ROYAL HOLLOWAY, UNIVERSITY OF LONDON

DOCTORAL THESIS

Finiteness and Cubulative Properties of Algebraic Bieri-Strebel Groups

Author:
Lewis MOLYNEUX

Supervisor:
Prof. Brita NUCINKIS

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy

in the

Department of Mathematics

January 19, 2026

Declaration of Authorship

I, Lewis MOLYNEUX, declare that this thesis titled, “Finiteness and Cubulative Properties of Algebraic Bieri-Strebel Groups” and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

“Groups really are fundamental. . . What could be more universal than symmetry?”

Grant Sanderson, "Group theory, abstraction, and the 196,883-dimensional monster"

ROYAL HOLLOWAY, UNIVERSITY OF LONDON

Abstract

School of Engineering, Physical and Mathematical Sciences

Department of Mathematics

Doctor of Philosophy

Finiteness and Cubulative Properties of Algebraic Bieri-Strebel Groups

by Lewis MOLYNEUX

Thompson Groups are of perennial interest in the field of Geometric Group Theory, producing unique and elegant results in many areas of study. These properties often extend to the much larger class of Bieri-Strebel groups. In particular, finiteness properties are a point of interest among all Thompson-Like groups.

This thesis will examine finiteness properties among multiple different classes of Bieri-Strebel groups, starting from the more basic finiteness properties such as $\mathbf{F}_n$ and proceeding to more complex finiteness properties such as the BNSR invariant. We will employ primarily geometric methods, including Bestvina-Brady Morse Theory, to calculate these properties in general for large classes of groups, as well as implementing a new technique (initially published in [MNSR24]) in order to link the BNSR invariant of a group with that of certain finite-index subgroups. We will also use these geometric techniques to demonstrate a class of Bieri-Strebel groups can act geometrically on a CAT(0) cube complex.

Acknowledgements

This project could have never come to fruition without the kindness and support of the people around me. Primarily I must thank my supervisor, Brita Nucinkis, for four years of limitless patience. My absence from academia prior to starting this project has meant I had a lot of knowledge gaps she was always willing to help me fill in, even when the answers were probably obvious to anyone else.

I would also like to thank Yuri Santos Rego, Matt Zaremsky, and Pep Burillo who have all been incredible mentors over the last four years. In particular, Pep's introduction to Thompson Groups in my first year set my path as a mathematician, and I would not have been able to publish my work were it not for the hard work of Yuri as a co-author to my first paper and Matt as the editor. All three have been invaluable to helping me develop the ideas in this thesis, and it would not exist in this form without them.

I owe a lot to all the postgraduates I shared my time with at Royal Holloway, and in particular I'd like to thank Pavel who discussed ideas with me frequently and often pointed me to what I needed to advance my research, and Tom and Vicki, who frequently arranged social activities for the other postgraduates to enjoy. I also wish to thank Sofiya for not only helping me understand my own ideas, but for always being a calming presence during otherwise stressful conferences.

Finally, I'd like to acknowledge my parents and my brother, who not only supported me while I was searching for PhD projects, but continued to support me throughout the entire process, and always pushed me to stretch further than I thought possible. Without them, I may never have gone back to university, let alone travelled to all the places I've been in the last four years.

Contents

List of Figures

This thesis is dedicated to the memory of
Dr John Harold Molyneux
(The original Dr Molyneux)
and
Mrs Carole Jean Molyneux

Chapter 1

Finiteness Properties

Finiteness properties are a core concept in geometric group theory. Although it is possible to define many different kinds of finiteness property, many of them fit into one of two different categories: homotopical finiteness conditions, and homological finiteness conditions. We shall concern ourselves with both types. Initially we shall discuss a common example from each type: The homotopical finiteness condition F_n, and the homological finiteness condition FP_n.

1.1 F_n

Definition 1.1.1. A CW-complex X is a topological space constructed according to the following process ([Hat02], page 5):

- Begin with a discrete set X^0. The elements of X^0 are the 0-cells of the complex.
- We form the n-skeleton X^n from the $n-1$-skeleton X^{n-1} and a collection of n-disks D^n_α with attaching maps $\phi_\alpha : S^{n-1}_\alpha \to X^{n-1}$ as the quotient space $X^{n-1} \bigsqcup_\alpha D^n_a$ with the boundary δD^n_α of D^n_α identified with $\phi_\alpha(SD^n_\alpha)$ pointwise (each point $x \in D^n_\alpha$ is identified with $\phi_\alpha(x) \in X^{n-1}$).
- If we terminate this process at some finite value n, then $X = X^n$, otherwise $X = \bigcup_n X^n$.

Definition 1.1.2. For a group G, a G-CW-Complex is a CW-complex X with a G-action such that each $g \in G$ when considered as the homeomorphism $g : X \to X$, then g is an automorphism of X. That is to say for each k-cell $x \in X^k$, $g(x)$ is also a k-cell. This means that the action of G on X preserves the cell structure. ([Geo08], page 84).

Definition 1.1.3. For a group G, an Eilenberg-Maclane Space, classifying space, or $K(G,1)$, is a topological space where the first homotopy group is isomorphic to G and all other homotopy groups are trivial ([Hat02], 1.B).

It is worth noting that a $K(G,1)$ is not unique to a group. A group may have many distinct classifying spaces. However, all of these spaces will be the same up to homotopy. We will also need to consider the universal cover of a $K(G,1)$, referred to as an EG.

Definition 1.1.4. A group G has the homotopical finiteness property F_n if there exists a $K(G,1)$ X such that X has a finite n-skeleton. That is to say, a finite number of cells of dimension n or lower. The space may have infinitely many cells of dimension higher than n. We say a group has the property F_∞ if it has F_n for all n, and we say it has the property F if it has a $K(G,1)$ with a finite number of cells in any dimension. [Wal65]

F_n is a generalisation of some very intuitive properties to discuss regarding infinite groups. A group having the F_1 property is equivalent to the group being finitely generated, and F_2 is equivalent to being finitely presented. There exist groups of type F_n but not type F_{n+1} for all n, as well as groups that are of type F_∞ but not type F

1.2 FP_n

Definition 1.2.1. An exact sequence is a sequence of modules M_k over a ring R and homomorphisms $\delta_k : M_k \to M_{k-1}$ such that the image of δ_k in M_{k-1} is the kernel of δ_{k-1}.

Definition 1.2.2. A projective resolution of a module M is an exact sequence

$$... \to P_k \to P_{k-1} \to ... \to P_2 \to P_1 \to M \to 0$$

such that each P_k is a projective module. [**10**; Bro82]

Definition 1.2.3. A module A has the homological finiteness property FP_n over a ring R if there exists a projective resolution of A of the form

$$... \to P_n \to P_{n-1} \to ... \to P_2 \to P_1 \to P_0 \to R \to 0$$

where P_k are finitely generated projective R modules, for $k \leq n$.

A group G has the homological finiteness property $FP_n(R)$ if R is of type FP_n over the group ring RG. A group G has the property $FP_\infty(R)$ if it has $FP_n(R)$ for all n, which is to say it has a projective resolution with the modules P_i finitely generated and projective for all i, and it has the $FP(R)$ property if there exists a projective resolution of the form

$$0 \to P_n \to P_{n-1} \to ... \to P_2 \to P_1 \to RG \to R \to 0$$

where P_k are finitely generated projective RG modules, for $k \leq n$.

For purposes of brevity, this thesis will use FP_n, FP_∞, and FP as shorthand for $FP_n(\mathbb{Z})$, $FP_\infty(\mathbb{Z})$, and $FP(\mathbb{Z})$. Similarly to F_n, there exist groups of type FP_n but FP_{n+1} for all n, as well as groups that are of type FP_∞ but not type FP. For example, Bestvina and Brady ([Bes],example 6.3) constructed a group that is FP_n but not FP_{n+1} from the kernel of a homomorphism from certain right angled artin groups G to the ring of integers $\mathbb{Z}$. Furthermore, Thompson's group (defined in 2.1.1) was shown by Brown and Geoghegan to be FP_∞ but not FP [BG84]

The finiteness properties F_n and FP_n are closely related. F_n implies FP_n for all n, with F_∞ and F implying FP_∞ and FP respectively. This is all implied by the fact that we can create a projective resolution for a group G over Z based on an EG for G. If this EG has a finite number of orbits of $k - cells$ (equivalent to the associated $K(G,1)$ having a finite number of k-cells), then the module P_k in the projective resolution will be finitely generated. Thus, if a group G has the property F_n, then we can generate a projective resolution with the first n modules finitely generated, and so the group also has the property FP_n.

This implication does not generally work in reverse. The exception is that FP_1 implies F_1 ([Bro87], page 197). Otherwise, FP_n does not imply F_n. In fact, for any $n = 2$, there exist groups of type FP_n but not F_n [Bes]. However, we do know that F_2 and FP_n implies F_n.

1.3 BNSR Invariants

Alongside finiteness properties such as F_n and FP_n, we can construct more complex properties that, alongside being group invariants for the purposes of distinguishing different groups, provide more information about group structure. This thesis is particularly interested in the BNSR invariant, or sigma (Σ) invariant.

Definition 1.3.1. The character sphere of a group G, written $S(G)$ or occasionally $\Sigma^0(G)$, is the space $Hom(G;\mathbb{R})/\sim$ of homomorphisms from G to the additive group of real numbers $\mathbb{R}$, modulo the equivalence relation where, for $a, b \in Hom(G;\mathbb{R})$, $a \sim b$ if there exists a positive real number r such that $a(g) = r * b(g)$ for all $g \in G$. ([BS92], Chapter 1).

For a group G, the space $Hom(G;\mathbb{R})$ is isomorphic to $\mathbb{R}^n$, where n is the number of free generators in the abelianization G_{ab}. This is also known as the torsion-free rank of G_{ab}, or $r_0(G_{ab}$ ([BS92], lemma 1.1). Each equivalence class of $\sim$ can be imagined as a ray from the origin to a point at infinity. As such, the quotient space is created by selecting a point from each of these rays, creating a sphere centred at the origin.

1.3.1 Homotopical BNSR Invariants

Definition 1.3.2. For a group G with generating set S, the Cayley graph $\Gamma(G, S)$ is the graph with vertex set $V(\Gamma) = G$ and edge set $E(\Gamma) := \{g, g' \in G | \exists s \in S^{\pm} \, st \, g' = gs\}$ [Löh17].

To define the homotopical BSNR invariant, we start by constructing a labelled Cayley graph, Γ_χ, by selecting a representative character χ from one of the equivalence relations in our character sphere. As every vertex in Γ is a group element, we can label each vertex g with the value $\chi(g)$. From here, we will form the "top half" labelled Cayley graph $\Gamma_{\chi \geq 0}$ by taking the induced subgraph of Γ that includes only vertices with a label greater than or equal to 0.

Definition 1.3.3. The first homotopical BNSR invariant of a finitely generated group G is defined as the following[BS92]

$$\Sigma^1(G) := \{[\chi] \in S(G) | \Gamma_{\chi \geq 0} \text{ is connected}\}$$

We can see that the choice of representative character for a given equivalence class is arbitrary for the purposes of calculating the BNSR invariant. Indeed, as characters in each equivalence class are the same up to multiplication of a positive real number, there can be no element of G that is mapped to a positive number by χ and a nonpositive number by χ', assuming $\chi \sim \chi'$, and vice-versa.

The first BNSR invariant can be generalised in a similar way to how finitely generated can be generalised into F_n. However, this will require a higher dimensional space than the Cayley graph to calculate.

Definition 1.3.4. For a group G, a contractible CW-complex X is an EG if G has a free, cocompact action on X via deck transformations. That is to say, G acts on X by homeomorphisms that map k-cells to k-cells and if it fixes any cell, it fixes it pointwise.

It is possible to construct an EG for a group G by taking a $K(G,1)$ for G and finding the universal cover. Similarly, one can reach a $K(G,1)$ from an EG by quotienting the space by the specified G action.

While the restriction that the action be cocompact isn't strictly necessary for our purposes, it is typical in situations where one is exploring similar properties in both homotopy and homology, as requiring the action to be cocompact results in possible models for EG being restricted to "simpler" models, in the sense that they have fewer cells. Thus, were we to then wish to analyse the homology of this space, it would be simpler to work with.

In order to calculate higher order BNSR invariants. We will need to create a "top half" of our EG X in a similar manner to the "top half" Cayley graph we created before. This is not as straightforward as when working with the Cayley graph, as a Cayley graph has the group G as a vertex set, while an EG does not necessarily have that property. Instead, we establish the "height" of each vertex by choosing a height function $h : X^{(0)} \to \mathbb{R}$ that is equivariant with our character χ. That is to say that, given the group action $G \times X \to X$, we must have that $h(g \cdot x) = h(x) + \chi(g)$. Outside of this stipulation, our choice of h is arbitrary, as shown in [BS92]. We can then form the "top half" of X by first defining the 0-skeleton as

$$X^{(0)}_{h \geq r} := \{x \in X^{(0)} | h(x) \geq r\}$$

We then define the k-skeleta iteratively as

$$X^{(k)}_{h \geq r} := \{x \in X^{(k)} | \text{all faces of x are in } X^{(k-1)}_{h \geq r}\}$$

Definition 1.3.5. For an F_n group G, the n-th BNSR invariant is [BS92]

$$\Sigma^n(G) := \{[\chi] \in S(G) | \exists\ EG\ X \text{ and } \chi\text{-equivariant height function } h \text{ st } X_{h \geq r} \text{ is } n\text{-connected}\}$$

We consider $[\chi]$ to be in $\Sigma^\infty(G)$ if $[\chi] \in \Sigma^n(G)$ for all n .

This initial definition can be difficult to work with, especially when attempting to show a character is not in the Sigma invariant. However, by slightly softening the requirement for the "top-half" to be n-connected, we can create a condition that applies to all EG for a given group G.

Definition 1.3.6. ([Ren88], Kapitel II, Definition 3.6). For $X_{h \geq r}$ as described above, we say that $X_{h \geq 0}$ is essentially n-connected for $n \in \mathbb{Z}_{\geq -1}$ if there is a real number d such that the map between the j-th homotopy groups $i_j : \pi_j(X_{h \geq r}) \to \pi_j(X_{h \geq r-d})$ induced by the inclusion map $i : X_{h \geq r} \to X_{h \geq r-d}$ is the zero map for all $j \leq k$.

Put another way, an essentially connected filtration of a space may have homotopy, but if the filtration is extended by a finite amount, that homotopy vanishes. This is useful for calculating the BNSR invariant for the following reason.

Citation 1.3.7. *([Ren88], Kapitel IV, Satz 3.4). Let G be a group of type F_n and let X be a model for EG with G-finite n-skeleton. Let $\chi : G \to \mathbb{R}$ be a nontrivial character and $h : X \to \mathbb{R}$ be a height function equivariant with χ. Then*

$$[\chi] \in \Sigma^n(G) \iff X_{h \geq r} \text{ is essentially } n\text{-connected}$$

This replaces the condition in 1.3.5 with a condition that can be applied to any EG, in particular allowing us to use an EG that has a not essentially connected filtration as proof that $[\chi] \notin \Sigma^n(G)$.

1.3.2 Homological BNSR Invariants

Analagous to the difference between F_n and $FP_n(R)$, we can define a homological BNSR invariant using the projective resolution. As the character χ is a map from G to $\mathbb{R}$, we can easily define $G_\chi := \{g \in G | \chi(g) \geq 0\}$. We can extend this concept to the group ring $\mathbb{Z}G$, creating the subring $\mathbb{Z}G_\chi$.

Definition 1.3.8. For an FP_n group G and a $\mathbb{Z}G$ module A, the n-th homological BNSR invariant is defined as follows [BR88]

$$\Sigma^n(G; A) := \{[\chi] \in S(G) | A \text{ is of type } FP_n \text{ over the ring } \mathbb{Z}G_\chi\}$$

We say that $[\chi]$ is in $\Sigma^\infty(G; A)$ if it is in $\Sigma^n(G; A)$ for all n.

We may relate the homotopical and homological BNSR invariants (and in particular the homological BNSR invariant over the ring of integers $\mathbb{Z}$) in the following way:

$$\Sigma^1(G) = \Sigma^1(G; \mathbb{Z})$$
$$\Sigma^n(G) = \Sigma^2(G) \cap \Sigma^n(G; \mathbb{Z}) \; \forall n \geq 2$$

Alongside being a group invariant, the homological and homotopical BNSR invariants share a powerful property for calculating the finiteness properties of certain subgroups.

Lemma 1.3.9. *For an F_n group G with $G' \leq H \leq G$ the following are equivalent [BS92]:*

- *H has the finiteness property F_k.*
- *For all $\chi \in Hom(G, \mathbb{R})$ such that $H \in ker(\chi)$, $[\chi] \in \Sigma^k(G)$.*

Analogously for the homological invariant: For an $FP_n(R)$ group G with $G/G' \leq H \leq G$ the following are equivalent:

- *H has the finiteness property $FP_k(R)$.*
- *For all $\chi \in Hom(G, \mathbb{R})$ such that $H \in ker(\chi)$, $[\chi] \in \Sigma^k(G; R)$.*

This result is not only powerful as a tool for understanding the finiteness properties of various subgroups of a given group, but can also be used to calculate the

BNSR invariant of a group if the finiteness properties of certain subgroups are already known.

1.4 BNSR invariants and Finite Index

The following is the result of collaborative work between Lewis Molyneux, Brita Nucinkis, and Yuri Santos Rego. Work contributed in its entirety by the other authors will be labelled

In discussions of BNSR invariants of Bieri-Strebel groups, it has proven useful to be able to derive knowledge of the BNSR invariant of a group from knowledge of the invariant of a finite index subgroup. Unfortunately, such a connection is nontrivial. For example, the group

$$D_\infty := \langle x, y | x^2 = y^2 = 1 \rangle$$

has $\mathbb{Z}$ as an index 2 subgroup (generated by the element xy). However, D_∞ has a finite abelianization and therefore only the trivial character (sending all elements to 0), while $\mathbb{Z}$ has two nontrivial equivalence classes of characters in its character sphere. As such, we can see that finite index alone does not provide a simple correspondence between characters for a group and a subgroup. There are further restrictions we can introduce to create such a correspondence.

Theorem 1.4.1. *Let G be a group of type F_n, $H \subseteq G$ and $|G : H| < \infty$. Let $i : H \hookrightarrow G$ be the inclusion map. If $r_0(G_{ab}) = r_0(H_{ab})$, then $i^* : S(G) \to S(H)$ is a well defined homeomorphism and for all n we have that $i^*(\Sigma^n(G)) = \Sigma^n(H)$.*

Theorem 1.4.2. *Let A be a $\mathbb{Z}G$ module of type FP_n, $H \subseteq G$ and $|G : H| < \infty$. Let $i : H \hookrightarrow G$ be the inclusion map. If $r_0(G_{ab}) = r_0(H_{ab})$, then $i^* : S(G) \to S(H)$ is a well defined homeomorphism and for all n we have that $i^*(\Sigma^n(G; A)) = \Sigma^n(H; A)$.*

These two theorems were originally published in [MNSR24]. All work towards these theorems by the author was undertaken during the PhD course as part of the research towards the PhD.

The proof of these two theorems comprises two main parts: The proof that finite index and equivalence of torsion-free rank is sufficient for i^* to be a well defined homeomorphism, and the proof that finite index and equivalence of torsion-free rank is sufficient for $i^*(\Sigma^n(G)) = \Sigma^n(H)$.

Lemma 1.4.3. *Suppose G is a finitely generated group, let $H \subseteq G$, and write $\pi : G \twoheadrightarrow G_{ab}$ for the canonical projection and $i : H \hookrightarrow G$ for the inclusion. The following hold:*

- *If $|G : H| < \infty$ then the map $i^* : Hom(G; \mathbb{R}) \to Hom(H; \mathbb{R})$ induced by the inclusion i is injective.*
- *If the image $\pi(H)$ is infinite, then there exists a nontrivial morphism $e : Hom(H; \mathbb{R}) \to Hom(G; \mathbb{R})$. That is to say, any character ψ of $H \subseteq G$ corresponds to a character $e(\psi)$ of G and the image $e(Hom(H; \mathbb{R}))$ in $Hom(G; \mathbb{R})$ is a nonzero subspace.*

Part 2 of this lemma was originally observed by Kochloukova and Vidussi [KV23]. However, their proof assumes that the characters in G are extensions of characters from a subgroup H. This lemma makes no such assumption, and allows for $e(\psi)$ to not be a valid extension of ψ. As such, it might be the case that $i^* e(\psi) \neq \psi$.

Proof. Part 1: Suppose for contradiction that i^* is not injective. Therefore there exists a nonzero character $\chi \in Hom(G;\mathbb{R})$ such that $i^*(\chi) = \chi_{|H} = 0$. As χ is nonzero, there must exist $g \in G$ such that $\chi(g) \neq 0$, but as $\chi(H) = 0$, we can see that $g \notin H$. Furthermore, we can conclude that $g^n \notin H$ for all $n \in \mathbb{N}$, as

$$\begin{aligned} g^n \in H &\implies \chi(g^n) = 0 \\ &\iff n\chi(g) = 0 \\ &\iff \chi(g) = 0 \end{aligned}$$

which would contradict our assumption that $\chi(g) \neq 0$. We can go on to say that $g^i \neq g^j$ for all $i < j \in \mathbb{N}$, as $g^i = g^k \implies g^{j-i} = 1 \implies g^{j-1} \in H$, which is a contradiction. As such, each g^n is a distinct coset of H, so H cannot be finite index, contradicting our assumption. Hence, i^* must be injective.

The following proof was developed by Yuri Santos Rego as part of collaborative work on [MNSR24]

Part 2: Consider the finite dimensional $\mathbb{Q}$-vector space $V = G_{ab} \otimes_{\mathbb{Z}} \mathbb{Q}$. Since the image $\pi(H) \subseteq G_{ab}$ is infinite, the set $\pi(H)$ must contain some torsion-free element and thus $\pi(H) \otimes_{\mathbb{Z}} \mathbb{Q}$ contains a partial basis for V. Label this partial basis $B' = \{\bar{h_1}, ..., \bar{h_m}\}$, where each $\bar{h_i}$ is the image in G_{ab} of some h_i in H. From this partial basis we construct a full basis of V as $B = \{\bar{h_1}, ..., \bar{h_m}, \bar{g_{m+1}}, ..., \bar{g_r}\}$, with $\bar{g_j}$ being the image of some $g_j \in G$. Since the image of characters of a group factors through the abelianization, we may define

$$e(\psi)(g) := \sum_{i=1}^{m} a_i \psi(h_i)$$

where the a_x with $x \in B$ are the co-ordinates of the image of G in V using the basis B. e is thus a homomorphism from $Hom(H;\mathbb{R})$ to $Hom(G;\mathbb{R})$ and because $\pi(H) \subseteq G_{ab}$ is infinite and G is finitely generated, the induced map $H \to \pi(H) \otimes_{\mathbb{Z}} \mathbb{R} \cong \mathbb{R}^m$ gives a nontrivial character, labelled $\psi \in Hom(H,\mathbb{R})$, by projecting onto the line spanned by a nonzero vector of $\pi(H) \otimes_{\mathbb{Z}} \mathbb{R}$. By construction, the character $e(\psi)$ is also nontrivial. □

We can now proceed to the first major part of theorems 1.4.1 and 1.4.2, that is the connection between $S(G)$ and $S(H)$ when $|G : H| < \infty$ and G_{ab} and H_{ab} have the same torsion-free rank.

Proposition 1.4.4. *Let G be a finitely generated group with $H \subseteq G$ and $|G : H| < \infty$. The following are equivalent:*

- $r_0(G_{ab}) = r_0(H_{ab})$.
- $i^* : Hom(G;\mathbb{R}) \to Hom(H;\mathbb{R})$ *is an isomorphism of vector spaces.*
- *The assignment* $i^*([\chi]) := [\chi_{|H}]$ *is defined on all character classes* $[\chi] \in S(G)$, *and the corresponding map* $i^* : S(G) \to S(H)$ *is a homeomorphism.*
- *Every character* ψ *in* $Hom(H;\mathbb{R})$ *admits a lift* ψ' *in* $Hom(G;\mathbb{R})$ *such that* $\psi'_{|H} = \psi$ *and* $\psi \neq 0 \iff \psi' \neq 0$.

Proof. The equivalence of (1), (2) and (3) are consequences of 1.4.3 part 1 and of the fact that $dim_{\mathbb{R}}(Hom(\Gamma;\mathbb{R})) = r_0(\Gamma)$ for any group Γ. As we have $|G : H| < \infty$ we know i^* is a monomorphism by 1.4.3. As such, all we need is for i^* to be surjective, but as i^* is injective, we know $dim_{\mathbb{R}}(Im(i^*)) = r_0(G_{ab})$, and as $r_0(G_{ab}) = r_0(H_{ab})$, $Im(i^*)$ must cover all dimensions of $Hom(H;\mathbb{R})$ and thus i^* is an isomorphism. $i^* : S(G) \to S(H)$ is a homeomorphism as a consequence of $i^* : Hom(G;\mathbb{R}) \to Hom(H;\mathbb{R})$ being an isomorphism.

From here, we can connect (4) to (2) by observing that the function $e : Hom(H;\mathbb{R}) \to Hom(G;\mathbb{R})$ given by $e(\psi) = \psi'$ (as described in the proof of 1.4.3 part 2) is a right inverse of i^*. □

Example 1.4.5. The extension map $e : Hom(H;\mathbb{R}) \to Hom(H,\mathbb{R})$ mentioned in the proof of 1.4.4 can be explicitly constructed. Let $\{x_1, ..., x_n\}$ be a generating set for G and say $r_0(G_{ab}) = r_0(H_{ab}) = k \leq n$. Without loss of generality we may assume a subset of abelianized generators $\{\bar{x}_1, ..., \bar{x}_k\}$ generates $(G_{ab})_0$, the torsion free part of G_{ab}. Since $|G : H| < \infty$, for each $i \in \{1, ..., n\}$ there exists $\alpha_i \in \mathbb{N}$ such that $x_i^{\alpha_i} \in H$ (If this were not the case, then $x_i^n H$ would be an infinite set of cosets of H, contradicting finite index). Using the fact that each $\bar{x}_i$ has infinite order in G_{ab}, we can conclude that $0 \neq \alpha_i \bar{x}_i \in (H_{ab})_0$ for all $i = 1, ..., k$. Let $\alpha = \mathrm{lcm}\{\alpha_1, ..., \alpha_k\}$. Given a character $\psi : H \to \mathbb{R}$, we define the lift $e(\psi) = \psi' : G \to \mathbb{R}$ by:

$$\psi'(x_i) = \frac{1}{\alpha}\psi(x_i^{\alpha}) \text{ for all } i \in \{1, ..., n\}$$

This is well defined for all x_i as $x_i^{\alpha_i} \in H$, and as we have defined the character on the generators of G, it naturally extends to all elements of G.

The second stage of the proof is to demonstrate the connection between $\Sigma^n(G)$ and $\Sigma^n(H)$. In principle, we wish to show that $[\chi] \in Sigma^n(G) \iff [\chi_{|H}] \in \Sigma^n(H)$, and $[\psi] \in \Sigma^n(H) \iff [\psi'] \in \Sigma^n(G)$. As discussed in 1.3.1 and 1.3.2, these invariants are defined through topological spaces and projective resolutions. What we shall demonstrate is that, given a space or resolution suitable for calculating $\Sigma^n(G)$ or $\Sigma^n(G;A)$ respectively, we can construct a suitable space or resolution for calculating $\Sigma^n(H)$ or $\Sigma^n(H;A)$, and vice-versa.

To finish the proof of 1.4.2, we can cite a result from Meier, Meinert and VanWyk.

Citation 1.4.6. *([MMV98], Proposition 9.3) Suppose that we have $H \subseteq G$ such that $|G : H| < \infty$ and A is a $\mathbb{Z}G$ module of type FP_n. Further suppose that $\chi : G \to \mathbb{R}$ restricts to a nonzero homomorphism of H. Then*

$$[\chi_{|H}] \in \Sigma^n(H;A) \iff [\chi] \in \Sigma^n(G;A)$$

Proof of Theorem 1.4.2. 1.4.4 demonstrates that, for a group G with finite index subgroup H such that $r_0(G_{ab}) = r_0(H_{ab})$, any character class $[\chi] \in S(G)$ has a nontrivial restriction $[\chi_{|H}] \in S(H)$, and any character $\psi \in S(H)$ has a suitable lift $\psi' \in S(G)$. 1.4.6 shows that, as long as the restrictions and lifts exist, then they are in the same invariants. The theorem is thus an immediate consequence of these two results. □

To prove 1.4.1, we can use the nature of an EG as a space that is acted upon cocompactly by a group G to use the same space as an EH for a finite index subgroup H.

Proposition 1.4.7. *Let G be a group of type F_n with $H \subseteq G$ such that $|G : H| < \infty$ and $r_0(G) = r_0(H)$. The following holds*

$$[\chi] \in \Sigma^n(G) \implies [\chi_{|H}] \in \Sigma^n(H)$$

$$[\psi] \in \Sigma^n(H) \implies [\psi'] \in \Sigma^n(G)$$

where ψ' is the lift of the character $\psi \in Hom(H;\mathbb{R})$ to $Hom(G;\mathbb{R})$ as described in 1.4.4.

Proof. Starting with the first implication, consider a model X for EG with G-finite n-skeleton (as G is F_n, this must exist). Now suppose $[\chi] \in \Sigma^n(G)$. From the definition of the BNSR-invariant, we know that $X_{h_\chi \geq 0}$ is essentially n-connected for some height function h_χ that is equivariant with χ. Since H is finite index in G, X is also a model for EH with finite H-skeleton, and $h_\chi = h_{\chi_{|H}}$, so $X_{h_{\chi_{|H}} \geq 0}$ is the same as $X_{h_\chi \geq 0}$, and is therefore essentially n-connected. Hence $[\chi_{|H}] \in \Sigma^n(H)$. We now assume $[\psi] \in \Sigma^n(H)$. We choose a model X for EH such that X is a simplicial complex with G-finite n-skeleton and one G-orbit of zero cells that we may label with elements of G. As $|G : H| < \infty$, we know such a model must exist. We now fix a set $T = t_0, ..., t_{m-1}$ of coset representatives of H in G, such that $t_0 = 1$, and construct a ψ-equivariant height function $h_\psi : X \to \mathbb{R}$ on the vertices of X as follows. For any $\gamma \in H$ we set $h_\psi(\gamma) = \psi(\gamma)$, and for each $t_i \in T$ we set $h_\psi(t_i) = 0$. As the cosets of H partition G, we may write any element of G uniquely as $g = t_i\gamma$. When using the height function, we therefore get

$$h_\psi(g) = h_\psi(t_i) + h_\psi(\gamma) = \psi(\gamma)$$

We then extend this function piecewise linearly to the entire n-skeleton on X. As we have $[\psi] \in \Sigma^n(H)$ by assumption, we know by 1.3.7 that $X_{h \geq r}$ is essentially n-connected. The final step is to show that this connectivity is preserved when changing to a height function $h_{\psi'}$ corresponding to ψ', the lift of ψ inside $Hom(G;\mathbb{R})$. The complicating factor being that $\psi'(t_i)$ is not necessarily equal to 0. We will take $d = \mathbf{min}\{\psi'(t_i)\}$. Note that, as $t_0 = 1$, $d \leq 0$. We claim that, for every $g \in G$, $h_\psi(g) \geq 0$ if and only if $\psi'(g) \geq d$. To demonstrate this, we write $g = t_i\gamma$ as above, since $h_\psi(g) = \psi(\gamma)$ and $\psi'(g) = \psi'(t_i\gamma) = \psi'(t_i) + \psi(\gamma)$, we get

$$h_\psi(g) \geq 0 \iff \psi(\gamma) \geq 0 \iff \psi'(t_i) + \psi(\gamma) \geq d + 0 \iff \psi'(g) \geq d$$

This implies that the 0-skeleton of $X_{h_\psi \geq 0}$ is precisely the same as the 0-skeleton of $X_{h_{\psi'} \geq d}$. From the definition of the "top half" given in 1.3.1, we can see that $X_{h_{\psi'} \geq d}$ must be essentially n-connected, and by 1.3.7, $[\psi'] \in \Sigma^n(G)$. □

Proof of Theorem 1.4.1. Similarly to the proof of 1.4.2, 1.4.4 guarantees the existence of character lifts and restrictions between $S(G)$ and $S(H)$, and 1.4.7 demonstrates that, for corresponding characters $\chi \in S(G)$ and $\chi_{|H} \in S(H)$

$$[\chi] \in \Sigma^n(G) \iff [\chi_{|H}] \in \Sigma^n(H)$$

And that the same holds for $\psi \in S(H)$ and its lift $\psi' \in S(G)$. The theorem is therefore a consequence of both these results. □

Chapter 2

Bieri-Strebel Groups

Bieri-Strebel groups emerge as a natural generalisation of Thompson's group and the related Brown-Thompson groups. Many properties are shared between Thompson's group and Bieri-Strebel groups and yet more properties of Bieri-Strebel groups are defined by how they differ from Thompson's group. With this in mind, we will define and discuss Thompson's group in order to discuss Bieri-Strebel groups in comparison.

2.1 Thompson's Group

Definition 2.1.1. Thompson's group **F** is the group of piecewise-linear, orientation-preserving homeomorphisms of the unit interval $[0, 1]$ (denoted I) such that:

- All gradients are in the multiplicative group $\langle 2 \rangle$.
- There are finitely many breakpoints separating the linear piece.
- All breakpoints fall in $I \bigcap \mathbb{Z}[\frac{1}{2}]$.

Thompson's group **T** and Thompson's group **V** are related groups that are defined over the unit circle and the cantor set respectively. They are not used in this thesis, but an interested reader can learn more in [CFP96]

Thompson's groups have many unique properties. **V** was the first known finitely presented infinite simple group, and **F** has a simple, not finitely generated commutator subgroup. **F** is most commonly written with the following presentation [Bur]

$$\mathbf{F} = \langle x_0, x_1, ... | x_j x_i = x_i x_{j+1} \text{ for all } i < j \rangle \tag{2.1}$$

F has the F_∞ property and is therefore finitely presented, but the finite presentation is cumbersome and not often used in results. It can be seen in [Bur].

An element of Thompson's group **F** can be considered as a function mapping the interval I onto itself. As they are orientation preserving homeomorphisms, they must always map 0 to 0 and 1 to 1, and all slopes must be positive. Considering Thompson group elements in this manner can lead to interesting dynamical properties that can be derived through analysing these functions. See, for example [HM23].

2.1.1 Tree-Pair Diagrams

A frequently useful property of Thompson's group **F** (and many other Thompson-like groups) is the ability to express elements of the group as diagrams consisting of an ordered pair of rooted binary trees. This property can be derived from the expression of elements as functions on the interval.

