ALGANT Master Thesis in Mathematics

Stallings Foldings

Jon Merladet Urigüen

Supervisors:

Remke Kloosterman & Pascal Weil & Montse Casals

Academic year 2021-2022
June 2022

Contents

Introduction **iii**

1 Introduction to the classical study of free groups **1**
- 1.1 Free Groups . . . 1
- 1.2 The universal property of free groups and its main consequence 2
- 1.3 The rich complexity of free groups . . . 3
- 1.4 Nielsen basis . . . 6

2 X-graphs and the lattice of subgroups of a free group **13**
- 2.1 Graphs, X-graphs and subgroups of free groups . . . 13
- 2.2 The subgroup problem . . . 17
- 2.3 Core graphs and subgroup graphs . . . 18
- 2.4 Construction of the subgroup graph
  The finitely generated case . . . 20
- 2.5 Spanning trees, geodesic trees and free bases . . . 27
- 2.6 Subgroups graphs and index
  A motivation for the infinitely generated case . . . 34
- 2.7 The infinitely generated case . . . 36

3 The lattice of subgroups of free groups **39**
- 3.1 Additional results on graphs and subgroups: Rank and intersections . . . 40
- 3.2 Overgroups . . . 41

A A technical result on finite index subgroups **45**

References **47**

Introduction

Group Theory is a branch of Mathematics that studies one of the most basic and fundamental structures in Algebra: a set with an operation, that is associative, has an identity and each element has an inverse. It is then of no surprise that Group Theory has attracted many mathematicians, both for its charm and the broad variety of applications it has in other areas of Mathematics and other sciences such as Physics and Chemistry.

This thesis is mainly about *free groups*, which are a specific class of groups. Although it may seem that we are restricting ourselves to a small branch of Group Theory, the structure of free groups contains all the information about all groups, since, in a way, any group is isomorphic to a quotient of a free group by some normal subgroup of it. This means that any group can be represented by a free group adding some relations between its generators.

We will see in Chapter 1 that free groups can be thought of as an abstract group generated by a set of symbols X, in which all elements are uniquely written as a reduced word on X, i.e., concatenations of elements of X and their formal inverses X^{-1} that do not contain xx^{-1} or $x^{-1}x$ as subwords. We denote this group by $F(X)$, and we say that it is the *free group on* X. The unique property about $F(X)$ is that there is no trivial relation between the elements of X, i.e., the only reduced word representing the identity is the empty word. This implies for example that its elements only commute with powers of themselves, there are no torsion elements, they are infinite groups. This is why they are called free groups, and X is called a *free basis*.

When one starts studying free groups, one soon realizes that this "freedom" leads to rich complexity. Indeed, in a free group $F(X)$, denoting the cardinality of X as the *rank* of $F(X)$, one may want to think about the rank as the dimension of the structure, but there is a reason why it is not called like that. In a free group of rank 2, one may find subgroups of not only rank greater than 2, but even of infinite rank! It also has an infinite number of subgroups with pairwise trivial intersection which are isomorphic to itself. This is very counter-intuitive, specially if one is used to other algebraic structures such as finite dimensional vector spaces, in which proper subspaces always have less dimension.

But as in all sciences, digging into the complex unknown is where most of the rewards are. Many great mathematicians studied free groups and successfully proved many properties about them: Jakob Nielsen initiated the study of free groups from an algebraic approach in 1924, proving many basic properties of free groups, in particular he established the fact that subgroups of a finitely generated free group are again free. Max Dehn generalized this property to infinitely generated free groups, and Otto Schreier provided an algebraic proof. This result is now known as the Nielsen–Schreier Theorem.

In this thesis, we focus on the study of free groups and we present two different methods to study them. In Chapter 1, we introduce the classic approach using word combinatorics. With this approach, we will see some basic facts, mostly proven by Nielsen in 1924, and another one, not so basic, about ascending chains of subgroups (Takahasi's theorem). In the other chapters, we introduce modern tools from topology and graph theory, first introduced by John R. Stallings in 1983. We contrast the two approaches by providing alternative proofs of the same results, such as Nielsen's theorem and Takahasi's theorem.

In Chapter 2, we will introduce the notion of X-*graphs*, foldings and subgroup graphs, which are the tools that allow us to translate all questions about the lattice of free groups into properties of graphs. With these we will show the versatility of this technique: We will reprove and sometimes extend some of the statements of Chapter 1, and also give techniques to answer questions about subgroups of free groups. We will be able to determine the index of any finitely generated subgroup of a free group; determine if two finitely generated subgroups are isomorphic, or if one is contained in the other, etc.

Finally, in Chapter 3, we will reprove Takahasi's theorem using Stallings' techniques, thus showing the versatility and beauty of Stallings foldings.

Chapter 1

Introduction to the classical study of free groups

Let us briefly recall the definition of a free group.

1.1 Free Groups

Let X be a set of symbols. The free group on X is the set of *reduced words* on X; these are the *words* on $X \cup X^{-1}$ (i.e. concatenations of elements in $X \cup X^{-1}$) which do not contain subwords of the form xx^{-1} or $x^{-1}x$ for $x \in X$. The group structure is the following.

· For w_1 and w_2 two reduced words on X, the operation $w_1 \cdot w_2$ is the reduction (cancellation of all subwords of the form xx^{-1} and $x^{-1}x$ for $x \in X$) of the concatenation w_1w_2 of the two words, which we denote by $\overline{w_1w_2}$ (we sometimes say that $\overline{w_1w_2}$ is the reduced form of w_1w_2 with respect to X). So $w_1 \cdot w_2 = \overline{w_1w_2}$.

· The identity element is the empty word, denoted by ϵ.

· The inverse of a reduced word $w = y_1 \cdots y_n \in F(X)$, where each $y_j \in X \cup X^{-1}$ for $j = 1, \dots, n$, is $w^{-1} = y_n^{-1} \cdots y_1^{-1}$.

Although basic, it is not straightforward that this operation is well defined and that it defines a group, for that one can find the details in Theorem 1 in [2].

Example 1.1.1. Let $X = \{a, b\}$, then $F(X)$ consists of all words on $\{a, b, a^{-1}, b^{-1}\}$ which do not contain aa^{-1}, $a^{-1}a$, bb^{-1} or $b^{-1}b$ as a subword (from now on we write $F(a, b)$ instead of $F(\{a, b\})$). Notice that $F(a, b)$ can also be generated by a and ab, and in fact, one can see that all elements can be written as a unique combination of a and ab. We will later see that this implies that both $\{a, b\}$ and $\{a, ab\}$ are what we will call *free bases* of

$F(X)$. Furthermore, the transformation $\{a, b\} \to \{a, ab\}$ is what we will call a *Nielsen transformation.* It is then apparent that free groups have more than one possible free basis (in fact repeating this argument one obtains an infinite number of bases of cardinality 2).

1.2 The universal property of free groups and its main consequence

As the name indicates, free groups are free objects in the category of groups. We show next that they satisfy the universal property and exhibit the main consequence of it: Every group is a quotient of a free group.

Let X be a set and let $F(X)$ be the free group on X. Let H be an arbitrary group and let $\psi\colon X \to H$ be an arbitrary map. Then we can extend ψ to a homomorphism of groups $\Psi\colon F(X) \dashrightarrow H$ making the following diagram commutative:

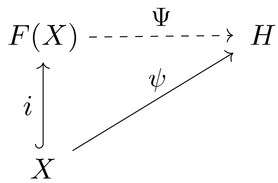

where i denotes the inclusion map from X to $F(X)$.

Indeed, for a reduced group word w on X with reduced form

$$w = x_1^{\epsilon_1} \cdots x_n^{\epsilon_n}, \tag{1.1}$$

where $x_i \in X$ and $\epsilon_i = \pm 1$ for $i = 1, \dots, n$, we define the image of w under Ψ as

$$\Psi(w) = \psi(x_1)^{\epsilon_1} \psi(x_2)^{\epsilon_2} \cdots \psi(x_n)^{\epsilon_n}. \tag{1.2}$$

It is easy to see that this map Ψ defines a homomorphism (see [2]).

This is in fact a universal property, i.e., it characterizes free groups. So if a group G with a generating set S satisfies it, we say that G is *free* on S, and S is a free basis of G (or that S freely generates G). This is equivalent to saying that all elements of G can be uniquely written as a reduced word on S, or that non-trivial reduced words on S correspond to a non-trivial element of G. We say that the cardinality of a free basis of G is the *rank* of G, which we will prove to be well defined in the next section (as we still do not know if different free bases of the same group can have different cardinalities).

Now here is the aforementioned main consequence of this property. Let G be a group with a generating set S, and let X be the set of symbols

$$X = \{x_s \mid s \in S\}. \tag{1.3}$$

Then by the universal property, the map $\psi\colon X \to G$, that sends x_s in X to s in $S \subset G$ extends to a homomorphism Ψ making the diagram below commutative.

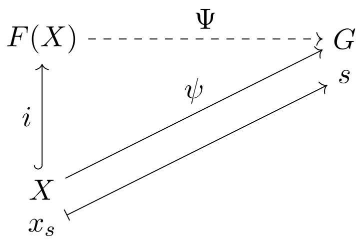

Since by definition $\Psi(X) = \psi(X) = S$, we have that $\Psi(F(X)) = \Psi(\langle X \rangle) = \langle S \rangle = G$, and thus, Ψ is clearly surjective. So the nice application is that by the First Isomorphism Theorem of groups, we have that

$$\frac{F(X)}{\ker(\Psi)} \simeq G. \tag{1.4}$$

In conclusion, this means that all groups are isomorphic to a quotient of a free group, so as mentioned in the introduction, all information of a group is contained in the lattice of a certain free group.

1.3 The rich complexity of free groups

Let us see some properties of free groups that we mentioned in the introduction which show how big and complicated free groups can be. In this section, we use word combinatorics to prove these properties and in chapters 2 and 3 we will show how we can obtain neat and simplified proofs of some of these properties using Stallings techniques.

Lemma 1.3.1. *The free group $F(a, b)$ of rank 2 contains a free subgroup of countable rank.*

Proof. For $n \in \mathbb{N}$, let x_n denote the word $a^{-n}ba^n$. Then we claim that the set $Y = \{x_n \mid n \in \mathbb{N} \cup \{0\}\}$ freely generates a free subgroup of $F(a, b)$. Indeed, let us see that a non-trivial reduced word w on Y is non-trivial in $F(a, b)$. Assume that w has the following reduced form with respect to Y:

$$w = a^{-i_1}b^{\epsilon_1}a^{i_1}a^{-i_2}b^{\epsilon_2}a^{i_2} \cdots a^{-i_n}b^{\epsilon_n}a^{i_n}, \tag{1.5}$$

where for each $j = 1, \ldots, n$, we have that $i_j \in \mathbb{N}$ and $\epsilon_j \in \{-1, +1\}$. Since w is reduced on Y, we have that either $i_j \neq i_{j+1}$ or $\epsilon_j + \epsilon_{j+1} \neq 0$. But w need not be reduced on $\{a, b\}$.

For two consecutive b's to cancel, it must happen that $i_j = i_{j+1}$ and $\epsilon_j = -\epsilon_{j+1}$ for some $j = 1, \ldots, n-1$. But by the characterization of w being reduced on Y, one of the equalities must be false. So we see that no b^{ϵ_j}'s are cancelled with each other, at most, they are brought next to one another. This implies that ϵ can not be the reduced form of w with respect to $\{a, b\}$, as we wanted. □

This exercise already shows how apparently messy can things go inside free groups. But at least we have that two free bases of a free group have the same cardinality, thus making the definition of rank of a free group well defined.

Theorem 1.3.2. *Let $F(X)$ and $F(Y)$ be two free groups on X and Y respectively. If X and Y are finite or countable sets, then*

$$F(X) \simeq F(Y) \iff |X| = |Y|. \tag{1.6}$$

Proof. Let us assume first that $|X| = |Y|$. Then take a bijective map $f\colon X \to Y \subseteq F(Y)$. By the universal property of free groups, f extends to a group homomorphism $F\colon F(X) \to F(Y)$. By definition of free basis, and since X maps bijectively to Y, we have that F is an isomorphism of groups, as we wanted.

Let us now assume that $F(X) \simeq F(Y)$. Again by the universal property of free groups, all possible homomorphisms

$$f\colon F(X) \longrightarrow C_2, \tag{1.7}$$

where C_2 denotes the cyclic group of order 2, are given by choosing an arbitrary image of each element of X.

This implies that there are $2^{|X|}$ possible homomorphisms. Then since $F(X) \simeq F(Y)$, we have that $2^{|X|} = 2^{|Y|}$. In the case that $|X|$ and $|Y|$ are finite, we can say that $|X| = |Y|$. Furthermore, we can say in general that X and Y must be either both infinite or both finite sets.

In the case that both X and Y are infinite countable sets, we will make use of the following claim, which we will prove later.

Claim 1.3.3. *A countable union of countable sets is countable.*

Now, we have that $F(X) \simeq F(Y)$ implies $|F(X)| = |F(Y)|$. So we will show that for any infinite countable set Z, we have that $|F(Z)| = |Z|$, thus finishing the proof. Indeed, it is clear that $|Z| \leq |F(Z)|$ since we can embed Z in $F(Z)$, so let us see that $|F(Z)| \leq |Z|$.

Recall that $F(Z)$ are reduced words on Z of a certain length, so

$$F(Z) \subset \bigcup_{i \geq 0} F(Z)_i, \tag{1.8}$$

where $F(Z)_i$ denotes the words (reduced and non reduced) on $Z \cup Z^{-1}$ of length i for $i = 0, 1, \dots$. So if we prove that $F(Z)_i$ is countable for every i, we are done by the claim.

