Topological methods in the free group

January 24, 2018

1 Free groups

1.1 Definition

Let S be a (possibly infinite) set, which we call the "alphabet". Elements of S are called "letters". We add for each element $s \in S$ a letter $\bar{s}$, and denote $\bar{S}$ the set $\{\bar{s} \mid s \in S\}$. We take the convention that $\bar{\bar{s}} = s$.

Definition 1.1: *A word in S is a finite sequence $u_1 \dots u_n$ of elements of $S \cup \bar{S}$. A word is said to be reduced if for every $i = 1, \dots, n-1$, we have $u_i \neq \bar{u}_i$. If $w = u_1 \dots u_n, w' = v_1 \dots v_m$ are two words, their concatenation is $u_1 \dots u_n v_1 \dots v_m$. We denote $F(S)$ the set of reduced words on S.*

Remark 1.2: *The concatenation of two reduced words $w = u_1 \dots u_n, w' = v_1 \dots v_m$ is not necessarily reduced - it could be that $u_n = \bar{v_1}$. In this case, we erase the pair $u_n v_1$ to get the strictly smaller word $u_1 \dots u_{n-1} v_2 \dots v_m$. Repeat if necessary until you get a reduced word: this reduced word is denoted ww' and is called the concatenation-reduction of w with w'.*

Note that we have $ww' = u_1(\dots(u_{n-1}(u_n(v_1 \dots v_m))))$.

Lemma 1.3: *The set $F(S)$ together with the concatenation reduction operation is a group.*

Proof. The hard thing to check is associativity. We use a trick: to each element $u \in S \cup \bar{S}$, we associate the permutation L_u of $F(S)$ given by $L_u(w) = uw$. Note that $L_{\bar{u}} \circ L_u = L_u \circ L_{\bar{u}} = \mathrm{Id}$.

More generally, if $g \in F(S)$ with $g = u_1 \dots u_n$ we define $L_g = L_{u_1} \circ \dots \circ L_{u_n}$. This gives a map $L : F(S) \to \mathcal{S}(F(S))$.

This map is injective: if $v_1 \neq v_2$, then $L_{v_1}(\epsilon) = v_1 \neq v_2 = L_{v_2}(\epsilon)$ so $L_{v_1} \neq L_{v_2}$.

Now if $u \in S \cup \bar{S}$, it is easy to check that $L_u \circ L_g = L_{ug}$ (if $u_1 \neq \bar{u}$ this is immediate, if $u_1 = \bar{u}$ we get $ug = u_2 \dots u_n$ and $L_u \circ L_g = L_u \circ L_{\bar{u}} \circ L_{u_2} \circ \dots \circ L_{u_n}$ - since $L_{\bar{u}} = L_u^{-1}$ we are done).

More generally if $g' \in F(S)$ with $g' = v_1 \dots v_m$ we have $L_{g'g} = L_{g'} \circ L_g$. Indeed, $gg = v_1(\dots(v_{m-1}(v_m g)))$ so

$$L_{gg'} = L_{v_1} \circ (\dots \circ (L_{v_{m-1}} \circ (L_{v_m} \circ L_g))) = L_{g'} \circ L_g$$

Since $\circ$ is associative in $\mathcal{S}(F(S))$, the identity $L_{g'g} = L_{g'} \circ L_g$ immediately gives us associativity of concatenation-reduction in $F(S)$.

The neutral element for concatenation-reduction is the empty sequence e. The inverse of $g = u_1 \dots u_n$ is $\bar{u_n} \dots \bar{u_1}$. This finishes the proof. $\square$

Notation: Sometimes instead of $\bar{s}$ we write s^{-1} and instead of $\bar{S}$ we write S^{-1}.

Definition 1.4: *We call $F(S)$ the free group on S.*

Lemma 1.5: *(Universal property) Let S be a set, and $F(S)$ the free group on S. For any group G, and any choice of elements $(g_s)_{s \in S}$, there is a unique morphism $h : F(S) \to G$ such that $h(s) = g_s$ for all $s \in S$.*

Proof. If h is to be a morphism, we must have $h(\bar{s}) = h(s^{-1}) = h(s)^{-1} = g_s^{-1}$. We denote g_s^{-1} by $g_{\bar{s}}$. Similarly, for any reduced word $s_1 \dots s_n$, we must have $h(s_1 \dots s_n) = g_{s_1} \cdot \dots \cdot g_{s_n}$. This proves uniqueness of h.

Let us now see that it is a morphism: let $w = s_1 \dots s_n$ and $w' = t_1 \dots t_m$ be reduced words. Suppose that j is such that $ww' = s_1 \dots s_{n-j} t_{j+1} \dots t_m$ - in particular, for each $0 \leq i < j$ we have $s_{n-i} = \bar{t}_{i+1}$. Then by definition $h(ww') = g_{s_1} \cdot \dots \cdot g_{s_{n-j}} \cdot g_{t_{j+1}} \cdot \dots \cdot g_{t_m}$. On the other hand,

$$h(w) = g_{s_1} \cdot \dots \cdot g_{s_n} = g_{s_1} \cdot \dots \cdot g_{s_{n-j}} \cdot g_{s_{n-j+1}} \cdot \dots \cdot g_{s_{n-j}}$$

$$h(w') = g_{t_1} \cdot \dots \cdot g_{t_m} = g_{\bar{s}_n} \cdot \dots \cdot g_{\bar{s}_{n-j+1}} \cdot g_{t_{j+1}} \cdot \dots \cdot g_{t_m}$$

so we get equality. □

Definition 1.6: *We say the reduced word v in S* ***represents*** *the element $h(v)$. In other words, the element represented by $v = s_1 \dots s_n$ is $g_{s_1} \cdot \dots \cdot g_{s_n}$*

We often blur the distinction between s and g_s, and denote the latter simply by s, so that the reduced word $v = s_1 \dots s_n$ (which is an element in $F(S)$) represents the element $s_1 \dots s_n$ (of G!!). We will even drop the $\cdot$ to make things even more confusing...

Remark 1.7: *The morphism h is surjective iff $\hat{S} = \{g_s \mid s \in S\}$ is a generating set for G. The morphism h is injective iff no non trivial reduced word represents the identity element.*

Definition 1.8: *If the morphism $h : F(S) \to G$ is an isomorphism, we say that G is free on $\{g_s \mid s \in S\}$.*

Remark 1.9: *G is free on $\{g_s \mid s \in S\}$ iff $\hat{S}$ generates G and no nontrivial reduced word on S represents the identity element.*

Remark 1.10: *In fact there are three possible ways to define free groups*

1. *The constructive way: given a set S, build $F(S)$ just as we did.*
2. *By the universal property: category theory tells us that in the category of groups there is a unique object which satisfies the property given in Lemma 1.5, we call it the free group.*
3. *The botanical point of view: if you find in nature a group G which admits a generating set S such that no nontrivial reduced products in the elements of S and their inverses is trivial, call it free on S.*

1.2 Bases and rank

The free group $F(S)$ is of course free on S - the morphism given by Lemma 1.5 is just the identity.

But there are many other sets $T \subseteq F(S)$ such that $F(S)$ is free on T.

Example 1.11: Let $S = \{a, b\}$. Let $T = \{\alpha, \beta\}$ and let $h : F(T) \to F(S)$ be the unique morphism such that $h(\alpha) = a$ and $h(\beta) = ba$.

By universal property of $F(S)$ there is a morphism $f : F(S) \to F(T)$ given by $a \mapsto \alpha$ and $b \mapsto \beta\alpha^{-1}$. Now $g \circ f : F(S) \to F(S)$ fixes both a and b - by uniqueness in the universal property, it must be the identity. Similarly $f \circ g = \mathrm{Id}$ so we see that f, g are in fact isomorphisms.

We conclude that $F(S)$ is free on $\{a, ab\}$.

Definition 1.12: *If G is free on a set $S \subseteq G$ we call S a basis of G.*

Remark 1.13: *By Remark 1.9, S is a basis if it generates and there is no non trivial relation between the elements - this should remind you of the definition of the basis of a vector space.*

Lemma 1.14: *Let S, S' be sets. Then $F(S)$ is isomorphic to $F(S')$ iff $|S| = |S'|$.*

Proof. Suppose $|S| = |S'|$. Take h extending $S \to S'$ bijection and h' extending $S' \to S$ inverse bijection, then $h' \circ h$ is a group morphism $F(S) \to F(S)$ which extends the identity on S: by the Lemma this must be the identity (using uniqueness), hence h and h' are isomorphisms.

Other direction: count the number of morphisms $G \to \mathbb{Z}/2$: each of the $2^{|S|}$ choices of image for S gives a unique morphism by universal property, and every morphism is obtained in this way. Thus $2^{|S|}$ morphisms. If G isomorphic to G', there are exactly as many morphisms $G' \to \mathbb{Z}/2$ as $G \to \mathbb{Z}/2$, hence $|S| = |S'|$ (for infinite cardinals this requires the generalized continuum hypothesis or remark that card of G equals that of $|S|$). ☐

Corollary 1.15: *Any two bases of G have the same cardinality.*

Definition 1.16: *The rank of a free group is the cardinality of a basis.*

Warning. Bases of free groups are not as well-behaved as bases of vector spaces.

1. Not every group admits a basis, only free groups!
2. Not every generating set contains a basis.
3. If S is a subset of G which is free (i.e. no nontrivial reduced word on S represents the identity element) but not generating, it cannot in general be extended to a basis of G.
4. A free group of rank k may have free subgroups of rank $n > k$, indeed of infinite rank!

1.3 Free factors

Definition 1.17: *Let G be a free group, and H a subgroup of G. We say H is a free factor of G if there exists a basis $S = s_1, \ldots, s_n$ of G such that $H = \langle s_1, \ldots, s_k \rangle$ for some $k \leq n$.*

Note that in this case, H is free on $\{s_1, \ldots, s_k\}$ - by definition it generates H, and no nontrivial word on the set S_H represents the identity element.

1.4 Some questions on free groups

Question 1: *Given a set $S \subseteq G$, is there an algorithm to decide whether S is a basis? or whether it can be extended to a basis?*

We will see in Corollary 5.13 that in fact any subgroup of a free group is free.

Question 2: *Let H be a finitely generated subgroup of a free group G which is generated by a finite set of elements of G.*

1. *Is there an algorithm to find a basis for H from these elements?*
2. *Is there an algorithm which given an element h decides whether $h \in H$?*
3. *Can we tell whether H is normal in G?*
4. *Can we tell whether H has finite index in G?*
5. *Can we tell whether H is a free factor in G?*

Some more questions

Question 3: *If H, K are finitely generated subgroups of F, is $H \cap K$ finitely generated? Can we bound its rank?*

2 Graphs

We recall definitions of a graph to fix notations, but we will use freely notions and results from elementary graph theory.