Definition 2.1.2. ([CFP96], section 2). A partition of the interval I is a set $\{x_0, ..., x_n\}$ such that

$$0 = x_0 < x_1 < ... < x_{n_1} < x_n = 1$$

Each interval $[x_i, x_{i+1}]$ is an interval of the partition, and a partition is dyadic if $x_i \in I \bigcap \mathbb{Z}[\frac{1}{2}]$ for all $0 \leq i \leq n$. A dyadic partition is standard if each interval may be written in the form $[\frac{a}{2^b}, \frac{a+1}{2^b}]$ with $a < 2^b - 1$

It is easy to see that an element of f may be uniquely defined by its breakpoints, as the function must be piecewise linear, so the parts of the function between breakpoints is merely a set of linear gradients that connect consecutive breakpoints. Furthermore, each breakpoint is a pair of dyadic numbers, so we may express a set of breakpoints as an ordered pair dyadic partitions, each containing the same number of intervals.

If we wished to express an element of **F** visually using this breakpoint form, we could draw the two partitions of the interval, with breakpoints marked, and then indicate which interval in the domain is mapped to each interval in the codomain, as in 2.1. Note that, when working in **F**, indicating where each interval is mapped is not strictly necessary. As the function is an orientation preserving homeomorphism of I, we know that 0 must be mapped to 0, 1 must be mapped to 1, and that the function must be continuous. From this alone, we may conclude that the first interval of the domain must be mapped to the first interval in the codomain, the second to the second, and so on.

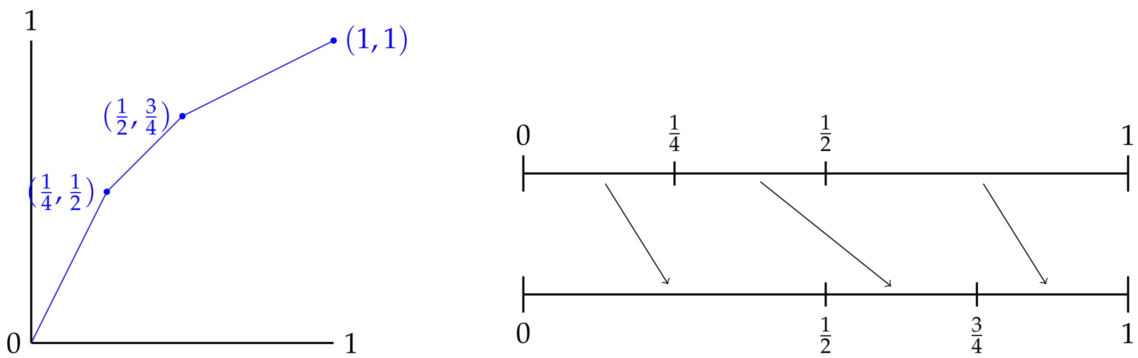

FIGURE 2.1: The same element of **F**, expressed as both a function $f : I \to I$ and as a pair of partitions.

We may derive the tree-pair representation of an element from its partition representation by first rewriting the partitions as standard dyadic partitions. This may introduce redundant breakpoints, that is to say breakpoints that have slopes of the same gradient of either side, but this does not change the function in any way and we may always find an ordered pair of standard dyadic partitions for each element of **F** ([CFP96], lemma 2.2). We may then express a standard dyadic partition as a rooted binary tree in the following way:

- Each node of the tree represents a standard dyadic interval.
- The root node represents the interval I.
- For a node with corresponding interval $[\frac{a}{2^b}, \frac{a+1}{2^b}]$, the two nodes directly beneath it represent $[\frac{2a}{2^b}, \frac{2a+1}{2^{b+1}}]$ (on the left) and $[\frac{2a+1}{2^{b+1}}, \frac{a+1}{2^b}]$ (on the right).

In order for the structure of the tree to be meaningful, which is to say that a caret has a consistent meaning across the tree, and the depth of a leaf to correspond to the length of the interval it represents in the interval partition, the types of interval divides must be consistent. When we talk about tree pairs for more generalised Thompson groups in 2.3.2, we will introduce more flexibility. For now, we require that each caret represent the same division of an interval into two equal parts.

As any element of **F** can be represented as an ordered pair of standard dyadic partitions of I, and each partition can be represented as a rooted binary tree, we can represent any element of **F** as an ordered pair of rooted binary trees

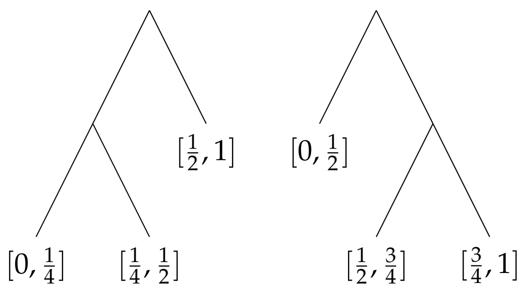

FIGURE 2.2: The element from 2.1, now presented as a tree-pair. Note how the nodes in the trees are positioned similarly to the breakpoints in the partitions.

Just as we have that the domain and codomain partitions must have an equal number of intervals, we require that each tree in a tree-pair diagram have the same number of end points (or leaves).

2.1.2 Redundant Carets

An important thing to note is that, while an element of **F** can be represented with a tree-pair diagram, this representation is not unique. Indeed, there are an infinite number of tree-pairs that represent each element of **F**. This may be observed in partition pairs as well. As mentioned in 2.1.1, when constructing a partition pair, we may introduce a redundant breakpoint by bisecting an interval in the domain and the interval it is mapped to in the codomain. Introducing this breakpoint creates two new intervals in both the domain and the codomain, each half the length of the interval that spawned it. As such, the gradient of the function is unchanged. Were we to place such a breakpoint on the graph of the function, we would see it as a point with the same gradient either side.

When we refer to a caret while discussing tree-pair diagrams, we refer to a subgraph consisting of one parent node, its direct descendants, and the edges connecting them. Carets are the building blocks of our tree-pairs, as each one represents a breakpoint in the corresponding partition. Just as we may introduce a redundant breakpoint by bisecting an interval in the domain and the corresponding interval in the codomain, so too may we create a redundant caret by adding a caret to one of the leaves of the domain tree, and to its corresponding leaf in the codomain. When presented with a tree-pair diagram, we may detect redundant carets on the lowest level of the trees in the following way:

- Number each leaf on the domain tree from left to right, and repeat on the codomain tree.
- If a caret in the domain tree has the same numbered leaves as a caret from the codomain tree, then those carets form a pair of redundant carets.

We shall refer to a tree-pair diagram with no redundant carets as a reduced diagram. We may consider an equivalence class on the set of tree-pair diagrams where two diagrams are considered equal if it is possible to construct one diagram from the other via the addition and removal of redundant carets. We can then consider the set of equivalence classes of tree pairs, modulo this equivalence relation.

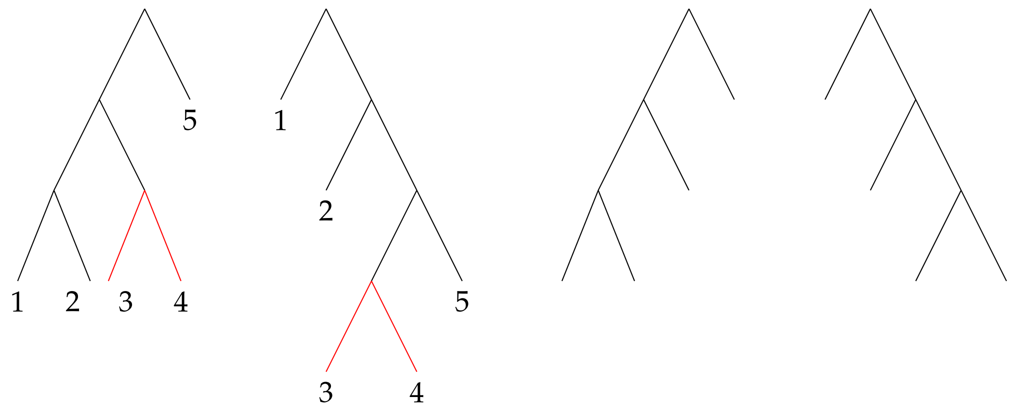

FIGURE 2.3: A tree-pair diagram with a redundant caret, highlighted in red, and the equivalent reduced diagram.

Lemma 2.1.3. *([Bur], proposition 1.2.5). There is precisely one reduced diagram within each equivalence class of tree-pair diagrams.*

Proof. This proof is split into two parts: proving the existence of at least one reduced diagram in each equivalence class, and proving the uniqueness of that reduced diagram.

Existence: For each equivalence class, take an arbitrary diagram contained within. If the diagram is reduced, then we are done. Otherwise, it must contain at least one pair of redundant carets. We remove all redundant carets and check again. Repeating this process will result in a reduced diagram for each equivalence class.

Uniqueness: Suppose there exists 2 distinct reduced diagrams A, B. As A and B are distinct, they must have different carets, and as they are both reduced, these different carets must be non-redundant. As such, to transform A into B or vice versa would require the addition or removal of non-redundant carets, which means that A and B cannot be in the same equivalence class. □

With each equivalence class containing exactly one reduced diagram, each distinct reduced diagram representing a distinct element of **F**, and each element of **F** being representable as a tree-pair diagram, we can conclude that the set of equivalence classes of tree-pair diagrams is precisely the set of elements of **F**. Pushing this concept to its natural conclusion, we can also define a composition of equivalence classes of tree pair diagrams that corresponds precisely with the binary operation of function composition in **F**, allowing the group of equivalence classes of tree-pair diagrams to be precisely the same group as **F**.

Definition 2.1.4. ([Bur], section 1.3) Composition of equivalence classes of tree-pair diagrams is performed in the following way:

- Take the reduced representatives of each equivalence class.
- Add redundant carets to each tree-pair until the codomain tree of the first tree pair is identical to the domain tree of the second tree pair.
- The product of the composition is the equivalence class containing the tree pair with the domain tree of the first tree pair as its domain tree and the codomain tree of the second tree pair as its codomain tree.

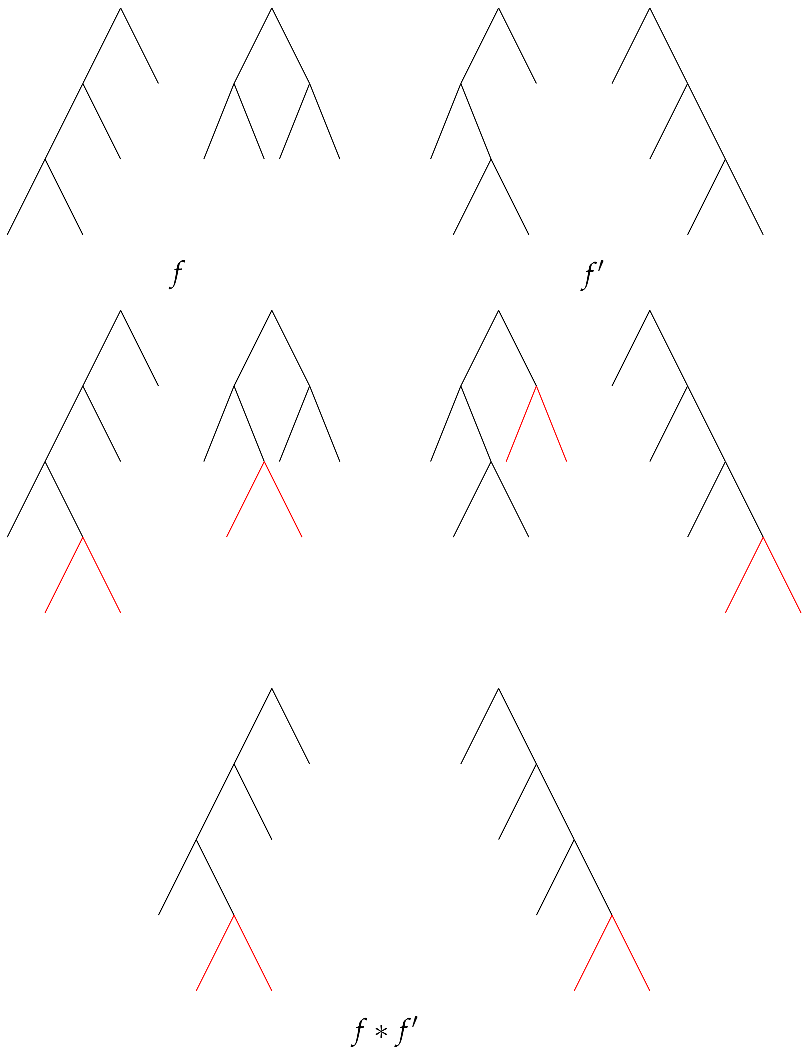

FIGURE 2.4: Demonstration of composition of tree-pairs.

It is important to highlight that it is always possible to add redundant carets to make the middle two trees identical in step 2. As shown in ([Bur], Lemma 1.3.1), superimposing any two trees will create a tree that can be constructed from either with the addition of redundant carets, and the addition of redundant carets can be done to any tree-pair an arbitrary number of times without leaving the equivalence class. As such, the method of tree-pair composition described in 2.1.4 is possible with any two elements of **F**.

Being able to consider **F** as the group of equivalence classes of tree pair diagrams as well as the group of homeomorphisms is a powerful tool. Tree-pair diagrams are the basis for a construction of the infinite presentation of **F** ([Bur], Thm 2.1.1) and, as will be seen in 2.2, are useful in the construction of geometric spaces upon which the group can act.

2.1.3 $\mathbf{F}_n$.

While there are many possible ways to generalise Thompson's group **F**, one that provides a useful middle ground between Thompson's group and the Bieri-Strebel groups is the family of Brown-Thompson groups $\mathbf{F}_n$. Many concepts from Brown-Thompson groups will carry over to Algebraic Bieri-Strebel groups and many of the concepts we have discussed and shall discuss for **F** carry over to the $\mathbf{F}_n$ groups. It is therefore prudent to define and discuss the Brown-Thompson groups before proceeding to the Bieri-Strebel Groups.

Definition 2.1.5. For $n \in \mathbb{N}, n \geq 2$, the Higman-Thompson group $\mathbf{F}_n$ is the group of piecewise-linear, orientation preserving homeomorphisms of the unit interval $[0, 1]$ (commonly written as I) such that:

- All gradients are in $\langle n \rangle$.
- There are finitely many breakpoints between the slopes.
- All breakpoints fall in $I \bigcap \mathbb{Z}[\frac{1}{n}]$.

as with **F**, we can define related groups $\mathbf{T}_n$ and $\mathbf{V}_n$ over the unit circle and the cantor set.

Evidently, the Brown-Thompson groups are a natural generalisation of Thompson's group. It can be seen that Thompson's group **F** is the same as the Brown-Thompson group $\mathbf{F}_2$. The Brown-Thompson groups have a similar presentation to **F**.

$$\mathbf{F}_n = \langle x_0, x_1, ..., | x_j x_i = x_i x_{j+n-1} \text{ for } j > i \rangle \tag{2.2}$$

Similar to **F**, the Brown-Thompson groups $\mathbf{F}_n$ can be represented with both partition pairs and tree-pair diagrams. The primary difference between tree-pairs representing different Brown-Thompson groups is the shape of the carets. In **F**, each time we subdivide an interval, we split it into 2 equal pieces. This concept generalises to $\mathbf{F}_n$, where we subdivide each interval into N equal pieces. This means that, in the tree pair diagrams for $\mathbf{F}_n$, each caret has n legs. That is to say that each node either has 0 direct descendants (and is therefore a leaf), or n direct descendants.

Other than this difference, tree-pairs function identically in $\mathbf{F}_n$ to how they function in **F**. Redundant carets may be detected in the same way. We still form the same equivalence classes of tree-pairs, and tree-pair composition functions in the exact same way. As such, many properties for **F** derived from tree-pair diagrams can be analogised to $\mathbf{F}_n$, even if those properties are not the same. The presentation in 2.2 is a good example of this, as it is similar to the presentation for **F**, but changes as n changes.

2.1.4 Characters for $\mathbf{F}_n$

In 5, we will calculate the BNSR invariant for a set of generalised Thompson groups. As mentioned in 1.3, this will require us to know the characters of those groups. These will also be calculated in Chapter 5. To provide the necessary context for those calculations, we will discuss the characters of **F** and $\mathbf{F}_n$ here.

As discussed in 1.3.1, the size of the character sphere, and thus the number of characters we need to find, is dependant on the abelianization of the group in question. Thus, our first goal is to calculate the abelianization of **F**.

Lemma 2.1.6. *([BGK10], section 1.4) The abelianization $\mathbf{F}_{ab}$ is isomorphic to $\mathbb{Z}^2$.*

Proof. Taking the abelianization of the group with the presentation 2.1, we can use the infinite family of relations to conclude the following for $j \geq 1$ (using addition notation for composition in the abelianization):

$$\begin{aligned} \bar{x}_j + \bar{x}_i &= \bar{x}_i + \overline{x_{j+1}} \\ \bar{x}_j &= \bar{x}_i - \bar{x}_i + \overline{x_{j+1}} \\ \bar{x}_j &= \overline{x_{j+1}} \end{aligned} \tag{2.3}$$

Via a simple induction, we can therefore see that $\bar{x}_j = \bar{x}_k$ for all $j, k \geq 1$. This means that $\mathbf{F}_{ab}$ has two free generators, $\bar{x}_0$ and $\bar{x}_1$, and no torsion generators, meaning it must be isomorphic to $\mathbb{Z}^2$ □

We now know that $r_0(\mathbf{F}) = 2$, and so we would expect there to be two linearly independent characters in $\mathbb{F}$. To determine what these characters are, we can cite Bieri and Strebel.

Citation 2.1.7. *([BS92], Chapter IV, section 3). For G a group of piecewise-linear homeomorphisms of the interval, the character χ_0, the gradient of the slope around the point 0, and the character χ_1, the gradient of the slope around the point 1, will always be linearly independent characters in $Hom(G, \mathbb{R})$.*

When working with these two characters, we will use $log_n(\chi)$ for the group $\mathbf{F}_n$, as this allows the characters to be additive under group composition, rather than multiplicative. As $r_0(\mathbf{F}) = 2$, these two characters must generate the entirety of $Hom(\mathbf{F}, \mathbb{R})$.

Moving on to $\mathbf{F}_n$, we can use a similar technique to 5.1.1 to conclude the following:

Citation 2.1.8. *([Zar17], Section 3.2) The abelianization of $\mathbf{F}_n$ is isomorphic to $\mathbb{Z}^n$, and as such $r_0(\mathbf{F}_n) = n$.*

From 2.1.7, we know that two of these characters will be χ_0 and χ_1, the slopes at 0 and 1. The rest of the characters for these groups all function in a similar way to each other. When considering the groups $\mathbf{F}_n$ for $n \geq 3$, we can observe orbits of breakpoints. Expressing each breakpoint in $\mathbf{F}_n$ as $\frac{a}{n^b}$, each orbit consists of breakpoints such that $a \cong k \mod n-1$ for some fixed k. That is to say that there exists an orbit of breakpoints where we can express each breakpoint in the form $\frac{1}{n^b}$, another orbit of breakpoints where we can express each breakpoint as $\frac{2}{n^b}$ and so on. We can see why these orbits are disjoint most easily in the tree-pair representation. We will use $\mathbf{F}_3$ as an example. Taking an arbitrary tree pair and a pair of corresponding leaves in the tree pair, we can adjust the map by adding a caret to the left of the domain leaf and to the right of the codomain leaf, shifting the corresponding leaf on the codomain tree. However, as we are adding three-legged carets, we are replacing one leaf with three, netting two new leaves each time and thus shifting the corresponding leaf two to the right, as in 2.5. As we can only perform these shifts 2 at a time, the odd numbered leaves can never map to the even numbered leaves, so these form the two separate orbits.

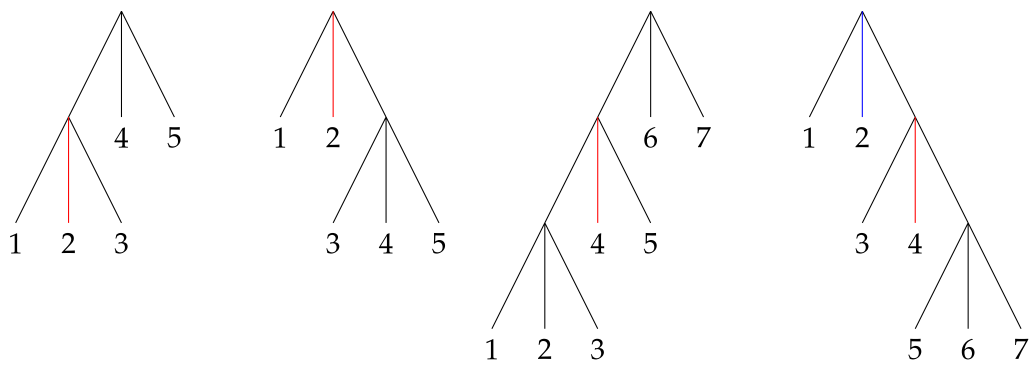

FIGURE 2.5: A tree pair in $\mathbf{F}_3$, and a different tree pair constructed by adding a caret to each tree. Note how the highlighted leaf in the domain tree is mapped to a different leaf in the codomain tree in the new construction.

Using these orbits of breakpoints, we can construct the remaining characters of $\mathbf{F}_n$ in the following way. Consider an individual breakpoint in an element $f \in \mathbf{F}_n$. Each breakpoint has a slope entering the breakpoint and a slope leaving. We can assign a value to each breakpoint equal to the difference in gradient between the exit slope γ_+ and the entry slope γ_-. We will adjust this slightly and say that $\gamma = log_n(\gamma_+) - log_n(\gamma_-)$, as this will allow our character to be additive. For orbits of breakpoints $X_1, ..., X_{n-1}$, we calculate the characters as the following ([Zar17], Definition 3.2)

$$\psi_i(f) = \sum_{x \in X_i} \gamma(x)$$

While this is an infinite sum, almost all terms are 0, as any invisible breakpoints x have the same exit and entry slope by definition, so $\gamma(x)$ is 0 for all invisible breakpoints.

Important to note here is that this defines a unique character for each of the $n-1$ orbits of breakpoints for $\mathbf{F}_n$. Adding those to the two characters provided by 2.1.7 provides us with $n+1$ characters. However, from 2.1.8, we would expect $\mathbf{F}_n$ to have n linearly independent characters. This implies that there is some linear dependence between the set of characters $\chi_0, \chi_1, \psi_1, ..., \psi_n$. In fact, we can deduce the relationship between these characters by considering the changes in slope gradient.

Take an arbitrary element $f \in \mathbf{F}_n$ and consider $\chi_1(f) - \chi_0(f)$, the difference between the log of the slope at 0 and the log of the slope at 1. If this is nonzero, then it implies that the slope at 0 is different to the slope at 1, so the slope must change somewhere in the function. As the function is piecewise-linear, this can only happen at breakpoints. The ψ characters detect whenever the function changes slope at a breakpoint, and $\sum_{i=1}^{n-1} \psi_i$ is the net total that the slope has changed when considering every breakpoint, and is therefore equal to the difference between the log of the final slope (at 1) and the log of the initial slope (at 0). As such we have

$$\chi_1(f) - \chi_0(f) = \sum_{i=1}^{n-1} \psi_i$$

which gives us linear dependence between our $n+1$ characters. We will typically exclude ψ_{n-1} when talking about the characters of $\mathbf{F}_n$ due to this, as $Hom(\mathbf{F}, \mathbb{R})$ is spanned by the remaining characters.

2.2 The Stein-Farley Cube Complex

An incredibly important concept in the understanding of Thompson groups and of their finiteness properties in particular, is the Stein-Farley Cube Complex. An early version of this complex was developed by Brown in [Bro87] in order to prove that **F** and associated groups had the F_∞ property. It was later adapted by Stein in [Ste92] and Farley in [Far03] in order to prove further properties of Thompson Groups. Chapter 3 is all about adapting this complex for use with Bieri-Strebel groups, so we will take the time to explain its construction within the context of the Brown-Thompson groups $\mathbf{F}_n$.

2.2.1 Brown's Complex

The following is adapted from ([Bro87], Section 4) and [Zar19].

We define an n-ary forest as a non-empty ordered set of rooted n-ary trees. Just as we constructed tree pairs in 2.1.1, we will construct the space of forest-tree pairs: ordered pairs that consist of one n-ary forest and one rooted n-ary tree, where the forest has the same number of leaves as the tree. By convention, we will consider the forest on the left and the tree on the right and will write such a pair as (F, T).

Similarly to how we worked with tree pairs in 2.1.2, we will introduce an equivalence relation such that two forest-tree pairs will be equivalent if one may be constructed from the other via the addition or removal of redundant carets, with redundant carets defined as they were for tree-pairs. From here, we can form the equivalence classes $[F, T]$ by taking the set of all forest-tree pairs modulo the equivalence relation.

From here, we will impose a partial order on the set of equivalence classes X in the following way: we define a split as a transformation of a forest-tree pair that changes the forest by deleting the top caret of a tree contained within that forest, replacing the tree with two trees, each of which was a subtree of the original tree.

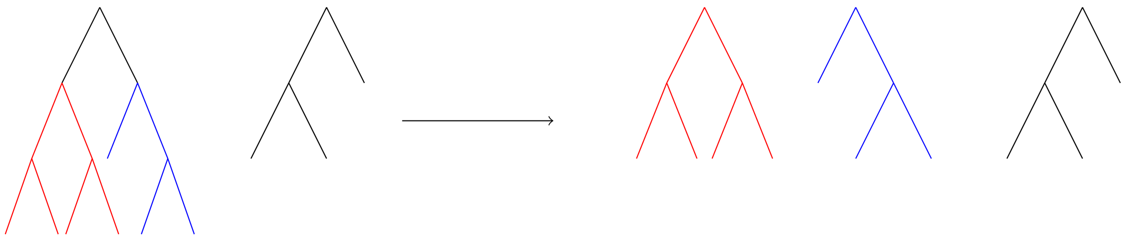

FIGURE 2.6: A split in the first tree of a two tree forest, with the subtrees of the split tree highlighted.

From here, we impose the partial order by saying that $[F,T] \leq [F',T']$ if $[F',T']$ can be produced from $[F,T]$ via performing any number of splits. We can see this is a partial order as it fulfills the following criteria:

- Reflexivity: $[F,T] \leq [F,T]$ as you may perform 0 splits to reach $[F,T]$ from $[F,T]$;
- Antisymmetry: Splits necessarily increase the number of trees in a forest. If $[F,T] \leq [F',T'] \leq [F,T]$, then F has at least as many trees as F' and F' has at least as many trees as F. Thus the number of splits performed to reach one from the other is 0 and therefore F and F' are the same forest;
- Transitivity: If there is a path of splits A that constructs $[F_A,T_A]$ from $[F,T]$, and a path of splits B that constructs $[F_B,T_B]$ from $[F_A,T_A]$, then the composed path AB constructs $[F_B,T_B]$ from $[F,T]$. Thus $[F,T] \leq [F_A,T_A] \leq [F_B,T_B] \implies [F,T] \leq [F_B,T_B]$.

Having created the poset X. We also wish to show that it is directed, that is to say, for any two elements $x,x' \in X$, there is an element y such that $x \leq y, x' \leq y$. We may show this by considering the forest-tree pairs. Taking two forest-tree pairs $[F,T]$ and $[F',T']$, we can construct the forest-tree pair $[F'',T'']$ in the following way:

We begin by constructing T'' similarly to how we perform tree-pair composition: superimposing T and T' will create T'', which may be constructed from T and T' via the addition of redundant carets, and as we are working with equivalence classes, we can make the forest-tree pairs $[\bar{F},T''] \cong [F,T]$ and $\bar{[F']},T''] \cong [F',T]$. As $\bar{F}$ and $\bar{F}'$ are each in a forest-tree pair with T'', we know they both must have the same number of leaves as T'', and thus the same number of leaves as each other. We will then perform splits on $\bar{F}$ and $\bar{F}'$ until each consists of the same number of trivial trees (trees with no carets). We will label this trivial forest F''. As we can reach F'' from a sequence of splits from both $\bar{F} and \bar{F}'$, we must have that $[F'',T''] \geq [\bar{F},T''] \cong [F,T]$ and $[F'',T''] \geq [\bar{F}',T''] \cong [F',T']$, and thus the poset is directed.

From here, we may construct the geometric realisation $\tilde{X}$ of the poset X in the following way:

- The 0 cells, or vertices, of the space $\tilde{X}$ are the elements of the poset X.
- There is a 1-cell, or edge, joining two vertices $\tilde{x}, \tilde{x'} \in \tilde{X}^0$ if $x \leq x'$ (or $x' \leq x$) in the poset.
- The 0-cells $\tilde{x_0}, ..., \tilde{x_n}$ form an n-simplex if we have $x_0 \leq ... \leq x_n$ in the poset.

This constructs $\tilde{X}$ as a simplicial complex. $\mathbf{F}_n$ has a free action on this complex via tree pair composition. Treating $[F, T]$ as a tree pair allows us to compose it with a tree pair $[T_1, T_2]$ in $\mathbf{F}_n$, giving us a right-$\mathbf{F}_n$ action on the complex. Unfortunately, $\tilde{X}$ is not $\mathbf{F}_n$ finite, so we cannot draw finiteness properties directly from this complex. In [Bro87], Brown uses Brown's criterion (discussed in more detail in chapter 3) to determine finiteness properties for $\mathbf{F}_n$ from this complex. In addition to other geometric setup, Brown uses the following result regarding posets that will be useful to us later

Citation 2.2.1. *[Qui73] If a poset X is directed, then the geometric realisation $\tilde{X}$ is contractible, that is to say that all of its homotopy groups are trivial.*

This allows us to conclude that $\tilde{X}$ is contractible. Stein and Farley's adaptations to Brown's complex will maintain this property, as it remains exceedingly useful for calculating finiteness properties for $\mathbf{F}_n$. Our adaptations of Brown's complex and of the Stein-Farley cube complex in 3 and 4 will also preserve this property.

2.2.2 Stein's Complex

The following is adapted from [Ste92] and [Zar19].

Stein adapted Brown's complex through a process of simplification that maintained the structure of the space, including the action of $\mathbf{F}_n$ on the space and its homotopical properties. The 0-skeleton of Stein's complex is identical to the 0-skeleton of Brown's complex. The 1-skeleton is significantly reduced.

Stein's complex still uses splits to determine the 1-skeleton, but instead of two vertices being connected if there is any sequence of splits to construct one from the other, we restrict to "elementary" sequences of splits. A sequence of splits is elementary if each tree of the forest is split only once. That is to say that once a tree has been split, none of its subtrees may be split in an elementary sequence. If a forest-tree pair $[F', T']$ can be constructed from an elementary sequence of splits on the forest-tree pair $[F, T]$, then we write $[F, T] \preceq [F', T']$. We can see that $[F, T] \preceq [F', T'] \implies [F, T] \leq [F', T']$, but $\preceq$ is not an equivalence relation, as it is not transitive.

The n-skeleton is determined similarly to Brown's complex, but for a sequence of elements $x_0 \leq ... \leq x_n$ to be an n simplex in Stein's complex, we require that $x_i \preceq x_j$ for all $i \leq j \leq n$. Despite this removing the vast majority of n-simplices in the complex, we can still show that the space is contractible by building back up to Brown's complex from Stein's complex

Lemma 2.2.2. *([Zar19], Proposition 4) Stein's complex $\tilde{X}'$ is homotopy equivalent to Brown's complex $\tilde{X}$, and is therefore contractible.*

Proof. We begin by considering the maximal elementary sequence of splits of a forest. Given a forest F, there is a unique forest F' obtained by splitting each non-trivial tree of F, and not splitting any subtrees of those trees. We call this the elementary core of F, or F'. Now consider $x, z \in X$ such that $x \leq z$ but $x \npreceq z$. We consider the subset of X defined as the following:

$$(x, z) := \{y \in X | x < y < z\}$$

For each element $y \in (x, z)$, we know there is a sequence of splits $\bar{y}$ that produces y from x. We will remove from $\bar{y}$ all splits on subtrees of x, creating the elementary sequence $\bar{y}'$ that constructs the element y' from x. Clearly $x \preceq y' < y < z$, and therefore $y' \in (x,z)$. We will construct the map $\sigma(y)(x,z) \to (x,z) : y \mapsto y'$. We can see σ is a poset map and that $\sigma_{|Im(\sigma)}$ is the identity map on $Im(\sigma)$, so by ([Qui73], section 1.5), we can conclude that the subspace $\widetilde{(x,z)} \subset \widetilde{X}$, the geometric realisation of the poset (x,z), is contractible.

From here, our goal is to show that we can build X out of X' by gluing in missing simplices, and that any time we glue in a simplex, the relative link we glue along is contractible, meaning the new simplex does not change the homotopy. If we can build X out of X' without changing its homotopy, then X and X' must be homotopy equivalent, and thus X' is contractible. The missing simplices can be grouped into subcomplexes of the form

$$\widetilde{[x,z]} := \{y \in \widetilde{X|x \leq} y \leq z\}$$

where $x \npreceq z$, as the subcomplexes where $x \preceq z$ are already in $\widetilde{X'}$. The relative link we glue $\widetilde{[x,z]}$ along is

$$\widetilde{[x,z)} \bigcup \widetilde{(x,z]} := \{y \in \widetilde{X|x \leq} y < z\} \bigcup \{y \in \widetilde{X|x <} y \leq z\}$$

From here, we can see $\widetilde{[x,z)} \cup \widetilde{(x,z]}$ is the suspension of $\widetilde{[x,z]}$, as it is just $\widetilde{[x,z]}$ with the addition of the points $\widetilde{x}$ and $\widetilde{z}$ and the connecting k-simplices. As $\widetilde{(x,z)}$ is contractible, and the suspensions of contractible spaces are themselves contractible, we know the relative link is contractible. Thus constructing X out of X' does not change homotopy, so they must be homotopy equivalent. □

2.2.3 Farley's Complex

The following is adapted from [Far03] and [Zar19].

Stein's complex was further adapted by Farley for use with diagram groups [Far03]. Farley was not only able to use the complex for the F_n finiteness properties previously discussed, but was also able to use it to prove cubulative properties for diagram groups.

Farley's adaptation is in many ways a continuation of Stein's. The primary change is the further removal of 1-cells and the merging of k-cells. While Stein's complex has an edge between any two vertices $\widetilde{x}$ and $\widetilde{x'}$ such that $x \preceq x'$, Farley reduces this so that there are only edges between $\widetilde{x}$ and $\widetilde{x'}$ if x' may be constructed from x by a single split.

The key realisation of construction is that the 1-skeleton of the space $\widetilde{(\succeq x)} := \{y \in \widetilde{X|x \preceq} y\}$ is a boolean lattice, which can be seen when considering the trees in the forest F, where $x = [F,T]$, have 2 states in the elements of $(\succeq x)$, they can be split or not split. As such each element of $(\succeq x)$ can be recognised by which trees of F have been split to construct them, forming a boolean lattice. As such, all cells in the complex can be seen as cubes, making it a cube complex.

Definition 2.2.3. ([Sch], 2.1) Take C a set of cubes such that for each element $c \in C$, $c \cong [0,1]^k$ for some $k \in \mathbb{N}$. Take S a set of isometries such that each $\sigma \in S$ is an isometry $\sigma : F \to F'$, where F, F' are faces of cubes $c, c' \in C$. We restrict S such that there are no maps $\sigma : F \to F'$ where F, F' are faces of the same cube, and such that for any two cubes $c, c' \in C$, there is at most one map $\sigma \in S$ mapping a face of c to a face of c' or vice versa. Then the space $X = (\sqcup_{c \in C} c)/$ is a cube complex, where is the equivalence relation generated by $\{x\ \sigma(x) | \sigma \in S, x \in dom(x)\}$.