One can see $F(Z)_i$ as the direct product of $Z \cup Z^{-1}$ with itself i times, so for i greater than 0, we can write $F(Z)_i$ as the following countable union

$$F(Z)_i = \bigcup_{z \in Z \cup Z^{-1}} F(Z)_{i-1} \times \{z\}. \tag{1.9}$$

Then since $F(Z)_0 = \{\epsilon\}$ is countable and $Z \cup Z^{-1}$ too, we can clearly see by induction, again by the claim, that all $F(Z)_i$ are countable, and thus we are done. □

Proof of the claim. Let I be a countable index set, which we may assume that is $\mathbb{N}$. And let $Z_i = (z_i^1, z_i^2, z_i^3, \dots)$ be countable sets for $i \in \mathbb{N}$, then

$$\bigcup_{i \in \mathbb{N}} Z_i = \{z_1^1, z_1^2, z_2^1, z_1^3, z_2^2, z_3^1, \dots, z_j^k, \dots\}, \tag{1.10}$$

where the next term of z_j^k is z_{j+1}^{k-1} if $k > 1$ and z_1^{j+1} otherwise. So the union is countable. □

Remark. Theorem 1.3.2 can be extended to sets X and Y of arbitrary cardinality. To prove it, one should first generalize the claim to countable unions of sets of greater cardinality, but this would require more machinery. Either way, in this work we will only consider free groups of countable rank.

From now on, we denote by F_n the free group of rank n, where n is a natural number.

With this theorem and the previous lemma we have the following fact: A free group F of countable rank greater than or equal to 2 contains an infinite number of subgroups that have trivial intersection and that are isomorphic to F itself. Indeed, take the free group of rank 2, then we know there is a subgroup of infinite rank inside, from which we can obtain an infinite number of subgroups with the same rank as F with trivial intersection, which by the theorem above are isomorphic to the original group.

Let us construct some subgroups of $F(a, b)$ in a similar fashion as in Lemma 1.3.1. Let H_1 be the subgroup of rank 2 generated by the free basis $\{x_0, x_1\}$ with $x_0 = b$ and $x_1 = b^a$. Then define the subgroup $H_2 = \langle x_1, x_1^{x_0} \rangle$, which is contained in H_1 and is also of rank 2 (one can view it as a subgroup of $F(x_0, x_1)$, where we treat x_0 and x_1 as symbols). Repeating this process we obtain the chain

$$F(a, b) > H_1 = \langle b, b^a \rangle > H_2 = \langle b^a, (b^a)^b \rangle > H_3 = \langle (b^a)^b, ((b^a)^b)^{b^a} \rangle > \cdots .$$

We have just constructed an infinite strictly descending chain of subgroups of rank 2. One may wonder if we can construct a similar ascending chain of subgroups of rank 2. But one soon encounters a problem: Intuitively, each time we consider a bigger subgroup, the words tend to be shorter (this is more evident when considering subgroups of rank 1, i.e., cyclic subgroups). We can see this in reverse in our ascending chain example, the words grow longer as we consider smaller subgroups. So we can expect that maybe this phenomenon is impeding us to construct ascending chains of bounded rank. Note that if we allow the subgroups to have arbitrary rank, we can construct an ascending chain from any subgroup of infinite rank.

Let us first formalize this notion of size of a word.

Definition 1.3.4. Let $F(X)$ be the free group on a set of symbols X and let $w = y_1 y_2 \cdots y_n$ be a reduced word on X, where $y_i \in X \cup X^{-1}$ for $i = 1, \dots, n$.
We then say that n is the *word length* of w with respect of X, and we write $|w|_X = n$.

Let $F(X)$ be the free group on $X = \{a, b\}$. Then products of elements in $\{a, b, a^{-1}, b^{-1}\}$ have greater word length with respect to X than any of its components (as long as we do not concatenate an element with its inverse). If this were to be true with for an arbitrary free basis of a general subgroup of $F(X)$, then we could explain why we cannot construct infinite ascending chains of rank 2. But not all free bases satisfy this. Indeed, take the free basis $X' = \{a, ab\}$ of $F(X)$, then $|a^{-1} \cdot ab|_X = |b|_X = 1$ but $|ab|_X = 2$. So this word length growing condition is sometimes satisfied and sometimes not.

1.4 Nielsen basis

Definition 1.4.1. Let $F(X)$ be a free group, and let H be a subgroup of $F(X)$. Then we say that a subset S of H is a *Nielsen basis* of H with respect to X if S generates H and S satisfies the following three conditions:

$(N0)$ $\forall\ s_1,\ s_2 \in S$, we have that $s_1 \cdot s_2 \neq \epsilon$,

$(N1)$ $\forall\ s_1,\ s_2 \in S \cup S^{-1}$ such that $s_1 \neq s_2^{-1}$, we have the inequalities $|s_1 \cdot s_2|_X \geq |s_1|_X$ and $|s_1 \cdot s_2|_X \geq |s_2|_X$,

$(N2)$ $\forall\ s_1,\ s_2,\ s_3 \in S \cup S^{-1}$ such that $s_1 \neq s_2^{-1} \neq s_3$, we have the inequality $|s_1 \cdot s_2 \cdot s_3|_X > |s_1|_X - |s_2|_X + |s_3|_X$.

Examples 1.4.2. Let w be a reduced word on X, then $\{w\}$ is a Nielsen basis of $\langle w \rangle$ with respect to X. Indeed, all conditions can be proven using the following fact: The word w has some reduced form $v^{-1}uv$ such that w^2

has reduced form $v^{-1}uuv$. And thus, we have that $|w^n|_X = 2|v|_X + n|u|_X$. Note that with this we can prove that non trivial free groups are torsion free.

Furthermore, the set $\{a^{-n}ba^n \mid n \in \mathbb{N} \cup \{0\}\}$ introduced in Lemma 1.3.1 is also a Nielsen basis of the subgroup of $F(a, b)$ it generates.

Remark 1.4.3. Condition $(N0)$ is equivalent to $S \cap S^{-1} = \emptyset$. Condition $(N1)$ is equivalent to saying that no more than half of each word is cancelled in the product $s_1 \cdot s_2$. Condition $(N3)$ is equivalent to saying that at least one letter of s_2 is not cancelled in the product $s_1 \cdot s_2 \cdot s_3$. The three conditions combined imply that a word w on $S \cup S^{-1}$ with non trivial reduced form on S

$$s_1 \cdots s_n, \tag{1.11}$$

has greater word length with respect to X than each of the elements of S that contribute to generate w, i.e,

$$|w|_X > |s_1|_X, \ldots, |s_n|_X. \tag{1.12}$$

It is left to the reader to check that these facts about Nielsen basis are true. For the last fact, we suggest checking that $|s_1 \cdots s_n|_X > |s_1 \cdots s_{n-1}|_X$ and $|s_1 \cdots s_n|_X > |s_2 \cdots s_n|_X$.

Our goal is to obtain a Nielsen basis of any finitely generating subgroup of $F(X)$. To achieve this, we will need some transformations that modify a given generating set so that it still generates the subgroup.

Definition 1.4.4. Let A be a set of elements of a free group $F(X)$, and denote by H the subgroup generated by them. Then the following operations are called *Nielsen Transformations*:

$(T1)$ Substitute $a \in A$ with a^{-1}.

$(T2)$ For $a,\ b \in A$ such that $a \neq b$, substitute b with ab.

$(T3)$ Remove $a \in A$ if $a = \epsilon$.

It is easy to check that after applying one of these transformations to a generating set A, we obtain another generating set. Furthermore, if A is a free basis of H, then applying transformations $(T1)$ or $(T2)$ results in another free basis of H.

Note that combining these Nielsen transformations, we can also substitute b with ba, ba^{-1} and $a^{-1}b$ (we call these operations "variants ' of $(T2)$).

As it will be useful in the next theorem, we present here the notion of *initial* and *terminal* subword. For w and v two reduced words in X, we say that w has initial subword v when there exist a reduced word v' such that w has reduced form vv'. We say that w has terminal subword v if w^{-1} has

v^{-1} as a initial subword. Furthermore, we define the *left half* of a word w to be the initial subword $L_h(w)$ of w of length $\lfloor \frac{|w|_X+1}{2} \rfloor$.

In the following theorem, we will repeatedly apply Nielsen transformations to a set A. To avoid confusion, notice that after each transformation, the set A is modified, but each time the notation A is used, it denotes the transformed set, not the initial one.

Theorem 1.4.5. *Let H be a finitely generated subgroup of a free group $F(X)$, then any finite generating subset of H can be transformed using Nielsen transformations into a Nielsen basis.*

Proof. Let $A = \{h_1, \dots, h_n\}$ be a generating set of H. Before continuing, notice that we can assume that X is a finite set by not considering the elements of X such that neither them nor their inverses appear in any of the reduced forms of the elements of A. Now proceed to Step 1.

Step 1: Apply transformations $(T2)$ and $(T3)$ repeatedly to A until it satisfies $(N0)$.

Step 2: If A satisfies $(N1)$, go to Step 3. Otherwise, we can apply transformation $(T2)$ (or one of its variants) to A to substitute one element of A with another word that has less word length with respect to X. Note that we can always apply $(T2)$ because the case $|w \cdot w|_X < |w|_X$ is not possible by Example 1.4.2.

If the transformed set does not satisfy $(N0)$, go to Step 1. If the transformed set satisfies $(N1)$, go to Step 3, otherwise, repeat Step 2.

Since $\sum_{h_i \in A} |h_i|_X$ is finite, we deduce that A will satisfy $(N1)$ after a finite number of repetitions of Step 2.

Define a total order on $X \cup X^{-1}$ (which is possible since X is finite), which in turn extends to a lexicographical order $<$ on $F(X)$. Now we define a new order $\prec$ of $F(X)$: For $w_1,\ w_2 \in F(X)$, then $w_1 \prec w_2$ if either $\min\{L_h(w_1),\ L_h(w_1^{-1})\} < \min\{L_h(w_2),\ L_h(w_2^{-1})\}$ or the two minima are equal and $\max\{L_h(w_1),\ L_h(w_1^{-1})\} < \max\{L_h(w_2),\ L_h(w_2^{-1})\}$. It is easy to see that $\prec$ is a partial order. Furthermore, we claim the following, which we will prove later.

Claim 1.4.6. *In the context of this theorem, the order $\prec$ is a total order.*

Step 3: If A satisfies $(N2)$, we are done. Otherwise, we will use the total order $\prec$ to deduce that we can apply transformations to A so that it satisfies $(N2)$ at some point. Let (v_1, v_2, v_3) be a triplet of elements in $A \cup A^{-1}$ which does not satisfy condition $(N2)$ (we call it a *breaking triplet*). Then we know that v_2 must be completely cancelled in the product of the three. So the components of the triplet must have reduced form

$$v_1 = u_1 x^{-1},\ v_2 = xy,\ v_3 = y^{-1} u_3, \tag{1.13}$$

for some $u_1,\ x,\ y,\ u_3 \in F(X)$. Furthermore, since the set A already satisfies condition $(N1)$, we have by Remark 1.4.3 the inequalities $|x|_X \leq 1/2|xy|_X$ and $|y|_X \leq 1/2|xy|_X$, which implies that $|x|_X = |y|_X$. By the same remark, we can also deduce that $|x^{-1}|_X \leq |u_1|_X$ and $|y^{-1}|_X \leq |u_3|_X$, which implies that

$$L_h(u_1x^{-1}) = L_h(u_1y) \text{ and } L_h(u_3^{-1}y) = L_h(u_3^{-1}x^{-1}). \tag{1.14}$$

Then, since $x \neq y^{-1}$, we have two options: We have that $x < y^{-1}$, in which case we apply transformation $(T2)$ to substitute v_3 with $v_2 \cdot v_3 = xu_3$ (note that using Equation (1.14), it can be seen that $v_2 \cdot v_3 \prec v_3$). Or we have that $y^{-1} < x$, in which case we apply a variant of transformation $(T2)$ to substitute v_1 with $v_1 \cdot v_2 = u_1y$ (note that using Equation (1.14), it can be seen that $v_1 \cdot v_2 \prec v_1$).

We note that both transformations do not change the sum $\sum_{h_i \in A} |h_i|_X$, but we do not know if $(N1)$ is still satisfied. So if $(N1)$ is not satisfied, go to Step 2. Otherwise, go to Step 4.

Step 4: If A satisfies $(N2)$, we are done. Otherwise, go to Step 3.

Now let us see why this algorithm must terminate, thus transforming A into a Nielsen basis.

On the one hand, note that the sum $\sum_{h_i \in A} |h_i|_X$ is finite, and Step 3 and 4 do not increase such sum. So the number of times Step 3 lead us to Step 2 is finite, which means that we can assume that we never return to Step 2. Let n be the number of elements of A at this point (observe that since we no longer repeat Step 2, the cardinality of A will not change).

On the other hand, after each repetition of Step 3, we reduce the order of some element of A with respect to the order $\prec$. Since $\prec$ is a total order, and the number of subsets of n elements with the same sum $\sum_{h_i \in A} |h_i|_X$ is finite, we can see that we can only repeat Step 3 a finite number of times, so after a finite number of repetitions, we must have transformed A into a Nielsen basis. □

Proof of the claim. Let w_1 and w_2 be two distinct reduced words of $F(X)$. Assume that $w_1 \not\prec w_2$, then let us prove that $w_2 \prec w_1$. The statement $w_1 \not\prec w_2$ is equivalent to saying that either $\min\{L_h(w_1),\ L_h(w_1^{-1})\} > \min\{L_h(w_2),\ L_h(w_2^{-1})\}$ or the two minima are equal and $\max\{L_h(w_1),\ L_h(w_1^{-1})\} \geq \max\{L_h(w_2),\ L_h(w_2^{-1})\}$. In the first case, we clearly have that $w_2 \prec w_1$, as we wanted. In the second case, we can exclude the case of the two maxima being equal, since otherwise, the two minima and the two maxima being equal imply that $w_1 = w_2$, contrary to our hypothesis. Then, since the inequality of the two maxima is strict, we can deduce that $w_2 \prec w_1$. □

Here is a nice corollary of this theorem, first proved by Jacob Nielsen.

Corollary 1.4.7 (Nielsen Subgroup Theorem). *Let H be a finitely generated subgroup of a free group $F(X)$, then H is free of finite rank.*

Proof. Let A be a finite generating set of H, then by the previous theorem, we can transform A into a Nielsen Basis. This basis generates H and by remark 1.4.3, non-trivial reduced words on A have word length (with respect to X) greater than 0, and thus, do not correspond to the empty word ϵ. □

Now with this machinery, let us see how one would classically prove that ascending chains of subgroups of bounded rank must terminate.