Definition 2.1: *A graph* Γ *is given by two sets* E, V *(edges and vertices), and two functions* $\bar{\cdot} : E \to E$ *and* $\iota : E \to V$ *such that* $\bar{\bar{e}} = e$ *and* $\bar{e} \neq e$*. The vertex* $\iota(e)$ *is called the initial vertex of* e*, we define* $\tau : E \to V$ *by* $\tau(e) = \iota(\bar{e})$*, we call* $\tau(e)$ *the terminal vertex of* e*. An orientation is a choice of subset* E^+ *of* E *such that for any* $e \in E$*,* $|E^+ \cap \{e, \bar{e}\}| = 1$*.*

Definition 2.2: *The geometric realization of a graph* Γ *is the one dimensional CW-complex* $\hat{\Gamma}$ *with* 0*-skeleton* $X^0 = V$ *and collection of* 1*-cells given by* $\{[0,1] \times \{e\} \mid e \in E^+\}$ *(where* $E+$ *is an orientation), with attaching maps* $\phi_e : \{0,1\} \to X^0$ *given by* $\phi_e(0) = \iota(e)$ *and* $\phi_e(1) = \iota(\bar{e})$*. In other words,* $\hat{\Gamma}$ *is the topological space obtained by quotienting* $V \sqcup \bigsqcup_{e \in E^+} [0,1] \times \{e\}$ *by the equivalence relation generated by*

- $(0, e) \sim \iota(e)$*;*
- $(1, e) \sim \tau(e)$

(with the quotient topology)

Example 2.3: Consider the n-arc graph A_n. Then $\hat{A}_n$ is homeomorphic to $[0, n]$. Or to $[0, 1]$.

Definition 2.4: *A maps between graphs is a pair of functions* $V \to V, E \to E$ *preserving the structure. (We usually use the same letter, say* f*, to denote both graphs).*

Remark 2.5: *Given* $f : \Gamma \to \Delta$ *a map of graphs, and* $E^+(\Delta)$ *an orientation for* Δ*, the set* $\{e \in E(\Gamma) \mid f(e) \in E^+(\Delta)\}$ *an orientation of Gamma (if* $f(e) \in E^+(\Delta)$*, then* $f(\bar{e}) = \overline{f(e)}$ *is not in* $E^+(\Delta)$*, and vice versa).*

Remark 2.6: *If* $f : \Gamma \to \Delta$ *is a map between graphs, there is a continuous map* $\hat{f} : \hat{\Gamma} \to \hat{\Delta}$ *between the geometric realization of the graphs sending the interval representing an edge* e *in* Γ *homeomorphically onto the interval representing* $f(e)$ *in* $\hat{Delta}$*.*

Unless stated otherwise, all graphs are finite.

3 Paths and fundamental groups

Definition 3.1: *Let* v *be a vertex in a graph* Γ*. A path* p *at* v *of length* $n = |p|$ *is a finite (possibly empty) sequence* $e_1, \ldots, e_n$ *in* $E(\Gamma)$ *such that* $\iota(e_1) = v$ *and* $\iota(e_{i+1}) = \tau(e_i)$ *for all* $i \in \{1, \ldots, n-1\}$*. The initial vertex of* p *is* $\iota(e_1) = v$*, its terminal vertex is* $\tau(e_n)$*. If they are the same call* p *a circuit. If the sequence is empty call* p *the constant path at* v*.*

If p, q *paths such that initial vertex of* q *is final vertex of* p *we can define the concatenation* pq *of* p *and* q*, it has length* $|pq| = |p| + |q|$*.*

Remark 3.2: *A path of length* n *can be thought of as a graph map* $A_n \to \Gamma$*.*

Definition 3.3: *Geometric realization of a path = continuous map* $[0, n] \to \hat{\Gamma}$ *realizing the graph map* $A_n \to \Gamma$*.*

Definition 3.4: *A round trip is a path of the form* $e, \bar{e}$*. If* p *contains a round trip, we can erase it to obtain a path* p' *of length* $|p'| = |p| - 2$*, we say that* p' *is obtained from* p *by elementary reduction. If* p *doesn't contain any round trip say it is reduced.*

Remark 3.5: *Any edge path can be transformed into a reduced edge path.*

Definition 3.6: *The equivalence relation on paths generated by elementary reduction is called homotopy and denoted by* $\simeq$*.*

Remark 3.7: *Every path is homotopic to a reduced path.*

Exercise 3.8: If two paths are homotopic in the combinatorial sense above, any two geometric realizations of the paths are homotopic relative endpoints in the usual topological sense.

Lemma 3.9: *Concatenation of paths is compatible with homotopy, that is, if $p \simeq p'$ and $q \simeq q'$ then $pq \sim p'q'$.*

Proof. Immediate - if $p = e_1 \dots e_n$ elementary reduces to $p' = e_1 \dots e_{i-1} e_{i+2} \dots e_n$, then $pq = e_1 \dots e_n f_1 \dots f_m$ elementary reduces to $p'q$. □

Denote by $\pi(\Gamma)$ the set of homotopy equivalence of paths in Γ, see concatenation as a (partially defined) operation on $\pi(\Gamma)$ with neutral elements the constant paths.

Lemma 3.10: *This operation is associative (concatenation of paths is associative). For any element $[p] \in \pi(\Gamma)$ where $p = e_1 \dots e_n$, there is an inverse $[p]^{-1} = [\bar{p}]$ where $\bar{p}$ is the path $\bar{e}_n \dots \bar{e}_1$.*

Definition 3.11: *The subset $\pi_1(\Gamma, v)$ of $\pi(\Gamma)$ of homotopy classes of paths starting and ending at a fixed vertex v forms a group under this operation, called the (combinatorial) fundamental group of Γ based at v.*

Example 3.12:

1. If Γ is a tree, there is a unique reduced path between any two given points (standard result of graph theory). In particular, there is a unique reduced path between v and itself - the constant path. This implies that every circuit is homotopic to the constant path at v, so $\pi_1(\Gamma)$ is trivial.

2. If $\Gamma = C_n$ is the cycle on n edges, any reduced circuit is homotopic to a power of the circuit $p = e_1 \dots e_n$. Two different powers are not homotopic - this is harder to prove formally, but we will soon prove an even stronger result - see Theorem 3.16. Thus $\pi_1(C_n) \simeq \mathbb{Z}$.

Lemma 3.13: *If $f : \Gamma \to \Delta$ is a graph map, the map $f_* : \pi_1(\Gamma, v) \to \pi_1(\Delta, f(v))$ given by $f([e_1 \dots e_n]) = [f(e_1) \dots f(e_n)]$ is well defined, and it is a group homomorphism.*

Proof. If $f : \Gamma \to \Delta$ is a graph map, the images by f of two homotopic paths are homotopic, so f_* is well defined. To se it is a morphism is immediate. □

Remark 3.14: *It is possible to show that the combinatorial fundamental group we defined is isomorphic to the topological fundamental group of the geometric realization via the obvious map.*

It requires some covering theory though, and we haven't talked about that yet. It also requires to show that any topological path is homotopy equivalent to the geometric realization of a path in the sense above.

Remark 3.15: *It is a standard result from graph theory that any connected graph admits maximal subtrees (via Zorn's lemma). A maximal subtree contains all the vertices of the graph.*

Proposition 3.16: *Let Γ be a connected graph, let T be a maximal subtree in Γ and let v be a vertex of Γ. Let $\{e_i, \bar{e}_i\}_{i \in I}$ be the edges of $\Gamma - T$, and for each i let p_i, q_i be the unique path in T from v to $\iota(e_i), \tau(e_i)$ respectively.*

The fundamental group $\pi_1(\Gamma, v)$ is free on the set $\{[p_i e_i q_i] \mid i = 1, \dots, k\}$.

Proof. Let $\alpha_i = [p_i e_i q_i] \in \pi_1(\Gamma, v)$. We consider $\mathbb{F}(S)$ the free group on $S = \{s_i \mid i \in I\}$, and we define a homomorphism $h : \mathbb{F}(S) \to \pi_1(\Gamma, v)$ by setting $h(s_i) = \alpha_i$ (universal property of the free group). We show h is surjective: let p be a circuit in Γ based at v: we prove by induction on $|p|$ that p is homotopic to a circuit which is the concatenation of circuits of the form $p_i e_i q_i$. Write $p = p' e_l r$ or $p = p' \bar{e}_l r$ where r is a path contained entirely in T, and $l \in I$. Clearly p is homotopic to $p' \bar{p}_l p_l e_l q_l \bar{q}_l r$ (respectively $p' p_l \bar{p}_l \bar{e}_l \bar{q}_l q_l r$). By induction hypothesis, p' is homotopic to a product of circuits of the form $p_i e_i q_i, \bar{p}_i \bar{e}_i \bar{q}_i$.

Now define $f : \pi_1(\Gamma, v) \to \mathbb{F}(S)$ as follows: if p is a circuit based at v, consider the finite sequence $e_{i_1}^{\epsilon_1}, \dots, e_{i_t}^{\epsilon_t}$ of edges of $\Gamma - T$ appearing in p (where $\epsilon_i \in \{\pm 1\}$ and e_i^{-1} denotes $\bar{e}_i$). We let $f_([p]) = s_{i_1}^{\epsilon_1} \dots s_{i_t}^{\epsilon_t}$ (thought of as a product in $\mathbb{F}(S)$). To see this is well defined, suppose p is elementary homotopic to p via the removal of a subsequence $e\bar{e}$. If e is in T, clearly $f([p]) = f([p'])$, if e is not in

T then $f([p'])$ can be obtained from $f([p])$ by erasing a subproduct ss^{-1}, but this means that in the free group $f([p]) = f([p'])$. The map f clearly preserves concatenation so it is a morphism.

It is easy to see that $f \circ h = \mathrm{Id}$. Since h is surjective, this means f, h are isomorphisms, which proves the result. □

Remark 3.17: *If we know that the combinatorial fundamental group is isomorphic to the topological fundamental group, we can use the fact that the quotient map $\Gamma \to \Gamma/T$ which collapses the maximal subtree to a point is a homotopy equivalence, then show that the fundamental group of a graph with one vertex and k loops is free on the elements corresponding to the k loops (note that this is exactly the form of Γ/T).*

Corollary 3.18: *Let v be a vertex of the graph Γ. Any circuit p in Γ based at v is homotopic to a unique reduced path.*

Proof. If p is reduced, note that it is uniquely defined by the (ordered) list of the edges of $\Gamma - T$ it contains, and moreover, this list is reduced (i.e. doesn't contain $e\bar{e}$). Now if p, p' are homotopic, then $f([p]) = f([p'])$, which means that this list of edges is the same for p and p' (there is a unique reduced word representing any element of the free group). □

Example 3.19: Consider the free group $\mathbb{F}(a_1, \dots, a_k)$ on k letters. We think of it as the fundamental group of the rose graph R_k with k petals and unique vertex u.