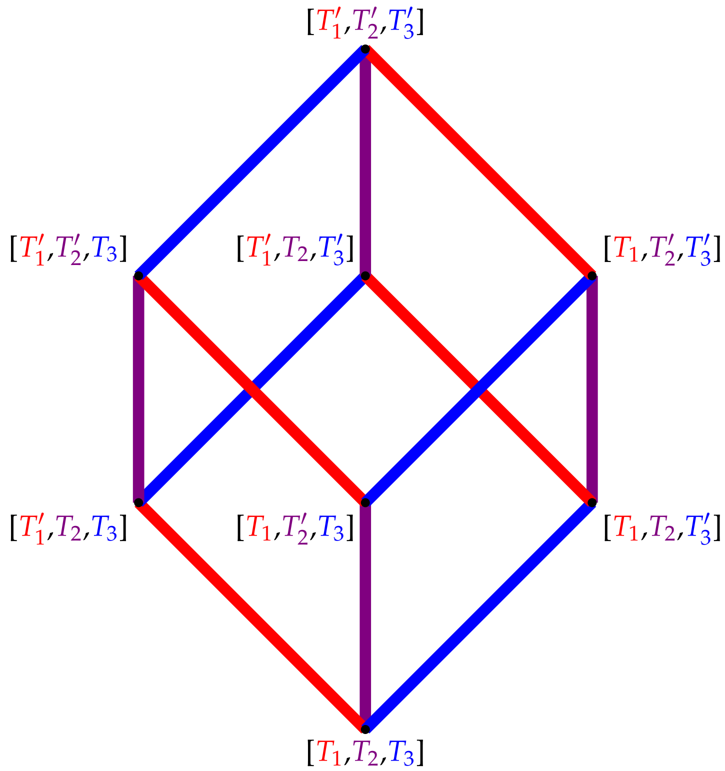

FIGURE 2.7: The lattice of a cube constructed by different possible split sequences of a forest with three trees.

Citation 2.2.4. *([Sch],2.10) A cubical complex X is a complete geodesic metric space, with the metric d such that $d(x_1, x_2) = 1$ where x_1, x_2 are any two distinct vertices in X joined by an edge, if X is either finite dimensional of locally finite.*

Definition 2.2.5. [Gro87] Let (X, d) be a geodesic metric space. If, for any geodesic triangle $\triangle$ in X with comparison triangle $\bar{\triangle}$ in $(\mathbb{R}^2, \bar{d})$, and any two points $x, y \in \triangle$, we have that $d(x, y) \leq \bar{d}(\bar{x}, \bar{y})$, then X is a $CAT(0)$ space.

Lemma 2.2.6. *([Far03], Theorem 1) The Stein-Farley Cube complex $\widetilde{X''}$ is a $CAT(0)$ cube complex.*

To prove that Farley's complex is a $CAT(0)$ space, we will need to use tools that rely on Farley's complex being a geodesic metric space and cube complex, so we will prove this first.

Lemma 2.2.7. *The Stein-Farley Cube complex $\widetilde{X''}$ is a cube complex and a complete geodesic metric space.*

Proof. Consider the height function f on $\tilde{X}''^{(0)}$, where $f(x)$ is the number of trees in the forest of the forest-tree pair of x. By construction, each n-cell in $\tilde{X}''$ has a unique vertex x such that $f(x)$ is the lowest among vertices incident to that cell. As discussed at the beginning of this subsection, we can consider the n-cells above each vertex x as an n-cube, where $n \leq f(x)$. As we can consider all n-cells in this way, each n-cell in $\tilde{X}''$ is an n-cube, and thus $\tilde{X}''$ is a cube complex.

We can see that $\tilde{X}''$ is locally finite as each vertex x has $f(x)$ vertices x' that are joined by edges to x such that $f(x) < f(x')$ (one for each tree in the forest of x) and $f(x) - 1$ vertices x'' connected by edges to x such that $f(x'') < f(x)$ (one for each possible merge of two consecutive trees in the forest of x). As $f(x)$ is always finite, the number of edges incident to x is always finite, and thus $\tilde{X}''$ is locally finite. 2.2.4 then confirms that this is a complete geodesic metric space. □

To prove 2.2.6, we will have to introduce two new concepts for working with CW-complexes. Both of these concepts will prove useful for proofs in later chapters.

Definition 2.2.8. ([Sch], 2.14) For X a geodesic metric space and cube complex, the link $lk(v, X)$ of a vertex $v \in X$ is the spherical complex $\{x \in X | d(x, v) = \epsilon\}$ for some $0 < \epsilon < 1$, with induced structure from X.

The link of a vertex x detects each n-cell incident to x. Each one appears as an $n-1$ simplex in the link, and incidence between these cells is preserved. This makes the link a powerful tool for detecting connectivity locally in a CW-complex.

Definition 2.2.9. A simplicial complex X is considered *flag* if every finite subset of $X^{(0)}$ that is pairwise joined by edges forms the 0-skeleton for a simplex in X. Another way of saying this is that, any time the 1-skeleton for a simplex appears in X, that simplex must also appear.

Citation 2.2.10. *[BH99] If X is a locally finite, simply connected cubical complex and, for each vertex $x \in X^{(0)}$, $lk(x)$ is a simplicial flag complex, then X is a $CAT(0)$ space.*

Before we prove 2.2.6, we must also borrow an observation from Brown that was originally developed for Brown's complex, but applies to both Stein's and Farley's complexes as well.

Citation 2.2.11. *([Bro87], Lemma 4.18) Consider an n-ary forest F, and a set of merges of n consecutive trees of that forest $Y_1, ..., Y_k$. Each merge Y_i may be considered as an n-tuple of consecutive trees. The merges $Y_1, ..., Y_k$ have a lower bound (that is to say, a forest that all of them may construct through repeated merges) if and only if the n-tuples of consecutive trees are pairwise disjoint. Furthermore, they have a highest upper bound that is reached from F by performing each merge Y_i in any order.*

Proof of 2.2.6. We know from 2.2.7 that $\tilde{X}''$ is a cube complex and locally finite. Furthermore, we know $\tilde{X}''$ is simply connected as it is contractible, as the Stein complex $\tilde{X}'$ is contractible and $\tilde{X}''$ is homotopy equivalent to the Stein complex.

We now need to show that the link $lk(x)$ is flag for each vertex x. We consider the link in the following way: Any two ascending edges (edges to vertices x' such that $f(x) < f(x'))$ have a 2-cell between them, as each edge represents a split of a different tree in the forest of x. Furthermore, any set of k ascending edges forms part of the 1-skeleton of a k-cell for the same reason. As such, any time a 1-skeleton of a simplex S appears in $lk(x)$ where each vertex s represents an ascending edge from x, we know the simplex S will appear in $lk(x)$ as well.

Similarly, if the 1-skeleton for a simplex S appears where all the vertices represent descending edges, then we know that the merges represented by those descending edges are pairwise disjoint, as each pair of edges have a 2 -cell between them, which implies a lower bound for those two merges, which implies disjoint by 2.2.11. As all of these edges are pairwise disjoint, they must have a lower bound by 2.2.11, and as such there must be an n-cell between them, so the simplex S must appear in $lk(x)$.

Finally, if we have an ascending edge and a descending edge with a 2-cell between them, we know that the split must be disjoint from the merge. Consider the vertex x', which has the forest produced from the forest of x by the split represented by the ascending edge. The existence of the 2-cell (of which x' is the unique highest vertex) implies a lower bound for the merge represented by the descending edge from x and the inverse of the split represented by the ascending edge of x. By 2.2.11) these merges must be disjoint, and thus the split and merge from the forest of x must be disjoint as well. Thus if there is the 1-skeleton of an $n-1$-simplex S in $lk(x)$, then any ascending and descending edges represented by vertices in S must have disjoint merges and splits, and so we can perform all the merges to reach the unique lowest vertex of the n cube that would imply the existence of S in $lk(x)$, and then construct the cube by performing all the splits and inverse merges in any order. □

When a group is described as "cubulated", it is generally meant that the group has a free, proper action on a finite dimensional $CAT(0)$ cube complex. $\mathbf{F}_n$'s action on the Stein-Farly Cube Complex does not fulfill this criteria, as the complex itself is not finite-dimensional. However, these groups are still frequently described as cubulated due to this action. Wise and Jankievicz have described a group with a similar action as "curiously cubulated" [JW21]. We will use this term to indicate that these actions are different to what is often used by the term "cubulated".

2.3 Bieri-Strebel Groups

Definition 2.3.1. ([BS16], page ii) For an interval of the real numbers I, a multiplicative subgroup of the group of positive real numbers P and a $\mathbb{Z}[P]$ submodule of the real numbers A, we define the Bieri-Strebel group $G(I, A, P)$ as the group of piecewise-linear, orientation preserving homeomorphisms of I such that:

- All gradients are in P.
- There are finitely many breakpoints between the slopes.
- All breakpoints fall in $I \bigcap A$.

Thompson's group $\mathbf{F}$ is the group $G([0,1], \mathbb{Z}[\frac{1}{2}], \langle 2 \rangle)$ and $\mathbf{F}_n$ is the group $G([0,1], \mathbb{Z}[\frac{1}{n}], \langle n \rangle)$.

2.3.1 Algebraic Numbers and Subdivision Polynomials

While Bieri-Strebel groups are very general in their basic definition, it helps to limit the scope of discussion in order to develop more specific results. As such, when we discuss Bieri-Strebel groups, we will generally mean Algebraic Bieri-Strebel groups.

Definition 2.3.2. For a positive algebraic number β, the algebraic Bieri-Strebel group $\mathbf{F}_\beta$ is the Bieri-Strebel group $G([0,1], \mathbb{Z}[\beta], \langle \beta \rangle)$.

Algebraic Bieri-Strebel groups are particularly useful to study as groups of partition pairs, as discussed in 2.1.1. This is due to algebraic numbers association with polynomials, which we may use to build partitions.

Definition 2.3.3. A subdivision polynomial is a polynomial of the form $a_n x^n + a_{n-1}x^{n-1} + ... + a_1 x - 1$, where a_i are all in $\mathbb{Z}_{\geq 0}$.

Lemma 2.3.4. *([Win], Lemma 2.2.1) Any nontrivial subdivision polynomial $P(x)$ has exactly one real root greater than 0, and that root lies between 0 and 1.*

Proof. We can clearly see that $P(0) = -1$ and that $P(1) > 0$ as long as we don't have $a_i = 0\ \forall i$ or that exactly one a_i equals 1 and all else equal 0 (these cases are trivial and can be disregarded). As $P(x) : \mathbb{R} \to \mathbb{R}$ is a continuous function, there must be some value $0 < \beta < 1$ such that $P(\beta) = 0$.

To show that β is unique, we will suppose there exists $\beta' \neq \beta$. Without loss of generality, assume $\beta < \beta'$. As a_i are all non-negative (excluding again the trivial case where $a_i = 0$ for all i), we know that $P(\beta) < P(\beta')$. As such, they cannot be equal and $P(\beta') \neq 0$. □

Please note that in [Win], Winstone expresses his subdivision polynomials in the form $x^n - a_1 x^{n-1} + ... + a_{n-1}x + a_n$. This has two uses for Winstone. This has uses in his case, but is not our prefered method of expressing subdivision polynomials. The primary difference this creates is that our subdivision polynomials have a unique root between 0 and 1, while Winstone's polynomials have a unique root greater than 0.

The reason we call polynomials of this form "subdivision polynomials" is that they form a method of subdividing an interval into $\sum_{i=1}^{n} a_i$ intervals, each with a length in $\langle \beta \rangle$. This is extremely useful for assembling partition pair representations of elements of $\mathbf{F}_\beta$. This can be achieved by writing $P(\beta) = 0$, where β is the positive real root of our subdivision polynomial P. As each subdivision polynomial has a -1 constant term, we can rewrite this equation as $a_n\beta^n + a_{n-1}\beta^{n-1} + ... + a_1\beta = 1$. As we will generally choose 1 as the length of our interval when constructing Algebraic Bieri-Strebel Groups, we can interpret this equation as the sum of the length of these segments (which all have length expressible in the form β^k for some $k \in \mathbb{N}$) is equal to 1. Thus, we can divide the interval into $\sum_{k=1}^{n} a_k$ segments, comprised of a_n segments of length β^n, a_{n-1} segments of length β^{n-1} and so on. We can find such a partition for any β that is the positive real root of a subdivision polynomial.

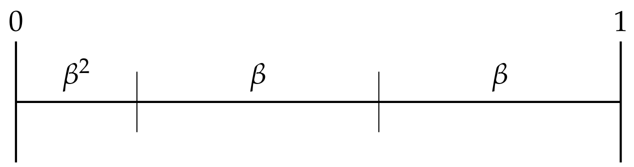

FIGURE 2.8: A partition of the unit interval based on the subdivision polynomial $x^2 + 2x - 1$

This method of partition can be applied to any of the subintervals created by the previous partition in a manner similar to the repeated bisection seen when forming partition pairs for **F**. As long as the same polynomial partition is consistently used, all created subintervals will have length in $\langle\beta\rangle$. This is extremely useful when forming elements of $\mathbf{F}_n$ as the gradient of the piece of any element that maps β^i to β^j will be β^{j-i}, which will always be in $\langle\beta\rangle$, the slope group for $\mathbf{F}_\beta$.

Applying the methods of subdivision polynomials to our previous examples allows us to see Algebraic Bieri Strebel groups as a natural generalisation of $\mathbf{F}$ and $\mathbf{F}_n$. As partitions in $\mathbf{F}$ split an interval into 2 equal parts, and partitions in $\mathbf{F}_n$ split an interval into n equal parts, we can assign them the subdivision polynomials $2x-1$ and $nx-1$ respectively. Thus we can see $\mathbf{F}_n$ as the "linear" Bieri-Strebel groups, that is to say the algebraic Bieri-Strebel groups with linear subdivision polynomials.

2.3.2 Winstone's Tree Pairs

An important result for our understanding of Algebraic Bieri-Strebel groups, and those with quadratic subdivision polynomials in particular, was the development of tree-pair presentations for these groups. Tree-pairs were originally introduced for the group with subdivision polynomial x^2+x-1, commonly written as $\mathbf{F}_\tau$, by Burillo, Nucinkis and Reeves in [BNR21]. This paper has three major conclusions for the development of tree-pair representations for algebraic Bieri-Strebel groups. The first is that the different lengths of intervals, all of which are of the form β^k for the group $\mathbf{F}_\beta$, correspond to different depths on the tree, and as partitions in Bieri-Strebel groups create intervals of different lengths, the carets representing those partitions should have legs reaching to different depths.

The Second realisation follows on from the first. As partitions are no longer creating multiple intervals of the same length, the order that those intervals appear in is important. To represent this in tree pairs, we have to make use of multiple different carets, which order their legs of different length in different ways.

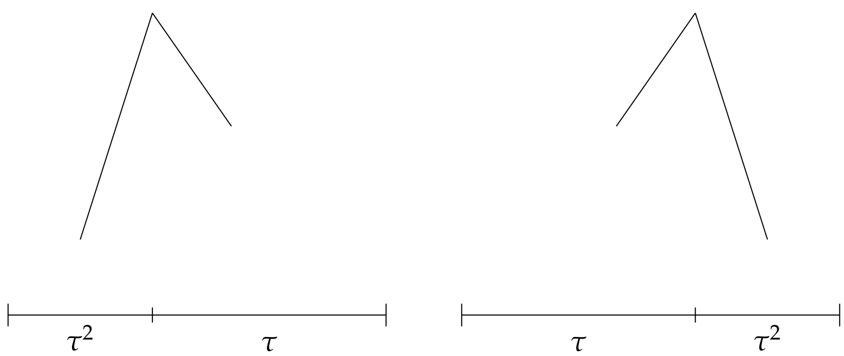

FIGURE 2.9: The two caret types for the Bieri-Strebel group corresponding to x^2+x-1, known as $\mathbf{F}_\tau$, and the interval partitions they each represent. We consider the long leg at depth 2 and the short leg at depth 1.

The final realisation is an extension of the second. With two different types of carets with which to build trees, it is now possible to build two trees with different carets that represent the same interval partition. As such, when we wish to construct equivalence classes of tree-pair diagrams as we did in 2.1.2, we have to include in the equivalence relation the ability to switch between equivalent subtrees. We refer to pairs of equivalent trees built with different carets as caret relations.

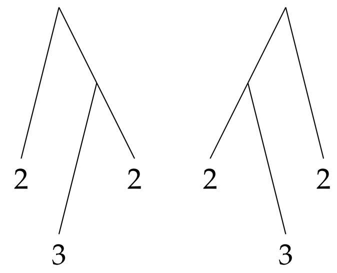

FIGURE 2.10: Two equivalent trees built from carets in the $\mathbf{F}_\tau$ tree-pair presentation.

Burillo, Nucinkis and Reeves are able to use this tree-pair representation to construct an infinite presentation for $\mathbf{F}_\tau$. The presentation contains two different infinite families of generators, corresponding to the two different caret types in the tree-pair representation. The presentation also has a similar set of relations to the **F** presentation in 2.1, and also has a set of relations corresponding to the tree equivalence depicted in 2.10. Combining all of this, the presentation constructed in [BNR21] can be written as

$$\mathsf{F}_\tau = \langle x_0, y_0, x_1, y_1, ... | a_j b_i = b_i a_{j+1} \text{ for } a, b \in \{x, y\}, i < j; x_i x_{i+1} = y_i^2 \rangle \tag{2.4}$$

When considering x_i and y_i as PL-homeomorphisms, we may write them in the following way

$$\begin{aligned} x_i(n) &= \begin{cases} n & \text{for } 0 \leq n \leq 1 - \tau^i, \\ \tau^{-2}n - \tau^{-1}(1-\tau^i) & \text{for } 1 - \tau^i \leq n \leq 1 - \tau^i + \tau^{i+4}, \\ n + \tau^{i+3} & \text{for } 1 - \tau^i + \tau^{i+4} \leq n \leq 1 - \tau^{i+1}, \\ \tau n + \tau^2 & \text{for } 1 - \tau^{i+1} \leq n \leq 1, \end{cases} \\ y_i(n) &= \begin{cases} n & \text{for } 0 \leq n \leq 1 - \tau^i, \\ \tau^{-1}n - \tau^{-1}(1-\tau^i) & \text{for } 1 - \tau^i \leq n \leq 1 - \tau^{i+1}, \\ \tau n + \tau^2 & \text{for } 1 - \tau^{i+1} \leq n \leq 1. \end{cases} \end{aligned} \tag{2.5}$$

Burillo, Nucinkis and Reeves were also able to reduce this infinite presentation down to a finite presentation containing 4 generators and 10 relations. As with the finite presentation for **F**, this presentation is cumbersome to use in comparison to the infinite presentation, but can be seen at ([BNR21], Section 4).

In his thesis [Win], Winstone was able to generalise these results to the quadratic Bieri-Strebel groups, that is to say the algebraic Bieri-Strebel groups with a quadratic associated subdivision polynomial. In doing so, Winstone developed two theorems that will be important for this thesis, as well as an infinite presentation for a subset of these groups. Quite possibly the most important result from [Win] is the following:

Citation 2.3.5. *([Win], Theorem 1.2.3) For a quadratic Bieri-Strebel group $\mathbf{F}_\beta$ with subdivision polynomial of the form $ax^2 + bx - 1$, $\mathbf{F}_\beta$ has a well defined tree-pair representation if and only if $a \leq b$.*

When considering tree-pair representations of other quadratic Bieri-Strebel groups, we have to consider the rapid increase of caret types as the coefficients of the subdivision polynomial increase. As depicted in 2.9, the subdivision polynomial dictates the possible interval partitions, and therefore the possible caret types. The important observation is that the number of intervals in a partition for $\mathbf{F}_\beta$ is equal to $a+b$, where $\mathbf{F}_\beta$ has the subdivision polynomial ax^2+bx-1. We would expect any caret to have a legs of length 2 and b legs of length 1 (to correspond to the a intervals of length β^2 and the b intervals of length β in a partition for $\mathbf{F}_\beta$. As such, there are $\binom{a+b}{a}$ possible caret types in the tree pair representation for $\mathbf{F}_\beta$.

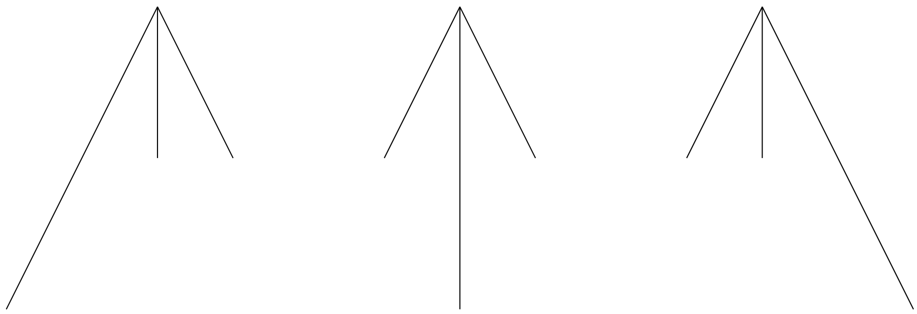

FIGURE 2.11: The three possible caret types in $\mathbf{F}_\beta$ with subdivision polynomial x^2+2x-1.

Winstone's second major result simplifies the large variety of caret types for quadratic Bieri-Strebel groups, making tree-pair representations much more manageable.

Citation 2.3.6. *([Win], Remark 35) There exist two caret types in each well-defined tree pair representation of a quadratic Bieri-Strebel group such that for any tree within that tree-pair representation, an equivalent tree may be constructed using only carets of those two types*

This result is not just important for simplifying the tree-pair representations, but also leads into the general presentation for quadratic Bieri-Strebel groups that have tree-pair representations. Reducing each tree pair representation once again allows us to generate $\mathbf{F}_\beta$ from two infinite families of generators. Using these caret types, and the caret relations between them, Winstone was able to construct infinite presentations for a large number of quadratic Bieri-Strebel groups.

Citation 2.3.7. *([Win], Theorem 1.2.5) Let $\mathbf{F}_\beta$ be an algebraic Bieri-Strebel group with associated subdivision polynomial ax^2+bx-1. If $a \leq b$, then $\mathbf{F}_\beta$ has the infinite presentation:*

$$\boldsymbol{F}_\beta = \langle x_0, y_0, x_1, y_1, ... | R_1, R_2 \rangle \tag{2.6}$$

where R_1 and R_2 are the relations:

$$\begin{aligned} &R_1 : f_j g_i = g_i f_{j+a+b-1} \text{ for } g, f \in \{x, y\}, i < j \\ &R_2 : x_{i+a} x_{i+a+1} ... x_{i+2a-1} x_i = y_i y_{i+1} ... y_{i+a-1} y_i \text{ for all } i \geq 0 \end{aligned} \tag{2.7}$$

If we so desired, we could reduce these infinite presentations to finite presentations. However, we will be using geometric methods to prove that each of these groups have the F_∞ property in chapter 3, which will imply the existence of a finite presentation. We have no immediate use for the finite presentations and so will leave their existence implied rather than explicitly writing them out.

For quadratic Bieri-Strebel groups with subdivision polynomial $ax^2 + bx - 1$, $a > b$, not much is known outside of generalities for all Bieri-Strebel groups or other large subsets of such groups. Winstone ([Win], Theorem 1.2.4) concludes that each such group has an element g such that g cannot be represented either as a partition pair or as a tree pair. Owen Tanner ([Tan23], Theorem 3) was able to show that all such groups are finitely generated.

Of particular interest are the groups with polynomial of the form $(n-1)nx^2 + x - 1$ for $n \geq 2$. This polynomial can be factorised as $(nx-1)((n-1)x+1)$ and therefore has the unique positive root $\frac{1}{n}$. By the definition of Algebraic Bieri-Strebel group given in 2.3.2, we would expect the group with this subdivision polynomial to have slopes in $\langle \frac{1}{n} \rangle$ and breakpoints in $\mathbb{Z}[\frac{1}{n}]$, which would make it the same as $\mathbf{F}_n$. Notably, $\mathbf{F}_n$ can be expressed as the algebraic Bieri-Strebel group with subdivision polynomial $nx - 1$, which when used to compute interval divisions and tree pairs in the style of Winstone produces the tree-pairs described in 2.1.3. This leads us to make the following conjecture:

Conjecture 2.3.8. *A well defined tree-pair for an algebraic Bieri-Strebel group $\boldsymbol{F}_\beta$ can only be derived from an associated subdivision polynomial $P(x)$ if $P(x)$ is the minimal polynomial with the root β among polynomials of the form $a_n x^n + ... + a_1 x - 1$, $a_i \in \mathbb{Z}_{\geq 0}$.*

Among linear and quadratic subdivision polynomials, we can see evidence for this conjecture. In particular, we know that all linear subdivision polynomials take the form $nx - 1$ for some x (with the polynomial having root $\frac{1}{n}$. Should a quadratic subdivision polynomial have the root $\frac{1}{n}$, then by the factor theorem, that polynomial must have $(x - \frac{1}{n})$ as a factor, which can be expressed as $nx - 1$ when working with integer polynomials. We can now use what we know about the general form of subdivision polynomials to determine facts about other factors of the quadratic polynomial. Since the polynomial is quadractic, it can only be factorised as the product of two linear polynomials, meaning we can write

$$ax^2 + bx - 1 = (nx-1)(cx+d) \tag{2.8}$$

Immediately, we can conclude that $d = 1$, as $-1 * 1 = -1$. Similarly, we can conclude $c > 0$, as otherwise $a \leq 0$, which would mean the quadratic polynomial would not be a subdivision polynomial. That leaves us with the two equations $a = nc$ and $b = n - c$. As n and c are both positive integers, we can therefore conclude that $a > b$. As such, any time a quadratic subdivision polynomial $ax^2 + bx - 1$ shares its unique positive root with a linear subdivision polynomial $nx - 1$, we have that $a > b$ and it therefore falls outside the set of subdivision polynomials for which Winstone derived tree pairs.

We note that this statement cannot be an if and only if. That is to say that it is not necessarily the case that a subdivision polynomial that does not share a root with a sudivision polynomial of lesser degree will have a tree pair representation. As a brief example, the polynomial $3x^2 + x - 1$ has the root $\frac{\sqrt{13}-1}{6}$, which is clearly irrational and therefore cannot be the root of any linear polynomial, but by 2.3.5, this subdivision polynomial does not have a well defined tree pair representation.

2.3.3 Higher Order Polynomials

For algebraic Bieri-Strebel groups with higher order subdivision polynomials, little is known outside of general results for Bieri-Strebel groups, such as Tanner's finite generation result [Tan23]. A caret relation is known for the group corresponding to the polynomial $x^3 + x - 1$, but it is unknown if this allows a well-defined tree-pair representation.

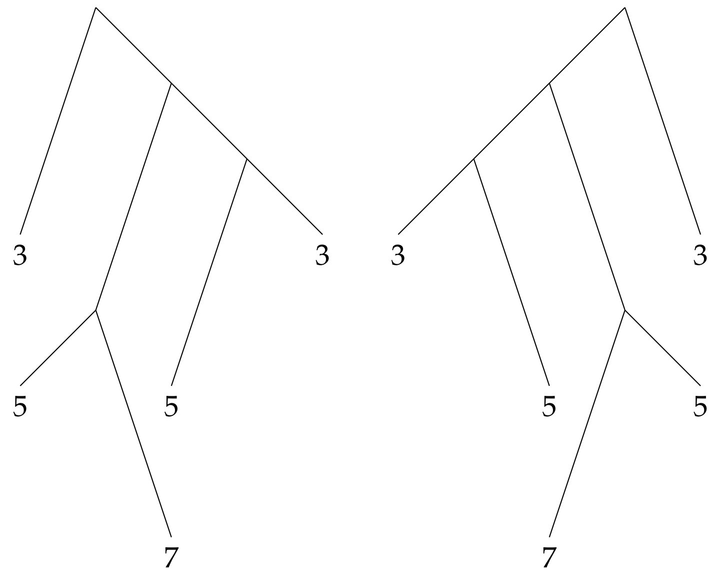

FIGURE 2.12: The caret relation for $x^3 + x - 1$. It differs from quadratic caret relations in that each tree uses both caret types.

While determining well-defined tree pairs for higher order Bieri-Strebel groups is an as-yet unsolved problem, we are able to determine some cases where a tree-pair representation based on the subdivision polynomial does not work.

Theorem 2.3.9. *The Bieri-Strebel group with subdivision polynomial $ax^{2n} + bx^n - 1$ cannot have a well defined tree-pair representation.*

To begin to prove this, we must first issue a correction to a theorem in Winstone's Thesis

Citation 2.3.10. *([Win], Theorm 2.3.3) For all $0 < p \in \mathbb{Z}[\beta]$, where β is the root of an n-th degree polynomial, we can write:*

$$p = b_0 + b_1\beta + ... + b_{n_1}\beta^{n-1}$$

where $b_i \in \mathbb{Z}_{\geq 0}$ for all i

Counterexample. Consider the polynomial $x^4 + x^2 - 1$, with the root $\sqrt{\tau}$. We will take $p = 1 - \sqrt{\tau}$. To work within Winstone's framework (where he uses subdivision polynomials with reciprocal root to ours, as noted in 2.3.1), we take $\beta = \sqrt{\tau}^{-1}$, which means $p = 1 - \beta^{-1}$, and as we have that $\beta^4 - \beta^2 - 1 = 0$, we can say that $\beta^{-1} = \beta^3 - \beta$, and therefore $p = 1 - \beta^3 + \beta$.

Following the methodology of Winstone's proof, the next step is to write the coefficients of p as a vector $\underline{p}$ and create the matrix A based on the polynomial $x^4 - x^2 - 1$

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}; \underline{p} = \begin{pmatrix} -1 \\ 0 \\ 1 \\ 1 \end{pmatrix}$$

To simplify a later segment, we will calculate

$$A\underline{p} = \begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \end{pmatrix}$$

The core claim of Winstone's proof is that there is $\mathbf{N}$ such that for $N \geq \mathbf{N}$, $A^N\underline{p}$ has only positive entries. We will use the constructed example to show this is not the case. First, we will calculate

$$A^4 = \begin{pmatrix} 2 & 0 & 1 & 0 \\ 0 & 2 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

Here we wish to emphasise the structure of the matrix, with alternating entries in each row and column being empty. We will consider a general matrix of this form and multiply it by A:

$$\begin{pmatrix} a & 0 & b & 0 \\ 0 & c & 0 & d \\ e & 0 & f & 0 \\ 0 & g & 0 & h \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & a & 0 & b \\ c+d & 0 & c & 0 \\ 0 & e & 0 & f \\ g+h & 0 & g & 0 \end{pmatrix}$$

Which has a similar alternating pattern, but all the nonzero entries are now zero and all the zero entries are now nonzero. Multiplying this matrix by A again gives us.

$$\begin{pmatrix} 0 & a & 0 & b \\ c+d & 0 & c & 0 \\ 0 & e & 0 & f \\ g+h & 0 & g & 0 \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} a+b & 0 & a & 0 \\ 0 & c+d & 0 & c \\ e+f & 0 & e & 0 \\ 0 & g+h & 0 & g \end{pmatrix}$$

which is of the same structure as A^4. As such, we can see that A^{2N} will have the same structure as A^4 for $N \geq 2$, an A^{2N+1} will have the same structure as A^5, for all $N \geq 2$. We can now show that neither of these structures can create a vector $A^N\underline{p}$ such that all entries in the vector are nonnegative. We can simplify by using $A\underline{p}$ as the vector, from which we can clearly see:

$$\begin{pmatrix} a & 0 & b & 0 \\ 0 & c & 0 & d \\ e & 0 & f & 0 \\ 0 & g & 0 & h \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \end{pmatrix} = \begin{pmatrix} b \\ -d \\ f \\ -h \end{pmatrix}$$

$$\begin{pmatrix} 0 & a & 0 & b \\ c & 0 & d & 0 \\ 0 & e & 0 & f \\ g & 0 & h & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \end{pmatrix} = \begin{pmatrix} -b \\ d \\ -f \\ h \end{pmatrix}$$

As A is a nonnegative matrix, all entries in A^N will be nonnegative for all N, and so $A^N \underline{p}$ will always contain two negative numbers for $N \geq 5$. □

It is straightforward to see that the alternating nature of this counterexample will generalise to any polynomial of the form $ax^4 + bx^2 - 1$, as the resulting polynomial A will just be

$$\begin{pmatrix} 0 & 1 & 0 & 0 \\ a & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ b & 0 & 0 & 0 \end{pmatrix}$$

and as such we have

$$A^4 = \begin{pmatrix} a+b^2 & 0 & b & 0 \\ 0 & a+b^2 & 0 & b \\ ab & 0 & a & 0 \\ 0 & ab & 0 & a \end{pmatrix}$$

which has the same structure as A^4 from the counterexample. As such, the alternating patten will repeat. We can then construct a $\underline{p}$ such that $A^N \underline{p}$ always has a negative number.

What is less immediately obvious is that this pattern extends to any $a_n x^n + ... + a_1 x - 1$ with the condition that the set of $\{i_1, ..., i_k\}$ where a_{i_j} is nonzero is not coprime. In this case, we have that A^N has the same structure as $A^N + k$, where $k = \mathsf{hcf}(i_1, ..., i_k)$. However, this can be seen with elementary matrix multiplication.

The next step is to show the following:

Lemma 2.3.11. *For any β as the root of a quadratic subdivision polynomial $ax^2 + bx - 1$, $\sqrt{\beta}$ is not in $\mathbb{Z}[\beta]$.*

Proof. By definition, elements of $\mathbb{Z}[\beta]$ have the form $c_0 + c_1\beta + c_2\beta^2 + \ldots$. However, because β is the root of the quadratic polynomial $ax^2 + bx - 1$, we can write $\beta^2 = \frac{1}{a} - \frac{b}{a}\beta$. As such, any term with a β^2 or higher power can be substituted, allowing us to express any element in $\mathbb{Z}$ in the form $c + d\beta$.