Theorem 1.4.8 (Takahasi, 1951). *Let X be a finite set, let $F(X)$ be the free group on X and let M be a positive integer. If we have the following chain of groups*

$$K_1 \leq K_2 \leq \cdots \leq K_i \leq \cdots , \tag{1.15}$$

where K_i are subgroups of $F(X)$ of rank at most M, then for some $N \in \mathbb{N}$, we have that $K_j = K_{j+1}$ for all $j \geq N$, i.e., the chain stabilizes.

Proof. By Theorem 1.4.5, we know that there is a Nielsen basis for each subgroup K_i with respect to X. We also know that the length (with respect to X) of a word generated by a Nielsen basis is greater or equal than the length (with respect to X) of each of the generators that contribute to generate such word (see remark 1.4.3). The main idea of this proof is that the word length of the generators will somehow decrease as i increases, but we have to take into account some subtleties.

For $i = 1, 2, \ldots$, let $T_i = \{w_1^i, w_2^i, \ldots, w_{M_i}^i\}$ be a Nielsen basis for each K_i, where $M_i \leq M$. Then we have that

$$K_1 \leq \cdots \leq K_i = \langle w_1^i, w_2^i, \ldots, w_{M_i}^i \rangle \leq K_{i+1} = \langle w_1^{i+1}, w_2^{i+1}, \ldots, w_{M_{i+1}}^{i+1} \rangle \leq \cdots ,$$

which implies that each generator w_j^i of K_i can be written as a reduced word on T_{i+1}.

We know that the word length (with respect ot X) of each generator w_j^i of K_i is greater than or equal to the word length (with respect to X) of some elements of the Nielsen basis $\{w_1^{i+1}, \ldots, w_{M_{i+1}}^{i+1}\}$ of K_{i+1}, but it need not be true that all generators of K_{i+1} have smaller word length to some of K_i, as there could be generators of K_{i+1} that do not contribute to generate any element of K_i (i.e. the subgroup obtained by removing such generators still contains K_i). So to use the notion of word length, we will have to somewhat ignore these generators.

Step 1. For $i = 2, 3, \ldots,$ let T_i' be the least subset of T_i such that $K_{i-1} \leq \langle T_i' \rangle$ (put $T_1' = T_1$).

Now by the property of Nielsen bases mentioned above and the way we constructed these sets, we have the sequence of positive integers

$$\max_{w \in T'_1}\{|w|_X\} \geq \max_{w \in T'_2}\{|w|_X\} \geq \cdots,$$

which must stabilize.

Since there is only a finite number of words of a fixed length (with respect to X) in $F(X)$ (X is finite by hypothesis), there must be an infinite subsequence $\{T'_{j_i}\}_{i \in \mathbb{N}}$, with $0 < j_1 < j_2 < \cdots$ integers, in which a word denoted by v_1 is contained in T'_{j_i} for all $i = 1, 2, \ldots$. Now consider the correspondent subchain of the original chain

$$K_{j_1} \leq K_{j_2} \leq K_{j_3} \cdots, \tag{1.16}$$

with their correspondent Nielsen bases T_{j_1}, T_{j_2}, $T_{j_3}, \ldots$, which all contain the element v_1. If these Nielsen bases only consists of v_1, we have finished (since it would imply that $K_{j_1} = K_{j_2} = \cdots$). Otherwise, we can assume that all these Nielsen basis contain more elements apart from v_1, and we go to Step 2.

Step 2. Repeat Step 1 but now with the chain $K_{j_1} \leq K_{j_2} \leq \cdots$ with their respective Nielsen bases T_{j_1}, T_{j_2}, $T_{j_3} \ldots$. We will obtain a sequence of sets $\{T''_{j_i}\}$, where each T''_{j_i} is the least subset of T_{j_i} such that $K_{j_{i-1}} \leq \langle T''_{j_i} \rangle$. All of these sets must contain v_1 by the way they are constructed, and by the assumption of the previous paragraph, they must contain other words apart from v_1. Now, similarly as before, we consider the sequence of positive integers

$$\max_{w \in T''_{j_1} \setminus \{v_1\}}\{|w|_X\}, \ \max_{w \in T''_{j_2} \setminus \{v_1\}}\{|w|_X\}, \ \max_{w \in T''_{j_3} \setminus \{v_1\}}\{|w|_X\}, \ \cdots$$

where we have removed v_1 from consideration. Then we can see that this sequence is also decreasing. Indeed, for $i = 2, 3, \ldots$, we know that each element of $T''_{j_{i+1}} \setminus \{v_1\}$ contributes to generate an element of T''_{j_i}, but it cannot be v_1, since each element of T''_{j_i} is obtained with a unique word on $T''_{j_{i+1}}$, so v_1 is only obtained with v_1.

Then as in Step 1, this sequence must stabilize, so we can find a subchain of the chain $K_{j_1} \leq K_{j_2} \leq \cdots$, in which both v_1 and v_2 are contained in their respective Nielsen bases. If v_1 and v_2 already generate all these subgroups, we are done, otherwise we go to the next step.

This process cannot go further than Step M, as M is the bound on the rank of the subgroups of the chain. So at some point we must find an infinite subchain of the original chain with exactly the same Nielsen basis, which implies that the original chain stabilizes as we wanted. $\square$

Chapter 2

X-graphs and the lattice of subgroups of a free group

We will now start discussing the more recent approach to the study of free groups. The main idea is to represent a free group $F(X)$ as the fundamental group of a bouquet of $|X|$ circles, and subgroups of $F(X)$ as other graphs. This approach is centred in the notion of graphs, so we begin this section introducing the main concepts of graph theory that we will use.

Now let us start discussing the main structure that we will use to study free groups.

2.1 Graphs, X-graphs and subgroups of free groups

As we mentioned, we are interested in representing a subgroup of a free group with a graph. We first recall the notion of a graph, and then we relate it to a free group via the fundamental group.

We present below a formal definition of a standard *graph* to fix the notation.

Definition 2.1.1. A graph is a quadruplet $(V, E, (\cdot)^{-1}\colon E \to E, \iota\colon E \to V)$, where V is said to be the set of vertices, E the set of edges, the edge $(\cdot)^{-1}(e)$ the reverse of e (which we denote by $(e)^{-1}$) and $\iota(e)$ the initial vertex of e, that satisfy the following two conditions:

(i) $e^{-1} \neq e$,

(ii) $(e^{-1})^{-1} = e$.

Finally, we denote $\iota(e^{-1})$ by $\tau(e)$ and we call it the terminal vertex of e.

We can give a graphical representation of a graph as a set of points and arrows connecting them. Let us see an example.

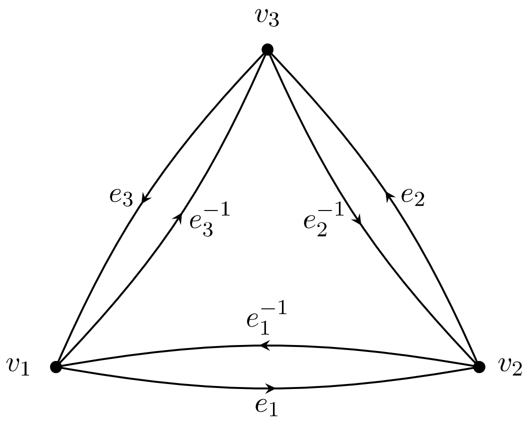

To the left we have the picture that represents the graph given by $V = \{v_1, v_2, v_3\}$, $E = \{e_1, {e_1}^{-1}, e_2, {e_2}^{-1}, e_3, {e_3}^{-1}\}$ and $\iota(e_1) = v_1$, $\iota({e_1}^{-1}) = v_2$, $\iota(e_2) = v_2$, $\iota({e_2}^{-1}) = v_3$, $\iota(e_3) = v_3$, $\iota({e_3}^{-1}) = v_1$.

Although this is a valid representation of a graph, we would like to work with a simpler representation. Drawing just one edge from each pair $\{e, e^{-1}\}$ defines an *oriented visual representation* of the graph, where we draw the arrow corresponding to each chosen edge, hiding but not forgetting the reverse edges. From here on, when we count the number of edges on a given graph, we always mean the number of pairs of edges, or the number of arrows drawn in an oriented visual representation. In the case above, choosing e_1, e_2 and e_3 results in the oriented visual representation below.

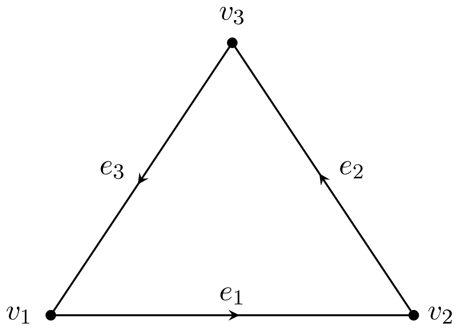

Now for a graph $(V, E, (\cdot)^{-1} \colon E \to E, \iota \colon E \to V)$, we say that a path p is a finite sequence of edges in E in which the terminal vertex of an edge in the sequence is the initial vertex of the next edge. So

$$p = e_1 \cdots e_n, \tag{2.1}$$

where $\tau(e_i) = \iota(e_{i+1})$ and $e_i \in E$ for $i = 1, \ldots, n-1$. We say that the reverse path of p is the path $e_n^{-1} \cdots e_1^{-1}$, which we denote by p^{-1}.

Note that when representing a graph with an oriented visual representation, it does not mean that paths must follow the direction of the arrows drawn. The edges allowed in a path are always taken from the whole set of edges E.

We say that the number of edges in a path is the length of such path, and we denote such number by $|p|$. We say that the initial and terminal vertex of p are $\iota(e_1)$ and $\tau(e_n)$ respectively, and if $\iota(e_1) = \tau(e_n) = v$ for some vertex v, we say that p is a *circuit* with respect to v. If a path contains no subpaths of the form ee^{-1}, we say that the path is reduced. The set of reduced circuits with respect to a vertex v is called *the fundamental group of* Γ *with respect to* v, and it is denoted by $\pi_1(\Gamma, v)$.

Analogously to the operation of free groups, the fundamental group of a graph is in fact a group (as the name implies), in which the multiplication $q \cdot q'$ of two circuits q and q' consists of concatenating and reducing, i.e., removing subpaths of the form ee^{-1} until the circuit qq' is reduced (note that these removals do not change the initial and terminal vertex). The proof that this operation is well defined is analogous to the one in free groups. We will see later that this similarity is no coincidence.

Similarly as in free groups, we denote the reduced path obtained after reducing a non reduced path p by $\overline{p}$.

Remark 2.1.2. Note that we can view graphs as the topological quotient space (with the quotient topology) obtained from the disjoint union of closed bounded segments of the real line, whose endpoints may be identified between themselves. Then the classical notion of the fundamental group of a topological space coincides with our notion defined above.

Now we are interested in obtaining subgroups of free groups from graphs, and for that we introduce the concepts of X-*graphs* and *languages*.

Definition 2.1.3. [X-graph] For a set of symbols X, we say that an X-graph is a graph $(V, E, (\cdot)^{-1}, \iota)$ together with a *labelling map* $\mu\colon E \to X \cup X^{-1}$ such that $\mu(e^{-1}) = (\mu(e))^{-1}$.

We say that $\mu(e)$ is the *label* of the edge e. Now for a path $p = e_1 \cdots e_n$, we define the label of p as the following word on $X \cup X^{-1}$:

$$\mu(p) = \mu(e_1)\mu(e_2)\cdots\mu(e_n). \tag{2.2}$$

The label of a path in an X-graph is a word on $X \cup X^{-1}$, but it is not necessarily reduced. Indeed, if there is a vertex v with two edges e_1 and e_2 with the same label x such that $\iota(e_1) = \iota(e_2)$, as shown below, then the path $p = e_1^{-1}e_2$ has label $x^{-1}x$.

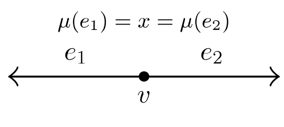

Figure 2.1: Local view of an oriented visual representation of a graph with two edges with the same label and initial vertex (the end vertex of the edges may be v itself).

To fix this we need the X-graph to be *folded*.

Definition 2.1.4. Let $\Gamma = (V, E, (\cdot)^{-1}, \iota)$ together with a labelling map $\mu\colon E \to X \cup X^{-1}$ be an X-graph. Then we say that Γ is folded when for all $a \in X \cup X^{-1}$ and for all $v \in V$, there is at most one edge with initial vertex v and label a.

Now, if Γ is a folded X-graph, then one can easily see that the label of a reduced path in Γ is a reduced word on X.

As mentioned in the introduction of this chapter, now we will obtain a subgroup of $F(X)$ from a folded X-graph Γ. We do this by taking the image of $\pi_1(\Gamma, v)$ under the map

$$\begin{aligned} \mu\colon \pi_1(\Gamma, v) &\to F(X) \\ q &\mapsto \mu(q), \end{aligned}$$

where $\mu(q)$ is the label of the circuit q, as defined in Definition 2.1.3. Note that we need Γ to be folded so that this map is well defined.

One can easily see that this map is in fact a homomorphism of groups, so the image of $\pi_1(\Gamma, v)$ under μ is a subgroup of $F(X)$, as we wanted.

Definition 2.1.5. Let Γ be an X-graph and v one of its vertices. Then we say that the *language of* Γ *with respect to* v is the image of the homomorphism

$$\mu\colon \pi_1(\Gamma, v) \to F(X), \tag{2.3}$$

and we denote it by $L(\Gamma, v)$.

Note that since Γ is folded, then the map μ is clearly injective (the label of a non trivial reduced path is a non trivial reduced word), and thus

$$\pi_1(\Gamma, v) \simeq L(\Gamma, v) \leq F(X). \tag{2.4}$$

Remark 2.1.6. The injectivity of μ for folded graphs means that there is at most one reduced circuit with a certain label, but in fact, one can generalize this to reduced paths, i.e., there is at most one reduced path with initial vertex v and with a certain label: Indeed, suppose there are two reduced paths p_1 and p_2 with the same label and starting at v. Then the circuit $p_1^{-1}p_2$ maps to $\mu(p_1)^{-1}\mu(p_2) = \epsilon$, and thus, the path reduced form of $p_1^{-1}p_2$ is the empty circuit, which implies $p_1 = p_2$ (since both paths are reduced, the path reductions in the sequence $p_1p_2^{-1}$ start at the middle where p_1 meets p_2^{-1}).