If H is a subgroup of $\mathbb{F}(a_1, \dots, a_k)$ generated by elements $h_1, \dots, h_l$, there is a unique homomorphism from the free group $\mathbb{F}(s_1, \dots, s_l)$ on l elements to $\mathbb{F}(a_1, \dots, a_k)$ which sends s_i on h_i.

We build a graph of maps f as follows: take the graph R_l with central vertex v, and subdivide the i-th petal in $l(h_i)$ edges (where $l(h_i)$ denotes the length of h_i as a reduced word) - call Δ the resulting graph. Note that $\pi_1(\Delta, v)$ is still free of rank l - identify it with the group $\mathbb{F}(s_1, \dots, s_l)$.

Define the map $f : \Delta \to R_k$ to send the i-th petal to the reduced circuit corresponding to h_i. Then the induced morphism $f_* : \pi_1(\Delta, v) = \mathbb{F}(s_1, \dots, s_l) \to \mathbb{F}(a_1, \dots, a_k) = \pi_1(R_k, u)$ is exactly the homomorphism above.

In particular, $H = f_*(\mathbb{F}(s_1, \dots, s_l))$.

4 Immersions, covering and foldings

Definition 4.1: *Let v be a vertex in a graph Γ. The star of v in Γ is $\mathrm{St}(v, \Gamma) = \{e \in E(\Gamma) \mid \iota(e) = v\}$. A map of graphs $f : \Gamma \to \Delta$ induces a map $f_v : \mathrm{St}(v, \Gamma) \to \mathrm{St}(f(v), \Delta)$. The map f is called an immersion if for any vertex v, f_v is injective. It is called a covering if for any vertex v, $f_v : \mathrm{St}(v, \Gamma) \to \mathrm{St}(f(v), \Delta)$ is bijective.*

Example 4.2: Non injective immersion of a circle. Covering: identity, several copies, connected nontrivial example.

Remark 4.3: *The image by an immersion of a reduced path is a reduced path.*

Definition 4.4: *If $e, e' \in \mathrm{St}(v)$ for some vertex v are such that $e' \neq \bar{e}$, the pair (e, e') is called admissible. The graph $\Gamma/[e = e']$ obtained from Γ by identifying e with e', $\tau(e)$ with $\tau(e')$ and $\bar{e}$ with $\bar{e}'$ is called the folding of (e, e'). The quotient map $\Gamma \to \Gamma/[e = e']$ is called the folding map.*

Remark 4.5: *There are two kinds of foldings - if $\tau(e) \neq e'$, and if $\tau(e) = \tau(e')$.*

Remark 4.6: *If $f : \Gamma \to \Delta$ is a map of graphs, and e, e' are edges of Γ such that $\iota(e) = \iota(e')$ and $f(e) = f(e')$ then f factors through the folding map $\Gamma \to \Gamma/[e = e']$.*

Proposition 4.7: *Any graph map between finite graphs can be written as composition of foldings plus one immersion.*

Proof. Let $f : \Gamma \to \Delta$ be a graph map. If f is an immersion, we are done. If not, it identifies two edges e, e' with $\iota(e) = \iota(e')$, so it factors through the folding map $\Gamma \to \Gamma/[e = e']$. Repeat. This terminates since folding reduces the number of edges in a graph. □

Lemma 4.8: *If $f : \Gamma \to \Delta$ is a folding map, the induced morphism $f_* : \pi_1(\Gamma, u) \to \pi_1(\Delta, f(u))$ is surjective.*

Proof. Let $\Delta = \Gamma/[e = e']$ where $\iota(e) = \iota(e')$. Let $p = e_1 \dots e_m$ be a reduced circuit at $f(u)$. We consider successively the edges of p. If $e_i \neq f(e)$ and $e_i \neq \overline{f(e)}$, there is a unique edge $\tilde{e}_i$ such that $f(\tilde{e}_i) = e_i$, and if $e_i = f(e)$ or $e_i = \overline{f(e)}$ there are two possible choices of preimages.

If $\tau(e) = \tau(e')$, any choice of edges $\tilde{e}_1 \dots \tilde{e}_m$ such that $f(\tilde{e}_i) = e_i$ will form a circuit $\tilde{p}$ at u whose image by f is p, proving the result.

Thus we assume now that $\tau(e) = \tau(e')$. Let i be an index for which $e_i \in \{f(e), \overline{f(e)}\}$.

If $e_i = f(e)$, consider the edge e_{i+1} : since p is reduced, $e_{i+1} \neq \overline{f(e)}$. Hence there is a unique edge $\tilde{e}_{i+1}$ such that $f(\tilde{e}_{i+1}) = e_{i+1}$. Now choose $\tilde{e}_i \in \{e, e'\}$ such that $\iota(\tilde{e}_{i+k}) = \tau(\tilde{e}_{i+k-1})$.

If $e_i = \overline{f(e)}$, we consider $\tau(\tilde{e}_{i-1})$: it must be equal either to $\iota(\bar{e})$ or to $\iota(\bar{e}')$. In the first case we set $\tilde{e}_i = \bar{e}$, in the second case $\tilde{e}_i = \bar{e}'$.

It is easy to check that $\tilde{p} = \tilde{e}_1 \dots \tilde{e}_m$ is a circuit based at u. □

Corollary 4.9: *Let $f : \Gamma \to \Delta$ be a map between finite graphs, write it as a composition of foldings $f_1, \dots, f_r$ and an immersion $j : \Gamma' \to \Delta$. Then $f_*(\Gamma) = f_*(\Gamma')$.*

Proof. We have $f = j \circ f_r \circ \dots \circ f_1$, so $f_* = j_* \circ (f_r)_* \circ \dots \circ (f_1)_*$. Since all the $(f_i)_*$ are surjective, $\mathrm{Im}(f_*) = \mathrm{Im}(j_*)$. □

Example 4.10: In Example 3.19, where $H = \langle h_1, \dots, h_l \rangle \leq \mathbb{F}(a_1, \dots, a_k)$, the map $f : \Delta \to R_k$ can be factored through foldings as above to get an immersion f' of a graph $\Delta' = \Delta_H$ into R_k. We have $f'(\Delta_H) = f(\Delta)$ so this immersion "represents" H, in the sense that $\pi_1(f'(\Delta_H), f'(v')) = H$.

In fact, we will see in the sequel that the morphisms induced by immersions are injective on the fundamental group, so the graph Δ_H has fundamental group isomorphic to H. This will help us for example to find a basis for H.

Example 4.11: $\mathbb{F} = \mathbb{F}(a, b)$ and $H = \langle a^3 b, \bar{a} b a b, a^2 \bar{b} a \rangle$

5 Lifting lemmas

Again we present here the combinatorial theory to be self contained, though if we thought of the geometric realizations we could simply apply general covering theory.

5.1 Path and homotopy lifting

Lemma 5.1: *(Path lifting lemma) If $f : \Gamma \to \Delta$ is an immersion, v a vertex in Γ and p a path in Δ with initial vertex $f(v)$ then there exists at most one path $\tilde{p}$ in Γ with initial vertex v.*

If f is a covering, such a path always exists.

Proof. Suppose $p = e_1 \dots e_m$. By injectivity of f_v, there is at most one edge $\tilde{e}_1$ in $\mathrm{St}(v, \Gamma)$ such that $f_v(\tilde{e}_1) = e_1$. Moreover, $f(\tau(\tilde{e}_1)) = \tau(e)_1$. Now there is at most one edge in $\mathrm{St}(\tau(\tilde{e}_1))$ such that $f(\tilde{e}_2) = e_2$, and so on. The path $\tilde{p} = \tilde{}_1 \dots \tilde{e}_m$, if it exists, is the only possible lift.

If star maps are known to be also surjective, the lift is easily seen to exist. □

Lemma 5.2: *(Homotopy lifting lemma) If $f : \Gamma \to \Delta$ is an immersion, v a vertex in Γ and p, q paths in Δ with initial vertex $f(v)$. If p, q are homotopic then two lifts $\tilde{p}, \tilde{q}$ with initial vertex v are homotopic.*

Proof. We show this in the case where $q = e_1, \dots, e_n$ is obtained from $p = e_1, \dots, e_i, e, \bar{e}, e_{i+1}, \dots e_n$ by an elementary reduction which consists in erasing the round trip $e\bar{e}$ in p. (The general case then follows easily).

In $\tilde{p}$, the lift $\tilde{e}$ is followed by an edge $\tilde{e}'$ with initial vertex $\tau(\tilde{e})$, which maps by f onto $\bar{e}$. But f is a graph map, so $f(\bar{\tilde{e}}) = \overline{f(\tilde{e})} = \bar{e}$, and on the other hand $\bar{\tilde{e}}$ has initial vertex $\tau\tilde{e}$ just like e'. By injectivity of $f_{\tau(\tilde{e})}$, we must have $e' = \bar{\tilde{e}}$. Thus $\tilde{p}$ contains the round trip $\tilde{e}\bar{\tilde{e}}$.

If we erase this round-trip, it is easy to see that we get a lift of q, which by uniqueness must be $\tilde{q}$. □

Corollary 5.3: *If $f : \Delta \to \Gamma$ is an immersion the morphism $f_* : \pi_1(\Delta, u) \to \pi_1(\Gamma, f(u))$ is injective.*

Proof. Suppose $f_*([p]) = [f(p)] = [f(q)] = f_*([q])$, this means $f(p)$ and $f(q)$ are homotopic. Now that p and q are lifts of $f(p)$ and $f(q)$ at u, so by homotopy lifting, they are also homotopic - hence $[p] = [q]$. □

An immediate and quite amazing corollary is the following:

Corollary 5.4: *Any finitely generated subgroup of a finitely generated free group is free.*

Note that this is also true if we drop finite generatedness...

Proof. Let H be the subgroup of $\mathbb{F}(a_1, \dots, a_k)$ generated by $h_1, \dots, h_l$. By the previous sections, we know how to find a graph Δ_H and an immersion $f : \Delta_H \to R_k$ such that $\pi_1(f'(\Delta_H), v) = f'_*(\pi_1(\Delta_H, v)) = H$. But now we know that f_* is injective, so it gives an isomorphism between $\pi_1(\Delta_H, v)$ and H. Since fundamental groups of graphs are free, we get that H is free. □

5.2 Reading off a basis

In fact, this gives us more. We know how to find a basis for $\pi_1(\Delta_H, u)$: pick a maximal subtree, and take a circuit based at u for each pair $e, \bar{e}$ of edges which do not lie in the maximal subtree. The image of these circuits by f_* gives a basis for H.