Now, suppose $\sqrt{\beta} \in \mathbb{Z}[\beta]$. We can then write $\sqrt{\beta} = c + d\beta$. From the property of the square root, we can then write

$$
\begin{aligned}
(c+d\beta)(c+d\beta) &= \beta \\
c^2+2cd\beta+d^2\beta^2 &= \beta \\
c^2+2cd\beta+d^2(\frac{1}{a}-\frac{b}{a}\beta) &= \beta \\
c^2+\frac{d^2}{a}+(2cd-\frac{d^2b}{a})\beta &= \beta
\end{aligned} \tag{2.9}
$$

As β is irrational, we know that the only way for this equation to be true is if the constant coefficients on each side are equal, and so are the β coefficients. This leaves us with the simultaneous equations

$$
\begin{aligned}
c^2+\frac{d^2}{a} &= 0 \\
2cd-\frac{d^2b}{a} &= 1
\end{aligned} \tag{2.10}
$$

From the first equation c^2 and d^2 are squares of numbers in $\mathbb{Z}$, and therefore must be positive, and a is positive by the definition of a subdivision polynomial (as the subdivision polynomial in question is quadratic, $a \neq 0$ and so the equation is well defined). As such, both terms of the sum must be nonnegative and so it can only hold if $c = d = 0$. However, setting $c = d = 0$ in the second equation results in the equation reducing to $0 = 1$. As such we have a contradiction and so our assumption that $\sqrt{\beta} \in \mathbb{Z}$ must be incorrect. □

We can also add this brief corollary:

Corollary 2.3.12. *For any β as the root of a quadratic subdivision polynomial ax^2+bx-1, $n \in \mathbb{N}$, $\sqrt[2n]{\beta}$ is not in $\mathbb{Z}[\beta]$.*

Proof. $\mathbb{Z}[\beta]$ is closed multiplicatively, so if $\sqrt[2n]{\beta}$ was in $\mathbb{Z}[\beta]$, that would imply $(\sqrt[2n]{\beta})^n = \sqrt{\beta}$ is in $\mathbb{Z}[\beta]$, which we know is not true by 2.3.11. □

In the discussion of Winstone's tree pairs in 2.3.2, we mentioned that carets in the tree-pair representation of the Bieri-Strebel group with subdivision polynomial $ax^2 + bx-1$ will have a legs of length 2 and b legs of length 1. This is derived from the polynomial itself, using it to partition the unit interval. As such, we will begin by assuming that tree-pairs for higher order polynomials will be derived similarly.

Consider the tree pair representation formed from the non-coprime power subdivision polynomial $P_1 = P(x) = a_n x^{kn} + a_{n-1}x^{k(n-1)} + ... + a_1 x^k - 1$, where $k \in \mathbb{N}$. Based on 2.3.2, we would expect carets in this tree-pair representation to have a_n legs of length k_n, a_{n-1} legs of length $k(n-1)$ and so on. We now consider the polynomial $P_2 = P(\sqrt[k]{x}) = a_n x^n + a_{n-1}x^{n-1} + ... + a_1 x - 1$. The carets in the tree-pair representation for this polynomial will have a_n legs of length n, a_{n-1} legs of length $n-1$ and so on. From here, we can consider the map i from the set of carets for $P(x)$ to the set of carets for $P(\sqrt[k]{x})$, which maps each caret to the caret obtained by dividing the length of each leg by k. We can determine the number of possible carets in the caret set for $P(x)$ by determining the possible distinct orderings for the set of legs each caret has, which is a_n legs of length kn, a_{n-1} legs of k_{n-1}, and so on. Similarly, the set of possible carets for $P(\sqrt[k]{x})$ is determined by the possible orderings of a_n carets of length n, $a_n - 1$ legs of length $n-1$ etc. As such, the set of legs being ordered for

$P(x)$ consists of the same number of subsets $S[P_1]_i$ containing legs of the same length as the set of legs being ordered for $P(\sqrt[k]{x})$, and these subsets can be ordered in such a way that $|S[P_1]_i| = |S[P_2]_i|$ for all i. Thus the number of possible orderings for legs on the $P(x)$ carets is the same as the number of possible orderings for the $P(\sqrt[k]{x})$, so the map is surjective, and if two carets in $P(x)$ are mapped to the same caret in $P(\sqrt[k]{x})$, then they must have the same ordering of legs, and are thus the same caret, and so the map is injective.

We now induce a map i^* to tree pairs in the tree pair representation of $P(x)$ by applying i to each caret in the tree pair individually, preserving adjacency between carets. The path to each leaf from the root of the tree can be broken down into legs for carets, and if i^* applied to tree pairs just applies the map i to each caret in the tree, then the depth of each leaf will be mapped from n to $\frac{n}{k}$.

We consider the path from the root of a tree to an arbitrary leaf in a tree pair diagram in the Bieri-Strebel group with $P(x)$ as subdivision polynomial. each leg in this path is of length $k * b$ for some $1 \leq b \leq n$. Thus the full path can be expressed as $\Sigma_{i=1}^{m} k * b_i$ where m is the number of legs in the path. When we apply $i*$ to this tree pair, each leg in this path is mapped from a leg of length $k * b_i$ to a leg of length b_i. Thus the length of this path is just $\Sigma_{i=1}^{m} b_i$. As each b_i is a natural number, so is $\Sigma_{i=1}^{m} b_i$, thus the length of the path from the root to the leaf is a natural number, so the depth of each leaf must be a natural number.

Consider now the preimage under i^* of a tree pair $[T_1, T_2]$ in the space of tree pair diagrams for $P(\sqrt[k]{x})$. The depth of the ith leaf in $[T_1, T_2]$ has depth d_i and so any tree pair $[T_1', T_2']$ in the preimage of $[T_1, T_2]$ under i must be such that the depth of the ith leaf is of depth $k * d_i$. The equivalence relation between tree pair relations (in particular the caret relations) dictate that two tree pairs are in the same equivalence class if they have the same leaf depths in the same order, and so the map i^* must map equivalence classes to equivalence classes and cannot map two different equivalence classes into the same equivalence class. We can see that it is surjective by considering the map i_*, that multiplies each leaf depth by k. Clearly $i^* i_*$ is the identity map on the tree-pair representation for $P(\sqrt[k]{x})$, so for each $[T, T]$ in the tree pair representation for $P(\sqrt[k]{x})$ there must be a tree pair $[T', T']$ in the tree pair representation for $P(x)$ such that $i_*([T, T]) = [T', T']$, $i^*([T', T']) = [T, T]$, and hence i^* is surjective.

We can also see that i^* is a homomorphism with regard to tree pair composition. We can see this by performing tree-pair composition simultaneously in $P(x)$ and $P(\sqrt[k]{x})$. Whenever we add a caret c while working in $P(x)$, we add $i(c)$ to the same leaf in $P(\sqrt[k]{x})$. Thus i^* is an isomorphism between the equivalence classes in the tree-pair representations for $P(x)$ and $P(\sqrt[k]{x})$. This leads us to the following lemma:

Lemma 2.3.13. *consider β as the unique positive root of $P(x) = a_n x^n + ... + a_1 x - 1$ and $\sqrt[k]{\beta}$ as the root of $P(x^k) = a_n x^{kn} + ... + a_1 x^k - 1$. Suppose the tree pair representation based of $P(x)$ is a well-defined tree pair representation for $\mathbf{F}_\beta$. Either $\mathbf{F}_{\sqrt[k]{\beta}} \cong \mathbf{F}_\beta$, or the tree-pair based on $P(x^k)$ is not a well-defined tree-pair representation for $\mathbf{F}_{\sqrt[k]{\beta}}$.*

From here, we wish to construct an element we know is in $\mathbf{F}_{\sqrt[2n]{\beta}}$ but cannot be in $\mathbf{F}_\beta$. We will construct an element that is in $\mathbf{F}_{\sqrt{\beta}}$, and by a similar argument to the proof of 2.3.12, must be in $\mathbf{F}_{\sqrt[2n]{\beta}}$ as well.

We first create a partition pair which we know cannot be in $\mathbf{F}_\beta$. From 2.3.11, we know that if a partition pair contains $\sqrt{\beta}$ as a nontrivial breakpoint, then that partition pair cannot be an element of $\mathbf{F}_\beta$. From here, the simplest solution would be to have $\sqrt{\beta}$ as the sole breakpoint in the first partition and then have $1-\sqrt{\beta}$ as the sole breakpoint of the second partition. The issue with this is that the resulting map would map an interval of length $\sqrt{\beta}$ to an interval of length $1-\sqrt{\beta}$. This would require $\frac{1-\sqrt{\beta}}{\sqrt{\beta}}$ to be in $\langle\sqrt{\beta}\rangle$, which is challenging to prove. Instead, we will construct a partition pair that has only slopes in $\langle\beta\rangle$. Since $\langle\beta\rangle \subseteq \langle\sqrt{\beta}\rangle$, we can then guarantee our partition pair is in $\mathbf{F}_{\sqrt{\beta}}$.

We create our partition pair in the following way. First we take two copies of the interval with $\sqrt{\beta}$ as the sole breakpoint. From here, we subdivide the interval of length $\sqrt{\beta}$ in each partition in different ways. Each interval will now be divided into a intervals of length $\beta^2 * \sqrt{\beta}$ and b intervals of length $\beta * \sqrt{\beta}$, where β is the root of ax^2+bx-1, implying

$$\begin{aligned} a\beta^2 + b\beta &= 1 \\ a\beta^2 * \sqrt{\beta} + b\beta * \sqrt{\beta} &= \sqrt{\beta} \end{aligned} \tag{2.11}$$

The ordering of the subintervals in the partitions is arbitrary, aside from 2 points. First, the sole interval of length $1-\sqrt{\beta}$ must be the ith interval in each partition for some i, and the partitions must not be identical.

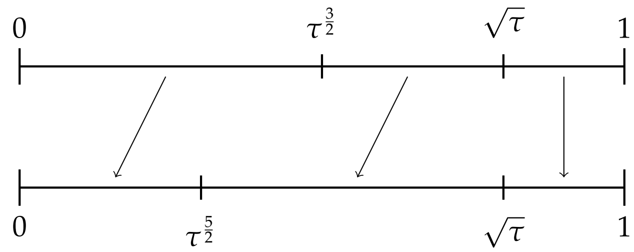

FIGURE 2.13: An example of a partition pair in $\mathbf{F}_{\sqrt{\tau}}$ with the breakpoint $\sqrt{\tau}$ excluding it from $\mathbf{F}_\tau$. The slopes in this map have gradients τ, τ^{-1} and 1 from left to right

As seen in the example 2.13, maps constructed in this way have gradients in $\langle\sqrt{\beta}\rangle$ (indeed, they are in $\langle\beta\rangle$), but as they have $\sqrt{\beta}$ as a breakpoint, these maps are in $\mathbf{F}_{\sqrt{\beta}}$ but not in $\mathbf{F}_\beta$. Thus, we can conclude that $\mathbf{F}_\beta$ is not the same group as $\mathbf{F}_{\sqrt{\beta}}$ and so the tree pair representation based on $P(x^2)$ is not a well-defined tree pair representation for $\mathbf{F}_{\sqrt{\beta}}$ by 2.3.13. Similarly, this element is in $\mathbf{F}_{\sqrt[2n]{\beta}}$, and so $P(x^{2n})$ is not a well-defined tree pair representation for $\mathbf{F}_{\sqrt[2n]{\beta}}$.

Finally, we need to preclude the possibility of a different tree pair representation. For this we rely on our counterexample to 2.3.10. What we have shown with the counterexample (and its generalisation) is that we can form an interval in a partition in $\mathbf{F}_{\sqrt{\beta}}$ with a length that cannot be expressed in the form $a_4\beta^{k+4} + a_3\beta k + 3 + a_2\beta^{k+2} + a_1\beta^{k+1}$ via repeated subdivision (the substitution used in Winstone's proof). However, were we constructing a tree to represent this partition, such an interval would have to be expressed in this way. As such, these intervals cannot be expressed as part of a tree pair presentation and thus we have our proof that Bieri-Strebel groups with associated polynomial $ax^4 + bx^2 - 1$ cannot have a well defined tree-pair representation. Combining this with 2.3.12 is our proof of 2.3.9.

While this seems challenging to generalise in its entirety, we would like to offer this conjecture regarding non-coprime power polynomials:

Conjecture 2.3.14. *The Bieri-Strebel group with associated subdivision polynomial* $a_n x^{kn} + a_{n-1} x^{k(n-1)} + ... + a_1 x^k - 1, k > 1$ *with not all* $a_i = 0$ *does not have a well-defined tree-pair representation.*

Chapter 3

F_∞ and a Cubulation for Bieri-Strebel Groups

3.1 F_∞ for Bieri-Strebel Groups

In this section, we will demonstrate the F_∞ property for certain algebraic Bieri-Strebel groups. This will be accomplished via the construction of a complex similar to Brown's complex for Thompson-Higman groups 2.2.1, followed by the application of Brown's criterion on the complex. This method is similar to the method used by Cleary [Cle95]. The primary difference between our approach and Cleary's is that Cleary defines his complex through partition pairs and has to prove that all elements of the groups he works with (in particular, Cleary wrote his proof regarding the Bieri-Strebel group with subdivision polynomial $x^2 + 2x - 1$) can be represented by partition pairs. Our complex is constructed through forest-tree pairs (making our complex closer to Brown's complex) which allows us to take advantange of Winstone's tree-pair theorem 2.3.5 to prove the F_∞ property for many Bieri-Strebel groups at once. The ultimate goal is therefore to prove the following theorem:

Theorem 3.1.1. *If $\mathbf{F}_\beta$ is an algebraic Bieri-Strebel group with a well-defined tree-pair presentation, then $\mathbf{F}_\beta$ is of type F_∞.*

and by combining it with Winstone's tree-pair theorem, the following corollary:

Corollary 3.1.2. *If $\mathbf{F}_\beta$ is an algebraic Bieri-Strebel group with subdivision polynomial $ax^2 + bx - 1$, $a \leq b$, then $\mathbf{F}_\beta$ has the F_∞ property.*

3.1.1 Brown's Criterion

Brown's criterion is an alternative method of proving the F_∞ property for groups rather than using the definition provided in 2.1.3. It is most useful when one is able to construct a complex for a given group to act on, but the complex is not finite in all dimensions (in the case of a $K(G, 1)$), or not finite in all dimensions when quotiented by the G-action (in the case of an EG).

Definition 3.1.3. ([Bro87], section 2) A filtration of a G complex X is a family $\{X_\alpha\}_{\alpha \in D}$ of subcomplexes of X such that:

- Each subcomplex X_α is invariant under the action of G (ie: $g \circ x \in X_\alpha$ for all $g \in G, x \in X_\alpha$

- D is a directed set. That is to say, D has the preorder $\leq$ such that any two elements have an upper bound (as mentioned when discussing partial orders in 2.2.1).
- For all $\alpha, \beta \in D$ such that $\alpha \leq \beta$, we have that $X_\alpha \subseteq X_\beta$.
- $X = \bigcup_{\alpha \in D} X_\alpha$

Brown initially wrote his criterion as the following:

Citation 3.1.4. *([Bro87], Theorem 2.2). Let X be a G-CW-complex such that $\pi_i(X) = 0$ for $i < n$ and that the stabiliser of any p-cell $x \in X$ is of type FP_{n-p}. Take a filtration $\{X_\alpha\}_{\alpha \in D}$ such that each X_α has a finite n-skeleton mod G. Then G is of type FP_n if and only if the direct system of reduced homology groups $\{\tilde{H}_i(X_\alpha)\}$ has a trivial direct limit.*

Brown also presents a corollary to this criterion, which is more directly applicable to our method, and is indeed the result he uses when proving F_∞ for the Thompson-Higman groups.

Citation 3.1.5. *([Bro87], Corollary 3.3) Let X be a contractible G-complex such that the stabiliser of every cell is of type F_∞. Let $\{X_j\}_{j \geq 1}$ be a filtration such that each X_j is finite mod G. Suppose the connectivity of the quotient space X_{j+1}/X_j tends to ∞ as j tends to ∞. Then G is of type F_∞.*

Note that the filtration in this is stricter than requiring a directed set. Instead of a preorder, the filtration in this corollary requires a full order and uses the set of natural numbers $\mathbb{N}$ as an index set for the set of subcomplexes forming the filtration. This results in a full order of subcomplexes such that $X_i \subset X_j$ whenever $i \leq j$.

3.1.2 Construction of the Brown-Cleary Complex

The Brown-Cleary complex is an adaptation of both Brown's complex for Higman-Thompson groups (as described in ([Bro87], section 4)) and Cleary's complex for Bieri-Strebel groups (as described in [Cle95]), which itself was an adaptation of Brown's complex). It is more similar structurally to Cleary's complex, but uses an adapted form of Brown's forest-tree pairs. We shall construct these forest-tree pairs from a generalisation of Winstone's tree pairs for quadratic Bieri-Strebel groups. For the algebraic Bieri-Strebel group $\mathbf{F}_\beta$, each tree in a tree-pair representation for $\mathbf{F}_\beta$ will be built from carets in the set of caret types, C_β. The caret set C_β will have the following assumptions:

- Each leg of a caret will have a length in the set $\mathbb{N}$ of natural numbers.
- All the caret types $c \in C_\beta$ will have the same number of legs of a given length k, and thus the same total number of legs. Caret types shall be differentiated by the ordering of legs of different length.

From here, the tree-pair representation shall be constructed as in 2.3.2: we consider any ordered pair of rooted trees built with carets from C_β a valid tree-pair as long as both trees have the same number of leaves. We will then quotient out by the equivalence relation that considers two tree-pairs equivalent if one may be built from the other via the addition or removal of redundant carets and/or the replacement of a subtree with an equivalent subtree via caret relations.

From here, we shall form the set of forest-tree pairs X similarly to Brown. We will consider a forest to be any finite ordered set of rooted trees built with carets from the caret set C_β. We will then consider the set of all pairs consisting of one forest and one tree where the forest and the tree have the same number of leaves, modulo the above equivalence relation once more.

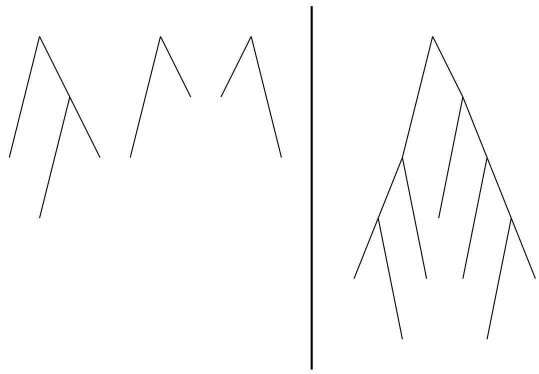

FIGURE 3.1: An example of a forest-tree pair in the complex for the group $\mathbf{F}_\tau$.

As with Brown's complex, we wish to add a partial order to this set of forest-tree pairs to create a poset. While based on the partial order used by Brown, we make adjustments to account for the different lengths of legs in our carets, which require us to treat subtrees as existing at different depths depending on the length of the leg to which they are attached.

Definition 3.1.6. A basic expansion of a tree contained in the forest of a forest-tree pair is performed by splitting the tree by deleting the top caret, turning each maximal proper subtree into a distinct tree (preserving their order).

We will now define two maps from the space of forest-tree pairs to $\mathbb{N}$ that will prove useful for working with the space.

Definition 3.1.7. The map $t : X \to \mathbb{N}$ maps an equivalence class of forest-tree pairs $[F, T]$ to the number of trees contained in the forest F.

The map $t : X \to \mathbb{N}$ maps an equivalence class of forest-tree pairs $[F, T]$ to the number of basic expansions required to obtain (F, T) from a tree-pair (considered as a forest-tree pair with only one tree in its forest)

We can see that t is a well-defined map as the addition or removal of redundant carets cannot add or remove trees to the forest F. Any added redundant caret must be added to an existing leaf of F, and removing redundant carets can only reduce a tree in F to the trivial tree, not remove it entirely. Additionally, the number of expansions required to obtain (F, T) from a tree-pair is directly correlated with the number of trees in F. Each basic expansion replaces 1 tree with k trees, where k is the number of legs on each caret in the caret set C_β, resulting in a net increase of $k - 1$ trees. In order to make e well defined, we will restrict X to only forest-tree pairs obtainable from a tree pair via a finite sequence of basic expansions. With this restriction applied, we can see that $t([F, T]) = (k + 1)e([F, T]) - 1$.

We can now define the relation $\leq$ on X in the following way:

Definition 3.1.8. For the equivalence classes of forest-tree pairs $[F, T], [F', T'] \in X$, we have that $[F, T] \leq [F', T']$ if there is a finite path of basic expansions a such that performing a on $[F, T]$ will create $[F'T']$.

Similarly to how we work with equivalence classes of tree pairs, we may add or remove redundant carets to forest-tree pairs and we may replace subtrees with equivalent subtrees (as in 2.3.2) and still remain within the same equivalence class. As such, the path a may include not just basic expansions, but also the addition and removal of redundant carets and the exchanging of equivalent subtrees. It is worth noting that the only type of action in a that can affect the number of trees in the forest F is a basic expansion.

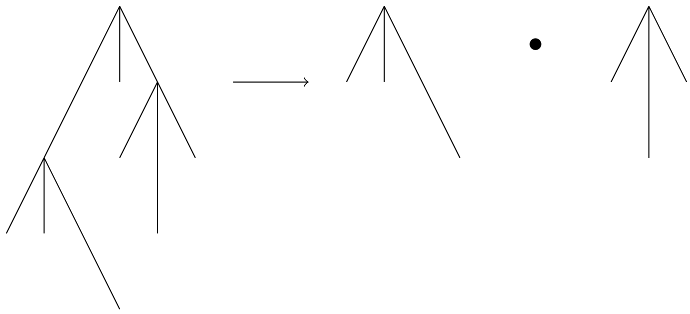

FIGURE 3.2: An example of a split performed on a tree with label 1. Note that the empty subtree attached to the middle leg of the top caret becomes a trivial tree after the split.

Lemma 3.1.9. *The relation $\leq$ is a partial order on the set X*

Proof. Three properties need to be fulfilled by a relation to make it a partial order: reflexivity, transitivity, and antisymmetry.

1. Reflexivity: The forest-tree pair $[F, T]$ can be reached from itself via the empty path of expansions 0, thus $[F, T] \leq [F, T]$ and so the relation is reflexive.

2. Transitivity: If we have $[F, T] \leq [F', T'] \leq [F'', T'']$, then there is a path of expansions a that constructs $[F', T']$ from $[F, T]$ and a path of expansions b that constructs $[F'', T'']$ from $[F', T']$. We can make the path ab that performs the expansions of a followed by the expansions of b. ab therefore constructs $[F'', T'']$ from $[F, T]$. Thus $[F, T] \leq [F'', T'']$ and so $\leq$ has the transitive property.

3. Antisymmetry: To prove antisymmetry we will consider a function from X to $\mathbb{N}$. $t([F, T])$ is simply the number of trees in F. Consider the basic expansions of $[F, T]$. If an expansion is performed on any tree, then it will increase $t([F, T])$ by an amount determined by the caret set C_β, but $t([F, T])$ must increase by at least 1.

Suppose $[F, T] \leq [F', T'] \leq [F, T]$. Therefore there is a sequence of basic expansions a that constructs $[F', T']$ from $[F, T]$ and a path b that constructs $[F, T]$ from $[F', T']$. As each basic expansion must increase t, we know $t([F, T]) \leq t([F', T']) \leq t([F, T])$. This immediately implies that $t([F, T]) = t([F', T'])$. $[F, T]$ must be constructible from $[F', T']$ by a sequence of basic expansions, but the only sequence of basic expansions that does not increase t is the trivial sequence with no expansions. As such, $[F, T] = [F', T']$. Hence, our relation is antisymmetric. ☐

As we have shown that $\leq$ is a partial order on the set X, we can consider $(X, \leq)$ a partially ordered set, or poset. As with Brown's complex in 2.2.1, we now wish to show this poset is directed.

Lemma 3.1.10. *The poset $(X, \leq)$ is directed*

Proof. To show this poset is directed, we must show that for any two elements $[F_1, T_1]$ and $[F_2, T_2]$, there exists an element $[F^*, T^*]$ such that $[F, T] \leq [F^*, T^*]$ and $[F, T] \leq [F^*, T^*]$. We begin by constructing T^* as the minimal tree that can be constructed from both T and T' by adding redundant carets and exchanging equivalent subtrees via caret relations. Assuming our tree-pair representation is well defined, such a tree always exists by ([Win], Lemma 2.7.9). Adding redundant carets will alter the forests in the forest tree pairs as well, leaving us with $[F_1', T^*] = [F_1, T_1]$ and $[F_2', T^*] = [F_2, T_2]$ (while the forest-tree pairs have changed, they are each in the same equivalence class they started in, as adding redundant carets and caret relations cannot move a forest-tree pair outside its equivalence class).

From here we will perform all possible basic expansions on F_1' and F_2' without adding redundant carets. The order we perform these expansions is arbitrary, but as an example we shall say that we will perform a basic expansion on the leftmost tree that is not a trivial tree. Repeating this process will construct from both forests F^*, the trivial n forest, which comprises of n trivial trees. We know that n is the same number in both cases, as both F_1' and F_2' are in forest tree pairs with T^*, which means they both have the same number of leaves as T^*, and the trivial n forest must have n leaves. As such we know we can construct $[F^*, T^*]$ by performing a series of basic expansions on $[F_1', T^*]$ and $[F_2', T^*]$. As such, we know $[F_1, T_1] = [F_1', T^*] \leq [F^*, T^*]$ and $[F_2, T_2] = [F_2', T^*] \leq [F^*, T^*]$ and so the poset must be directed. □

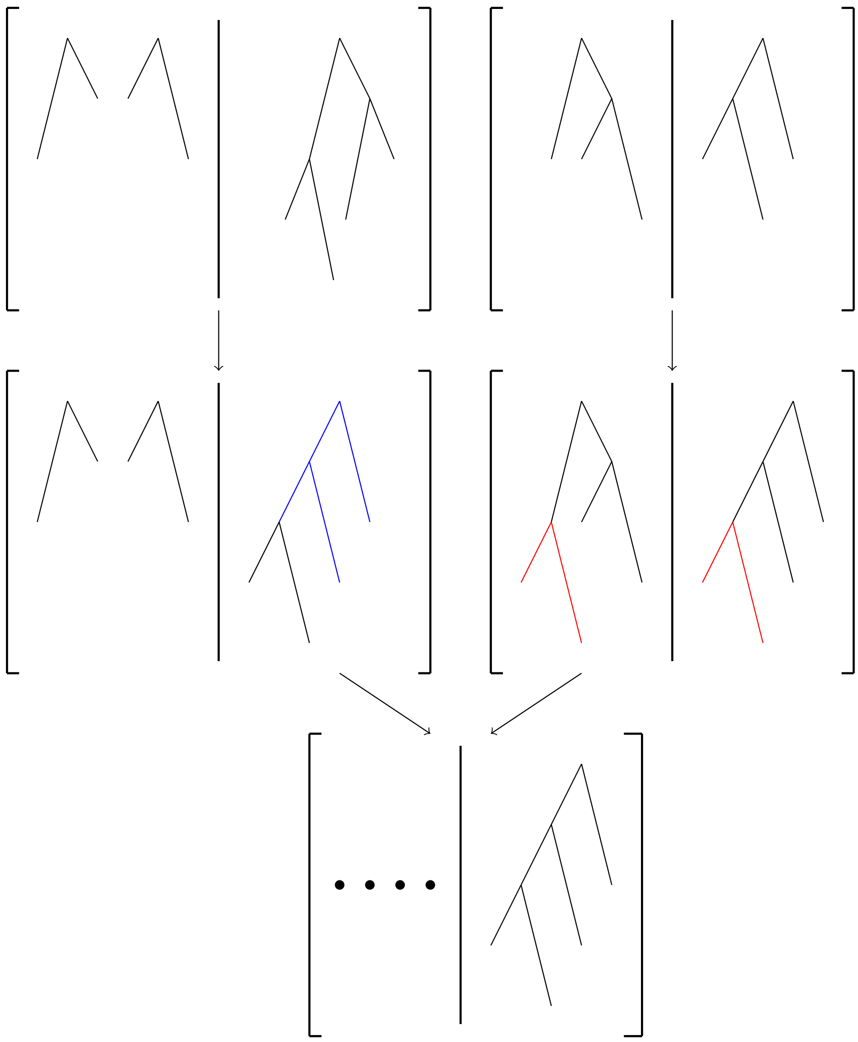

FIGURE 3.3: An example of finding the element of X greater than two given elements, including a subtree exchange (blue) and the addition of a redundant caret (red).

As mentioned beneath 3.1.7, we have restricted the set X to only include forest-tree pairs that can be constructed from tree-pairs (considered here as forest trees where the forest only contain one tree) via a finite sequence of basic expansions. This will restrict X in two main ways. First, it will limit X to only containing forest-tree pairs $[F, T]$ where F contains exactly $(k-1)n+1$ trees for $n \in \mathbb{N}$, where k is the number of legs on a caret in C_β. Secondly, it will guarantee that the only forest-tree pairs without a possible reverse expansion (hereafter called contractions) will be those with only one tree in their forest. This is what allows the function $e : X \to \mathbb{N}$, which maps a forest-tree pair to the number of expansions required to create it from a tree pair, to be a well defined function.

Furthermore, we shall limit the form of each trees by limiting the variety of carets to express those trees. Our biggest tool for reducing possible expansions is 2.3.6. As discussed in 2.3.2, the Bieri-Strebel group $\mathbf{F}_\beta$ with subdivision polynomial $ax^2 + bx - 1$, there are $\binom{a+b}{a}$ caret types in C_β, which clearly tends to ∞ caret types as $a + b$ tends to ∞. Using 2.3.6 reduces the number of caret types we need to consider to just 2 for all quadratic Bieri-Strebel groups, while not affecting the poset X in any way. This is because any two forest-tree pairs $[F, T]$ and $[F', T']$ where either T or the tree T_i in F (ie: the ith tree from the left) have the same leaf depths in the same order but are built with different carets are in the same equivalence class, and so fall into the same element in the poset. Restricting to building with only two caret types reduces the size of each equivalence class but, as shown by Winstone in 2.3.6, does not remove any equivalence class in its entirety.

The general shape of the two carets we reduce C_β to is fixed. The first caret type (called type a from now on) always has legs 1 to a be length 2 with all other legs length 1, and the second caret type (called type b from now on) has legs $a + 1$ to $2a$ as length 2 with all other legs length 1. This is clearly well defined for groups with well defined tree pairs, as by 2.3.5, the quadratic Biei-Strebel groups with well-defined tree-pairs have subdivision polynomial $ax2 + bx - 1$, $a \leq b$ and so the total legs on each caret is $a + b \leq 2a$.

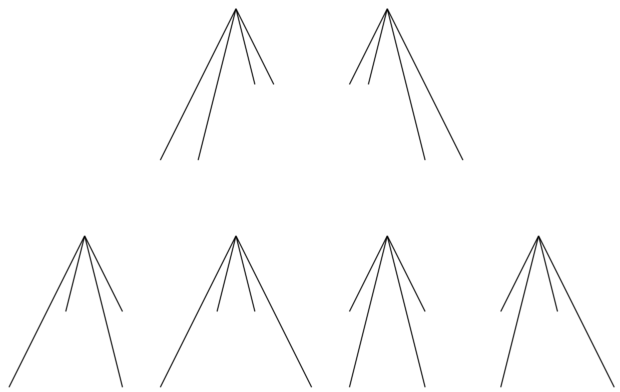

FIGURE 3.4: Though there are 6 possible carets for the subdivision polynomial $2x^2 + 2x - 1$, only the top two are necessary to represent any possible partition pair with a tree pair.

We now take the geometric realisation $\tilde{X}$ of the poset $(X, \leq)$. As mentioned in 2.2.1, the geometric realisation of a poset is a simplicial complex where each vertex is labelled with an element $x \in X$ and we insert an n-simplex between the vertices labelled by $x_0, ..., x_n$ where $x_0 \leq x_1 \leq ... \leq x_n$.

From 2.2.1, we know that the geometric realisation of a directed poset is a contractible simplicial complex, and so $\tilde{X}$ is contractible. We can define a right action $\tilde{X} \times \mathbf{F}_\beta \to \tilde{X}$ where $\mathbf{F}_\beta$ is the Bieri-Strebel group defined by tree pairs built with the caret set C_β. We define this action by taking a forest-tree pair $[F, T] \in X$ and a tree-pair $[T_1, T_2] \in F_\beta$ and performing tree-pair composition on these two objects (in the order $[F, T][T_1, T_2]$). Recalling the mechanics of tree-pair composition described in 2.1.4, can perform the first step by adding redundant carets and performing caret relation exchanges to transform (F, T) into (F', T') and (T_1, T_2) into (T', T_2'), then the result of the composition is the equivalence class $[F', T_2']$. To determine the kernel of the action of $\mathbf{F}_\beta$ on $\tilde{X}$, we have to identify the equivalence classes $[T_1, T_2]$ such that $[F', T_2'] = [F, T]$. We already know that (F, T) is in the same equivalence class as (F', T'), as (F', T') was obtained from (F, T) only via the addition of redundant carets and the exchanging of equivalent subtrees. Thus we can determine that $[T_1, T_2]$ is in the kernel of the action when $[F', T'] = [F', T_2']$, which immediately implies that

$T' = T'_2$, which implies that $[T', T'_2] = [T', T']$, which is the identity element of $\mathbf{F}_\beta$. Hence, if an element f of $\mathbf{F}_\beta$ maps a vertex x of $\tilde{X}$ to itself, then f is the identity element of $\mathbf{F}_\beta$. As such, the action of F_β on $\tilde{X}$ is free, in that the stabiliser of any point in $\tilde{X}$ is trivial.

We will now impose the filtration $\{X_{e(x)\leq n}\}_{n\in\mathbb{Z}_{\geq 0}}$, where each subcomplex is the complete subcomplex containing all $x \in X$ such that $e(x) \leq n$ for some $n \geq 0$.

Lemma 3.1.11. *Each subcomplex $\tilde{X}_{e(x)\leq n}$ is finite mod $\mathbf{F}_\beta$.*

Proof. We first consider the subcomplex $\tilde{X}_{e(x)\leq 0}$. The only forest-tree pairs where $e(x) = 0$ are those that have exactly 1 tree in their forest ie: tree-pairs. This means $\tilde{X}_{e(x)\leq 0}$ taken with the action of $\mathbf{F}_\beta$ is precisely the tree-pair representation for $\mathbf{F}_\beta$. As $\mathbf{F}_\beta$ is transitive when acting on itself, $\tilde{X}_{e(x)\leq 0}$ is trivial mod $\mathbf{F}_\beta$.

We now consider $\tilde{X}_{e(x)\leq 1}$. As $e : \tilde{X} \to \mathbb{Z}_{\geq 0}$, for each vertex $x \in tildeX_{e(x)\leq 1}$, we either have $e(x) = 0$ or $e(x) = 1$. We know that that all x such that $e(x) = 0$ are contained within a single orbit, so we turn our attention to when $e(x) = 1$. As mentioned beneath 3.1.7, $t(x) = (k+1)e(x) - 1$, and so all x such that $e(x) = 1$ must have $t(x) = k$, where k is the number of legs on a caret in C_β. Consider two distinct vertices $x_1, x_2 \in \tilde{X}$, with associated forest-tree pairs $[F_1, T_1]$ and $[F_2, T_2]$ respectively. Both F_1 and F_2 contain exactly k trees, and we may consider each tree seperately. As described in 2.1.4, given any two trees in a well defined tree-pair representation, we can find a tree that contains trees equivalent to both trees (via caret relations) as rooted subtrees. Considering the trees in the forests F_1 and F_2 are ordered, we can add redundant carets and perform caret relation exchanges on $[F_1, T_1]$ and $[F_2, T_2]$ until we obtain (F'_1, T'_1) and (F'_2, T'_2), where the first tree of F'_1 is the same as the first tree of F'_2. We can repeat this process on each other tree in F_1 and F_2 until we obtain (F'', T_1'') in the equivalence class of $[F_1, T_1]$, and (F'', T_2'') in the equivalence class of $[F_2, T_2]$. From here, we can observe that the tree pair equivalence class $[T_1'', T_2'']$ acting on the right of $[F'', T_1''] = [F_1, T_1]$ will produce $[F'', T_2''] = [F_2, T_2]$. $[T_1'', T_2'']$ is an equivalence class of ordered tree pairs and is therefore an element of $\mathbf{F}_\beta$. Thus for any two forest-tree pairs $[F_1, T_1]$ and $[F_2, T_2]$ such that F_1 has the same number of trees as F_2, there is an element of $\mathbf{F}_\beta$ that maps $[F_1, T_1]$ to $[F_2, T_2]$ under the action of $\mathbf{F}_\beta$ on $\tilde{X}$. To conclude this proof, we merely need to observe that, as $e : \tilde{X} \to \mathbb{Z}_{\geq 0}$, $e(X)$ takes a finite number of values in any $\tilde{X}_{e(x)\leq n}$. We have just shown that each $\tilde{X}_{e(x)=i}$ is a single orbit for the action of $\mathbf{F}_\beta$ on $\tilde{X}$ for any $i \in \mathbb{Z}_{\geq 0}$, so there must be a finite number of of orbits of $\mathbf{F}_\beta$ on $\tilde{X}_{e(x)\leq n}$. □

Before we can begin to apply Brown's criterion to the complex, we must demonstrate the connectivity of the space.