In conclusion, we have that a folded X-graph represents a subgroup of $F(X)$. In the next section, we will work on the reverse, i.e. obtaining a unique X-graph from a given subgroup of a free group. From this graph we will be able to obtain information about the subgroup. This will allow us to translate all the information in the lattice of free groups to X-graphs and vice versa.

Let $F(X)$ be the free group on the set of symbols X, and let H be a subgroup of $F(X)$. We want to assign to H a folded X-graph $\Gamma(H)$ so that $L(\Gamma(H), v) = H$. Before starting the constructions, let us see how "relations" between X-graphs give us information between the subgroups they represent.

2.2 The subgroup problem

Below we have two X-graphs Γ_1 and Γ_2, where $X = \{a, b\}$.

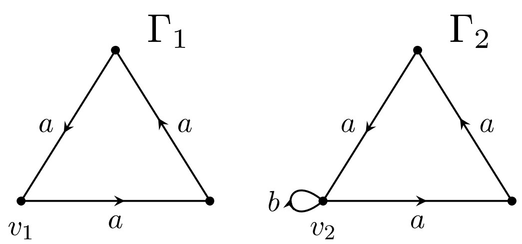

Figure 2.2: Visual representations of two different graphs Γ_1 and Γ_2 such that $L(\Gamma_1, v_1) \leq L(\Gamma_2, v_2)$.

It is clear that Γ_1 can be embedded into Γ_2, and that any circuit with its respective label in Γ_1 can also be found in Γ_2 with the same label, so in fact, we have that $L(\Gamma_1, v_1) \leq L(\Gamma_2, v_2)$. This can be generalized by considering *maps of X-graphs.*

Definition 2.2.1. Let Γ_1 and Γ_2 be two X-graphs with vertices v_1 and v_2 in Γ_1 and Γ_2 respectively. Then we call map of X-graphs a map f that sends vertices to vertices and edges to edges that respects incidence and labels, i.e, for all edges e in Γ_1, we have that

$$f(\iota(e)) = \iota(f(e)),\ f(e^{-1}) = f(e)^{-1} \text{ and } \mu(f(e)) = \mu(e). \tag{2.5}$$

If f sends v_1 to v_2, and we want to emphasize it, we denote it by $f\colon (\Gamma_1, v_1) \to (\Gamma_2, v_2)$. Otherwise, we write $f\colon \Gamma_1 \to \Gamma_2$.

We say that two X-graphs Γ_1 and Γ_2 are isomorphic as X-graphs if there are maps of X-graphs $f_1\colon \Gamma_1 \to \Gamma_2$ and $f_2\colon \Gamma_2 \to \Gamma_1$, such that

(i) $f_2 \circ f_1 = id_{\Gamma_1}$,

(ii) $f_1 \circ f_2 = id_{\Gamma_2}$,

where for a graph Γ, id_Γ denotes the identity map $id_\Gamma\colon \Gamma \to \Gamma$ that fixes all vertices and edges. In this case, we say that both f_1 and f_2 are isomorphisms of X-graphs.

Furthermore, we say that a map of X-graphs f is a *monomorphism* of X-graphs if it is injective on both the set of vertices and the set of edges.

Similarly, we say that f is an *epimorphism* of X-graphs if it is surjective on both the set of vertices and the set of edges.

Observe that f is an isomorphism of X-graphs if and only if it is both a monomorphism and an epimorphism of X-graphs.

It is then easy to see that if there is a map of folded X-graphs $f\colon (\Gamma_1, v_1) \to (\Gamma_2, v_2)$, then $L(\Gamma_1, v_1) \leq L(\Gamma_2, v_2)$ (see Proposition 2.3.2 for more details). As a consequence, if Γ_1 and Γ_2 are isomorphic as X-graphs, then $L(\Gamma_1, v_1) = L(\Gamma_2, v_2)$.

Now the question is this: For a given subgroup H of $F(X)$, is there a unique folded X-graph $\Gamma(H)$ up to isomorphism such that $L(\Gamma(H), v_1) = H$? The answer is no. In the next section we will see an example, and how to achieve uniqueness up to isomorphism.

2.3 Core graphs and subgroup graphs

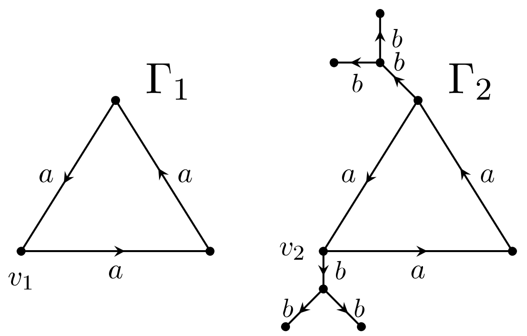

Figure 2.3: Visual representations of two different folded X-graphs Γ_1 and Γ_2 such that $L(\Gamma_1, v_1) = L(\Gamma_2, v_2) = \langle a^3 \rangle$.

As we can see in Figure 2.3, there is more than one X-graph representing the same subgroup of $F(X)$, so we will impose some conditions on the graphs to construct a correspondence between X-graphs and subgroups of free groups. In particular, we will define *core graphs*, in which there are no unnecessary trees 'hanging' on the graph.

Definition 2.3.1. Let Δ be a folded X-graph and v one of its vertices. Then we denote the union of all reduced circuits in Δ with respect to v by $\mathrm{Core}(\Delta, v)$, i.e.,

$$\mathrm{Core}(\Delta, v) = \cup\{q \mid q \text{ is a reduced circuit in } \Delta \text{ with respect to } v\}. \quad (2.6)$$

If $\mathrm{Core}(\Delta, v) = \Delta$, we say that Δ is a *core graph* with respect to v.

It is then clear that in general, we have that $L(\Delta, v) = L(\mathrm{Core}(\Delta, v), v)$. Note that being a core graph with respect to v is equivalent to all vertices except v having more than one edge leaving them.

We next show that for a subgroup H of $F(X)$, there is a unique up to isomorphism folded core X-graph Γ with respect to one of its vertices v such that $L(\Gamma, v) = H$. We denote this unique X-graph by $\Gamma(H)$ and we call it the *subgroup graph* of H with respect to X or X-subgroup graph. For reasons that we will see later, we will write 1_H instead of v to reference the vertex the subgroup graph is based at.

Let us first prove the uniqueness up to isomorphism part. For that we need the following proposition.

Proposition 2.3.2. *Let H_1 and H_2 be two subgroups of a free group $F(X)$ of finite rank, and let Γ_1 and Γ_2 be two folded core X-graphs with respect to vertices 1_{H_1} and 1_{H_2} respectively, such that $L(\Gamma_1, 1_{H_1}) = H_1$ and $L(\Gamma_2, 1_{H_2}) = H_2$. Then $H_1 \leq H_2$ if and only if there is a map of X-graphs $f\colon \Gamma_1 \to \Gamma_2$ such that $f(1_{H_1}) = 1_{H_2}$. In which case, the map is unique.*

Proof. Suppose we have a map $f\colon \Gamma_1 \to \Gamma_2$. Let $h \in H_1 = L(\Gamma_1, 1_{H_1})$, so there is a unique (by Remark 2.1.6) reduced circuit q_1 in Γ_1 with respect to 1_{H_1} with label h. Then the image of q_1 under f is the unique reduced circuit q_2 in Γ_2 with respect to $f(1_{H_1}) = 1_{H_2}$ with the same label h (unique by Remark 2.1.6). So $h = \mu(q_2) \in L(\Gamma_2, 1_{H_2}) = H_2$, and thus, $H_1 \leq H_2$. We also have that f is unique: Each reduced circuit can only be sent to the unique circuit with the same label, and core graphs are the union of such reduced circuits.

For the converse, assume that $H_1 \leq H_2$. We will define f by sending the terminal vertex of reduced paths starting at 1_{H_1} in Γ_1 to the terminal vertex of the respective paths starting at 1_{H_2} in Γ_2 with the same label. Then f will act on the edges so that the path in Γ_1 gets sent to the path in Γ_2. Let us see how this works in detail.

For each vertex u_1 in Γ_1, let p_1 be a reduced path in Γ_1 from 1_{H_1} to u_1 with label $\mu(p_1) = w$, then we claim that there is a path p_2 in Γ_2 with the same label w. Indeed, we have that $\Gamma_1 = \mathrm{Core}(\Gamma_1, 1_{H_1})$, so p_1 is a sub-path of some reduced circuit q_1 with respect to 1_H, and thus, the label of q_1 is a reduced word $w \cdot s \in L(\Gamma_1, 1_{H_1}) = H_1$ for some reduced word $s \in F(X)$. Then since $H_1 \leq H_2$, there must be a reduced circuit q_2 in Γ_2 with the same label $w \cdot s$, and its first $|p_1|$ edges define the unique path p_2 that we wanted.

Denoting the end vertex of the path p_2 by u_2, we define the image of u_1 under f to be u_2. We also map the edges of p_1 to the respective edges of p_2 (i.e., the nth edge of p_1 is mapped to the nth edge of p_2). To define the image of all edges of Γ_1, we will repeat this argument with all paths from 1_{H_1} to u_1. But we must first check that f is well defined, and it is sufficient to show that it is well defined on the vertex set, i.e., u_2 must not depend on the choice of p_1. So, let p_1' be another path from 1_{H_1} to u_1 with label $\mu(p_1') = w' \in F(X)$, and let p_2' be the correspondent path in Γ_2 with the same label, starting at 1_{H_2} and ending a some vertex u_2'. Then $p_1(p_1')^{-1}$ is a circuit with respect to 1_H, and thus

$$\overline{\mu(p_1(p_1')^{-1})} = \overline{w \cdot (w')^{-1}} \in H_1 \leq H_2. \tag{2.7}$$

Denoting by k the word $\overline{w \cdot (w')^{-1}}$, and by r be the reduced circuit in Γ_2 with respect to 1_{H_2} with label k, we have that the path $\overline{rp_2'}$ has label w and goes from 1_{H_2} to u_2'. Then since p_2 and $\overline{rp_2'}$ have the same label w, and they are reduced paths starting at 1_{H_2}, we have that $p_2 = \overline{rp_2'}$ and thus, the end vertices of the two paths are the same, i.e., $u_2 = u_2'$ as we wanted. □

Now we have the tools to prove that the subgroup graph is unique up to isomorphisms of X-graphs.

Corollary 2.3.3 (Subgroup graph)**.** *Let H be a subgroup of a free group $F(X)$. Let Γ_1 and Γ_2 be two folded core X-graphs with respect to v_1 and v_2 respectively, such that $L(\Gamma_1, v_1) = H = L(\Gamma_2, v_2)$. Then (Γ_1, v_1) is isomorphic to (Γ_2, v_2) as X-graphs.*

Proof. By Theorem 2.3.2, we know there are unique maps of X-graphs $\phi_1 \colon \Gamma_1 \to \Gamma_2$ and $\phi_2 \colon \Gamma_2 \to \Gamma_1$ such that $\phi_1(v_1) = v_2$ and $\phi_2(v_2) = v_1$. The compositions $\phi_2 \circ \phi_1$ and $\phi_1 \circ \phi_2$ are the unique maps of X-graphs from (Γ_1, v_1) to (Γ_1, v_1), therefore, we have that $\phi_2 \circ \phi_1 = 1_{\Gamma_1}$ and $\phi_2 \circ \phi_1 = 1_{\Gamma_2}$ as we wanted. □

2.4 Construction of the subgroup graph The finitely generated case

We have defined the notion of subgroup graphs, and we have proved that if it exists, it is unique. In this section we will give a construction for the finitely generated subgroup case, which will result in a finite graph. Before giving the construction for the infinitely generated case, we will work on some examples which will lead to an important proposition (see Proposition 2.6.3). In the proof of this proposition, we will find the motivation for the construction of the infinitely generated case.

We will start the procedure of constructing the subgroup graph with a specific relevant subgroup $H = \langle a^{-1}ba, a^{-2}ba^2 \rangle \leq F(a, b)$ (see Lemma 1.3.1 of Chapter 1), while we explain the process in general. Note that we will have to solve some difficulties on the way.

Let H be a finitely generated subgroup of the free group $F(X)$, so that $H = \langle h_1, \ldots, h_n \rangle$. We first define the following X-graph Γ.

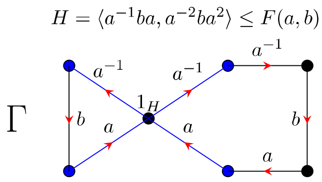

Define a vertex 1_H and then for $i = 1, \ldots, n$ we construct a disjoint circuit with respect to 1_H so that the label of the circuit is h_i (on our example we have two circuits with labels $a^{-1}ba$ and $a^{-2}ba^2$). It is clear that the language of Γ with respect to 1_H is H.

To the left we have an oriented visual representation of the X-graph Γ constructed with the circuits $a^{-1}ba$ to the left and $a^{-2}ba^2$ to the right. At first glance, we see two edges leaving the vertex 1_H with the same label a^{-1}, so the X-graph Γ obtained is not folded. Furthermore, the two edges below that are entering the vertex 1_H have label a, so their respective inverse edges (that are not shown in the representation) are also an obstruction for Γ being folded. To fix this, we need to somehow iteratively *fold* Γ without changing its language with respect to 1_H.

We say that an elementary folding consists of defining another graph $\Gamma/\!\sim\, = (V/\!\sim_V, E/\!\sim_E, [(\cdot)^{-1}], [\iota])$, where:

(i) The equivalence relation $\sim_E$ on the set of edges E is defined by:

$$e \sim_E e' \text{ and } e^{-1} \sim_E e'^{-1} \text{ if } \iota(e) = \iota(e') \text{ and } \mu(e) = \mu(e').$$

The equivalence class of an edge $e \in E$ with the relation $\sim_E$ is denoted by $[e]$.

