilde{w}\rceil$ to be the cyclic word obtained after replacing every occurrence of x in $\tilde{w}$ with $\tilde{\phi}(x)$ but prior to performing free/cyclic reduction. Figure 3 presents an illustration of this notation.*

The following sets of maps of conjugacy classes will be of central importance in what follows:

DEFINITION 1.2.7. *Let the set of* **permutation maps** *be defined as*

$$\tilde{\Pi} = \{\tilde{\pi} \mid \pi \in \Pi\}$$

and take the set of **basic shift maps** *as*

$$\tilde{\Psi} = \{\tilde{\psi} \mid \psi \in \Psi\}.$$

DEFINITION 1.2.8. *The* **permutation equivalence relation** $\sim_\Pi$ *on* F_2 *is defined so that for any two elements $u, v \in F_2$, $u \sim_\Pi v$ holds precisely when there is some $\pi \in \Pi$ for which $\pi u = v$. The fact that this gives an equivalence relation is*

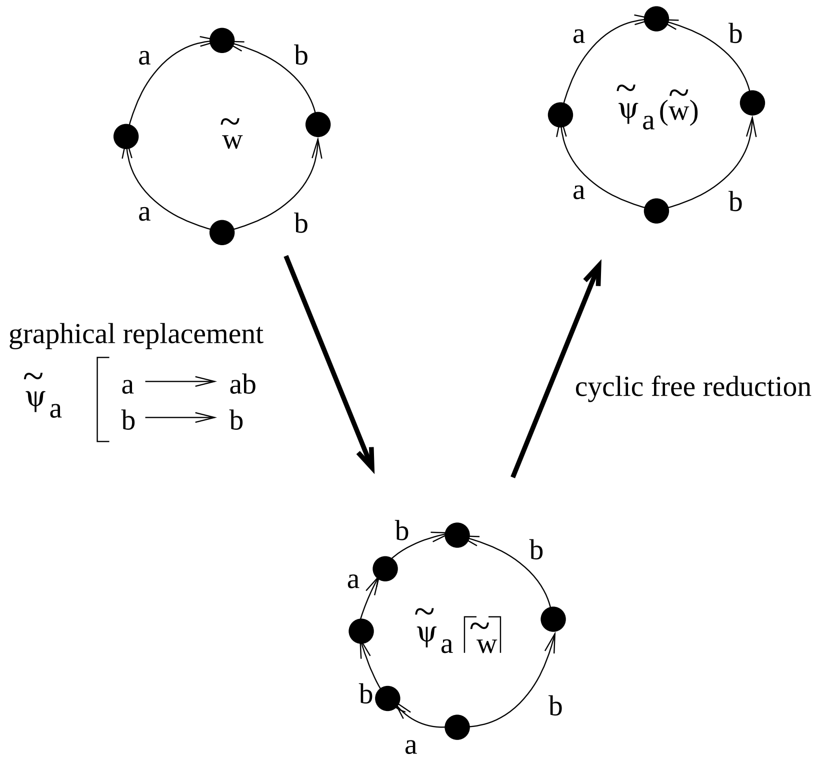

Figure 3: $\tilde{\phi}_a([aaBB])$ versus $\tilde{\phi}_a\lceil [aaBB] \rceil$.

obvious, since the permutation automorphisms include the identity and are closed under composition and inversion.

Definition 1.2.9. *The* **shift relation** $\sim_\Psi$ *is defined so that for* $u, v \in F_2$, $u \sim_\Psi v$ *holds precisely when there is some* $\psi \in \Psi$ *for which* $\psi u = v$. *The shift relation is not an equivalence relation, since in general, the composition of two basic shifts is not a basic shift.*

Since the elements of Π and Ψ are automorphisms, Observation 1.2.5 (pp. 4) applies, and the corresponding relations $\sim_\Pi$ and $\sim_\Psi$ factor through the relation $\mathcal{J}$ (See Definition 1.2.1) in the obvious way to give the equivalence relation $\approx_\Pi$ and relation $\approx_\Psi$ on $\tilde{F}_2$.

Specifically, we obtain:

Definition 1.2.10. *The* **permutation equivalence relation** $\approx_\Pi$ *on* $\tilde{F}_2$ *is defined so that for any two elements* $u, v \in \tilde{F}_2$, $u \approx_\Pi v$ *holds precisely when there is some* $\tilde{\pi} \in \tilde{\Pi}$ *for which* $\tilde{\pi} u = v$.

Definition 1.2.11. *The* **shift relation** $\approx_\Psi$ *on* $\tilde{F}_2$ *is defined so that for* $u, v \in \tilde{F}_2$, $u \approx_\Psi v$ *holds precisely when there is some* $\tilde{\psi} \in \tilde{\Psi}$ *for which* $\tilde{\psi} u = v$.

Remark 1.2.12. *Since automorphisms* ${}_x\psi$ *and* ψ_x *differ only by a conjugation, it is straightforward to verify that for every* $x \in X \cup X^{-1}$,

$$ {}_x\tilde{\psi}^i \quad \equiv_{\tilde{F}_2} \quad \tilde{\psi}^i_x, $$

are identical maps of conjugacy classes.

By Remark 1.2.12 (pp. 6) that $\tilde{\Psi} = \{\tilde{\psi}_a, \tilde{\psi}_b, \tilde{\psi}_A, \tilde{\psi}_B, {}_a\tilde{\psi}, {}_b\tilde{\psi}, {}_A\tilde{\psi}, {}_B\tilde{\psi}\}$ has only 4 distinct members (while in contrast, $|\Psi| = 8$). For concreteness, we take

$$\tilde{\Psi} = \{\tilde{\psi}_x \mid x \in X \cup X^{-1}\}, \tag{1}$$

and to avoid confusion, in what follows we will never use the names ${}_x\tilde{\psi}$ (i.e. left subscripts) to identify the elements of $\tilde{\Psi}$. Combining the above equality with Remark 1.1.4 (pp. 3), we see that

$$\tilde{\psi}_x^{-i} \quad \equiv_{\tilde{F}_2} \quad \tilde{\psi}_{x^{-1}}^{i}$$

for all $i \in \mathbb{Z}$ and x in $X \cup X^{-1}$.

In subsequent sections, we shall investigate the structure of Whitehead's Graph Γ modulo $\mathcal{J}$, which we name in the next definition:

DEFINITION 1.2.13. *The* **Automorphism Graph**

$$\Omega = (\tilde{F}_2\ ,\ \ \approx_\Pi \cup \approx_\Psi)$$

is the combinatorial object whose vertices are conjugacy classes of F_2, where every pair of vertices $u, v \in \tilde{F}_2$ is connected by a directed edge if and only if u and v are related by $\approx_\Pi$ or $\approx_\Psi$. The directed edges of Ω are equipped by a labelling

$$\ell : \tilde{F}_2 \times \tilde{F}_2 \to 2^{\tilde{\Pi} \cup \tilde{\Psi}}$$

which satisfies $\phi \in \ell(u, v) \Leftrightarrow \phi u = v$.

The Automorphism Graph Ω is simply $\Gamma/\mathcal{J}$, i.e. the graph obtained by taking Γ and identifying all vertices that are related by the conjugacy equivalence relation on F_2 (See Definition 1.2.1). Moreover, observe that if $u, v \in F_2$ and $u \sim_{\mathcal{J}} v$, then for any $\phi \in Aut(F_2)$, $\phi(u) \sim_{\mathcal{J}} \phi(v)$. This is simply because $Inn(F_2) \triangleleft Aut(F_2)$. In the case where $\phi \in (\Phi \cup \Psi)$ this means that we (i) identify vertices u and v (ii) identify vertices $\phi(u)$ and $\phi(v)$, and (iii) the edges $(u, \phi(u))$ with $(v, \phi(v))$ — both of which are labelled by ϕ. In the quotient graph Ω, the vertex $u\mathcal{J} = v\mathcal{J}$ is connected to the vertex $\phi(u)\mathcal{J} = \phi(v)\mathcal{J}$ by an edge labelled with the map of conjugacy classes $\tilde{\phi}$.

Since every conjugation is an automorphism, Remark 1.1.2 (pp. 3) implies that u and v are in the same connected component of Γ if and only if $u\mathcal{J}$ and $v\mathcal{J}$ are in the same connected component of $\Omega = \Gamma/\mathcal{J}$. We shall hereafter focus on elucidating the structure of Ω.

REMARK 1.2.14. *In this work, we adhere to several conventions regarding depictions of the automorphism graph Ω (see Figure 4). First, vertices of Ω are always labelled by (an arbitrary) minimal-length representative of the corresponding conjugacy class. A directed edge labelled by $\tilde{\psi}_x$ presents the same information as the reverse edge labelled by $\tilde{\psi}_{x^{-1}}$ (where x is in $\{a, b, A, B\}$). Our second convention seeks to circumvent this ambiguity: edges are never labelled by $\tilde{\psi}_x$ when x is a negated generator (i.e. A or B). When reading a composite automorphism (along some path in Ω) it is implicitly assumed that the reader will invert the automorphisms along any edges that are traversed in the direction opposite to their depicted orientation. Finally, in light of equation (1), there is a natural bijection between $\tilde{\Psi}$ and $X \cup X^{-1}$. For brevity, whenever an edge e is assigned a label $\ell(e) = \tilde{\psi}_x \in \tilde{\Psi}$*

for some $x \in X \cup X^{-1}$, we shall simply label e by the letter x in any figure in which it appears.

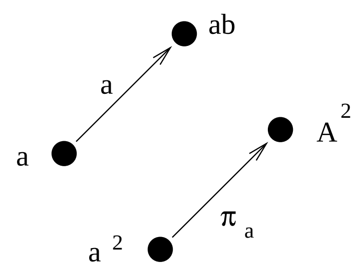

Figure 4: A small part of the automorphism graph Ω.

EXAMPLE 1.2.15. *As a concrete example then, in the automorphism graph Ω, the conjugacy classes of a, ab, a^2, A^2 are represented by four distinct vertices. The pairs of elements (a, ab) and (a^2, A^2) are connected by directed edges, labelled by a and $\tilde{\pi}_a$. These edges are present precisely because $\tilde{\psi}_a([a]) = [ab]$ and $\tilde{\pi}_a([a^2]) = [A^2]$. Figure 4 shows the corresponding region of the Automorphism Graph Ω.*

1.2.2. *Permutation Equivalence.* We will also consider the quotient graph

$$\Omega^* = (\tilde{F}_2/{\approx_\Pi}\ ,\ \ \approx_\Psi).$$

which is obtained by collapsing all vertices in Ω that are related by $\approx_\Pi$, and identifying any resulting parallel edges.

Since every permutation automorphism is in fact an automorphism, it follows from Remark 1.1.2 (pp. 3) that u and v are in the same connected component of Ω if and only if $u{\approx_\Pi}$ and $v{\approx_\Pi}$ are in the same connected component of $\Omega^* = \Omega/{\approx_\Pi}$. In what follows we will consider the structure of Ω^* as well as Ω.

1.3. Shelling Orbits in Ω and Ω^*.

DEFINITION 1.3.1. *For each conjugacy class u in $\tilde{F}_2$, let $Orb(u)$ denote the* **orbit** *of u under the action of $\tilde{\psi}$ where $\psi \in Aut(F_2)$:*

$$Orb(u) = \{\tilde{\psi}(u) \mid \psi \in Aut(F_2)\}.$$

Clearly, $\tilde{F}_2$ is partitioned into such orbits. We denote by $Orb^*(u)$ the image of $Orb(u)$ under the projection map induced by $\approx_\Pi$.

DEFINITION 1.3.2. *Given a conjugacy class u in $\tilde{F}_2$ we define the set of its* **minimal representatives** *as*

$$Min(u) = \{w \in Orb(u) \mid \forall v \in Orb(u), |w| \leqslant |v|\}.$$

Every linear order on X extends to a linear lexicographic order on $F(X)$. Here we will take $B < A < a < b$ and shall denote the lexicogaphically smallest element of $Min(u)$ as $\bar{u}$.

DEFINITION 1.3.3. *A conjugacy class u in $\tilde{F}_2$ is called* **minimal** *(in its orbit) if u is in $Min(u)$.*

Definition 1.3.4. *Let u be a conjugacy class in $\tilde{F}_2$. The* **level section of u**, *denoted $A(u)$, is the set of elements in $Orb(u)$ which have the same length as u:*

$$A(u) = \{v \in Orb(u) \ \mid \ |v| = |u|, \psi \in Aut(F_2)\}.$$

The image of $A(u)$ under the projection map induced by $\approx_\Pi$ is denoted as $A^*(u)$.

Definition 1.3.5. *For each $n \in \mathbb{N}$, we define the* **level n section** *of Ω (resp. Ω^*) to be the subgraph induced by vertices representing conjugacy classes of length n. We denote this subgraph as Ω_n (resp. Ω^*_n). A subgraph G of Ω is called a* **level subgraph** *if G is a subgraph of Ω_n for some natural number n.*

Definition 1.3.6. *Let u be a conjugacy class in $\tilde{F}_2$. We denote the subgraph of Ω induced by vertices of $A(u)$ as Ω_u. The* **level neighborhood of u**, *denoted $B(u)$, is the connected component of Ω_u which contains vertex u. The image of $B(u)$ under the projection map induced by $\approx_\Pi$ is denoted as $B^*(u)$.*

Figure 5 (on the subsequent page) metaphorically depicts the relationship between u, and the sets $B(u)$, $A(u)$, $Orb(u)$, Ω_n, when $|u| = n$.

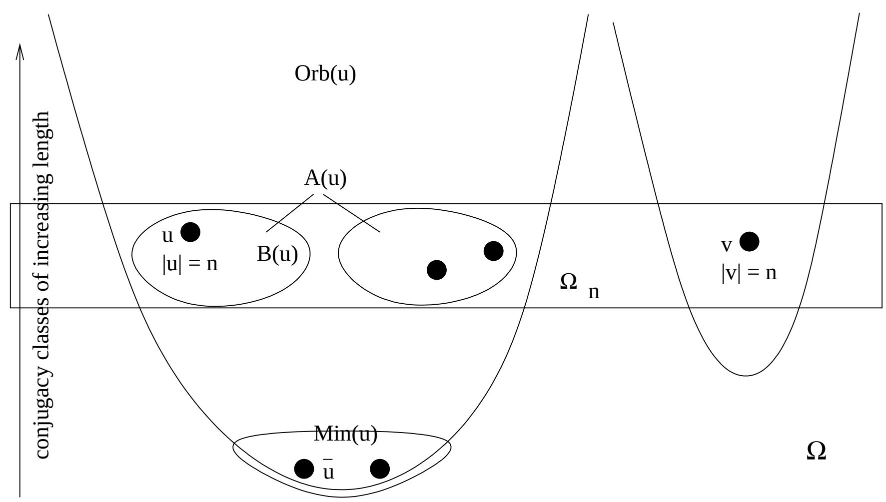

Figure 5: How u, $B(u)$, $A(u)$, $Min(u)$, $Orb(u)$ and Ω_n lie inside the Automorphism Graph Ω.

While, in general, Ω_u need not be a connected graph, we have:

Theorem 1.3.7 (Whitehead). *For any conjugacy class in u in $\tilde{F}_2$, if u is minimal then $A(u) = B(u)$.*

This is the content of Whitehead's theorem [27, 28], and is perhaps most easily understood in terms of the Diamond Property, given that the Whitehead automorphisms form a confluent rewriting system for automorphically conjugate elements of a free group.

1.4. Overview of the Result. The main result of this thesis can be stated now:

Theorem 1.4.1 (Main Theorem). *There are uniform constants C, N such that for any u in $\tilde{F}_2$ if $|u| > N$ and $|B(u)| > C$, then $|B(u)|$ is at most $8|u| - 40$.*

A more precise statement and proof of this theorem appears on page 74 of this manuscript. When this theorem is combined with Theorem 1.3.7 (pp. 9), it follows that if u is minimal, then $|A(u)|$ is bounded by $8|u| - 40$. This gives an affirmative answers to the conjecture of Miasnikov and Shpilrain in [17].

We will see that the bound of the previous theorem is tight, in the sense that for all integers n sufficiently large, there exists an element u_n having length n with the property that $|B(u_n)|$ is precisely $8|u_n| - 40$.

In addition, along the way, we will obtain considerable information about the combinatorial structure of the graphs induced by $B(u)$ and $Orb(u)$, for an arbitrary conjugacy class u in $\tilde{F}_2$. This information will be crucial to sharper analysis of algorithms for testing automorphic conjugacy in F_2.

1.4.1. *Outline of the Approach.* The proof of Theorem 1.4.1 is carried out in the following manner. We will begin by characterizing particular graphs which cannot occur in Ω^*_n. Once a sufficiently rich collection of these **graphs** has been proven to be **forbidden**, they will act as structural obstructions to the growth of spanning trees inside any connected component of Ω^*_n.

We will then give a description of maximal trees which do not contain any of the forbidden graphs. For n sufficiently large (bigger than some uniform constant N), it will be shown that **maximal legal trees** which avoid the forbidden graphs must fall into one of two categories:

(i) Trees that are chains having fewer than $n - 5$ vertices (with a very particular simple and predictable edge label structure).
(ii) Trees that have a small number of vertices (fewer than some uniform constant C).

We will conclude that any spanning tree of a connected component of Ω^*_n must be a subtree of one of the afforementioned maximal trees. It follows that for an arbitrary conjugacy class u in $\tilde{F}_2$, $B(u)$ must have a spanning tree that is a subgraph of one of the above maximal legal trees. Thus, corresponding to each of the above cases, $B(u)$ either has (i) very simple chain-like structure or (ii) is very small. Finally, we will "pull back" the structural characterization of Ω^*_n to obtain information about Ω_n, and then ultimately, about Γ.

The remainder of this document is organized as follows: In Section 2, a general combinatorial framework is established that is required to carry out proofs about forbidden graphs in Ω. Then in Section 3, this framework is employed to interpret candidate subgraphs in the automorphism graph as combinatorial constraints on conjugacy classes of F_2. Showing a subgraph is forbidden then amounts to showing that the corresponding system of combinatorial constraints is infeasible. In Section 4, algorithmic issues are discussed and the analysis of Whitehead's algorithm is tightened using the results of the prior sections. Section 5 surveys some of the computational tools that were used to explore the automorphism graph Ω^*. Finally, we conclude in Section 6.

2. Combinatorial Groundwork

As noted, our first goal is to characterize particular graphs that are forbidden in Ω^*_n. We begin by defining the universe of graphs which may or may not occur.

2.1. Hypothetical Subgraphs.

DEFINITION 2.1.1. *A* **hypothetical subgraph** *(of Ω) is a directed graph G (i)* without vertex labels, *(ii) whose edges are labelled by subsets of $\tilde{\Pi} \cup \tilde{\Psi}$, where (iii) each vertex of G has the property that every element of $\tilde{\Pi} \cup \tilde{\Psi}$ appears in the label of at most one of its incident outgoing edges.*

DEFINITION 2.1.2. *We say that a hypothetical subgraph G of Ω is* **realized** *(by a subgraph G' of Ω) if there exists an directed graph isomorphism of $\sigma : G \to G'$ which is an* **edge-label embedding**, *i.e. for which $\ell(e) \subseteq \ell(\sigma(e))$ for all $e \in E[G]$.*

DEFINITION 2.1.3. *A hypothetical subgraph G of Ω is* **realized as a level subgraph** *of Ω if it is realized by by a subgraph G' of Ω_n for some n.*

EXAMPLE 2.1.4. *Figure 6 shows a hypothetical graph G that is a chain of three vertices connected by edges labelled $\tilde{\phi}_a$ (Note that in the illustration we follow the conventions outlined in Remark 1.2.14, and label the edges simply by a). Then G is realized in Ω. In particular, we show two (of the many) realizations of G: the one on the left is a realization of G as a level subgraph, while the one on the right is not.*

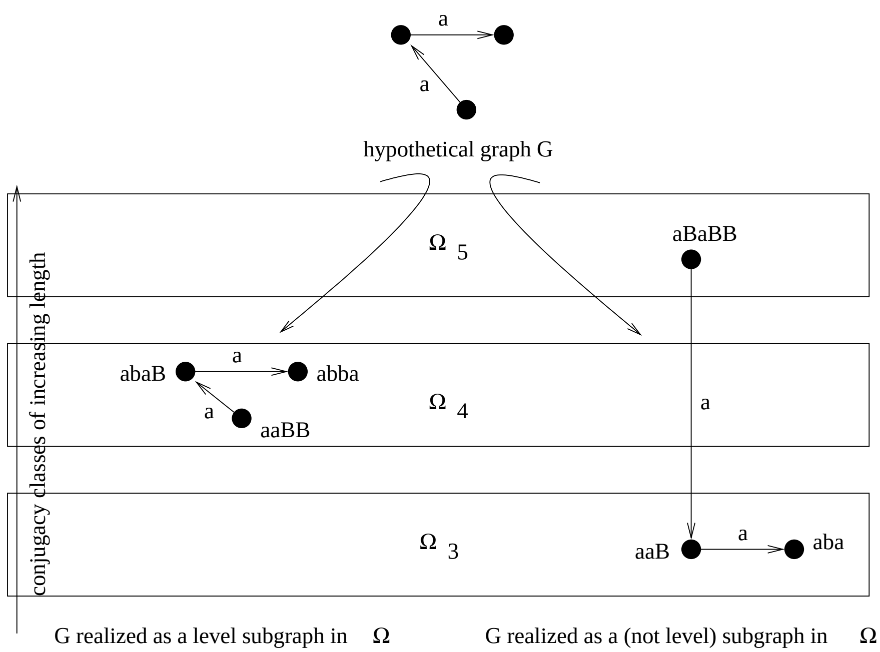

Figure 6: An example of how a hypothetical graph G might be realized in Ω.

DEFINITION 2.1.5. *A* **hypothetical subgraph** *(of Ω^*) is a directed graph G^** without vertex labels *whose edges are labelled by subsets of $\tilde{\Psi}$.*

DEFINITION 2.1.6. *We say that a hypothetical subgraph G^* of Ω^* is* **realized** *in Ω^* if*

(i) *G^* is realized by some subgraph G' of Ω, where additionally*
(ii) *no pair of vertices of G' are related by $\approx_\Pi$.*

In this event, we say that G^* is realized by the subgraph $G'/\approx_\Pi$ of Ω^*.

DEFINITION 2.1.7. *A hypothetical subgraph G^* of Ω^* is* **realized in a level subgraph** *of Ω^* if it is realized by a subgraph $G'/\approx_\Pi$ of Ω^*_n for some n.*

EXAMPLE 2.1.8. *Returning to Example 2.1.4, we see that although the conjugacy classes $aaBB$, $abaB$, $abba$ realize G in Ω, they* do not *realize G in Ω^*, since two of the conjugacy classes are related by $\approx_\Pi$. Specifically, since $aaBB \approx_\Pi abba$, this violates condition (ii) of Definition 2.1.6. Figure 7 shows that G is still realized as a level subgraph of Ω^*, albeit in ${\Omega^*}_6$, by the conjugacy classes $aaBBBB$, $abaBBB$, $abbaBB$. It is a straightforward exercise to verify that no two of these conjugacy length* 6 *classes is related by $\approx_\Pi$.*

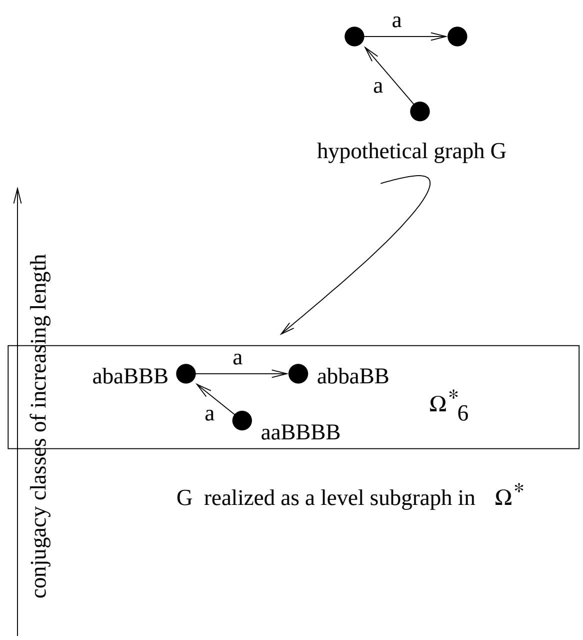

Figure 7: An example of how a hypothetical graph G might be realized as a level subgraph of in Ω^*.

REMARK 2.1.9. *It follows immediately from the previous definitions that if G^* is not realized in Ω, then G^* is not realized in Ω^*. However, if G^* is realized in Ω, this does not necessarily imply that G^* is realized in Ω^*, as the former is a priori a weaker condition. We will see an example of such a graph G^* later in Proposition 3.2.5 (pp. 35).*

2.2. Combinatorial Equations. Given a hypothetical subgraph G (of Ω), the assertion that "G is realized as a level subgraph in Ω_n for some n" entails strong constraints on the conjugacy classes that may appear as its vertices in any realization of G.

In what follows, corresponding to particular hypothetical subgraphs G (of Ω) and vertex $w \in V[G]$, we will derive a **system of combinatorial constraints** on O_w, which we will denote as $\Lambda_{G,w}$. The constraints $\Lambda_{G,w}$ will be constructed so as to have the following property:

If graph G' is any level subgraph of Ω that is isomorphic to G via a edge-label embedding $\sigma : G \to G'$, then the conjugacy class $\sigma(w)$ satisfies $\Lambda_{G,w}$.

NOTE 2.2.1. *The system $\Lambda_{G,w}$ will in general, not be a sufficient set of constraints; that is, if a conjugacy class w_0 satisfies $\Lambda_{G,w}$ this will not imply that some level neighborhood H of w_0 in Ω is isomorphic to G. Finding a necessary and sufficient system of constraints might be too difficult in general; fortunately for our objectives it is enough to find just a necessary set of constraints. Our strategy is then as follows: show that $\Lambda_{G,w}$ is infeasible, and thus conclude that no level subgraph H of Ω is isomorphic to G, i.e. that G is not realized as a level subgraph of Ω.*

To give a formal definition of the combinatorial constraints which constitute $\Lambda_{G,w}$, we need to introduce the next two functions.

DEFINITION 2.2.2. *The* **counting function**

$$\sharp : F_2 \times \tilde{F}_2 \to \mathbb{N}$$

is defined so that $\sharp(g, \tilde{w})$ is the number of distinct (possibly overlapping) subsegments that are found to be labelled by g within the cycle graph $O_{\tilde{w}}$ in a clockwise reading.[1] *For succinctness, we will frequently use the* **symmetrized counting function**

$$\circledsharp : F_2 \times \tilde{F}_2 \to \mathbb{N}$$

whose values are defined by

$$\circledsharp(g, \tilde{w}) \quad = \quad \sum_{\epsilon = \pm 1} \sharp(g^{\epsilon}, \tilde{w}).$$

EXAMPLE 2.2.3. *Let $\tilde{w}$ be the conjugacy class of $babAABBaaB$, whose cycle graph is shown in the figure below.*

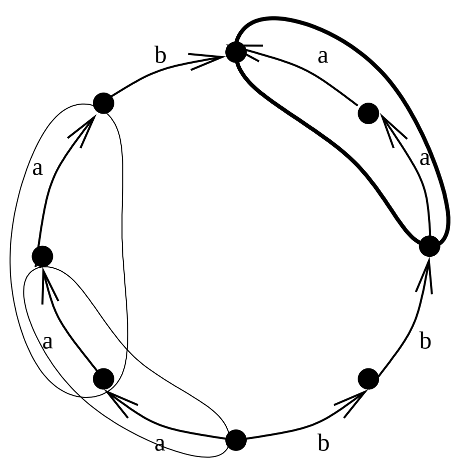

[1] The reader is advised that this is *not* the same as the number of cyclically reduced words $w' \in F_2$ whose conjugacy class is $\tilde{w}$ and that begin with subword g. This is most easily seen in the case when w is a proper power of a cyclically reduced word g.

Let $g = aa$. Then

$$\begin{aligned} \sharp(g, \tilde{w}) &= \sharp(aa, [babAABBaaB]) = 2, \text{ and} \\ \text{\textcircled{$\sharp$}}(g, \tilde{w}) &= \text{\textcircled{$\sharp$}}(aa, [babAABBaaB]) = 3. \end{aligned}$$

This can be seen by considering the cycle graph in Figure 2.2.3 and noting that it contains precisely two distinct occurrences of aa and one occurrence of AA. Note that in computing $\sharp(g, \tilde{w})$ and $\text{\textcircled{$\sharp$}}(g, \tilde{w})$, the cycle graph $O_{\tilde{w}}$ is always constructed using the cyclically reduced form of the word w.

Using $\sharp$ and $\text{\textcircled{$\sharp$}}$, a system of combinatorial equations $\Lambda_{G,w}$ will be constructed relative to a distinguished vertex w of G. Specifically, $\Lambda_{G,w}$ will be a system of equations in one variable w,

$$\begin{aligned} \lambda_1(w) &= 0 \\ \lambda_2(w) &= 0 \\ &\vdots \\ \lambda_k(w) &= 0 \end{aligned}$$

with each combinatorial equation $\lambda_j(w)$ $(j = 1, \ldots, k)$ having the form

$$\sum_{i=0}^{m_i} n_j \cdot \sharp(g_i, w)$$

where $m_j \in \mathbb{N}$, $n_i \in \mathbb{Z}$, $g_i \in F_2$ are constants depending on G. A **solution** of $\Lambda_{G,w}$ is a conjugacy class $w_0 \in \tilde{F}_2$ such that $\lambda_j(w_0) = 0$ for every $j = 1, \ldots, k$.

In the next section (subsections 3.2 and 3.4) we will construct $\Lambda_{G,w}$ for very particular classes of graphs and show that these systems of combinatorial equations are infeasible, hence proving that the corresponding graphs are not realized as level subgraphs of Ω. First, however, we need to develop some combinatorial groundwork describing the structure of conjugacy classes (viewed as cycle graphs), and quantifying how the function $\sharp$ behaves when composed with powers of a basic shift $\tilde{\psi}_x \in \tilde{\Psi}$. The remainder of this section is devoted to these objectives.

2.3. Basic Properties of $\sharp$. We seek to quantify how $\sharp$ behaves when composed with powers of a basic shift $\tilde{\psi}_x \in \tilde{\Psi}$. To start, we show that applying $\tilde{\psi}_x$ to a conjugacy class w cannot alter the number of occurrences of x or x^{-1} in w.

LEMMA 2.3.1 (Stability Lemma). *For all $w \in \tilde{F}$, $x \in X \cup X^{-1}$, $\epsilon \in \{-1, +1\}$ and $i \in \mathbb{Z}$,*

$$\sharp(x^\epsilon, \tilde{\psi}_x^i w) \quad = \quad \sharp(x^\epsilon, w)$$

PROOF. *Base case* when $i = 1$: Starting with the cycle graph O_w, we apply $\tilde{\psi}_x$ graphically (i.e. without performing cyclic free reduction), and refer to the resulting cycle graph as $\tilde{\psi}_x(O_w)$. Clearly, when read clockwise, the number of edges that are labelled $x^{\pm\epsilon}$ in $\tilde{\psi}_x(O_w)$ is the same as the number in O_w, namely $\sharp(x^{\pm\epsilon}, w)$. Suppose, towards contradiction, that $\sharp(x^\epsilon, \tilde{\psi}_x w) < \sharp(x^\epsilon, w)$. Then in any cancellation diagram that describes the transformation of $\tilde{\psi}_x(O_w)$ into the cycle graph $O_{\tilde{\psi}_x w}$ there must be some occurrence of x^ϵ and $x^{-\epsilon}$ which mutually cancel. We fix a cancellation diagram and use this to select the mutually cancelling occurrences of x^ϵ, $x^{-\epsilon}$ that are closest together in the graph $\tilde{\psi}_x(O_w)$. It follows that $\tilde{\psi}_x(O_w)$ is a cycle graph labelled by a word $x^\epsilon \cdot u \cdot x^{-\epsilon} \cdot v$ where either u or v:

(i) Consists entirely of edges labelled by $\hat{x}$, $\hat{x}^{-1}$, and
(ii) Freely reduces to the trivial word.

Suppose u satisfies (i) and (ii). Then $\tilde{\psi}_x(O_w)$ contains $x^\epsilon \cdot u \cdot x^{-\epsilon}$ as a subsegment which freely reduces to 1, and the graphic pre-image of this segment is a subsegment of O_w whose length is at least 2 and yet freely reduces to 1. This contradicts the fact that cycle graph O_w was labelled by the cyclically reduced word w. The situation when v satisfies (i) and (ii) is completely analogous.

Inductive step: Suppose $\sharp(x^\epsilon, \tilde{\psi}_x^{i-1}w') = \sharp(x^\epsilon, w')$ for all $w' \in \tilde{F}_2$. Then taking $w = \tilde{\psi}_x^{i-1}w'$ and appealing to the previous argument, we conclude that

$$\sharp(x^\epsilon, \tilde{\psi}_x^{i}w') \quad = \quad \sharp(x^\epsilon, \tilde{\psi}_x^{i-1}w'),$$

where by inductive hypothesis, the right-hand side is $\sharp(x^\epsilon, w')$. By induction on i, the lemma is proved for $i \geqslant 0$.

By Remark 1.1.4 (pp. 3), $\tilde{\psi}_x^{-i} \equiv \tilde{\psi}_{x^{-1}}^{i}$, so applying the previous argument to $\tilde{\psi}_{x^{-1}}$, the lemma is proved for $i \leqslant 0$. □

OBSERVATION 2.3.2. *Take the basic shift automorphism ψ_a and let w be an arbitrary conjugacy class in $\tilde{F}_2$. An occurrence of $\psi_a^{-1}a = aB$ might arise in $\tilde{\psi}_a(w)$ in one of two ways: to see this, consider a cancellation diagram that describes the transformation of $\tilde{\psi}_a(O_w)$ into the cycle graph $O_{\tilde{\psi}_a w}$. Then aB can arise in $O_{\tilde{\psi}_a w}$ because of either*

(1) *an occurrence of aBB in O_w, or*
(2) *an occurrence of aBA in O_w. Note that aBA is stable under $\tilde{\psi}_a$ and so will always gives rise to an occurrence of aB under $\tilde{\psi}_a$.*

In other words, $\psi_a^{-1}a$ arises in $\tilde{\psi}_a(w)$ either because of an occurrence of $aBB = \psi_a^{-2}a$ in w, or because of an occurrence of $aBA = (\psi_a^{-1}a)a^{-1}$ in w.

The prior observation is merely an illustration of the following general result:

LEMMA 2.3.3. *Let w in $\tilde{F}_2$ be a conjugacy class, and ψ_x a basic shift. Comparing the structure of conjugacy classes w and $\tilde{\psi}_x(w)$ we find that every occurrence of $\psi_x^{-k}x^\epsilon$ in $\tilde{\psi}_x(w)$ arises in one of two ways.*

(1) *an occurrence of $\psi_x^{-(k+1)}x^\epsilon$ in w, or*
(2) *an occurrence of $[(\psi_x^{-k}x)x^{-1}]^\epsilon$ in w.*

PROOF. Let $x \in X^\delta$, $\delta \in \{+1, -1\}$. Consider the case when $\epsilon = +1$; then $\psi_x^{-k}x = x\hat{x}^{-\delta k}$. The subsequent argument is illustrated in Figure 8.