Example 5.5: $\mathbb{F} = \mathbb{F}(a, b)$ and $H = \langle a^3b, \bar{a}bab, a^2\bar{b}a \rangle$. If we choose a spanning tree for Δ_H, we get a basis of H:

gives us a basis $\{a^3b, \bar{b}\bar{a}\bar{b}a\}$. We can choose a different spanning trees to get a different basis, for example $\{a^2\bar{b}a, \bar{a}bab\}$.

5.3 Membership problem

Recall we asked the following question:

Question 4: *Given a finite generating set $\{h_1, \dots, h_l\}$ of a subgroup $H = \langle h_1, \dots, h_l \rangle$ of $\mathbb{F}(a_1, \dots, a_k)$, and an element $w \in \mathbb{F}(a_1, \dots, a_k)$, can we decide whether $w \in H$?*

Lemma 5.6: *If $f : \Delta \to \Gamma$ is an immersion, the image in $\pi_1(\Gamma, f(u))$ of $\pi_1(\Delta, u)$ by f_* is precisely the set of homotopy classes of reduced circuits based at $f(u)$ which admit lifts in (Δ, u), and whose lifts are circuits at u.*

Proof. Consider a reduced circuit p in Γ based at $f(u)$. Then $[p]$ is in the image of f_* iff there exists a circuit $\hat{p}$ in Δ based at u such that $[f(\hat{p})] = [p]$.

Suppose there exists $\hat{p}$ such that $f(\hat{p}) = p$, then $f_*([\hat{p}]) = [p]$ so $[p]$ is in the image of f_*.

Suppose now conversely that $[p]$ is in the image of f_* - there exists a circuit $\hat{p}$ which wlog we may assume reduced, such that $f_*([\hat{p}]) = [f(\hat{p})] = [p]$. Since f is an immersion, $f(\hat{p})$ is also reduced. Since it is homotopic to the reduced circuit p, they must in fact be equal. □

Note that if we drop the assumption that p is reduced, it may be that p doesn't have a lift.

Now this gives us a way to answer the question. If Δ_H is a graph and $f : \Delta_H \to R_k$ is an immersion representing H, we see that w is in H iff the circuit in R_k corresponding to w lifts to a circuit in Δ_H.

Example 5.7: $\mathbb{F} = \mathbb{F}(a, b)$ and $H = \langle a^3b, \bar{a}bab, a^2\bar{b}a \rangle$ we had Δ_H:

Does $w = a$ belong to H? What about $w = a^2\bar{b}a$? And $w = b$?

5.4 More lifting. Existence of coverings

A way of interpreting the path lifting lemma: draw the commutative diagram with A_n. Here is a generalization.

Lemma 5.8: *(General lifting lemma) Suppose $f : \Gamma \to \Delta$ is a covering, $g : \Theta \to \Delta$ a map of graphs with Θ connected, and u, v vertices of Γ, Θ such that $g(v) = f(u)$.*

There exists $\tilde{g} : \Theta \to \Gamma$ such that $f \circ \tilde{g} = g$ iff $g_*(\pi_1(\Theta, v)) \subseteq f_*(\pi_1(\Gamma, u))$. Moreover, if $\tilde{g}$ exists, it is unique.

Proof. Define $\tilde{g}$ on a vertex w by: 1. choosing a path p from v to w (Θ connected), 2. lifting the path $g(p)$ to a path $\tilde{p}$ with initial point u (use path lifting lemma), 3. setting $\tilde{g}(w)$ to be the endpoint of this lift.

By definition $f \circ \tilde{g}(w) = g(w)$.

Note that if $\tilde{g}$ exists, then $\tilde{g}(p)$ is a lift of $g(p)$ starting at u. By uniqueness of lifts, $\tilde{g}(w)$ must be the endpoint of this unique lift - this proves uniqueness.

To see this is well defined, suppose q is another path from v to w. Then $p\bar{q}$ is a circuit in Θ. By hypothesis $g_*([p\bar{q}]) = f_*([r])$ for some circuit r based at u. Now $f(r)$ is homotopic to $g(p\bar{q})$ so by the homotopy lifting lemma r and the unique lift of $g(p\bar{q})$ are homotopic, in particular they have the same endpoints. Now by uniqueness, the lift of $g(p\bar{q})$ is of the form $\tilde{p}\hat{\bar{q}}$ where $\tilde{p}$ is the unique lift of $g(p)$ starting at u and $\hat{\bar{q}}$ is the unique lift of $\bar{q}$ starting at the terminal point of $\tilde{p}$. But by uniqueness of the lift of $g(p\bar{q})$, we get that $\hat{\bar{q}}$ ends at u. This implies that $\bar{\hat{\bar{q}}}$ is the unique lift of q starting at u, in other words, $\tilde{q}$. Hence $\tilde{q}$ and $\tilde{p}$ end at the same point, so $\tilde{g}$ is well defined.

Same for definition on edges. □

Lemma 5.9: *(Existence of universal coverings) If Δ is a connected graph, and v a vertex, there exists a covering $f : \tilde{\Delta} \to \Delta$ with $\tilde{\Delta}$ a tree.*

Example 5.10: Example of $\mathbb{F}_2$. See intuition for the general case.

Proof. We define $\tilde{\Delta}$ as follows: its vertex set is $\{[p] \mid p \text{ a path in } \Delta \text{ starting at } v\}$, its edge set is $\{([p], e) \mid e \in E(\Delta) \text{ with } \tau(p) = \iota(e)\}$ where $([p], e)$ has initial vertex $[p]$, and $\overline{([p], e)} = ([pe], \bar{e})$ (this is indeed an involution).

The covering map f is defined by $f([p]) = \tau(p)$ and $f(([p], e)) = e$. It is easy to see it is indeed a covering.

To see that $\tilde{\Delta}$ is a tree, let $([p_1], e_1), \dots, ([p_n], e_n)$ be a reduced non empty circuit in $\tilde{\Delta}$: we have $[p_{i+1}] = [p_i e_i]$ for each i, so we get $[p_j] = [p_1 e_1 \dots e_{j-1}]$. Moreover, we have $[p_n e_n] = [p_1]$, so we get $[p_1 e_1 \dots e_n] = [p_1]$. We may assume p_1 is reduced, so erasing roundtrips in $p_1 e_1 \dots e_n$ should give us p_1. Note that there are no round trips in $e_1 \dots e_n$, otherwise this gives a round trip in $([p_1], e_1), \dots, ([p_n], e_n)$. So the only possibility is that the last edge of p_1 is $\bar{e}_1$ - but this implies that $e_n = \overline{e_1}$. Write $e_1 \dots e_n = e_1 \dots e_{j-1} e_j \dots e_r \overline{e_{j-1}} \dots \overline{e_1}$ such that $p = p' \overline{e_{j-1}} \dots \overline{e_1}$, and the last edge in p' is distinct from $\overline{e_j}$. The path $p' e_j \dots e_r \overline{e_{j-1}} \dots \overline{e_1}$ is reduced, and it is homotopic to p - we must have that $e_j \dots e_r$ is empty, but this implies that there is a round trip in $e_1 \dots e_n$, a contradiction. □

In Exercise set 2, you will prove that any two such coverings are isomorphic.

Proposition 5.11: *The fundamental group $\pi_1(\Delta, u)$ acts on $\tilde{\Delta}$ by $[q] \cdot [p] = [qp]$ for $[q] \in \pi_1(\Delta, u)$ and $[p] \in V(\tilde{\Delta})$ and $[q] \cdot ([p], e) = ([qp], e)$ for $([p], e) \in E(\tilde{\Delta})$.*

Moreover, we have $f([q] \cdot [p]) = f([p])$ for all $[p] \in V(\tilde{\Delta})$ and $f([q] \cdot ([p], e)) = f(([p], e))$ for all $([p], e) \in E(\tilde{\Delta})$.

Lemma 5.12: *(Existence of coverings) If Δ is a connected graph, v a vertex and H a subgroup of $\pi_1(\Delta, v)$, there exists a covering $f : \Gamma \to \Delta$ with Γ connected with a vertex u such that $f(u) = v$, such that $f_*(\pi_1(\Gamma, u)) = H$.*

Any such two coverings are isomorphic.

The index of H in $\pi_1(\Delta, v)$ is the cardinality of $f_H^{-1}(v)$.

The proof I gave in the lecture was slightly different, I will try to update this soon.

Proof. We define an equivalence relation $\sim_H$ by saying that two points $[p], [p']$ in the universal cover $\tilde{\Delta}$ defined above are equivalent if p, p' have the same terminal point and $[p\bar{p}'] \in H$, that is, . Two edges $([p], e)$ and $([p'], e')$ are equivalent if $e = e'$ and $[p] \sim_H [p']$. Let Δ_H be the quotient of $\tilde{\Delta}$ by this equivalence relation. Note that if two vertices (respectively edges) of $\tilde{\Delta}$ are equivalent, they have the same image under $f : \tilde{\Delta} \to \Delta$, hence f factors through the quotient $g : \tilde{\Delta} \to \Delta_H$ as $f = f_H \circ g$. It is easy to se that f_H is a covering.

Recall that the basepoint u in $\tilde{\Delta}$ is the homotopy class of the constant path at v, denoted by 1_v. Thus the basepoint $g(u)$ in Δ_H is the equivalence class in $\sim_H$ of $u = [1_v]$. A circuit p in Δ based at v lifts in $\tilde{\Delta}$ to a path starting at u whose endpoint is $[p]$, hence it lifts in Δ_H to a path starting at $g(u)$ whose endpoint is $g([p])$, i.e. the equivalence class of $[p]$ under $\sim_H$. This is equal to $g(u)$ iff $[p] \in H$. The image in $\pi_1(\Delta, v)$ of $\pi_1(\Delta_H, g(u))$ is precisely the set of circuits based at v whose lifts in $(\Delta_H, g(u))$ are also circuits, thus we just proved that this image is exactly H.