(ii) The equivalence relation $\sim_V$ on the set of vertices V is defined by:

$$\tau(e) \sim_V \tau(e') \text{ if } e \sim_V e',$$

and we say that two related vertices $v \sim_V v'$ are *glued*.

The equivalence class of an edge $v \in V$ with the relation $\sim_V$ is denoted by $[v]$.

(iii) The inverse of an edge in the new graph $\Gamma/\!\sim$ is defined by the map

$$\begin{aligned} [(\cdot)^{-1}]\colon E/\!\sim_E &\longrightarrow E/\!\sim_E \\ [e] &\longmapsto [e]^{-1} = [e^{-1}], \end{aligned}$$

where e^{-1} is the inverse edge of e in the original graph Γ.

(iv) The initial vertex of an edge in the new graph $\Gamma/\!\sim$ is defined by the map

$$\begin{aligned} [\iota]\colon E/\!\sim_E &\longrightarrow V/\!\sim_V \\ [e] &\longmapsto [\iota](e) = [\iota(e)]. \end{aligned}$$

We leave the notation for the end vertex of an edge as τ, so $\tau([e]) = [\iota]([e]^{-1})$ is the end vertex of the edge $[e]$, not to be confused with $\tau(e)$. Do note that $\tau([e]) = [\iota]([e]^{-1}) = [\iota(e^{-1})] = [\tau(e)]$.

Below we show an example, where we apply this procedure to the X-graph Γ above.

Γ

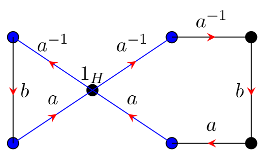

$\Gamma/\!\sim$

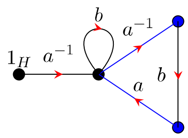

Note that this new graph $\Gamma/\!\sim$ might still not be folded (see lower graph to the left), but repeating this process must result in a folded X-graph Γ' in a finite number of steps since each elementary folding reduces the number of edges, and there is a finite number of them.

So to obtain the subgroup graph of a finitely generated subgroup as we started above, we need to see that each elementary folding does not change the language.

Lemma 2.4.1. *Let Γ be an X-graph and v one of its vertices. Then the resulting X-graph $\Gamma/\!\sim$ after an elementary folding satisfies*

$$L(\Gamma/\!\sim, [v]) = L(\Gamma, v). \tag{2.8}$$

Proof. It is straightforward to see that we can define a map of X-graphs $f\colon (\Gamma, v) \to (\Gamma/\!\sim, [v])$, sending each vertex and edge in Γ to its respective equivalence class in $\Gamma/\!\sim$. So we have that $L(\Gamma, v) \leq L(\Gamma/\!\sim, [v])$.

For the converse inclusion, we need to see that we can 'lift' any reduced circuit p' with respect to $[v]$ in $\Gamma/\!\sim$ with label w to a reduced circuit p with respect to v in Γ labelled by a word with reduced form equal to w.

We will first claim a similar statement for paths, which we will prove later.

Claim. *In the context of this lemma, for u and v vertices of Γ, we can 'lift' any reduced path p' from $[u]$ to $[v]$ in $\Gamma/\!\sim$ with label w to a reduced path p from u to v in Γ labelled by a word with reduced form equal to w.*

Then, we can "lift' any reduced circuit p' with respect to $[v]$ in $\Gamma/\sim$ to a circuit p with respect to v in Γ with the same label. So we can conclude that $L(\Gamma/\sim, [v]) \leq L(\Gamma, v)$, and thus, we have that $L(\Gamma/\sim, [v]) = L(\Gamma, v)$ as we wanted. $\square$

Now let us prove the claim.

Proof of the claim. Let us proceed by induction on the length of p'. If p' has length 0, then it is the empty path starting at $[v]$ of $\Gamma/\sim$, so we lift p' to the empty path at v.

Now consider the case of length $n > 0$, i.e., let us lift a path p' of length n from $[u]$ to $[v]$ in $\Gamma/\sim$ to a path p from u to v in Γ. By definition of $\Gamma/\sim$, we know that the path p' of $\Gamma/\sim$ is of the form

$$p' = [p_1][p_1] \ldots [p_n], \tag{2.9}$$

where $[p_i]$ is the equivalence class of some edge p_i of Γ, $[\iota(p_1)] = [u]$ and $[\tau(p_n)] = [v]$.

Then, assume by induction that we can find a path σ from u to $\iota(p_n)$ in Γ that corresponds to the subpath $[p_0] \ldots [p_{n-1}]$ from $[u]$ to $\tau([p_{n-1}]) = [\iota(p_n)]$ of p'. This means that we can compose σ and p_n. The path $\sigma\ p_n$ already has the same label as p', it starts at u and it ends at some vertex in the equivalence class $[\tau(p_n)] = [v]$, but it need to be v. Below we have the situation in the two cases: To the right, $[v]$ is of cardinality one; to the left, $[v]$ is of cardinality greater than one.

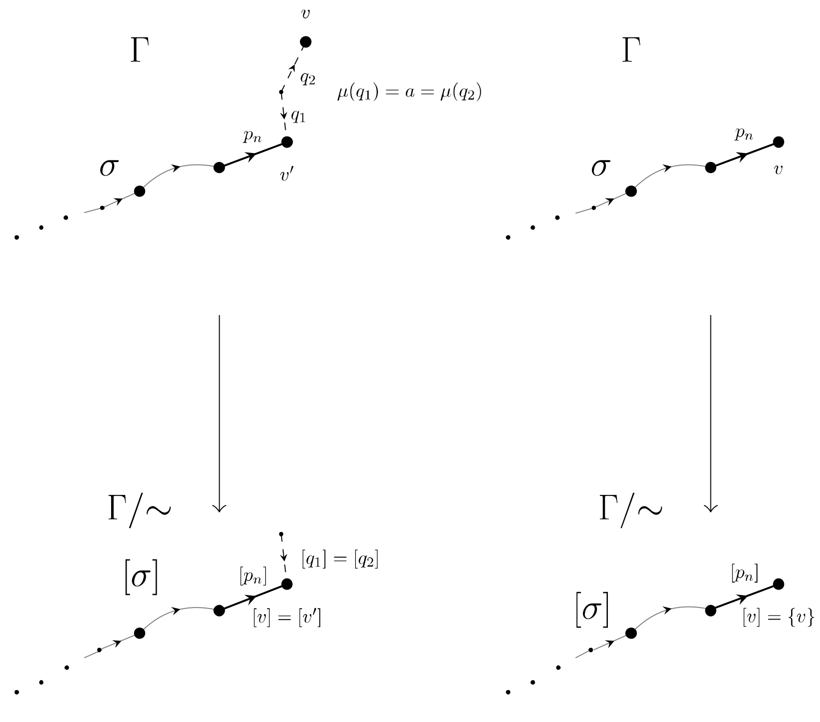

If $\tau([p_n]) = [v]$ has cardinality one, there is nothing to choose, as $\sigma\ p_n$ is already a path that ends at the only vertex in $\tau([p_n]) = \{\tau(p_n)\} = \{v\}$.

Otherwise, we have the situation on the left. If we wanted the 'lifted' path to end at the vertex v but instead $\sigma\ p_n$ ends at v' such that $[v] = [v']$, then as shown on the picture, we can join v' and v with the path $q_1^{-1}q_2$ with label $a^{-1}a$ for some $a \in X$ (the existence of this path is given by the definition of $v' \sim v$). So finally, the path $\sigma\ p_n\ q_1^{-1}q_2$ is a reduced path that we wanted: It 'lifts' p', starts at u and ends at v. □

Let us resume our search for the subgroup graph of a finitely generated subgroup $F(X)$. We continue our demonstration with the example $H = \langle a^{-1}ba, a^{-2}ba^2\rangle \leq F(a,b)$.

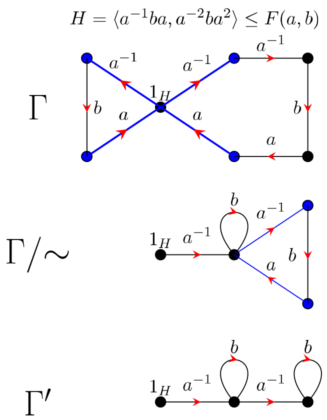

We started with the subgroup $H = \langle h_1, \ldots, h_n\rangle$, we defined the X-graph Γ with some disjoints circuits (upper graph in the picture to the left), we realized that we needed to fold some edges (blue edges), so we now finish by applying elementary foldings to Γ until we obtain a folded X-graph Γ' (lower graph to the left).

At the end we have a folded X-graph Γ' such that $L(\Gamma', 1_H) = H$, and now the question is this: Is Γ' a core X-graph with respect to 1_H?

The answer is yes. Indeed, note that we can define a surjective map of X-graphs $F\colon \Gamma \to \Gamma'$ (by surjective map of X-graph we mean that it is surjective on both vertices and edges) by sending vertices and edges to their respective class of equivalences. As we have defined Γ as a union of reduced circuits with reduced words as their respective labels and F is surjective, we have that Γ' is the union of the images under F of these circuits. Furthermore, since, by definition, a map of X-graphs respects labels, we have that Γ' is the union of circuits with reduced words as labels, and thus, these circuits are reduced themselves and Γ' is a folded core X-graph as we wanted.

We put below the conclusion as a lemma and a corollary for reference.

Lemma 2.4.2. *Let H be a finitely generated subgroup of $F(X)$, then there exists a unique up to isomorphisms finite folded core X-graph $\Gamma(H)$ with respect to a vertex 1_H such that*

$$L(\Gamma(H), 1_H) = H.$$

As an easy consequence, we have the following algorithm to solve the membership problem in a free group.

Corollary 2.4.3. *There exists an algorithm which, given a finitely generated subgroup $H = \langle h_1, \ldots, h_n \rangle$ of a free group $F(X)$ and a reduced word w on X, decides whether w is in H or not.*

Proof. Construct the finite subgroup graph $\Gamma(H)$ of H. Next, if w has reduced form

$$w = y_1 y_2 \cdots y_n, \tag{2.10}$$

with respect to X, then we try to construct a circuit with respect to 1_H with label w edge by edge.

We start at 1_H, and we look for an edge with initial vertex 1_H and label y_1. If there is no such edge, then w is not in H. If there is such an edge e_1, we repeat the search at the vertex $\tau(e_1)$ with the label y_2. We continue until finding a path p with n edges, initial vertex 1_H and label w. If this path is a circuit, then w is in H, otherwise, w is not in H. □

Let us see some examples with subgroups of $F(X)$, where $X = \{a, b\}$.

$H_1 = \langle b^2, ab^{-1}, a^2 \rangle$ $\Gamma(H_1)$

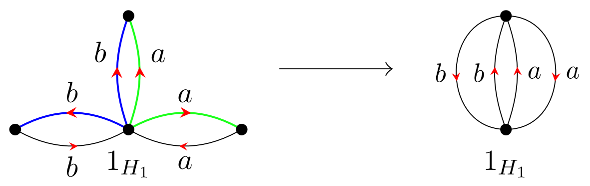

$H_2 = \langle a^2, b, aba^{-1} \rangle$ $\Gamma(H_2)$

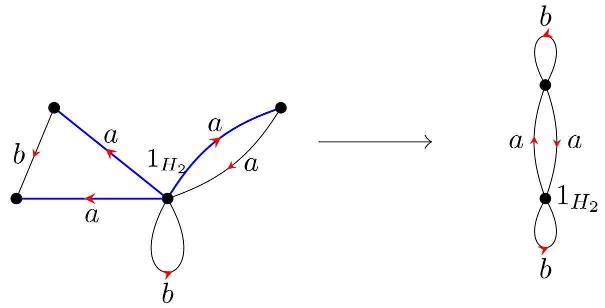

$H_3 = \langle b^2, ab \rangle$ $\Gamma(H_3)$

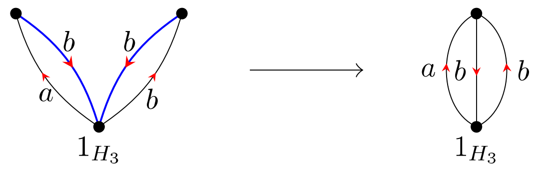

Notice that the first two subgroups H_1 and H_2 have finite index in $F(X)$. To see that, it is convenient to show first that H_1 and H_2 are normal subgroups in $F(X)$. For that, one should check this by doing the conjugates of the generators of H_1 and H_2 with the elements in $X \cup X^{-1}$ and see whether the result is in the respective subgroup or not.

But now, instead of checking if we can write the conjugates as a product of the generators, we can do it effortlessly using the subgroup graphs! To see that a given conjugate h^g is in the subgroup H_i, we just check if we can construct a reduced circuit with respect to 1_{H_i} that has label h^g.

Once we know that they are normal, one can deduce that the quotients $F(X)/H_1$ and $F(X)/H_2$ are isomorphic to the group C_2 of order 2.

Indeed, in $F(X)/H_1$, the coset $[a]$ of a is equal to the coset $[b]$ of b (since $ab^{-1} \in H_1$), which implies that $F(X)/H_1$ is generated by $[a]$. Furthermore, $[a]$ has order 2, since $a^2 \in H_1$ but $a \notin H_1$. So $F(X)/H_1 = C_2$. Similarly, in $F(X)/H_2$, the coset $[b]$ is the identity of the group, which implies that

$F(X)/H_2$ is generated by $[a]$. Furthermore, $[a]$ has order 2, since $a^2 \in H_2$ but $a \notin H_2$. So $F(X)/H_2 = C_2$.

In contrast, the third subgroup $H_3 = \langle b^2, ab \rangle$ has infinite index, as all cosets $H_3 a^n$ are different for $n = 0, 1, \ldots$, since we can not construct reduced circuits labelled by these powers in $\Gamma(H_3)$).