As described in Definitions 1.2.4 (pp. 4) and 1.2.6 (pp. 5), we denote the word obtained by graphically applying ψ_x to w without performing any reduction as $\psi_x\lceil w\rceil$. In other words, $\psi_x\lceil w\rceil$ is the word obtained by replacing every occurrence of x in w by $x\hat{x}^\delta$. Denote the freely reduced form of $\psi_x\lceil w\rceil$ as $\psi_x(w)$.

By hypothesis, $\psi_x(w)$ contains $x\hat{x}^{-\delta k}$. Fix any such occurrence α_2. Since $\psi_x(w)$ was obtained from $\psi_x\lceil w\rceil$ by free reduction, α_2 has a preimage α_1 in $\psi_x\lceil w\rceil$. Since $\psi_x\lceil w\rceil$ was obtained by graphically applying ψ_x to w, α_1 has a preimage α_0 in w.

By Lemma 2.3.1 (pp. 14), α_0, α_1 and α_2 all begin with x. Since ψ_x maps $x \mapsto x\hat{x}^\delta$, it follows that $\psi_x\lceil w\rceil$ must be $x\hat{x}^\delta\hat{x}^{-\delta(k+1)}$. From this, we get that $\alpha_0 = x \circ \alpha_0'$, where ψ_x maps α_0' to a word beginning with $\hat{x}^{-\delta(k+1)}$. It follows that α_0' is either $\hat{x}^{-\delta(k+1)}$ or $\hat{x}^{-\delta k}x^{-1}$. Hence α_0 is either $x\hat{x}^{-\delta(k+1)}$ or $x\hat{x}^{-\delta k}x^{-1}$. The former is $\psi_x^{-(k+1)}x^\epsilon$, while the latter is $(\psi_x^{-k}x)x^{-1}$. This proves the case when

$\epsilon = +1$. The case when $\epsilon = -1$ is completely analogous, and can be carried out using the same figure reflected along the vertical axis. $\square$

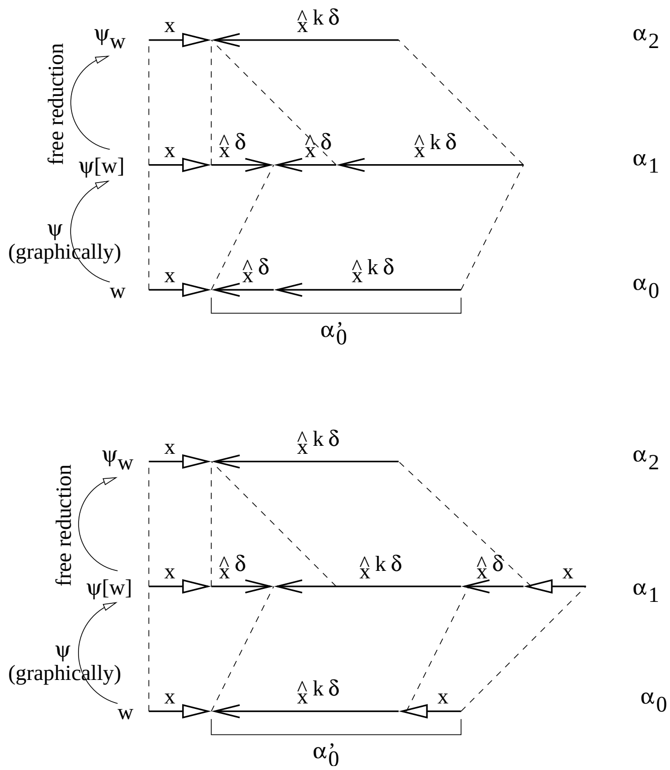

Figure 8: Ways in which $\psi_x^{-k} x$ may arise in $\tilde{\psi}_x(w)$.

The previous lemma suggests the next three definitions:

DEFINITION 2.3.4. *For each $x \in X \cup X^{-1}$, $\epsilon \in \{+1, -1\}$, and $k \in \mathbb{N}$, $k \geqslant 1$, the two words*

$$S_{x,k,\epsilon} \stackrel{def}{=} [(\psi_x^{-k} x) x^{-1}]^{\epsilon}$$

$\epsilon = \pm 1$ *are called* $\psi_{\mathbf{x}}$**-stable words of weight k**. *Note that $\psi_x S_{x,k,\epsilon} = S_{x,k,\epsilon}$.*

The table below lists ψ_x-stable words of weight k, for $x \in X \cup X^{-1}$, $\epsilon \in \{+1, -1\}$.

x	$\epsilon = +1$	$\epsilon = -1$
a	aB^kA	ab^kA
b	bA^kB	ba^kB
A	AB^ka	Ab^ka
B	BA^kb	Ba^kb

DEFINITION 2.3.5. *Given any $w \in \tilde{F}_2$, $x \in X \cup X^{-1}$ and $k \in \mathbb{Z}$, define*

$$s(w, x, k) = \begin{cases} \sharp((\psi_x^{-k} x) x^{-1}, w) + \sharp(x(\psi_x^{-k} x^{-1}), w) & \textit{if } k \geqslant 1 \\ 0 & \textit{otherwise.} \end{cases}$$

The quantity $s(w,x,k)$ counts the total number of occurrences of $\psi_x^{-k}x$ and $\psi_x^{-k}x^{-1}$ in $\tilde{\psi}_x(w)$ that arise because of some ψ_x-stable subword (of weight k) inside w. Note that for a fixed natural number k, the quantity $s(w,x,k)$ is expressible as a finite expression in terms of the function $\sharp$.

The following derived expression will also be useful.

DEFINITION 2.3.6. *Given any $w \in \tilde{F}_2$, $x \in X \cup X^{-1}$ and $i,j \in \mathbb{Z}$ define*

$$s(w,x,i,j) \quad = \quad \sum_{k=i}^{i+j-1} s(w,x,k)$$

The quantity $s(w,x,i,j)$ counts the total number of ψ_x-stable words of weight k, where $i \leqslant k < i+j$. Note that for a fixed natural numbers i and j, the quantity $s(w,x,i,j)$ is expressible as a finite expression in terms of the function $\sharp$.

The next lemma formalizes a combinatorial fact about basic shift automorphisms, which will be used many times as the foundation for inductive arguments involving the $\sharp$ function.

LEMMA 2.3.7 (Exponent Transfer Lemma). *For all $w \in \tilde{F}_2$, $x \in X \cup X^{-1}$ and $i,j \geqslant 0$,*

$$\circledsharp(\psi_x^{-i}x, \tilde{\psi}_x^j w) \quad = \quad s(w,x,i,j) + \circledsharp(\psi_x^{-i-j}x, w)$$

PROOF. For each $x \in X^\delta$, $\epsilon, \delta \in \{+1,-1\}$, the definition of $\psi_x \in \Psi$ implies that

$$\psi_x^{-i}x^\epsilon = (x\hat{x}^{-i\delta})^\epsilon,$$

and that for all $j \geqslant 0$

$$\tilde{\psi}_x^j : \left\{ \begin{array}{lll} x^\epsilon & \mapsto & (x\hat{x}^{j\delta})^\epsilon \\ \hat{x} & \mapsto & \hat{x} \end{array} \right. .$$

It follows from Lemma 2.3.3 (pp. 15) that every occurrence of $(x\hat{x}^{-i\delta})^\epsilon$ in $\tilde{\psi}_x^j w$ corresponds either to:

(1) an occurrence of $(x\hat{x}^{(-i-j)\delta})^\epsilon$ in w, or
(2) an occurrence of $[(\psi_x^{-k}x)x^{-1}]^\epsilon$ in w, for some $k = i, \ldots, i+j-1$.

But $(x\hat{x}^{(-i-j)\delta})^\epsilon$ is the same as $\psi_x^{-i-j}x^\epsilon$, so the number of occurrences of type (1) is

$$\sum_{\epsilon=\pm 1} \sharp(\psi_x^{-i-j}x^\epsilon, w).$$

On the other hand, the number of occurrences of type (2) is

$$\sum_{k=i}^{i+j-1} s(w,x,k),$$

which by definition equals $s(w,x,i,j)$. The lemma is proved. □

Since we are interested in the interplay between the action of $Aut(F)$ on F and the length function on F, it is natural to ask about the "rate of change of length" of a conjugacy class w, under powers of a basic shift $\tilde{\psi}_x$. The last two lemmas of the previous section can be used to determine a closed form expression for this quantity.

LEMMA 2.3.8 (Rate of Change of Length Lemma). *For every* $w \in \tilde{F}_2$, $x \in X \cup X^{-1}$, $\tilde{\psi}_x \in \tilde{\Psi}$ *and* $i \geqslant 1$,

$$|\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| \quad = \quad \textcircled{\sharp}(x, w) - 2[s(w, x, 1, i-1) + \textcircled{\sharp}(\psi_x^{-i} x, w)].$$

PROOF. Let $w' = \tilde{\psi}_x^{i-1} w$. Then $|\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| = |\tilde{\psi}_x(w')| - |w'|$. Now for any $w' \in \tilde{F}_2$ it is straightforward to verify that

$$\begin{aligned} |\tilde{\psi}_x w'| &= \sharp(\hat{x}, w') + \sharp(\hat{x}^{-1}, w') \\ &\quad +2[\sharp(x, w') + \sharp(x^{-1}, w')] \\ &\quad -2[\sharp(\psi_x^{-1} x, w') + \sharp(\psi_x^{-1} x^{-1}, w')] \\ |w'| &= \sharp(x, w') + \sharp(x^{-1}, w') + \sharp(\hat{x}, w') + \sharp(\hat{x}^{-1}, w'). \end{aligned}$$

So it follows that

$$\begin{aligned} |\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| &= \sharp(x, \tilde{\psi}_x^{i-1} w) + \sharp(x^{-1}, \tilde{\psi}_x^{i-1} w) \\ &\quad -2[\sharp(\psi_x^{-1} x, \tilde{\psi}_x^{i-1} w) + \sharp(\psi_x^{-1} x^{-1}, \tilde{\psi}_x^{i-1} w)] \end{aligned} \tag{2}$$

Since $i \geqslant 1$, the Stability Lemma 2.3.1 (pp. 14) yields

$$\sharp(x^\epsilon, \tilde{\psi}_x^{i-1} w) \quad = \quad \sharp(x^\epsilon, w)$$

for $\epsilon = \pm 1$, while the Rate of Change of Length Lemma 2.3.7 (pp. 17) yields

$$\sum_{\epsilon = \pm 1} \sharp(\psi_x^{-1} x^\epsilon, \tilde{\psi}_x^{i-1} w) \quad = \quad s(w, x, 1, i-1) + \sum_{\epsilon = \pm 1} \sharp(\psi_x^{-i} x^\epsilon, w).$$

Substituting into equation (2) above, the proof follows. □

In the case where $i = 1$, the previous proposition takes a particularly simple form:

COROLLARY 2.3.9. *For all* $w \in \tilde{F}_2$ *and* $x \in X \cup X^{-1}$

$$|\psi_x w| - |w| \quad = \quad \textcircled{\sharp}(x, w) - 2\textcircled{\sharp}(\psi_x^{-1} x, w)$$

PROOF. By Lemma 2.3.8 (pp. 18), we know that for all $w \in \tilde{F}_2$, $\tilde{\psi}_x \in \tilde{\Psi}$ and $i \geqslant 1$

$$|\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| \quad = \quad \textcircled{\sharp}(x, w) - 2[s(w, x, 1, i-1) + \textcircled{\sharp}(\psi_x^{-i} x, w)].$$

But when $i = 1$, $s(w, x, 1, 0) = 0$, and so the corollary follows. □

2.4. Eliminating ψ_x-Stable Subwords. Notice that if $\psi = \psi_x$ is a basic shift automorphism, then subwords of a conjugacy class w that are ψ_x-stable *do not* influence whether $|\tilde{\psi} w| = |w|$. On the other hand, ψ_x-stable words are being counted by the s term in the statement of the Rate of Change of Length Lemma 2.3.8 (pp. 18):

$$|\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| \quad = \quad \textcircled{\sharp}(x, w) - 2[s(w, x, 1, i-1) + \textcircled{\sharp}(\psi_x^{-i} x, w)].$$

Informally speaking, what we are about to do is eliminate the ψ_x-stable words from w (since they are not relevant to the question of whether $|\tilde{\psi} w| = |w|$) and then simplify the expression of the Rate of Change of Length Lemma 2.3.8 (pp. 18). Let us proceed to do this now, formally.

Definition 2.4.1. *For each $x \in X \cup X^{-1}$, let us take*

$$Z_x \stackrel{def}{=} \{z_{x,k} \mid k \in \mathbb{N}\}$$

to be a new set of letters. Let $Z = Z_a \cup Z_b \cup Z_A \cup Z_B$. Put $F^+ \stackrel{def}{=} F_2 * F(Z)$ to be the free product of $F(\{a,b\})$ and $F(Z)$.

To start, a close inspection of the table of ψ_x-stable words (pp. 16) reveals

Observation 2.4.2. *For any fixed $x \in X \cup X^{-1}$*

(a) *One ψ_x-stable word cannot be a proper subword of another ψ_x-stable word.*
(b) *Two distinct ψ_x-stable words cannot overlap inside w.*

These two observations make the maps κ_x below well-defined.

Definition 2.4.3. *For each $x \in X \cup X^{-1}$, we define a map $\kappa_x : S_{x,k,\epsilon} \mapsto z^{\epsilon}_{x,k}$, where $S_{x,k,\epsilon}$ is one of the ψ_x-stable words of weight k in Definition 2.3.4 (pp. 16).*

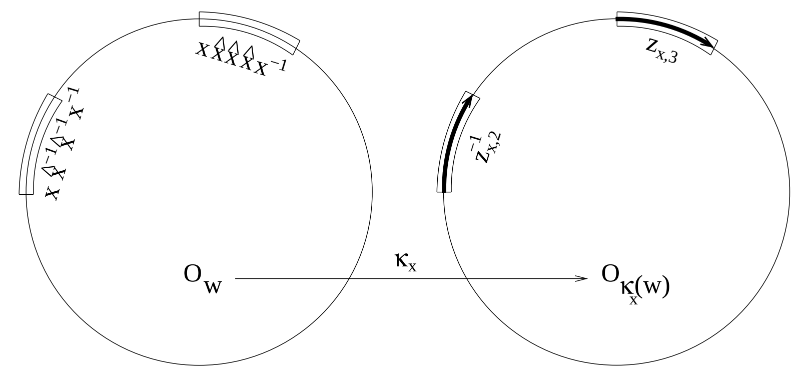

Figure 9: How κ_x acts on a conjugacy class w.

The map κ_x replaces the two ψ_x-stable words of weight k by new letters $z^{\pm 1}_{x,k}$ (see Figure 9). It follows that κ_x is a map from the conjugacy classes of F_2 to the conjugacy classes of $F_2 * F(Z_x) \subset F^+$. It will turn out that κ_x has several useful properties which will make it beneficial to consider $\kappa_x(\tilde{w})$ in place of $\tilde{w}$ at various points in our investigation of the structure of level neighborhoods in Ω.

To express these properties formally, we shall first extend the $\sharp$ function in the obvious way to all of $F^+ \times \tilde{F}^+$. As in Definition 2.2.2 (pp. 13), for any $g \in F^+$ and $\tilde{w} \in \tilde{F}^+$, we define $\sharp(g, \tilde{w})$ to be the number of distinct (possibly overlapping) subsegments that are found to be labelled by g within the cycle graph $O_{\tilde{w}}$ in a clockwise reading.

Next, we extend $s(w, x, i, j)$ to all $w \in F^+$, $x \in X \cup X^{-1}$, $i, j \in \mathbb{N}$. The definition of $s(w, x, i, j)$ is as in Definition 2.3.6 (pp. 17) except that we use the extended $\sharp$ function defined above.

Finally, we extend the basic shifts ψ_x (pp. 3) to all of F^+ by making them act identically on the free factor $F(Z)$.

Lemma 2.4.4. *For each $w \in F_2$, $x \in X \cup X^{-1}$, and $j \in \mathbb{N}$*

$$s(\kappa_x(\tilde{w}), x, 1, j) \quad = \quad 0.$$

PROOF. In $\kappa_x(\tilde{w})$, all ψ_x-stable subwords of weight k have been replaced by $z_{x,k}^{\pm 1}$. By definition, $s(\kappa_x(\tilde{w}), x, 1, j)$ is the sum of $s(\kappa_x(\tilde{w}), x, k)$ for $k = 1, \ldots, j$. But $s(\kappa_x(\tilde{w}), x, k)$ counts ψ_x-stable subwords of $\kappa_x(\tilde{w})$ having weight k, and so is 0. The lemma is proved. □

LEMMA 2.4.5. *For each $w \in F_2$ and $x \in X \cup X^{-1}$*

$$|\psi_x \kappa_x(\tilde{w})| - |\kappa_x(\tilde{w})| \quad = \quad |\psi_x \tilde{w}| - |\tilde{w}|$$

PROOF. If u is a ψ_x-stable word, then $\psi_x(u) = u$. Moreover, u begins with x and ends with x^{-1}. Combining this with Lemma 2.3.1 (pp. 14), we see that ψ_x-stable subwords in $\tilde{w}$ remain unaffected by application of ψ_x. Similarly, because ψ_x was extended to act identically on the free factor $F(Z)$ of F^+, all occurrences of $z_{x,k}^{\pm 1}$ also remain unaffected by application of ψ_x.

Since κ_x replaced ψ_x-stable words of weight k by new letters $z_{x,k}^{\pm 1}$, it follows that ψ_x-stable subwords in $\tilde{w}$ partition the cycle graph $O_{\tilde{w}}$ in precisely the same way that occurrences of $z_{x,k}^{\pm 1}$ partition the cycle graph $O_{\kappa_x(\tilde{w})}$. Since κ_x does not alter anything other than ψ_x-stable words of $\tilde{w}$, it follows that $\psi_x \kappa_x(\tilde{w})$ is the same as $\kappa_x \psi_x(\tilde{w})$. The assertion is proved. □

The next proposition is the "simplified" version of Lemma 2.3.8 (pp. 18), where ψ_x-stable words of weight k have been replaced by new free variables $z_{x,k}^{\pm 1}$.

PROPOSITION 2.4.6. *For each $w \in F_2$, $i \in \mathbb{N}$ and $x \in X \cup X^{-1}$,*

$$|\psi_x^i \tilde{w}| - |\psi_x^{i-1} \tilde{w}| \quad = \quad \textcircled{$\sharp$}(x, \kappa_x(\tilde{w})) - 2\textcircled{$\sharp$}(\psi_x^{-i} x, \kappa_x(\tilde{w}))$$

The reader may wish to compare it with the statement of Corollary 2.3.9 (pp. 18) which asserted that: *For all $w \in \tilde{F}_2$ and $x \in X \cup X^{-1}$*

$$|\psi_x w| - |w| \quad = \quad \textcircled{$\sharp$}(x, w) - 2\textcircled{$\sharp$}(\psi_x^{-1} x, w)$$

PROOF. (Proposition 2.4.6) By Lemma 2.4.5 (pp. 20)

$$|\tilde{\psi}_x^i \tilde{w}| - |\tilde{\psi}_x^{i-1} \tilde{w}| \quad = \quad |\tilde{\psi}_x^i \kappa_x(\tilde{w})| - |\tilde{\psi}_x^{i-1} \kappa_x(\tilde{w})|$$

Applying Lemma 2.3.8 (pp. 18) to $\kappa_x(w)$ we get

$$|\tilde{\psi}_x^i \kappa_x(\tilde{w})| - |\tilde{\psi}_x^{i-1} \kappa_x(\tilde{w})| \quad = \quad \textcircled{$\sharp$}(x, \kappa_x(\tilde{w})) - 2[s(\kappa_x(\tilde{w}), x, 1, i-1) + \textcircled{$\sharp$}(\tilde{\psi}_x^{-i} x, \kappa_x(\tilde{w}))].$$

By Lemma 2.4.4 (pp. 19)

$$s(\kappa_x(\tilde{w}), x, 1, i-1) \quad = \quad 0.$$

This completes the proof. □

Frequently, we will translate a hypothetical subgraph G (and a vertex $w \in V[G]$) into constraints on $\kappa_x(w)$ for some $x \in X^{\pm}$ (rather than constraints on w itself). In other words, we shall be constructing systems of combinatorial equations $\Lambda_{G,w}^x$ with the property that

> If graph G' is any level subgraph of Ω that is isomorphic to G via an edge-label embedding $\sigma : G \to G'$, then the conjugacy class $\kappa_x(\sigma(w))$ satisfies $\Lambda_{G,w}^x$.

Our strategy will remain the same: to show that $\Lambda_{G,w}^x$ is infeasible, and hence conclude that G is not realized as level subgraphs of Ω.

Observation 2.4.7. *The following facts will be used later, and can be immediately verified by inspecting the table of ψ_x stable subwords on pp. 16. For any fixed x in $X \cup X^{-1}$:*

(i) *Every ψ_x-stable word has length at least* 3.
(ii) *A ψ_x-stable word and a $\psi_{x^{-1}}$-stable word can overlap in* 0 *or* 1 *letters.*
(iii) *A ψ_x-stable word and a $\psi_{\hat{x}^{-1}}$-stable word can overlap in* 0 *or* 2 *letters.*
(iv) *A ψ_x-stable word and a $\psi_{\hat{x}}$-stable word can overlap in* 0 *or* 2 *letters.*

2.5. x-Decompositions. To be able to carry out arguments about whether or not a conjugacy w in $\tilde{F}_2$ class can satisfy a given system of combinatorial equations, we will need some machinery to describe the internal stucture of w, or more honestly, the structure of the cycle graph O_w. In this section, we introduce the x-decomposition of a conjugacy class $w \in \tilde{F}_2$, for each $x \in X \cup X^{-1}$.

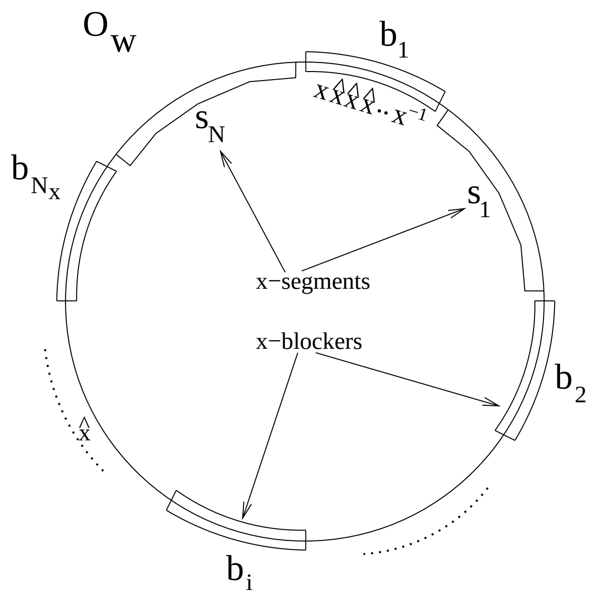

Figure 10: An x-decomposition of w.

Definition 2.5.1. *Given a conjugacy class $w \in \tilde{F}_2$ and a letter $x \in X \cup X^{-1}$ the* **x-decomposition** *of w is a decomposition of the cycle graph O_w in which all ψ_x-stable words are distinguished. Each occurrence of a ψ_x-stable word is referred to as an* **x-blocker**. *More specifically, an occurrence of $S_{x,k,+1}$ is referred to as a* **positive x-blocker of weight k** *and an occurrence of $S_{x,k,-1}$ is called a* **negative x-blocker of weight k**. *The number of x-blockers in w is denoted $N_x = N_x(w)$, and their weights will be denoted $b_0^x(w), \ldots, b_{N_x-1}^x(w)$ respectively. Abusing the notation, we will also refer to the ith blocker (considered as a subgraph of O_w) as b_i^x. When x is clearly understood from the context, we will omit the superscript x and simply refer to the blockers and their weights as $b_0(w), \ldots, b_{N_x-1}(w)$. We denote the number of* $\mathbf{x}$*-blockers in w having weight k, as $N_x^k = N_x^k(w)$. So $N_x(w) = \sum_{k \geqslant 1} N_x^k(w)$, where*

$$N_x^k(w) \quad = \quad \sharp(x\hat{x}^k x^{-1}, w), \text{ and}$$

Maximal sequences of edges in O_w which lie entirely between two consecutive x-blockers will be called **x-segments**. *The x-segment between x-blocker i and x-blocker $(i+1) \mod N_x$ will be referred to as the ith x-segment, and will be denoted s_i^x. When x is clearly understood from the context, we will omit the superscript x and simply refer to the segments as $s_0(w), \ldots, s_{N_x-1}(w)$. Figure 10 illustrates an x-decomposition of a conjugacy class w. Note that items (a) and (b) in Observation 2.4.2 (pp. 19) make the notion of an* **x-decomposition** *well-defined.*

The proof of the following lemma is immediate, since every x-blocker contains precisely one occurrence of $\psi_{x^{-1}}x^{\epsilon}$ and precisely two occurrences of $x^{\pm}$.

LEMMA 2.5.2 (κ_x Lemma).

$$\begin{aligned} \#(\psi_{x^{-1}}x, w) &= N_x(w) + \#(\psi_{x^{-1}}x, \kappa_x(w)) \\ \#(x, w) &= 2N_x(w) + \#(x, \kappa_x(w)) \end{aligned}$$

LEMMA 2.5.3 (Demarcator Structure Lemma). *Given a conjugacy class $w \in \tilde{F}_2$, fix any $x \in X \cup X^{-1}$. In an x-decomposition of w, consider two adjacent x-blockers b_i, b_{i+1} witnessed in some fixed clockwise reading of the cycle graph O_w. Then, within the intervening segment s_i, all occurrences of x^{-1} must precede any occurrences of x.*

PROOF. Suppose the lemma is false, towards contradiction. Then in s_i there is at least one occurrence of x which precedes some occurrence of x^{-1}. Consider the occurences of x and x^{-1} having this property that are closest together inside s_i. Then these delineate a x-blocker inside s_i. But this contradicts the definition of b_{i+1} as the next blocker after b_i in the x-decomposition of w. $\square$

The previous Lemma 2.5.3 indicates that the structure of an x-segment must necessarily be as shown below in Figure 11.

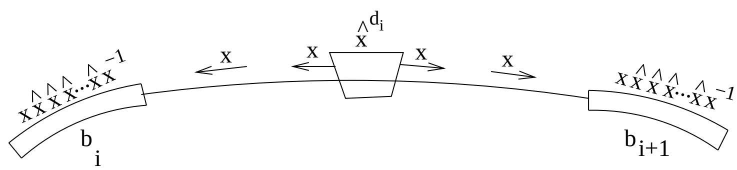

Figure 11: Demarcator between two x-blockers in the x-decomposition of w.

This leads us to the next definition.

DEFINITION 2.5.4. *Fix an x-decomposition of a conjugacy class w. Consider the last occurence of x^{-1} in s_i and the first occurrence of x in s_i. Then non-cancellation between these two letters implies that there must be at least one occurrence of $\hat{x}$ or $\hat{x}^{-1}$ between them. The subsequence $\hat{x}^{d_i}$ or $(\hat{x}^{-1})^{d_i}$ between the last occurence of x^{-1} in s_i and the first occurrence of x in s_i is called the ith* **x-demarcator** *and is said to have weight d_i. Abusing the notation, we will also refer to the ith x-demarcator (considered as a subgraph of O_w) as d_i.*

The previous definition immediately provides the following useful lemmas:

LEMMA 2.5.5 (Plus-1 Lemma). *Given a conjugacy class $w \in \tilde{F}_2$, and fixed x in $X \cup X^{-1}$,*

$$\#(x, w) \geqslant \sum_{i=1}^{N_{\hat{x}}(w)}(b_i + d_i) \geqslant \max\left(\sum_{i=1}^{N_{\hat{x}}(w)}(b_i^{\hat{x}} + 1),\ \sum_{i=1}^{N_{\hat{x}}(w)}(d_i^{\hat{x}} + 1)\right)$$

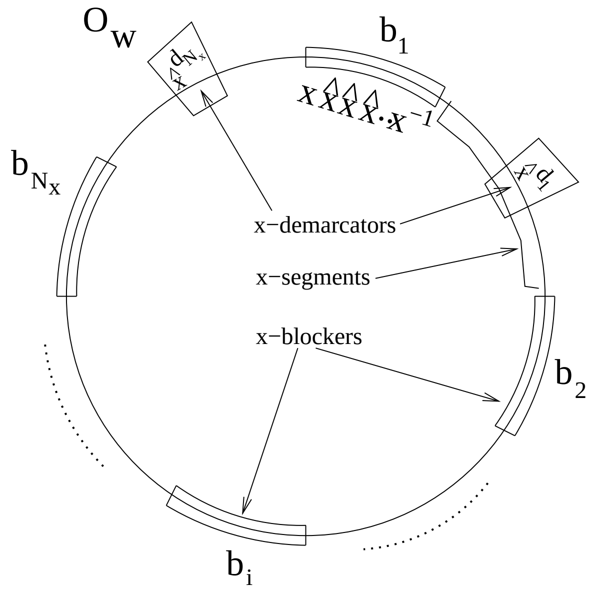

Figure 12: Demarcators between blockers in an x-decomposition of w.

Proof. If $N_{\hat{x}}(w)$ is 0 the statement is trivial. Otherwise, between every consecutive pair of $\hat{x}$-blockers $b_i^{\hat{x}}$ and $b_{i+1}^{\hat{x}}$ the intervening $\hat{x}$-segment $s_i^{\hat{x}}$ contains an $\hat{x}$-demarcator of length $d_i^{\hat{x}}$. Since $b_i^{\hat{x}}, d_i^{\hat{x}} \geqslant 1$, the lemma follows. □

Lemma 2.5.6 (Demarcator Lemma). *Given a conjugacy class $w \in \tilde{F}_2$, and fixed $x \in X \cup X^{-1}$,*

$$\text{\textcircled{$\sharp$}}(x, w) \quad \geqslant \quad \textstyle\sum_{k=1}^{|w|-3} (k+1) N_{\hat{x}}^k(w)$$

Proof. By Lemma 2.5.5 (pp. 22), $\text{\textcircled{$\sharp$}}(x, w) \geqslant \sum_{i=1}^{N_{\hat{x}}(w)} (b_i + 1)$. Regrouping this summation according to blocker weight k yields the summation $\sum_{k=1}^{\infty} (k+1) N_{\hat{x}}^k(w)$. Since all blockers have weight at most $|w| - 3$, the summation has finite support bounded by $k \leqslant |w| - 3$. □

For each positive $k_0 \in \mathbb{N}$, we have

Lemma 2.5.7 (Tail Lemma at k_0). *Given a conjugacy class $w \in \tilde{F}_2$, and fixed $x \in X \cup X^{-1}$,*

$$\text{\textcircled{$\sharp$}}(x, w) \quad \geqslant \quad k_0 \text{\textcircled{$\sharp$}}(x^{k_0} \hat{x}^{-1}, w) + \textstyle\sum_{k=1}^{k_0 - 1} (k+1) N_{\hat{x}}^k(w)$$

Proof. We associate with each occurrence of $x^{\pm}$ in w the maximal segment of the form $x^{\pm k}$ in which it lies. If we consider just occurrences of $x^{\pm}$ which lie inside $\hat{x}$-blockers of weight $k < k_0$ and the corresponding $\hat{x}$-demarcators, then by the Demarcator Lemma 2.5.6 (pp. 23) we have that

$$\text{\textcircled{$\sharp$}}(x, w) \quad \geqslant \quad \sum_{k=1}^{k_0 - 1} (k+1) N_{\hat{x}}^k(w) \tag{3}$$

This counting however, does not consider occurrences of $x^{\pm}$ which lie inside maximal segments of the form $x^{\pm k}$ for $k \geqslant k_0$. Of these, consider occurrences of x

which lie inside $x^k\hat{x}^{-1}$ and occurrences of x^{-1} which lie inside $\hat{x}x^{-k}$. Since $k \geqslant k_0$, $\hat{x}x^{-k}$ and $x^k\hat{x}^{-1}$ cannot overlap with an $\hat{x}$-blocker of weight $k < k_0$. By Lemma 2.5.3 (pp. 22), $\hat{x}x^{-k}$ and $x^k\hat{x}^{-1}$ cannot overlap with any $\hat{x}$-demarcators. Each occurrence of $\hat{x}x^{-k}$ and $x^k\hat{x}^{-1}$ thus contributes at least k_0 occurrences of $x^{\pm}$ that were not counted in expression (3). This completes the proof. $\square$

LEMMA 2.5.8 (Subword Lemma). *Given a conjugacy class $w \in \tilde{F_2}$, and fixed words g, h in F_2. If g is a graphic subword of h then*

$$\text{\textcircled{$\sharp$}}(g, w) \quad \geqslant \quad \text{\textcircled{$\sharp$}}(h, w)$$

PROOF. Every occurrence of $h^{\pm}$ contains an occurrence of $g^{\pm}$. $\square$

LEMMA 2.5.9 (Squeeze Lemma). *Given a conjugacy class $w \in \tilde{F_2}$, and fixed words g, h in F_2. If g is a graphic subword of h satisfying*

$$\text{\textcircled{$\sharp$}}(h, w) \quad \geqslant \quad \text{\textcircled{$\sharp$}}(g, w) + \sum_{i=1}^{k} \text{\textcircled{$\sharp$}}(f_i, w)$$

for some elements $f_1, \ldots, f_k$ in F_2, then

$$\text{\textcircled{$\sharp$}}(f_i, w) \quad = \quad 0$$

for $i = 1, \ldots, k$.

PROOF. By hypothesis, $\text{\textcircled{$\sharp$}}(h, w) \geqslant \text{\textcircled{$\sharp$}}(g, w)$. So by the Subword Lemma 2.5.8 (pp. 24), $\text{\textcircled{$\sharp$}}(h, w) = \text{\textcircled{$\sharp$}}(g, w)$. Since $\text{\textcircled{$\sharp$}}(f_i, w) \geqslant 0$ for $i = 1, \ldots, k$, it follows that $\text{\textcircled{$\sharp$}}(f_i, w)$ must all be identically 0. $\square$

The next two definitions are required in order to be able to express further inequalities satisfied by $\text{\textcircled{$\sharp$}}$.