The preimage of v in $\tilde{\Delta}$ is $\{[p] \mid p \text{ a circuit based at } v\}$ which is in bijection with $\pi_1(\Delta, v)$, so its preimage in $\tilde{\Delta}$ is the quotient of this set by the action of H given by $h \cdot [p] = [q_h p]$ where $h = [q_h]$. Thus it is in bijection with the set of cosets of H in $\pi_1(\Delta, v)$. □

We can now generalize Corollary 5.4

Corollary 5.13: *Any subgroup of a free group is free.*

Proof. Let $\mathbb{F}$ be a free group on a set S, and let H be a subgroup of $\mathbb{F}$. See $\mathbb{F}$ as the fundamental group of the rose R on $|S|$ petals. By the previous proposition, there exists a covering $(\Delta_H, u) \to (R, v)$ such that $f_*(\pi_1(\Delta_H, u)) = H$. By Proposition 5.3, we get that $\pi_1(\Delta_H, u)$ is isomorphic to H. But $\pi_1(\Delta_H, u)$ is the fundamental group of a graph, hence it is free. $\square$

6 Core graph associated to a finitely generated subgroup

Recall our algorithm to find a basis for a fg subgroup given by a finite set of generators.

Question 5: *Does the ultimate graph we get depend on how we choose to fold?*

We will see it doesn't. For this, we give the following definition.

Definition 6.1: *Let $Y = R_k$ be the rose on k petals, whose fundamental group we identify to $\mathbb{F}(a_1, \ldots, a_k)$. Let $H = \langle h_1, \ldots, h_l \rangle$. Let $f_H : (Y_H, v_H) \to (Y, v)$ be the covering corresponding to H. The core graph Γ_H of Y_H is the union of all the reduced circuits at v_H.*

Remark 6.2: *This is equivalent to saying that Γ_H is the union of images of finitely many reduced edge paths circuits representing generators of H.*

We want to prove that the core Γ_H is exactly the graph Δ_H we obtained by folding the rose with l petals, and that the immersion $f\,|_{\Delta_H}\colon \Delta_H \to R_k$ is exactly the restriction of the covering map $Y_H \to R_k$ to Γ_H. For this, the following lemma will be useful:

Lemma 6.3: *Let $j : \Delta \to \Gamma$ be an immersion between connected graphs. Suppose that $j_* : \pi_1(\Delta, u) \to \pi_1(\Gamma, j(u))$ is surjective. Then j must be injective.*

Proof. Suppose not. Then there exists v, w be distinct vertices of Δ such that $j(v) = j(w)$. If p is a reduced path in Δ between v and w, we may assume up to changing v, w that no two vertices on p have the same image. Now $j(p)$ is a reduced circuit in Γ, it is even cyclically reduced (its first and last edge are not inverses one of the other). If we write $p = p_1p_2$ where p_1, p_2 are reduced, note that $j(p_1)$ has a unique lift p_1' at w. Now p_2p_1' is a reduced path which is not a circuit, otherwise $\iota(p_2) = \tau(p_1')$ but $\iota(p_2) = \tau(p_1)$ so this would mean that $\overline{p_1}, \overline{p_1'}$ are two distinct lifts of $j(\bar{p_1})$ at the same point, a contradiction.

Now choose a shortest path q from u to p, suppose it meets p at some vertex x which divides p in p_1p_2. Up to replacing p by p_2p_1' as above, we may assume $x = v$. The path qp is reduced and not a circuit.

Now $j(qp)j(\bar{q})$ is a reduced circuit based at $j(u)$, so its homotopy class is an element of $\pi_1(\Gamma)$. By surjectivity of j_* there exists a circuit s based at u, which we may assume to be reduced, such that $j(s)$ is homotopic to $j(qp)j(\bar{q})$. Now both $j(s)$ (as the image of a reduced path by an immersion) and $j(qp)$ are reduced, so they are in fact equal. This is a contradiction since the lift of $j(pq)j(\bar{q})$ at u is not a circuit. $\square$

We can now prove equivalence of the core and the folded rose.

Lemma 6.4: *Let $Y = R_k$ be the rose on k petals, whose fundamental group we identify to $\mathbb{F}(a_1, \ldots, a_k)$. Let $H\langle h_1, \ldots, h_m \rangle$ and let $f : \hat{R}_l \to R_k$ be the graph map constructed in Example 4.10 such that $f_* : \mathbb{F}(b_1, \ldots, b_l) \to \mathbb{F}_k$ is the morphism defined by $f_*(b_i) = h_i$. Factor f as $j \circ f_r \circ \ldots \circ f_1$ where the f_i are foldings and j is an immersion $\Delta_H \to R_k$.*

Then Δ_H is isomorphic to the core Γ_H of the covering Y_H of R_k corresponding to H, and via this isomorphism j is simply the restriction of the covering map $Y_H \to R_k$.

Proof. Note that $j_*(\pi_1(\Delta_H)) = H$ so by the general lifting lemma, j lifts to a graph map $\tilde{j} : \Delta_H \to Y_H$. It is easy to see that $\tilde{j}$ must be an immersion. Note that $\tilde{j}(\Delta_H)$ is the core Γ_H since it is the union of lifts of reduced circuits in R_k corresponding to the generators of H. By Lemma 6.3, $\tilde{j}$ induces an isomorphism of graphs between Δ_H and Γ_H, and $j = f \circ \tilde{j}$ as required. $\square$

Corollary 6.5: *The graph Δ_H obtained does not depend on the sequence of foldings chose, nor on the generating set chosen for H to start with.*

We also show

Lemma 6.6: *$Y_H - \Gamma_H$ is a forest, each tree of which intersects Γ_H in a single vertex.*

In other words, one can construct Γ_H from Y_H by "chopping off" hanging trees.

Proof. If $Y_H - \Gamma_H$ is not a forest, one of its connected components is not a tree so it contains a reduced circuit. This reduced circuit gives a reduced circuit at v_H which is not in the core.

If one of the components intersects Γ_H in two vertices, there is a reduced path in it joining the two vertices, and this together with some paths in Γ_H to the two vertices gives a reduced circuit not contained in Γ_H, a contradiction. □

Here is one last characterization of the core:

Lemma 6.7: *The core is the largest connected finite subgraph of Y_H which contains v_H and has no valence 1 vertices (apart possibly the basepoint).*

Proof. If there are two such subgraphs, their union still satisfies the same properties, hence if a maximal one exists, it is unique. The core is connected and finite, contains v_H and has no valence 1 vertex except possibly for the basepoint (each vertex is in a reduced circuit based at v_H). To see that it is maximal, suppose by contradiction that it was contained in a bigger subgraph C with all the right properties. Then $C - \Gamma_H$ lives in $Y_H - \Gamma_H$ which is a forest each of whose tree intersects Γ_H in a single vertex, hence C can be obtained from Γ_H by attaching finite trees. But this contradicts the fact that C has no valence 1 vertices, unless $\Gamma_H = C$. □

Remark 6.8: *Let Δ_H be the core associated to a finitely generated subgroup H, and $j : \Delta_H \to R_n$ be the corresponding immersion. If the base vertex v of Δ_H has degree 1, we can remove the corresponding edge: the graph obtained has at most one vertex of degree 1. Repeat until you get a graph Δ_H^0 all of whose vertices have degree at least 2. The graph Δ_H^0 also immerses in R_n, and Δ_H is the union of Δ_H^0 with a path p joining v to a vertex u of Δ_H^0. The image of $\pi_1(\Delta_H^0, u)$ by j_* is exactly $[j(p)]H[j(p)]^{-1}$.*

Thus up to replacing H by a conjugate, we may assume that the core graph has no vertices of valence 1.

7 Finite index subgroups

Proposition 7.1: *Let H be a subgroup of $\mathbb{F}_n$. Then H has finite index in $\mathbb{F}$ iff the core Δ_H is equal to the covering space Y_H associated to H.*

Proof. We saw that the index of H in $\mathbb{F}_n$ is the cardinality of the preimage of the base vertex of R_n by the covering map $f : Y_H \to R_n$.

If $\Delta_H = Y_H$, in particular Y_H is finite, so H has finite index.

Conversely, if H has finite index, Y_H is finite ($f^{-1}(u)$ contains all the vertices of Y_H). But we know that the core Δ_H is the largest connected finite subgraph of Y_H containing the base vertex v and without valence 1 vertices (except possibly for the base vertex itself). Hence if $Y_H \neq \Delta_H$ it must have valence 1 vertices other than v. But this contradicts the fact that $f : Y_H \to R_n$ is a covering. □

Example 7.2: Consider the subgroup $H = \langle ab\bar{a}^2, a^2b\bar{a}, \bar{a}b\bar{a}, b\bar{a}^3\rangle$. We compute its core:

We see that the map $\Delta_H \to R_n$ is a covering, so H has finite index in $\mathbb{F}(a,b)$. In fact, the index is 3.

The index determines the rank of the subgroup.

Proposition 7.3: *If H has index k in $\mathbb{F}_n$, then* $\mathrm{rk}(H) = k(n-1)+1$.

Proof. We know that H is the fundamental group of the covering space Y_H which has exactly k vertices (since all the vertices are in $f^{-1}(u)$) and kn edges (each vertex has star of cardinality $2n$ since f is a covering, and each edge appears exactly twice in the union of all the stars).

A tree on k vertices has $k-1$ edges, so for any choice of maximal subtree for Y_H there are exactly $kn-(k-1) = k(n-1)+1$ edges outside of it - this implies that H has rank $k(n-1)+1$. □

We now prove

Proposition 7.4: *Fix $m \in \mathbb{N}^*$. There are finitely many subgroups of $\mathbb{F}_n$ with index m.*

Proof. Each subgroup corresponds to a finite cover $Y \to R_n$ where Y has m vertices. There are finitely many ways of drawing a covering (we need to draw mn colored edges joining the m vertices in such a way that the star of each vertex is complete). □

Corollary 7.5: *Any finite index subgroup of $\mathbb{F}_n$ contains a subgroup K which has finite index and is normal in $\mathbb{F}_n$.*

Proof. If H has index m, any conjugate gHg^{-1} also has index m (if $h_1H,\dots,h_mH$ are the cosets of H, then the $(gh_ig^{-1})gHg^{-1}$ are the cosets of gHg^{-1}). Now the intersection of all the conjugates of H is a normal subgroup. But it is in fact only a finite intersection of subgroups of finite index, hence it has itself finite index.

[Recall that if K has finite index in a group G and H is another subgroup, then $H \cap K$ has finite index in H: let $k_1,\dots,k_t$ be such that $H \subseteq \cup_i k_iK$, wlog we may assume that $k_i \in H$ for all i, so any element in H can be written as $h = k_ik$ for some $k \in K$, but then $k = k_i^{-1}h \in H$ so we have that $H \subseteq \bigcup_i k_i(K\cap H)$. In particular, if H itself has finite index, $K \cap H$ has finite index in G.] □

Recall that given a subgroup H of $\mathbb{F}_n = \mathbb{F}(a_1,\dots,a_n)$, it is not true in general that one can extend a basis of H to a basis of $\mathbb{F}_n$. (If this is the case, one says that H is a free factor of $\mathbb{F}_n$)

The following shows that we can always pass to a finite index subgroup in which this is true.