Lemma 3.1.12. *Suppose $y_1, ..., y_n$ are distinct simple contractions of the point $x \in X$. Then $y_1, ..., y_n$ have a lower bound in the poset, that is to say an element y such that $y \leq y_i \forall 1 \leq i \leq n$ if and only if each contraction y_i is disjoint from the other y_i. If the set $y_1, ..., y_k$ does have a lower bound, it has a greatest lower bound z. That is to say that $\forall y$ such that $y \leq y_i$ for $1 \leq i \leq n$, we have $y \leq z$.*

Proof. $\Leftarrow$: Suppose $y_1, ..., y_k$ are distinct, disjoint simple contractions of $x \in X$. Then performing each contraction on x in any order will result in z, which is a lower bound for the set $y_1, ..., y_k$ as $z \leq y_i$ for all $i \in \{1, ..., k\}$ by any path of retractions from x to z such that y_i is the first contraction. $\Rightarrow$: suppose the set of contractions $y_1, ..., y_k$

has a lower bound w. From the definition 3.1.6, we consider these contractions in the following way. We may take a set of l consecutive trees, where l is the number of legs on a caret in C_β. We can then attach a top caret to merge these trees into a single tree with label 1.

Suppose we have two distinct intersecting contractions on the forest F. This means there is at least one tree in F that falls in both contractions. Contractions are performed over a set of trees in F, but if we perform one of these contractions, then any tree that falls in the intersection will no longer be a tree (it will instead be a maximal subtree of the tree produced by the contractions). As such, we cannot perform 2 intersecting type 2 contractions sequentially and thus they do not have a lower bound.

Hence, we cannot form a lower bound for the set of contractions $y_1, ..., y_k$ if there is any intersection between any two contractions, and thus if there is a lower bound then the set must be pairwise disjoint.

We again assume that the set of contractions $y_1, ..., y_k$ has a lower bound w. We now know that the set must be pairwise disjoint, and thus we can form the lower bound z by performing each contraction in any order. Suppose $z < w$. That would imply that there is a path of contraction from w to z. But we know $w \leq x$, and so w must sit on a path of expansions from z to x. However, we reach z from x by performing each of the contractions y_i and no others. So for $z < w < x$, there must be a contraction y_i that is not performed from x to reach w, so w cannot be a lower bound for the contractions $y_1, ..., y_k$. Hence we can conclude that z is a highest lower bound for the contractions $y_1, ..., y_k$. □

The remainder of the proof of 3.1.1 is very similar to the proof found in ([Bro87], Section 4). 3.1.12 is analogous to Brown's Lemma 4.18, and the rest of the proof proceeds as in Brown. It is presented here for completeness.

Lemma 3.1.13. *For any $x \in \tilde{X}$, the complex $\tilde{X}_{<x}$ is homotopy equivalent to the simplicial complex $\Sigma(x)$ defined as follows*

- *The vertices of $\Sigma(x)$ are the distinct simple contractions of the forest of x.*
- *A set of vertices spans a simplex if and only if the contractions are pairwise disjoint.*

Proof. Let K be the space $\tilde{X}_{<x}$. For any $x' < x$ take $K_{x'}$ as the subcomplex $\tilde{X}_{\leq x'}$. This space contains x' and precisely all points beneath it. x' forms a simplex with every chain of points $x_0 < ... < x_k < x'$, and as such the space $\tilde{X}_{\leq x'}$ is a cone, and therefore contractible. We can now write $K = \bigcup K_{x'}$, where each x' is a simple contraction of x.

Consider a collection $\{x_i\}$ of simple contractions of x. We can see that if $\bigcap K_{x_i}$ is nonempty, it is because there is a point beneath each x_i. In other words, the contractions x_i have a lower bound. From this, we can see that $\bigcap K_{x_i}$ will consist of the lower bound z and all points beneath it, and is therefore just K_z. Thus the intersection of any K_{x_i} is contractible. From 3.1.12, we can therefore see K as the space of simple contractions of x, with an n-simplex spanning a set of n disjoint simple contractions. □

Lemma 3.1.14. *For any integer l, there is an integer $\mu(l)$ such that if the element x has a forest tree pair $[F, T]$ such that F contains $\mu(l)$ trees, then $\Sigma(x)$ is l connected.*

Proof. We consider $\Sigma(x)$ in the following way. Each set of l consecutive trees in the forest F can have two different merges applied to them (using one of the two different caret types). As such, there are $2(f-k+1)$ vertices in $\Sigma(x)$, which we can group into pairs of merges that share a merge set (that is to say, merge the same set of consecutive trees). Each pair is disconnected, but both vertices in each pair are adjacent to the same vertices in $\Sigma(x)$. From this, we can immediately see that $\Sigma(x)$ is connected as long as F contains at least $3k-1$ trees, as this means that any set of the trees of F overlaps with at most one of the two end sets (ie: the set that contains the first k trees, and the set that contains the last k trees). As such, any vertex in $\Sigma(x)$ will be adjacent to one of these two pairs of contractions, and the end set contractions do not overlap and are therefore adjacent.

We can then perform an induction in the following way. Assume that the forest containing n trees is m connected. We consider the forest containing $n+2k-1$ trees. In this forest, for each set of n consecutive trees, there exists a disjoint set of k consecutive trees (if the set of n consecutive trees has fewer than k trees to the left of it, then it has at most $k-1$ trees to the left of it, which means it has at least k trees to the right of it, forming a disjoint set of k consecutive trees). As such, for each m connected subset of $\Sigma(x)$, there is a pair of points in $\Sigma(x)$ that suspends that subset. Hence, any $m+1$ spheres are suspended and the forest of $n+2k-1$ trees is $m+1$ connected. By induction, we can say that the forest with $3k-1+(m-1)(2k-1)$ trees is m connected. □

Proof of 3.1.1. To construct $\tilde{X}_{\leq h+1}$ from $\tilde{X}_{\leq h}$, we add all vertices x such that $h(x) = h+1$, and each one cones off the subspace $\tilde{x}_{<x}$. Given the increasing connectivity of the spaces $\tilde{x}_{<x}$ from 3.1.13 and 3.1.14, the connectivity of the quotient space $\tilde{X}_{\leq h+1}/\tilde{X}_{\leq h}$ goes to ∞ as h goes to ∞. Hence, by 3.1.5, the group $\mathbf{F}_\beta$ is F_∞. □

3.2 Cubulation of Bieri-Strebel Groups

As the complex we have constructed is similar to Brown's complex, we can perform similar adaptations to it as described in 2.2.3. However, our scope is much more limited. We intend to prove the following:

Theorem 3.2.1. *The quadratic Bieri-Strebel group F_β with subdivision polynomial $x^2 + nx - 1$ is able to act without a fixed point on a $CAT(0)$ cube complex.*

Similar to Farley, our intention is to restrict Brown's complex by only restricting the equivalence relation $\leq$ to the relation $\preceq$, where $[F,T] \preceq [F',T']$ if $[F',T']$ can be constructed from $[F,T]$ with a single simple expansion. As in 2.2.3, this relation is not an equivalence relation, as it is not transitive.

A complicating factor in applying this method to Bieri-Strebel groups is the presence of caret relations. In $\mathbf{F}$, there is only one caret type, in particular this means that each tree will always produce the same ordered set of trees when we apply a basic expansion. Multiple different caret types mean that this is not true in the complex we have constructed for Bieri-Strebel groups. As seen in the proof of 3.1.11, it is possible to use redundant carets and caret relations to exchange any given caret in a tree for any other caret in the caret set C_β. As both caret relations and the addition of redundant carets do not change the equivalence class of a forest-tree pair, it is possible to exchange the top caret of any tree for any other caret type without changing the equivalence class of the forest-tree pair. Because of this, any tree in our forest can expand in $|C_\beta|$ different ways. This presents an immediate obstruction to applying Farley's method of constructing a $CAT(0)$ cube complex, as Farley's method relies on each tree having a single possible expansion. Our primary goal in this construction is therefore to produce a canonical expansion for each tree, without disrupting the connectivity of the space.

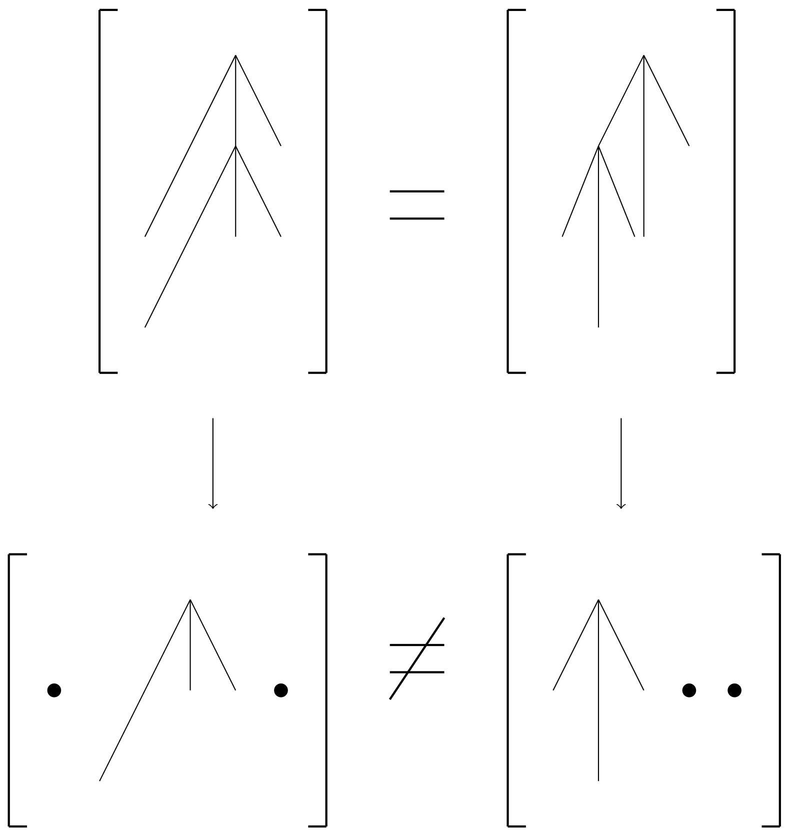

FIGURE 3.5: Two trees that are interchangeable via caret relations, but expand into different forests.

To assist in the construction of a cube complex, we wish to introduce a canonical form to any tree that may appear in our forest, such that the top caret (and indeed all other carets) are fixed for a given interval partition. To start, we may cite the following result from Jason Brown

Citation 3.2.2. *([Bro18], section 6.1) For a quadratic Bieri-Strebel group with subdivision polynomial of the form $x^2 + ax - 1$, we may write any element using Winstone's generators (as in 2.3.7) such that there are no y_i^{-1} terms for any $i \in \mathbb{Z}_{\geq 0}$.*

When translated to the language of tree-pair diagrams, this means there are no y-carets on the right hand tree in the tree pair. We then apply the bijection $i : \mathbf{F}_\beta \to \mathbf{F}_\beta$, $g \mapsto g^{-1}$. Considering the composition of tree pairs described in 2.1.4, it is fairly straightforward to see that $[T_1, T_2][T_2, T_1] = [T_1, T_1] = 1$ and so $[T_1, T_2]^{-1} = [T_2, T_1]$. Jason Brown's normal form grants us a tree pair diagram for every element of $\mathbf{F}_\beta$ such that the right hand tree countains only x-carets, and applying the bijection i gives us a tree pair diagram for every element of $\mathbf{F}_\beta$ such that the left hand tree contains only x-carets.

Beneath 3.1.7, we restricted the space X (and by extension $\tilde{X}$) such that a tree pair $[F, T]$ is in X only if it can be obtained from a tree pair in the tree pair representation of $\mathbf{F}_\beta$ by a finite sequence of basic expansions. We now further restrict the space to $\tilde{X}'$ by only including forest-tree pairs obtained by expanding tree pairs in this seminormal form with only x-carets in the left hand tree. We can use caret relations and redundant carets to convert any forest-tree pair (F, T) into a forest-tree pair in seminormal form with a generalisation of ([BNR21], proposition 2.1). Notably this does not leave the equivalence class $[F, T]$. Thus no equivalence class of X is eliminated in its entirety by this restriction.

The action of $\mathbf{F}_\beta$ on X and $\tilde{X}$ is also worth considering under this restriction. In particular, if we compose the forest-tree pair $[F, T]$ with the tree pair $[T_1, T_2]$, where T_1 contains y-carets, then we may be forced to add redundant y-carets to $[F, T]$ in order to complete the composition. This can be managed by applying ([BNR21], proposition 2.1) again once the composition has been concluded. Even if the composition has y-carets in its forest, there exists a forest-tree pair in the same equivalence class that does not.

With all this in mind, we can consider what effect this restriction has on expansions and contractions. In $\tilde{X}$, we had two forest-tree pairs in the same equivalence class that, upon expanding the same tree, produced two different forests (see 3.5 for an example). This is due to the fact that the caret relations in $\mathbf{F}_\beta$ change not just a single caret, but multiple maximal subtrees of the caret. As such, if we use a caret relation to exchange the top caret of a tree in F, then expand it, this will result in different trees than if we had expanded it without the caret relation. However, in $\tilde{X}'$, we only have one type of caret, so we may not apply caret relations. As such, the top caret of each tree is fixed, and therefore each tree will always produce the same trees when expanded. Similarly, a set of trees can only be contracted by joining each tree to an x-caret, so each set of trees can only be contracted on one way as well.

With the space $\tilde{X}'$ defined, if we wish to prove 3.2.1, we must show that $\mathbf{F}_\beta$ acts without a fixed point on $\tilde{X}'$ and that $\tilde{X}'$ is a $CAT(0)$ cube complex.

Lemma 3.2.3. *$\tilde{X}'$ is a cube complex.*

Proof. We will begin by showing that the cell above each $x \in \tilde{X}'$ is a cube. Consider the point x as a forest-tree pair $[F, T]$. If forest F in the reduced form of the pair (F, T) has no trivial trees, then we can simply take the cell $[x, \bar{x}]$ which has as 0-cells all elements y such that $x \preceq y$, where $\preceq$ is Stein's relation from 2.2.2. As this is an operation we can perform once on each tree in the forest, every combination of expanded and unexpanded trees forms a boolean lattice (as in 2.2.3). Now suppose our reduced forest-tree pair contains a trivial tree in the forest. Our forest may only contain x-carets. As such, when we add a redundant caret in order to expand the trivial tree, it must be an x-caret. Hence, there is only one possible expansion for trivial trees as well.

We now consider the intersections of such cells. Suppose the intersection of two cells in $\tilde{X}'$ contains at least two distinct 0-cells (if not, then the intersection is either trivial, or contains exactly 1 0-cell and is therefore a 0-cube). Using the height function $e : \tilde{X}'^0 \to \mathbb{Z}$ that maps a forest-tree pair to the number of expansions required to construct the forest-tree pair from a tree pair, we take any 0-cell x with $e(x)$ as low as possible, and any cell y such that $e(y)$ is as high as possible.

First we consider that the 0-cells (considered as forest-tree pairs) have some sort of poset relation. Say WLOG that $x \le y$. Then there is a path of simple expansions (or potentially a double expansion for a trivial tree) that constructs y from x. This path must be contained in both of the intersecting cubes. By construction of the cubes, all other orderings of the simple expansions that construct y from x (and the double expansion that forms a 2-cell above a forest with a trivial tree) are paths through both cubes, and so must appear in the intersection as well. As such, the intersection of two cubes forms a cube.

We now consider that x and y have no poset relation. Consider each point in a cube as a set of n boolean values, where n is the number of trees in the forest of its least vertex (or 1 greater for each trivial tree). In this case, we can interpret the poset relation as $x \le y$ if for each boolean value x_i in x that is equal to 1, y_i must also be equal to 1. If a cube contains both x and y, it must also contain $z = x$ **OR** y, where each boolean value is 1 if it is 1 in either x or y. Clearly $y \ge z$, and as z is in both cubes, it must be in the intersection, which means that $e(y)$ was not maximal for the intersection. Due to this contradiction, we can discard this as a possibility. □

Lemma 3.2.4. *$\mathbf{F}_\beta$ acts on $\tilde{X}$ without a fixed point.*

Proof. We use the argument from 2.2.1, as it adapts to this complex without significant changes. We have an element of $\mathbf{F}_\beta$, represented with the equivalence class of tree-pairs $[T_1, T_2]$, act on the right of a forest-tree pair $[F, T]$ via tree-pair composition, producing the forest-tree pair $[F', T']$. If $[F', T'] = [F, T]$, then $T_1 = T_2$ and $[T_1, T_2] = 1_{\mathbf{F}_\beta}$, hence only the identity fixes any point of $\tilde{X}'$. □

Lemma 3.2.5. *$\tilde{X}'$ is contractible.*

Proof. Consider the poset X. Even with the identification of subforests produced by expanding equivalent trees, the poset is still directed. For any two elements $x, y \in X$, we can find an element $z \ge x, y$ by using the process in 3.1.10.

We now need to show that we may transform $\tilde{X}$ into $\tilde{X}'$ without disrupting the connectivity of the space. We proceed in the manner of Stein and Farley 2.2.3.

First, we consider Stein's relation for forest-tree pairs, where $[F, T] \preceq [F', T']$ if $[F'T']$ can be constructed from $[F, T]$ by applying 1 simple expansion to any number of trees in the forest F. We now consider the space

$$(x, z) = \{y \in X | x < y < z\}$$

as with Stein's complex, we can find y' as the greatest element such that $x \preceq y' < y < z$, which means $y' \in (x, z)$. We then once again use the map $\sigma : (x, z) \to (x, z); y \mapsto y'$ and ([Qui73], section 1.5) to show that $\widetilde{(x, z)}$ is contractible. Finally, as with Stein's complex, we will replace simplices in X with subspaces $\widetilde{[x, z]}$ where $[x, z] = \{y \in X | x \le y \le z\}$. We express these as $\widetilde{[x, z)} \bigcap \widetilde{(x, z]}$ which is the suspension of $\widetilde{(x, z)}$ between the points x and z, and as we know $\widetilde{(x, z)}$ is contractible, its suspension must be as well.

To transform this complex into a cube complex, we simply replace each subcomplex $[x, \bar{x}]$ with an n-cube, where n is the number of trees in the forest of x. We know $[x, \bar{x}]$

is contractible, and an n-cube is certainly contractible as well. As such, we can see $\tilde{X}'$ is contractible. □

Lemma 3.2.6. *$\tilde{X}'$ is a $CAT(0)$ complex.*

Proof. considering 2.2.10, to show $\tilde{X}'$ is a $CAT(0)$ complex, we need to show the following

1. $\tilde{X}'$ is a cube complex.
2. $\tilde{X}'$ is simply connected.
3. For each vertex x in $\tilde{X}'$, the link $lk(x)$ is a flag simplicial complex.

For (1), we can refer to 3.2.3, and for (2) we can refer to 3.2.5, as $\tilde{X}'$ being contractible implies $\tilde{X}'$ being simply connected. We can straightforwardly see that $\tilde{X}'$ is locally finite through a similar argument as in 2.2.7. That is to say, as the number of vertices adjacent to a vertex x associated with forest-tree pair $[F, T]$ is limited by the number of trees in F, and the number of trees in F is always finite, we know each vertex is only adjacent to a finite number of other vertices, and as such $\tilde{X}'$ is locally finite. We can now cite 2.2.4 to demonstrate that $\tilde{X}'$ is a complete geodesic metric space, as such, our notion of links (as defined in 2.2.8 is well defined on $\tilde{X}'$. As such, the only thing that remains to be shown is that each vertex x has $lk(x)$ as a flag simplicial complex. We therefore wish to show that if there exists the 1-skeleton for a k-simplex in $lk(x)$, then that k-simplex is also in $lk(x)$.

We consider $lk(x)$ in the following way. Each vertex in $lk(x)$ is an edge in $\tilde{X}'$ and therefore implies an adjacent vertex. An edge between any two vertices v, v' in $lk(x)$ implies the existence of a 2-cell containing x and the adjacent vertices in $\tilde{X}'$.

We first consider 1-skeleta contained entirely in the ascending link $lk^{\uparrow}(x)$, based on the height function $h(x)$ which counts the number of expansions to construct x from an element of $\mathbf{F}_\beta$ (considered as an element of X'). Any vertex in $lk^{\uparrow}(x)$ implies a basic expansion of the forest of x. As discussed in the proof of 3.2.3, any two basic expansions in the forest of x have a common expansion and can therefore be part of a 2-cube with x as its lowest point. As such, the ascending link is an $n+i$-simplex, where n is the number of trees in the forest of x, and i is the number of trivial trees (which may be expanded two different ways, but those two ways have an upper limit). Any 1-skeleta of a k-simplex in $lk^{\uparrow}(x)$ would be the skeleta of a subsimplex of the $n+i$-simplex, and therefore the k-simplex exists.

We now consider 1-skeleta contained within $lk^{\downarrow}(x)$. Each vertex in $lk^{\downarrow}(x)$ implies a simple contraction, and two vertices v, v' have an edge connecting them if they have a lower bound, which they only have if they are disjoint contractions. As such, all the vertices in the 1-skeleta of a k-simplex in $lk^{\downarrow}(x)$ are pairwise disjoint and as such have a common lower bound. This common lower bound implies the existence of a k-cube in $\tilde{X}'$, which then implies the k-simplex in $lk^{\downarrow}(x)$.

We now consider the 1-skeleton K of a k-simplex in $lk(x)$ that is not wholly contained in $lk^{\uparrow}(x)$ or $lk^{\downarrow}(x)$. As all edges in $\tilde{X}'$ imply a basic expansion or contraction, any vertex x' adjacent to x must have that $h(x') \neq h(x)$. As such, all vertices in $lk(x)$ are either in $lk^{\uparrow}(x)$ or $lk^{\downarrow}(x)$. Therefore, the 0-skeleton of K can be considered as $(K^{(0)} \cap lk^{\uparrow}(x)) \cup (K^{(0)} \cap lk^{\downarrow}(x))$. If there is a 1-cell connecting $(K \cap lk^{\uparrow}(x))$ and $(K \cap lk^{\downarrow}(x))$, then this implies the existence of a 2-cell with an expansion $\uparrow$ of x to $x^{\uparrow}$ and a contraction $\downarrow$ of x to $x^{\downarrow}$ on its boundary. This further implies the parallel expansion where we perform $\uparrow$ on $x^{\downarrow}$ and the parallel contraction where we perform $\downarrow$ on $x^{\uparrow}$. This is only possible if $\uparrow$ and $\downarrow$ are disjoint.

As such, for the 1-skeleton of a simplex K in $lk(x)$, we have that each vertex k of K represents a distinct expansion or contraction of K and as these vertices are pairwise adjacent, they must also be pairwise disjoint. As such, x sits on the cube C such that the bottom vertex c of C is found from x by performing all contractions in $\{k\}$ in any order. All other vertices of C can be obtained from c by by first replacing the contractions of $K^{(0)}$ with their inverse expansions, then taking a subset of $K^{(0)}$ and performing those expansions in any order. The existence of C implies the existence of K in $lk(x)$. □

Proof of 3.2.1. Combining 3.2.3, 3.2.4 and 3.2.6 is our proof for 3.2.1. □

Chapter 4

Bestvina-Brady Morse Theory

Something that has proven useful in the calculation of BNSR invariants of many groups, in particular for Thompson-like groups, has been the application of Bestvina Brady Morse Theory (see, for example, [WZ15] and [Zar17]). Bestvina-Brady morse theory is a technique for determining the connectivity of a quotient space Y_1/Y_2 for certain filtrations of an affine cell complex Y. With a well-constructed cell-complex, we may apply a suitable height function to certain layers of a filtration in order to calculate the BNSR invariant using the definition 1.3.5.

4.1 Preliminaries

Definition 4.1.1. ([WZ15], 1.1) An affine cell complex is a complex constructed from euclidean polytopes. More formally, a space X is an affine cell complex if it is the quotient space of the disjoint union of a set of euclidean polytopes C modulo an equivalence relation with two criteria.

- Every polytope in C is mapped into X injectively.
- If two polytopes have an interior point identified then the entire interior of each polytope is identified isometrically.

Definition 4.1.2. [WZ15], 1.5) A Morse function (h,s) is a map $X \to \mathbb{R} \times \mathbb{R}$ such that both h and s are affine height functions on the affine cell complex X, the function s takes only finitely many values on vertices of X, and there is an $\epsilon \in \mathbb{R}_{>0}$ such that for any pair of adjacent vertices $x, x' \in X$, we either have that $|h(x) - h(x')| \geq \epsilon$ or that $h(x) = h(x')$ and $s(x) \neq s(x')$.

A Morse function with two functions like this is an adaptation of the definition of Morse functions provided in [BB97], which has a single height function that is discrete on vertices of the space. When working with such functions, we generally treat h as the "primary" height function of the Morse function, with s allowing us to distinguish between two adjacent vertices that are mapped to the same value under h. This allows us to consider a greater range of functions for our height function h. We call the height defined by a Morse function the "refined" height. If possible, it is best to choose a "secondary" height function s such that any two adjacent vertices in X are mapped to different values by s, as this makes the second criteria for a Morse function much easier to fulfil.

It is worth noting that it is possible to encode refined height into a single number and as such express our refined height function as a single function $h' : X \to \mathbb{R}$. For example, we could have $h'(x) = h(x) + \alpha s(x)$, where $\alpha \leq \epsilon \sup\{s(x) -$

$s(x')|x, x'\mathtt{adjacent}\}$. This forces the greatest possible difference in s between adjacent vertices to effect the value of h' by a smaller amount than the smallest possible difference non-zero difference in h between adjacent vertices, Thus achieving the same effect as the refined height as defined in 4.1.2. We have chosen to define our morse functions in this way for a pragmatic purpose. The results in chapter 5 we will be proving using Morse theory build on the work of Witzel and Zaremsky in [WZ15] and [Zar17], and will directly cite useful results from both these papers. We have elected to use the same definition of Morse function as Witzel and Zaremsky in order to aid clarity when we incorporate these results into our own proofs.

A concept important to Bestvina-Brady Morse Theory is links. We have previously defined the link of a vertex in 2.2.8, and so will not restate it here. However, we are able to combine links with height functions to develop new concepts.

Definition 4.1.3. For a vertex x in an affine cell-complex X, and a Morse function (h,s), the ascending link $lk^{(h,s)\uparrow}(x)$ is the intersection of an n-sphere (where n is the greatest dimension of any cell x is incident to) of radius $0 < \epsilon \ll 1$ and the subcomplex $X_{>(h,s)(x)}$ of cells γ that contain x such that $h(x) = \inf\{h(y)|y \in \gamma\}$ and $s(x) = \inf\{s(y)|y \in \gamma, h(y) = h(x)\}$.

Similarly, for the same complex and Morse function, we define the descending link $lk^{(h,s)\downarrow}(x)$ as the intersection of an n-sphere of radius $0 < \epsilon \ll 1$ and the subcomplex $X_{<(h,s)(x)}$ of cells γ that contain x such that $h(x) = \sup\{h(y)|y \in \gamma\}$ and $s(x) = \sup\{s(y)|y \in \gamma, h(y) = h(x)\}$

By the definition 4.1.2 we have provided of Morse function, every vertex adjacent to a given vertex x in an affine cell complex must be deemed higher or lower than it by a Morse function. This means that every edge incident to x will appear as a vertex in either the ascending or descending link. We can express this as $lk(x)^{(0)} = lk^{(h,s)\uparrow}(x)^{(0)} \bigcup lk^{(h,s)\downarrow}(x)^{(0)}$. However, that doesn't mean we can say

$$lk(x) = lk^{(h,s)\uparrow}(x) \bigcup lk^{(h,s)\downarrow}(x)$$

as any k-cell with $k > 1$ that includes at least 0-cell adjacent to x that is higher than x and at least one 0-cell adjacent to x that is lower than x, then neither the ascending link nor the descending link will contain it.

Definition 4.1.4. For a vertex x in an affine cell-complex X, and a Morse function (h,s), the ascending star $st^{(h,s)\uparrow}(x)$ is the subcomplex of X consisting of all cells γ such that x is the vertex of minimal refined height in γ. Similarly, for the same complex and Morse function, the ascending star $st^{(h,s)\uparrow}(x)$ is the subcomplex of X consisting of all cells γ such that x is the vertex of maximal refined height in γ.

4.2 The Morse Lemma

Lemma 4.2.1. *([WZ15], lemma 1.7) Let X be an affine cell complex, (h,s) be a Morse function and $X_{a\leq h\leq b}$ denote the full subcomplex of X supported on vertices x such that $s \leq h(x) \leq b$ for some numbers a, b in the codomain of h. Take p, q, r in $\mathbb{R} \bigcup \{\pm\infty\}$ such that $p \leq q \leq r$. If for every vertex $x \in X_{q<h\leq r}$ the descending link $lk^{(h,s)\downarrow}_{p\leq h}(x)$ is at least $k-1$ connected then the quotient space $X_{p\leq h\leq r}/X_{p\leq h\leq q}$ is k connected. If for every vertex $x \in X_{p\leq h<q}$ the ascending link $lk^{(h,s)\uparrow}_{h\leq r}(x)$ is at least $k-1$ connected then the quotient space $X_{p\leq h\leq r}/X_{q\leq h\leq r}$ is k connected.*

Proof. We only need to prove the first statement, as the second statement is the same as the first if we exchange the Morse function (h,s) for $(-h,-s)$. First, we will show that we can assume $r < \infty$. Suppose $r = \infty$. In this case, the subcomplex $X_{p \leq h \leq r}$ is just the "top half" complex $X_{p \leq h}$. The compactness of spheres is a standard result in topology which states that any n-sphere is compact for any dimension n. In particular, this means that all spheres are closed and bounded. As all spheres are bounded, this means there exists some finite r' such that $X_{p \leq h \leq r'}$ contains all the homotopy spheres of $X_{p \leq h}$. As such, $X_{p \leq h \leq r'}$ and $X_{p \leq h}$ are homotopy equivalent. Homotopy equivalent spaces are fully equivalent for the purposes of calculating the quotient space, and so whenever we see $X_{p \leq h}$ we may instead substitute $X_{p \leq h \leq r'}$, thus assuming r finite.

We also wish to assume that $r - q \leq \epsilon$, where epsilon is as used in the definition of Morse function 4.1.2. For this, we will use an inductive argument to show that any case where $r - q > \epsilon$ is equivalent to the base case where $r - q \leq \epsilon$. We begin with the base case. As $r - q \leq \epsilon$, and the distance between any two adjacent vertices is at least ϵ, we know that only one "layer" of vertices can exist in the space $X_{q < h \leq r}$. In particular no vertices in $X_{q < h \leq r}$ are adjacent, and they are not adjacent to any higher vertices in $X_{p < h \leq r}$. We consider the inclusion of $X_{p < h \leq q}$ into $X_{p < h \leq r}$. Any vertex x not in $X_{p < h \leq q}$ must be in $X_{q < h \leq r}$. By assumption, x has a $k - 1$ connected descending link, which means the space for any l-cell for $l \leq k$ with x as its uppermost cell must be filled by an l cell. As such, the inclusion of $X_{p < h \leq q}$ into $X_{p < h \leq r}$ induces isomorphisms in the first k homotopy groups of $X_{p < h \leq q}$ onto the first k homotopy groups of $X_{p < h \leq r}$. As the first k homotopy groups are all we care about with regard to demonstrating the connectivity of the quotient space $X_{p \leq h \leq r} / X_{p \leq h \leq q}$, we may continue to the general case. Considering a case where $r - q \leq n\epsilon$, we use the inductive assumption to induce isomorphisms in the first k homotopy groups of $X_{p < h \leq q}$ and $X_{p < h \leq q + n - 1\epsilon}$ via the inclusion. We thus only have to consider the inclusion of $X_{p < h \leq q + n - 1\epsilon}$ into $X_{p < h \leq q + n\epsilon}$. However, this is equivalent to the $r - q \leq \epsilon$ inclusion, which we can see by relabelling $q = q + n - 1\epsilon$ and $r = q + n\epsilon$. As such we know we may induce an isomorphism from the first k homotopy groups of $X_{p < h \leq q + n - 1\epsilon}$ to the first k homotopy groups of $X_{p < h \leq q + n\epsilon}$. As such, with the inductive step, we can induce isomorphisms from $X_{p < h \leq q}$ to $X_{p < h \leq q + n\epsilon}$. Any finite r will have $r - q \leq n\epsilon$ for some n, and as previously discussed, if r is infinite, we can substitute it with some finite r'. As such, any case can be reduced to the case where $r - q \leq \epsilon$.

Moving forward with the assumption that $r - q \leq \epsilon$, our goal is to construct a well order $\preceq$ on the vertices of $X_{q < h \leq r}$ such that the quotient space

$$S_{\preceq v} / S_{\prec v} := X_{p < h \leq q} \cup \bigcup_{w \preceq v} st_{p \leq h}^{(h,s)\downarrow}(w) / X_{p < h \leq q} \cup \bigcup_{w \prec v} st_{p \leq h}^{(h,s)\downarrow}(w) \tag{4.1}$$

is at least k-connected. Effectively, we are constructing an order to add vertices from $X_{q < h \leq r}$ into $X_{p < h \leq q}$ in order to construct $X_{p < h \leq r}$ in such a way that, each time we add a vertex we confirm that the connectivity of the quotient is maintained. Once we have concluded adding vertices, we will know that the quotient $X_{p < h \leq r} / X_{p < h \leq q}$ is at least k-connected. We choose any order $\preceq$ on the vertices v of $X_{q < h \leq r}$ such that $s(v) < s(v') \Rightarrow v \prec v'$. Consider the construction of $S_{\preceq v}$ from $S_{\prec v}$. We introduce the vertex v and with it each cell in $st_{p \leq h}^{(h,s)\downarrow}(v)$. This cones off the intersection of $S_{\prec v}$ and the boundary of the the star of v, $\delta st(v)$, with v being the cone point. We claim that $S_{\prec v} \cap \delta st(v)$ is precisely the boundary of $st_{p \leq h}^{(h,s)\downarrow}(v)$ in $X_{p \leq h}^{(h,s) \leq (h,s)(v)}$ (which we will write as B), and that B is homeomorphic to $lk^{(h,s)\downarrow}_{p\leq h}(w)$. If this is true, then we already know $lk^{(h,s)\downarrow}_{p\leq h}(w)$ is $k-1$ connected by assumption, so B and therefore $S_{\prec v} \cap \delta st(v)$ is $k-1$ connected as well. As such, the construction of $S_{\preceq v}$ from $S_{\prec v}$ consists of coning off a $k-1$ connected space and thus the quotient space 4.1 is at least $k-1$ connected.