DEFINITION 2.5.10. *Given two words g, h in F_2 we say that g and h* **overlap** *if*

- *There is a word f in F_2, with $|f| < |g| + |h|$, such that both h and g are subwords of f.*

We say that g and h **overlap properly** *if, in addition*

- *g is not a subword of h, and*
- *h is not a subword of g,*

DEFINITION 2.5.11. *Given a fixed word g in F_2 and $x \in X \cup X^{-1}$, we say that g is* **x-blocker-immune** *if neither g nor g^{-1} can properly overlap with an x-blocker. The word g is called* **x-demarcator-immune** *if neither g nor g^{-1} can properly overlap with an x-demarcator. Finally, the word g is called* **self-immune** *if both g and g^{-1} both cannot properly overlap with both g and g^{-1}.*

LEMMA 2.5.12 (Immunity Lemma). *Given a conjugacy class $w \in \tilde{F_2}$, $x \in X \cup X^{-1}$, and a fixed word g in F_2 that is self-immune and $\hat{x}$-demarcator-immune. Then*

$$\text{\textcircled{$\sharp$}}(x, w) \quad \geqslant \quad \sum_{k=1}^{|w|-3} (k+1) N_{\hat{x}}^{k}(w) + \left[\text{\textcircled{$\sharp$}}(x, g) \cdot \text{\textcircled{$\sharp$}}(g, \kappa_{\hat{x}}(w))\right].$$

PROOF. The number of occurrences of $x^{\pm}$ which occur either inside an $\hat{x}$-blocker or inside an $\hat{x}$-demarcator is $\sum_{k=1}^{|w|-3}(k+1)N_{\hat{x}}^{k}(w)$. Of the remaining occurrences of $x^{\pm}$, some number occur inside words of the form $g^{\pm}$. The total number of occurrences of g in w is $\textcircled{\sharp}(g,w)$, but of these we consider only a subset, namely the $\textcircled{\sharp}(g,\kappa_{\hat{x}}w)$ occurrences which do not overlap with any $\hat{x}$-blockers. Since g is $\hat{x}$-demarcator-immune, no occurences of g from this subset can overlap with any $\hat{x}$-demarcator. We have shown that there are $\textcircled{\sharp}(g,\kappa_{\hat{x}}w)$ occurrences of g which do not overlap with a $\hat{x}$-blocker or $\hat{x}$-demarcator. Since g is self-immune, each occurrence of $g^{\pm}$ contributes $\textcircled{\sharp}(x,g)$ distinct occurrences of $x^{\pm}$. This introduces an additional $\textcircled{\sharp}(x,g)\cdot\textcircled{\sharp}(g,\kappa_{\hat{x}}(w))$ ocurrences of $x^{\pm}$, yielding the expression in the statement of the lemma. □

LEMMA 2.5.13. *Given any conjugacy class $w \in \tilde{F_2}$, and $x \in X \cup X^{-1}$,*

$$\textcircled{\sharp}(\psi_{x^{-1}}^{-1}x, w) \quad \geqslant \quad N_x(w)$$

PROOF. Fix any x-demarcator d_i. Now consider the following two boundaries:

i. the boundary of d_i with the last occurence of x^{-1} in s_i (or the boundary of d_i with b_{i+1} if there are no occurences of x^{-1} in s_i).
ii. the boundary of d_i with the first occurrence of x in s_i (or the boundary of d_i with b_i if there are no occurences of x in s_i).

Then either [i] or [ii] must lie inside a word $\psi_{x^{-1}}^{-1}x^{\epsilon}$. □

The truncation and extension functions introduced next will permit us to express more inequalities satisfied by $\textcircled{\sharp}$.

DEFINITION 2.5.14. *The* **truncation functions** $t_l, t_r : F_2 \to \{a, b, A, B\}$ *as follows:*

$$\begin{aligned} t_l(w) &= \textit{the first symbol in } w, \\ t_r(w) &= \textit{the last symbol in } w. \end{aligned}$$

where w is assumed to be written in freely reduced form.

DEFINITION 2.5.15. *The* **extension functions** $e_l, e_r : F_2 \to (F_2)^3$ *are defined as:*

$$\begin{aligned} e_l(w) &= \{c \circ w \mid c \in X \cup X^{-1} \backslash \{t_l(w)^{-1}\}\}, \\ e_r(w) &= \{w \circ c \mid c \in X \cup X^{-1} \backslash \{t_r(w)^{-1}\}\}, \end{aligned}$$

where, for concreteness, the sets are ordered lexicographically.

LEMMA 2.5.16 (Extension Lemma). *Given a conjugacy class $w \in \tilde{F_2}$, $x \in X \cup X^{-1}$, and a fixed word g in F_2 for which $|w| > |g|$,*

$$\sharp(g,w) \quad = \quad \textstyle\sum_{p \in e_r(g)} \sharp(p,w) \quad = \quad \textstyle\sum_{q \in e_l(g)} \sharp(q,w)$$

PROOF. Since $|w| > |g|$, any occurrence of g in w occurs as a proper prefix of some word in $e_r(g)$ and as a suffix of some word in $e_l(g)$. The lemma is proved. □

LEMMA 2.5.17 (Symmetry Lemma). *Given a conjugacy class $w \in \tilde{F_2}$, $x \in X \cup X^{-1}$. Suppose that every occurrence of $x^{\pm}$ occurs inside an x-blocker. Then*

i. *Every occurrence of $x^{\pm}$ occurs inside an x^{-1}-blocker.*
ii. *Every occurrence of $\hat{x}^{\pm}$ occurs inside an x-blocker or an x-demarcator.*

PROOF. If every occurrence of $x^{\pm}$ occurs inside an x-blocker, then w is of the form

$$w \quad = \quad (x\hat{x}^{b_1}x^{-1})\hat{x}^{d_1} \cdots (x\hat{x}^{b_i}x^{-1})\hat{x}^{d_i} \cdots (x\hat{x}^{b_n}x^{-1})\hat{x}^{d_n}$$

So every occurrence of $\hat{x}^{\pm}$ occurs inside an x-blocker or an x-demarcator. Moreover, we can rearrange the parenthesization to emphasize that

$$w \quad = \quad (x^{-1}\hat{x}^{d_1}x)\cdots x^{-1}\hat{x}^{d_{i-1}}x)\hat{x}^{b_i}(x^{-1}\hat{x}^{d_i}x)\cdots \hat{x}^{b_n}(x^{-1}\hat{x}^{d_n}x)\hat{x}^{b_1},$$

which shows that in the x^{-1}-decomposition for w, every occurrence of $x^{\pm}$ occurs inside an x^{-1}-blocker. □

2.6. Asymptotic Behavior of Basic Shifts. The next proposition can be interpreted as stating that for any basic shift ψ_x and conjugacy class w, the length of $\psi_x^i w$ (as a function of i) has, in some sense, a first derivative which is non-decreasing:

PROPOSITION 2.6.1 (Non-Decreasing First Derivative). *For all $w \in \tilde{F}_2$ $\tilde{\psi} \in \tilde{\Psi}$, and $i \geqslant j > 0$*

$$|\tilde{\psi}^i w| - |\tilde{\psi}^{i-1}w| \geqslant |\tilde{\psi}^j w| - |\tilde{\psi}^{j-1}w|$$

PROOF. Suppose $\tilde{\psi} = \tilde{\psi}_x$, for some $x \in X \cup X^{-1}$. By Lemma 2.3.8 (pp. 18)

$$\begin{array}{rcl} |\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1}w| & = & \sharp(x,w) + \sharp(x^{-1},w) \\ & & -2[s(w,x,1,i-1) + \sum_{\epsilon=\pm1}\sharp(\psi_x^{-i}x^{\epsilon},w)] \\ |\tilde{\psi}_x^j w| - |\tilde{\psi}_x^{j-1}w| & = & \sharp(x,w) + \sharp(x^{-1},w) \\ & & -2[s(w,x,1,j-1) + \sum_{\epsilon=\pm1}\sharp(\psi_x^{-j}x^{\epsilon},w)] \end{array} \tag{4}$$

Now by hypothesis $j \leqslant i$. Then, since $\psi_x^{-j}x^{\epsilon}$ is a subword of $\psi_x^{-i}x^{\epsilon}$ (for $\epsilon = \pm 1$),

$$\sharp(\psi_x^{-i}x^{\epsilon},w) \quad \leqslant \quad \sharp(\psi_x^{-j}x^{\epsilon},w)$$

Consider the set of ψ_x-stable subwords of w having weight k, $1 \leqslant k \leqslant j-1$. It follows from Definition 2.3.6 (pp. 17) that this is a subset of the set of ψ_x-stable subwords of w having weight k', $1 \leqslant k' \leqslant i-1$. It follows immediately that

$$s(w,x,1,i-1) \quad \geqslant \quad s(w,x,1,j-1).$$

It is straightforward to verify that $\sum_{\epsilon=\pm1}\sharp(\psi_x^{-j}x^{\epsilon},w) - \sharp(\psi_x^{-i}x^{\epsilon},w)$ equals

$$\begin{array}{c} \sum_{k=j}^{i-1}\sharp(w,(\psi_x^{-k}x)x^{-1}) + \sharp(w,x(\psi_x^{-k}x^{-1})) + \\ \sharp(w,(\psi_x^{-k}x)x) + \sharp(w,x^{-1}(\psi_x^{-k}x^{-1})). \end{array} \tag{5}$$

Using Definition 2.3.6 (pp. 17) we compute $s(w,x,1,i-1) - s(w,x,1,j-1)$ to be

$$\sum_{k=j}^{i-1}\sharp(w,(\psi_x^{-k}x)x^{-1}) + \sharp(w,x(\psi_x^{-k}x^{-1})), \tag{6}$$

It follows immediately from expressions (5) and (6) above that

$$\sum_{\epsilon=\pm1}\sharp(\psi_x^{-j}x^{\epsilon},w) - \sharp(\psi_x^{-i}x^{\epsilon},w) \quad \geqslant \quad s(w,x,1,i-1) - s(w,x,1,j-1).$$

Finally, combining this inequality with expression (4), the proposition follows. □

Having shown that $|\tilde{\psi}^i w|$ is a function with a non-decreasing first derivative, we would like next to show that for i sufficiently small this derivative is non-positive, while for i sufficiently large, it is non-negative—this would permit us to conclude that $|\tilde{\psi}^i w|$ has a unique minimum value. Towards this end, we introduce

DEFINITION 2.6.2. *For each $x \in X \cup X^{-1}$, define $\rho_x : \tilde{F}_2 \to 2^{\mathbb{N}}$, so that*

$$\rho_x(w) \quad = \quad \{i \geqslant 0 \mid \sharp(\psi_x^{-i}x, w) + \sharp(\psi_x^{-i}x^{-1}, w) + s(w, x, i-1) \neq 0\}$$

for each $w \in \tilde{F}_2$. Since $|\psi_x^{-i}x| = i+1$ and all ψ_x-stable words having weight $i-1$ are of length $i+1$, it follows that $\rho_x(w) \subseteq \{0, 1, \ldots, |w|-1\}$.

The next lemma implies that the derivative of $|\tilde{\psi}^i w|$ has a zero, and hence that $|\tilde{\psi}^i w|$ attains a unique minimum value.

LEMMA 2.6.3 (Asymptote Lemma). *For all $w \in \tilde{F}_2$, and $x \in X \cup X^{-1}$, there exists a constant $c = c_x(w)$ no greater than $|w|$, such that for all $i \in \mathbb{Z}$*

$$\begin{aligned} [\rho_x(w) = \emptyset \ \vee \ i > \max\rho_x(w)] \quad &\Rightarrow \quad |\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| = c \\ [\rho_{x^{-1}}(w) = \emptyset \ \vee \ j < -\max\rho_{x^{-1}}(w)] \quad &\Rightarrow \quad |\tilde{\psi}_x^{j+1} w| - |\tilde{\psi}_x^{j} w| = -c \end{aligned}$$

PROOF. For $w = 1$, the lemma is trivially true, with $c = 0$. So assume $|w| \geqslant 1$. We begin by proving the first implication. By Lemma 2.3.8 (pp. 18),

$$\begin{aligned} |\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| \quad = \quad & \sharp(x, w) + \sharp(x^{-1}, w) \\ & -2[s(w, x, 1, i-1) + \sum_{\epsilon = \pm 1} \sharp(\psi_x^{-i} x^{\epsilon}, w)]. \end{aligned}$$

When $i > \max\rho_x(w)$, the definition of ρ_x implies that $\sharp(\psi_x^{-i}x^{\pm 1}, w) = 0$, and hence

(7) $\quad |\tilde{\psi}_x^i w| - |\tilde{\psi}_x^{i-1} w| = \sharp(x, w) + \sharp(x^{-1}, w) - 2s(w, x, 1, i-1) \overset{def}{=} c_x(w).$

Now, by definition

$$s(w, x, 1, i-1) = \sum_{k=1}^{i-1} s(w, x, k)$$

and since $i > \max\rho_x(w)$, it follows from the definition of ρ_x that $s(w, x, k) = 0$ for all $k > \max\rho_x(w)$. Thus

$$s(w, x, 1, i-1) = \sum_{k=1}^{\max\rho_x(w)} s(w, x, k) = \sum_{k=1}^{|w|-1} s(w, x, k)$$

We have shown that if $i > \max\rho_x(w)$, then in fact $c_x(w)$ as defined in (7) above, depends only on w and is independent of i. This proves the first implication.

To prove the second implication, we apply the previous argument to $\tilde{\psi}_{x^{-1}}$, and conclude that if $j' > \max\rho_{x^{-1}}(w)$ then $\sharp(\psi_{x^{-1}}^{-j'} x^{\epsilon}, w) = 0$, and hence

$$|\tilde{\psi}_{x^{-1}}^{j'} w| - |\tilde{\psi}_{x^{-1}}^{j'-1} w| = \sharp(x^{-1}, w) + \sharp(x, w) - 2s(w, x, 1, j'-1) = c_x(w).$$

By an analogous argument to the one above one sees that $s(w, x, 1, j'-1)$ is independent of j', provided $j' > \max\rho_{x^{-1}}(w)$. By Remark 1.1.4 (pp. 1.1.4), $\tilde{\psi}_{x^{-1}}^{j'} \equiv \tilde{\psi}_x^{-j'}$, so substituting $j = -j'$, we conclude that when $j < -\max\rho_{x^{-1}}(w)$,

$$|\tilde{\psi}_x^{j+1} w| - |\tilde{\psi}_x^{j} w| = -c_x(w)$$

which proves the second implication of the lemma. □

The previous lemmas show that derivative of $|\tilde{\psi}_x^i w|$ is non-decreasing, and its value becomes a non-positive constant for i sufficiently small and a non-negative constant for i sufficiently large; hence $|\tilde{\psi}_x^i w|$ has a unique minimum value. These conclusions are summarized visually by Figure (13), and were already reported, in some form, by Whitehead in [27, 28]. We now continue the analysis further.

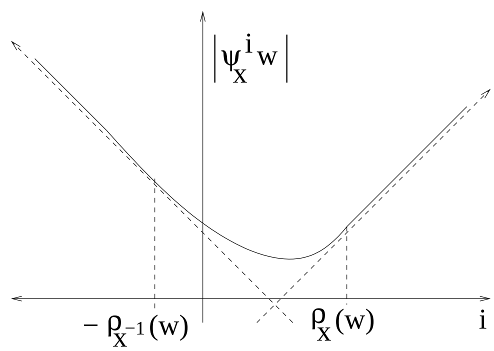

Figure 13: How $|\psi_x^i w|$ changes with i.

3. The Structure Within Levels

In remainder of this section, we will investigate the structure of level neighborhoods $B(u)$ (resp. $B^*(u)$) in Ω (resp. Ω^*). Most of our results will be built from Lemma 2.3.8 (pp. 18) and Proposition 2.6.1 (pp. 26). We will describe the combinatorial conditions on a conjugacy class $\tilde{w}$ under are necessary in order for $|\tilde{\psi}\tilde{w}| = |\tilde{w}|$ for special classes of automorphisms ψ.

3.1. Level x^k-Chains. In the previous section, Lemma 2.6.3 (pp. 27) showed that $|\tilde{\psi}^i w|$ attains a unique minimum value as i is varied. However, this minimum value need not be attained at a unique value of i; rather, the minimum value may be achieved over a long contiguous segment of integer values. This leads us to

DEFINITION 3.1.1. *For each $x \in X \cup X^{-1}$, $k \in \mathbb{N}$, If a conjugacy class $w \in \tilde{F}_2$ satisfies*

$$\forall i, j \in \{0, 1, \ldots, k\} \text{ distinct}, \quad \tilde{\psi}_x^i w \neq \tilde{\psi}_x^j w$$
$$\forall i \in \{0, 1, \ldots, k\}, \quad |\tilde{\psi}_x^i w| = |w|$$

Then w is called a **level $\mathbf{x^k}$-chain**. *The set of all x^k-chains in $\tilde{F}_2$ as $\mathcal{C}_{x^k}$.*

A level x^k-chain is thus a conjugacy class which lies at the start of a (non self-intersecting) path of length k in the graph Ω, whose vertices are distinct conjugacy classes that are all of the same length, and that can be traversed by k applications of $\tilde{\psi}_x$ (see Figure 14).

EXAMPLE 3.1.2. *Recalling Example 2.1.4 (pp. 11), we note that the conjugacy class $w = [aaBB]$ is a level a^2-chain, since $\tilde{\psi}_a(w) = abaB$ and $\tilde{\psi}_a^2(w) = abba$ are all of length 4 and are distinct conjugacy classes.*

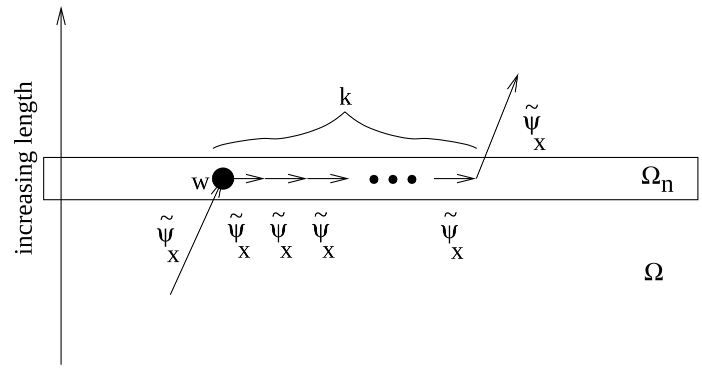

Figure 14: A level x^k-chain, w

Note that the previous definition does not preclude the possibility that a level x^k-chain may pass through distinct elements of $\tilde{F}_2$ that are related by permutation maps. Thus while a level x^k-chain is guaranteed to not self-intersect in the graph Ω, it may induce a self-intersecting trail in quotient Ω^*. This leads us to

DEFINITION 3.1.3. *A level x^k-chain w is called* **simple** *if it additionally satisfies*

$$\forall i, j \in \{0, \ldots, k\} \text{ distinct}, \quad \tilde{\psi}_x^i w \not\approx_\Pi \tilde{\psi}_x^j w$$

We denote the set of all simple x^k-chains in $\tilde{F}_2$ as $\mathcal{S}_{x^k} \subset \mathcal{C}_{x^k}$.

A simple level x^k-chain is thus an x^k-chain which remains non-self intersecting when it is projected into Ω^* via $\approx_\Pi$ (see Figure 15).

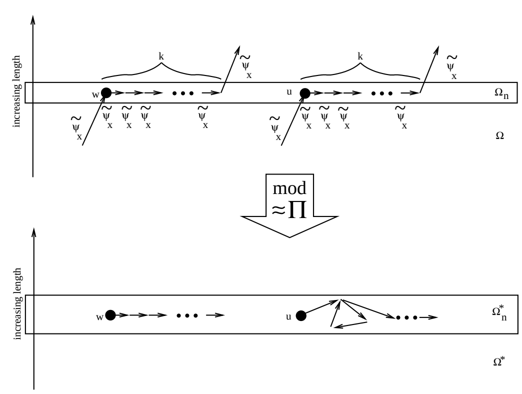

Figure 15: Two level x^k-chains: w is simple, while u is not.

EXAMPLE 3.1.4. *Recalling Example 2.1.8 (pp. 12), we note that the conjugacy class $w = [aaBB]$ is* not *a level* simple *a^2-chain, since $\tilde{\psi}_a^2(w) = abba$ is permutation-equivalent to the conjugacy class w. However, $w_1 = [aaBBBB]$ is a level* simple *a^2-chain, since $\tilde{\psi}_a(w_1) = abaBBB$ and $\tilde{\psi}_a^2(w_1) = abbaBB$ are all of length* 6 *and are distinct conjugacy classes, no two of which are related by a permutation map.*

3.2. Large-scale Obstructions. In this section we begin by describing a system of combinatorial equations which captures the constraints on the internal structure of conjugacy classes that are level x^k-chains—recall that we view a conjugacy class w in terms of the corresponding cycle graph O_w.

LEMMA 3.2.1. *Given any $w \in \tilde{F}_2$, $x \in X \cup X^{-1}$, $\tilde{\psi}_x \in \tilde{\Psi}$, and $k \in \mathbb{N}$,*

$$\begin{aligned} w \in \mathcal{C}_{x^k} \quad \Leftrightarrow \quad & \tfrac{1}{2}\textcircled{\sharp}(x, \kappa_x(w)) = \\ & \textcircled{\sharp}(\psi_x^{-1}x, \kappa_x(w)) = \\ & \textcircled{\sharp}(\psi_x^{-k}x, \kappa_x(w)) \end{aligned}$$

PROOF. By Proposition 2.6.1 (pp. 26) we know that

$$(\forall i = 1, \dots, k) \ \ |\tilde{\psi}_x^i w| = |w| \quad \Leftrightarrow \quad |\tilde{\psi}_x w| - |w| = 0 \text{ and } |\tilde{\psi}_x^k w| - |\tilde{\psi}_x^{k-1} w| = 0$$

By Lemma 2.3.8 (pp. 18)

$$\begin{aligned} |\tilde{\psi}_x w| - |w| = 0 \quad &\Leftrightarrow \quad \textcircled{\sharp}(x, \kappa_x(w)) - 2\textcircled{\sharp}(\psi_x^{-1}x, \kappa_x(w)) = 0 \\ |\tilde{\psi}_x^k w| - |\tilde{\psi}_x^{k-1} w| = 0 \quad &\Leftrightarrow \quad \textcircled{\sharp}(x, \kappa_x(w)) - 2\textcircled{\sharp}(\psi_x^{-k}x, \kappa_x(w)) = 0. \end{aligned}$$

Combining the previous two equations completes the proof of the lemma. □

The previous lemma describes the combinatorial equations $\Lambda^x_{G,w}$ which must be satisfied by $\kappa_x(w)$ if w is indeed an x^k-chain. The results of the lemma are leveraged in the subsequent theorem to obtain an upper bound on the length of (both simple and non-simple) level x^k-chains that may occur in Ω.

In essence then, the Theorems 3.2.2 (pp. 30) and 3.2.4 (pp. 31) below gives a characterization of certain forbidden graphs which are never realized in Ω_n and Ω_n^*. Figure (16) depicts these forbidden graphs. Because the sizes of these graphs is a not constant with respect to n, we refer to them as **large-scale obstructions**.

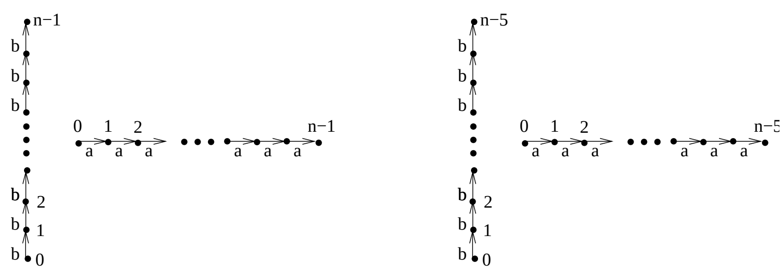

Figure 16: Large-scale obstructions: two graphs that cannot be realized in Ω_n (left), and two graphs that cannot be realized in Ω_n^* (right).

THEOREM 3.2.2. *For any $w \in \tilde{F}_2$ and $x \in X^\delta$, $\delta \in \{-1, +1\}$*

$$\textit{(I)} \qquad |w| \geqslant 2, \ \textit{and } w \in \mathcal{C}_{x^k} \quad \Rightarrow \quad k \leqslant |\kappa_x(w)| - 2$$

PROOF. Part (I). By Lemma 3.2.1 (pp. 30), we know that

$$(8) \qquad (\forall i = 1, \dots, k) \ \ |\tilde{\psi}_x^i w| = |w| \quad \Leftrightarrow \quad \begin{aligned} & \tfrac{1}{2}\textcircled{\sharp}(x, \kappa_x(w)) = \\ & \textcircled{\sharp}(\psi_x^{-1}x, \kappa_x(w)) = \\ & \textcircled{\sharp}(\psi_x^{-k}x, \kappa_x(w)) \end{aligned}$$

If $k = 0$, then the lemma is trivially true since $|w| \geqslant 2$. If $k > 0$, then $w \in \mathcal{C}_{x^k}$ implies $\psi_x w \neq w$, and hence $\text{\textcircled{$\sharp$}}(x, \kappa_x(w)) \neq 0$. From this integer inequality it follows that

$$\text{\textcircled{$\sharp$}}(x, \kappa_x(w)) \quad \geqslant \quad 2$$

and thus, by appealing to (8), we conclude that

$$\gamma \stackrel{def}{=} \text{\textcircled{$\sharp$}}(\psi_x^{-k} x, \kappa_x(w)) \quad \geqslant \quad 1.$$

We have shown that the total number of x, x^{-1} in w is 2γ. Of these, the number of x, x^{-1} which fall inside a subword of the form $\psi_x^{-k} x$ or $\psi_x^{-k} x^{-1}$ can be no more than γ, since

$$\begin{aligned} \sharp(x, \tilde{\psi}_x^{-k} x) = \sharp(x^{-1}, \tilde{\psi}_x^{-k} x^{-1}) &\quad = \quad 1 \\ \sharp(x^{-1}, \tilde{\psi}_x^{-k} x) = \sharp(x, \tilde{\psi}_x^{-k} x^{-1}) &\quad = \quad 0. \end{aligned}$$

Then, since $|\psi_x^{-k} x| = |\psi_x^{-k} x^{-1}| = k + 1$, it follows that

$$|\kappa_x(w)| \geqslant \gamma(k+1) + \gamma = \gamma(k+2) \tag{9}$$

Suppose (toward contradiction) that $k > |\kappa_x(w)| - 2$, then (9) would imply $\gamma < 1$. Then it must be that $\gamma = 0$, and hence $\frac{1}{2}[\text{\textcircled{$\sharp$}}(x, \kappa_x(w)) = \gamma = 0$. But then $\psi_x\, w = w$, which is a contradiction. So it must be that $k \leqslant |\kappa_x(w)| - 2$. We have shown that $w \in \mathcal{C}_{x^k} \Rightarrow k \leqslant |\kappa_x(w)| - 2$. □

The following classification of edge pairs in O_w will be used later.

DEFINITION 3.2.3. *Given a conjugacy class $w \in \tilde{F}_2$ and two distinct edges e_1 and e_2 of O_w that do not lie inside any x-blocker. Then e_1 and e_2 are precisely one of the following relationships (see Figure 17).*

1. *UC_x (Unified Circular with respect to x) if $N_x(w) = 0$.*
2. *UL_x (Unified Linear with respect to x) if $N_x(w) > 0$ and both e_1, e_2 lie in some x-segment s_i.*
3. *DL_x (Disjoint Linear with respect to x) if $N_x(w) > 0$ and both e_1, e_2 lie in distinct x-segments s_i, s_j, $i \neq j$.*

THEOREM 3.2.4. *For any $w \in \tilde{F}_2$ and $x \in X^\delta$, $\delta \in \{-1, +1\}$*

$$\begin{array}{lllll} (I^*) & |w| \geqslant 2, \text{ and } w \in \mathcal{C}_{x^k} & \Rightarrow & k \leqslant |w| - 2 \\ (II^*) & |w| \geqslant 10, \text{ and } w \in \mathcal{S}_{x^k} & \Rightarrow & k \leqslant |w| - 6 \end{array}$$

PROOF. Part (I*). By Theorem 3.2.2 (pp. 30), we know that $k \leqslant |\kappa_x(w)| - 2$. Since $|\kappa_x(w)| \leqslant |w|$, the assertion follows.

Part (II*). Since $\mathcal{S}_{x^k} \subset \mathcal{S}_{x^{k+1}}$ it suffices to show that $w \in \mathcal{S}_{x^{|w|-5}}$ and $|w| \geqslant 10$ leads to a contradiction.

If $k = |w| - 5$ and $|w| \geqslant 10$, then equation (9) imples

$$\gamma \leqslant \frac{|\kappa_x(w)|}{|w| - 3} \leqslant \frac{|w|}{|w| - 3} \leqslant \frac{10}{7} < 2.$$

So $\gamma \in \{0, 1\}$. By the previous argument, γ cannot be 0.

In the case $\gamma = 1$, we know the following about the structure of w:

(A) There are $2\gamma = 2$ occurrences of x, x^{-1} in $\kappa_x(w)$.

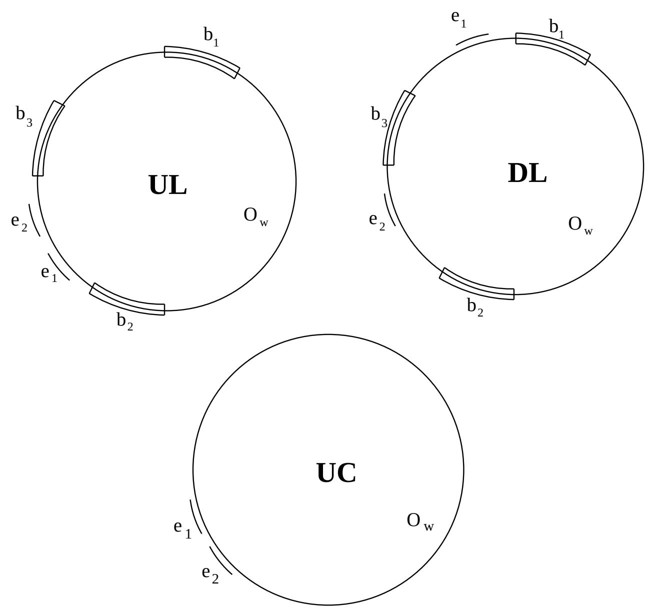

Figure 17: Edges e_1 and e_2 in $UL_x(w_1)$, in $DL_x(w_1)$, and $UC_x(w_2)$ relationships, respectively.

Now there are three possible configurations in which these occurrences of $x^{\pm}$ can occur inside $\kappa_x(w)$, as per Definition 3.2.3 (pp. 31). Each of these configurations are treated in turn; the Unified Circular case starting at pp. 32, the Unified Linear case starting at pp. 33, and the Disjoint Linear case starting at pp. 35.

The Unified Circular Case:

In this case, since there are no x-blockers in w, $\kappa_x(w) = w$. Now since $\gamma = 1$, we note that

(A*) Now if there is one occurrence of x and one occurrence of x^{-1} then there is an x-blocker, a contradiction. So we know that either there are two occurrences of x, or there are two occurrences of x^{-1} in $\kappa_x(w)$.

(B) There are $\gamma = 1$ occurrences of $\psi_x^{-k}x^{\epsilon}$, and $\gamma = 1$ occurrences of $\psi_x^{-1}x^{\epsilon}$, for some $x \in X^{\delta}$ and $\epsilon, \delta \in \{+1, -1\}$. Since $\psi_x^{-1}x^{\epsilon}$ is a subword of $\psi_x^{-k}x^{\epsilon}$, there are no occurrences of $\psi_x^{-1}x^{\epsilon}$ except the one that occurs as a subword of the occurrence of $\psi_x^{-k}x^{\epsilon} = (x\hat{x}^{-k\delta})^{\epsilon}$.

The facts (A) and (B) force the cyclic word w be of one of the following types:

$$\begin{array}{llcl}
\text{Type 1}: & x \in X & w & = & x\hat{x}^{-k} \cdot \hat{x}^{-p} \cdot x \cdot \hat{x}^{q} \\
\text{Type 2}: & x \in X^{-1} & w & = & x\hat{x}^{k} \cdot \hat{x}^{p} \cdot x \cdot \hat{x}^{-q} \\
\text{Type 3}: & x \in X & w & = & \hat{x}^{k}x^{-1} \cdot \hat{x}^{-q} \cdot x^{-1} \cdot \hat{x}^{p} \\
\text{Type 4}: & x \in X^{-1} & w & = & \hat{x}^{-k}x^{-1} \cdot \hat{x}^{q} \cdot x^{-1} \cdot \hat{x}^{-p}
\end{array}$$

where $p, q \in \mathbb{N}$ satisfy

$$(k+1) + p + 1 + q = |w|. \tag{10}$$

Note that since $k = |w| - 5$, equation (10) implies that $p + q = 3$.

Let us define

$$i = \begin{cases} \frac{k+p-q}{2} - 1 & \text{if } 2 \mid (k+p-q) \\ \left\lfloor \frac{k+p-q}{2} \right\rfloor & \text{otherwise.} \end{cases}$$

$$j = \begin{cases} \frac{k+p-q}{2} + 1 & \text{if } 2 \mid (k+p-q) \\ \left\lceil \frac{k+p-q}{2} \right\rceil & \text{otherwise.} \end{cases}$$

Clearly i and j are distinct. To show that i and j in $\{0, \ldots, k\}$ we note that since $p+q=3$, it follows that $-3 \leqslant p-q \leqslant 3$. From this we may conclude

$$\begin{array}{rcccccl} \frac{k+p-q}{2} - 1 & \geqslant & \frac{k-3}{2} - 1 & = & \frac{|w|}{2} - 5 & \geqslant & 0 \\ \frac{k+p-q}{2} + 1 & \leqslant & \frac{k+3}{2} + 1 & = & \frac{|w|}{2} & \leqslant & k, \end{array}$$

where the last two inequalities follow from our assumption that $k+5 = |w| \geqslant 10$. We will now show that i and j satisfy $\tilde{\psi}_x^i w \approx_\Pi \tilde{\psi}_x^j w$, thereby contradicting the fact that $w \in \mathcal{S}_{x^k}$.