Proposition 7.6: *(Marshall Hall's theorem) Let H be a finitely generated subgroup of $\mathbb{F}_n = \mathbb{F}(a_1,\dots,a_n)$. Then there exists a finite index subgroup G of $\mathbb{F}_n$ such that $H \leq G$, and any basis of H extends to a basis of G.*

Proof. Let Δ_H be the core of H, and $j : \Delta_H \to R_n$ be the associated immersion. If it is a covering, H has finite index in $\mathbb{F}_n$ and we are done. If this is not a covering, we will add edges to Δ to extend it to a covering $\Gamma_H \to R_n$ in the following way: for each vertex v of Δ_H we need to add the "missing edges" in the star of v.

For each edge e of R_n, there is at most one edge in each star of vertex of Δ_H which is mapped to e. Moreover, the number of edges in Δ_H mapped to e and to $\bar{e}$ are the same. Suppose that some star $\mathrm{St}(v,\Delta_H)$ does not contain an edge mapped to $e \in E(R_n)$. Then there must be a vertex u of Δ_H so that $\mathrm{St}(u,\Delta_H)$ does not contain an edge mapped to $\bar{e}$. In this case, we add a pair of opposite edges $e',\bar{e}'$ to Δ_H with $\iota(e') = v$ and $\iota(\bar{e}') = u$, and we extend j by setting $j(e') = e$. We repeat until there are no missing edges, we get a finite graph Γ_H and a covering map $J : \Gamma_H \to R_n$.

Now $\pi_1(\Gamma_H, v)$ corresponds to a finite index subgroup G of $\mathbb{F}_n$. Moreover, Γ_H was obtained from Δ_H by adding edges, thus a maximal subtree of Δ_H is also a maximal subtree of Γ_H. This shows that there is a basis of H which extends to a basis of G. □

We can show something stronger:

Proposition 7.7: *Let H be a finitely generated subgroup of $\mathbb{F}_n$. Let $g_1, \ldots, g_k$ be elements of $\mathbb{F}_n$ which do not lie in H. There exists a finite index subgroup G of $\mathbb{F}_n$ which contains H and such that $g_1, \ldots, g_k$ do not lie in G.*

Proof. Let Δ_H be the core of H. Build a new graph Λ by attaching to Δ_H k arcs of lengths $|g_1|, \ldots, |g_k|$ respectively, and extend the immersion $\Delta_H \to R_n$ to a graph map $\Lambda \to R_n$ by mapping these arcs to the circuit representing the g_i's.

We factor f through a sequence of foldings until we get an immersion $j : \Lambda' \to R_n$. Note that Δ_H embeds in Λ', and that the arcs we added in Λ are not mapped to circuits in Λ'.

We extend $\Lambda \to R_n$ to a covering $\hat{\Lambda} \to R_n$ by adding single edges as in the previous proof. The fundamental group $G = \pi_1(\hat{\Lambda}, v)$ contains H as a free factor as before, and the elements g_i are not contained in G. □

Corollary 7.8: *Free groups are residually finite, i.e. for any non trivial element $g \in \mathbb{F}$ there exists a morphism θ from $\mathbb{F}$ to a finite group such that $\theta(g) \neq 1$.*

Proof. Let $H = 1$, by the previous proposition there is a finite index subgroup G of $\mathbb{F}$ which does not contain g. By Corollary 7.5 we may assume K is normal. The morphism $\mathbb{F} \to \mathbb{F}/K$ does not kill g, and $\mathbb{F}/K$ is finite. □

In fact we have proved something stronger

Corollary 7.9: *For any finitely generated subgroup H of $\mathbb{F}$, and any element g of $\mathbb{F}$ which are not in H there is a morphism θ from $\mathbb{F}$ to a finite group such that $\theta(g) \notin \theta(H)$.*

8 Homomorphisms between finitely generated free groups

Now consider a homomorphism $\phi : \mathbb{F}(b_1, \ldots, b_l) \to \mathbb{F}(a_1, \ldots, a_n)$, given by the images $\phi(b_j)$ of the generators of $\mathbb{F}_l$.

We can represent ϕ as a graph map $f : \hat{R}_l \to R_n$, where $\hat{R}_l$ is the subdivided rose. We can then factor f as $j \circ f_s \circ \ldots \circ f_1$ where the maps f_i are foldings and j is an immersion $\Delta_{\mathrm{Im}\phi} \to R_n$.

Proposition 8.1: *The morphism ϕ is injective iff there are no folds on edges with the same endpoints. It is surjective iff j is bijective.*

Proof. We have $\phi = f_* = j_* \circ (f_s)_* \circ \ldots \circ (f_1)_*$.

We saw that the morphism $(f_i)_*$ is always surjective, and in Exercise 2 you proved that it is an isomorphism iff the edges in the folded pair have distinct endpoints. Finally, we saw that the morphism associated to an immersion is always injective. This proves the first statement.

If j is an isomorphism of graphs then j_* is an isomorphism of groups, in particular it is surjective and hence so is ϕ. Conversely, if ϕ is surjective then j_* is surjective. We proved in Lemma 6.3 that if j_* is surjective then j is globally injective. To see j must also be surjective, note that if j_* is surjective then $H = \mathrm{Im}(\phi) = \mathbb{F}_n$, so the covering $Y_H \to R_n$ corresponding to H is in fact the identity. Thus Δ_H is equal to Y_H which is just R_n, this proves surjectivity of j. □

We also show

Proposition 8.2: *Let $\phi : \mathbb{F}(b_1, \ldots, b_l) \to \mathbb{F}(a_1, \ldots, a_n)$ be a morphism. There exists a basis $\beta_1, \ldots, \beta_l$ of $\mathbb{F}(b_1, \ldots, b_l)$ such that for some index k in $\{1, \ldots, l\}$ we have that $\phi \mid_{\langle \beta_1, \ldots, \beta_k \rangle}$ is injective and $\phi(\beta_j) = 1$ for $j = k+1, \ldots, n$.*

Proof. As usual, we represent ϕ by $f : \hat{R}_l \to R_n$. Note that a folding map corresponding to an admissible pair (e, e') with $\tau(e) = \tau(e')$ (folding of the second type) does not produce new admissible pairs. Hence when we write $f = j \circ f_s \circ \ldots \circ f_1$, we may assume that the foldings of the second type are all at the end of the sequence, i.e. that there exists some index t such that $f_1, \ldots, f_t$ are of the first type, and $f_{t+1}, \ldots, f_s$ are of the second type.

Consider the graph Λ obtained once all the foldings of the first type have been performed. We know that the morphism $(f_t)_* \circ \ldots \circ (f_1)_*$ is in fact an isomorphism between $\pi_1(\hat{R}_l, u)$ and $\pi_1(\Lambda, v)$. Thus, it is enough to find a basis for $\pi_1(\Lambda, v)$ with the required properties under the morphism $(j \circ f_s \circ \ldots \circ f_{t+1})_*$.

Now consider the equivalence class on edges of Λ generated by $e \sim e'$ iff e and e' are folded by one of the f_j with $j > t$. If we pick a maximal subtree Λ_0 of Λ, it contains at most one edge in every such equivalence class. It also gives us a basis $\{\gamma_{e_1}, \ldots, \gamma_{e_l}\}$ for $\pi_1(\Lambda, v)$, where each e_i is an edge of $\Lambda - \Lambda_0$. Note that if $e_i \sim e_m$, then $(j \circ f_s \circ \ldots \circ f_{t+1})_*(\gamma_{e_i}) = (j \circ f_s \circ \ldots \circ f_{t+1})_*(\gamma_{e_m})$.

Without loss of generality, we may assume that $k \leq m$ is such that $e_1, \ldots, e_k$ contains exactly one representative of each equivalence class. We set $\beta_1 = \gamma_{e_1}, \ldots, \beta_k = \gamma_{e_k}$, and for each $m > k$, if $i \leq k$ is such that $e_m \sim e_i$, we let $\beta_m = \gamma_{e_m}\gamma_{e_i}^{-1}$. Clearly $\beta_1, \ldots, \beta_l$ is still a basis for $\pi_1(\Lambda, v)$. Moreover, for each $m > k$ we have $(j \circ f_s \circ f_1)_*(\beta_m) = 1$. Finally, if Λ' is the union of Λ_0 with the edges $e_1, \ldots, e_k$ and their opposites, it is easy to see that $j \circ f_s \circ \ldots \circ f_{t+1}$ restricts to an immersion on Λ'. Hence the induced morphism $(j \circ f_s \circ \ldots \circ f_{t+1})_*$ is injective on $\pi_1(\Lambda', v) = \langle \beta_1, \ldots, \beta_k \rangle$. ☐

9 Normal subgroups

Proposition 9.1: *Let $f : Y \to X$ be a covering map between connected graphs. The action of the automorphism group of the covering is transitive on the set $f^{-1}(u)$ iff the subgroup $H = f_*(\pi_1(Y, v))$ is normal in $\pi_1(X, u)$.*

Proof. Suppose the action is transitive on $f^{-1}(u)$. Let $[p] \in \pi_1(X, u)$ and $[q] \in H$. There is a unique lift $\tilde{q}$ of q at v, and this lift is a circuit since $[q] \in H$.

Denote by y the end vertex of the unique lift of p based at v - we have $f(y) = u$. By transitivity of the action of $\mathrm{Aut}(f)$ on $f^{-1}(u)$, there is an automorphism of the covering sending v to the vertex y. This automorphism sends $\tilde{q}$ to a lift of q at y which is a circuit. We thus get that there is a lift of $[pq\bar{p}]$ at v which is a circuit, hence $[p][q][p]^{-1}$ lies in H. This shows H is normal.

For the other direction, suppose that H is normal. Then the action of $\pi_1(X, u)$ on the universal covering $\tilde{X} \to X$ factors through the covering map $h : \tilde{X} \to Y$: indeed, if $[p_1]$ and $[p_2]$ are two vertices of $\tilde{X}$ that are identified by h, there must be an element $[q]$ of H such that $[p_2] = [q] \cdot [p_1] = [qp_1]$. Let now $[p] \in \pi_1(X, u)$: we want to prove that $h([p] \cdot [p_1]) = h([p] \cdot [p_2])$. But $[p] \cdot [p_2] = [pqp_1] = [pq\bar{p}pp_1] = [pq\bar{p}] \cdot ([p] \cdot [p_1])$, and since H is normal, $[pq\bar{p}] = [p][q][p]^{-1}$ lies in H, so $[p] \cdot [p_1]$ and $[p] \cdot [p_2]$ have the same image by h.