We can see that $S_{\prec v} \cap \delta st(v) \subseteq B$ as intersecting $\delta st(v)$ with $S_{\prec v}$ removes any vertex w such that $s(w) > s(v)$. As $S_{\prec v} \cap \delta st(v)$ is a full subcomplex of $\delta st(v)$, in order to show $B \subseteq S_{\prec v} \cap \delta st(v)$, we need to show that any vertex w adjacent to v with $(h,s)(w) < (h,s)(v)$ is in $S_{\prec v}$. If $h(w) < h(v)$ then $h(w) \leq h(v) - \epsilon \leq r - \epsilon \leq q$, and hence $w \in X_{p\leq h\leq q} \subseteq S_{\prec v}$. If $s(w) < s(v)$, then $w \prec v$ by the definition of $\preceq$ and hence $w \in S_{\prec v}$. As such, B is precisely $S_{\prec v} \cap \delta st(v)$. We can see B is homeomorphic to $lk^{(h,s)\downarrow}_{p\leq h}(w)$ by the definition of lk and st. Hence the quotient 4.1 is at least k connected. □

The Morse lemma is a powerful tool for determining connectivity. However, it is not an if and only if criterion. As such, we cannot use it to demonstrate that a space is not k-connected or essentially k-connected, which proves necessary when we wish to demonstrate that a character χ is not in the k-th sigma invariant. Thankfully, there is a corollary to the Morse lemma that provides an if and only if criterion on the essential connectivity of a space.

Corollary 4.2.2. *([WZ15], Observation 1.8) Let an $m-1$ connected affine cell complex X be equipped with a Morse function $(h,s) : X \to \mathbb{R} \times \mathbb{R}$ and assume that all ascending links are $m-2$ connected. Then the filtration $\{X_{t\leq h}\}_{h\in\mathbb{R}}$ is essentially $(m-1)$-connected if and only if $X_{p\leq h}$ is $m-1$ connected for some p, if and only if all $X_{p'\leq h}$ are $m-1$ connected for all $p' \leq p$.*

Proof. Ascending links in X are $(m-2)$ connected by assumption. As such we can apply the Morse lemma 4.2.1 to determine that the quotient space $X_{p\leq h}$ $X_{q\leq h}$ is $m-1$ connected for any $p < q$. In particular, this implies the inclusion maps $\pi_k(X_{q\leq h}) \hookrightarrow \pi_k(X_{p\leq h})$ are isomorphisms for $k \leq m-1$ and surjective when $k = m-1$. These maps would be the trivial maps if and only if the homotopy groups of $X_{p\leq h}$ are trivial, which is equivalent to $X_{p\leq h}$ being $(m-1)$ connected. Hence the filtration is essentially $(m-1)$ connected if and only if $X_{p\leq h}$ is $(m-1)$ connected for some p, equivalently all $p' \leq p$ by the definition of essential connectivity (see 1.3.6). □

Chapter 5

BNSR Invariants for Quadratic Bieri-Strebel Groups

5.1 An Alternate Calculation for $\mathbf{F}_\tau$

The following section is the result of collaborative work between Lewis Molyneux, Brita Nucinkis, and Yuri Santos Rego. Work contributed in its entirety by the other authors will be labelled.

While our primary calculation of BNSR invariants will follow Zaremsky's method [Zar17], we may employ an alternate method for $\mathbf{F}_\tau$, the Bieri-Strebel group with subdivision polynomial $x^2 + x - 1$. We may combine the finite index theorems 1.4.1 and 1.4.2 with the method of Bieri, Geoghegan and Kochloukova for calculating the sigma invariant of $\mathbf{F}$ in order to calculate the BNSR invariant.

5.1.1 Abelianization and The Character Sphere

We begin with some initial facts. We may cite ([BNR21], chapter 5) as a source for the abelianization of $\mathbf{F}_\tau$, which is $\mathbb{Z}^2 \oplus \mathbb{Z}/2\mathbb{Z}$. As such $r_0(\mathbf{F}_\tau) = \mathbb{Z}^2$ and therefore $S(\mathbf{F}_\tau) = S^1$. This means that we expect to be able to form a basis of $Hom(\mathbf{F}_\tau, \mathbb{R})$ from 2 linearly independent characters. We can now cite Bieri and Strebel ([BS92], Chapter 3 section 3) which tells us that any group G of PL-homeomorphisms of the interval will have a character χ_0 such that $\chi_0(f) = ln(f'(0))$ and a character χ_1 such that $\chi_1(f) = ln(f'(1))$ for all $f \in G$. As all characters in $Hom(\mathbf{F}_\tau, \mathbb{R})$ fall into an equivalence class in the character sphere, we can choose to represent the equivalence classes $[\chi_0]$ and $[\chi_1]$ with any functions of the form $log_k(f'(0)$ and $log_k(f'(1))$ respectively, where $k \in \mathbb{R}_{\geq 0}$. In the case of $\mathbf{F}_\tau$, we will choose $\chi_0(f) = log_\tau(f'(0))$ and $\chi_1(f) = log_\tau(f'(1))$. As all slopes of functions in $\mathbf{F}_\tau$ are in $\langle \tau \rangle$, this will result in $\chi_0(f), \chi_1(f) \in \mathbb{Z}$ for all $f \in \mathbf{F}_\tau$.

We can demonstrate the linear independence of these two characters by selecting two elements f, g of $\mathbf{F}_\tau$ such that $\chi_0(f) = 1, \chi_0(g) = 0$ and $\chi_1(f) = 0, \chi_1(g) = 1$. As an example, we choose f and g as the following elements of $\mathbf{F}_\tau$.

$$f(x) = \begin{cases} \tau x & \text{for } 0 \leq x \leq \tau^2 \\ \tau^{-1}x - \tau^2 & \text{for } \tau^2 \leq x \leq \tau \\ x & \text{for } \tau \leq x \leq 1 \end{cases}$$

$$g(x) = \begin{cases} x & \text{for } 0 \leq x \leq \tau^2 \\ \tau^{-1}x - \tau^3 & \text{for } \tau^2 \leq x \leq \tau \\ \tau x + \tau^2 & \text{for } \tau \leq x \leq 1. \end{cases}$$

We can simply read off the slope at 0 and 1 for each of these elements, so it is easy to see that they fulfil our requirements for f and g, thus demonstrating linear independence. As we have two linearly independent characters of $Hom(\mathbf{F}_\tau, \mathbb{R})$, and we know from the abelianization that $Hom(\mathbf{F}_\tau, \mathbb{R}) \cong \mathbb{R}^2$ we know that these characters must span $Hom(\mathbf{F}_\tau, \mathbb{R})$, ie: we can write any character of $\mathbf{F}_\tau$ as $a\chi_0 + b\chi_1$, with $a, b \in \mathbb{R}$.

5.1.2 The Subgroup K

Important to our calculation of the BNSR invariant of $\mathbf{F}_\tau$ is the finite index subgroup $K \subseteq \mathbf{F}_\tau$. Using the generating set from the presentation of $\mathbf{F}_\tau$ given in 2.4, we define K as the following subgroup

$$K = \langle x_0, x_1, y_1, x_2, y_2 ... \rangle$$

As such, K is the subgroup generated by all generators of $\mathbf{F}_\tau$ except for y_0.

Lemma 5.1.1. *([MNSR24], proposition 4.2)* $|\mathbf{F}_\tau : K| = 2$ *and* $K_{ab} \cong \mathbb{Z}^2 \oplus \mathbb{Z}/2\mathbb{Z}$

Proof. To demonstrate $|\mathbf{F}_\tau : K| = 2$, we claim $\mathbf{F}_\tau = K \sqcup y_0 K$. We shall use the normal form for elements of $\mathbf{F}_\tau$ introduced by Burillo, Nucinkis and Reeves ([BNR21], Theorem 7.3). We may write any element $f \in \mathbf{F}_\tau$ in the form

$$g = x_0^{i_0} y_0^{\epsilon_0} x_1^{i_1} y_1^{\epsilon_1} ... x_n^{i_n} y_n^{\epsilon_n} x_m^{-j_m} x_{m-1}^{-j_{m-1}} x_0^{j_0}$$

with $i_0, ..., i_n, j_0, ..., j_m \in \mathbb{Z}_{\geq 0}$ and $\epsilon_0, ..., \epsilon_n \in \{0, 1\}$. Using the relations from 2.4, any time we have $x_k^{i_k} y_0$ with $k \geq 1$, we may rewrite it as $y_0 x_{i+1}$. As such any coset fK, $f \in \mathbf{F}_\tau$ we can write $fK = x_0^{i_0} y_0^{\epsilon_0} K$. Cleary $fK = K$ when $\epsilon_0 = 0$, but we need to check the case when $\epsilon_0 = 1$. In this case, we can repeat the following calculation

$$\begin{aligned} x_0^{i_0} y_0 &= x_0^{i_0 - 1} x_0 y_0 \\ &= x_0^{i_0 - 1} x_0 x_1 x_1^{-1} y_0 \\ &= x_0^{i_0 - 1} y_0^2 x_1^{-1} y_0 \\ &= x_0^{i_0 - 1} y_0^2 x_1^{-1} y_0 \\ &= x_0^{i_0 - 1} y_0^2 y_0 x_2^{-1} \\ &= x_0^{i_0 - 1} y_0 y_0^2 x_2^{-1} \\ &= x_0^{i_0 - 1} y_0 x_0 x_1 x_2^{-1} \end{aligned} \tag{5.1}$$

This allows us to shift an x_0 to the right of the y_0 one at a time, until we are left with $y_0 K$. Hence each coset of K is either K or $y_0 K$, and as cosets of a subgroup are disjoint and partition the group, we know that $= K \sqcup y_0 K$, and hence $|\mathbf{F}_\tau : K| = 2$.

To calculate the abelianization of K, we will take our presentation for K and abelianize the generators. For clarity, we will write abelian composition additively. Considering the relations carried over from the presentation 2.4, we can see that $\bar{x}_j + \bar{x}_i = \bar{x}_i + x_{\bar{j+1}}$ immediately reduces to $\bar{x}_j = x_{\bar{j+1}}$ for all $j \geq 1$. Similarly, $\bar{y}_j = y_{\bar{j+1}}$ for all $j \geq 1$. This reduces the generating set to $\bar{x}_0, \bar{y}_0, \bar{x}_1$ and $\bar{y}_1$. We can eliminate $\bar{y}_1$ with the relation $2\bar{y}_1 = \bar{x}_1 + \bar{x}_2 = 2\bar{x}_1$, implying $\bar{y}_1 = \bar{x}_1$. This leaves us with 3 generators.

$\bar{x_0}$ and $\bar{x_1}$ are free abelian generators, but we can see that $\bar{y_0}$ is of order 2 via the relation $2\bar{y_0} = \bar{x_0}\bar{x_1}$. The abelian group with these generators is $\mathbb{Z}^2 \oplus \mathbb{Z}/2\mathbb{Z}$, and hence $K_{ab} \cong \mathbb{Z}^2 \oplus \mathbb{Z}/2\mathbb{Z}$. □

A concept important to Bieri, Geoghegan and Kochloukova's calculation of the BNSR invariant of **F** is HNN-extensions. As we are adapting their calculation for $\mathbf{F}_\tau$, we will also need to understand HNN extensions.

Definition 5.1.2. Let H be a group and $\sigma : H \to H$ be a monomorphism. An ascending HNN extension with base H is a group given by the presentation

$$H*_{t,\sigma} = \langle H, t | tht^{-1} = \sigma(h)\, \forall h \in H \rangle$$

We will now introduce the subgroup $\mathbf{F}_\tau[1]$, defined as the subgroup of $\mathbf{F}_\tau$ generated by $\{x_1, y_1, x_2, y_2, ...\}$. We can see that this subgroup is isomorphic to $\mathbf{F}_\tau$ via the map γ that sends x_i to x_{i-1} and y_i to y_{i-1}. Indeed, the map $\gamma_n : \mathbf{F}_\tau[n] \to \mathbf{F}_\tau$ that sends x_i to x_{i-n} and y_i to y_{i-n} is an isomorphism and as such all $\mathbf{F}_\tau[n]$ are isomorphic to $\mathbf{F}_\tau$, and therefore have the F_∞ property.

Lemma 5.1.3. *The subgroup $K \subseteq \mathbf{F}_\tau$ is isomorphic to the ascending HNN extension $\mathbf{F}_\tau[1]*_{t,\gamma^{-1}}$, where γ maps x_i to x_{i-1} and y_i to y_{i-1}.*

Proof. From the definition of an ascending HNN extension 5.1.2 we can see that $L = \mathbf{F}_\tau[1]*_{t,\gamma^{-1}} = \langle \mathbf{F}_\tau[1], t | tx_it^-1 = x_{i+1}, ty_it^-1 = y_{i+1} \rangle$. For any element f in $\mathbf{F}_\tau[1]$, we may write it in its normal form (inherited from the normal form of $\mathbf{F}_\tau$), then conjugate it by t. This is the same as conjugating all of the generators that make up f by t. As the conjugation of a generator by t is defined, the conjugation of any element of $\mathbf{F}_\tau[1]$ is defined.

We claim that the following map is a group isomorphism

$$\begin{aligned} \phi : \mathbf{F}_\tau[1]*_{t,\gamma^{-1}} &\to K \\ x_i &\mapsto x_i\, \forall i \geq 1 \\ y_i &\mapsto y_i\, \forall i \geq 1 \\ t &\mapsto x_0 \end{aligned} \tag{5.2}$$

We can immediately see that ϕ is surjective, as it maps onto every generator of K. Furthermore, the HNN relations $tx_it^-1 = x_{i+1}, ty_it^-1 = y_{i+1}$ are recreating the relations $x_ix_0 = x_0x_{i+1}$ and $y_ix_0 = x_0y_{i+1}$ from K, and all other relations of K are the same as those inherited from $\mathbf{F}_\tau[1]$. Hence ϕ is a group homomorphism. All that remains to be shown is that ϕ is injective.

To show ϕ is injective, we will begin by noting that $\mathbf{F}_\tau[1]$ is a subgroup of L, and thus any element of $\mathbf{F}_\tau[1] \subseteq L$ can be written in the normal form inherited from $\mathbf{F}_\tau$. Now consider an element $w \in ker(\phi) \subseteq L$. As L is an HNN extension, we may write w in the form

$$w = g_0 t^{\epsilon_1} g_1 t^{\epsilon_2} ... g_{n-1} t^{\epsilon_{n-1}} g_n$$

where $g \in \mathbf{F}_\tau[1]$ and $\epsilon_i \in \{\pm 1\}$. Using the HNN relations for L, we can rewrite w by using the following two relations

$$g_i t^{-1} = t^{-1} g_i' \text{for some } g_i' \in \mathbf{F}_\tau[1]$$
$$t g_i = g_i' t \text{for some } g_i' \in \mathbf{F}_\tau[1]$$

Through these relations we may rewrite w into the form $w = t^{-a} g' t^b$ where $g' \in \mathbf{F}_\tau[1]$, $a, b \in \mathbb{Z}_{\geq 0}$. Now consider $\phi(w) = \phi(t^{-a})\phi(g')\phi(t^b)$. By assumption, $w \in ker(\phi)$ and therefore $\phi(w) = 1$. Suppose $\phi(w)$ is in normal form. Then $g' = 1, a = b$ and therefore $w = 1$. Now suppose $\phi(w)$ is not in normal form. We may rewrite $\phi(w)$ as $x_0^a \phi(g) x_0^{-b}$. Using the x_0 relations we may shift the x_0^{-b} to the left of $\phi(g)$. This leaves us with $\phi(w) = x_0 a - b\phi(g[b])$, where $g[b]$ is g but with each x_i and y_i is replaced with x_{i+b} and y_{i+b} respectively. This is now a normal form. As such, we can conclude $a = b$ and $g[b] = 1$. $g[b] = 1$ immediately implies $g = 1$. As such $w = 1$ and thus each $w \in ker(\phi)$ is the trivial element and ϕ is injective. Hence ϕ is an isomorphism. □

Here, we wish to cite the work of Bieri, Geoghegan and Kochloukova, which will allow us to leverage our representation of K as an ascending HNN extension into a method of calculate the BNSR invariant.

Citation 5.1.4. *([BGK10], Theorem 2.1) Let G decompose as an ascending HNN extension $H*_{\sigma,t}$. Let χ be a character such that $\chi(H) = 0$, $\chi(t) = 1$.*

1. *If H is of type F_n, then $[\chi] \in \Sigma^n(G)$.*
2. *If H is of type FP_n over a ring R, then $[\chi] \in \Sigma^n(G; R)$*
3. *If H is finitely generated and γ is not surjective, then $[-\chi] \notin \Sigma^1(G)$*

For Thompson's group $\mathbf{F}$, Bieri, Geoghegan and Kochloukova are able to show that $\mathbf{F}$ can be written as an ascending HNN extension with base group $\mathbf{F}[1]$ (defined analogously to $\mathbf{F}_\tau[1]$) and so are able to use 5.1.4 to determine the BNSR invariant of $\mathbf{F}$ directly. Our goal is to instead use 5.1.4 to determine properties of the BNSR invariant of K, then use 1.4.1 and 1.4.2 to imply properties about the BNSR invariant of $\mathbf{F}_\tau$.

5.1.3 Calculating the Invariant

Theorem 5.1.5. *Given $\mathbf{F}_\tau$ with the presentation 2.4, and the characters $\chi_0, \chi_1 \in Hom(\mathbf{F}_\tau, \mathbb{R})$ such that $\chi_0(x_0) = 1, \chi_0(x_1) = 1$ and $\chi_1(x_0) = 0, \chi_1(x_1) = 0$, then the BNSR invariant of $\mathbf{F}_\tau$ is as follows*

1. $\Sigma^1(\mathbf{F}_\tau) = \Sigma^1(\mathbf{F}_\tau, \mathbb{Z}) = S(\mathbf{F}_\tau) \setminus \{[-\chi_0], [-\chi_1]\}$
2. $\Sigma^\infty(\mathbf{F}_\tau) = \Sigma^\infty(\mathbf{F}_\tau, \mathbb{Z}) = \Sigma^2(\mathbf{F}_\tau) = \Sigma^1(\mathbf{F}_\tau) \setminus \{-a\chi_0 - b\chi_1] | a, b > 0\}$

The proof for 5.1.5 will be broken down into 3 main parts. These sections of the proof relate to 3 groupings of points of $S(\mathbf{F}_\tau)$: The single points $[\chi_0], [\chi_1], [-\chi_0]$ and $[-\chi - 1]$ (discussed in 5.1.6; the set of points $\{-a\chi_0 - b\chi_1] | a, b > 0\}$, described as the short interval (discussed in 5.1.11); and the set of points $\{a\chi_0 + b\chi_1] | b > 0\} \cup \{a\chi_0 + b\chi_1] | a > 0\}$, described as the long interval (discussed in 5.1.8).

One more useful tool for calculating the BNSR invariant of $\mathbf{F}_\tau$ is the automorphism $\mu : \mathbf{F}_\tau \to \mathbf{F}_\tau$, where for f an element of $\mathbf{F}_\tau$ considered as a PL-homeomorphism, $\mu(f)$ is f conjugated by the function $t \mapsto 1 - t$. By inspection, we can see that $\mu(f)'(0) = f'(1)$ and $\mu(f)'(1) = f'(0)$ for all $f \in \mathbf{F}_\tau$. As such, $\chi_0(f) = \chi_1(\mu(f))$ and

$\chi_1(f) = \chi_0(\mu(f))$. We may use automorphisms in conjunction with BNSR invariants in the following way

$$\begin{aligned} & f \in \Sigma^N(\mathbf{F}_\tau) \\ \iff & \mu(f) \in \Sigma^N(\mu(\mathbf{F}_\tau)) \\ \iff & \mu(f) \in \Sigma^N(\mathbf{F}_\tau) \end{aligned} \tag{5.3}$$

We will refer to the automorphism μ and its application to BNSR calculation as μ symmetry, and it is analogous to the v symmetry used by Bieri, Geoghegan and Kochloukova for **F** ([BGK10], Section 1.4).

Lemma 5.1.6. *Let χ_0 and χ_1 be as described in 5.1.5. Then*

$$[\chi_0], [\chi_1] \in \Sigma^\infty(\boldsymbol{F}_\tau) and [-\chi_0], [-\chi_1] \notin \Sigma^1(\boldsymbol{F}_\tau)$$

$$[\chi_0], [\chi_1] \in \Sigma^\infty(\boldsymbol{F}_\tau; \mathbb{Z}) and [-\chi_0], [-\chi_1] \notin \Sigma^1(\boldsymbol{F}_\tau; \mathbb{Z})$$

Proof. By inspection of 2.5, we can see that $x_i'(0) = y_i'(0) = 0$ for all $i \geq 1$. As such, we can see $\chi_0(\mathbf{F}_\tau[1]) = 0$. Similarly, we can see that $\chi_0(x_0) = 2$. As such, $\frac{1}{2}\chi_0$ is a character such that $\frac{1}{2}\chi_0(\mathbf{F}_\tau[1]) = 0$ and $\frac{1}{2}\chi_0(x_0) = 1$. From 5.1.3 we know that the subgroup K is isomorphic to the ascending HNN extension $\mathbf{F}_\tau[1]*_{t,\gamma^{-1}}$ with t being mapped to x_0 by the isomorphism. From this we may apply 5.1.4 to determine that $[\chi_0] \in \Sigma^\infty(K)$, as $\mathbf{F}_\tau[1]$ has the F_∞ property. We may also determine that $[-\chi_0] \notin Sigma^1(K)$, as γ^{-1} is not surjective on $\mathbf{F}_\tau[1]$. We can then use 1.4.1 to determine $[\chi_0] \in \Sigma^\infty(\mathbf{F}_\tau)$ and $[-\chi_0] \notin Sigma^1(\mathbf{F}_\tau)$. An application of μ-symmetry then gets us $[\chi_1] \in \Sigma^\infty(\mathbf{F}_\tau)$ and $[-\chi_1] \notin Sigma^1(\mathbf{F}_\tau)$. We may then use the relations between the homotopical and homological BNSR invariants discussed in 1.3.2 to conclude that $[\chi_0], [\chi_1] \in \Sigma^\infty(\mathbf{F}_\tau; \mathbb{Z})$ and $[-\chi_0], [-\chi_1] \notin \Sigma^1(\mathbf{F}_\tau; \mathbb{Z})$.

As a brief corollary, we may apply 1.4.1 and 1.4.2 to determine that $Sigma^n(K)$ must have the same shape as $\Sigma^n(\mathbf{F}_\tau)$ in particular concluding that $[\chi_1] \in \Sigma^\infty(K)$ and $[-\chi_1] \notin \Sigma^1(K)$, as well as that $[\chi_0], [\chi_1] \in \Sigma^\infty(K; \mathbb{Z})$ and $[-\chi_0], [-\chi_1] \notin \Sigma^1(K; \mathbb{Z})$. □

For the next section of the proof, we will need to cite another result of Bieri, Geoghegan and Kochloukova relating to HNN extensions

Citation 5.1.7. *([BGK10], Theorem 2.3) Let G decompose as an ascending HNN extension $H*_{\phi,t}$. Let $\chi : G \to \mathbb{R}$ be a character such that $\chi|_H$ is non-trivial. If H is of type F_∞ and $[\chi|_H] \in \Sigma^\infty(H)$, then $[\chi] \in \Sigma^\infty(G)$.*

Lemma 5.1.8. *Let χ_0 and χ_1 be as described in 5.1.5. Let χ be an arbitrary character in $Hom(\boldsymbol{F}_\tau, \mathbb{R})$. Then*

$$\chi \in \{a\chi_0 + b\chi_1 | b > 0\} \cup \{a\chi_0 + b\chi_1 | a > 0\} \Rightarrow [\chi] \in \Sigma^\infty(\boldsymbol{F}_\tau)$$

$$\chi \in \{a\chi_0 + b\chi_1 | b > 0\} \cup \{a\chi_0 + b\chi_1 | a > 0\} \Rightarrow [\chi] \in \Sigma^\infty(\boldsymbol{F}_\tau; \mathbb{Z})$$

Proof. We begin by considering a character $\chi \in Hom(K, \mathbb{R})$ and $H = \mathbf{F}_\tau[1]$ as a subgroup of K. We claim that $\chi(x_1) > 0 \Rightarrow \chi|_H \in [\chi_1] \in S(H)$. From 1.4.4 we know that $i* : Hom(\mathbf{F}_\tau, \mathbb{R}) \to Hom(K, \mathbb{R}); \chi \mapsto \chi|_K$ is an isomorphism of vector spaces, and thus we may write any character $\chi \in Hom(K, \mathbb{R})$ as $\chi = a\chi_0|_k + b\chi_1|_K$, $a, b \in \mathbb{R}$. Consider $\chi_H = a\chi_0|_H + b\chi_1|_H$. As discussed in the proof of 5.1.6, $\chi_0|_H$ is trivial. Thus $\chi_H = b\chi_1|_H$. This means that for $\chi \in Hom(K, \mathbb{R})$, $\chi\_H \in [\chi_1|_H]$ or $\chi\_H \in [-\chi_1|_H]$. As χ_1 is defined such that $\chi_1(x_1) = 1$ any χ such that $\chi(x_1) > 0$ must be in $[\chi_1|_H]$ by the definition of the equivalence class $[\chi_1|_H]$.

We now recall that $H = \mathbf{F}_\tau[1]$ is isomorphic to $\mathbf{F}_\tau$ via the isomorphism $\gamma : \mathbf{F}_\tau[1] \to \mathbf{F}_\tau$ that maps x_i to x_{i-1}. This induces the homeomorphism $\gamma* : S(\mathbf{F}_\tau[1]) \to S(\mathbf{F}_\tau)$ onto the character spheres, which sends $[\chi_1|_H]$ to $[\chi_1]$. As $[\chi_1] \in \Sigma^\infty(\mathbf{F}_\tau)$, $[\chi_1|_H] \in \Sigma^\infty(H)$. We can now apply 5.1.7 to conclude that $\chi(x_1) > 0 \Rightarrow [\chi] \in \Sigma^\infty(K)$. Applying 1.4.1 gets us $\chi(x_1) > 0 \Rightarrow [\chi] \in \Sigma^\infty(\mathbf{F}_\tau)$. Thus we can see that $\chi \in \{a\chi_0 + b\chi_1]|b > 0\} \Rightarrow [\chi] \in \Sigma^\infty(\mathbf{F})$.

From here, we may apply μ symmetry once again. The homeomorphism $\mu* : S(\mathbf{F}_\tau) \to S(\mathbf{F}_\tau)$ induced by the automorphism μ will map the half-circle interval $\{[a\chi_0 + b\chi_1]|b > 0\}$ to the half circle interval $\{a\chi_0 + b\chi_1]|a > 0\}$ and thus we can conclude $\chi \in \{a\chi_0 + b\chi_1]|a > 0\} \Rightarrow [\chi] \in \Sigma^\infty(\mathbf{F})$, thus proving the lemma for the homotopical BNSR invariant. We may apply this to the homological invariant by once again using the relations from 1.3.2. □

All that remains to be determined of the BNSR invariant of $\mathbf{F}_\tau$ (both homotopical and homological) is the short interval of characters of the form $-a\chi_0 - b\chi_1]; a, b > 0$. In order to perform this last part of the calculation, we will need to cite two more results from Bieri, Geoghegan and Kochloukova.

Citation 5.1.9. *([BGK10], Corollary 1.2) The kernel of a nonzero discrete character χ has type FP_n over the ring R if and only if both $[\chi]$ and $[-\chi]$ are in $\Sigma^n(G, \mathbb{R})$.*

Citation 5.1.10. *Suppose G is a group that contains no nonabelian free subgroups and is of type FP_2 over a ring R. Let $\tilde{\chi} : G \to \mathbb{R}$ be a nonzero discrete character. Then G decomposes as an ascending HNN extension $H*_{t,\gamma}$ where H is a finitely generated subgroup of $ker(\tilde{\chi})$ and $\tilde{\chi}(t)$ generates $im(\tilde{\chi})$.*

Lemma 5.1.11. *Let R be a ring, then $\chi \in \{-a\chi_0 - b\chi_1]|a, b > 0\} \Rightarrow \chi \notin \Sigma^2(\mathbf{F}_\tau; R)$.*

Proof. Initially, we can see that the discrete characters are dense in the short interval and, as the interval is open (since it is missing the end points $[-\chi_0]$ and $[-\chi_1]$), we only need to show that the discrete characters are not in $\Sigma^2(\mathbf{F}_\tau; R)$ (see for example, [BGK10], Proposition 2.9). Take a discrete character $\chi \in Hom(\mathbf{F}_\tau, \mathbb{R})$ written in the form $\chi = a\chi_0 + b\chi_1$ with $a, b \in \mathbb{Q} \setminus \{0\}$. Using the elements $f, g \in F$ from the section 5.1.1, we can an element $t \in \mathbf{F}_\tau$ such that

$$\chi_0(t) = mb\,\chi_1(t) = -ma$$

for some m in $\mathbb{Q} \setminus \{0\}$. This gives us that $\chi(t) = 0$. Since χ has a discrete image in $\mathbb{R}$ and $a \neq 0$, we have that there must exist a t_0 such that $|\chi_0(t_0)|$ is minimal among all elements of $\mathbf{F}_\tau$ that fulfill the properties of t. In particular, $\chi_0(t_0)$ is nonzero.

Let $G = ker(\chi)$. Since the abelianization of $\mathbf{F}_\tau$ is $\mathbb{Z}^2 \oplus \mathbb{Z}/2\mathbb{Z}$, we have that $G = \langle \sqrt{\mathbf{F}'_\tau}, t_0 \rangle = sqrt\mathbf{F}'_\tau \rtimes \langle t_0 \rangle$, where $sqrt\mathbf{F}'_\tau := \{f \in \mathbf{F}_\tau | f^n \in \mathbf{F}'_\tau \text{for some } n\}$. By inspection, we can see that $\chi_0|_G$ is a discrete, nonzero character that is trivial on the subgroup $sqrt\mathbf{F}'_\tau$ (as all characters are) and that $im(\chi_0|G)$ is generated by $\chi_0(t_0)$, since $|\chi_0(t_0)|$ is minimal among t.

Suppose that G has type FP_2 over a ring R. By 5.1.10, we may decompose G as the HNN extension $H*_{t,\gamma}$, where H is some finitely generated subgroup of $\sqrt{\mathbf{F}'_\tau}$. As H is a subgroup of $\sqrt{\mathbf{F}'_\tau}$, we know that all elements h in H must have that $\chi_0(h) = 0$, $\chi_1(h) = 0$. As such, the slopes $h'(0)$ and $h'(1)$ must be 0. That is to say that no elements of H have support in some region around 0 and 1. As H is finitely generated, there is a generator h_0 of H such that h_0 has a break point with minimal value (ie: closest to 0) among all elements of H. Similarly, there must also exist a generator h_1 with a breakpoint of maximal value (ie:, closest to 0). As such, there exists a value ϵ'' such that all elements of H are supported in the interval $[\epsilon'', 1-\epsilon'']$. Similarly, we know t_0 is an element of $\mathbf{F}_\tau$, so by the definition of $\mathbf{F}_\tau$ given in 2.3.1, t_0 has finitely many breakpoints. As such, there is a value ϵ' such that all breakpoints of t_0 are in the interval $[\epsilon', 1-\epsilon']$, and so t_0 is linear on the intervals $[0,\epsilon']$ and $[1-\epsilon', 1]$. Take $\epsilon = \mathbf{min}\{\epsilon', \epsilon''\}$. By construction, H is supported on $[\epsilon, 1-\epsilon]$ and t_0 is linear on the intervals $[0,\epsilon]$ and $[1-\epsilon, 1]$.

Since we have $\sqrt{\mathbf{F}'_\tau} \rtimes \langle t_0 \rangle = G \cong H*_{t,\gamma}$, we can say that $\sqrt{\mathbf{F}'_\tau} = \bigcup_{n\geq 1} t^n H t^{-n}$. As such, for each $f \in \sqrt{\mathbf{F}'_\tau}$, there is some n such that $t^{-n} f t^n$ is in H. Hence $t^{-n} f t^n$ is supported in the interval $[\epsilon, 1-\epsilon]$. From here, we can see that any f in $\sqrt{\mathbf{F}'_\tau}$ must be supported in $[t_0^n(\epsilon), t_0^n(1-\epsilon)]$ for some n. As $\sqrt{\mathbf{F}'_\tau}$ has support in $(0,1)$, there must be a subsequence of $\{t_0^n(\epsilon)\}_{n\in\mathbb{N}}$ that converges to 0, and similarly there must be a subsequence of $\{t_0^n(1-\epsilon)\}_{n\in\mathbb{N}}$ that converges to 1. However, we know t_0 is linear outside the interval $[\epsilon, 1-\epsilon]$. As such, in order for $\{t_0^n(\epsilon)\}_{n\in\mathbb{N}}$ to approach 0, $t_0(\epsilon) < \epsilon$ and so $t'_0 < 1$ on the interval $[0,\epsilon]$. Similarly, for $\{t_0^n(1-\epsilon)\}_{n\in\mathbb{N}}$ to approach 1, we must have $t_0(1-\epsilon) > 1-\epsilon$ and so $t'_0 > 1$ on the interval $[1-\epsilon, 1]$. Hence, for our discrete character $\chi = a\chi_0 + b\chi_1$, we must have $ab < 0$. We concluded this after assuming that G has type FP_2 over a ring R. This leaves us with the implication

$$\chi = a\chi_0 + b\chi_1,\ a, b \in \mathbb{Q} \text{ and } ker(\chi) \text{is of type} FP_2 \Rightarrow ab < 0$$

The contrapositive of this states that for a discrete character $\chi = a\chi_0 + b\chi_1$, $ab > 0 \Rightarrow ker(\chi)$ is not of type FP_2. By 5.1.9, this means that either $[\chi]$ or $[-\chi]$ is not in $\Sigma^2(\mathbf{F}_\tau; R)$. However, we already know from 5.1.8 that $\chi = a\chi_0 + b\chi_1$ is in $\Sigma^\infty(\mathbf{F}_\tau; R)$ for $a, b > 0$. Hence, $\chi = a\chi_0 + b\chi_1$ is not in $\Sigma^2(\mathbf{F}_\tau; R)$ for $a, b < 0$.