We begin by considering the case when w is Type (1). If $2 \mid (k+p-q)$ then

$$\begin{aligned} \tilde{\psi}_x^i w &= x\hat{x}^{-k-p+i} x\hat{x}^{q+i} \\ &= x\hat{x}^{-(k+p-\frac{k+p-q}{2}+1)} x\hat{x}^{(q+\frac{k+p-q}{2}-1)} \\ &= x\hat{x}^{-(\frac{k+p+q}{2}+1)} x\hat{x}^{(\frac{k+p+q}{2}-1)} \end{aligned}$$

$$\begin{aligned} \tilde{\psi}_x^j w &= x\hat{x}^{-k-p+j} x\hat{x}^{q+j} \\ &= x\hat{x}^{-(k+p-\frac{k+p-q}{2}-1)} x\hat{x}^{(q+\frac{k+p-q}{2}+1)} \\ &= x\hat{x}^{-(\frac{k+p+q}{2}-1)} x\hat{x}^{(\frac{k+p+q}{2}+1)}. \end{aligned}$$

Otherwise, $2 \not\mid (k+p-q)$ and

$$\begin{aligned} \tilde{\psi}_x^i w &= x\hat{x}^{-k-p+i} x\hat{x}^{q+i} \\ &= x\hat{x}^{-(k+p-\lfloor\frac{k+p-q}{2}\rfloor)} x\hat{x}^{(q+\lfloor\frac{k+p-q}{2}\rfloor)} \\ &= x\hat{x}^{-(\frac{k+p+q}{2}+\frac{1}{2})} x\hat{x}^{(\frac{k+p+q}{2}-\frac{1}{2})} \end{aligned}$$

$$\begin{aligned} \tilde{\psi}_x^j w &= x\hat{x}^{(-k-p+j} x\hat{x}^{q+j} \\ &= x\hat{x}^{-(k+p-\lceil\frac{k+p-q}{2}\rceil)} x\hat{x}^{(q+\lceil\frac{k+p-q}{2}\rceil)} \\ &= x\hat{x}^{-(\frac{k+p+q}{2}-\frac{1}{2})} x\hat{x}^{(\frac{k+p+q}{2}+\frac{1}{2})}. \end{aligned}$$

In both these situations, the length-preserving automorphism $\tilde{\pi}_{\hat{x}} \in \tilde{\Pi}$ satisfies

$$\tilde{\pi}_{\hat{x}}(\tilde{\psi}_x^i w) = \tilde{\psi}_x^j w$$

so $\tilde{\psi}_x^i w \approx_\Pi \tilde{\psi}_x^j w$, and hence $w \notin \mathcal{S}_{x^{|w|-5}}$, a contradiction. Thus, for words of Type (1) of length $\geqslant 10$, $w \in \mathcal{S}_{x^k}$ implies $k \leqslant |w| - 6$.

The proofs for words of Type (2), (3) and (4) are analogous, using the same definitions for i and j.

Thus if $w \in \tilde{F}_2$ is of length $\geqslant 10$ and $w \in \mathcal{S}_{x^k}$, then $k \leqslant |w| - 6$.

The Unified Linear Case:

In the unified linear case, the $2\gamma = 2$ occurrences of $x^\pm$ occur inside the same x-segment. For concreteness, suppose they lie in the ith x-segment, i.e. between x-blockers $b_i(w)$ and $b_{i+1}(w)$.

It follows then from the Demarcator Structure Lemma 2.5.3 (pp. 22), that either

- a_1 Two occurrences of x lie on the same side of the x-demarcator within the x-segment.
- a_1 Two occurrences of x^{-1} lie on the same side of the x-demarcator within the x-segment.
- b. One occurrence of x^{-1} and one occurrence of x lie on opposite sides of the x-demarcator within the x-segment.

The possible configurations (a_1), (a_2) and (b) are shown in Figure 3.2.

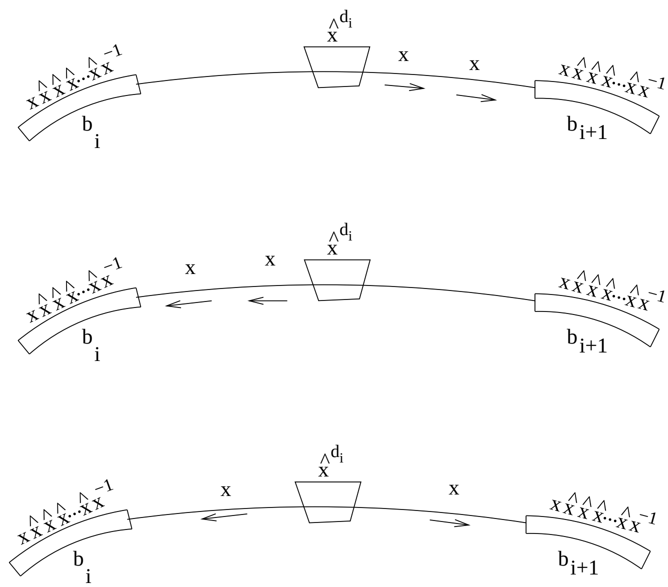

Figure 18: The unified linear case: Configurations a_1, a_1, and b.

Consider configuration (a_1). Then two occurrences of x must occur in a subword of the segment of the form $x\hat{x}^{-\lambda_1}x\hat{x}^{-\lambda_2}$ where $|\lambda_2|$ is maximal. First, note that λ_1, λ_2 must be non-negative, since otherwise $|\tilde{\psi}_x w| > |w|$. Now $|w| \geqslant 2 + \lambda_1 + \lambda_2 + (b_i + 2) + d_i$, and $b_i, d_i \geqslant 1$. It follows that $\lambda_1 + \lambda_2 \leqslant |w| - 6$. Hence λ_1 and λ_2 are each $\leqslant |w| - 6$. It follows that $|\tilde{\psi}_x^{|w|-5} w| > |w|$. In other words, $w \notin \mathcal{C}_{x^{|w|-5}}$, and hence $w \notin \mathcal{S}_{x^{|w|-5}}$.

Consider configuration (a_2). Then two occurrences of x^{-1} must occur in a subword of the segment of the form $\hat{x}^{\lambda_1}x^{-1}\hat{x}^{\lambda_2}x^{-1}$ where $|\lambda_1|$ is maximal. First, note that λ_1, λ_2 must be non-negative, since otherwise $|\tilde{\psi}_x w| > |w|$. Now $|w| \geqslant 2 + \lambda_1 + \lambda_2 + (b_i + 2) + d_i$, and $b_i, d_i \geqslant 1$. It follows that $\lambda_1 + \lambda_2 \leqslant |w| - 6$. Hence λ_1 and λ_2 are each $\leqslant |w| - 6$. It follows that $|\tilde{\psi}_x^{|w|-5} w| > |w|$. In other words, $w \notin \mathcal{C}_{x^{|w|-5}}$, and hence $w \notin \mathcal{S}_{x^{|w|-5}}$.

Consider configuration (b). Then the Demarcator Structure Lemma 2.5.3 (pp. 22) implies that the occurrence of x^{-1} precedes the demarcator and the occurrence of x succeeds the demarcator. It follows that x^{-1} must occur in a subword of the segment having the form $\hat{x}^{\lambda_1}x^{-1}$ and x must occur in a subword of the segment having the form $x\hat{x}^{-\lambda_2}$, where $|\lambda_1|$ and $|\lambda_2|$ are maximal. First, note that λ_1, λ_2

must be non-negative, since otherwise $|\tilde{\psi}_x w| > |w|$. Now $|w| \geqslant 2 + \lambda_1 + \lambda_2 + (b_i + 2) + d_i$, and $b_i, d_i \geqslant 1$. It follows that $\lambda_1 + \lambda_2 \leqslant |w| - 6$. Hence λ_1 and λ_2 are each $\leqslant |w| - 6$. It follows that $|\tilde{\psi}_x^{|w|-5} w| > |w|$. In other words, $w \notin \mathcal{C}_{x^{|w|-5}}$, and hence $w \notin \mathcal{S}_{x^{|w|-5}}$.

The Disjoint Linear Case:

In the disjoint linear case, the $2\gamma = 2$ occurrences of $x^{\pm}$ occur inside the different x-segments. For concreteness, suppose one lies in the ith x-segment, i.e. between x-blockers $b_i(w)$ and $b_{i+1}(w)$, while the second lies in the jth x-segment, i.e. between x-blockers $b_j(w)$ and $b_{j+1}(w)$. The occurrences of $x^{\pm}$ must be opposed by $\hat{x}^{\lambda_1}$ and $\hat{x}^{\lambda_2}$, where $|\lambda_1|, |\lambda_2|$ are maximal. Now $|w| \geqslant 2 + \lambda_1 + \lambda_2 + (b_i + 2) + d_i + (b_j + 2) + d_j$, and $b_i, d_i, b_j, d_j \geqslant 1$. It follows that $\lambda_1 + \lambda_2 \leqslant |w| - 10$. Hence λ_1 and λ_2 are each $\leqslant |w| - 10$. It follows that $|\tilde{\psi}_x^{|w|-5} w| > |w|$. But if $w \notin \mathcal{C}_{x^{|w|-5}}$, then $w \notin \mathcal{S}_{x^{|w|-5}}$.

Having considered the Unified Circular case the Unified Linear case, and the Disjoint Linear case, we conclude that always $|w| \geqslant 10$ implies $w \notin \mathcal{S}_{x^{|w|-5}}$. □

To see that the bound on the length of (not necessarily simple) x^k-chains given in Part (I) of Theorem 3.2.4 (pp. 31) cannot be improved, consider the conjugacy class w of the word $a^2 B^k$ in F, and take $\tilde{\psi} = \tilde{\psi}_a$. Then $\tilde{\psi}_a^i w = a^{2+i} B^{k-i}$ are distinct for all $i = 1, \ldots, k$ and have the same length as w. The next proposition shows that the bound on the length of simple x^k-chains given in Part (II) of Theorem 3.2.4 (pp. 31) also cannot be improved.

PROPOSITION 3.2.5. *For all $k \geqslant 10$ and $x \in X \cup X^{-1}$*

$$\mathcal{S}_{x^{k-6}} \neq \emptyset$$

PROOF. Let $w = a^n bABAbb \in \tilde{F}_2$, with $n = k - 6$. We show that $w \in \mathcal{S}_{B^{k-6}}$.

Clearly, $\tilde{\psi}_B^i w = a^{n-i} bABA^{1+i} bb$. So for $i = 0, \ldots, |w| - 5$, the elements $\tilde{\psi}_B^i w$ are all distinct, having length equal to $|w|$.

Suppose (towards contradiction) that there are distinct $i, j \in \{0, 1, \ldots, |w| - 5\}$ such that $\pi a^{n-i} bABA^{1+i} bb = a^{n-j} bABA^{1+j} bb$ for some $\pi \in \Pi$. Since $i \neq j$, π is non-trivial. Since $n \geqslant 4$, for any i either a^{n-i} or A^{1+i} (or both) is of length > 2. Since a^{n-j} and A^{1+j} are the only two uniformly labelled subsegments whose length can exceed 2, it follows that $\pi : a \mapsto A$. But then, since $\sharp(b, w) \neq \sharp(B, w)$, it follows that $\pi : b \mapsto b$ and $\pi : B \mapsto B$. Hence, $\pi = \pi_a$. But $\pi_a(a^{n-i} bABA^{1+i} bb) = A^{n-i} baBa^{1+i} bb$ is *not* conjugate to $a^{n-j} bABA^{1+j} bb$, since the latter contains a subword ABA while the former does not. We have arrived at a contradiction.

Thus, conjugacy class $w = a^n bABAbb$ $(n \geqslant 4)$ is a witness to the existence of B^k-chains that precisely meet the bound of Theorem 3.2.4 (pp. 31). We remark that examples of such maximal-length x^k-chains can be easily constructed for all $x \in X \cup X^{-1}$. □

3.3. Level σ-Chains. In section 3.1, we considered basic shift maps $\tilde{\psi}_x$ for $x \in X \cup X^{-1}$. These gave rise to the notion of level x^k-chains and simple level x^k-chains in Ω. Now, to deal with small-scale obstructions (i.e. obstruction whose size is a constant independent of n) in a similar manner, we generalize x^k-chains as follows.

DEFINITION 3.3.1. *Let $\sigma = \sigma_{|\sigma|} \sigma_{|\sigma|-1} \cdots \sigma_2 \sigma_1$ be a freely reduced word in F_2, where , $\sigma_i \in X \cup X^{-1}$ for $i = 1, \ldots, n$. The composite shift automorphism ψ_σ is* *defined as*

$$\psi_\sigma \stackrel{def}{=} \psi_{\sigma_{|\sigma|}} \cdots \psi_{\sigma_2}\psi_{\sigma_1} : F_2 \to F_2$$

As usual, $\bar{\psi}_\sigma$ *is defined to be the induced map on the conjugacy classes* $\tilde{F}_2$.

DEFINITION 3.3.2. *Let* $\sigma \in F_2$ *be a freely reduced word. A* **level σ-chain** *is a conjugacy class* $w \in \tilde{F}_2$ *satisfying*

$$\forall i, j \in \{1, \ldots, |\sigma|\} \text{ distinct}, \;\; \tilde{\psi}_{\sigma_i \cdots \sigma_1} w \neq \tilde{\psi}_{\sigma_j \cdots \sigma_1} w \neq w$$
$$\forall i \in \{0, 1, \ldots, k\}, \;\; |\tilde{\psi}_{\sigma_i \cdots \sigma_1} w| = |w|$$

We denote the set of all σ*-chains in* $\tilde{F}_2$ *as* $\mathcal{C}_\sigma$.

DEFINITION 3.3.3. *A* σ*-chain is called* **simple** *if it additionally satisfies*

$$\forall i, j \in \{1, \ldots, k\} \text{ distinct}, \;\; \tilde{\psi}_{\sigma_i \cdots \sigma_1} w \not\approx_\Pi \tilde{\psi}_{\sigma_j \cdots \sigma_1} w \not\approx_\Pi w$$

and the set of all simple σ*-chains in* $\tilde{F}_2$ *is denoted* $\mathcal{S}_\sigma$.

DEFINITION 3.3.4. *Let* $\sigma \in F_2$ *be a freely reduced word. We say that* σ *is* **realized as a simple level chain** *in* Ω^* *if there exists a conjugacy class* $w \in \mathcal{S}_\sigma$.

Let $\sigma \in F_2$ *be a freely reduced word. We say that* σ *is realized as a* **proper** *simple level chain in* Ω^* *if there exists word* $\sigma' \in F_2$ *which properly contains* σ *as a subword and a conjugacy class* $w \in \mathcal{S}_{\sigma'}$.

LEMMA 3.3.5 (Chain Inversion Lemma). *Let* $\sigma \in F_2$ *be a freely reduced word.*

$$\mathcal{S}_\sigma = \emptyset \quad \Leftrightarrow \quad \mathcal{S}_{\sigma^{-1}} = \emptyset.$$

PROOF. Suppose $w \in \mathcal{S}_\sigma$. Then let $w' = \tilde{\psi}_\sigma w$. Then $w' \in \mathcal{S}_{\sigma^{-1}}$. The argument for the reverse implication is identical since inversion is idempotent. □

LEMMA 3.3.6 (Alphabet Symmetry Lemma). *Let* $\sigma \in F_2$ *be a freely reduced word.*

$$\mathcal{S}_\sigma = \emptyset \quad \Leftrightarrow \quad \mathcal{S}_{\hat{\sigma}} = \emptyset.$$

PROOF. Fix P to be a proof that $w \in \mathcal{S}_\sigma$. Let $\hat{P}$, the formal object obtained from P by changing all a symbols to b, all b symbols to a, all A symbols to B, and all B symbols to A. The $\hat{P}$ is a proof that $\hat{w} \in \mathcal{S}_{\hat{\sigma}}$. The argument for the reverse implication is identical since $\hat{\pi}$ is idempotent. □

3.4. Small-scale Obstructions. In contrast to Theorem 3.2.4 (pp. 31) of the previous section, here we describe forbidden subgraphs of Ω_n^* whose size is constant (independent of n). We refer to such subgraphs as **small-scale obstructions**. Specifically, we prove that each of the graphs depicted in Figure 19 can not be realized in Ω_n^*. Proving that each of these graphs is forbidden will once again involve (i) translating the structure of the graph into a system of combinatorial equations that are necessarily satisfied by one of its conjugacy classes w, and (ii) proving that this set of combinatorial equations is infeasible. Each of these small-scale obstructions is dealt with in turn.

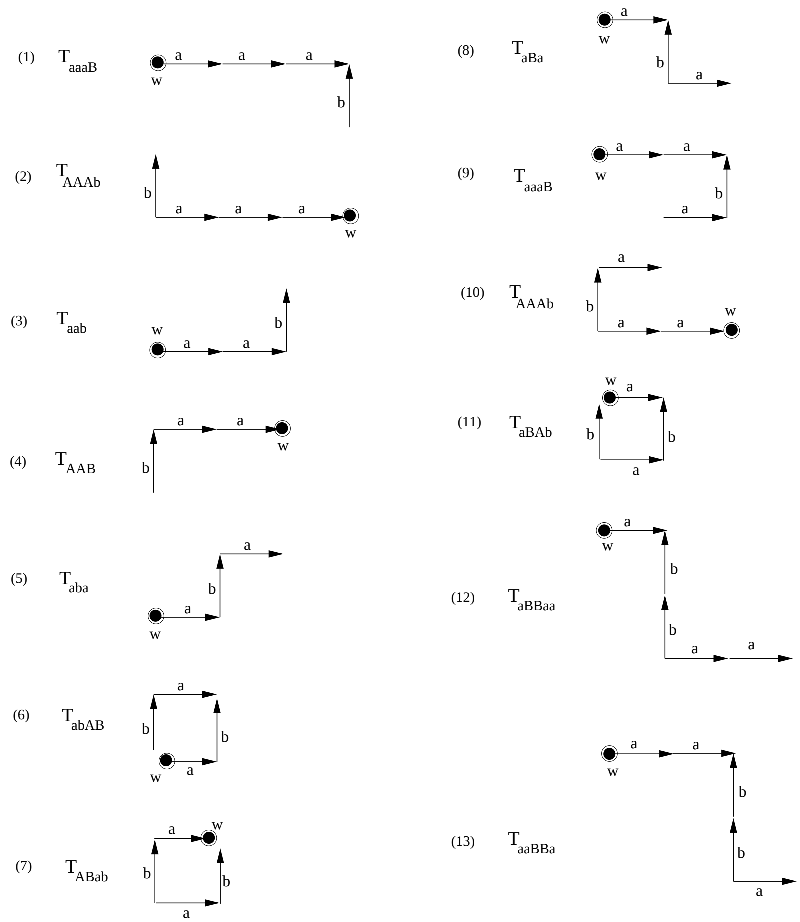

Figure 19: A menagerie of forbidden subgraphs, none of which can be realized in Ω_n^*.

3.4.1. *Obstruction 1: The Forbidden Graph* T_{aaaB}. The vertices of Obstruction T_{aaaB} are named:

$$\begin{aligned} v_1 &= \phi_a(v_0) \\ v_2 &= \phi_a(v_1) \\ v_3 &= \phi_a(v_2) \\ v_4 &= \phi_B(v_3). \end{aligned}$$

The graph T_{aaaB} is depicted in figure 20.

We shall now deduce a set of constraints on the structure of the vertex v_0.

Constraints from v_1. Since $|v_0| = |\phi_a(v_0)| = |v_1|$, it follows from Corollary 2.3.9 (pp. 18) that

$$\boxed{\sharp(a, v_0) = 2\sharp(aB, v_0).} \tag{11}$$

(1) T_{aaaB}

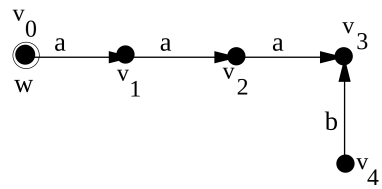

Figure 20: The Forbidden Graph T_{aaaB}.

Constraints from v_2. Since $|v_1| = |\phi_a(v_1)| = |v_2|$, it follows from Corollary 2.3.9 (pp. 18) that

$$\textcircled{\sharp}(a, v_1) \quad = \quad 2\textcircled{\sharp}(aB, v_1).$$

We will now compute $\textcircled{\sharp}(a, v_1)$ and $2\textcircled{\sharp}(aB, v_1)$, respectively. First, let us consider $\textcircled{\sharp}(a, v_1)$. Since v_1 is $\phi_a(v_0)$, it follows that

$$\textcircled{\sharp}(a, v_1) \quad = \quad \textcircled{\sharp}(a, v_0).$$

Now let us consider $2\textcircled{\sharp}(aB, v_1)$. Since v_1 is $\phi_a(v_0)$, it follows that

$$2\textcircled{\sharp}(aB, v_1) \quad = \quad 2\textcircled{\sharp}(aBA, v_0) + 2\textcircled{\sharp}(aB^2, v_0).$$

Thus $|v_1| = |\phi_a(v_1)| = |v_2|$ implies the following constraint on v_0:

$$\boxed{\textcircled{\sharp}(a, v_0) \quad = \quad 2\textcircled{\sharp}(aB^2, v_0) + 2\textcircled{\sharp}(aBA, v_0).} \tag{12}$$

Constraints from v_3. Since $|v_2| = |\phi_a(v_2)| = |v_3|$, it follows from Corollary 2.3.9 (pp. 18) that

$$\textcircled{\sharp}(a, v_2) \quad = \quad 2\textcircled{\sharp}(aB, v_2).$$

We will now compute $\textcircled{\sharp}(a, v_2)$ and $2\textcircled{\sharp}(aB, v_2)$, respectively. First, let us consider $\textcircled{\sharp}(a, v_2)$. Since v_2 is $\phi_a(v_1)$, it follows that

$$\textcircled{\sharp}(a, v_2) \quad = \quad \textcircled{\sharp}(a, v_1).$$

Since v_1 is $\phi_a(v_0)$, it follows that

$$\textcircled{\sharp}(a, v_1) \quad = \quad \textcircled{\sharp}(a, v_0).$$

To summarize, we have shown that

$$\textcircled{\sharp}(a, v_2) \quad = \quad \textcircled{\sharp}(a, v_0).$$

Now let us consider $2\textcircled{\sharp}(aB, v_2)$. Since v_2 is $\phi_a(v_1)$, it follows that

$$2\textcircled{\sharp}(aB, v_2) \quad = \quad 2\textcircled{\sharp}(aBA, v_1) + 2\textcircled{\sharp}(aB^2, v_1).$$

Since v_1 is $\phi_a(v_0)$, it follows that

$$\begin{aligned} 2\textcircled{\sharp}(aBA, v_1) &= 2\textcircled{\sharp}(aBA, v_0) \\ 2\textcircled{\sharp}(aB^2, v_1) &= 2\textcircled{\sharp}(aB^3, v_0) + 2\textcircled{\sharp}(aB^2A, v_0). \end{aligned}$$

To summarize, we have shown that

$$2\textcircled{\sharp}(aB, v_2) \quad = \quad \begin{array}{l} [\, 2\textcircled{\sharp}(aB^3, v_0) + 2\textcircled{\sharp}(aB^2A, v_0) \\ \quad + 2\textcircled{\sharp}(aBA, v_0) \,] \end{array} .$$

Thus $|v_2| = |\phi_a(v_2)| = |v_3|$ implies the following constraint on v_0:

(13)
$$\boxed{\textcircled{$\sharp$}(a, v_0) \quad = \quad \begin{array}{l}[\,2\textcircled{$\sharp$}(aB^3, v_0) + 2\textcircled{$\sharp$}(aB^2A, v_0) \\ \quad + 2\textcircled{$\sharp$}(aBA, v_0)\,]\end{array} \quad .}$$

Constraints from v_4. Since $|v_3| = |\phi_B(v_3)| = |v_4|$, it follows from Corollary 2.3.9 (pp. 18) that

$$\textcircled{$\sharp$}(B, v_3) \quad = \quad 2\textcircled{$\sharp$}(BA, v_3).$$

We will now compute $\textcircled{$\sharp$}(B, v_3)$ and $2\textcircled{$\sharp$}(BA, v_3)$, respectively. First, let us consider $\textcircled{$\sharp$}(B, v_3)$. Since v_3 is $\phi_a(v_2)$, it follows that

$$\textcircled{$\sharp$}(B, v_3) \quad = \quad \textcircled{$\sharp$}(a, v_2) - 2\textcircled{$\sharp$}(aB, v_2) + \textcircled{$\sharp$}(b, v_2).$$

Since v_2 is $\phi_a(v_1)$, it follows that

$$\begin{aligned} \textcircled{$\sharp$}(b, v_2) &= -2\textcircled{$\sharp$}(aB, v_1) + \textcircled{$\sharp$}(a, v_1) + \textcircled{$\sharp$}(b, v_1) \\ \textcircled{$\sharp$}(a, v_2) &= \textcircled{$\sharp$}(a, v_1) \\ -2\textcircled{$\sharp$}(aB, v_2) &= -2\textcircled{$\sharp$}(aBA, v_1) - 2\textcircled{$\sharp$}(aB^2, v_1). \end{aligned}$$

Since v_1 is $\phi_a(v_0)$, it follows that

$$\begin{aligned} \textcircled{$\sharp$}(b, v_1) &= \textcircled{$\sharp$}(b, v_0) + \textcircled{$\sharp$}(a, v_0) - 2\textcircled{$\sharp$}(aB, v_0) \\ 2\textcircled{$\sharp$}(a, v_1) &= 2\textcircled{$\sharp$}(a, v_0) \\ -2\textcircled{$\sharp$}(aB^2, v_1) &= -2\textcircled{$\sharp$}(aB^3, v_0) - 2\textcircled{$\sharp$}(aB^2A, v_0) \\ -2\textcircled{$\sharp$}(aB, v_1) &= -2\textcircled{$\sharp$}(aBA, v_0) - 2\textcircled{$\sharp$}(aB^2, v_0) \\ -2\textcircled{$\sharp$}(aBA, v_1) &= -2\textcircled{$\sharp$}(aBA, v_0). \end{aligned}$$

To summarize, we have shown that

$$\textcircled{$\sharp$}(B, v_3) \quad = \quad \begin{array}{l}[\,3\textcircled{$\sharp$}(a, v_0) - 2\textcircled{$\sharp$}(aB^2A, v_0) \\ \quad - 2\textcircled{$\sharp$}(aB, v_0) - 2\textcircled{$\sharp$}(aB^2, v_0) \\ \quad - 2\textcircled{$\sharp$}(aB^3, v_0) + \textcircled{$\sharp$}(b, v_0) \\ \quad - 4\textcircled{$\sharp$}(aBA, v_0)\,]\end{array} \quad .$$

The contraints deduced at vertex v_3 can now be used to simply this expression. Specifically, since

$$\textcircled{$\sharp$}(a, v_0) = \begin{array}{l}[\,2\textcircled{$\sharp$}(aB^3, v_0) + 2\textcircled{$\sharp$}(aB^2A, v_0) \\ \quad + 2\textcircled{$\sharp$}(aBA, v_0)\,]\end{array} \quad ,$$

it follows that

$$\begin{array}{r}[\,3\textcircled{$\sharp$}(a, v_0) - 2\textcircled{$\sharp$}(aB^2A, v_0) \\ - 2\textcircled{$\sharp$}(aB, v_0) - 2\textcircled{$\sharp$}(aB^2, v_0) \\ - 2\textcircled{$\sharp$}(aB^3, v_0) + \textcircled{$\sharp$}(b, v_0) \\ - 4\textcircled{$\sharp$}(aBA, v_0)\,]\end{array} \quad = \quad \begin{array}{l}[\,-2\textcircled{$\sharp$}(aBA, v_0) - 2\textcircled{$\sharp$}(aB^2, v_0) \\ \quad + 2\textcircled{$\sharp$}(a, v_0) + \textcircled{$\sharp$}(b, v_0) \\ \quad - 2\textcircled{$\sharp$}(aB, v_0)\,]\end{array} \quad .$$

The contraints deduced at vertex v_2 can now be used to simply this expression. Specifically, since

$$\textcircled{$\sharp$}(a, v_0) = 2\textcircled{$\sharp$}(aB^2, v_0) + 2\textcircled{$\sharp$}(aBA, v_0),$$

it follows that

$$\begin{array}{l} [\, -2\textcircled{\sharp}(aBA, v_0) - 2\textcircled{\sharp}(aB^2, v_0) \\ \qquad + 2\textcircled{\sharp}(a, v_0) + \textcircled{\sharp}(b, v_0) \\ \qquad\qquad - 2\textcircled{\sharp}(aB, v_0)\,] \end{array} = \textcircled{\sharp}(a, v_0) + \textcircled{\sharp}(b, v_0) - 2\textcircled{\sharp}(aB, v_0).$$

The contraints deduced at vertex v_1 can now be used to simply this expression. Specifically, since

$$\textcircled{\sharp}(a, v_0) = 2\textcircled{\sharp}(aB, v_0),$$

it follows that

$$\textcircled{\sharp}(a, v_0) + \textcircled{\sharp}(b, v_0) - 2\textcircled{\sharp}(aB, v_0) = \textcircled{\sharp}(b, v_0).$$

So, we see that

$$\textcircled{\sharp}(B, v_3) = \textcircled{\sharp}(b, v_0).$$

Now let us consider $2\textcircled{\sharp}(BA, v_3)$. Since v_3 is $\phi_a(v_2)$, it follows that

$$2\textcircled{\sharp}(BA, v_3) = 2\textcircled{\sharp}(A^2, v_2) + 2\textcircled{\sharp}(BA, v_2).$$

Since v_2 is $\phi_a(v_1)$, it follows that

$$\begin{aligned} 2\textcircled{\sharp}(BA, v_2) &= 2\textcircled{\sharp}(BA, v_1) + 2\textcircled{\sharp}(A^2, v_1) \\ 2\textcircled{\sharp}(A^2, v_2) &= 2\textcircled{\sharp}(AbA, v_1). \end{aligned}$$

Since v_1 is $\phi_a(v_0)$, it follows that

$$\begin{aligned} 2\textcircled{\sharp}(BA, v_1) &= 2\textcircled{\sharp}(A^2, v_0) + 2\textcircled{\sharp}(BA, v_0) \\ 2\textcircled{\sharp}(A^2, v_1) &= 2\textcircled{\sharp}(AbA, v_0) \\ 2\textcircled{\sharp}(AbA, v_1) &= 2\textcircled{\sharp}(Ab^2A, v_0). \end{aligned}$$

To summarize, we have shown that

$$2\textcircled{\sharp}(BA, v_3) = \begin{array}{l} [\, 2\textcircled{\sharp}(BA, v_0) + 2\textcircled{\sharp}(AbA, v_0) \\ \quad + 2\textcircled{\sharp}(Ab^2A, v_0) + 2\textcircled{\sharp}(A^2, v_0)\,] \end{array}.$$

Thus $|v_3| = |\phi_B(v_3)| = |v_4|$ implies the following constraint on v_0:

$$\boxed{\textcircled{\sharp}(b, v_0) = \begin{array}{l} [\, 2\textcircled{\sharp}(BA, v_0) + 2\textcircled{\sharp}(AbA, v_0) \\ \quad + 2\textcircled{\sharp}(Ab^2A, v_0) + 2\textcircled{\sharp}(A^2, v_0)\,] \end{array}.} \tag{14}$$

Having determined the constrains on v_0 entailed by each of the vertices in T_{aaaB}, we are now ready to prove the following proposition.

PROPOSITION 3.4.1. *T_{aaaB} is not realized as a level subgraph in Ω_n^* for $n \geqslant 4$.*

PROOF. Suppose, towards contradiction that T_{aaaB} is realized. Then there exists a conjugacy class $w \in \mathcal{S}_{aaaB}$.

Appealing to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 2$

$$\textcircled{\sharp}(BA, v_0) + \textcircled{\sharp}(A^2, v_0) = \textcircled{\sharp}(a, v_0) - \textcircled{\sharp}(aB, v_0).$$

Applying this to simplify the constraints (14),

$$\textcircled{\sharp}(b, v_0) = 2\textcircled{\sharp}(a, v_0) + 2[-\textcircled{\sharp}(aB, v_0) + \textcircled{\sharp}(aBa, v_0) + \textcircled{\sharp}(aB^2a, v_0)] \ . \tag{15}$$

Appealing to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 4$

$$\textcircled{\sharp}(aB^3, v_0) + \textcircled{\sharp}(aB^2A, v_0) = \textcircled{\sharp}(aB^2, v_0) - \textcircled{\sharp}(aB^2a, v_0)$$

Applying this to simplify the constraints (13),

$$\textcircled{\sharp}(a, v_0) \quad = \quad [\, 2\textcircled{\sharp}(aB^2, v_0) - 2\textcircled{\sharp}(aB^2a, v_0) + 2\textcircled{\sharp}(aBA, v_0)\,] \ .$$

Appealing again to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 3$

$$\textcircled{\sharp}(aB^2, v_0) + \textcircled{\sharp}(aBA, v_0) = \textcircled{\sharp}(aB, v_0) - \textcircled{\sharp}(aBa, v_0)$$

so

(16) $$\textcircled{\sharp}(a, v_0) \quad = \quad [\, 2\textcircled{\sharp}(aB, v_0) - 2\textcircled{\sharp}(aBa, v_0) - 2\textcircled{\sharp}(aB^2a, v_0)\,] \ .$$

Combining equations (15) and (16), we see that

(17) $$\textcircled{\sharp}(b, v_0) \quad = \quad \textcircled{\sharp}(a, v_0)$$

Now, by the Tail Lemma at $k_0 = 3$, we know that

(18) $$\textcircled{\sharp}(b, v_0) \quad \geqslant \quad 3\textcircled{\sharp}(aB^3, v_0) + 3\textcircled{\sharp}(aB^2A, v_0) + 2\textcircled{\sharp}(aBA, v_0).$$

On the other hand, the constraints (13) deduced at vertex v_3 together with (17) tell us that

(19) $$\textcircled{\sharp}(b, v_0) \quad = \quad [\, 2\textcircled{\sharp}(aB^3, v_0) + 2\textcircled{\sharp}(aB^2A, v_0) + 2\textcircled{\sharp}(aBA, v_0)\,] \ .$$

Combining equations (18) and (19), we see that

$$0 \quad \geqslant \quad \textcircled{\sharp}(aB^3, v_0) + \textcircled{\sharp}(aB^2A, v_0).$$

and hence that $\textcircled{\sharp}(aB^3, v_0) = \textcircled{\sharp}(aB^2A, v_0) = 0$.

The constraints (13) then reduce to stating that

$$\textcircled{\sharp}(a, v_0) \quad = \quad 2\textcircled{\sharp}(aBA, v_0).$$

It follows that every occurrence of $a^{\pm}$ occurs inside an a-blocker. This implies that $\phi_a(w) = w$, and hence that v_0 and v_1 coincide. Thus $w \notin \mathcal{S}_{aaaB}$. □

Appealing to the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36), we obtain the following immediate corollary.

COROLLARY 3.4.2 (Obstruction 1). *The sets $\mathcal{S}_{aaaB}$, $\mathcal{S}_{bAAA}$, $\mathcal{S}_{bbbA}$, and $\mathcal{S}_{aBBB}$ contain no conjugacy classes of length $\geqslant 4$.*

3.4.2. *Obstruction 2: The Forbidden Graph T_{AAAb}.* The vertices of Obstruction T_{AAAb} are named: $v_1 = \phi_A(v_0)$, $v_2 = \phi_A(v_1)$, $v_3 = \phi_A(v_2)$, $v_4 = \phi_b(v_3)$. The graph T_{AAAb} is depicted in figure 21.

(2) T_{AAAb}

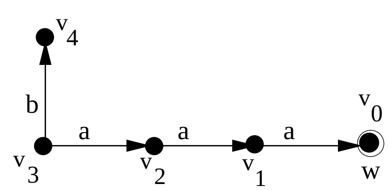

Figure 21: The Forbidden Graph T_{AAAb}.

Following the same type of process that was carried out for T_{aaaB} (see pp. 37-40), we can derive a set of constraints on the structure of the vertex v_0.