Hence if $v' \in f^{-1}(u)$, let r be a path in Y from v to v: then $f(r)$ is a circuit based at u in X, so $[f(r)] \in \pi_1(X, u)$. Now if $\tilde{r}$ is the lift of r at $\tilde{u}$ (where $h(\tilde{u}) = v$), its endpoint u' is sent to v' by h. Moreover we have $h(\tilde{r}) = r$ thus via the action of $\pi_1(X, u)$ on $f^{-1}(u)$, the element $[f(r)]$ sends v to v'. ☐

Remark 9.2: *This means that the covering is homogeneous, i.e. it "looks the same" from every lift of the base vertex. In fact one can show that H is normal iff the automorphism group of the covering is transitive on every fiber.*

Using this we can get the following results on finitely generated normal subgroups in free groups.

Proposition 9.3: *If H is a finitely generated normal subgroup of $F = \mathbb{F}(a_1, \ldots, a_n)$, then either H has finite index, or H is trivial.*

Proof. We assume that H is non trivial. Consider the covering $f : Y_H \to R_n$ of the rose R_n corresponding to H, denote by v the basepoint of Y_H. Since H is finitely generated, we can define its core Γ_H, and it is a finite subgraph of Y_H.

We will prove that Y_H is finite, so the index of H in F is finite as desired. Let D denote the diameter of the core Γ_H. If Y_H is infinite there exists a vertex $y \in V(Y_H)$ which is at distance at least $D + 1$ from the core Γ_H (Y_H is a locally finite graph). Thus the ball of radius D around y is a tree, which is not the case of the ball of radius D around v - the group of automorphisms of the covering $f : Y_H \to R_n$ is not transitive on the inverse image $f^{-1}(u)$ of the unique vertex u of R_n - this contradicts normality of H. ☐

We also get

Proposition 9.4: *Let H be a finitely generated subgroup of $\mathbb{F}_n$. Then H is normal iff $\Delta_H \to R_n$ is a covering, and the automorphism group of the core graph Δ_H (together with colouring of edges representing the map $\Delta_H \to R_n$) is transitive on vertices.*

Proof. Suppose H is normal. In the proof of the preceding Proposition, we showed that the covering Y_H associated to H must be a finite graph. By Proposition 7.1 we get that in fact $Y_H = \Delta_H$.

Now we know that the group of automorphisms of the covering map $f : Y_H = \Delta_H \to R_n$ acts transitively on $f^{-1}(u)$, but since u is the unique vertex of R_n, all the vertices of Δ_H are in $f^{-1}(u)$.

Conversely, suppose that $\Delta_H \to R_n$ is a covering and that the group of automorphisms of the colored graph Δ_H is transitive on vertices - clearly, Δ_H is the covering associated to H, and by Proposition 9.1, we get that H is normal. □

Example 9.5: Which of the following subgroups are normal?

1. $H = \langle\rangle$ in $\mathbb{F}(a, b, c)$
2. $H = \langle\rangle$ in $\mathbb{F}(a, b)$

10 Intersection of finitely generated subgroups

Definition 10.1: *Let f_1, f_2 be maps of graphs $\Gamma_i \to \Lambda$ for $i = 1, 2$. A graph Γ_0 with maps $g_i : \Gamma_0 \to \Gamma_i$ is called the pullback of f_1, f_2 if $f_1 \circ g_1 = f_2 \circ g_2$, and for any graph Δ and any pair $h_1 : \Delta \to \Gamma_1, h_2 : \Delta \to \Gamma_2$ of graph maps for which $f_1 \circ h_1 = f_2 \circ h_2$, there is a unique graph map $\phi : \Delta \to \Gamma_0$ such that $h_i = g_i \circ \phi$.*

Remark 10.2: *If the pullback exists, it is unique up to isomorphism. Indeed, any such other pullback Γ_0' with maps h_1', h_2', we get graph maps $\phi : \Gamma_0' \to \Gamma_0$ and $\phi' : \Gamma_0 \to \Gamma$ with all the right commutations properties. But by uniqueness we get that $\phi \circ \phi' = \mathrm{Id}$.*

Lemma 10.3: *Let G_0 be given by*

- $V(G_0) = \{(v_1, v_2) \mid v_i \in V(\Gamma_i)$ *such that* $f_1(v_1) = f_2(v_2)\}$
- $E(G_0) = \{(e_1, e_2) \mid e_i \in E(\Gamma_i)$ *such that* $f_1(e_1) = f_2(e_2)\}$;
- $\iota((e_1, e_2)) = (\iota(e_1), \iota(e_2))$;
- $\overline{(e_1, e_2)} = (\bar{e}_1, \bar{e}_2)$.

is a graph, and together with the maps g_1, g_2 given by the projections to the first and second coordinates, it is the pullback for the pair f_1, f_2.

Proof. To check it is a graph is an easy exercise. Check that g_1, g_2 are graph maps - easy. Commutation of the diagram is easy. Now if $h_1 : \Delta \to \Gamma_1, h_2 : \Delta \to \Gamma_2$ also satisfy these properties, a graph map $\phi : \Delta \to G_0$ such that $h_i = g_i \circ \phi$ must be of the form $\phi(v) = (h_1(v), h_2(v))$ and $\phi(e) = (h_1(e), h_2(e))$ - this proves existence and uniqueness of ϕ. □

Example 10.4: Let $\mathbb{F} = \mathbb{F}(a, b)$ and $H_1 = \langle a^2, ba\rangle, H_2 = \langle ba, b^3 a\bar{b}a\rangle$. We first find the cores of H_1, H_2:

Then we construct the pullback

Remark 10.5: *As the example above shows, the pullback graph is not necessarily connected even if the graphs themselves are connected.*

Proposition 10.6: *If f_1, f_2 are immersions representing the subgroups H_1, H_2 of $F = \pi_1(\Lambda, u)$, then g_1, g_2 are immersions and the image by $(f_i \circ g_i)_*$ of the fundamental group of the connected component Γ_0^0 of Γ_0 containing the basepoint is the subgroup $H_1 \cap H_2$.*

Proof. Let us check that g_1, g_2 are immersions. Suppose e, e' is an admissible pair of edges of the pullback Γ_0 such that $g_1 = g_1' \circ f$ where $f : \Gamma_0 \to \Gamma_0'$ is the folding map associated to e, e'. Since $f_2 \circ g_2 = f_1 \circ g_1$, and f_2 is an immersion, we must have $g_2(e) = g_2(e')$, hence $g_2 = g_2' \circ f$. Thus we get $f_2 \circ g_2' = f_1 \circ g_1'$. By the universal property of Γ_0 there exists a map $h : \Gamma_0' \to \Gamma_0$ such that $g_i' = g_i \circ h$ thus $g_i = g_i' \circ f = g_i \circ (h \circ f)$ but by uniqueness of the nmorphism ϕ in the definition of the pullback we must have $h \circ f = \mathrm{Id}$ which proves that f is an isomorphism of graph maps - a contradiction.

Now let $a \in \pi_1(\Gamma_0^0, u_0)$: its image by the embedding $(f_i \circ g_i)_*$ lies in the image of the embedding $(f_i)_*$, i.e. in H_i for $i = 1, 2$. Conversely if $a \in H_1 \cap H_2$, it is represented by a reduced circuit α in Γ based at u of length m, and it lifts to a reduced circuit α_i of length m in Γ_i based at u_i for each one of $i = 1$ and $i = 2$. Let C_m be the circuit graph on m edges: we have immersions $j_i : C_m \to \Gamma_i$ for $i = 1, 2$. By the universal property of the pullback, we get a graph map $\phi : C_m \to \Gamma_0$ such that $h_i = g_i \circ \phi$. This shows precisely that the circuits α_1, α_2 have a common lift at u_0 in Γ_0^0, which proves that $a \in \pi_1(\Gamma_0^0, u_0)$. □

We get the following immediate corollary

Corollary 10.7: *(Howsons theorem) The intersection of two finitely generated subgroups of a free group is finitely generated.*

Proof. If Γ_1, Γ_2 are finite graphs, their pullback is also finite. □

Proposition 10.8: *(Hannah Neumann inequality) If H_1, H_2 are finitely generated subgroups of a free group $\mathbb{F}$, we have*

$$\mathrm{rk}(H_1 \cap H_2) - 1 \leq 2(\mathrm{rk}(H_1) - 1)(\mathrm{rk}(H_2) - 1)$$

Proof. First note that we can assume without loss of generality that the rank of the ambiant free group is 2 (embed $\mathbb{F}$ in $\mathbb{F}_2$). Build the core graphs Γ_1, Γ_2 associated to H_1, H_2 together with the corresponding immersions f_1, f_2. By Proposition 10.6, the core graph of $H_1 \cap H_2$ is the connected component Γ_0^0 of the pullback Γ_0 of f_1, f_2 which contains the basepoint, and there are immersions $g_i : \Gamma_0^0 \to \Gamma_i$ for $i = 1, 2$ such that $f_1 \circ g_1 = f_2 \circ g_2$.

Note that the degree of each vertex in Γ_i is at most 4. Moreover, by Remark 6.8 up to conjugating both H_1 and H_2 by some element g we may assume that the core graph of $H_1 \cap H_2$ has no vertices of valence 1. But there are immersions $g_i : \Gamma_0^0 \to \Gamma_i$ for $i = 1, 2$, so the base vertices in Γ_1, Γ_2 also have degree at least 2.

Now recall that the rank of the fundamental group of a finite graph is $m - n + 1$ where m is the number of edges and n the number of vertices. In a graph where the degree of the vertices is bounded by 4 we have that

$$m - n + 1 = (n_1 + 2n_2 + 3n_3 + 4n_4)/2 - (n_1 + n_2 + n_3 + n_4) + 1 = -n_1/2 + n_3/2 + n_4 + 1$$

where n_i is the number of vertices of degree i.