Now we can start seeing the following pattern. The first two subgroup graphs satisfy this condition: Adding an extra edge labelled with some symbol in $X \cup X^{-1}$ results in a non-folded X-graph. In contrast, the third subgroup graph does not satisfy it (we can add an edge labelled by a^{-1} from 1_{H_3} to the other vertex), and it happens to be of infinite index.

We will see that this condition is actually necessary and sufficient for a finitely generated subgroup to have finite index. Here is a hint of why this condition is necessary: If it is not satisfied by a given subgroup graph $\Gamma(H)$ with respect to 1_H for H a subgroup of $F(X)$, then the resulting X-graph Γ' obtained after adding a 'missing' edge (without adding a new vertex) is also folded and core with respect to 1_H. Thus, the new graph is the subgroup graph of some $H' \leq F(X)$ that strictly contains H. We are not able to see this yet, but the reason H is of infinite index in $F(X)$ is because it is of infinite index in H'.

To see this, we are interested in answering the following question: What is the relation between H and H'? We will obtain the answer in the next section.

2.5 Spanning trees, geodesic trees and free bases

Until now, we do not have a tool to obtain a generator basis of a subgroup from its X-graph. In fact, so far we have only seen that there is a subgroup associated to a folded X-graph, but we do not have any further information about the subgroup. That is why we would like to be able to obtain a free basis or a Nielsen basis from a folded X-graph.

To answer all these questions and more, we will use a concept of topology: contractible spaces. In particular, we know that if we 'contract' a contractible subspace of a given topological space, the resulting space is homotopically equivalent to the original one, and thus, their respective fundamental groups are isomorphic. This will prove to be extremely useful in X-graphs. Let us first explore what the equivalent notion to contractible subspaces are in the context of graphs.

Definition 2.5.1. Let $\Gamma = (V, E, (\cdot)^{-1}, \iota)$ be a graph, and let Δ be a subset of $E \cup V$. We say that Δ is a *subgraph* of Γ, if we can restrict $(\cdot)^{-1}$ and ι to Δ, i.e., we have that Δ is a graph with the induced structure.

We say that Δ is a *forest* if its fundamental group with respect to any vertex is trivial, which means that there are no cycles. If furthermore, Δ is connected, then we say that Δ is a tree. So a forest is a disjoint union of trees.

Remark 2.5.2. Note that in a tree T, there is a unique reduced path between any two vertices. So if we denote the unique reduced path from u to v by $[u, v]_T$, then from v to u the unique reduced path $[v, u]_T$ is equal to $[u, v]_T^{-1}$.

If we are going to contract a tree inside a graph Γ to simplify it, then we are interested in finding a *spanning tree* T inside the graph (that is a tree which contains all the vertices of the graph). This way, when contracting such spanning tree, we will end up with a bouquet of a given number of loops (maybe an infinite number), whose fundamental group is free. Furthermore, we would also like to construct spanning *geodesic* trees with respect to a vertex v, in which the unique path $[v, v']_T$ between v and any other vertex v' in the tree is a geodesic path in Γ, that is, there is no shorter path in Γ between v and v'. Geodesic trees are of interest because when contracting them, we will obtain Nielsen bases.

Note that with the notion of geodesic paths, we can define a distance in a connected graph: The distance between two vertices is n when there is a geodesic path of length n between them.

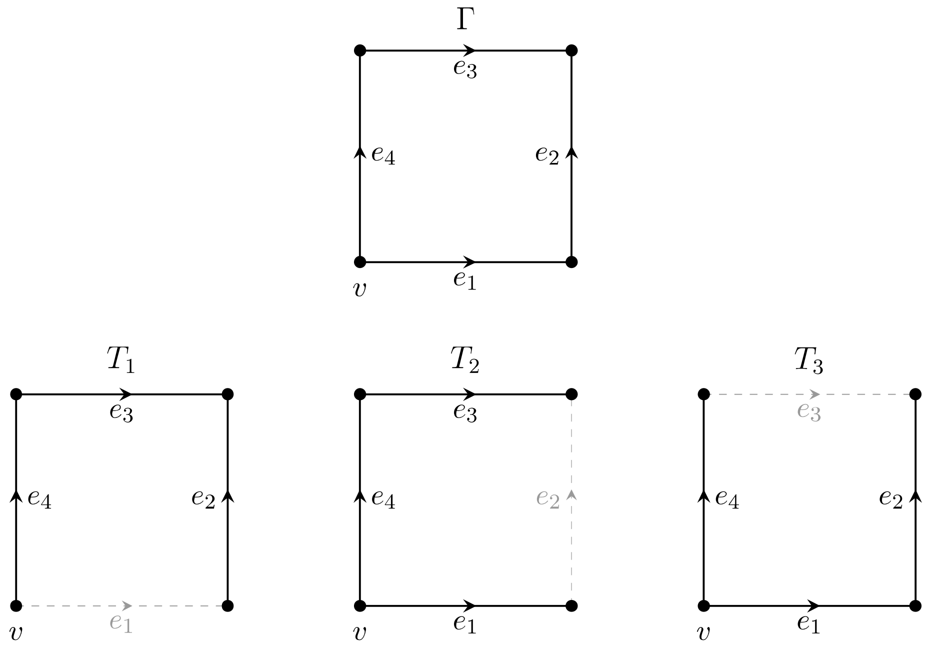

Figure 2.4: Visual representations of a graph Γ and three spanning trees T_1, T_2 and T_3 of Γ, such that T_2 and T_3 are geodesic with respect to v but T_1 is not.

Lemma 2.5.3. *Let Γ be a connected graph and let v be a vertex in Γ. Then there exists a geodesic tree with respect to v.*

Proof. Since the graph is connected, every vertex v' of Γ is at some finite distance from v, so we will construct a geodesic tree by iteratively adjoining all vertices at a given distance from v.

Start with the tree $T_0 = \{v\}$, which is clearly geodesic. Now assume that we have constructed a chain of geodesic trees

$$T_1 \subseteq T_2 \subseteq \cdots \subseteq T_{n-1}, \tag{2.11}$$

where T_i contains all vertices which are connected to v by a path of length of at most i, for $i = 1, \ldots, n-1$.

Then define T_n by adjoining each vertex v' of Γ which is not in T_{n-1}, but there exists a path p in Γ of length n from v to v'. Also add the last edge of the path p (and its inverse edge) to T_n, so that T_n is connected (there may be more that one such path of length n, but it is important to adjoin the last of edge of only one of them). It is easy to see that T_n is also a geodesic tree, and it contains all vertices which are at a distance of at most n from v.

In this manner, we obtain an infinite chain of geodesic trees (if Γ is finite, this chain stabilizes), and all vertices of Γ must be contained in one of them.

Thus, the union $\cup_{i\in\mathbb{N}} T_i$, is a geodesic tree, and it contains all vertices of Γ, and thus, it is a spanning geodesic tree. □

Remark 2.5.4. This procedure always results in a geodesic tree. But if we want arbitrary spanning trees, just construct the chain by using a non maximality argument, i.e., if T_{n-1} is not maximal, there exists a bigger tree T_n that contains it. Then the union must be maximal.

Remark 2.5.5. Note that in the case that Γ is a finite graph, we can count the number of edges in a oriented visual representation of a maximal subtree T of Γ. We started with one vertex, and then we added pairs of an edge and its terminal vertex. It is then clear that the number of vertices in T is exactly one more than the number of edges in the oriented visual representation of it. So if T is maximal, its orientation representations have $|V| - 1$ edges, where V is the set of vertices of Γ.

Now let us define rigorously what contracting a subtree of a graph is.

Definition 2.5.6. Let $\Gamma = (V, E, (\cdot)^{-1}, \iota)$ be a connected graph and $T = (V', E', (\cdot)^{-1}, \iota)$ be a subtree such that $V' \subseteq V$ and $E' \subseteq E$. Then we define the graph

$$\Gamma/T = (V \setminus V' \cup \{*\}, E \setminus E', (\cdot)^{-1}, \iota'), \tag{2.12}$$

where $(\cdot)^{-1}$ is induced from Γ. And where we define $*$ to be the new initial vertex of all those edges whose initial vertices where in V', i.e., we define the map $\iota' \colon E \setminus E' \to E \setminus E'$, such that the image of $e \in E \setminus E'$ is

$$\iota'(e) = \begin{cases} \iota(e) & \text{if } \iota(e) \notin V', \\ * & \text{if } \iota(e) \in V'. \end{cases}$$

We say that Γ/T is the quotient of Γ relative to T.

Let us see the usefulness of this construction if the fundamental group of Γ/T with respect to $*$ is isomorphic to the fundamental of Γ with respect to v.

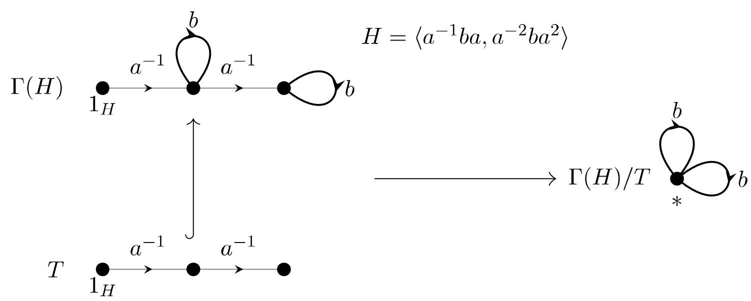

To the left we have the X-graph $\Gamma(H)$, where H denotes the subgroup $\langle a^{-1}ba, a^{-2}ba^2\rangle$ of $F(a,b)$ and the maximal subtree T. The quotient is a bouquet of two circles (not counting the reverse loops), and is no longer a folded X-graph. But we know that its fundamental group is isomorphic to F_2. Then we would have that

$$H = L(\Gamma, 1_H) \simeq \pi_1(\Gamma, 1_H) \simeq \pi_1(\Gamma/T, *) = F_2, \tag{2.13}$$

which means that H is actually free of rank 2. We already knew this from Lemma 1.3.1, but this method is neater.

Furthermore, once we understand how to construct the isomorphism $\pi_1(\Gamma, 1_H) \simeq \pi_1(\Gamma/T, *)$, we can find a free basis of H by computing the preimage of a free basis of $\pi_1(\Gamma/T, *) \simeq F_2$.

Now let us finally prove the second isomorphism of Equation (2.13).

Theorem 2.5.7. *Let $\Gamma = (V, E, (\cdot)^{-1}, \iota)$ be a connected folded X-graph, and let T be a maximal subtree in Γ. Then*

$$\pi_1(\Gamma, v) \simeq \pi_1(\Gamma/T, *) \simeq F_{|S|}, \tag{2.14}$$

where v is one of the vertices of Γ, $*$ is the only vertex in the quotient Γ/T and S is the set of edges of Γ which are not in T (we only consider one orientation of each edge to be in S).

In addition, the set $Y_T = \{\mu([v, \iota(e)]_T e[\tau(e), v]_T) \mid e \in S\}$ is a free basis of $L(\Gamma, v)$. In furthermore, the tree T is a spanning geodesic tree, Y_T is a Nielsen basis.

Proof. First note that in this context, a reduced circuit q with respect to v has the unique reduced form

$$q = p_0 e_0 p_1 e_1 \cdots p_n e_n p_{n+1}, \tag{2.15}$$

where p_i is a reduced path contained in T for $i = 0, \ldots, n+1$ and e_j are edges outside of T for $j = 0, \ldots n$.

So we can define the map

$$\begin{aligned} f\colon \pi_1(\Gamma, v) &\longrightarrow \pi_1(\Gamma/T, *) \\ q &\longmapsto f(q) = e_0 \cdots e_n. \end{aligned}$$

By remark 2.5.2, it is straightforward to see that any reduced circuit is uniquely determined by the edges outside of T where it passes through, and it cannot pass through one of these edges e and e^{-1} consequently (as it would not be reduced otherwise, since the reduced path in between these two edges is the empty one). Thus the map is injective. For surjectivity, it is enough to see that this map is a homomorphism, since the reduced circuit $[v, \iota(e)]_T e[\tau(e), v]_T$ is mapped to e for an edge e outside of T, and $\pi_1(\Gamma/T, *)$

is generated by those edges. Note that once we proof this, we will have that in fact, the set Y_T of labels of the paths of the form $[v, \iota(e)]_T e[\tau(e), v]_T$ with $e \in S$ is a free basis of $L(\Gamma, v)$.

Let q and q' be two reduced circuits in Γ written as

$$q = p_0 e_0 p_1 e_1 \cdots p_n e_n p_{n+1} \text{ and } q' = p'_0 e'_0 p'_1 e'_1 \cdots p'_m e'_m p'_{m+1}. \tag{2.16}$$

Then $f(q \cdot q')$ will have reduced form $e_0 \cdots e_i e'_j \cdots e'_m$ that represent the edges outside of T that are left after reducing the concatenation qq'. One can easily see that this means that $e_{i+1} \cdots e_n = (e'_0 \cdots e'_{j-1})^{-1}$, so we have that

$$f(q \cdot q') = e_0 \cdots e_i e'_j \cdots e'_m = \overline{e_0 \cdots e_n e'_0 \cdots e'_m} = f(q) \cdot f(q'),$$

and thus f is an isomorphism, which proves the first isomorphism of the theorem.

Finally, let us see that Y_T is a Nielsen basis when T is also a geodesic tree. We already have that Y_T generates $L(\Gamma, v)$, so we must check conditions $(N0)$, $(N1)$ and $(N2)$. We will check the Nielsen conditions with the lengths of the reduced circuits, since their label have the same length (as words).

Condition $(N0)$ is evidently satisfied since $S \cap S^{-1} = \emptyset$.

To see that condition $(N1)$ is satisfied, we first note that for each $e \in S$, we have that $[v, \iota(e)]_T e$ and $e[\tau(e), v]_T$ are paths from v to $\tau(e)$ and from $\iota(e)$ to v respectively. So since T is a geodesic tree, we have that $||[v, \iota(e)]_T| - |[\tau(e), v]_T|| \leq 1$. Then denoting by p_e the reduced path $[v, \iota(e)]_T e[\tau(e), v]_T$, we have that

$$[v, \iota(e)]_T e|, |e[\tau(e), v]_T| \geq \frac{1}{2}|p_e| = \frac{1}{2}|\mu(p_e)|_X. \tag{2.17}$$

Now let $e, f \in S \cup S^{-1}$ such that $e \neq f^{-1}$. Then, the circuit $p_e \cdot p_f$ has reduced form

$$[v, \iota(e)]_T e[\tau(e), \iota(f)]_T f[\tau(f), v]_T. \tag{2.18}$$

Now since $[v, \iota(e)]_T e$ and $f[\tau(f), v]_T$ are not cancelled, we have that no more than half of p_e and p_f are cancelled, which we know implies $(N1)$.