In particular, from the fact that $|v_0| = |\phi_A(v_0)| = |v_1|$, we can deduce that

$$(20) \qquad \text{\textcircled{$\sharp$}}(A, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(AB, v_0).$$

Then, from the fact that $|v_1| = |\phi_A(v_1)| = |v_2|$, we can deduce that

$$(21) \qquad \text{\textcircled{$\sharp$}}(A, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(AB^2, v_0) + 2\text{\textcircled{$\sharp$}}(ABa, v_0).$$

Then, from the fact that $|v_2| = |\phi_A(v_2)| = |v_3|$, we can deduce that

$$(22) \qquad \text{\textcircled{$\sharp$}}(A, v_0) \quad = \quad \begin{array}{l} [\,2\text{\textcircled{$\sharp$}}(AB^2a, v_0) + 2\text{\textcircled{$\sharp$}}(ABa, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(AB^3, v_0)\,] \end{array} .$$

Then, from the fact that $|v_3| = |\phi_b(v_3)| = |v_4|$, we can deduce that

$$(23) \qquad \text{\textcircled{$\sharp$}}(B, v_0) \quad = \quad \begin{array}{l} [\,2\text{\textcircled{$\sharp$}}(bA, v_0) + 2\text{\textcircled{$\sharp$}}(AB^2A, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(A^2, v_0) + 2\text{\textcircled{$\sharp$}}(ABA, v_0)\,] \end{array} .$$

Having determined the constrains on v_0 entailed by each of the vertices in T_{AAAb}, we are now ready to prove the following proposition.

Proposition 3.4.3. *The graph T_{AAAb} cannot be realized as a subgraph of Ω_n^* for $n \geqslant 4$.*

Proof. Suppose, towards contradiction that T_{AAAb} is realized. Then there exists a conjugacy class $w \in \mathcal{S}_{AAAb}$.

Appealing to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 2$,

$$\text{\textcircled{$\sharp$}}(aB, v_0) + \text{\textcircled{$\sharp$}}(a^2, v_0) = \text{\textcircled{$\sharp$}}(a, v_0) - \text{\textcircled{$\sharp$}}(ab, v_0).$$

Applying this to constraints (23) from v_4,

$$(24) \quad \text{\textcircled{$\sharp$}}(b, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(a, v_0) - 2\text{\textcircled{$\sharp$}}(ab, v_0) + 2\text{\textcircled{$\sharp$}}(aba, v_0) + 2\text{\textcircled{$\sharp$}}(ab^2a, v_0)] \; .$$

Appealing to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 4$,

$$\text{\textcircled{$\sharp$}}(AB^3, v_0) + \text{\textcircled{$\sharp$}}(AB^2a, v_0) = \text{\textcircled{$\sharp$}}(AB^2, v_0) - \text{\textcircled{$\sharp$}}(AB^2A, v_0)$$

Applying to constraints (22) from v_3,

$$\text{\textcircled{$\sharp$}}(a, v_0) \quad = \quad [\,2\text{\textcircled{$\sharp$}}(AB^2, v_0) - 2\text{\textcircled{$\sharp$}}(AB^2A, v_0) + 2\text{\textcircled{$\sharp$}}(ABa, v_0)\,] \; .$$

Appealing again to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 3$,

$$\text{\textcircled{$\sharp$}}(AB^2, v_0) + \text{\textcircled{$\sharp$}}(ABa, v_0) = \text{\textcircled{$\sharp$}}(AB, v_0) - \text{\textcircled{$\sharp$}}(ABA, v_0),$$

so in fact

$$(25) \quad \text{\textcircled{$\sharp$}}(a, v_0) \quad = \quad [\,2\text{\textcircled{$\sharp$}}(AB, v_0) - 2\text{\textcircled{$\sharp$}}(ABA, v_0) - 2\text{\textcircled{$\sharp$}}(AB^2A, v_0)\,] \; .$$

Equation (24) implies that

$$\text{\textcircled{$\sharp$}}(b, v_0) - 2\text{\textcircled{$\sharp$}}(a, v_0) + 2\text{\textcircled{$\sharp$}}(ab, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(aba, v_0) + 2\text{\textcircled{$\sharp$}}(ab^2a, v_0)$$

while equation (25) implies that

$$-\text{\textcircled{$\sharp$}}(a, v_0) + 2\text{\textcircled{$\sharp$}}(ba, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(aba, v_0) + 2\text{\textcircled{$\sharp$}}(ab^2a, v_0).$$

Combining these expressions we see that

$$(26) \qquad \text{\textcircled{$\sharp$}}(b, v_0) - 2\text{\textcircled{$\sharp$}}(ba, v_0) \quad = \quad \text{\textcircled{$\sharp$}}(a, v_0) - 2\text{\textcircled{$\sharp$}}(ab, v_0)$$

Appealing to the Extension Lemma 2.5.16 (pp. 25), we know that if $|v_0| \geqslant 2$,

$$\begin{array}{rcl} \text{\textcircled{$\sharp$}}(b, v_0) - \text{\textcircled{$\sharp$}}(ba, v_0) & = & \text{\textcircled{$\sharp$}}(bb, v_0) + \text{\textcircled{$\sharp$}}(bA, v_0) \\ \text{\textcircled{$\sharp$}}(a, v_0) - \text{\textcircled{$\sharp$}}(ab, v_0) & = & \text{\textcircled{$\sharp$}}(aa, v_0) + \text{\textcircled{$\sharp$}}(aB, v_0), \end{array}$$

Where by definition $\text{\textcircled{$\sharp$}}(bA, v_0) = \text{\textcircled{$\sharp$}}(aB, v_0)$. We can simplify (26) to

$$\text{\textcircled{$\sharp$}}(bb, v_0) + \text{\textcircled{$\sharp$}}(ab, v_0) \quad = \quad \text{\textcircled{$\sharp$}}(aa, v_0) + \text{\textcircled{$\sharp$}}(ba, v_0).$$

Appealing again to the Extension Lemma 2.5.16 (pp. 25), when $|v_0| \geqslant 2$,

$$\text{\textcircled{$\sharp$}}(b, v_0) + \text{\textcircled{$\sharp$}}(Ab, v_0) \quad = \quad \text{\textcircled{$\sharp$}}(a, v_0) + \text{\textcircled{$\sharp$}}(Ba, v_0).$$

It follows that

$$\text{\textcircled{$\sharp$}}(b, v_0) \quad = \quad \text{\textcircled{$\sharp$}}(a, v_0).$$

Now, by the Tail Lemma at $k_0 = 3$, we know that

$$\text{\textcircled{$\sharp$}}(b, v_0) \quad \geqslant \quad 3\text{\textcircled{$\sharp$}}(AB^3, v_0) + 3\text{\textcircled{$\sharp$}}(AB^2a, v_0) + 2\text{\textcircled{$\sharp$}}(ABa, v_0). \tag{27}$$

Combining (27) with the constraints (22)

$$\text{\textcircled{$\sharp$}}(b, v_0) - \text{\textcircled{$\sharp$}}(a, v_0) \quad \geqslant \quad \text{\textcircled{$\sharp$}}(AB^2a, v_0) + \text{\textcircled{$\sharp$}}(AB^3, v_0).$$

Since $\text{\textcircled{$\sharp$}}(b, v_0) = \text{\textcircled{$\sharp$}}(a, v_0)$, we see that

$$\text{\textcircled{$\sharp$}}(AB^2a, v_0) = \text{\textcircled{$\sharp$}}(AB^3, v_0) \quad = \quad 0$$

Constraints (22) then reduce to stating that

$$\text{\textcircled{$\sharp$}}(a, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(aBA, v_0).$$

It follows that every occurrence of $a^{\pm}$ occurs inside an a-blocker. This implies that $\phi_a(w) = w$, and hence that v_0 and v_1 coincide. Thus $w \notin \mathcal{S}_{AAAb}$. □

Appealing to the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36), the following is an immediate corollary of the previous proposition.

COROLLARY 3.4.4 (Obstruction 2). *The sets $\mathcal{S}_{AAAb}$, $\mathcal{S}_{Baaa}$, $\mathcal{S}_{BBBa}$, $\mathcal{S}_{Abbb}$ contain no conjugacy classes of length $\geqslant 4$.*

3.4.3. *Obstruction 3: The Forbidden Graph T_{aab}.* The vertices of Obstruction T_{aab} are named: $v_1 = \phi_a(v_0)$, $v_2 = \phi_a(v_1)$, $v_3 = \phi_b(v_2)$. The graph T_{aab} is depicted in figure 22.

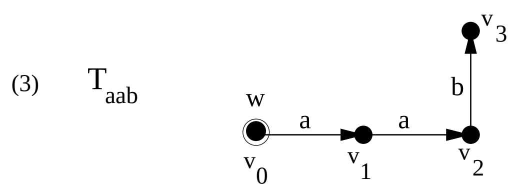

Figure 22: The Forbidden Graph T_{aab}.

Following the same type of process that was carried out for T_{aaaB} (see pp. 37-40), we can derive a set of constraints on the structure of the vertex v_0.

In particular, from the fact that $|v_0| = |\phi_a(v_0)| = |v_1|$, we can deduce that

$$\text{\textcircled{$\sharp$}}(a, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(aB, v_0). \tag{28}$$

Then, from the fact that $|v_1| = |\phi_a(v_1)| = |v_2|$, we can deduce that

$$\text{\textcircled{$\sharp$}}(a, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(aB^2, v_0) + 2\text{\textcircled{$\sharp$}}(aBA, v_0). \tag{29}$$

Then, from the fact that $|v_2| = |\phi_b(v_2)| = |v_3|$, we can deduce that

$$(30) \qquad \textcircled{\sharp}(b, v_0) \;=\; \begin{aligned}[t] &[\, 2\textcircled{\sharp}(ab^2A, v_0) + 2\textcircled{\sharp}(abA, v_0) \\ &\quad + 2\textcircled{\sharp}(b^3A, v_0)\,] \end{aligned} .$$

Having determined the constrains on v_0 entailed by each of the vertices in T_{aab}, we are now ready to prove the following proposition.

PROPOSITION 3.4.5. *T_{aab} is not realized as a level subgraph of Ω^*.*

PROOF. Suppose, towards contradiction that T_{aab} is realized. Then there exists a conjugacy class $w \in \mathcal{S}_{aab}$.

Note that b^3A is both a-demarcator-immune and self-immune. So by the Immunity Lemma 2.5.12 (pp. 24),

$$(31) \qquad \textcircled{\sharp}(b, v_0) \;\geqslant\; 2\textcircled{\sharp}(abA, v_0) + 3\textcircled{\sharp}(ab^2A, v_0) + 3\textcircled{\sharp}(b^3A, v_0).$$

Combining (31) and (30), it follows that

$$0 \;\geqslant\; \textcircled{\sharp}(ab^2A, v_0) + \textcircled{\sharp}(b^3A, v_0),$$

and hence that $\textcircled{\sharp}(ab^2A, v_0) = \textcircled{\sharp}(b^3A, v_0) = 0$. We have shown that

$$(32) \qquad \textcircled{\sharp}(b, v_0) \;=\; 2\textcircled{\sharp}(abA, v_0)$$

By the Demarcator Lemma 2.5.6 (pp. 23), we know that

$$(33) \qquad \textcircled{\sharp}(b, v_0) \;=\; 2\textcircled{\sharp}(abA, v_0) + \sum_{k=3}^{|w|-3} (k+1)\textcircled{\sharp}(ab^kA, v_0).$$

Combining (32) and (33), we see that $\textcircled{\sharp}(ab^kA, v_0) = 0$, for all $k \geqslant 2$. In other words, all a-blockers must have weight 1.

Since aB^2 is a-demarcator-immune and self-immune, and all a-blockers are of weight 1, the Immunity Lemma 2.5.12 (pp. 24) gives us that

$$(34) \qquad \textcircled{\sharp}(b, v_0) \;\geqslant\; 2\textcircled{\sharp}(abA, v_0) + 2\textcircled{\sharp}(aB^2, \kappa_a(v_0)).$$

But since all a-blockers have weight 1, all occurrences of aB^2 in v_0 must occur inside $\kappa_a(v_0)$. Thus, we have

$$(35) \qquad \textcircled{\sharp}(b, v_0) \;\geqslant\; 2\textcircled{\sharp}(abA, v_0) + 2\textcircled{\sharp}(aB^2, v_0).$$

Combining this with (32), we can conclude that $\textcircled{\sharp}(aB^2, v_0) = 0$. Applying this in turn to (29), we see that

$$(36) \qquad \textcircled{\sharp}(a, v_0) \;=\; 2\textcircled{\sharp}(abA, v_0).$$

It follows that every occurrence of $a^{\pm}$ occurs inside an a-blocker. This implies that $\phi_a(w) = w$, and hence that v_0 and v_1 coincide. Thus $w \notin \mathcal{S}_{aab}$. $\square$

Appealing to the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36), the following is an immediate corollary of the previous proposition.

COROLLARY 3.4.6 (Obstruction 3). *$\mathcal{S}_{aab}$, $\mathcal{S}_{BAA}$, $\mathcal{S}_{bba}$, and $\mathcal{S}_{ABB}$ are empty.*

(4) T_{AAB}

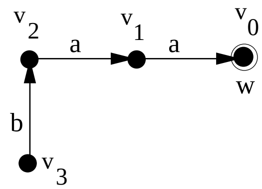

Figure 23: The Forbidden Graph T_{AAB}.

3.4.4. *Obstruction 4: The Forbidden Graph T_{AAB}.* The vertices of Obstruction T_{AAB} are named: $v_1 = \phi_A(v_0)$, $v_2 = \phi_A(v_1)$, $v_3 = \phi_B(v_2)$. The graph T_{AAB} is depicted in figure 23.

Following the same type of process that was carried out for T_{aaaB} (see pp. 37-40), we can derive a set of constraints on the structure of the vertex v_0.

In particular, from the fact that $|v_0| = |\phi_A(v_0)| = |v_1|$, we can deduce that

$$\textcircled{\sharp}(A, v_0) \;\; = \;\; 2\textcircled{\sharp}(AB, v_0). \tag{37}$$

Then, from the fact that $|v_1| = |\phi_A(v_1)| = |v_2|$, we can deduce that

$$\textcircled{\sharp}(A, v_0) \;\; = \;\; 2\textcircled{\sharp}(ABa, v_0) + 2\textcircled{\sharp}(AB^2, v_0). \tag{38}$$

Then, from the fact that $|v_2| = |\phi_B(v_2)| = |v_3|$, we can deduce that

$$\textcircled{\sharp}(B, v_0) \;\; = \;\; \begin{array}{l} [\; 2\textcircled{\sharp}(aBA, v_0) + 2\textcircled{\sharp}(aB^2A, v_0) \\ \quad + 2\textcircled{\sharp}(B^3A, v_0) \;] \end{array} . \tag{39}$$

Having determined the constrains on v_0 entailed by each of the vertices in T_{AAB}, we are now ready to prove the following proposition.

PROPOSITION 3.4.7. *The graph T_{AAB} cannot be realized as a level subgraph of Ω^*.*

PROOF. Suppose, towards contradiction that T_{AAB} is realized. Then there exists a conjugacy class $w \in \mathcal{S}_{AAB}$.

Note that B^3A is both a-demarcator-immune and self-immune. So by the Immunity Lemma 2.5.12 (pp. 24),

$$\textcircled{\sharp}(b, v_0) \;\; \geqslant \;\; 2\textcircled{\sharp}(aBA, v_0) + 3\textcircled{\sharp}(aB^2A, v_0) + 3\textcircled{\sharp}(b^3A, v_0). \tag{40}$$

Combining (40) and (39), it follows that

$$0 \;\; \geqslant \;\; \textcircled{\sharp}(aB^2A, v_0) + \textcircled{\sharp}(B^3A, v_0),$$

and hence that $\textcircled{\sharp}(aB^2A, v_0) = \textcircled{\sharp}(B^3A, v_0) = 0$. We have shown that

$$\textcircled{\sharp}(b, v_0) \;\; = \;\; 2\textcircled{\sharp}(aBA, v_0) \tag{41}$$

By the Demarcator Lemma 2.5.6 (pp. 23), we know that

$$\textcircled{\sharp}(b, v_0) \;\; = \;\; 2\textcircled{\sharp}(aBA, v_0) + \sum_{k=3}^{|w|-3} (k+1)\textcircled{\sharp}(aB^kA, v_0). \tag{42}$$

Combining (41) and (42), we see that $\textcircled{\sharp}(aB^kA, v_0) = 0$, for all $k \geqslant 2$. In other words, all a-blockers have weight 1.

The Tail Lemma 2.5.7 (pp. 23) at $k_0 = 2$ tells us that

$$\textcircled{\sharp}(b, v_0) \;\; \geqslant \;\; 2\textcircled{\sharp}(ABa, v_0) + 2\textcircled{\sharp}(AB^2, v_0). \tag{43}$$

Combining (38) and (43), we see that $\text{\textcircled{$\sharp$}}(b, v_0) \geqslant \text{\textcircled{$\sharp$}}(a, v_0)$. On the other hand, $\text{\textcircled{$\sharp$}}(a, v_0) \geqslant 2\text{\textcircled{$\sharp$}}(aBA, v_0)$. But $\text{\textcircled{$\sharp$}}(aBA, v_0) = \text{\textcircled{$\sharp$}}(b, v_0)$. Thus, we conclude that $\text{\textcircled{$\sharp$}}(b, v_0) = \text{\textcircled{$\sharp$}}(a, v_0)$.

Appealing to (41), we see that

$$\text{\textcircled{$\sharp$}}(a, v_0) \quad = \quad 2\text{\textcircled{$\sharp$}}(aBA, v_0)$$

It follows that every occurrence of $a^{\pm}$ occurs inside an a-blocker. This implies, by the Symmetry Lemma 2.5.17 (pp. 25) that every occurrence of $a^{\pm}$ occurs inside an A-blocker. But then, $\phi_A(w) = w$, and hence that v_0 and v_1 coincide. Thus $w \notin \mathcal{S}_{AAB}$. $\square$

Appealing to the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36), the following is an immediate corollary of the previous proposition.

Corollary 3.4.8 (Obstruction 4). *$\mathcal{S}_{AAB}$, $\mathcal{S}_{baa}$, $\mathcal{S}_{BBA}$, and $\mathcal{S}_{abb}$ are empty.*

3.5. Obstruction Rewriting Rules, Part I. We will use the following lemma to devise a graph rewriting rule that can be used to generate new forbidden graphs from the ones we have found so far.

Lemma 3.5.1 (The $x\hat{x}^{-1}x$ Relation Lemma). *For all $x \in X \cup X^{-1}$ and $w \in \tilde{F}_2$,*

$$\psi_{x\hat{x}^{-1}x}(w) \approx_{\Pi} w.$$

Proof. When $\delta = +1$, the composite map $\psi_{x\hat{x}^{-1}x}$ takes

$$\begin{array}{rcl} x & \mapsto & x\hat{x}x^{-1} \\ \hat{x} & \mapsto & x^{-1}. \end{array}$$

Taking $\pi_0 \in \Pi$ to be the map

$$\begin{array}{rcl} x & \mapsto & \hat{x}^{-1} \\ \hat{x} & \mapsto & x, \end{array}$$

we see that $\pi_0\psi_{x\hat{x}^{-1}x}(w) = w$ and hence that $\psi_{x\hat{x}^{-1}x}(w) \approx_{\Pi} w$.

When $\delta = -1$, the composite map $\psi_{x\hat{x}^{-1}x}$ takes

$$\begin{array}{rcl} x & \mapsto & \hat{x}^{-1} \\ \hat{x} & \mapsto & \hat{x}x\hat{x}^{-1}. \end{array}$$

Taking $\pi_1 \in \Pi$ to be the map

$$\begin{array}{rcl} x & \mapsto & \hat{x} \\ \hat{x} & \mapsto & x^{-1}, \end{array}$$

we see that $\pi_1\psi_{x\hat{x}^{-1}x}(w) = w$ and hence that $\psi_{x\hat{x}^{-1}x}(w) \approx_{\Pi} w$.

This completes the proof of the lemma. $\square$

The $x\hat{x}^{-1}x$-Relation Lemma 3.5.1 (pp. 46) proved above can be leveraged to provide a graph-rewriting rule for obstructions. A graph-rewriting rule τ is a deterministic procedure which produces a new hypothetical graphs from old hypothetical graphs.

Definition 3.5.2. *A graph-rewriting rule τ is conservative if for every hypothetical subgraph T of Ω^*, T cannot be realized as a level subgraph of Ω^* if and only if τT cannot be realized as a level subgraph of Ω^*.*

Conservative graph-rewriting rules allow us to enlarge the set of forbidden graphs from ones already known.

We define graph-rewriting schema τ_x ($x \in X \cup X^{-1}$) which act on the hypothetical subgraphs of Ω^*. Before τ_x can be made to act on a hypothetical graph T, it must be parametrized by a suitable triple of vertices $p, v, q \in V[T]$.

DEFINITION 3.5.3. *Given a hypothetical subgraph T of Ω^*, fix $x \in X^{\pm}$, and let v, p, q be three distinct vertices in $V[T]$. If v, p, q satisfy the following conditions:*

- *v, p, q are all distinct,*
- *$\tilde{\psi}_x(p) = v$,*
- *$\tilde{\psi}_{\hat{x}}(p) = q$,*
- *$(v, \tilde{\psi}_{\hat{x}^{-1}}(v)) \notin E[T]$,*
- *$(q, r) \in E[T] \Rightarrow r = p$;*

*then we say that the triple (v, p, q) are a τ_x-**pivot** in T.*

The graph-rewriting transformation $\tau_{x[p,v,q]}$ acts on T as follows:

DEFINITION 3.5.4. *Given a hypothetical subgraph T of Ω^*, fix $x \in X^{\pm}$, and (v, p, q) a triple of vertices from $V[T]$. We define the graph $\tau_{x[v,p,q]}(T)$ as follows: If (v, p, q) are not a τ_x-pivot, put $\tau_{x[v,p,q]}(T) = T$. Otherwise, let $\tau_{x[v,p,q]}(T)$ be given by:*

$$V[\tau_{x[v,p,q]}(T)] \quad = \quad V[T] \cup \{u\} \backslash \{q\}$$

where u is a new vertex representing $\tilde{\psi}_{\hat{x}^{-1}}(v)$, and take

$$E[\tau_{x[v,p,q]}(T)] \quad = \quad E[T] \cup (u, v) \backslash (p, q),$$

where the new edge (u, v) is labelled by $\hat{x}^{-1}$ (thereby signifying that $u = \tilde{\psi}_{\hat{x}^{-1}}(v)$).

The operation of $\tau_{a[p,v,q]}$ and $\tau_{b[v,p,q]}$ on T is depicted in Figure 24. In the figure, boxed/outline edges are used to depict where edges are required to not be present.

REMARK 3.5.5. *Given a hypothetical subgraph T of Ω^*, a fixed $x \in X^{\delta}$ ($\delta = \pm 1$), and three vertices v, p, q from $V[T]$, it is easy to verify from Figure 24 that the action of τ_x is invertible. Specifically, for every $x \in X^{\pm}$ we have $\tau^{-1}_{x[p,v,q]} = \tau_{x^{-1}[v,p,\tilde{\psi}_{x^{-1}}(v)]}$.*

The next proposition shows that the previously defined graph rewriting rules τ_x ($x \in X \cup X^{-1}$) are conservative.

PROPOSITION 3.5.6 (Obstruction Rewriting Rule 1). *Given a connected tree $T = (V, E)$ and $p, v, q \in V$, $x \in X^{\pm}$, the tree T is forbidden in Ω^*_n if and only if $\tau_{x[p,v,q]}(T)$ is forbidden in Ω^*_n.*

PROOF. If $\tau_{x[p,v,q]}(T) = T$ the statement is trivial. Otherwise, suppose that T is not forbidden in Ω^*_n. Let Ω_T be a minimal induced subgraph of Ω^*_n which contains an isomorphic copy of T. Without loss of generality, we may take Ω_T to be the graph induced by the isomorphic copy of T. Moreover, we identify T with the isomorphic spanning tree of Ω_T.

Now, since $q = \psi_{\hat{x}x^{-1}\hat{x}}(u)$, by Lemma 3.5.1 (pp. 46), we know that $q \approx_{\Pi} u$. It follows that vertices v and q coincide in Ω_T. Since (u, v) is not in Ω_T and (p, q) is not a cut edge of T, the operation of adding (u, v) and removing (p, q) gives rise

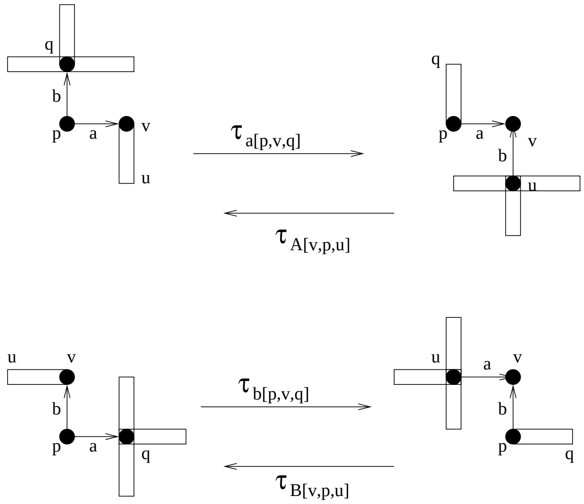

Figure 24: Graph rewriting rules τ_x $(x \in X \cup X^{-1})$.

to just a different spanning tree of Ω_T. It follows that $\tau_{x[p,v,q]}(T)$ is a subgraph of Ω_T, and hence a subgraph of Ω_n^*.

To see the reverse, take $T' = \tau_{x[p,v,q]}(T)$. If T' occurs in Ω_n^* then by the previous argument, so must $\tau_{x^{-1}[v,p,u]}(T')$. But by Remark 3.5.5,

$$\tau_{x^{-1}[v,p,u]} = \tau^{-1}_{x[p,v,q]}.$$

So it follows that T occurs in Ω_n^*. □

3.6. More Small-Scale Obstructions. We will make use of the graph-rewriting rules introduced in the previous section to demonstrate another forbidden graph.

3.6.1. *Obstruction 5: The Forbidden Graph T_{aba}.* The vertices of Obstruction T_{aba} are named: $v_1 = \phi_a(v_0)$, $v_2 = \phi_b(v_1)$, $v_3 = \phi_a(v_2)$. The graph T_{aba} is depicted in figure 25.

(5) $\mathrm{T}_{\mathrm{aba}}$

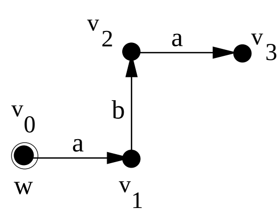

Figure 25: The Forbidden Graph T_{aba}.

Following the same type of process that was carried out for T_{aaaB} (see pp. 37-40), we can derive a set of constraints on the structure of the vertex v_0.

In particular, from the fact that $|v_0| = |\phi_a(v_0)| = |v_1|$, we can deduce that

$$\textcircled{\sharp}(a, v_0) = 2\textcircled{\sharp}(aB, v_0). \tag{44}$$

Then, from the fact that $|v_1| = |\phi_b(v_1)| = |v_2|$, we can deduce that

$$\textcircled{\sharp}(b, v_0) = 2\textcircled{\sharp}(b^2A, v_0) + 2\textcircled{\sharp}(abA, v_0). \tag{45}$$

Then, from the fact that $|v_2| = |\phi_a(v_2)| = |v_3|$, we can deduce that

$$
\text{\textcircled{$\sharp$}}(a, v_0) \;=\; \begin{array}{l}
[\, 2\text{\textcircled{$\sharp$}}(a^2BA, v_0) + 2\text{\textcircled{$\sharp$}}(baBA, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(a^2B^2, v_0) + 2\text{\textcircled{$\sharp$}}(baB^2, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(aBaB^2, v_0) + 2\text{\textcircled{$\sharp$}}(aBaBA, v_0) \,]
\end{array}. \tag{46}
$$

Our goal is to show that T_{aba} cannot be realized as a level subgraph of Ω^*, i.e. that $\mathcal{S}_{aba} = \emptyset$. We will show something slightly weaker, namely that if $\sigma \in F_2$ is a word which contains aba as a proper subword, then $\mathcal{S}_\sigma$ contains no conjugacy class of length $\geqslant 5$. To do this, we will need the following lemma.

LEMMA 3.6.1. *If a conjugacy class w of length $\geqslant 5$ is in $\mathcal{S}_{aba}$, then*

$$\text{\textcircled{$\sharp$}}(aBa, w) = \text{\textcircled{$\sharp$}}(BaB, w) = 0$$

PROOF. Appealing to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 5$,

$$
\begin{aligned}
\text{\textcircled{$\sharp$}}(baBA, v_0) + \text{\textcircled{$\sharp$}}(baB^2, v_0) &= \text{\textcircled{$\sharp$}}(baB, v_0) - \text{\textcircled{$\sharp$}}(baBa, v_0), \\
\text{\textcircled{$\sharp$}}(a^2BA, v_0) + \text{\textcircled{$\sharp$}}(a^2B^2, v_0) &= \text{\textcircled{$\sharp$}}(a^2B, v_0) - \text{\textcircled{$\sharp$}}(a^2Ba, v_0) \\
\text{\textcircled{$\sharp$}}(aBaB^2, v_0) + \text{\textcircled{$\sharp$}}(aBaBA, v_0) &= \text{\textcircled{$\sharp$}}(aBaB, v_0) - \text{\textcircled{$\sharp$}}(aBaBa, v_0).
\end{aligned}
$$

Substituting into (46), we see that

$$
\text{\textcircled{$\sharp$}}(a, v_0) \;=\; \begin{array}{l}
[\, 2\text{\textcircled{$\sharp$}}(baB, v_0) - 2\text{\textcircled{$\sharp$}}(baBa, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(a^2B, v_0) - 2\text{\textcircled{$\sharp$}}(a^2Ba, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(aBaB, v_0) - 2\text{\textcircled{$\sharp$}}(aBaBa, v_0) \,]
\end{array}. \tag{47}
$$

Appealing to the Extension Lemma 2.5.16 (pp. 25), if $|v_0| \geqslant 3$,

$$\text{\textcircled{$\sharp$}}(baB, v_0) + \text{\textcircled{$\sharp$}}(a^2B, v_0) = \text{\textcircled{$\sharp$}}(aB, v_0) - \text{\textcircled{$\sharp$}}(BaB, v_0).$$

Substituting into (47), we see that

$$
\text{\textcircled{$\sharp$}}(a, v_0) \;=\; \begin{array}{l}
[\, 2\text{\textcircled{$\sharp$}}(aB, v_0) - 2\text{\textcircled{$\sharp$}}(BaB, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(baBa, v_0) - 2\text{\textcircled{$\sharp$}}(a^2Ba, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(aBaB, v_0) - 2\text{\textcircled{$\sharp$}}(aBaBa, v_0) \,]
\end{array}. \tag{48}
$$

Appealing to the Extension Lemma 2.5.16 (pp. 25) differently, if $|v_0| \geqslant 5$,

$$
\begin{aligned}
\text{\textcircled{$\sharp$}}(a^2B^2, v_0) + \text{\textcircled{$\sharp$}}(baB^2, v_0) &= \text{\textcircled{$\sharp$}}(aB^2, v_0) - \text{\textcircled{$\sharp$}}(BaB^2, v_0), \\
\text{\textcircled{$\sharp$}}(a^2BA, v_0) + \text{\textcircled{$\sharp$}}(baBA, v_0) &= \text{\textcircled{$\sharp$}}(aBA, v_0) - \text{\textcircled{$\sharp$}}(BaBA, v_0) \\
\text{\textcircled{$\sharp$}}(aBaB^2, v_0) + \text{\textcircled{$\sharp$}}(aBaBA, v_0) &= \text{\textcircled{$\sharp$}}(aBaB, v_0) - \text{\textcircled{$\sharp$}}(aBaBa, v_0).
\end{aligned}
$$

Substituting into (46), we see that

$$
\text{\textcircled{$\sharp$}}(a, v_0) \;=\; \begin{array}{l}
[\, 2\text{\textcircled{$\sharp$}}(aB^2, v_0) - 2\text{\textcircled{$\sharp$}}(BaB^2, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(aBA, v_0) - 2\text{\textcircled{$\sharp$}}(BaBA, v_0) \\
\quad + 2\text{\textcircled{$\sharp$}}(aBaB, v_0) - 2\text{\textcircled{$\sharp$}}(aBaBa, v_0) \,]
\end{array}. \tag{49}
$$

The constraints (44) deduced at vertex v_1 permit us to simplify (48) and (49) and obtain

$$
\begin{aligned}
\text{\textcircled{$\sharp$}}(aBaB, v_0) &= 2\text{\textcircled{$\sharp$}}(BaB, v_0) + 2\text{\textcircled{$\sharp$}}(baBa, v_0) + \\
&\quad\; 2\text{\textcircled{$\sharp$}}(a^2Ba, v_0) + 2\text{\textcircled{$\sharp$}}(aBaBa, v_0) \\
\text{\textcircled{$\sharp$}}(aBaB, v_0) &= 2\text{\textcircled{$\sharp$}}(aBa, v_0) + 2\text{\textcircled{$\sharp$}}(BaB^2, v_0) + \\
&\quad\; 2\text{\textcircled{$\sharp$}}(BaBA, v_0) + 2\text{\textcircled{$\sharp$}}(aBaBa, v_0)
\end{aligned} \tag{50}
$$

By the Subword Lemma 2.5.8 (pp. 24), we know that

$$\begin{aligned}\textcircled{\sharp}(aBaB, v_0) &\leqslant 2\textcircled{\sharp}(BaB, v_0)\\ \textcircled{\sharp}(aBaB, v_0) &\leqslant 2\textcircled{\sharp}(aBa, v_0).\end{aligned}$$

Combining with (50) we see that

$$\begin{aligned}2\textcircled{\sharp}(baBa, v_0) + 2\textcircled{\sharp}(a^2Ba, v_0) + 2\textcircled{\sharp}(aBaBa, v_0) &= 0,\\ 2\textcircled{\sharp}(BaB^2, v_0) + 2\textcircled{\sharp}(BaBA, v_0) + 2\textcircled{\sharp}(aBaBa, v_0) &= 0.\end{aligned} \tag{51}$$

Thus (50) reduces to

$$\textcircled{\sharp}(aBa, v_0) = \textcircled{\sharp}(aBaB, v_0) = \textcircled{\sharp}(BaB, v_0) \tag{52}$$

Suppose, towards contradiction, that $\textcircled{\sharp}(aBa, w) > 0$. Then fix some occurrence of aBa in w. Since $\textcircled{\sharp}(aBa, v_0) = \textcircled{\sharp}(aBaB, v_0)$, it follows that this occurrence must be part of an occurrence of $aBaB$ in w. Now, the next symbol (after this occurrence of $aBaB$) cannot be A or B, since (51) tells us that $\textcircled{\sharp}(BaB^2, v_0) = \textcircled{\sharp}(BaBA, v_0) = 0$. It follows that the next symbol must be a–i.e., this occurrence of $aBaB$ lies inside an occurrence of $aBaBa$. Repeating the argument inductively, for the second occurrence of aBa inside $aBaBa$, we conclude that w is of the form $(aB)^k$ for some $k \in \mathbb{N}$. But then, if $k > 0$ $|\tilde{\psi}_a(w)| < |w|$, contradicting that $|v_1| = |v_0|$. If $k = 0$, then $\tilde{\psi}_a(w) = w$, contradicting that v_1 and v_0 are distinct vertices. We have shown that $\textcircled{\sharp}(aBa, w) = 0$.