We may refer back to the relations between homotopical and homological BNSR invariants in 1.3.2 to exclude these characters from $\Sigma^2(\mathbf{F}_\tau)$ as well. □

Between the three lemmas 5.1.6, 5.1.8 and 5.1.11, the only point of ambiguity is whether points on the short interval are in $\Sigma^1(\mathbf{F}_\tau)$. However, we may cite Bieri and Strebel's work on groups of PL-homeomorphisms. In particular, ([BS16], Chapter IV, corollary 3.4) tells us that for any group of PL-homeomorphisms G, $\Sigma^1(G) = S(G) \setminus \{[\chi_0], [\chi_1]\}$, where χ_0 and χ_1 are the characters that measure the slope at 0 and the slope at 1, respectively. This tells us that the ambiguous short interval characters are in $\Sigma^1(\mathbf{F}_\tau)$. Combining this with our three lemmas is our proof of 5.1.5.

5.2 BNSR invariants via Morse Theory

5.2.1 Abelianizations and r_0 for Bieri-Strebel Groups

In order to calculate the BNSR invariants for quadratic Bieri-Strebel groups, we must first determine the equivalence classes of characters in $S(\mathbf{F}_\beta)$. The easiest way to do this is to find a spanning set of characters for $Hom(\mathbf{F}_\beta, \mathbb{R})$. Once we have a spanning set, we may express each character of $Hom(\mathbf{F}_\beta, \mathbb{R})$ as a linear combination of our spanning characters, and then intersect $Hom(\mathbf{F}_\beta, \mathbb{R})$ with a sphere of any radius around the origin to find class representatives for the equivalence class of $S(\mathbf{F}_\beta)$.

Considering the characters of $\mathbf{F}_n$ discussed in 2.1.4, we would like to apply similar methods of constructing characters to quadratic Bieri-Strebel groups. For the initial step, we may simply cite Bieri and Strebel's work on groups of piecewise-linear homeomorphisms. In particular, ([BS92], Chapter IV, Lemma 3.1) indicates that all quadratic Bieri Strebel groups have $\chi_0(f) = ln(f'(0))$ and $\chi_1(f) = ln(f'(1))$ as linearly independent characters in $Hom(\mathbf{F}_\beta), \mathbb{R}$.

Our next step is to show that characters similar to the ψ characters described in 2.1.4 can be constructed for quadratic Bieri-Strebel groups with well defined tree-pair representations. To begin, we can see that the group $\mathbf{F}_\beta$ with subdivision polynomial $ax^2 + bx_1$ will have $a + b - 1$ orbits of breakpoints. We make a similar argument to that for $\mathbf{F}_n$, as illustrated in 2.5. Any caret in the caret set C_β will have $a + b$ legs, and so the addition of any caret to tree in a tree pair for $\mathbf{F}_\beta$ will add $a + b - 1$ leaves to the tree (as it replaces 1 leaf with $a + b$ leaves). While the legs of different length might initially appear to split these orbits into smaller orbits, they provide no obstruction. Each breakpoint b sits in $\mathbb{Z}[\beta]$, and so can be expressed in the form $\Sigma_{i=1}^{n} a_n \beta^n$. We begin by considering a_1. If a_1 is greater than the number of legs of length 1 in a caret of C_β then we may rewrite b by using the identity $\beta = a\beta^3 + b\beta^2$ until it is less than or equal to the number of legs of length 1. We may then use a right hand caret (ie: a caret with all long legs on the right) as the top caret of a tree constructed to include b as a breakpoint. We then attach a right caret to the $a_1 + 1$th leg of the top caret and repeat this process with a_2. Repeating this process n times will construct a tree with b as a breakpoint. The argument from here is entirely analogous to the $\mathbf{F}_n$ case, allowing us to conclude there are $a + b - 1$ orbits of breakpoints for the group $\mathbf{F}_\beta$. Note that this result can also be found in ([Win], Theorem 4.8.6).

Having established that the group $\mathbf{F}_\beta$ has $a + b - 1$ orbits on the breakpoints $\mathbb{Z}[\beta] \cap (0, 1)$, we may define characters similar to the ψ characters defined for $\mathbf{F}_n$ in 2.1.4. We define the character ψ_i, $i \in 1, ..., a + b - 1$ in the following way: Let X_i be an orbit of breakpoints of F_β in $\mathbb{Z}[\beta] \cap [0, 1]$, with $x \in X_i$ a breakpoint. For $f \in F_\beta$, we may consider $\gamma_x(f) = log_\beta(f'_\geq(x)) - log_\beta(f'_\leq(x))$, where $f'_\geq(x)$ is the slope immediately to the right of the breakpoint x and $f'_\leq(x)$ is the slope immediately to the left. We then define the character $\psi_i(f) = \Sigma_{x \in X_i} \gamma_x(f)$.

Considering χ_0, χ_1 and the ψ_i together, we now have $a + b + 1$ characters for $\mathbf{F}_\beta$. However, we do not expect these characters to be linearly independent. In fact, we can demonstrate a linear dependence between them. Each $\gamma_x(f)$ measures the change in slope of f at the breakpoint x (expressed additively), and hence $\psi_i(f)$ expresses the net change in the slope of f across all breakpoints in the orbit X_i. We can therefore see that the sum $\Sigma_{i=1}^{a+b-1} \psi_i(f)$ will measure the net change in the slope of f over all breakpoints. As f is piecewise linear, we know the slope can only change at breakpoints, so the total net change in the slope of f between 0 and 1 is measured by $\Sigma_{i=1}^{a+b-1} \psi_i(f)$. This means we can write the linear dependence as $\Sigma_{i=1}^{a+b-1} \psi_i(f) = \chi_1(f) - \chi_0(f)$. Removing any one character from our set of $a + b + 1$ characters (we will typically choose ψ_{a+b-1}) will leave us with a linearly independent set of $a + b$ characters.

We now wish to show, where possible, that $r_0(\mathbf{F}_\beta) \cong \mathbb{Z}^{a+b}$. When this is the case, then we know from ([BS92], Chapter I, Lemma 1.1) that $Hom(\mathbf{F}_\beta, \mathbb{R}) \cong \mathbb{R}^{a+b}$ and is therefore spanned by our linearly independent set of $a+b$ characters. To begin, we may cite a result of Winstone regarding abelianizations of Bieri-Strebel groups.

Citation 5.2.1. *([Win], Theorem 4.8.11) For a quadratic bieri-Strebel group $\mathbf{F}_\beta$ with subdivision polynomial $2nx^2 + (2n+1) - 1$, we have that $\mathbf{F}_{\beta_{ab}} \cong \mathbb{Z}^{4n+1} \oplus \mathbb{Z}/(n+1)\mathbb{Z}$.*

This indicates that there are quadratic Bieri-Strebel groups where $r_0(\mathbf{F}_\beta) \cong \mathbb{Z}^{a+b}$. For these groups, we have a spanning set for $Hom(\mathbf{F}_\beta, \mathbb{R})$. We wish to generalise this result, and it is possible as, unlike Winstone, we are not interested in calculating the full abelianization, only the torsion free part. We wish to use Winstone's presentation 2.3.7to determine $r_0(\mathbf{F}_\beta)$.

Instead of working within the abelianization, we wish to instead apply a generic character χ to all elements, then analyse the group relations in order to determine the size of r_0. The advantage of this approach is that, as $\chi(f) \in \mathbb{R}\,\forall f \in \mathbf{F}_\beta$, we may apply the tools of linear algebra when examining the relations. For example, the group $\mathbf{F}_\tau$ has the relation $y_0^2 = x_0x_1$. Taking an arbitrary character of both sides of this relation gives us $2\chi(y_0) = \chi(x_0) + \chi(x_1)$. As both sides of the equation are just real numbers, we can divide through by 2 to determine that $\chi(y_0) = \frac{\chi(x_0)+\chi(x_1)}{2}$. As such, we know that the value of $\chi(y_0)$ can always be determined from the values of $\chi(x_0)$ and $\chi(x_1)$. Therefore, we do not need y_0 to generate $r_0(\mathbf{F}_\tau)$.

Examining Winstone's relations from 2.3.7, there are two types of relations. We first examine the R_1 relations of the form

$$f_jg_i = g_if_{j+a+b-1}\, f, g \in x, z, j > i$$

Applying an arbitrary character to both sides of the relation will get us $\chi_(f_j) + \chi(g_i) = \chi(g_i) + \chi(f_{j+a+b-1})$. Subtracting $\chi(g_i)$ from both sides will leave us with the relation $\chi_(f_j) = \chi(f_{j+a+b-1})$. As we require $j > i$ $i, j \geq 0$, this implies that for $k \geq a+b$, there is a $l \in 1, ..., a+b-1$ such that $\chi(x_k) = \chi(x_l)$ and $\chi(z_k) = \chi(z_l)$. As such, the character values of all elements of F_β can be determined wholly from the character values of the elements $\{x_0, z_0, x_1, z_1, ..., x_{a+b-1}, z_{a+b-1}\}$. From this, we know that $r_0(\mathbf{F}_\beta) \cong \mathbb{Z}^k$, where $k \leq 2(a+b)$.

We now examine the R_2 relations. While the infinite presentation of the group has infinite R_2 relations, we may reduce this down to those concerning the set of elements $\{x_0, z_0, x_1, z_1, ..., x_{a+b-1}, z_{a+b-1}\}$. For example, the relation $y_1^2 = x_1x_2$ is an R_2 relation in $\mathbf{F}_\tau$. Taking the character value of both sides we get $2\chi(y_1) = \chi(x_1) + \chi(x_2)$. However, we know that $\chi_(x_1) = \chi(x_2)$ from the R_1 relation $x_1x_0 = x_0x_2$. Hence we can rewrite the R_2 relation as $2\chi(y_1) = 2\chi(x_1)$, eliminating the x_2 term.

Generalising this to $\mathbf{F}_\beta$ gives us $a+b$ unique R_2 relations with elements from the set $\{x_0, z_0, x_1, z_1, ..., x_{a+b-1}, z_{a+b-1}\}$. We write these relations out as

$$\begin{gathered}\chi(x_a)+\chi(x_{a+1})+...+\chi(x_{2a})+\chi(x_0)=2\chi(z_0)+\chi(z_1)+...+\chi(z_{a-1})\\ \chi(x_{a+1})+\chi(x_{a+2})+...+\chi(x_{2a+1})+\chi(x_1)=2\chi(z_1)+\chi(z_2)+...+\chi(z_a)\\ ...\\ \chi(x_{a+b-1})+\chi(x_1)+...+\chi(x_a)+\chi(x_{b-1})=2\chi(z_{b-1})+\chi(z_{b+2})+...+\chi(z_1)\\ ...\\ \chi(x_{a-1})+\chi(x_a)+...+\chi(x_{2a-1})+\chi(x_{a+b-1})=2\chi(z_{a+b-1})+\chi(z_1)+...+\chi(z_{a-1})\end{gathered} \tag{5.4}$$

This leaves us with a system of linear equations, which we may approach using linear algebra. In particular, as all terms involving x_i generators are on the left of the equations, and all terms involving z_i are on the right, we may form the vectors

$$\underline{x}=\begin{pmatrix}\chi(x_0)\\ \chi(x_1)\\ ...\\ \chi(x_{a+b-1})\end{pmatrix}\ \underline{z}=\begin{pmatrix}\chi(z_0)\\ \chi(z_1)\\ ...\\ \chi(z_{a+b-1})\end{pmatrix}$$

which we can then use to write the system of equations as the matrix equation

$$A\underline{x}=B\underline{z}$$

If either A or B is invertible, then we may solve this matrix equation to determine each $\chi(x_i)$ solely in terms of $\chi(z_j)$, or each $\chi(z_i)$ solely in terms of $\chi(x_j)$. This would reduce the generating set for $r_0(\mathbf{F}_\beta)$ such that $r_0(\mathbf{F}_\beta)\cong\mathbb{Z}^{a+b}$. Based on the equations in 5.4, we can form the matrix B as the following $(a+b)\times(a+b)$ matrix

$$B=\begin{pmatrix}2&1&\dots&1&0&0&\dots&0\\ 0&2&1&\dots&1&0&\dots&0\\ \vdots&&&&&&&\vdots\\ 0&0&\dots&0&2&1&\dots&1\\ 0&1&0&\dots&2&1&\dots&1\\ \vdots&&&&&&&\vdots\\ 0&1&\dots&1&0&\dots&0&2\end{pmatrix}$$

Where the number of entries equal to 1 in each row is equal to $a-1$. We are primarily concerned with whether B is invertible, and as such, whether $det(B)\neq 0$. We may calculate the determinant of B down the first column, which is empty except from the 2 in the first row. As such, the determinant of B is just $2*det(B')$, where B' is the following $(a+b-1)\times(a+b-1)$ submatrix of B

$$B' = \begin{pmatrix} 2 & 1 & \dots & 1 & 0 & 0 & \dots & 0 \\ 0 & 2 & 1 & \dots & 1 & 0 & \dots & 0 \\ \vdots & & & & & & & \vdots \\ 0 & \dots & 0 & 2 & 1 & 1 & \dots & 1 \\ 1 & 0 & 0 & \dots & 2 & 1 & \dots & 1 \\ \vdots & & & & & & & \vdots \\ 1 & \dots & 1 & 0 & 0 & \dots & 0 & 2 \end{pmatrix}$$

In order for B' to be invertible, we require its rank to be $a + b - 1$. Fortunately, B' fits into a classification of matrices that have a known method for determining their ranks.

Definition 5.2.2. An $n \times n$ matrix A is a circulant matrix if the ith row is the first row with the entries transposed $i - 1$ columns to the right. In particular, for a set of values $\{c_0, ..., c_{n-1}\}$, a circulant matrix is a matrix of the form

$$\begin{pmatrix} c_0 & c_{n-1} & \dots & c_2 & c_1 \\ c_1 & c_1 & c_{n-1} & & c_2 \\ \vdots & c_1 & c_0 & \ddots & \vdots \\ c_{n-2} & & \ddots & \ddots & c_{n-1} \\ c_{n-1} & c_{n-2} & \dots & c_1 & c_0 \end{pmatrix}$$

For the circulant matrix A with entries $\{c_0, ..., c_{n-1}\}$, we define the polynomial $f_A(x) = c_0 + c_1 x + ... + c_{n-1} x^{n-1}$.

Citation 5.2.3. *([Ing56], Theorem 2.1) The rank of an $n \times n$ circulant matrix A with entries $\{c_0, ..., c_{n-1}\}$ is equal to $n - d$ where d is the degree of the polynomial $f_A(x), x^n - 1\}$. In particular, if these two polynomials are coprime, than A is invertible.*

The factors of $x^n - 1$ are all cyclotonic polynomials so, in order to show B' is invertible, we wish to show $f_{B'}(x)$ has no cyclotonic factors. There are two very simple cases we can show immediately. For $\mathbf{F}_\beta$ with subdivision polynomial $ax^2 + bx - 1$, if $a = 1$, then B' is just $2I$ for I the $(a + b) \times (a + b)$ identity matrix. This is clearly invertible, and checking via the polynomial confirms this as $f_{2I} = 2$, which clearly has no cyclotonic factors. We may perform a similar analysis for $a = 2$, for which $f_{B'}(x) = x^{n-1} + 2$. The roots of this are clearly the $n - 1$th roots of -2. As such, $x^{n-1} + 2$ shares no roots with $x^n - 1$, so cannot share any factors.

We can generalise this one more time to $a = 3$. For an $\mathbf{F}_\beta$ with $a = 3$, we have that $f_{B'} = x^{n-1} + x^{n-2} + 2$. We wish to show this is coprime with $x^n - 1$. The roots of $x^n - 1$ are all roots of unity, which can all be considered as complex numbers v with $|v| = 1$. For any v to be a root of $x^{n-1} + x^{n-2} + 2$, we must have that $v^{n-1} + v^{n-2} = -2$. As we know $|v| = 1$, this is only possible if $v^{n-1} = v^{n-2} = -1$. However, this immediately creates a contradiction as we get $v^{n-1} = v * v^{n-2} = v^{n-2}$ which implies $v = 1$. But if $v = 1$ then $v^k = 1$ for all k, and thus $v^{n-1} \neq 1, v^{n-2} \neq 1$. Hence no root of unity v is a root of $x^{n-1} + x^{n-2} + 2$ and thus $x^{n-1} + x^{n-2} + 2$ is coprime with $x^n - 1$.

Finally, there is another case we can consider, which is the case when $a + b - 1$ is prime. First, we note that $(x - 1)$ is never a factor of $f_{B'}$. All coefficients in $f_{B'}$ are positive, so $f_{B'}(1) > 0$ (in fact, $f_{B'}(1) \geq 2$, as the constant coefficient is always 2). As such, we may consider the polynomial

$$\frac{x^n - 1}{x - 1} = x^{n-1} + x^{n-2} + ... + x + 1$$

When $n-1$ is prime, then $\frac{x^n-1}{x-1}$ is a cyclotomic polynomial, and therefore irreducible. If $\frac{x^n-1}{x-1}$ is irreducible, then either $f_{B'} = \frac{x^n-1}{x-1}$, or $f_{B'}$ and $\frac{x^n-1}{x-1}$ are coprime, but we can see $f_{B'} \neq \frac{x^n-1}{x-1}$ as $f_{B'}$ always has a constant coefficient of 2 and $\frac{x^n-1}{x-1}$ always has a constant coefficient of 1. Therefore, $\frac{x^n-1}{x-1}$ and $f_{B'}$ must be coprime.

Combining the above with 5.2.1 gives us the following lemma:

Lemma 5.2.4. *For F_β a quadratic Bieri-Strebel group with subdivision polynomial $ax^2 + bx - 1$. IF we have one of the following criteria:*

- $a = 2n, b = 2n + 1$ *for some* n
- $a = 1, 2$ *or* 3
- $a + b - 1$ *is a prime number*

Then $r_0(F_\beta) = a + b$.

This combines with our knowledge of characters for $\mathbf{F}_\beta$ to grant us a spanning set for $Hom(\mathbf{F}_\beta, \mathbb{R})$.

Lemma 5.2.5. *Take F_β as in 5.2.4, with $\{X_0, ..., X_{a+b-1}\}$ the set of orbits of F_β on break-points of $(0, 1)$. Then the set of characters $\{\chi_0, \chi_1, \psi_0, ..., \psi_{a+b-2}\}$ is a linearly independent spanning set in $Hom(F_\beta, \mathbb{R})$, where the characters are defined as follows:*

- $\chi_0(f) = log_\beta(f'(0)), \chi_1(f) = log_\beta(f'(1))$
- $\psi_i(f) = \Sigma_{x \in X_i} \gamma_x(f)$, *where* $\gamma_x(f) = f'_{\geq}(x) - f'_{\leq}(x)$

5.2.2 Preliminary Results

The remainder of this chapter will be dedicated to calculating the BNSR invariant for quadratic Bieri-Strebel groups $\mathbf{F}_\beta$ with subdivision polynomial of the form $x^2 + ax - 1$. In order to calculate the BNSR invariant for $\mathbf{F}_\beta$, we will use the complex $\tilde{X}'$ discussed in 3.2. We discussed the shape of links for this space in 3.2.6. In simple terms, the descending link of a forest tree pair $[F, T]$ is a simplicial flag complex with a vertex for each distinct set of k consecutive trees in F, and any 2 vertices have an edge between them if there if their associated sets of trees have trivial intersection. As before, the forest F for any forest-tree pair $[F, T]$ in X will be composed entirely of x-carets, or carets that have all legs of length 2 on the left of the caret.

In order to apply Bestvina-Brady Morse theory to the rest of our calculation, we must understand the links of a 0-cell in our complex. From our definition of the complex in 3.1.2, we know that any 1-cell adjacent to a 0-cell represents either a basic expansion or basic contraction of the forest-pair associated with the 0-cell. When we define our Morse functions according to 4.1.2, we will use either $t(x)$ or $-t(x)$ as our secondary function, where $t(x)$ counts the number of trees in the forest of the forest tree pair associated to the 0-cell x, and is affinely extended to all other cells. As such, it is helpful to consider the ascending and descending links with regards to these functions. The ascending link with regards to $t(x)$ consists of all basic expansions of x, and the descending link consists of all basic contractions. When using $-t(x)$, these links are identical, but swapped.

Considering the ascending link of x first, we can refer to 2.2.6 to conclude that, as the cell above a 0-cell x is an n-cube for $n = t(x)$. From the definition of a link, we can say that the ascending link is the intersection of this cube with a sphere of radius $\epsilon < 1$ around x, which is an $n - 1$ simplex, and thus contractible.

We now consider the descending link of x. The structure of our space means that this ends up being similar to the matchings set up by Zaremsky in ([Zar17], Chapter 5). As such, our connectivity argument is similar to ([Zar17], Lemma 5.2).

Lemma 5.2.6. *The descending link of a point x with $2(n+1)k - n$ trees in the forest of the associated forest tree pair is n-connected.*

Proof. From 3.1.12, we know that a set of simple contractions in the descending link of a point x are connected if the contractions themselves are disjoint. We argue that a descending link for a point x with $3k - 1$ trees is connected in the following way. We consider the leftmost set of k consecutive trees and the rightmost set of k consecutive trees. Each of these sets has a single contraction associated with it. As the number of trees in the forest is greater than $2k - 1$, these two sets of trees are disjoint. Thus, the contraction across the leftmost set of trees is adjacent to the contraction on the rightmost set of trees.

Suppose now that there exists a set of consecutive trees that intersects with both the leftmost set of trees and the rightmost set. As our forest is $3k - 1$ trees, there are $k - 1$ trees in the centre of the forest that do not belong to either the leftmost set of trees or the rightmost set of trees. Any set of consecutive trees that intersects with both the leftmost set and the rightmost set would have to contain all of these trees, along with at least one tree from the leftmost set and one tree from the rightmost set (otherwise it wouldn't intersect with those sets). That means such a set must contain at least $k + 1$ trees. As such there are no sets of k consecutive trees that intersects with both the leftmost and rightmost set. Put another way, every set of k consecutive trees is disjoint from either the leftmost set or the rightmost set of k consecutive trees. Because of this, the contraction over a given set of k consecutive trees is adjacent to the contraction over the leftmost set, or it is adjacent to the contraction over the rightmost set.

We can now demonstrate the descending link is connected in the following way. For any two points in the descending link (representing two contractions of sets of k consecutive trees), we know that each of those two points must be adjacent to either both contractions of the leftmost set, or both contractions of the rightmost set. We also know that each contraction of the leftmost set is adjacent to both contractions of the rightmost set. As such, there are three possibilities for a path between two given points A and B in the descending link.

- Both A and B are adjacent to the contraction L of the leftmost set: the path is $A \to L \to B$.
- Both A and B are adjacent to the contraction R of the rightmost set: the path is $A \to R \to B$.
- one of A and B is adjacent to L and one is adjacent to R: WLOG, assume A is adjacent to L. The path is $A \to L \to R \to B$.

From here we can generalise by referring to ([Zar17], Lemma 5.2). Having demonstrated that our descending links have the same structure as Zaremsky's matching complexes, we may cite his result on connectivity. This gives us that the descending link of a forest with at least $(3k-1)+(n-1)(2k-1)$ trees is n-connected. We can simplify this to $(2n+1)k-n$. □

5.2.3 χ_0 and χ_1

We first consider the characters $-\chi_0$ and $-\chi_1$, where $\chi_0(g) = log_\tau(g'(0))$ and $\chi_1(g) = log_\tau(g'(1))$. This section will be dedicated to proving the following lemma:

Lemma 5.2.7. $[-\chi_0], [-\chi_1] \notin \Sigma^1(\mathbf{F}_\beta)$.

Note that $\tau < 1$, $[-\chi_0]$ contains the character $d_0(g) = g'(0)$ and $[-\chi_1]$ contains $d_1(g) = g'(1)$. In order to use these characters on the complex X (or any similar complexes), we must introduce height functions that are equivariant to these characters. This is relatively straightforward. When working with tree pairs, the character d_0 is proportional to the difference between the depth of the leftmost leaf of the left tree and the leftmost leaf of the right tree. Similarly, d_1 is proportional to the depth of the rightmost leaf of the left tree and the rightmost leaf of the right tree. We can extend this to the set of forest-tree pairs by applying the exact same logic, subtracting the depth of the leftmost leaf of the tree from the depth of the leftmost leaf of the forest will get us a height function equivariant with a character in $[-\chi_0]$, and subtracting the depth of the rightmost leaf of the tree from the depth of the rightmost leaf of the forest will get us a height function equivariant with a character in $[-\chi_1]$. We label these height functions h_0 and h_1 respectively.

We wish to work within a finite index subgroup of our group $\mathbf{F}_\beta$ and then apply 1.4.1 in order to calculate the invariant for $\mathbf{F}_\beta$. We consider the presentation given in 2.3.7 when $a = 1$. In particular, the presentation reduces to contain two kinds of relations.

- $b_j a_i = a_i b_{j+k-1}$, for $a, b \in \{x, y\}, i < j$
- $y_i y_{i+1} = x + i^2$

This observation also appears in ([Bro18], section 6.1). As such, we have $K_\beta \subseteq \mathbf{F}_\beta$ such that $K = \langle x_0, x_1, ..., y_1, y_2 ...\rangle$, as in 5.1.2. We may consider the space X_K as the maximal subspace of X consisting of vertices whose forest-tree pairs may be obtained from elements of K via a sequence of basic expansions (and forest-tree equivalences). In practice, this gives us forest-tree pairs with no y-carets on the leftmost branch of either the forest or the tree. Based on the normal form presented for $\mathbf{F}_\beta$ in ([Bro18], Section 7.2), we may construct elements of $\mathbf{F}_\beta$ (and therefore elements of K) such that there are no right handed carets in the left tree of the tree pair. Any expansions of these normal form trees will therefore only have left handed carets in the forest. For each equivalence class of forest-tree pairs, we will choose a reduced representative that has the following properties:

- The reduced representation will only have left-handed carets in its forest.
- The reduced representative will have a minimal number of carets with regards to the first property.

We can reduce the characters we need to consider by using μ-symmetry as in 5.1.3. As there is an automorphism of each $\mathbf{F}_\beta$ that maps $-\chi_0$ to $-\chi_1$ (and vice versa), $-\chi_0$

and $-\chi_1$ will have the same placement in the BNSR invariant. Thus, we only need to calculate the position in the invariant for $-\chi_0$.

We will now consider the space $Y_K = X_K^{2k-1 \leq t(x) \leq 3k-2}$, with the Morse function (h_0, t). Under this function, the ascending link will consist of any basic expansion on any tree in the forest except the leftmost (though only on vertices x such that $t(x) = 2k-1$. When $t(x) = 3k-2$, a basic expansion will lead to an x' where $t(x') = 4k-3$, and thus x' is not in Y_K), along with basic contractions on the leftmost set of k trees (which are only in Y_K if $t(x) = 3k-2$). On any point $x \in Y_K$, there is either a basic expansion or a basic contraction in the ascending link. Thus the ascending link is nonempty.

We now wish to show $Y_k^{h_0>0}$. For this we wish to follow the process established by Witzel and Zaremsky in ([WZ15], Section 5). As such, we will need to make use of the following result.

Citation 5.2.8. *([WZ15], Proposition 1.8) Let an $(m-1)$-connected affine cell complex X be equipped with a Morse function $(h,s) : X \to \mathbb{R} \times \mathbb{R}$ and assume that all ascending links are $(m-2)$-connected. Then the filtration $\{X^{h \geq t}\}_{t \in \mathbb{R}}$ is essentially $(m-1)$-connected if and only if $X^{h \geq p}$ is $(m-1)$-connected for some p, if and only if all $x^{h \geq p'}$ are each $(m-1)$-connected for all $p' leq p$.*

As we have already established that the ascending links in Y_K are nonempty, or -1-connected, we may use this lemma to determine that the complex $Y_K^{h_0>0}$ is not essentially connected if and only if $Y_k^{h_0>l}$ is not connected for some l. As the action of $\mathbf{F}_\beta$ is transitive, this reduces to showing that $Y_K^{h_0>0}$ is not connected.

Take a vertex y in $Y_K^{-\chi_0>0}$, and define $L(y)$ as the depth of the leftmost leaf of the tree of y, when y is considered by its reduced representative. Since we are in K, all carets on the leftmost branch of the tree of y are left-handed, and thus this is just 2 times the number of carets on this branch. Now consider a vertex y' adjacent to y. We wish to show that is y is adjacent to y' in $Y_K^{h_0>0}$, then $L(y) = L(y')$.

Firstly, in any case that does not expand or contract the leftmost tree of the forest of Y, no redundant carets will be introduced on the leftmost branch of Y either before or after the basic expansion or contraction, so we only need to consider a basic expansion or contraction of the leftmost tree. Suppose we perform a basic expansion. The only reason we would need to introduce a redundant caret is if the leftmost tree of the forest is trivial. However, as we are in $Y_K^{h>0}$, we know that the depth of the leftmost leaf of the forest must be greater than the depth of the leftmost leaf of the tree. We also know that the tree $y \in Y_K^{h_0>0}$ cannot be trivial, as $t \geq 2k-1$. Thus the depth of the leftmost leaf of the tree must be greater than 0, so the depth of the leftmost leaf of the forest must be greater than zero, so the leftmost tree of the forest cannot be trivial. Finally, we consider a basic contraction of the k leftmost trees of the forest. Again, this cannot produce a redundant caret. Each reduced representative only contains left handed carets in the forest, so we can only contract along a left caret. Thus the leftmost caret cannot be part of a hidden cancellation (as described in ([BNR21], section 7). Thus, if there was a possible cancellation that would imply y was not in reduced representative form, as was previously assumed. With no redundant carets to cancel, we can conclude that $L(y) = L(y')$ for adjacent vertices $y, y' \in Y_K^{h_0>0}$. As there are clearly multiple values $L(y)$ can take in $Y_K^{h_0>0}$, there must be multiple connected components in $Y_K^{h_0>0}$, thus it is not connected.

We can now apply 5.2.8 to conclude that $Y_K^{h_0>0}$ is not essentially connected. As such, $[-\chi_0] \notin \Sigma^1(K_\beta)$. We can use 1.3.5 to conclude that $[-\chi_0] \notin \Sigma^1(\mathbf{F}_\beta)$ and then use μ-symmetry to conclude $[-\chi_1] \notin \Sigma^1(\mathbf{F}_\beta)$.

5.2.4 The Long Interval

We will now apply Morse Theory to determine where the characters of the form $\chi = \alpha\chi_0 + \gamma\chi_1; \alpha > 0$ or $\gamma > 0$ sit in the BNSR invariant. In particular, we wish to prove the following lemma.

Lemma 5.2.9. *For χ of the form $\chi = \alpha\chi_0 + \gamma\chi_1; \alpha > 0$ or $\gamma > 0$, $\chi \in \Sigma^n(\boldsymbol{F}_\beta)\forall n$*

We will proceed in a manner similar to Witzel and Zaremsky in ([WZ15], section 4). We begin by reducing the cases we need to consider. We do this by using μ-symmetry (introduced in 5.1.3), allowing us to assume $\alpha > \gamma$, implying $\alpha > 0$. We then form the Morse function $(\chi'(x), -t(x))$, where $\chi'(x)$ is the height function on the complex X equivariant with the character χ. By the definition of BNSR invariants provided in 1.3.7, we need to find a complex X_n such that the top half $X_n^{\chi' \geq 0}$ is essentially n-connected. We shall select $X_n = X_{2 \leq t \leq m}$. By 3.1.14, we know there is an m such that X_m is n connected, and by 3.1.11, the space is cocompact with regards to the action of G. Thus, this is a suitable complex for calculating the n-th BNSR invariant. By 4.2.2, χ being in $\Sigma^n(G)$ is equivalent to the following lemma.

Lemma 5.2.10. *Let x be a vertex in X_n, then the ascending link $lk\uparrow(x)_{X_n}^{\chi'(x),-t(x)}$ is $n-1$ connected*

Proof. We first consider the 0-cells of $lk\uparrow(x)_{X_n}^{\chi'(x),-t(x)}$. These 0-cells can be created by both simple expansions and simple contractions in the forest of x. An expansion must strictly increase χ', as any expansion increases t and thus decreases $-t$. A contraction will be in the ascending link as long as it does not explicitly decrease χ.

Now, suppose $t(x) = m$. Any expansions of x would lead to a vertex x' such that $t(x') = m+1$, thus $x' \notin X_n$ and so the ascending link of x must consist only of basic contractions. Since we have $\alpha > 0$, we know that any merge including the leftmost tree of the forest of x will increase the slope at 0 and thus decrease χ' (recall $\chi_0 = (log_\beta(g'(0))$ for $\beta < 1$). Depending on b, a merge including the rightmost tree of the forest of x may increase, decrease or fix χ'. All other merges cannot change χ' but decrease t and thus will be in the ascending link. Thus the ascending link will consist of all contractions on the forest of $m-1$ or $m-1$ trees depending on b. In either case, for sufficiently high m, this will be n connected by 5.2.6.

We now consider the case where $t(x) < m-(k-1)$. The only possible splits in this ascending link is the split of the leftmost tree and potentially the split of the rightmost tree (all other splits will not effect χ' but will increase t). While every merge of k trees is included in the ascending link apart from the merge of the k leftmost trees (which would decrease χ') and possibly the merge of the k rightmost trees (which will increase, decrease or fix χ' depending on γ. As the merge of the k rightmost trees is not included, all other expansions or contractions in the ascending link are disjoint from the expansion of the leftmost tree, thus the vertices implied by all of these basic moves in the ascending link are all adjacent to the vertex implied by the expansion of the leftmost tree. Hence the vertex implied by the expansion of the leftmost tree is a cone point for the ascending link. Thus the ascending link is contractible and therefore n-connected.

Finally, suppose $m-k < t(x) < m$. At first glance, this case would be similar to the $t(x) < m-(k-1)$ case, using the expansion of the leftmost tree as a cone point. The issue is when $\gamma > 0$, which would cause the expansion of the rightmost tree to imply a vertex in the ascending link. In the previous case, this is fine as the edge between these two vertices would exist by the existence of the point $x' \in X_n$ reached by performing both basic expansions in either order. However, if $t(x) = m-(k-1)$, then $t(x') > m$, and therefore $x' \notin X_n$. However, the edge between these two vertices is the only edge that can't be included. So, we consider the contractible space L' where we cone off the space using the vertex from the expansion of the leftmost tree, then form the ascending link of x from L' by removing this edge along its relative link. AS this relative link is still at least n-connected, so is the ascending link of x. □

5.2.5 The Short Interval

Our next step is to determine the position of characters of the form $\chi = \alpha\chi_0 + \gamma\chi_1, ; \alpha, \beta < 0$. Our goal is to prove the following lemma.

Lemma 5.2.11. *for characters of the form $\chi = \alpha\chi_0 + \gamma\chi_1, ; \alpha, \gamma < 0$, $\chi \notin \Sigma^2(\boldsymbol{F}_\beta)$*

In order to do this, we will need to introduce the concept of a nerve complex.

Definition 5.2.12. ([ES52],page 234). Given an open cover $C := \{U_i | i \in I\}$ of a topological space X we construct the nerve complex $N(C)$ by first taking the set of 0-cells $\{U_i\}_{i\in I}$. We then insert an n-cell between the 0-cells $\{U_0, U_1, ..., U_n\}$ if the intersection $U_0 \bigcap U_1 \bigcap ... \bigcap U_n$ is nonempty.