To check condition $(N2)$, we see that for $e,\ f,\ g \in S \cup S^{-1}$ such that $e \neq f^{-1} \neq g$, we have that the reduced form of $p_e \cdot p_f \cdot p_g$ is

$$[v, \iota(e)]_T e[\tau(e), \iota(f)]_T f[\tau(f), \iota(g)]_T g[\tau(g), v]_T. \tag{2.19}$$

Since f is not cancelled, condition $(N2)$ is satisfied. □

Remark 2.5.8. By Remark 2.5.5, if Γ is a finite graph with n vertices, we know that an orientation representation of a maximal tree T has $n-1$ edges.

Then by construction, we have that Γ/T is a bouquet of $\frac{|E|}{2} - n + 1$ loops and thus $\pi_1(\Gamma/T, *) \simeq F_{\frac{|E|}{2}-n+1}$.

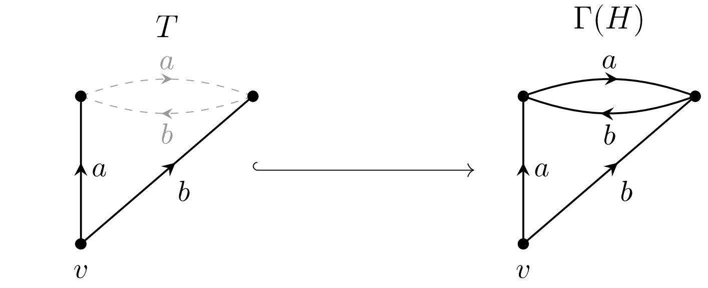

Figure 2.5: Visual representations of a subgroup graph $\Gamma(H)$ for some subgroup $H \leq F(a, b)$ and geodesic spanning tree T of $\Gamma(H)$. By Theorem 2.5.7, the set $Y_T = \{a^2b^{-1}, b^2a^{-1}\}$ is a Nielsen basis of H.

Corollary 2.5.9. *A free group of rank two contains a subgroup of infinite rank.*

Proof. The picture below shows that $H = \langle a^{-1}ba, a^{-2}ba^2, \ldots \rangle$ is free of infinite countable rank, and one can also deduce that the generators $a^{-i}ba^i$ form a free basis.

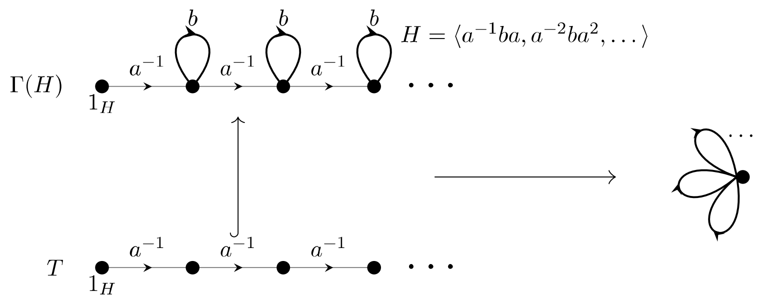

□

Now with this proof we can reprove one of the statements of the Introduction proved in Chapter 1, which results in a useful algorithm to find the rank and a free basis of a given finitely generated subgroup of a free group.

Corollary 2.5.10. *Let $F(X)$ be a free group of finite rank, and let H be a finitely generated subgroup. Then H is free. Furthermore, there is an algorithm to find a free basis.*

Proof. First, construct the finite subgroup graph $\Gamma(H)$. Secondly, choose a spanning tree T of $\Gamma(H)$. Finally, we have that H is freely generated by the labels of $[1_H, \iota(e)]_T \, e \, [\tau(e), 1_H]_T$ for $e \in E \setminus T$. □

If one goes back to the examples in the previous page, we obtain that the given generator bases were in fact free bases, as it was already know from Lemma 1.3.1 of Chapter 1. Although this was already proven in Chapter 1, we have obtained a much more elegant proof. Furthermore, we have reproven the existence of Nielsen basis for finitely generated subgroups.

Note that we are short of a construction of the subgroup graph for the infinitely generated case to proof that infinitely generated subgroups of $F(X)$ are free. And once we have this construction, we can immediately conclude that we can also find Nielsen basis for infinitely generated subgroups.

2.6 Subgroups graphs and index
A motivation for the infinitely generated case

Now recall the question we wanted to answer in the previous section: Given a X-subgroup graph $\Gamma(H)$ with respect to 1_H, to which we can add a 'missing' edge labelled with some element in $X \cup X^{-1}$ at some vertex, what is the index of H in $F(X)$?

Let $\Gamma(H')$ denote the new X-subgroup graph of a subgroup H' of $F(X)$ obtained after adding a missing edge to $\Gamma(H)$. We know that H' strictly contains H. So choosing the same spanning tree T in both graphs (since we can regard $\Gamma(H)$ as a subgraph of $\Gamma(H')$ with the same set of vertices), we obtain two bouquets of circles $\Gamma(H)/T$ and $\Gamma(H')/T$, one strictly contained in the other. This means that we can find a free basis of H' which strictly contains a free basis of H (we say that H is a *free factor* of H'). In particular, we have that

$$H = \langle h_1, \ldots, h_{n-1} \rangle \lneq \langle h_1, \ldots, h_{n-1}, h'_n \rangle = H', \tag{2.20}$$

where $h_1, \ldots, h_{n-1}, h'_n$ are some elements of $F(X)$ which form a free basis of H'. Finally, since no power of h'_n can be in H, we have that all such powers are in different cosets of H and thus, H is of infinite index in H'.

Let us conclude this section with a definition and a theorem to summarize and extend these results.

Definition 2.6.1. Let Γ be an X-graph. Then we say that Γ is X-*regular* when for each vertex u of Γ and each $y \in X \cup X^{-1}$, there is an edge e in Γ such that $\iota(e) = u$ and $\mu(e) = y$.

Remark 2.6.2. Do note that if Γ is a finite folded X-graph, and it is not X-regular, then we can add a 'missing' edge, i.e., there exist a symbol a in X and two vertices v and v' such that we can insert an edge with label a from v to v' so that the resulting graph is still folded. Indeed, denote A as

$$A = \{(u, v) \in V \times V \mid \exists\, e \in E,\ (u, v) = (\iota(e), \tau(e)),\ \mu(e) = a\}, \tag{2.21}$$

and note that the two projections of A into V are injective since Γ is folded. Then the two projections have the same cardinality, so if Γ is not X-regular, we can always insert new edges as we wanted. (Observe that this argument may fail if $|V|$ is an infinite set)

As an easy consequence of a X-graph Γ being regular and folded, we have the following: For any vertex v in Γ and any reduced word w on X, there exists a unique path in Γ with initial vertex v and label w.

Proposition 2.6.3. *Let H be a subgroup of $F(X)$, where X is a finite set. Then $|F(X)\colon H| = n < \infty$ if and only if $\Gamma(H)$ is an X-regular graph with n vertices.*

Proof. Suppose that $\Gamma(H)$ is a X-regular graph with n vertices, let us see that $|F(X)\colon H|$ equals to n.

Since what we want is to construct a correspondence between the vertices of $\Gamma(H)$ and the cosets of H, for each vertex u of $\Gamma(H)$, let us pick a reduced path p_u from 1_H to u (notice that two such paths are in the same coset of H), and denote its label by g_u (choose p_{1_H} to be the empty path with label ϵ). We claim that

$$F(X) = Hg_{1_H}\dot{\cup}Hg_{u_2}\dot{\cup}\cdots\dot{\cup}Hg_{u_n},$$

where $1_H, u_2, \ldots, u_n$ are the vertices of $\Gamma(H)$.

One must first see that any element of $F(X)$ is contained in one of these cosets. But that is clear since any word w in $F(X)$ corresponds to a unique reduced path p with label w starting at 1_H (this is possible because $\Gamma(H)$ is regular), so w is in the coset $Hg_{\tau(p)}$.

Secondly, we must see that for u, u' vertices of $\Gamma(H)$, we have that $Hg_u = Hg_{u'}$ if and only if $u = u'$. To see this, construct a path starting at u with the same label as $p_{u'}^{-1}$, i.e., with label $g_{u'}^{-1}$ (this path exists because $\Gamma(H)$ is X-regular), and denote its reverse path by $p'_{u'}$. Then since $\Gamma(H)$ is folded, we know that $p'^{-1}_{u'}$ ends at 1_H if and only if $p'_{u'} = p_{u'}$ (since $p'_{u'}$ and $p_{u'}$ would be two paths starting at 1_H with the same label). So $p_u p'^{-1}_{u'}$ starts and ends at 1_H if and only if $p'_{u'} = p_{u'}$ if and only if $u = u'$ (since the terminal vertices of $p'_{u'}$ and $p_{u'}$ are u and u' respectively). Therefore, we have that $Hg_u = Hg_{u'}$ if and only if $g_v g_{u'}^{-1} \in H$ if and only if $u = u'$, which is what we wanted.

Now for the converse, assume that $|F(X)\colon H| = n < \infty$. Then, since H is a subgroup of finite index and $F(X)$ is finitely generated, we know that H is finitely generated (by Lemma A.0.1 in the Appendix) and thus by Lemma 2.4.2, $\Gamma(H)$ is a finite graph. Now we already know that the graph must be X-regular, since otherwise, we showed that H would be a proper free factor of a subgroup H' of $F(X)$, and thus of infinite index in H' and consequently in $F(X)$.

Let us finish by seeing that the $\Gamma(H)$ has exactly n vertices. If it has n' vertices, then since we already proved that it is X-regular, the previous implication tells us that the index of H in $F(X)$ is n', which implies that $n = n'$ as we wanted. □

2.7 The infinitely generated case

Now we give a construction of $\Gamma(H)$ for any subgroup H of $F(X)$, where the vertices will be cosets of H, as the previous proposition suggested.

Let H be a subgroup of $F(X)$, then we define a X-graph Δ as follows.

The set of vertices is

$$V\Delta = \{Hg \mid g \in F(X)\}. \tag{2.22}$$

As for the set of edges; two vertices Hg and Hg' are connected by an edge labelled by $a \in X \cup X^{-1}$ if $Hga = Hg'$. In this case, we set the initial vertex of this edge to be Hg, and with the obvious reverse edge being the edge labelled by a^{-1} that starts at Hg'. Then we claim that the language of Δ with respect to the vertex H is the subgroup H.

Indeed, let $w \in F(X)$, and let q be the unique reduced path starting at the vertex H with label w (its existence is given since Δ is X-regular). Assume that the reduced form of w is

$$w = y_1 \cdots y_n, \tag{2.23}$$

where $y_i \in X \cup X^{-1}$.

Then by construction of Δ, the first vertex of the path q is H, the second vertex is Hy_1, the third Hy_1y_2 and the last is Hw. Thus we have that

$$w \in L(\Delta, H) \Leftrightarrow q \in \pi_1(\Delta, 1_H) \Leftrightarrow Hw = H \Leftrightarrow w \in H. \tag{2.24}$$

Now taking the core of this graph Δ with respect to H, we obtain a folded core graph with respect to 1_H whose language at 1_H is H, as we wanted.

Now let us finally complete Corollary 1.4.7 from Chapter 1, using the tools we have introduced in this chapter.

Theorem 2.7.1. *Let H be a subgroup of a free group $F(X)$ of finite rank. Then H is again a free group.*

Proof. Let $\Gamma(H)$ be the subgroup graph of H, and let T be a spanning tree of $\Gamma(H)$. Then by Theorem 2.5.7, we have that

$$H = L(\Gamma(H), 1_H) \simeq \pi_1(\Gamma(H), 1_H) \simeq \pi_1(\Gamma(H)/T, *) \simeq F_{|S|}, \tag{2.25}$$

where S is the set of edges of $\Gamma(H)$ which are not in T (we only consider one orientation of each edge to be in S). □

This proof shows the strength we gain with the tools developed in this chapter. We managed to easily extend Corollary 1.4.7 to any subgroup, furthermore avoiding the need to construct a Nielsen basis.

Chapter 3

The lattice of subgroups of free groups

We already have the main tools we wanted to introduce in this thesis. Now let us see what we can do with them.

Let us start with one peculiar alternative solution to Lemma A.0.1 of the Appendix, which says that finite index subgroups of finitely generated groups are finitely generated. In Chapter two, we used this exercise to prove Theorem 2.6.3, which motivated the construction of the subgroup graph in the general case. Interestingly, we can effortlessly prove this exercise using such construction. This is a good example of how a property of free groups tells us what happens in other groups.

Let $F(X)$ be a free group of finite rank $|X| = n$, and let H be a subgroup of finite index m. Then we know that its subgroup graph $\Gamma(H)$ is finite by construction, and thus by Theorem 2.5.7, we know that H is finitely generated. In fact, we can know more about the rank of H. Indeed, by Theorem 2.6.3, the subgroup graph $\Gamma(H)$ of H has exactly m vertices and it is X-regular. Thus we can count the number of edges (edges in an orientation representation): For each vertex, we have $2n$ edges leaving and entering it. But since we are only interested in counting one orientation of each pair of edges, we have $\frac{2n \cdot m}{2} = n \cdot m$ edges. So by Theorem 2.5.7, H has rank $n \cdot m - m + 1$.

Now let $G = \langle g_1, \ldots, g_n \rangle$ be an arbitrary group and let H be a subgroup of finite index m. To see that H is finitely generated we do the following: Using the universal property of free groups, we represent G as a quotient $F(x_1, \ldots, x_n)/N$. Then H corresponds to a subgroup $\tilde{H}$ of $F(x_1, \ldots, x_n)$ of index m, which we know is finitely generated (it has rank $n \cdot m - m + 1$). Thus H is finitely generated, as it is isomorphic to $\tilde{H}/N$.