A symmetric argument to the one given in the previous paragraph can be used to show that $\textcircled{\sharp}(bAb, w) = 0$. □

Having determined the constrains on v_0 entailed by each of the vertices in T_{AAAb}, and that $\textcircled{\sharp}(aBa, w) = \textcircled{\sharp}(BaB, w) = 0$, we make the following assertion.

Lemma 3.6.2. *If a conjugacy class w of length $\geqslant 5$ is in $\mathcal{S}_{aba}$, then*

$$\textcircled{\sharp}(a, w) = \textcircled{\sharp}(b, w).$$

Proof. Using the Extension Lemma to expand the constraints (28) deduced at vertex v_1, we see that if $|v_0| \geqslant 3$,

$$\textcircled{\sharp}(a, v_0) = 2\textcircled{\sharp}(baB, v_0) + 2\textcircled{\sharp}(aaB, v_0) + 2\textcircled{\sharp}(BaB, v_0).$$

Using the Extension Lemma on the other side, we get that if $|v_0| \geqslant 3$,

$$\textcircled{\sharp}(a, v_0) = 2\textcircled{\sharp}(aBA, v_0) + 2\textcircled{\sharp}(aBB, v_0) + 2\textcircled{\sharp}(aBa, v_0).$$

By Lemma 3.6.1 (pp. 49), $\textcircled{\sharp}(aBa, v_0) = 0$, so

$$\textcircled{\sharp}(a, v_0) = 2\textcircled{\sharp}(aBA, v_0) + 2\textcircled{\sharp}(aBB, v_0)$$

which is, by the constraints (29) deduced at vertex v_2, equal to $\textcircled{\sharp}(b, v_0)$. Thus $\textcircled{\sharp}(a, w) = \textcircled{\sharp}(b, w)$. □

Having shown that $\textcircled{\sharp}(aBa, w) = \textcircled{\sharp}(BaB, w) = 0$ and $\textcircled{\sharp}(a, w) = \textcircled{\sharp}(b, w)$, we are now ready to prove the following proposition. (The definition of "being realized as a *proper* σ-chain", which was given in Definition 3.3.4 on pp. 36).

Proposition 3.6.3. *The graph T_{aba} cannot be realized as a proper simple chain in a level subgraph of Ω_n^* if $n \geqslant 5$.*

Proof. To prove this, we will consider the ways in which T_{aba} might appear as a proper simple chain in an level subgraph of Ω^*. There are six ways:

L1. It occurs as a trailing subgraph of T_{aaba}.
L2. It occurs as a trailing subgraph of T_{Baba}.
L3. It occurs as a trailing subgraph of T_{baba}.
R1. It occurs as an leading subgraph of T_{abaa}.
R2. It occurs as an leading subgraph of T_{abaB}.
R3. It occurs as an leading subgraph of T_{abab}.

Each of these possibilities is depicted in Figure 26.

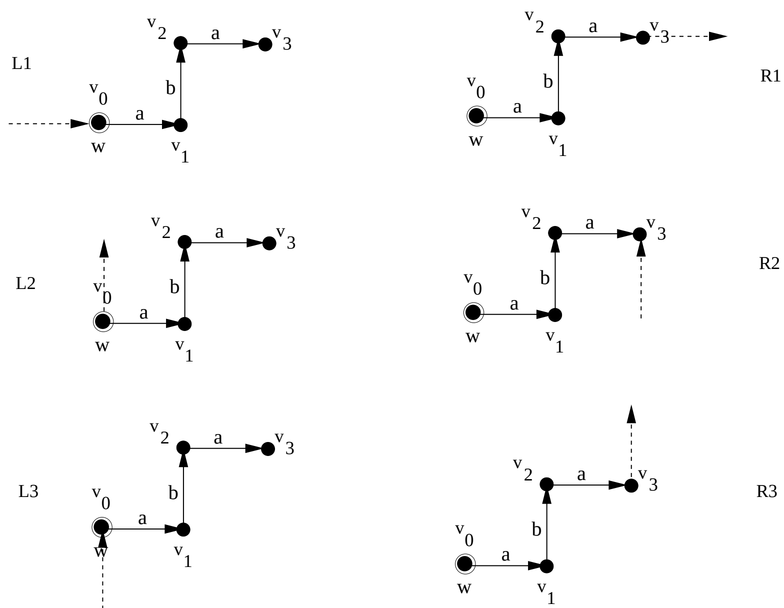

Figure 26: The six ways T_{aba} might appear as a proper simple chain.

The impossibility of each of these configurations will now be proven, in turn. The cases L1, L2, R1, and R2 will be shown to be impossible using the Obstruction Rewriting Rule 1 presented in Proposition 3.5.6 (pp. 47) and by appealing to already-known forbidden graphs. The remaining cases L3 and R3 will be proved separately using a combinatorial argument similar to the ones seen so far for showing the particular graphs are forbidden.

Case L1: The tree T_{aaba} contains the tree T_{aab} as a subgraph. Hence, T_{aaba} cannot be realized, since T_{aab} was shown to be forbidden in Corollary 3.4.6 (pp. 44). The argument is illustrated in Figure 27.

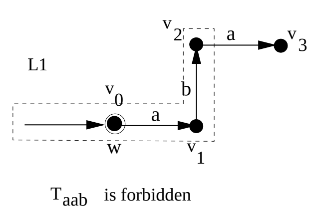

Figure 27: Using graph rewriting to show case L1 is forbidden.

Case L2: The tree T_{Baba} cannot be realized, since applying a graph rewriting rule transforms it into a graph which contains T_{bba}. The latter graph is forbidden, as a consequence of of Corollary 3.4.8 (pp. 46). The argument is illustrated in Figure 28.

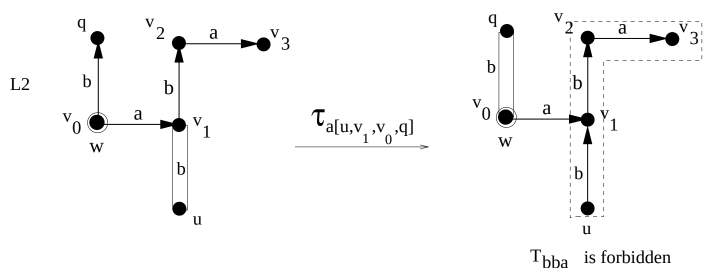

Figure 28: Using graph rewriting to show case L2 is forbidden.

Case R1: The tree T_{abaa} contains the tree T_{baa} as a subgraph. Hence, T_{aaba} cannot be realized, since T_{baa} was shown to be forbidden as a consequence of Corollary 3.4.8 (pp. 46). The argument is illustrated in Figure 29.

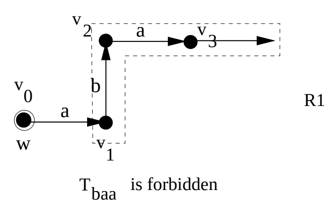

Figure 29: Using graph rewriting to show case R1 is forbidden.

Case R2: The tree T_{abaB} cannot be realized, since applying a graph rewriting rule transforms it into a graph which contains T_{abb}. The latter graph is forbidden, as a consequence of of Corollary 3.4.8 (pp. 46). The argument is illustrated in Figure 30.

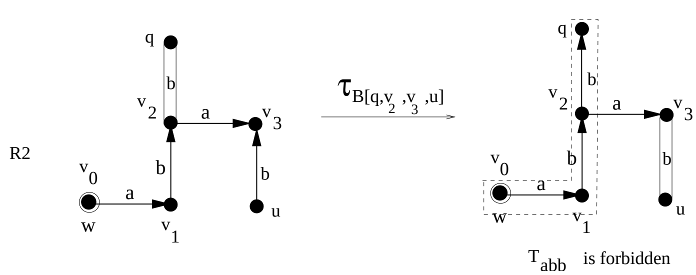

Figure 30: Using graph rewriting to show case R2 is forbidden.

It remains to consider the cases L3 and R3. These are symmetric, so we shall consider only the case R3. We shall show that the graph T_{abab} cannot arise as a level subgraph in Ω^*. We begin by deriving combinatorial conditions from the graph, in a manner similar to the analyses conducted earlier.

Subcase: The Forbidden Graph T_{abab}. The additional vertex of Obstruction T_{abab} is named: $v_4 = \phi_b(v_3)$.

Following the same type of process that was carried out for T_{aaaB} (see pp. 37-40), we are able to deduce from fact that $|v_3| = |\phi_b(v_3)| = |v_4|$, that

$$
(53) \qquad \text{\textcircled{$\sharp$}}(b, v_0) = \begin{array}{l} [\, 2\text{\textcircled{$\sharp$}}(abAb^2A^2, v_0) + 2\text{\textcircled{$\sharp$}}(abAbA, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(b^3AB, v_0) + 2\text{\textcircled{$\sharp$}}(abAB, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(abA^2, v_0) + 2\text{\textcircled{$\sharp$}}(b^2Ab^2A^2, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(ab^2AB, v_0) + 2\text{\textcircled{$\sharp$}}(b^3A^2, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(b^2Ab^2AB, v_0) + 2\text{\textcircled{$\sharp$}}(abAb^2AbA, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(ab^2AbA, v_0) + 2\text{\textcircled{$\sharp$}}(b^2Ab^2AbA, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(abAb^2AB, v_0) + 2\text{\textcircled{$\sharp$}}(b^3AbA, v_0) \\ \;+ 2\text{\textcircled{$\sharp$}}(ab^2A^2, v_0)\,] \end{array} .
$$

Appealing to the Extension Lemma 2.5.16 (pp. 25), we know that if $|v_0| \geqslant 5$,

$$
\begin{aligned}
\text{\textcircled{$\sharp$}}(b^3AB, v_0) + \text{\textcircled{$\sharp$}}(b^3A^2, v_0) &= \text{\textcircled{$\sharp$}}(b^3A, v_0) - \text{\textcircled{$\sharp$}}(b^3Ab, v_0) \\
\text{\textcircled{$\sharp$}}(abA^2, v_0) + \text{\textcircled{$\sharp$}}(abAB, v_0) &= \text{\textcircled{$\sharp$}}(abA, v_0) - \text{\textcircled{$\sharp$}}(abAb, v_0) \\
\text{\textcircled{$\sharp$}}(ab^2AB, v_0) + \text{\textcircled{$\sharp$}}(ab^2A^2, v_0) &= \text{\textcircled{$\sharp$}}(ab^2A, v_0) - \text{\textcircled{$\sharp$}}(ab^2Ab, v_0)
\end{aligned}
$$

Substituting into (53) and using the Extension Lemma 2.5.16 (pp. 25) repeatedly, we see that if $|v_0| \geqslant 5$,

$$
(54) \qquad \text{\textcircled{$\sharp$}}(b, v_0) = 2\text{\textcircled{$\sharp$}}(b^3A, v_0) + 2\text{\textcircled{$\sharp$}}(abA, v_0)
$$

Since b^3A is a-demarcator-immune, by Immunity Lemma 2.5.12 (pp. 24),

$$
(55) \qquad \text{\textcircled{$\sharp$}}(b, v_0) \geqslant 3\text{\textcircled{$\sharp$}}(b^3A, v_0) + 2\text{\textcircled{$\sharp$}}(abA, v_0)
$$

Combining (54) and (55), we see that $\text{\textcircled{$\sharp$}}(b^3A, v_0) = 0$. Thus,

$$
(56) \qquad \text{\textcircled{$\sharp$}}(b, v_0) = 2\text{\textcircled{$\sharp$}}(abA, v_0).
$$

Now appealing to Lemma 3.6.2 (pp. 50), we know that

$$
(57) \qquad \text{\textcircled{$\sharp$}}(a, v_0) = 2\text{\textcircled{$\sharp$}}(abA, v_0).
$$

It follows that every occurrence of $a^{\pm}$ occurs inside an a-blocker. This implies that $\phi_a(w) = w$, and hence that v_0 and v_1 coincide. Thus $w \notin \mathcal{S}_{abab}$.

This completes the case R3. The case L3 is completely symmetric and is omitted. Since each of the cases L1-L3 and R1-R3 were handled, we can conclude that T_{aba} cannot be realized as a proper simple chain in a level subgraph of Ω^* by any conjugacy class with length $\geqslant 5$.

Proposition 3.6.3 is proved. $\square$

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

COROLLARY 3.6.4 (Obstruction 5). *If σ in F_2 contains aba, ABA, bab or BAB as a proper subword, then $\mathcal{S}_\sigma$ contains no conjugacy class of length $\geqslant 5$.*

3.7. Obstruction Rewriting Rules, Part II. To find more forbidden subgraphs, we need some more sophisticated conservative graph rewriting rules that extend the rewriting rules presented in Definition 3.5.4 (pp. 47). We will use the following lemma as the foundation of these new extended rewriting rules.

LEMMA 3.7.1 (The $x^2\hat{x}^{-1}x\hat{x}^{-1}$ Relation Lemma). *For all $x \in X$ and $w \in \tilde{F}_2$,*

$$\begin{array}{rcl} \psi_{\hat{x}x\hat{x}^{-1}x^2}(w) & \approx_\Pi & w. \\ \psi_{x^2\hat{x}^{-1}x\hat{x}}(w) & \approx_\Pi & w. \end{array}$$

PROOF. By definition,

$$\psi_{\hat{x}x\hat{x}^{-1}x^2} \quad = \quad \psi_{\hat{x}} \circ \psi_x \circ \psi_{\hat{x}}^{-1} \circ \psi_x \circ \psi_x.$$

Thus,

$$\psi_{\hat{x}x\hat{x}^{-1}x^2} \quad = \quad \left\{ \begin{array}{lcl} x & \mapsto & x^{-1}\hat{x}x \\ \hat{x} & \mapsto & x^{-1}\hat{x}^{-1}x^{-1}\hat{x}x \end{array} \right.$$

But $x^{-1}\hat{x}x \sim_{\mathcal{J}} \hat{x}$ and $x^{-1}\hat{x}^{-1}x^{-1}\hat{x}x \sim_{\mathcal{J}} x^{-1}$. It follows that $\psi_{\hat{x}^{-1}x\hat{x}^{-1}x^2}(w) \approx_\Pi w$.

For the second assertion, note that by definition,

$$\psi_{x^2\hat{x}^{-1}x\hat{x}} \quad = \quad \psi_x \circ \psi_x \circ \psi_{\hat{x}}^{-1} \circ \psi_x \circ \psi_{\hat{x}}.$$

Thus,

$$\psi_{x^2\hat{x}^{-1}x\hat{x}} \quad = \quad \left\{ \begin{array}{lcl} x & \mapsto & \hat{x} \\ \hat{x} & \mapsto & \hat{x}^{-1}\hat{x}^{-1}x^{-1}xx \end{array} \right.$$

But $\hat{x}^{-1}\hat{x}^{-1}x^{-1}xx \sim_{\mathcal{J}} x^{-1}$. It follows that $\psi_{x^2\hat{x}^{-1}x\hat{x}^{-1}}(w) \approx_\Pi w$.

This completes the proof of the lemma. ☐

We define graph-rewriting schema ρ_x ($x \in X \cup X^{-1}$) which act on the hypothetical subgraphs of Ω^*. The schema ρ_x generalizes the previously defined schema τ_x. As before, for ρ_x to act on a hypothetical graph T, it must be suitably parametrized by a 3-tuple of vertices $p, v, q \in V[T]$.

DEFINITION 3.7.2. *Given a hypothetical subgraph T of Ω^*, fix $x \in X^{\pm}$), and let v, p, q be three distinct vertices in $V[T]$. If v, p, q satisfy the following conditions:*

- *v, p, q are all distinct,*
- *$\tilde{\psi}_x(p) = v$,*
- *$\tilde{\psi}_{\hat{x}}(p) = q$,*
- *$(v, \tilde{\psi}_{\hat{x}^{-1}}(v)) \notin E[T]$,*
- *$(q, \tilde{\psi}_{x^{-1}}(v))) \notin E[T]$.*

*then we say that the triple (v, p, q) are a ρ_x-**pivot** in T.*

DEFINITION 3.7.3. *Given a hypothetical subgraph T of Ω^*, fix $x \in X^\delta$ ($\delta = \pm 1$), and v, p, q a ρ_x-pivot in T. Define $M_x, M_{\hat{x}}$ be edge sets of cardinality $\leqslant 1$ as*

$$M_x \quad = \quad \left\{ \begin{array}{ll} \{(q, \tilde{\psi}_x(q)) \textit{ labelled by } x\ \} & \textit{if } (q, \tilde{\psi}_x(q)) \in E[T] \\ \emptyset & \textit{otherwise.} \end{array} \right.$$

$$M_{\hat{x}} \quad = \quad \left\{ \begin{array}{ll} \{(q, \tilde{\psi}_{\hat{x}}(q)) \textit{ labelled by } \hat{x}\ \} & \textit{if } (q, \tilde{\psi}_{\hat{x}}(q)) \in E[T] \\ \emptyset & \textit{otherwise.} \end{array} \right.$$

Let u be a new vertex representing $\tilde{\psi}_{\hat{x}^{-1}}(v)$, and hence not present in T. Define edge sets

$$\rho(M_x) = \begin{cases} \{(u, \tilde{\psi}_{\hat{x}^{-1}}(u)) \text{ labelled by } \hat{x}^{-1}\} & \text{if } S_x \neq \emptyset \\ \emptyset & \text{otherwise.} \end{cases}$$

$$\rho(M_{\hat{x}}) = \begin{cases} \{(u, \tilde{\psi}_{x^{-1}}(u)) \text{ labelled by } x^{-1}\} & \text{if } S_{\hat{x}} \neq \emptyset \\ \emptyset & \text{otherwise.} \end{cases}$$

Let $M(v,p,q) = M_x \cup M_{\hat{x}}$ and define $\rho M(v,p,q) = \rho(M_x) \cup \rho(M_{\hat{x}})$.

The graph-rewriting transformation $\rho_{x[p,v,q]}$ acts on T as follows:

Definition 3.7.4. *Given a hypothetical subgraph T of Ω^*, fix $x \in X^{\pm}$, and (v,p,q) a triple of vertices from $V[T]$. We define the graph $\rho_{x[v,p,q]}(T)$ as follows: If (v,p,q) are not a ρ_x-pivot, put $\rho_{x[v,p,q]}(T) = T$.*

Otherwise, let $\rho_{x[v,p,q]}(T)$ be given by:

$$V[\rho_{x[v,p,q]}(T)] = V[T] \cup \{u\} \backslash \{q\}$$

where u is a new vertex representing $\tilde{\psi}_{\hat{x}^{-1}}(v)$, and take

$$E[\rho_{x[v,p,q]}(T)] = E[T] \cup \rho M(v,p,q) \backslash M(v,p,q).$$

The operation of $\rho_{a[p,v,q]}$ and $\rho_{b[p,v,q]}$ on T (in is depicted in Figure 31. In the figure, boxed/outline edges are used to depict where edges are required to not be present.

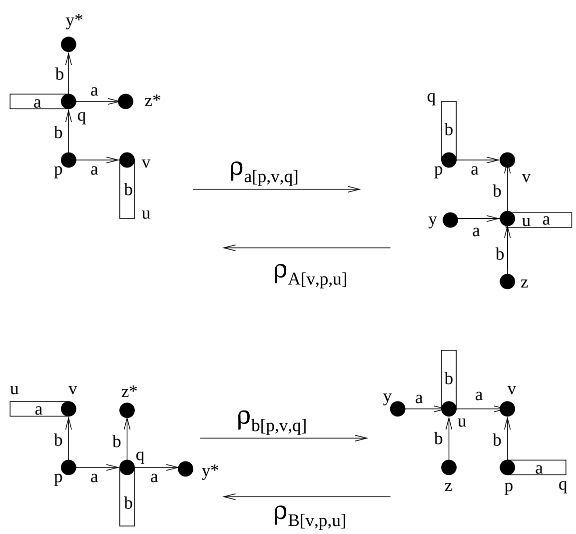

Figure 31: Extended graph rewriting rules ρ_x $(x \in X \cup X^{-1})$.

Remark 3.7.5. *It is straightforward to verify from Figure 31 that each of the operations is invertible. Specifically*

$$\rho^{-1}_{a[p,v,q]} = \rho_{A[v,p,u]}$$
$$\rho^{-1}_{b[p,v,q]} = \rho_{B[v,p,u]}.$$

The next proposition shows that the previously defined graph rewriting rules ρ_x ($x \in X \cup X^{-1}$) are conservative.

PROPOSITION 3.7.6 (Obstruction Rewriting Rule 2). *Given a tree $T = (V, E)$ and $v, p, q \in V$, $x \in X \cap X^{-1}$, the tree T is forbidden in Ω_n^* if and only if $\rho_{x[v,p,q]}(T)$ is forbidden in Ω_n^*.*

PROOF. If $\rho_{x[p,v,q]}(T) = T$ the statement is trivial. Otherwise, suppose that T is not forbidden in Ω_n^*. Let Ω_T be a minimal induced subgraph of Ω_n^* which contains an isomorphic copy of T. Without loss of generality, we may take Ω_T to be the graph induced by the isomorphic copy of T. Moreover, we identify T with the isomorphic spanning tree of Ω_T.

If $\tilde{\psi}_x(q) \in V[T]$ let us call this vertex z^*. If $\tilde{\psi}_{\hat{x}}(q) \in V[T]$ let us call this vertex y^*. If $\tilde{\psi}_{x^{-1}}(u) \in V[\rho_{x[p,v,q]}(T)]$ let us call this vertex y. If $\tilde{\psi}_{\hat{x}^{-1}}(u) \in V[\rho_{x[p,v,q]}(T)]$ let us call this vertex z. In other words, if these vertices are present they satisfy:

$$\begin{aligned} z^* &= \tilde{\psi}_x(q) \\ y^* &= \tilde{\psi}_{\hat{x}}(q) \\ y &= \tilde{\psi}_{x^{-1}}(u) \\ z &= \tilde{\psi}_{\hat{x}^{-1}}(u). \end{aligned}$$

Now, since $q = \psi_{\hat{x}x^{-1}\hat{x}}(u)$,

$$\begin{aligned} y^* &= \psi_{\hat{x}^2x^{-1}\hat{x}x}(y) \\ z^* &= \psi_{x\hat{x}x^{-1}\hat{x}^2}(z). \end{aligned}$$

By Lemmas 3.5.1 (pp. 46) and 3.7.1 (pp. 54), we know that

$$\begin{aligned} q &\approx_\Pi u \\ y &\approx_\Pi y^* \\ z &\approx_\Pi z^* \end{aligned}$$

It follows that three pairs of vertices v and q, y and y^*, z and z^* actually coincide in Ω_T.

Thus, adding edges $\rho M(v, p, q)$ and removing edges $M(v, p, q)$ merely gives rise to a different spanning tree of Ω_T. It follows that $\rho_{x[p,v,q]}(T)$ is a subgraph of Ω_T, and hence a subgraph of Ω_n^*.

To see the reverse, take $T' = \rho_{x[p,v,q]}(T)$. If T' occurs in Ω_n^* then by the previous argument, so must $\rho_{x^{-1}[v,p,u]}(T')$. But by Remark 3.7.5 (pp. 55),

$$\rho_{x^{-1}[v,p,u]} = \rho^{-1}_{x[p,v,q]}.$$

So it follows that T occurs in Ω_n^*. □

3.8. Still More Small-Scale Obstructions. We will make use of the graph-rewriting rules introduced in the previous sections to demonstrate more forbidden graphs.

3.8.1. *Obstruction 6: The Forbidden Graph T_{abAB}.* The vertices of Obstruction T_{abAB} are named $v_1 = \phi_a(v_0)$, $v_2 = \phi_b(v_1)$, $v_3 = \phi_A(v_2)$, $v_4 = \phi_B(v_3)$. The graph T_{abAB} is depicted in Figure 32.

We are now ready to prove the following proposition.

PROPOSITION 3.8.1. *The graph T_{abAB} cannot be realized as a proper simple chain inside Ω_n^* if $n \geqslant 5$.*

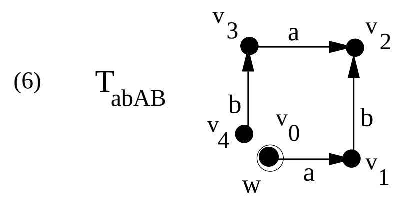

Figure 32: The Forbidden Graph T_{abAB}.

Proof. To prove this, we will consider the ways in which T_{abAB} might appear as a proper simple chain in an level subgraph of Ω^*. There are six ways:

L1. It occurs as an trailing subgraph of T_{aabAB}.
L2. It occurs as an trailing subgraph of T_{babAB}.
L3. It occurs as an trailing subgraph of T_{BabAB}.
R1. It occurs as a leading subgraph of T_{abABA}.
R2. It occurs as a leading subgraph of T_{abABB}.
R3. It occurs as a leading subgraph of T_{abABa}.

Each of these possibilities is depicted in Figure 33.

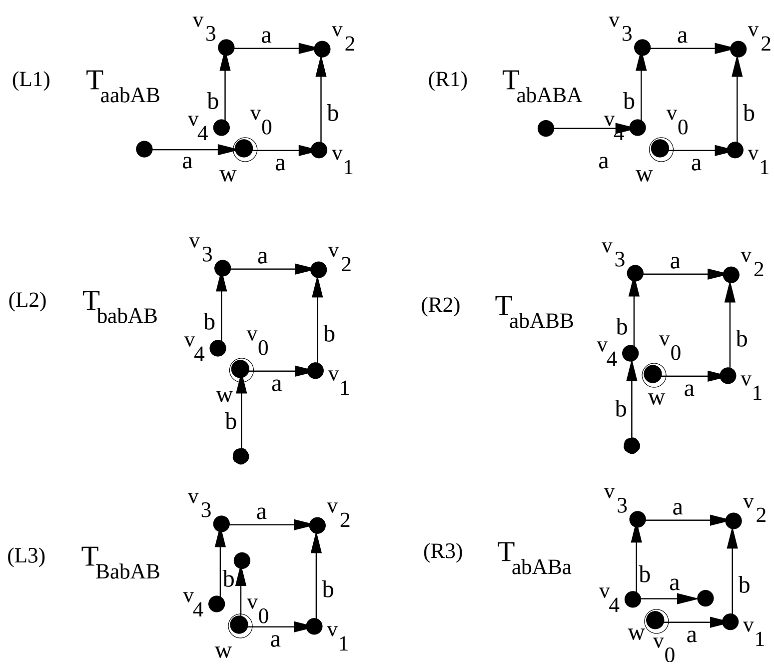

Figure 33: The six ways T_{abAB} might appear as a proper simple chain.

The impossibility of each of these configurations will now be proven, in turn. The cases L1, L2, R1, and R2 will be shown to be impossible using the Obstruction Rewriting Rule 1 presented in Proposition 3.5.6 (pp. 47) and by appealing to already-known forbidden graphs. The remaining cases L3 and R3 will be proved using the Obstruction Rewriting Rule 2 presented in Proposition 3.7.6 (pp. 56).

Case L1: The tree T_{aabAB} contains the tree T_{aab} as a subgraph. But T_{aab} cannot be realized, because of Corollary 3.4.6 (pp. 44). Hence, T_{aabAB} cannot be realized. The argument is illustrated in Figure 34.

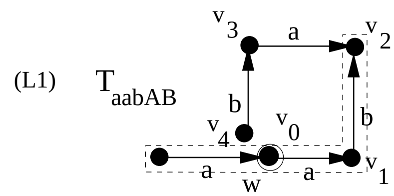

Figure 34: Why the L1 case (T_{aabAB}) is forbidden.

Case L2: The tree T_{babAB} contains the tree T_{bab} as a proper simple subchain. But T_{bab} cannot be realized, because of Corollary 3.6.4 (pp. 53). Hence, T_{babAB} cannot be realized. The argument is illustrated in Figure 35.

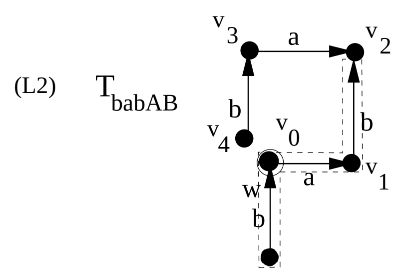

Figure 35: Why the L2 case (T_{babAB}) is forbidden.

Case R1: The tree T_{abABA} contains the tree T_{ABA} as a subgraph. But T_{ABA} cannot be realized, because of Corollary 3.6.4 (pp. 53). Hence, T_{abABA} cannot be realized. The argument is illustrated in Figure 36.

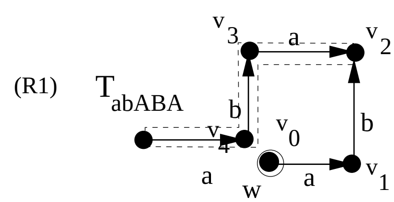

Figure 36: Why the R1 case (T_{abABA}) is forbidden.

Case R2: The tree T_{abABB} contains the tree T_{BBA} as a subgraph. But T_{BBA} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{abABB} cannot be realized. The argument is illustrated in Figure 37.

It remains to consider the L3 and R3 cases. To show that T_{abABa} and T_{BabAB}, we will use the extended graph rewriting rules τ_x and ρ_x described in Definitions 3.5.4 (pp. 47) and 3.7.4 (pp. 55) respectively.

Case L3: The tree T_{abABA} can be rewritten using two appliations of τ and one application of ρ to produce a graph which contains the tree T_{baa} as a subgraph. But T_{baa} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{abABA} cannot be realized. The argument is illustrated in Figure 38.

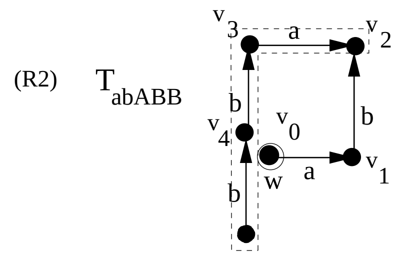

Figure 37: Why the R2 case (T_{abABA}) is forbidden.

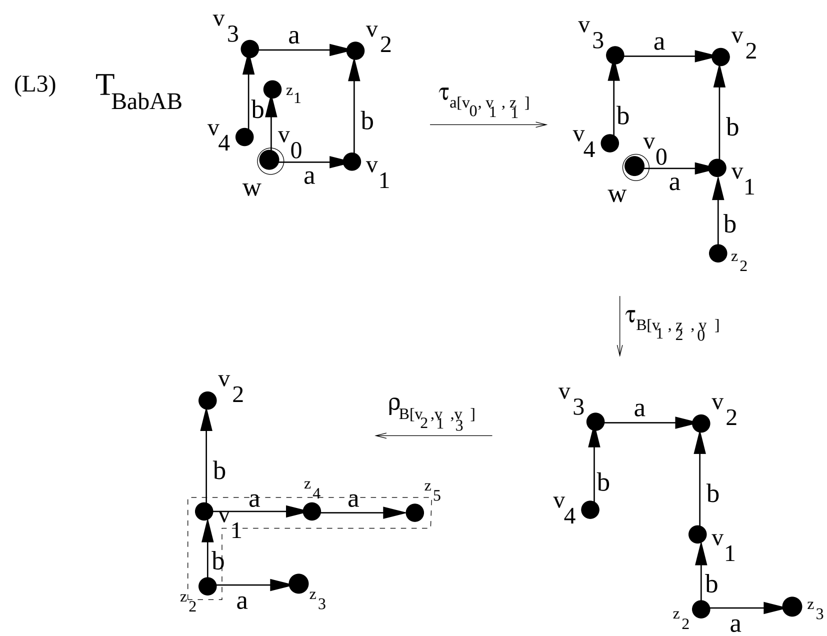

Figure 38: Why the L3 case (T_{abABA}) is forbidden.

Case R3: The tree T_{abABA} can be rewritten using two appliations of τ and one application of ρ to produce a graph which contains the tree T_{abb} as a subgraph. But T_{abb} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{abABA} cannot be realized. The argument is illustrated in Figure 39.

Since each of the cases L1-L3 and R1-R3 were handled, and the corresponding graphs shown to be forbidden, we can conclude that T_{abAB} cannot be realized as a proper simple chain in a level $n \geqslant 5$ subgraph of Ω^*.

Proposition 3.8.1 is proved. □

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

COROLLARY 3.8.2 (Obstruction 6). *If σ in F_2 contains $abAB$ or $baBA$ as a proper subword, then $\mathcal{S}_\sigma$ contains no conjugacy classes of length $\geqslant 5$.*

3.8.2. *Obstruction 7: The Forbidden Graph T_{ABab}.* The vertices of Obstruction T_{ABab} are named: $v_1 = \phi_A(v_0)$, $v_2 = \phi_B(v_1)$, $v_3 = \phi_a(v_2)$, $v_4 = \phi_b(v_3)$. The graph T_{ABab} is depicted in Figure 40.

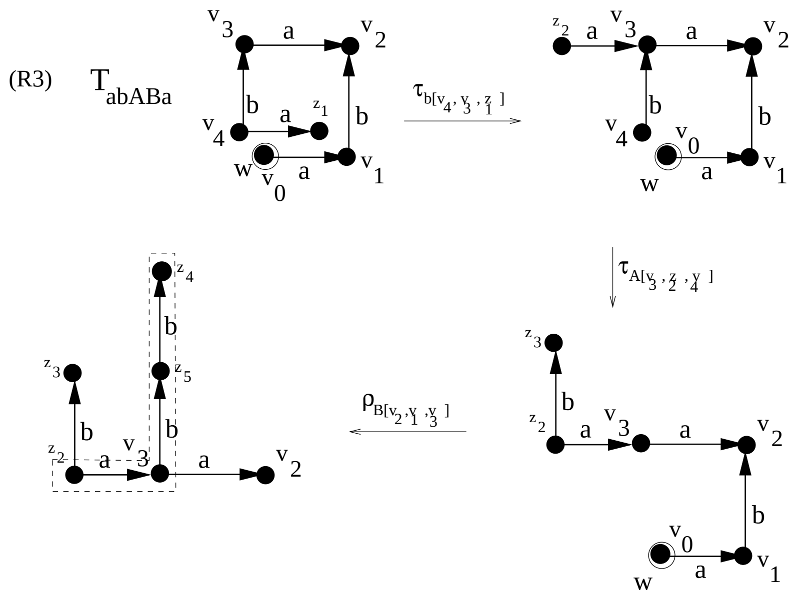

Figure 39: Why the R3 case (T_{babAB}) is forbidden.