We will then use the following result of Witzel and Zaremsky in order to replicate their method in calculating $\Sigma^2(G)$

Citation 5.2.13. *([WZ15], Lemma 6.2) Let a simplicial complex X be covered by connected subcomplexes $\{X_i\}_{i\in I}$. Suppose the nerve $N(\{X_i\}_{i\in I})$ is connected, but not simply connected. Then X is connected but not simply connected.*

In order to demonstrate that χ is not in $\Sigma^2(G)$, we will work within the subgroup K_β once again. As a brief reminder, this is the subgroup generated by the usual generating set of $\mathbf{F}_\beta$ (as described in 2.3.7) but without the y_0 generator. This will result in no right handed carets on the leftmost branch of either tree. We may use Jason Brown's normal form ([Bro18], section 7.2). to find a representative for each element in K (here considered an equivalence class of tree-pair diagrams) such that the left hand tree has no right hand carets either. As all vertices in X_K are associated with a forest-tree pair constructed by performing a finite sequence of basic expansions on a tree pair, then all forest-tree pairs in X_K will have only left-handed carets in the forest.

Take the height function h equivariant to χ defined as before. We now introduce the Morse function (h, t) and restrict to $Y_K = X_K^{3k-2\leq t\leq 6k-5}$. We first wish to show that all ascending links are connected. This will allow us to conclude that Y_K is connected via the Morse lemma. We first examine our character. As $\chi = a\chi_0 + b\chi_1$ and $a, b < 0$, we know that a merge of either the leftmost k trees or the rightmost k trees will increase the depth of the leftmost leaf of the forest or the rightmost leaf respectively. This will then increase h, and thus these merges are in the ascending link. No other merges can effect h and reduce t, and thus must be excluded. A split of the leftmost or rightmost tree would reduce χ and thus is excluded, but any other split fixes h and increases t, so must be included.

We now note that Y_K contains 4 "layers" of vertices, as the number of trees in a forest of a forest tree pair must be of the form $nk - (n-1)$. In, Y_K, this leaves the possibilities $\{3k-2, 4k-3, 5k-4, 6k-5\}$. We consider the smaller two first. as there are two layers above these, we can perform two basic expansions and remain in Y_K. Any two basic expansions are disjoint and thus adjacent in the ascending link. Furthermore, each of the two included basic contractions (second layer only, as a basic contraction in the lowest layer would remove us from Y_K) is disjoint to at least one basic expansion, as there are at least $k+2$ trees in each forest for $k \geq 2$. For the top two layers, we note that the contraction along the leftmost k trees and the contraction along the rightmost k trees are disjoint, as $3k - 2 \geq 2k$ for all $k \geq 2$. These contractions are therefore adjacent in the ascending link. Each expansion is disjoint from at least one of these contractions and so adjacent to it in the ascending link. Thus the ascending link is connected.

We wish to apply the result of Witzel and Zaremsky to this complex. In order to do so, we will split the complex into subcomplexes. For each 0-cell in Y, we shall take the reduced representative of the forest-tree pair, by which we mean the equivalent forest-tree pair with the smallest number of carets. On these representatives we will define the metrics $L(x)$ and $R(x)$. $L(x)$ is the depth of the leftmost leaf of the tree in the forest-tree pair associated with x (as used in 5.2.3, while $R(x)$ is the depth of the rightmost leaf of the tree. We will then take the complex and divide it into subcomplexes $Y_{F(x)=p}^{t(x)=q}$ for $F \in \{L, R\}$. For example, $Y_{L(x)=3}^{t(x)=6}$ would consist of all forest-tree pairs x such that the leftmost leaf of the tree is of depth 3 and the forest contains 6 trees.

We now restrict to $Y_K^{h(x)\geq 0}$. We can immediately conclude that this subcomplex is connected as all ascending links inf Y_K are connected. If a point x' adjacent to x appears in the ascending link of x, then $h(x') \geq h(x)$, thus $x \in Y_K^{h(x)\geq 0} \Rightarrow x' \in Y_K^{h(x)\geq 0}$. Thus the ascending links remain connected and the Morse lemma allows us to conclude $Y_K^{h(x)\geq 0}$ is connected. Our goal is thus to show it is not essentially simply connected. By 5.2.8, we only need to show that $Y_K^{h(x)\geq 0}$ is not simply connected.

We consider the nerve of $Y_K^{h(x)\geq 0}$, induced by our subcomplexes $Y_{F(x)=p}^{t(x)=q}$. We claim that each cell $c \in Y_K^{\chi(x)\geq 0}$ is contained within either $Y_{L(x)=i}$ or $Y_{R(x)=i}$ for some i. Take $x \in c$ as a vertex, and consider its forest-tree pair. As $t(x) > 1$, we know the tree of x must be non-trivial. Therefore the depth of the leftmost and rightmost leaves of this tree are both greater than zero. We know $h(x) > 0$, so the depth of either the leftmost or rightmost leaf of the forest must be greater than 0 (potentially both). In either case, we can make an argument similar to that made in 5.2.3 that $L(x') = L(x)$ or $R(x') = R(x)$ for any vertex x' adjacent to x. Thus, the cell c must be in either $Y_{L(x)=i}$ or $Y_{R(x)=i}$. We further note that $Y_{L(x)=i} \bigcap Y_{L(x)=j} = 0$ and $Y_{R(x)=i} \bigcap Y_{R(x)=j} = 0$ whenever $i \neq j$, as that would imply a leftmost or rightmost leaf having two leaf depths simultaneously, which is obviously not possible.

In order to apply 5.2.13 and demonstrate that $Y_K^{h(x)\geq 0}$ is not simply connected, we simply need to demonstrate the existence of a nontrivial loop in the nerve of $Y_K^{h(x)\geq 0}$. To do this, we may conceive of the following four forest-tree pairs.

- x_1, such that $L(x_1) = 4$ and $R(x_1) = 2$.
- x_2, such that $L(x_2) = 4$ and $R(x_2) = 3$.

- x_3, such that $L(x_3) = 6$ and $R(x_3) = 3$.
- x_4, such that $L(x_4) = 6$ and $R(x_4) = 2$.

As the depth values on the leftmost leaves are all even, each of these are constructable using only left-handed carets, and so are clearly in Y_K, and as long as the tree in the forest-tree pair has minimal depth on the leftmost and rightmost leaves (which is 2 in either case), each of these forest-tree pairs will be in $Y_K^{\chi(x)\geq 0}$. Examining the nerve, we can see that $Y_{L(x)=4}$ must be adjacent to $Y_{R(x)=2}$ due to the existence of x_1. Similarly, we have $Y_{L(x)=4}$ adjacent to $Y_{R(x)=3}$, $Y_{L(x)=6}$ adjacent to $Y_{R(x)=3}$ and $Y_{L(x)=6}$ adjacent to $Y_{R(x)=2}$. However, as previously discussed, we have that $Y_{L(x)=4}$ cannot be adjacent to $Y_{L(x)=6}$ nor can $Y_{R(x)=2}$ be equal to $Y_{R(x)=3}$. As such, this creates the nontrivial loop $Y_{L(x)=4} \to Y_{R(x)=3} \to Y_{L(x)=6} \to Y_{R(x)=2} \to Y_{L(x)=4}$ in the nerve of $Y_K^{h(x)\geq 0}$. We can now apply 5.2.13 in order to conclude $Y_K^{h(x)\geq 0}$ is not simply connected. Hence we can conclude $\chi \notin \Sigma^2(K)$. We can then use 1.4.1 to conclude $\chi \notin \Sigma^2(\mathbf{F}_\beta)$.

5.2.6 All other characters

Finally, we wish to consider other characters that can arise within the character sphere. From 5.2.4, we expect there to be k dimensions of $Hom(G, \mathbb{R})$ for each group G we are considering. Having already discussed χ_0 and χ_1, we consider the other class of characters in $Hom(G, \mathbb{R})$. These are the ψ_i characters discussed in 5.2.1. We consider a character ψ of the form $\psi = \Sigma_{i=2}^{k-1} r_i \psi_i$, $r_i \in \mathbb{R}$.

Lemma 5.2.14. *For a character $\psi = \Sigma_{i=2}^{k-1} r_i \psi_i$, $r_i \in \mathbb{R}$, $\psi \in \Sigma^n(\mathbf{F}_\beta) \forall n$.*

It will benefit us to be able to calculate the value of these characters from tree pair diagrams. Consider that the ith linear piece of an element of $\mathbf{F}_\beta$ is represented by the pair of leaves consisting of the ith leaf from the left tree in the tree pair and the ith leaf from the right tree. As a leaf of depth l represents a segment of length β^l, if the left leaf has depth l_1 and the right leaf has depth l_2, then the linear piece represented by this pair of leaves must have slope $\frac{\beta^{l_2}}{\beta^{l_1}} = \beta^{l_2 - l_1}$. Thus, for slope g, $log_\beta(g) = l_2 - l_1$. As $\beta < 1$, we choose instead to take $log_{\beta^{-1}}(g) = l_1 - l_2$, as now a greater value corresponds to a steeper slope. As mentioned in 5.2.1, the character ψ_i is the sum of the differences between these values at each breakpoint across an orbit of breakpoints, so we will need to incorporate the slopes on either side of a breakpoint. For this, we merely need to take two consecutive pairs of leaves and perform this calculation for each, then subtract one from the other. For two consecutive pairs of leaves l_i and r_i (the ith leaves in the left and right tree respectively) and l_{i+1} and r_{i+1} (the $i+1$th leaves), the difference in log of slope at the breakpoint between these two pieces is just $\delta_i = (l_{i+1} - r_{i+1}) - (l_i - r_i)$. Thus, the character $\psi_i = \Sigma\{\delta_m | 1 \leq m \leq r-2, m \cong i \text{mod } (k-1)\}$.

We can generate height functions $j_i : X \to \mathbb{R}$ by considering our forest-tree pairs as functions $x : [0, t] \to [0, 1]$, where t is the number of trees in the forest of x, with the function x being the function that corresponds to the forest tree pair just as an element in F_β is a function with a corresponding tree pair. We can thus extend our method for calculating the character value of an element via the tree pair onto our space of forest-tree pairs. The important detail to consider is the effect that basic expansions and contractions will have on this character value. We first note that basic expansions and contractions cannot change the tree of the forest-tree pair, except by the addition of redundant carets. Adding redundant carets cannot change the character value of a forest-tree pair. As a confirmation of this, we can observe that adding a redundant caret to the ith leaf of a forest-tree pair will generate k new leaves (at the cost of the original) with the first new leaf being 2 levels deeper than the original and the rest being one level deeper. However, these leaves are being generated on both the forest and the tree, in the same positions. Considering our calculation of j_i subtracts the leaf depth of the ith leaf of the tree from the right leaf of the forest, the net change is always 0. Thus we only need to consider the changes to the forest when considering the effect of basic expansions and contractions.

We can once again use reduced representatives for our equivalence classes of forest-tree pairs to ensure that the forest contains only left-handed carets. For a simple contraction that does not include the leftmost tree or the rightmost tree, we can see that this will increase the depth of the first (maximal) subtree of the new tree by 2, and all others by one, without affecting any other leaves. This means that the difference between the $i-1$th leaf and ith leaf will decrease by 2. As discussed in 2.1.4 and 5.2.1, the addition of any caret to this subtree will add precisely $k-1$ leaves to the forest overall, and so the orbit of the first leaf of the second subtree will remain stable no matter the composition of the first subtree. As the depth of this subtree was increased by one and the first subtree was increased by 2, the difference in depth between the last leaf of the first subtree and the first leaf of the second subtree is increased by 1. As all other subtrees have their depth increased by one, no other depth differences are changed until we get to the last leaf of the new tree. This leaf has had its depth increased by one while the first leaf to the right of the tree is unchanged (as are all other leaves outside the new tree). As such, the difference in depth between these leaves is increased by one.

If we take $l < k-1$ such that $i \cong l \operatorname{mod}(k-1)$, then we have that j_{l-1} is decreased by 2, and j_l is increased by one. We can also conclude that the breakpoint between the last leaf of the new tree and the first leaf to the right of the tree is in the same orbit as the breakpoint between the first and second subtrees of the new tree. If we suppose all the maximal subtree are trivial, then we can easily see that the first leaf is $k-1$ leaves away from the last leaf (as all carets have k leaves). We can then build this single caret tree up to any arbitrary tree containing only left carets by adding each caret in order. We know that the addition of carets does not change the orbit of any leaf (as we are adding a net of $k-1$ leaves each time we add a caret) so no matter the shape of the new tree, the breakpoint between the first subtree and the second will be in the same orbit as the breakpoint between the final leaf of the tree and the first leaf to the right of the tree. We know that the depth difference on each of these breakpoints increases by 1, so overall j_l must increase by 2.

We can use this understanding to consider what happens when the simple contraction contains the first or last tree. If it contains the first, then there is no earlier leaf to have a difference in depth with the first leaf of the new tree, so j_{l_1} does not change. The change in j_l happens as normal. Similarly, if the last tree is included in the contraction, then there is no difference in depth between the last leaf of the new tree and the first leaf to the right of it. However, we still get the difference in depth between the first two subtrees, so we have that j_{l-1} decreases by 2 and j_l increases by 1. As simple expansions are the inverse of simple contractions, we can calculate their effects on the character value as simply the inverse of the effect of a simple expansion. That is to say that a simple expansion of the leftmost tree decreases j_1 by 2, a simple expansion of the rightmost tree increases j_{l_1} by 2 and decreases j_l by 1, and all other basic expansions increase j_{l-1} by 2 and decrease j_l by 2.

All of this allows us to generate a height function j for our character ψ. We can then form Morse functions with these height functions and our function $t(x)$. In particular, we wish to take the Morse function $(j, t(x))$ and apply it to the complex $Y = X^{pk \le t(x) \le pk^2}$ for $p = 4n + 5$, following the method of ([Zar17], chapter 6). In order to apply the Morse lemma and demonstrate connectivity, our goal is to prove the following lemma.

Lemma 5.2.15. *take y a vertex in $Y^{j(y)>0}$. Then the ascending link of y is $(n-1)$-connected.*

Proof. As in ([Zar17] proposition 6.6), we split this into two cases: $p \ge t(y) \ge pk$ and $pk \ge t(y) \le pk^2$, and first consider the case where $p \ge t(y) \ge pk$. The basic expansion of the leftmost tree of y changes j_1, but our character ψ is expressed as a linear combination of $\psi_2, ..., \psi_{k-1}$, and so our height function j is a linear combination of $j_2, ..., j_{k-1}$. As such, j is fixed by the basic expansion of the leftmost tree, and any basic expansion increases t. Thus the expansion of the leftmost tree is always in the ascending link. Similarly, the contraction of the k leftmost trees fixes j but reduces t, so is always in the descending link. As the contraction of the k leftmost trees is the only basic move that is not disjoint with the expansion of the leftmost tree, the expansion of the leftmost tree forms a cone point in the ascending link of y. Thus the ascending link of y is contractible as long as this expansion can be included. This works whenever $p \ge t(y) \ge pk$ as performing all expansions will still remain in Y.

We now consider the case when $pk \le t(y) \le pk^2$. We first claim that performing the maximal number of disjoint basic contractions on y will not move us out of Y. Take q as the number of possible disjoint basic contractions. Each basic contraction is formed over k trees, so $t(y) \ge qk$. If we are able to perform all of these basic contractions concurrently and remain in $Y = X^{pk \le t(x) \le pk^2}$, then for the product y' of all of these contractions, we must have $t(y') \ge p$. Each time we perform a basic contraction, we merge k trees into 1 tree, resulting in a net loss of $k-1$ trees. As such, we can rewrite $t(y')$ as $t(y) - q(k-1)$. Thus, we wish to demonstrate $t(y) - q(k-1) \ge p$. Suppose $q \ge p$. Then we can use the fact that $t(y) \ge qk$ to say $t(y) - q(k-1) \ge qk - q(k-1) \ge q \ge p$. Now suppose $q \le p$. In that case $t(y) - q(k-1) \ge qk - p(k-1) = p$. So in both cases the equation holds.

We wish to construct a large simplex σ constructed from disjoint merges that we can guarantee are in the ascending link. We first exclude the merges on the leftmost and rightmost k trees. As described at the beginning of this section, the effects of all other merges on $j_i(y)$ is that j_l will increase by 2 and j_{l-1} will decrease by 2 for some $1 \le l \le k-1$, with all other j_i fixed by this merge. Also note that this is cyclic, in that there are merges that increase j_1 and decrease j_{k-1}. We write our height function j as $j(y) = 0j_1(y) + c_2 j_2(y) + ... + c_{k-1} j_{n-1}(y)$. Due to the redundancy relation between the basis characters for $\mathbf{F}_\beta$, We may write any height function j with no j_1 term, but we include it here as there are still merges that effect j_1. We now wish to find $1 \le s \le k-1$ such that $c_{i-1} < c_i$, again taken cyclically, in that $c_{k-1} < c_1$ is acceptable. We know that $c_1 = 0$. Either there exists $c_{i-1} < c_i$, or there is a chain $0 = c_1 \ge c_2 \ge ... \ge c_{k-1}$. This implies $c_{k-1} \le 0$. If $c_{k-1} = 0$, then $c_i = 0$ for all c_i and j is the trivial height function corresponding to the trivial character, and we need not consider it. Thus $c_{k-1} < c_1$ so in all cases we have some $c_{s-1} < c_i$. We know that a merge starting on a leaf numbered $l = s \bmod (k-1)$ we will reduce j_{s_1} by 2, increase j_s by 2 and fix all other j_i. As such, one in every $k-1$ merges will increase j overall (as the increase caused by j_s rising is greater than the decrease caused by j_{s-1} falling). However, as each merge is across k trees, these basic contractions are not disjoint, so we can only take every other such merge.

We now need to calculate how many basic contractions σ contains. In order to guarantee that we do not include the first tree in any basic contraction, we will ignore the first $k-1$ possible merges. We then add $0 \leq s-1 \leq n-2$ such that we are on the right cycle of merges. We then take a merge every $2(k-1)$ trees to construct our simplex. So our first merge is across the trees in positions $s+(k-1), s+(k-1)+1, ..., s+2(k-1)$. In order to avoid our basic contractions having an intersection, we then skip a possible merge and perform another merge on $s+3(k-1), s+3(k-1)+1, ..., s+4(k-1)$. We repeat this process until we reach the final possible merge, which we will say is on $s+(v-1)(k-1), s+(v-1)(k-1)+1, ..., s+v(k-1)$. In order for this final merge not to include the rightmost tree, we require that $s+v(k-1) < t-1$, and as we have declared that this is the final possible merge, we have that $s+(v+2)(k-1) \geq t-1$ (otherwise there would be room for a merge across $s+(v+1)(k-1), s+(v+1)(k-1)+1, ..., s+(v+2)(k-1)$). We have that $t \geq pk$ and $s \leq k-2$, so we can say

$$
\begin{aligned}
k-2+(v+2)(k-1) \geq s+(v+2)(k-1) &\geq t-1 \geq pk-1 \\
(v+2)(k-1) &\geq pk-1-k+2 \\
v(k-1)+2(k-1) &\geq (p-1)k+1 \\
v(k-1) &\geq (p-1)k+1-2k+2 \\
v &\geq \frac{(p-3)k+3}{k-1}
\end{aligned}
\tag{5.5}
$$

as we require v to be an integer, we will take the floor and say $v \geq \lfloor \frac{(p-3)k+3}{k-1} \rfloor$. We can then substitute in our value $p = 4n+5$ in order to get that $v \geq \lfloor \frac{(4n+2)k+3}{k-1} \rfloor \geq 4n+2$. As we are taking every other possible merge to construct our simplex, we can say that the number of vertices in our simplex is $\frac{v}{2}$ or just $2n+2$, making it a $2n-1$-simplex.

Finally, we wish to use this simplex to demonstrate the connectivity of the entire ascending link. In order to do this, we will need to borrow another result from Zaremsky.

Definition 5.2.16. ([Zar17], definition 6.4) Let Δ be a simplicial complex. two simplices $\rho_1, rho_2 \in \Delta$ are joinable to each other if there exists a simplex ρ such that $\rho_1, \rho_2 \subseteq \rho$. For a given simplex $\sigma \in \Delta$, we say that Δ is flag with respect to σ if whenever we have a simplex ρ and σ' a face of σ such that every vertex of ρ is joinable to every vertex of σ', then we have that ρ is joinable to σ'.

A useful fact that Zaremsky states in this definition is that a simplicial complex that is flag with respect to all its simplices, then it is flag.

Citation 5.2.17. *([Zar17], Lemma 6.5) Let Δ be a simplicial complex and let $k \in \mathbb{N}$. Suppose there exists an l-simplex σ such that Δ is flag with respect to σ. Further suppose that for all vertices $v \in \Delta$, v is joinable to some $(l-k)$ face of σ. Then Δ is $(\lfloor \frac{l}{k} \rfloor - 1)$ connected.*

In order to apply 5.2.16 to the ascending link of y, we first have to demonstrate that the ascending link is flag with respect to σ. Were we working in X and not Y, we would have that the complex is already flag (as the 1-skeleton of an i simplex would imply i pairwise disjoint basic moves, and if they are pairwise disjoint then they can all be performed sequentially in any order. As such, the i simplex implied by the 1-skeleton exists). We consider a simplex ρ in the ascending link of y joinable to some face of σ, label σ'. All vertices of σ, and by extension σ', are basic contractions. Hence, as long as each y' reached from a move in ρ is still in y, then any y' reached from a move in $\rho * \sigma'$ cannot have $t(y') > pk^2$, as moves in σ' can only reduce t, and cannot have $t(y') < pk$, as performing all disjoint contractions on any vertex y such that $pk \leq t(y) \leq pk^2$ does not cause you to leave Y, as discussed at the beginning of this section. As such, as long as the basic moves of ρ do not leave Y, then the whole simplex of $\rho * \sigma'$ from the ascending link of y in X is in the ascending link of y in Y, and so the ascending link of y is flag with respect to σ.

Our next goal is to show that all vertices v in the ascending link of y are joinable to at least all but 2 of the vertices of σ. This is relatively straightforward. There are two possibilities for vertices in the ascending link of y, they can either be basic expansions of a single tree or basic contractions of a set of k consecutive trees. If v is a basic expansion, then it can only overlap with a single basic contraction in σ, as the basic contractions of σ are all disjoint. Now suppose v is a basic contraction. There is a gap of $k-2$ trees between each basic contraction in σ. As such, it is possible for a basic contraction to contain one tree from one contraction of σ, one tree from a consecutive contraction, and the $k-2$ trees in between, thus intersecting with exactly 2 basic contractions in σ. As such, every vertex in the ascending link of y is adjacent to at least all but two vertices in σ. This implies each vertex is joinable to some $2n-3$ face of σ (recall that σ is a $2n-1$-simplex). We may now apply 5.2.16 to conclude that the ascending link of y is $(\lfloor \frac{2n}{2} \rfloor - 1)$-connected, and as such is $n-1$ connected, as required. □

Finally, we consider characters χ of the form $\chi = a\chi_0 + b\chi_1 + \Sigma_{i=1}^{k-1} c_i\psi_i$, where at least one of a and b and at least one of the c_i are non-zero. Instead of calculating this directly, we instead wish to reduce this case to the case for $\psi = \Sigma_{i=1}^{k-1} c_i\psi_i$.

Lemma 5.2.18. *Suppose that $\psi = \Sigma_{i=2}^{k-1} c_i\psi_i \in \Sigma^\infty(\boldsymbol{F}_\beta)$ for all ψ. Then for any $\chi = a\chi_0 + \psi + b\chi_1$, $\chi \in \Sigma^\infty(\boldsymbol{F}_\beta)$ For any ψ.*

Proof. We begin by restricting χ to $\mathbf{F}_\beta[1]$. As $\mathbf{F}_\beta[1]$ contains no x_0 or y_0 generators, we have that $\chi|_{\mathbf{F}_\beta[1]}$ is just $\psi + b\chi_1$. We then use the μ-symmetry automorphism to map $\chi|_{\mathbf{F}_\beta[1]}$ to χ' such that $\chi' = a\chi_0 + \psi'$ where ψ' is nontrivial (though this automorphism takes us out of $\mathbf{F}_\beta[1]$). We can then restrict χ' to $\mathbf{F}_\beta[1]$, and $\chi'|_{\mathbf{F}_\beta[1]}$ is just ψ'. By assumption, $\psi' \in \Sigma^\infty(\mathbf{F}_\beta[1])$, as it is isomorphic to $\mathbf{F}_\beta$. We can then apply 5.1.7 to show that $\chi' \in \Sigma(K_\beta[1])$, where $\mathbf{F}_\beta[1] \supseteq K_\beta[1] = \langle x_1, x_2, y_2, x_3, y_3 ... \rangle$ is the equivalent subset to K_β inside $\mathbf{F}_\beta[1]$. As $K_\beta[1]$ is finite index in $\mathbf{F}_\beta[1]$, we can apply 1.4.1 in order to show $\chi' \in \Sigma^\infty(\mathbf{F}_\beta)$. We then apply μ-symmetry to show that $\chi|_{\mathbf{F}_\beta[1]}$ is in $\Sigma^\infty(\mathbf{F}_\beta[1])$. Finally, we repeat the previous process again, using 5.1.7 and 1.4.1 to show $\chi \in \Sigma^\infty(\mathbf{F}_\beta)$. □

We can now close this section with the statement of the theorem.

Theorem 5.2.19. *Suppose $\boldsymbol{F}_\beta$ is a Bieri-Strebel group with corresponding subdivision polynomial of the form $x^2 + bx - 1$. Then the BNSR invariant is of the following form.*

- *The character space has a basis of $a+b-1$ independent characters. χ_0 and χ_1 correspond to the log (base β of the slope at 0 and 1 respectively. $\{\psi_i\}_{2\leq i\leq a+b-1}$ are each the sum of difference of the log of the slope to the left and right of each brakepoint, summed over all breakpoints in the ith orbit of breakpoints.*
- *The character classes $[-\chi_0]$ and $[-\chi_1]$ are not in $\Sigma^1(\boldsymbol{F}_\beta)$.*
- *Characters of the form $a\chi_0 + b\chi_1; a, b < 0$ are in $\Sigma^1(\boldsymbol{F}_\beta)$, but not $\Sigma^2(\boldsymbol{F}_\beta)$.*
- *all other characters are in $\Sigma^\infty(\boldsymbol{F}_\beta)$.*

Proof. Point (1) was proven in section 5.2.1, point (2) is the lemma 5.2.7, point (3) is the lemma 5.2.11, and point (4) is reached by combining the lemmas 5.2.9, 5.2.14 and 5.2.18. □

We shall also state a generalisation of this theorem as a conjecture

Conjecture 5.2.20. *5.2.19 holds for all Bieri-Strebel groups with quadratic subdivision polynomial and well-defined tree-pair representation. That is to say, all Bieri-Strebel groups with subdivision polynomial $ax^2 + bx - 1, a \leq b$.*

Bibliography

[BB97] Mladen Bestvina and Noel Brady. "Morse theory and finiteness properties of groups". eng. In: *Inventiones mathematicae* 129.3 (1997), pp. 445–470. ISSN: 0020-9910.

[Bes] "Morse theory and finiteness properties of groups". eng. In: *Inventiones mathematicae* 129.3 (1997), pp. 445–470. ISSN: 0020-9910.

[BG84] Kenneth S. Brown and Ross Geoghegan. "An infinite-dimensional torsion-free FP_∞ group". In: *Inventiones mathematicae* 77.2 (1984), pp. 367–381. ISSN: 1432-1297. DOI: 10.1007/BF01388451. URL: https://doi.org/10.1007/BF01388451.

[BGK10] Robert Bieri, Ross Geoghegan, and Dessislava Kochloukova. "The Sigma Invariants of Thompson's Group F". eng. In: *Groups, Geometry and Dynamics* 4.2 (2010), pp. 263–273. ISSN: 1661-7207.

[BH99] Martin R. Bridson and André. Haefliger. *Metric spaces of non-positive curvature*. eng. Die Grundlehren der mathematischen Wissenschaften 319. Berlin ; Springer, 1999. ISBN: 3540643249.

[BNR21] Jose Burillo, Brita Nucinkis, and Lawrence Reeves. "An Irrational-Slope Thompson's Group". In: *Publicacions Matematiques* 65.2 (2021), pp. 809–839. URL: https://ddd.uab.cat/record/248605. (accessed: 07.06.2022).

[BR88] Robert Bieri and Burkhardt Renz. "Valuations on Free Resolutions and Higher Geometric Invariants of Groups". eng. In: *Commentarii mathematici Helvetici* 63.1 (1988), pp. 464–497. ISSN: 0010-2571.

[Bro18] Jason Brown. "A class of PL-homeomorphism groups with irrational slopes". MA thesis. University of Melborne, 2018.

[Bro82] Kenneth S Brown. *Cohomology of groups*. eng. Graduate texts in mathematics ; 87. New York: Springer, 1982. ISBN: 0387906886.

[Bro87] Kenneth S. Brown. "Finiteness Properties of Groups". In: *Journal of Pure and Applied Algebra* 44.1 (1987), pp. 45–75. ISSN: 0022-4049. DOI: https://doi.org/10.1016/0022-4049(87)90015-6. URL: https://www.sciencedirect.com/science/article/pii/0022404987900156.

[BS16] Robert Bieri and Ralph Strebel. *On Groups of PL-homeomorphisms of the Real Line*. eng. 2016.

[BS92] Robert Bieri and Ralph Strebel. "Geometric Invariants of Discrete Groups". In: ((1992)). Preprint of a book (to be published by W. de Gruyter).

[Bur] José Burillo. *Introduction to Thompson's group F*. URL: https://web.mat.upc.edu/pep.burillo/F%20book.pdf (visited on 05/23/2024).

[CFP96] James W Cannon, William J Floyd, and Walter R Parry. "Introductory Notes on Richard Thompson's Groups". In: *Enseignement Mathematique* 42 (1996), pp. 215–256.

[Cle95] Sean Cleary. "Groups of Piecewise-Linear Homeomorphisms with Irrational Slopes". eng. In: *The Rocky Mountain Journal of Mathematics* 25.3 (1995), pp. 935–955. ISSN: 0035-7596.

[ES52] Samuel Eilenberg and Norman Steenrod. *Foundations of Algebraic Topology*. Princeton: Princeton University Press, 1952. ISBN: 9781400877492. DOI: doi:10.1515/9781400877492. URL: https://doi.org/10.1515/9781400877492.

[Far03] Daniel S. Farley. "Finiteness and CAT(0) properties of diagram groups". eng. In: *Topology (Oxford)* 42.5 (2003), pp. 1065–1082. ISSN: 0040-9383.

[Geo08] Ross Geoghegan. *Topological Methods in Group Theory*. eng. Graduate Texts in Mathematics ; 243. New York: Springer, 2008. ISBN: 9780387746111.

[Gro87] M. Gromov. "Hyperbolic Groups". In: *Essays in Group Theory*. Ed. by S. M. Gersten. New York, NY: Springer New York, 1987, pp. 75–263. ISBN: 978-1-4613-9586-7. DOI: 10.1007/978-1-4613-9586-7_3. URL: https://doi.org/10.1007/978-1-4613-9586-7_3.

[Hat02] Allen Hatcher. *Algebraic Topology*. eng. Cambridge ; New York: Cambridge University Press, 2002. ISBN: 052179160X.

[HM23] James Hyde and Justin Tatch Moore. "Subgroups of $\mathrm{PL}_+ I$ which do not embed into Thompson's group F". eng. In: *Groups, geometry and dynamics* 17.2 (2023), pp. 533–554. ISSN: 1661-7207.

[Ing56] A. W. Ingleton. "The Rank of Circulant Matrices". In: *Journal of the London Mathematical Society* s1-31.4 (1956), pp. 445–460. DOI: https://doi.org/10.1112/jlms/s1-31.4.445. eprint: https://londmathsoc.onlinelibrary.wiley.com/doi/pdf/10.1112/jlms/s1-31.4.445. URL: https://londmathsoc.onlinelibrary.wiley.com/doi/abs/10.1112/jlms/s1-31.4.445.

[JW21] Kasia Jankiewicz and Daniel T. Wise. "A curiously cubulated group". eng. In: *Geometriae dedicata* 212.1 (2021), pp. 17–19. ISSN: 0046-5755.

[KV23] Dessislava H. Kochloukova and Stefano Vidussi. "Higher dimensional algebraic fiberings of group extensions". eng. In: *Journal of the London Mathematical Society* 108.3 (2023), pp. 978–1003. ISSN: 0024-6107.

[Löh17] C. Löh. *Geometric Group Theory: An Introduction*. Universitext. Springer International Publishing, 2017. ISBN: 9783319722542. URL: https://books.google.co.uk/books?id=1AxEDwAAQBAJ.

[MMV98] John Meier, Holger Meinert, and Leonard VanWyk. "Higher generation subgroup sets and the [Sigma]-invariants of graph groups". ger. In: *Commentarii mathematici Helvetici* 73 (1998), pp. 22–. ISSN: 0010-2571.

[MNSR24] Lewis Molyneux, Brita Nucinkis, and Yuri Santos Rego. "The sigma invariants for the golden mean Thompson group". In: *The New York Journal of Mathematics* 30 (2024), pp. 532–549.

[Qui73] Daniel Quillen. "Higher algebraic K-theory: I". In: *Higher K-Theories*. Ed. by H. Bass. Berlin, Heidelberg: Springer Berlin Heidelberg, 1973, pp. 85–147. ISBN: 978-3-540-37767-2.

[Ren88] Burkhardt Renz. "Geometrische Invarianten und Endlichkeitseigenschaften von Gruppen". PhD thesis. Frankfurt, Germany: University of Frankfurt, 1988.

[Sch] Petra Schwer. *Lecture Notes on CAT(0) Cubical Complexes*. URL: https://www.uni-muenster.de/AGKramer/SeminarSS15/CubeCplx-Lecturenotes.pdf (visited on 01/19/2026).

[Ste92] Melanie Stein. "Groups of Piecewise Linear Homeomorphisms". eng. In: *Transactions of the American Mathematical Society* 332.2 (1992), pp. 477–514. ISSN: 0002-9947.

[Tan23] Owen Tanner. "Studying Stein's Groups as Topological Full Groups". eng. In: (2023).

[Wal65] C. T. C. Wall. “Finiteness Conditions for CW-Complexes”. eng. In: *Annals of mathematics* 81.1 (1965), pp. 56–69. ISSN: 0003-486X.

[Win] Nicholas Winstone. *Bieri-Strebel Groups With Irrational Slopes*. URL: https://pure.royalholloway.ac.uk/ws/portalfiles/portal/45198839/2017hollowaynwphd.pdf. (accessed: 11.04.2023).

[WZ15] Stefan Witzel and Matthew C. B Zaremsky. “The Σ-invariants of Thompson's group F via Morse theory”. eng. In: (2015).

[Zar17] Matthew Curtis Burkholder Zaremsky. “On the Σ-Invariants of Generalized Thompson Groups and Houghton Groups”. eng. In: *International mathematics research notices* 2017.19 (2017), pp. 5861–5896. ISSN: 1073-7928.

[Zar19] Matthew C. B Zaremsky. “A Short Account of Why Thompson's Group F is of Type F_∞”. eng. In: (2019).