Let us see some other properties of free groups. Some of which were proven in Chapter 1, but now we show a more modern proof.

3.1 Additional results on graphs and subgroups: Rank and intersections

We already know how to represent any subgroup of $F(X)$ with its respective subgroup graph. This allowed us to prove that subgroups of free groups are free; and given a finitely generated subgroup, we can easily compute its index. In this chapter we will work on answering more questions related to the lattice of free groups, such as: Given an extension of subgroups $H \leq K$, what more information can the map of the respective subgroups graph give us? Is H a free factor of K? Can we find an intermediate subgroup which is a factor of K? Or given a subgroup $H \leq F(X)$ of infinite index, can we find a bigger subgroup of finite index that contains H?

We will also prove a result on intersections of subgroups in free groups, giving also a graph representation of the intersection of two subgroups.

Let H and K be two subgroups of $F(X)$. Assume that we have a non-trivial reduced word w that is the label of the reduced paths p and q in $\Gamma(H)$ and $\Gamma(K)$ respectively. The idea is to construct a new graph which map to $\Gamma(H)$ and $\Gamma(K)$ and uniquely lifts the reduced path p and q. This can be done with the product.

Definition 3.1.1. Let $\Gamma = (V_\Gamma, E_\Gamma, \iota_\Gamma, (\cdot)^{-1}_\Gamma)$ and $\Delta = (V_\Delta, E_\Delta, \iota_\Delta, (\cdot)^{-1}_\Delta)$ be two X-graphs. We say that the *product-graph* of Γ and Δ is the graph $\Gamma \times \Delta = (V_\Gamma \times V_\Delta, E, \iota_\Gamma \times \iota_\Delta, (\cdot)^{-1}_\Gamma \times (\cdot)^{-1}_\Delta)$, such that: The set of vertices of $\Gamma \times \Delta$ is the direct product $V_\Gamma \times V_\Delta$, the set of edges E of $\Gamma \times \Delta$ is the set of pairs (e, e') in $E_\Gamma \times E_\Delta$ such that $\mu(e) = \mu(e')$, the initial vertex of an edge $(e, e') \in E \subseteq E_\Gamma \times E_\Delta$ is $(\iota_\Gamma(e), \iota_\Delta(e'))$ and the reverse of (e, e') is (e^{-1}, e'^{-1}).

If Γ and Δ are folded, one can easily prove that the product graph $\Gamma \times \Delta$ is also folded. But the product of two connected X-graphs need not be connected. Indeed, take for example a graph of two vertices connected by edges labelled by a and b, then the product of this graph by itself has 4 vertices but only a pair of edges and their inverses, so it cannot be connected.

Now let us see that the product of graphs represents the intersection of the corresponding subgroups.

Theorem 3.1.2. *Let H and K be two subgroups of $F(X)$ and let $\Gamma(H)$ and $\Gamma(K)$ be their respective subgroup graphs. Then the language of $\Gamma(H) \times \Gamma(K)$ with respect to $(1_H, 1_K)$ is $H \cap K$.*

Proof. Let us denote $L(\Gamma(H) \times \Gamma(K), (1_H, 1_K))$ by L. By construction of the product graph, we can clearly construct maps of X-graphs from $\Gamma(H) \times \Gamma(K)$ to $\Gamma(H)$ and $\Gamma(K)$ respectively, which will be the projections. This means that $L \subseteq H \cap K$. Now let w be a word contained in $H \cap K$, then we know that there are reduced circuits p and q in $\Gamma(H)$ and $\Gamma(K)$ respectively with label w. Again by construction of the product graph, we can see that we can lift p and q to a reduced circuit in $\Gamma(H) \times \Gamma(K)$ with the same label. So $L = H \cap K$ as we wanted. □

Corollary 3.1.3. *Let H and K be two finitely generated subgroups of a free group $F(X)$. Then $H \cap K$ is also finitely generated.*

Proof. By Lemma 2.4.2, the subgroup graphs $\Gamma(H)$ and $\Gamma(K)$ are finite X-graphs. Then by construction, $\Gamma(H) \times \Gamma(K)$ is also finite, and by Theorem 3.1.2 and Theorem 2.5.7, $H \cap K$ is finitely generated. □

3.2 Overgroups

Let G be a finite group, then G has a finite number of subgroups. Now let $F(X)$ be a free group of finite rank and let N be a normal subgroup of $F(X)$ such that $F(X)/N \simeq G$. Then there can only be a finite number of intermediate subgroups between N and $F(X)$, i.e., subgroups H such that $N \leq H \leq F(X)$. Let us try to understand the situation with the perspective of X-graphs.

Since $F(X)$ is finitely generated and N is of finite index, the subgroup graph $\Gamma(N)$ of N is a finite X-regular core graph with respect to 1_N. Then, an intermediate subgroup K is in correspondence to an X-regular finite core graph Δ with respect to a vertex v such that there exist a map f of X-graphs from $(\Gamma(N), 1_N)$ to (Δ, v). In this context, we have the following lemma.

Lemma 3.2.1. *The map $f\colon (\Gamma(N), 1_N) \to (\Delta, v)$ is an epimorphism.*

Proof. Since Δ is a core graph with respect to v, we have that

$$\Delta = \cup\{q \mid q \text{ a reduced circuit in } \Delta \text{ from } v \text{ to } v\}, \tag{3.1}$$

then each reduced circuit q in Δ is determined by its corresponding label μ_q in $F(X)$. Thus, we only need there to be a reduced path p_{μ_q} in $\Gamma(H)$ starting at 1_H with label μ_q (since we know that the image of p_{μ_q} will also have label μ_q, so it must be q). Now since $\Gamma(H)$ is an X-regular graph, we can always find such path. □

Now the question is this: If Γ is a finite X-graph, is there only a finite number of X-graphs (up to isomorphisms) such that there is an epimorphism from Γ to them? The answer is yes, let us see why.

Let Δ be a finite folded X-graph such that there is an epimorphism of X-graphs $f\colon \Gamma \to \Delta$, then one can define the following equivalence relation $\sim$ on Γ; Two vertices are related if their image by f is the same and two edges are related if their image is also the same. Then one can see that $\Gamma/\sim$ is a folded X-graph, and it is isomorphic to Δ. Therefore, since the number of possible equivalence relations on Γ is finite, the answer to the question above is indeed yes.

Now let us see a quick application of these quotients. For an arbitrary finitely generated subgroup H of $F(X)$, and a subgroup K which contains H, we know that there is a map f of X-graphs from $\Gamma(H)$ to $\Gamma(K)$. Then, by the previous discussion, we know that $f(\Gamma(H))$ is isomorphic to one of the quotients of $\Gamma(H)$, and the subgroup $\mu(f(\Gamma(H)))$ of $F(X)$ is a free factor of K. Given the apparent importance of these quotients, we will give them a name.

Definition 3.2.2. Let H be a subgroup of a free group $F(X)$. We say that K is an X-*principal overgroup* of H when there exist an epimorphism $f\colon (\Gamma(H), 1_H) \to (\Gamma(K), 1_K)$ of X-graphs. We call the X-*fringe* to the set of all X-principal overgroups of H, and we denote such set by $\mathcal{O}_X(H)$.

Given an arbitrary extension of subgroups $H \leq K$, and the map of X-graphs $f\colon (\Gamma(H), 1_H) \to (\Gamma(K), 1_K)$, we denote the image of $\Gamma(H)$ under f by $\Gamma_K(H)$, which is inside $\mathcal{O}_X(H)$.

Remark 3.2.3. When H is finitely generated, $\mathcal{O}_X(H)$ is finite and computable. It is sufficient to compute all quotients of $\Gamma(H)$ by all possible relations. See [3].

Now let us conclude this work by reproving the last theorem of Chapter 1, but with the tools introduced in this work.

Theorem 3.2.4 (Takahasi)**.** *Let X be a finite set, let $F(X)$ be the free group on X and let M be a positive integer. If we have the following chain of groups*

$$K_1 \leq K_2 \leq \cdots \leq K_i \leq \cdots , \tag{3.2}$$

where K_i are subgroups of $F(X)$ of rank at most M, then for some $N \in \mathbb{N}$, we have that $K_j = K_{j+1}$ for all $j \geq N$, i.e., the chain stabilizes.

Proof. Translating the problem to X-graphs, we need to check that the chain of maps of X-graphs

$$(\Gamma(K_1), 1_{K_1}) \xrightarrow{f_1} (\Gamma(K_2), 1_{K_2}) \xrightarrow{f_2} \cdots , \tag{3.3}$$

stabilizes (by stabilize we actually mean that after some term the maps are isomorphisms of X-graphs), where for $i = 1, \ldots,$ each map f_i is given by Proposition 2.3.2. Observe that for $i = 1, \ldots,$ we have that $f_{i+1} \circ f_i$ is the unique map of X-graphs from $(\Gamma(K_i), 1_{K_i})$ to $(\Gamma(K_{i+2}), 1_{K_{i+2}})$ (and similarly for longer compositions).

This means that we have the following induced chain of epimorphisms of X-graphs

$$\Gamma(K_1) \xrightarrow{f_1} \Gamma_{K_2}(K_1) \xrightarrow{f_2|_{\Gamma_{K_2}(K_1)}} \Gamma_{K_3}(K_1) \longrightarrow \cdots , \tag{3.4}$$

where each X-graph $\Gamma_{K_i}(K_1)$ is a quotient of $\Gamma(K_1)$, but also a quotient of the previous X-graph $\Gamma_{K_{i+1}}(K_1)$.

Since $\Gamma(K_1)$ is a finite graph, we can only quotient it a finite number of times. So the chain above stabilizes to some folded core X-graph Δ_1 with respect to some vertex v_1. Furthermore, by Proposition 2.3.2, we know that $K_1 \leq L(\Delta_1, v_1)$. So assuming that K_1 is not trivial, we have that $L(\Delta_1, v_1)$ is not trivial either.

This means that we can assume (by passing to a subchain) that all graphs in the chain (3.3) contain Δ_1, which is fixed by all the maps in the chain. If Δ_1 is not proper in any $\Gamma(K_i)$, for $i = 1, \dots,$ then we are done, otherwise, passing to a subchain, we can assume that it is proper in $\Gamma(K_1)$.

Now with these hypothesis, we can deduce that Δ_1 can not be proper in any of the consequent X-graphs of the chain of Equation (3.3). Indeed, take a vertex in Δ_1 with an edge leaving it which is in $\Gamma(K_1) \setminus \Delta_1$, then since all the X-graphs considered are folded, the image of this edge can not be contained in Δ_1.

This argument actually shows two things: First, that Δ_1 is proper in all the graphs as we wanted; Second, that if we repeat the same argument of taking quotients as in Equation (3.4), we will obtain a graph Δ_2, which properly contains Δ_1 and is contained in all graphs after some term in the chain. So we repeat the argument.

One can then see that this process can be repeated at most M times. Indeed, after the Mth repetition, we get for all $i = 1, \dots,$ the chain of folded core X-graphs

$$\Delta_1 \subsetneq \Delta_2 \subsetneq \cdots \subsetneq \Delta_M \subset \Gamma(K_i), \tag{3.5}$$

But in terms of subgroups of $F(X)$, this means that for all $i = 1, \dots,$ we have that

$$\{\epsilon\} < L(\Delta_1, v_1) < L(\Delta_2, v_2) < \cdots < L(\Delta_M, v_M) \leq K_i. \tag{3.6}$$

Furthermore, each of the first $M - 1$ terms of the sequence above is a proper free factor of the next term, which means that the rank of $L(\Delta_M, v_M)$ is at least M. But since $L(\Delta_M, v_M)$ is a free factor of K_i, which has rank at most M, $L(\Delta_M, v_M)$ can not be proper in K_i, and thus, we are finished. □

The End.

Appendix A

A technical result on finite index subgroups

Lemma A.0.1. *Let G be a finitely generated group and let H be a subgroup of finite index of G. Then H is finitely generated.*

Solution. Let us assume without loss of generality that G is generated by a finite set $S = \{g_1, \ldots, g_r\}$ such that $S = S^{-1}$. Let $\{Hc_1, Hc_2, \ldots, Hc_n\}$ be the set of right cosets of H (we assume that $c_1 = 1_G$). Then for $i = 1, \ldots, n$ and for $j = 1, \ldots, r$ we have that

$$c_i g_j \in H c_{s_{i_j}}, \text{ for some } s_{i_j} \in \{1, 2, \ldots, n\}. \tag{A.1}$$

Then we claim that the set $A = \{c_i g_j c_{s_{i_j}}^{-1}\}$, where i and j range from 1 to n and from 1 to r respectively, is a generating basis of H.

Indeed, let $h \in H \leq G$. Let us see that we can write h as a product of elements of A. Since $h \in G$, we can write it as

$$h = g_{l_1} \cdots g_{l_m}, \tag{A.2}$$

where $l_i \in \{1, \ldots, r\}$. Then multiplying by $c_1 = 1_G$ to the left on both sides of the equation above and applying Equation (A.1) repeatedly, we obtain

$$h = h_{l_1} \cdots h_{l_m} c, \tag{A.3}$$

where $h_{l_i} \in A$ for $i = 1, \ldots, m$ and $c \in \{c_1, \ldots, c_n\}$, which implies $c = 1_G$ since clearly we have that $c \in H$. $\square$

Bibliography

[1] John R. Stallings, Topology of finite graphs, *Inventiones mathematicae* 71 (1983) 551-565.

[2] Alexei Myasnikov, Free groups, (https://www.math.unl.edu/~mbrittenham2/classwk/990s08/public/myasnikov.1.free.groups.pdf), 2008.

[3] Alexei Myasnikov, Enric Ventura, Pascal Weil, Algebraic extensions in free groups, *Geometric group theory, Geneva and Barcelona conferences, Birkhaüser* (Ed.) (2007) 225-253. (https://arxiv.org/abs/math/0610880)

[4] Roger C. Lyndon, Paul E. Schupp, Combinatorial Group Theory, Reprint of the 1977 Edition.