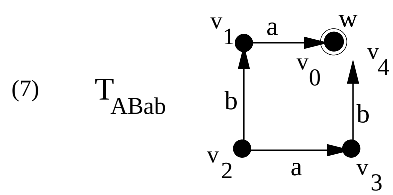

Figure 40: The Forbidden Graph T_{ABab}.

We are now ready to prove the following proposition.

PROPOSITION 3.8.3. *The graph T_{ABab} cannot be realized as a proper simple chain as a level subgraph in Ω_n^* for $n \geqslant 5$.*

PROOF. To prove this, we will consider the ways in which T_{ABab} might appear as a proper simple chain in an level subgraph of Ω^*. There are six ways:

L1. It occurs as an trailing subgraph of T_{aABab}.
L2. It occurs as an trailing subgraph of T_{bABab}.
L3. It occurs as an trailing subgraph of T_{BABab}.
R1. It occurs as a leading subgraph of T_{ABabA}.
R2. It occurs as a leading subgraph of T_{ABabB}.
R3. It occurs as a leading subgraph of T_{ABaba}.

Each of these possibilities is depicted in Figure 41.

The impossibility of each of these configurations will now be proven, in turn. The cases L1, L2, R1, and R2 will be shown to be impossible using the Obstruction Rewriting Rule 1 presented in Proposition 3.5.6 (pp. 47) and by appealing to already-known forbidden graphs. The remaining cases L3 and R3 will be proved using the Obstruction Rewriting Rule 2 presented in Proposition 3.7.6 (pp. 56).

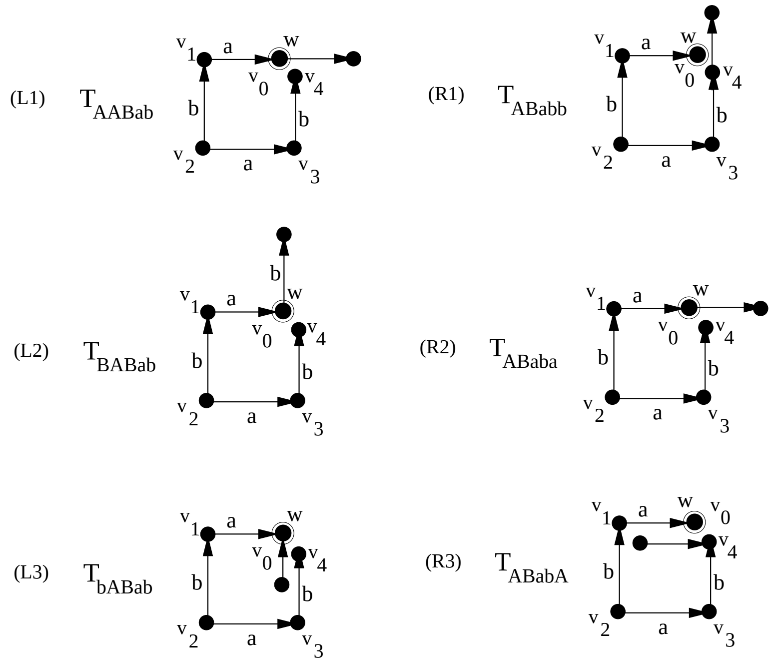

Figure 41: The six ways T_{ABab} might appear as a proper simple chain.

Case L1: The tree T_{aABab} contains the tree T_{aab} as a subgraph. But T_{aab} cannot be realized, because of Corollary 3.4.6 (pp. 44). Hence, T_{aABab} cannot be realized. The argument is illustrated in Figure 42.

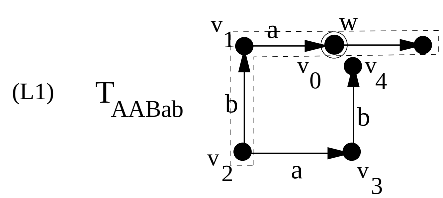

Figure 42: Why the L1 case (T_{aABab}) is forbidden.

Case L2: The tree T_{bABab} contains the tree T_{bab} as a proper simple subchain. But T_{bab} cannot be realized, because of Corollary 3.6.4 (pp. 53). Hence, T_{bABab} cannot be realized. The argument is illustrated in Figure 43.

Case R1: The tree T_{ABabA} contains the tree T_{ABA} as a subgraph. But T_{ABA} cannot be realized, because of Corollary 3.6.4 (pp. 53). Hence, T_{ABabA} cannot be realized. The argument is illustrated in Figure 44.

Case R2: The tree T_{ABabB} contains the tree T_{BBA} as a subgraph. But T_{BBA} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{ABabB} cannot be realized. The argument is illustrated in Figure 45.

It remains to consider the L3 and R3 cases. To show that T_{ABaba} and T_{BABab}, we will use the extended graph rewriting rules τ_x and ρ_x described in Definitions 3.5.4 (pp. 47) and 3.7.4 (pp. 55) respectively.

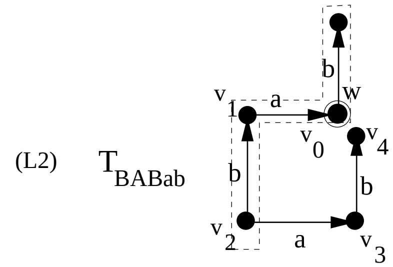

Figure 43: Why the L2 case (T_{bABab}) is forbidden.

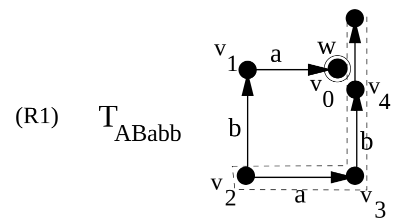

Figure 44: Why the R1 case (T_{ABabA}) is forbidden.

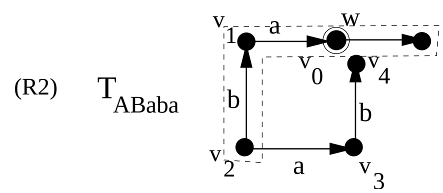

Figure 45: Why the R2 case (T_{ABabA}) is forbidden.

Case L3: The tree T_{ABabA} can be rewritten using two appliations of τ and one application of ρ to produce a graph which contains the tree T_{baa} as a subgraph. But T_{baa} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{ABabA} cannot be realized. The argument is illustrated in Figure 46.

Case R3: The tree T_{ABabA} can be rewritten using two appliations of τ and one application of ρ to produce a graph which contains the tree T_{abb} as a subgraph. But T_{abb} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{ABabA} cannot be realized. The argument is illustrated in Figure 47.

Since each of the cases L1-L3 and R1-R3 were handled, and the corresponding graphs shown to be forbidden, we can conclude that T_{ABab} cannot be realized as a proper simple chain in a level $n \geqslant 5$ subgraph of Ω^*.

Proposition 3.8.3 is proved. $\square$

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

COROLLARY 3.8.4 (Obstruction 7). *If σ in F_2 contains $ABab$ or $BAba$ as a proper subword, then $\mathcal{S}_\sigma$ contains no conjugacy classes of length $\geqslant 5$.*

3.8.3. *Obstruction 8: The Forbidden Graph T_{aBa}.* The vertices of Obstruction T_{aBa} are: $v_1 = \phi_a(v_0)$, $v_2 = \phi_B(v_1)$, $v_3 = \phi_a(v_2)$. The graph T_{aBa} is depicted in Figure 48.

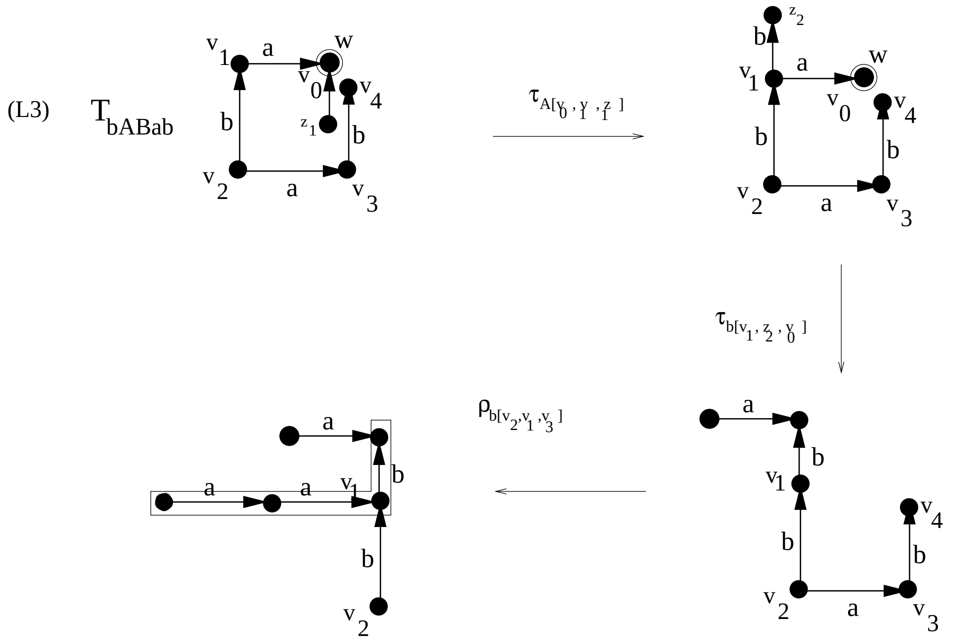

Figure 46: Why the L3 case (T_{ABabA}) is forbidden.

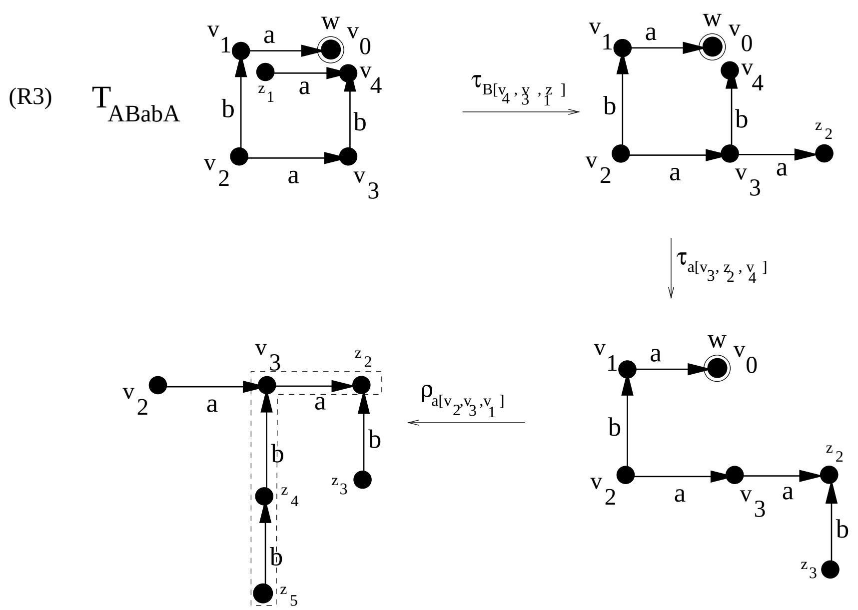

Figure 47: Why the R3 case (T_{bABab}) is forbidden.

We are now ready to prove the following proposition.

Proposition 3.8.5. *The graph T_{aBa} cannot be realized as a level subgraph of* Ω^*.

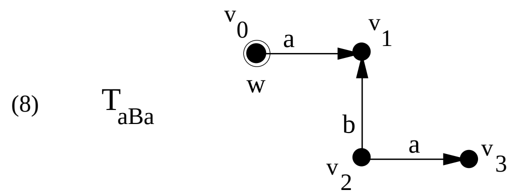

Figure 48: The Forbidden Graph T_{aBa}.

PROOF. Suppose T_{aBa} occurs as depicted in Figure 48. Then, by Lemma 3.5.1 (pp. 46), $v_0 \approx_\Pi v_3$. It follows that in Ω^* are the same vertex. Hence the subgraph T_{aBa} cannot occur be realized in Ω^* (at all), and hence certainly cannot be realized as a level subgraph of Ω^*. $\square$

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

COROLLARY 3.8.6 (Obstruction 8). $\mathcal{S}_{aBa} = \mathcal{S}_{AbA} = \mathcal{S}_{bAb} = \mathcal{S}_{BaB} = \emptyset$.

3.8.4. *Obstruction 9: The Forbidden Graph* T_{aaBA}. The vertices of Obstruction T_{aaBA} are named: $v_1 = \phi_a(v_0)$, $v_2 = \phi_a(v_1)$, $v_3 = \phi_B(v_2)$, $v_4 = \phi_A(v_3)$. The graph T_{aaBA} is depicted in Figure 49.

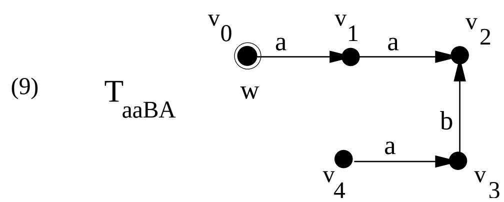

Figure 49: The Forbidden Graph T_{aaBA}.

We are now ready to prove the following proposition.

PROPOSITION 3.8.7. *T_{aaBA} cannot be realized as a level subgraph of Ω^*.*

PROOF. Suppose T_{aaBA} occurs as depicted in Figure 49. By applying the ρ graph rewriting rule once (see Figure 50), we can produce a graph which contains T_{abb}, which by Corollary 3.4.8 (pp. 46), cannot be realized as a level subgraph of Ω^*.

Since ρ is a conservative rewrite rule, we know that T_{aaBA} also cannot be realized as a level subgraph of Ω^*. $\square$

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

COROLLARY 3.8.8 (Obstruction 9). $\mathcal{S}_{aaBA} = \mathcal{S}_{abAA} = \mathcal{S}_{bbAB} = \mathcal{S}_{baBB} = \emptyset$.

3.8.5. *Obstruction 10: The Forbidden Graph* T_{AAba}. The vertices of Obstruction T_{AAba} are named: $v_1 = \phi_A(v_0)$, $v_2 = \phi_A(v_1)$, $v_3 = \phi_b(v_2)$, $v_4 = \phi_a(v_3)$. The graph T_{AAba} is depicted in Figure 51.

We are now ready to prove the following proposition.

PROPOSITION 3.8.9. *The graph T_{AAba} cannot be realized as a level subgraph of Ω^*.*

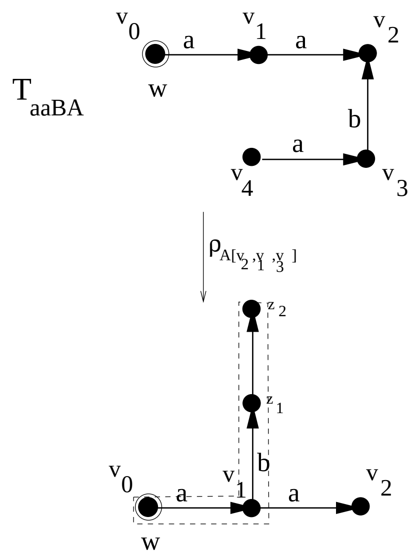

Figure 50: Rewriting T_{aaBA} using ρ to get a graph which is forbidden.

(10)

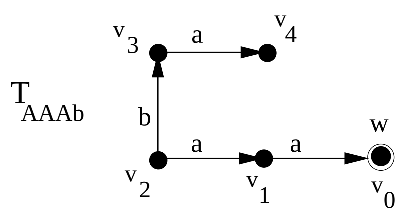

Figure 51: The Forbidden Graph T_{AAba}.

PROOF. Suppose T_{AAba} occurs as depicted in Figure 51. By applying the ρ graph rewriting rule once (see Figure 52), we can produce a graph which contains T_{baa}, which by Corollary 3.4.8 (pp. 46), cannot be realized as a level subgraph of Ω^*. Since ρ is a conservative rewrite rule, we know that T_{AAba} also cannot be realized as a level subgraph of Ω^*. $\square$

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

COROLLARY 3.8.10 (Obstruction 10). $\mathcal{S}_{AAba} = \mathcal{S}_{ABaa} = \mathcal{S}_{BBab} = \mathcal{S}_{BAbb} = \emptyset$.

3.8.6. *Obstruction 11: The Forbidden Graph T_{aBAb}.* The vertices of Obstruction T_{aBAb} are named: $v_1\phi_a(v_0)$, $v_2 = \phi_B(v_1)$, $v_3 = \phi_A(v_2)$, $v_4 = \phi_b(v_3)$. The graph T_{aBAb} is depicted in figure 53.

We are now ready to prove the following proposition.

PROPOSITION 3.8.11. *The graph T_{aBAb} cannot be realized as a proper simple chain as a level subgraph in Ω_n^* for $n \geqslant 5$.*

PROOF. To prove this, we will consider the ways in which T_{aBAb} might appear as a proper simple chain in an level subgraph of Ω^*. There are six ways:

L1. It occurs as an trailing subgraph of T_{aaBAb}.

L2. It occurs as an trailing subgraph of T_{BaBAb}.

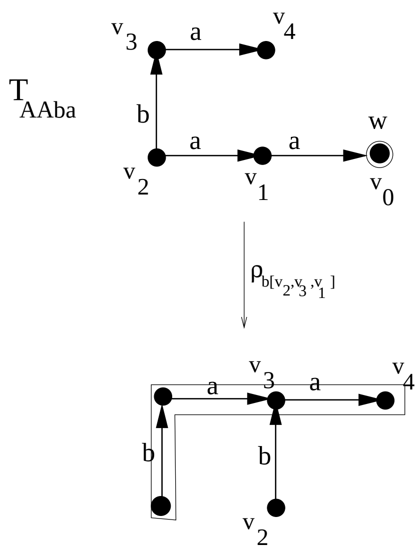

Figure 52: Rewriting T_{AAba} using ρ to get a graph which is forbidden.

(11)

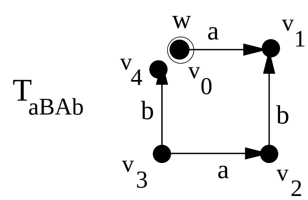

Figure 53: The Forbidden Graph T_{aBAb}.

L3. It occurs as an trailing subgraph of T_{baBAb}.
R1. It occurs as a leading subgraph of T_{aBAbA}.
R2. It occurs as a leading subgraph of T_{aBAba}.
R3. It occurs as a leading subgraph of T_{aBAbb}.

Each of these possibilities is depicted in Figure 54.

The impossibility of each of these configurations will now be proven, in turn, using the Obstruction Rewriting Rules 1 and 2 presented in Propositions 3.5.6 (pp. 47) and 3.7.6 (pp. 56), and by appealing to already-known forbidden graphs.

Case L1: Suppose the tree T_{aaBAb} is realized. Now T_{aaBAb} can be rewritten using a ρ transformation to yield a graph which contains the tree T_{aab} as a subgraph. But T_{aab} cannot be realized, because of Corollary 3.4.6 (pp. 44). Hence, T_{aaBAb} also cannot be realized. The argument is illustrated in Figure 55.

Case L2: The tree T_{BaBAb} contains the tree T_{BaB} as a proper simple subchain. But T_{BaB} cannot be realized, because of Corollary 3.8.6 (pp. 64). Hence, T_{BaBAb} cannot be realized. The argument is illustrated in Figure 56.

Case L3: Suppose the tree T_{baBAb} is realized. Now T_{baBAb} can be rewritten using a ρ transformation to yield a graph which contains the tree T_{baa} as a subgraph. But T_{baa} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{baBAb} also cannot be realized. The argument is illustrated in Figure 57.

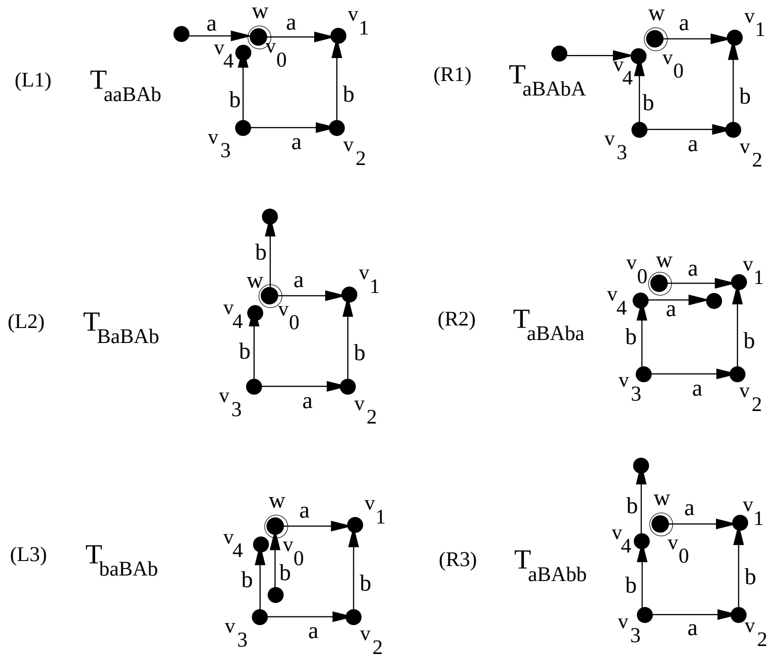

Figure 54: The six ways T_{aBAb} might appear as a proper simple chain.

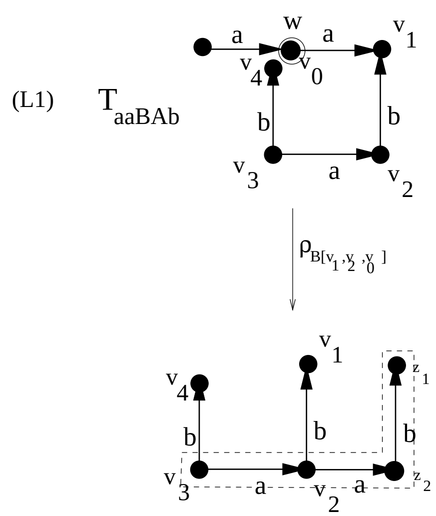

Figure 55: Why the L1 case (T_{aaBAb}) is forbidden.

Case R1: The tree T_{aBAbA} contains the tree T_{aBa} as a subgraph. But T_{aBa} cannot be realized, because of Corollary 3.8.6 (pp. 64). Hence, T_{aBAbA} also cannot be realized. The argument is illustrated in Figure 58.

Case R2: The tree T_{aBAba} contains the tree T_{BAba} as a subgraph. But T_{BAba} cannot be realized, because of Corollary 3.8.4 (pp. 62). Hence, T_{aBAba} also cannot be realized. The argument is illustrated in Figure 59.

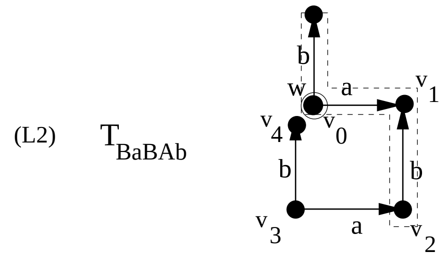

Figure 56: Why the L2 case (T_{baBAb}) is forbidden.

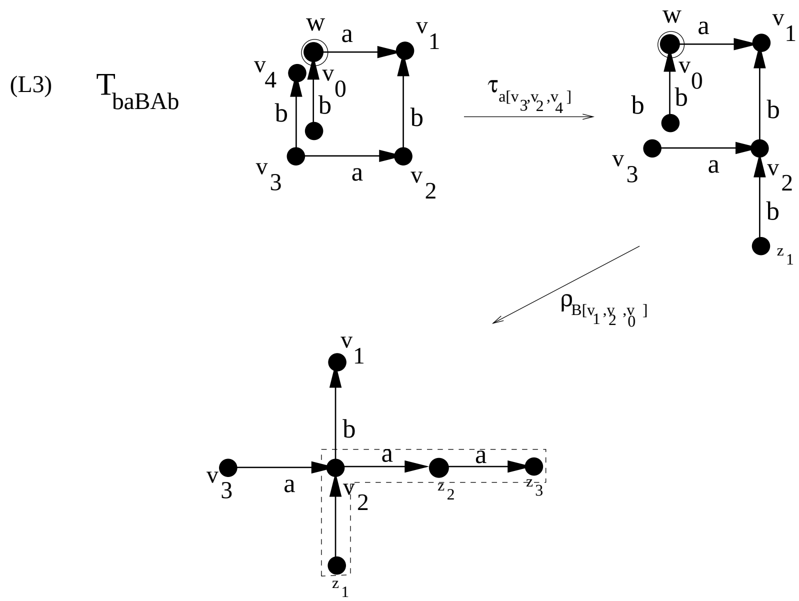

Figure 57: Why the L3 case (T_{aBAbA}) is forbidden.

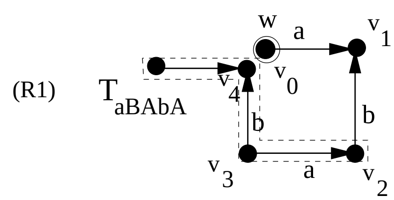

Figure 58: Why the R1 case (T_{aBAbA}) is forbidden.

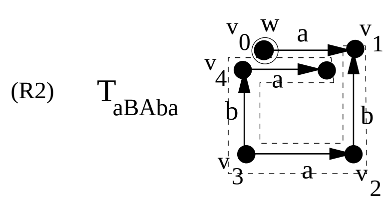

Figure 59: Why the R2 case (T_{aBAbA}) is forbidden.

Case R3: Suppose the tree T_{aBAbb} is realized. Now T_{aBAbb} can be rewritten using a ρ transformation to yield a graph which contains the tree T_{abb} as a subgraph. But T_{abb} cannot be realized, because of Corollary 3.4.8 (pp. 46). Hence, T_{aBAbb} also cannot be realized. The argument is illustrated in Figure 60.

(R3) T_{aBAbb}

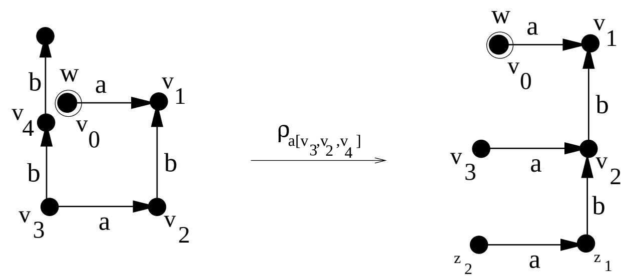

Figure 60: Why the R3 case (T_{baBAb}) is forbidden.

Since each of the cases L1-L3 and R1-R3 were handled, and the corresponding graphs shown to be forbidden, we can conclude that T_{aBAb} cannot be realized as a proper simple chain in a level subgraph of Ω_n^* for $n \geqslant 5$.

Proposition 3.8.11 is proved. $\square$

Appealing to the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36), the following is an immediate corollary of the previous proposition.

COROLLARY 3.8.12 (Obstruction 11). *If σ in F_2 contains aBAb, BabA, bABa or AbaB as a proper subword, then $\mathcal{S}_\sigma$ contains no conjugacy classes of length $\geqslant 5$.*

3.8.7. *Obstruction 12: The Forbidden Graph T_{aBBaa}.* The vertices of Obstruction T_{aBBaa} are denoted: $v_1 = \phi_a(v_0)$, $v_2 = \phi_B(v_1)$, $v_3 = \phi_B(v_2)$, $v_4 = \phi_a(v_3)$, $v_5 = \phi_a(v_4)$. The graph T_{aBAb} is depicted in Figure 61.

(12) T_{aBBaa}

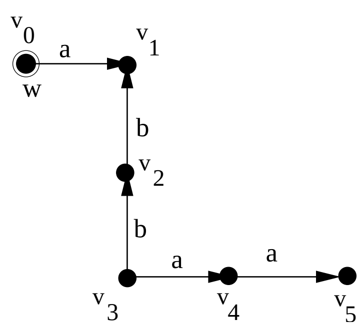

Figure 61: The Forbidden Graph T_{aBBaa}.

We are now ready to prove the following proposition.

PROPOSITION 3.8.13. *T_{aBBaa} cannot be realized in a level subgraph of Ω^*.*

PROOF. By applying a τ rewrite operation and then a ρ rewrite operation, we can transform T_{aBBaa} into a graph which contains T_{bba} as a subgraph. But T_{bba} is forbidden, by Corollary 3.4.6 (pp. 44). Thus, T_{aBBaa} is also forbidden. The argument is depicted in Figure 62. $\square$

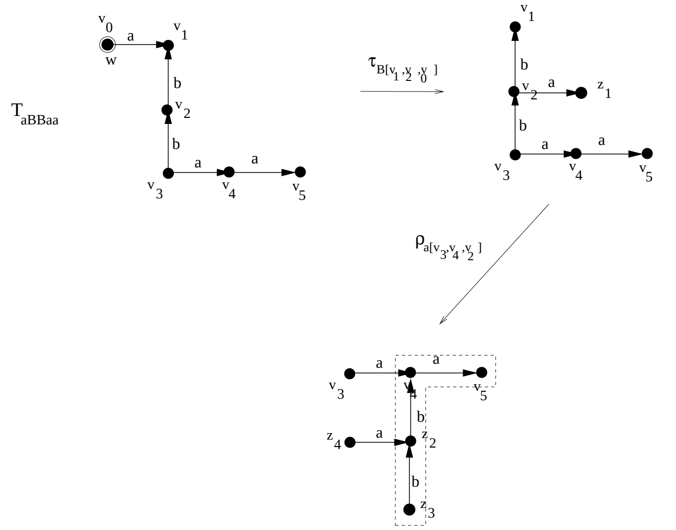

Figure 62: Why T_{aBBaa} is forbidden.

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

Corollary 3.8.14 (Obstruction 12). $\mathcal{S}_{aBBaa} = \mathcal{S}_{bAAbb} = \mathcal{S}_{AAbbA} = \mathcal{S}_{BBaaB} = \emptyset$.

3.8.8. *Obstruction 13: The Forbidden Graph T_{aaBBa}.* The vertices of Obstruction T_{aaBBa} are denoted: $v_1 = \phi_a(v_0)$, $v_2 = \phi_B(v_1)$, $v_3 = \phi_B(v_2)$, $v_4 = \phi_a(v_3)$, $v_5 = \phi_a(v_4)$. The graph T_{aBAb} is depicted in Figure 63.

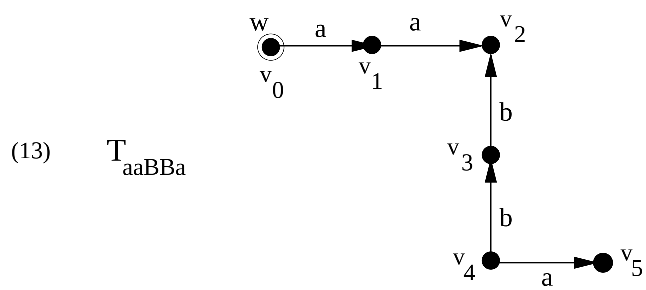

Figure 63: The Forbidden Graph T_{aaBBa}.

We are now ready to prove the following proposition.

Proposition 3.8.15. *The graph T_{aaBBa} cannot be realized in a level subgraph of Ω^*.*

Proof. By applying a τ rewrite operation and then a ρ rewrite operation, we can transform T_{aaBBa} into a graph which contains T_{abb} as a subgraph. But T_{abb} is forbidden, by Corollary 3.4.8 (pp. 46). Thus, T_{aaBBa} is also forbidden. The argument is depicted in Figure 64. $\square$

By the Chain Inversion Lemma 3.3.5 (pp. 36) and the Alphabet Symmetry Lemma 3.3.6 (pp. 36):

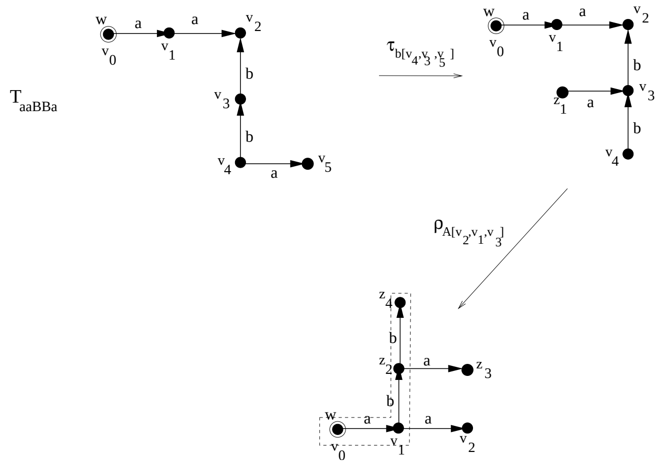

Figure 64: Why T_{aaBBa} is forbidden.

COROLLARY 3.8.16 (Obstruction 13). $\mathcal{S}_{aaBBa} = \mathcal{S}_{bbAAb} = \mathcal{S}_{AbbAA} = \mathcal{S}_{BaaBB} = \emptyset$.

3.9. Bounding the Size of Level Neighborhoods. In this section we put together all the information concerning small and large scale obstructions, and use this to control the structure of level neighborhoods in Ω^*.

The next theorem shows that is a sufficiently long conjugacy class w, if w is a σ-chain for a sufficiently long σ, then σ must be of the form x^k for some $x \in X \cup X^{-1}$.

THEOREM 3.9.1 (Length 5 Chains Theorem). *Suppose $w \in \tilde{F_2}$ has length $\geqslant 5$ and w is a σ-chain for some word $\sigma \in F_2$ where $|\sigma| = 5$. Then σ must be either aaaaa, bbbbb, AAAAA or BBBBB.*

PROOF. Let σ_0 be the initial prefix of σ, with $|\sigma_0| = 4$. We start by the determining possible values of σ_0. To do this, let us enumerate all possible $\sigma_0 \in F_2$ where $|\sigma_0| = 4$. If $w \in \tilde{F_2}$ is a σ-chain, then w is certainly also a σ_0-chain. The table below summarizes the possible and forbidden values for σ_0, and provides documentary evidence (Corollary number and page on which it was proved) to substantiate why particular values for σ_0 is disallowed.

Note: The values of σ_0 are organized into 4 tables, based on the first letter of σ_0. Within each table, the values of σ_0 are organized based on the 3-letter "base" prefix of σ_0.

len. 3 base	σ_0	Forbidden by Corollary #	σ_0	Forbidden by Corollary #	σ_0	Forbidden by Corollary #
aaa	$aaaa$	**POSSIBLE**	$aaab$	3.4.6, p.44	$aaaB$	3.4.2, p.41
aab	$aaba$	3.4.6, p.44	$aabb$	3.4.6, p.44	$aabA$	3.4.6, p.44
aaB	$aaBa$	3.8.6, p.64	$aaBA$	3.8.8, p.64	$aaBB$	**POSSIBLE**
aba	$abaa$	3.6.4, p.53	$abab$	3.6.4, p.53	$abaB$	3.6.4, p.53
abb	$abba$	3.4.8, p.46	$abbb$	3.4.8, p.46	$abbA$	3.4.8, p.46
abA	$abAb$	3.8.6, p.64	$abAA$	3.8.8, p.64	$abAB$	3.8.2, p.59
aBa	$aBaa$	3.8.6, p.64	$aBab$	3.8.6, p.64	$aBaB$	3.8.6, p.64
aBA	$aBAb$	3.8.12, p.69	$aBAA$	3.4.6, p.44	$aBAB$	3.6.4, p.53
aBB	$aBBa$	**POSSIBLE**	$aBBA$	3.4.8, p.46	$aBBB$	3.4.2, p.41
baa	$baaa$	3.4.8, p.46	$baab$	3.4.8, p.46	$baaB$	3.4.8, p.46
bab	$baba$	3.6.4, p.53	$babA$	3.6.4, p.53	$babb$	3.6.4, p.53
baB	$baBa$	3.8.6, p.64	$baBA$	3.8.2, p.59	$baBB$	3.8.8, p.64
bba	$bbaa$	3.4.6, p.44	$bbab$	3.4.6, p.44	$bbaB$	3.4.6, p.44
bbb	$bbba$	3.4.6, p.44	$bbbA$	3.4.2, p.41	$bbbb$	**POSSIBLE**
bbA	$bbAA$	**POSSIBLE**	$bbAb$	3.8.6, p.64	$bbAB$	3.8.8, p.64
bAb	$bAba$	3.8.6, p.64	$bAbA$	3.8.6, p.64	$bAbb$	3.8.6, p.64
bAA	$bAAA$	3.4.2, p.41	$bAAb$	**POSSIBLE**	$bAAB$	3.4.8, p.46
bAB	$bABa$	3.8.12, p.69	$bABA$	3.6.4, p.53	$bABB$	3.4.6, p.44
Aba	$Abaa$	3.4.8, p.46	$Abab$	3.6.4, p.53	$AbaB$	3.8.12, p.69
Abb	$Abba$	3.4.6, p.44	$AbbA$	**POSSIBLE**	$Abbb$	3.4.2, p.41
AbA	$AbAA$	3.8.6, p.64	$AbAb$	3.8.6, p.64	$AbAB$	3.8.6, p.64
AAb	$AAba$	3.8.10, p.65	$AAbA$	3.8.6, p.64	$AAbb$	**POSSIBLE**
AAA	$AAAA$	**POSSIBLE**	$AAAb$	3.4.4, p.43	$AAAB$	3.4.8, p.46
AAB	$AABa$	3.4.8, p.46	$AABA$	3.4.8, p.46	$AABB$	3.4.8, p.46
ABa	$ABaa$	3.8.10, p.65	$ABab$	3.8.4, p.62	$ABaB$	3.8.6, p.64
ABA	$ABAA$	3.6.4, p.53	$ABAb$	3.6.4, p.53	$ABAB$	3.6.4, p.53
ABB	$ABBa$	3.4.6, p.44	$ABBA$	3.4.6, p.44	$ABBB$	3.4.6, p.44
Baa	$Baaa$	3.4.4, p.43	$Baab$	3.4.6, p.44	$BaaB$	**POSSIBLE**
Bab	$Baba$	3.6.4, p.53	$BabA$	3.8.12, p.69	$Babb$	3.4.8, p.46
BaB	$BaBa$	3.8.6, p.64	$BaBA$	3.8.6, p.64	$BaBB$	3.8.6, p.64
Bab	$Baba$	3.6.4, p.53	$BabA$	3.8.12, p.69	$Babb$	3.4.8, p.46
BAA	$BAAA$	3.4.6, p.44	$BAAb$	3.4.6, p.44	$BAAB$	3.4.6, p.44
BAB	$BABa$	3.6.4, p.53	$BABA$	3.6.4, p.53	$BABB$	3.6.4, p.53
BBa	$BBaa$	**POSSIBLE**	$BBab$	3.8.10, p.65	$BBaB$	3.8.6, p.64
BBA	$BBAA$	3.4.8, p.46	$BBAb$	3.4.8, p.46	$BBAB$	3.4.8, p.46
BBB	$BBBa$	3.4.4, p.43	$BBBA$	3.4.8, p.46	$BBBB$	**POSSIBLE**

By examining the tables on the previous two pages, one may verify that σ_0 must be one of the following:

$$\{aaaa, bbbb, AAAA, BBBB, aaBB, bbAA,$$
$$AAbb, BBaa, aBBa, bAAb, AbbA, BaaB\}$$

Now we extend each of the above possible choices for σ_0 by one symbol to obtain possible choices for σ. Again, many of the possible choices are disallowed because they result in the presence of chains that we have already proved are not realized. The table below summarizes the possible choices that are allowed and forbidden for σ, and documents the evidence (Corollary number and page on which it was proved) to substantiate why particular values for σ are disallowed.

σ_0 base	σ	Forbidden by
$aaaa \to$	$aaaaa$	**POSSIBLE**
	$aaaab$	Cor. 3.4.6, p.44
	$aaaaB$	Cor. 3.4.2, p.41
$bbbb \to$	$bbbba$	Cor. 3.4.6, p.44
	$bbbbA$	Cor. 3.4.2, p.41
	$bbbbb$	**POSSIBLE**
$AAAA \to$	$AAAAA$	**POSSIBLE**
	$AAAAb$	Cor. 3.4.4, p.43
	$AAAAB$	Cor. 3.4.8, p.46
$BBBB \to$	$BBBBa$	Cor. 3.4.4, p.43
	$BBBBA$	Cor. 3.4.8, p.46
	$BBBBB$	**POSSIBLE**
$aaBB \to$	$aaBBa$	Cor. 3.8.16, p.71
	$aaBBA$	Cor. 3.4.8, p.46
	$aaBBB$	Cor. 3.4.2, p.41
$AAbb \to$	$AAbba$	Cor. 3.4.6, p.44
	$AAbbA$	Cor. 3.8.14, p.70
	$AAbbb$	Cor. 3.4.4, p.43
$BBaa \to$	$BBaaa$	Cor. 3.4.4, p.43
	$BBaab$	Cor. 3.4.6, p.44
	$BBaaB$	Cor. 3.8.14, p.70
$aBBa \to$	$aBBaa$	Cor. 3.8.14, p.70
	$aBBab$	Cor. 3.8.10, p.65
	$aBBaB$	Cor. 3.8.6, p.64
$bAAb \to$	$bAAba$	Cor. 3.8.10, p.65
	$bAAbA$	Cor. 3.8.6, p.64
	$bAAbb$	Cor. 3.8.14, p.70
$AbbA \to$	$AbbAA$	Cor. 3.8.16, p.71
	$AbbAb$	Cor. 3.8.6, p.64
	$AbbAB$	Cor. 3.8.8, p.64
$BaaB \to$	$BaaBa$	Cor. 3.8.6, p.64
	$BaaBA$	Cor. 3.8.8, p.64
	$BaaBB$	Cor. 3.8.16, p.71

Examination of the previous tables shows that σ must be one of the following either $aaaaa$, or $bbbbb$, or $AAAAA$, or $BBBBB$. This completes the proof of the theorem. □

THEOREM 3.9.2 (Level Structure Theorem). *For any u in $\tilde{F}_2$, if $|u| \geqslant 10$ and $|B^*(u)| > 4373$, then $B^*(u)$ is an x^k-chain having at most $|u| - 5$ vertices.*

PROOF. Denote the graph $B^*(u)$ as G. If G is an x^k-chain then the theorem follows immediately from Part (II) of Theorem 3.2.4 (pp. 31).

It remains to consider the case when G is not an x^k-chain. Every vertex in G has degree at most 4. It follows that the ball of radius 7 around a vertex contains at most $1 + 4 + 12 + 36 + 108 + 324 + 972 + 2916 = 4373$ distinct vertices. This means that a graph G with 4374 or more vertices is not contained in the ball of radius 7 about any vertex. Hence there must be two vertices u, v in G for which the shortest path between them in G has length $\geqslant 8$. Let us take u and v to be the vertices that are farthest apart in G, and denote the path between them as p.

Then the presence of p tells us that u is a σ-chain for some σ of length $|p| \geqslant 8 > 5$. By Theorem 3.9.1 (pp. 71), u is an $x_0^{|p|}$-chain for some $x_0 \in X^{\pm}$.

Let $\bar{p}$ be the maximal length x_0-chain in G which contains p as a subchain. Note that $\bar{p}$ is uniquely defined. Let $\bar{u}$, $\bar{v}$ be the start and end vertices of $\bar{p}$.

Suppose that there is a vertex w that is not on the path $\bar{p}$. Since G is connected, there is a shortest path q connecting w to $\bar{p}$. Suppose that q hits $\bar{p}$ at vertex z. We divide $\bar{p}$ into two parts: $\bar{p}_{\bar{u}}$ connects z to $\bar{u}$ by going along an initial segment of $\bar{p}$, and $\bar{p}_{\bar{v}}$ connects z to $\bar{v}$ by going along a final segment of $\bar{p}$.

We connect w with $\bar{u}$ via a path $p_{w,\bar{u}} = q \cdot \bar{p}_{\bar{u}}$. We connect w with $\bar{v}$ via a path $p_{w,\bar{v}} = q \cdot \bar{p}_{\bar{v}}$. Then either $p_{w,\bar{u}}$ or $p_{w,\bar{v}}$ must have length $\geqslant 1 + (|\bar{p}|/2) \geqslant 5$. If $|p_{w,\bar{u}}| \geqslant 5$, then by Theorem 3.9.1 (pp. 71) it follows that the labels on q are the same as the labels on $\bar{p}_{\bar{u}}$. But then w lies on $\bar{p}$, a contradiction. Likewise, if $|p_{w,\bar{v}}| \geqslant 5$, it follows that the labels on q are the same as the labels on $\bar{q}_{\bar{v}}$. But then w lies on $\bar{p}$, a contradiction. Thus, we conclude that there is no vertex w that is not on the path $\bar{p}$.

It follows that G is a chain graph labelled by x_0, and that $\bar{u}$ is a x_0^k-chain for some $x_0 \in X^{\pm}$ and $k \in \mathbb{N}$.

Part (II) of Theorem 3.2.4 (pp. 31) states that if $|u| \geqslant 10$ and u is an x^k-chain then k is at most $|u| - 5$. This completes the proof of the theorem. □

3.9.1. *Pulling Back to Ω and Γ.*

DEFINITION 3.9.3. *A graph G is called a* dense width w cylinder of length k *if it can be obtained by taking k disjoint complete graphs $H_1, \ldots, H_k$, where each H_i is a complete graph on w vertices ($i = 1, \ldots, k$), and connecting each vertex u in H_i to every vertex v in H_{i+1} ($i = 1, \ldots, k-1$). A graph G is called a* width w cylinder of length k *if it is a subgraph of a* dense width w cylinder of length k.

Main Theorem (Theorem 1.4.1) For any u in $\tilde{F}_2$, if $|u| \geqslant 10$ and $|B(u)| > 34984$, then $B(u)$ is width 8 cylinder of length at most $|u| - 5$ and has at most $8|u| - 40$ vertices.

PROOF. For any u in $\tilde{F}_2$, $|B(u)| \leqslant 8|B^*(u)|$, since $\approx_\Pi$ identifies at most 8 conjugacy classes to a single vertex in Ω^*. If $|B(u)| > 34984$ then $|B^*(u)| > 4373$ and by Theorem 3.9.2 (pp. 73) $B^*(u)$ is an x^k-chain having at most $|u| - 5$ vertices. But then $B(u)$ is width 8 cylinder of length at most $|u| - 5$, and so has at most $8|u| - 40$ vertices. □

The previous theorem can now be leveraged to give us information about Whitehead's graph Γ.

THEOREM 3.9.4 (Level Neighborhoods in Γ). *For any cyclically reduced word u in F_2, if $|u| > 4378$ then the connected level component of u in Γ has no more than $8|u|^2 - 40|u|$ vertices.*

PROOF. Suppose, towards contradiction, that the connected level component of u in Γ has more than $8|u|^2 - 40|u|$ vertices. Since $|u| \geqslant 4378$, it follows that $8|u|^2 - 40|u|$ exceeds $34984|u|$. So the connected level component of u in Γ has more than $34984|u|$ vertices. But the map from $\Gamma \to \Omega$ collapses at most $|u|$ vertices of Γ to a single conjugacy class in Ω, so it follows that the size of $B(\tilde{u})$ in Ω exceeds 34984. By Theorem 3.9.1 then $B(\tilde{u})$ is width 8 cylinder of length at most $|u| - 5$ and has at most $8|u| - 40$ vertices. Consider the pre-image of $B(\tilde{u})$ in Γ, restricted to vertices of length $|u|$. This subgraph of Γ is again just the connected level component of u in Γ. Since the map from $\Gamma \to \Omega$ collapses at most $|u|$ vertices of Γ to a single conjugacy class in Ω, the connected level component of u in Γ must have at most $8|u|^2 - 40|u|$ vertices. This is a contradiction. $\square$

4. Algorithmic Applications

In this section we explore the algorithmic applications of the the previously derived mathematical results. We begin with an overview of known algorithms for $\mathbf{CONJ}_n$ and $\mathbf{AUT\text{-}CONJ}_n$.

4.1. Algorithms for (Standard) Conjugacy. Recalling Definition 1.0.2 (pp. 2) where the decision problem $\mathbf{CONJ}_n(u,v)$ was first introduced. This involves taking as input an arbitrary pair $u, v \in F_n$ and determining if $\exists w \in F_n$ s.t. $w^{-1}uw = v$. The result that $\mathbf{CONJ}_n$ is decidable is, by now, folklore. The algorithm attributed to Greendlinger is as follows.

ALGORITHM $\mathbf{A}_{CONJ}$: Given u, v in a free group $F = F(X)$, $|X| = n$:

1. Construct two cycle graphs O_u and O_v having lengths $|u|$ and $|v|$ respectively. Write u clockwise on the edges of the first, and v along the second—these labelled graphs are called "circular words".
2. Now perform cyclic free reduction on these circular words, i.e. repeatedly contract all pairs of consecutive edges with labels x, x^{-1} or x, x^{-1} (for $x \in X$).
3. Upon termination of cyclic free reduction, check to see if the two circles graphs are equal graphs, as drawn.
   - If so, output 1. Halt.
   - Otherwise, output 0. Halt.

The next proposition shows that the above algorithm is correct:

PROPOSITION 4.1.1. *(Folklore)* $\mathbf{A}_{CONJ}$ *is a correct algorithm for* $\mathbf{CONJ}_n(u,v)$.

4.1.1. *Computational Complexity.* Let us consider the time complexity of the algorithm $\mathbf{A}_{CONJ}$ for standard conjugacy that was described in the previous section:

1. The construction of the cycle graphs can be done in time $O(|u|\log n + |v|\log n)$.
2. This reduction process terminates since the original words are of finite length, and their length strictly decreases at each reduction step. Cyclic free reduction $u \rightsquigarrow \tilde{u}$ and $v \rightsquigarrow \tilde{v}$ can be achieved in time $O(|u|\log n + |v|\log n)$.
3. To check to see if $\tilde{u}$ and $\tilde{v}$ are equal as graphs first check that $|\tilde{u}| = |\tilde{v}|$. If so, then fix a starting vertex p on $O_{\tilde{u}}$ and vary the start vertex q on $O_{\tilde{v}}$. Determine if the word read clockwise in $O_{\tilde{u}}$ starting from p is the same as the word read clockwise in $O_{\tilde{v}}$ from q. This can be done in time $O(|\tilde{u}||\tilde{v}|\log n)$. Since $|\tilde{u}| \leqslant |u|$ and $|\tilde{v}| \leqslant |v|$ and $|\tilde{u}| = |\tilde{v}|$ it follows that $|\tilde{u}|$ and $|\tilde{v}|$ are both less than $\min(|u|,|v|)$. So this stage of the algorithm works in time $O(\min(|u|^2,|v|^2)\log n)$.

The above analysis yields:

THEOREM 4.1.2. $\mathbf{A}_{CONJ}$ *decides* $\mathbf{CONJ}_n$ *in time* $O(\min(|u|^2,|v|^2)\log n)$ *time.*

COROLLARY 4.1.3. $\mathbf{A}_{CONJ}$ *decides* $\mathbf{CONJ}_2$ *in time* $O(\min(|u|^2,|v|^2))$ *time.*

4.2. Algorithms for Automorphic Conjugacy. In the previous section, we considered algorithms for (standard) conjugacy. In contrast, here we consider $\mathbf{AUT\text{-}CONJ}_n(u, v)$, the decision problem for automorphic conjugacy that was introduced in Definition 1.0.1 (pp. 2). In this problem, we are given $u, v \in F_n$ and are to detemine if there is an automorphism ϕ in $Aut(F_n)$ for which $\phi(u) = v$.

4.2.1. *Whitehead's Algorithm.* In 1936, J. H. C. Whitehead proved [27, 28] that $\mathbf{AUT\text{-}CONJ}_n$ is decidable. We describe the algorithm below following the illustration in Figure 65.

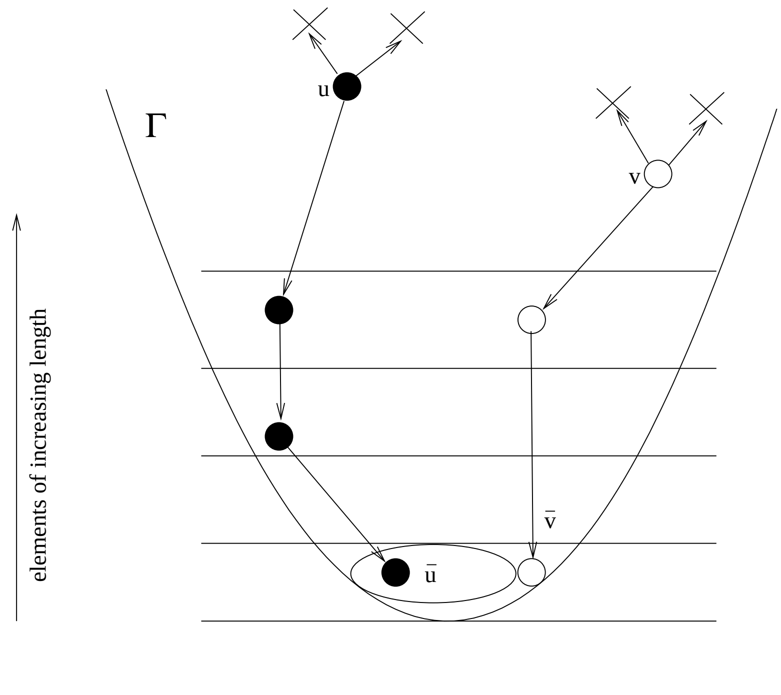

Figure 65: The operation of Whitehead's Algorithm.

ALGORITHM $\mathbf{A}_{AUT-CONJ}$: Given u, v in a free group $F = F(X)$, $|X| = n$:

1. Compute $\tilde{u}$ and $\tilde{v}$ and represent them as cyclic words.
2. Apply Whitehead automorphisms greedily to reduce the lengths of $\tilde{u}$ and $\tilde{v}$ until no further reduction in length is possible. Denote the minimal length forms as $\bar{u}$ and $\bar{v}$.
3. Check if $|\bar{u}| = |\bar{v}|$.
   - 3A. If not, output 0. Halt.
   - 3B. Otherwise, let $n = |\bar{u}| = |\bar{v}|$. Determine if $\bar{u}, \bar{v}$ are in the same connected component of Γ, as follows.
     - 3B-1. Construct via breadth-first search $\Gamma_{\bar{u}}$, the connected component of Γ which contains u in which all vertices have length $|\bar{u}|$.
     - 3B-2. Test whether $\bar{v}$ is in $\Gamma_{\bar{u}}$.
       - If so, output 1. Halt.
       - Otherwise, output 0. Halt.

The proof that Whitehead's Algorithm is correct is fairly lengthy and involved. For details, the reader is referred to the original papers of Whitehead [27, 28]m the classical texts of Magnus, Karrass and Solitar [14] and Lyndon and Schupp [13]. Another approach for correctness is show that the Whitehead automorphisms form a confluent rewriting system for automorphically conjugate elements of a free group, and then to appeal to the Diamond Property.

Here we will only address the time complexity of the original Whitehead's Algorithm. We will briefly consider improvements in Whitehead's Algorithm due Miasnikov and Shpilrain, and explore new improvements based on the mathematical results of previous sections.

4.2.2. *Computational Complexity.* Let us consider the time complexity of the algorithm $\mathbf{A}_{AUT-CONJ}$ for automorphic conjugacy that was described in the previous section:

1. The task of computing the cyclically reduced forms of u and v (i.e. $\tilde{u}$ and $\tilde{v}$ respectively) and representing them as cyclic words can be done in time $O(|u| \log n + |v| \log n)$.
2. There are $|W_n| = 2^n$ Whitehead automorphisms to consider when greedily reducing the lengths of $\tilde{u}$ and $\tilde{v}$ to find minimal length representatives $\bar{u}$ and $\bar{v}$. Consider the process applied to u: Then, each application of a Whitehead automorphism takes time $O(|u| \log n)$. Each application of a Whitehead automorphism is amortized against a non-trivial reduction in the length of u. It follows at most $|u|$ reductions can take place, and hence that it takes at most $O(2^n |u|^2 \log n)$ time to compute $\bar{u}$. Thus, this step of the procedure takes at most $O(2^n |u|^2 \log n + 2^n |v|^2 \log n)$ time.
3. We construct via breadth-first search from $\bar{u}$, the connected subgraph of the Whitehead graph induced by the vertices of the same length as $|\bar{u}|$. The edge density in this component is $O(2^n)$, since every vertex has at most $|W_n| = O(2^n)$ edges incident to it. If one can obtain a bound $f(m)$ on the size of the biggest level connected component in Γ_m, then (since $|\bar{u}| \leqslant |u|$ and $|\bar{v}| \leqslant |v|$) the running time of this step will become $O(2^n f(\min |u|, |v|))$. In the classical analysis, elementary combinatorics tells us that $f(m) = O(n^m)$ since this bounds the number of elements in F_n having length equal to m. Using different techniques, Miasnikov and Shpilrain determined a bounding function $f(m) = m^4$ *for the case when* $n = 2$. Now, it follows from our result Theorem 3.9.4 (pp. 74) that in the case of $n = 2$, one actually has a quadratic bounding function on the size of the biggest level connected component at level m in Γ. Thus, in the rank 2 case, this stage requires only $O(\min(|u|^2, |v|^2))$ time.

While the classical analysis yields:

THEOREM 4.2.1. $\mathbf{A}_{AUT-CONJ}$ *decides* $\mathbf{AUT\text{-}CONJ}_n$ *in time* $O(2^n n^{\min(|u|,|v|)})$.

for the case of F_2 we have shown:

COROLLARY 4.2.2. $\mathbf{A}_{AUT-CONJ}$ *decides* $\mathbf{AUT\text{-}CONJ}_2$ *in time* $O(\min(|u|^2, |v|^2))$.

Computational experiments by C. Sims and A. D. Miasnikov have long indicated that in practice, however, Whitehead's Algorithm runs in polynomial time. Until recently there was no formal analytic argument to explain this empirical fact. Then, in 1999 (to appear) Miasnikov and Shpilrain [17] obtained a polynomial bound for the running time of Whitehead's Algorithm in the case of F_2, and they hypothesized that Whitehead's Algorithm is polynomial for free groups of all ranks. Here, we have shown that $\mathbf{A}_{AUT-CONJ}$ decides **AUT-CONJ**$_2$ in the same time complexity as A_{CONJ} decides **CONJ**$_2$, namely in quadratic time (see Corollary 4.1.3 pp. 75). Since any algorithm for **AUT-CONJ**$_2$ is also an algorithm for **CONJ**$_2$, it is improbable that we will be able to find a faster algorithm to test automorphic conjugacy in F_2.

We remark that the structural description of the orbits of F_2 under the action of $Aut(F_2)$ makes it possible to devise algorithms which operate altogether differently from and surpass Whitehead's algorithm. This will be the subject of a separate publication.

5. Computational Tools

5.1. JIGGLE. The 1999 dissertation, "A Numerical Optimization Approach to General Graph Drawing" of Daniel Tunkelang (Carnegie Mellon University) [26] provides an excellent overview of graph drawing, an explanation and implementation of new methods to solve the general graph drawing problem. This implementation was extended by the author (as described later) to generate a visualization of the structure of $Aut(F2)$.

Whereas most prior research considered only special types of graphs, Tunkelang's methods address the general graph drawing problem. In addition to its generality, the approach presented in the dissertation improved on several aspects of graph drawing, including an improvement in performance over prior algorithms and an improvement in the quality of the resulting drawings. Tunkelang considers three basic aspects of graph drawing: drawing conventions, constraints, and aesthetics. Drawing conventions are basic specifications such as what space is being used for the drawing area (usually the plane $\mathbb{R}^2$) or what type of lines are being used between vertices.

JIGGLE constructs a physical model of a graph by considering vertices as massive positively charged spheres and reifying edge relationships as physical springs. The evolution of the system is then simulated under the influence of Hooke's Law, Coulomb's Law and Newton's Law (as well as other variants, See Figure 66). Tunkelang's principal contribution involves novel techniques for computing efficient approximations of the simulation, without which the approach would be computationally infeasible for all but the smallest graphs.

To avoid local minima in the energy surface, Tunkelang starts off by placing the vertices in a high-dimensional space, and the dimensions are gradually "squashed out", by applying gravity laterally along successive dimensions. Once the layout becomes ϵ thin in dimension k, the simulation projects its state into $k-1$ dimensional Eucliean space, and proceeds to apply gravity along dimension $k-1$. The process continues in this manner until a 2D layout is attained. The reader is referred to [25] for further details.

5.2. Extended Implementation for $Aut(F2)$. Much of the intuition behind this mathematical research was derived by playing with JIGGLE. The author has developed a set of tools for exploring the structure of $Aut(F)$ and other combinatorial objects related to free groups. This software is structured in three "modules":

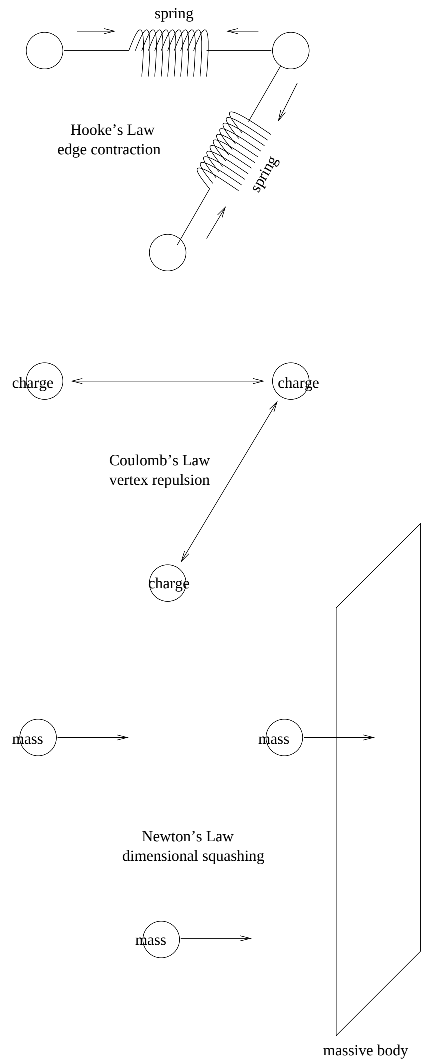

Figure 66: The three forces which act during JIGGLE graph layout.

- Combinatorial "back end", responsible for enumerating connected components and level neighborhoods in Γ, Ω and Ω^*.
- Visualization "front end", responsible for manipulating JIGGLE parameters and displaying the resulting drawings.
- Statistical "analyzer", responsible for detecting statistical patterns in terms of subgraph frequencies, symmetries and structure. This module operates both on the combinatorial graphs and on their drawings.

The role of this software as part of the mathematical laboratory cannot be overemphasized. The visualization front end was used to firstexamine long x^k chains and surrounding structures. The statistical analyzer was used to determine candidates for forbidden subgraphs (and to search for occurrences of counterexamples to hypothesized forbidden graphs).

And while these tools did not provide much information about *how to prove* the results of the previous sections, they served as a computational laboratory in which to make systematic observations and experiment upon the naturally occurring objects of interest: Γ, Ω and Ω^*. More details of this software will be made available in a separate publication.

6. Conclusions and Future Work

This work considered the structure of Whitehead's Graph Γ, introduced in 1936 by J. H. C. Whitehead as away to quantify relationships between the natural length function $|\ |$ of F_2, and the action of $Aut(F_2)$ on F_2.

Specifically, we provided a structural description of the level sections in Γ, by studying two natural quotients, Ω and Ω^*. These quotients required us to shift from the study of F_2 to the study of conjugacy classes in F_2. To facilitate the investigation, we introduced the notion of combinatorial equations on conjugacy classes. We provided techniques for mapping hypothesized subgraphs of Ω into systems of combinatorial equations—this mapping process produced systems of constraints with property that the infeasibility of the constraints implies non-occurrence of the subgraphs. By proving that there is a set of such forbidden subgraphs (and enlarging this set using rewriting rules) we were able to show that all sufficiently large level sections in Ω^* must be chains. Finally, we applied these results to improve the analysis of the classical Whitehead's Algorithm, showing that it tests automorphic conjugacy in F_2 in at most quadratic-time.

Below we list some questions that are presently under investigation, and are natural extensions of this work:

(1) **Find non-trivial lower bounds on the size of level neighborhoods for the automorphism graph of F_n, $n \geqslant 2$**
(2) **Is there a finite set of "obstruction" for level neighborhoods of the automorphism graph of F_n, $n \geqslant 3$**
(3) **Is there a recursive procedure to find "obstructions" for level neighborhoods of the automorphism graph of F_n, $n \geqslant 3$?**
(4) **Is there a recursive procedure to transform an "obstruction" graph into a system of combinatorial equations?**
(5) **Is there a recursive procedure to determine if a system of combinatorial equations has a solution? Describe the solution sets of such combinatorial equations.**

6.1. Acknowledgements. I would like thank Alexei Miasnikov for his help and comments in the development of this work. I am grateful to Janos Pach and Joseph Roitberg for reviewing this work, and my colleagues Toshiaki Jutsikawa, Denis Serbin, Delaram Kahrobaei for many spirited discussions. I am beholden to my friend and colleague Daniel Tunkelang for his expertise in automatic graph layout algorithms and for allowing me to use his doctoral thesis software to compute initial visual representations of the structure of the automorphism graph. I acknowledge invaluable support from the City University of New York. also thank Henry Dardy and Basil Decina of the U.S. Naval Research Laboratory in Washington D.C., and John Urban of Advanced Engineering Sciences at ITT Industries, for providing funding and resources toward my ongoing research in computer science and mathematics.

References

[1] G. Baumslag. *Topics in Combinatorial Group Theory [Chapter VII: Groups acting on Trees]*. Birkhauser, 1993.

[2] G. Baumslag, A. Myasnikov, and V. Remeslennikov. Algebraic geometry over groups. i: Algebraic sets and ideal theory. In *Journal of Algebra, vol.219 (1999), pp. 16-79*, 1999.

[3] D. J. Collins and H.Zieschang. *Combinatorial Group Theory and Fundamental Groups*. Springer Verlag, 1998.

[4] M. Coornaert, T. Delzant, and A. Papdopoulos. *Geometrie et theorie des groupes: Les groupes hyperboliques de Gromov (Lecture Notes in Matrhematics No. 1441)*. Springer Verlag, 1990.

[5] M. Coornaert and A. Papdopoulos. *Symbolic Dynamics and Hyperbolic Groupss (Lecture Notes in Matrhematics No. 1539)*. Springer Verlag, 1993.

[6] W. Dicks and M. Dunwoody. *Groups acting on Graphs (Cambridge Studies on Advanced Mathematics No. 17)*. Cambridge University Press, 1989.

[7] A. Duncan, N. D. Gilbert, and J. Howie. *Combinatorial and Geometric Group Theory (London Mathematical Society Lecture Notes No. 204) [A. Razborov article: On Systems of Equations in Free Groups]*. Cambridge University Press, 1995.

[8] D. Epstein. *Word Processing in Groups*. Jones and Bartlett Publishers, 1992.

[9] D. Evans. *Model Theory of Groups and Automorphism Groups (London Mathematical Society Lecture Notes No. 244) [I. Chiswell article: Non-standard Free Groups]*. Cambridge University Press, 1997.

[10] E. Ghys, A. Haefliger, and A. Verjovsky. *Group Theory from a Geometrical Viewpoint [J. Alonso et al. article: Notes on Hyperbolic Groups]*. World Scientific, 1991.

[11] R. I. Grigorchuk and P. F. Kurchanov. *Some Questions of Group Theory Related to Geometry*. Springer Verlag, 1998.

[12] M. Lothaire. *Combinatorics on Words*. Cambridge Mathematical Library, 1983.

[13] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*. Springer-Verlag, New York, 1977.

[14] W. Magnus, A. Karrass, and D. Solitar. *Combinatorial Group Theory*. Dover Publications, 1976.

[15] G. Makanin. Equations in a free group. In *Math. USSR Izvestiya, vol.21 (1983) no. 3, pp. 483-586*, 1983.

[16] G. Makanin. Decidability of the universal and positive theories of a free group. In *Math. USSR Izvestiya, vol.25 (1985) no. 1, pp. 75-88*, 1985.

[17] A. Miasnikov and V. Shpilrain. Automorphic orbits in free groups. *J. Algebra*, to appear.

[18] A. Myasnikov and O. Kharlampovich. Irreducible affine varieties over a free group. i: Irreducibility of quadratic equations and the nullstellensatz. In *Journal of Algebra, vol.200 No. 2 (1998), pp. 472-516*, 1998.

[19] A. Myasnikov and O. Kharlampovich. Irreducible affine varieties over a free group. ii: Systems in triangular quasi-quadratic form and description of residually free groups. In *Journal of Algebra, vol.200 No. 2 (1998), pp. 517-570*, 1998.

[20] A. A. Razborov. On systems of equations in a free group. In *Math. USSR Izvestiya, vol.25 (1985) no. 1, pp. 115-162*, 1985.

[21] J. S. Rose. *A Course on Group Theory.* Dover Publications, 1978.

[22] W. R. Scott. *Group Theory.* Dover Publications, 1987.

[23] J.-P. Serre. *Trees.* Springer-Verlag, 1980.

[24] C. Sims. *Computation with Finitely Presented Groups (Encyclopedia of Mathematics and its Applications, volume 48).* Cambridge University Press, 194.

[25] D. Tunkelang. Jiggle: Java interactive general graph layout environment. *Proc. 6th Int. Symp. Graph Drawing, volume 1547 of LNC*, pages 413–422, 1998.

[26] D. Tunkelang. *A Numerical Optimization Approach to General Graph Drawing.* PhD thesis, School of Computer Science, Carnegie Mellon University, 1999.

[27] J. H. C. Whitehead. On certain sets of elements in a free group. *Proc. London Math. Soc. II*, Ser. 41:48–56, 1936.

[28] J. H. C. Whitehead. On equivalent sets of elements in a free group. *Ann. of Math. II*, Ser. 37:782–800, 1936.

Index