Denote by n_i^j the number of vertices of valence i in Γ_j. In the pullback Γ_0 we have exactly $n_4^0 \leq n_4^1 n_4^2$ vertices of degree 4: indeed, the map induced by f_i on star of a vertex of degree 4 is bijective, hence two vertices of degree 4 give e vertex of degree 4 of the pullback. Now for vertices of degree 3, we have

$n_3^0 \leq n_4^1 n_3^2 + n_3^1 n_4^2 + n_3^1 n_3^2$. We get

$$\begin{aligned} \mathrm{rk}(H_1 \cap H_2) - 1 &= n_3^0/2 + n_4^0 \\ &\leq \frac{1}{2}(n_4^1 n_3^2 + n_3^1 n_4^2 + n_3^1 n_3^2) + n_4^1 n_4^2 \\ &\leq 2(n_4^1 + \frac{1}{2}n_3^1)(n_4^2 + \frac{1}{2}n_3^2) \\ &\leq 2(\mathrm{rk}(H_1) - 1)(\mathrm{rk}(H_2) - 1) \end{aligned}$$

which proves the claim. □

The Hannah Neumann conjecture, which held open for a long time, stated that the factor 2 in the above inequality can be dropped. It was proved in 2011 independently by Mineyev and by Friedman.

11 Automorphisms of free groups

We consider the group $\mathrm{Aut}(\mathbb{F}_n)$ of automorphisms of the free group $\mathbb{F}_n$ on n generators.

Theorem 11.1: *(Nielsen generators for $Aut(\mathbb{F}_n)$) Consider the set of automorphisms of $\mathbb{F}_n = \mathbb{F}(a_1, \dots, a_n)$ containing all the automorphisms $\sigma : \mathbb{F}_n \to \mathbb{F}_n$ defined in one of the following way on the generators*

(i) (Permutations of basis elements) $\sigma(a_i) = a_{s(i)}$ for some permutation s of $1, \dots, n$;

(ii) (Sign changes) $\sigma(a_i) \in \{a_i, a_i^{-1}\}$ for each i;

(iii) (Change of maximal tree) there exists an index i such that $\sigma(a_i) \in \{a_i, a_i^{-1}\}$ and for $j \neq i$ we have $\sigma(a_j) \in \{a_j, a_i^{\pm 1} a_j, a_j a_i^{\pm 1}, a_i^{\pm 1} a_j a_i^{\pm 1}\}$.

This set generates $Aut(\mathbb{F}_n)$.

Remark 11.2: *It is clear that the maps given in the list are automorphisms. We call them elementary automorphisms.*

Corollary 11.3: *The following set generates* $\mathrm{Aut}(\mathbb{F}_n)$*:*

(i) (Swapping two basis elements) $\sigma(a_i) = a_j$, $\sigma(a_j) = a_i$ and $\sigma(a_k) = a_k$ for all $k \neq i, k \neq j$;

(ii) (Sign changes) $\sigma(a_i) = a^{-1}$ and $\sigma(a_k) = a_k$ for all $k \neq i, k \neq j$;

(iii) (Multiplication) $\sigma(a_i) = a_i a_j$ for some $j \neq i$, and $\sigma(a_k) = a_k$ for all $k \neq i$.

Lemma 11.4: *Let Γ be a finite graph, and fix $v \in V(\Gamma)$. Let T, T' two spanning trees, and $\alpha_1, \dots, \alpha_n$, $\beta_1, \dots, \beta_n$ respectively be bases associated to T, T'. Then the automorphism $\sigma : \pi_1(\Gamma, v) \to \pi_1(\Gamma, v)$ given by $\alpha_i \mapsto \beta_i$ is a composition of elementary automorphisms.*

Proof. Note that $\alpha_i = [p_i e_i q_i]$ where e_i is an edge not in T,p_i is the unique path in T from v to $\iota(e_i)$ and q_i is the unique path in T from $\tau(e_i)$ to v. Similarly, $\beta_i = [p'_i f_i q'_i]$.

It is enough to consider the case where T' is obtained from T by 1. adding an edge e_i which is in T' but not in T and 2. removing an edge $f_j \in T - T'$ on the cyclically reduced loop of $[p_1 e_1 q_1]$ (there must be such an edge since T' is a tree, and we take it in the orientation given by this loop). Up to pre- and postcomposing σ by an automorphism of type (i) we may moreover assume $i = j = 1$.

Note now that $\beta_1 = [p'_1 f_1 q'_1] = [p_1 e_1 q_1] = \alpha_{e_1}$. Now for $i > 1$, we look at the various possibility for the paths p'_i, q'_i relative to the edge f, we see that $\beta_i \in \{\alpha_{e_1}^{\pm 1} \alpha_{e_i}, \alpha_{e_i} \alpha_{e_1}^{\pm 1}, \alpha_{e_1}^{\pm 1} \alpha_{e_i} \alpha_{e_1}^{\mp 1}\}$. □

We need another lemma

Lemma 11.5: *Let Γ be a finite graph, and let e, e' be a pair of edges such that $\iota(e) = \iota(e')$ and $\tau(e) \neq \tau(e')$. Let $f : \Gamma \to \Delta$ be the associated folding map.*

There exists spanning trees T_Γ, T_Δ of Γ, Δ and associated bases $\alpha_1, \dots, \alpha_n$ and $\beta_1, \dots, \beta_m$ such that f_* is of one of the following forms:

- $n = m$ *and* $f_*(\alpha_i) = \beta_i$ *for all* i.

- $f_*(\alpha_1) = \beta_1$, *and for any* $i > 1$ *we have* $f_*(\alpha_i) \in \{\beta_j, \beta_i^{\pm 1}\beta_j, \beta_j\beta_i^{\pm 1}, \beta_i^{\pm 1}\beta_j\beta_i^{\pm 1}\}$

Proof. Suppose first that $\tau(e) \neq \iota(e)$, $\tau(e') \neq \iota(e')$. Then we can choose T_Γ containing both e, e', and $T_\Delta = f(T_\Gamma)$ (this is still a tree, and by surjectivity of f it is a spanning tree). Then it is easy to see that f_* is of type (i).

Suppose now that $\tau(e) \neq \iota(e)$, but $\tau(e') = \iota(e')$. Choose T_Γ containing e (but obviously not e'), and $T_\Delta = f(T_\Gamma) - \{f(e')\}$. Then we see that if we choose the bases so that α_1 corresponds to e' and β_1 to $f(e')$, f_* is of type (ii). □

We can now prove Theorem 11.1:

Proof. Let $\phi : \mathbb{F}(a_1, \ldots, a_n) \to \mathbb{F}(a_1, \ldots, a_n)$ be an automorphism, we represent it by a graph map $f : \hat{R}_n \to R_n$ where $\hat{R}_n$ is the rose subdivided according to the images of the generators a_i as usual.

The map f can be written as $j \circ f_1 \circ \ldots \circ f_r$ where the $f_k : \Gamma_k \to \Gamma_{k+1}$ are foldings and j is an immersion, and because $f_* = \phi$ is both injective and surjective, in fact j is an isomorphism and none of the f_k fold edges which have the same endpoints.

For each folding f_k, apply Lemma 11.5 to get maximal subtrees T'_k of Γ_k and T_{k+1} of Γ_{k+1}, as well as associated bases, such that $(f_k)_*$ is of type (i) or (ii) (note that in general T_i and T'_i are distinct). We identify for each k the fundamental group $\pi_1(\Gamma_k, v_k)$ with $\mathbb{F}_n$ by denoting the basis associated to T_k by $a_1, \ldots, a_n$. Denote by $a_1^k, \ldots, a_n^k$ the basis associated to T'_k.

Thus there exists i such that the morphism $(f_k)_*$ sends a'_j to a_j if $j = i$, and to one of a_j, $a_i^{\pm 1}a_j$, $a_ja_i^{\pm 1}$, $a_i^{\pm 1}a_ja_i^{\pm 1}$ otherwise. Hence if σ is the automorphism defined by $a_j \mapsto a'_j$, we get that $h = (f_k)_* \circ \sigma$ is an elementary automorphism. By Lemma 11.4, we know that σ^{-1} is a product of elementary automorphisms, hence so is $(f_k)_* = h \circ \sigma^{-1}$.

Since j is an isomorphism of graphs, j_* is at most a permutation of the basis $a_1, \ldots, a_n$. This proves the result. □

12 Whitehead algorithm

This sectionreproduces arguments from Heusener and Weidmann "A remark on Whitehead's Lemma".

Definition 12.1: *Let* $\mathbb{F}$ *be the free group on* $a_1, \ldots, a_n$, *let* $g \in \mathbb{F}$. *We say that* a *is* ***primitive*** *if it can be extended to a basis of* $\mathbb{F}$, *i.e. there is a basis* $g = g_1, \ldots, g_n$.

Question 6: *How can we recognize if a word is primitive?*

A useful tool is to look in the abelianization of the group: consider the quotient map $q : \mathbb{F} \to \mathbb{Z}^n$ whose kernel is the subgroup generated by commutators of elements. If g is primitive in $\mathbb{F}$, then $q(g), q(g_2), \ldots, q(g_n)$ must be a basis for $\mathbb{Z}^n$.

Example 12.2: Let $g = [a, b] \in \mathbb{F}(a, b)$. Then $q(g) = 1$ so g cannot be primitive. Let $g = abab^{-1}a$, its image by q is $q(a)^3$ which is not part of a basis of $\mathbb{Z}^2$, hence g is not primitive.

However some elements which are not primitive may have image which is paert of a basis of $\mathbb{Z}^n$.

We will build an even finer test for primitivity of elements.

Definition 12.3: *Let* Γ *be a graph endowed with a graph map* $\Gamma \to R_n$ *(alternatively,* Γ *is oriented and labelled by elements of* $a_1, \ldots, a_n$*).*

We say a reduced word $w = a_{i_1}^{\epsilon_1} \ldots a_{i_m}^{\epsilon_m}$ *is readable in* Γ *if there is a path* $e_1 \ldots e_m$ *in* Γ *such that the label of* e_j *is* $a_{i_j}^{\epsilon_j}$.

The Whitehead graph $Wh(\Gamma)$ *associated to* Γ *is the graph on* $2n$ *vertices* v_i^+, v_i^- *for* $1 \leq i \leq n$, *with an edge* $(v_i^{\epsilon_i}, v_j^{\epsilon_j})$ *if and only if the word* $a_i^{\epsilon_i}a_j^{-\epsilon_j}$ *is readable in* Γ.

The Whitehead graph $Wh(w)$ *associated to a word* w *is the Whitehead graph of the cycle graph* C_m *labelled by the letters of* w.

The following was originally proved by Whitehead, we present the proof given by Heusener and Weidmann.

Theorem 12.4: *If w is primitive and cyclically reduced, then $Wh(w)$ is disconnected or has a cut vertex.*

(A cut vertex in a graph is a vertex whose complement is disconnected). Note that if some conjugate of w is primitive, then w itself is primitive.

To prove the theorem, we will need some intermediate results. We first note

Remark 12.5: *If Γ and Gamma′ are two oriented graphs labelled by $a_1, \ldots, a_n$, and $h : \Gamma \to \Gamma'$ is a graph map which respects orientation and labelling, then $Wh(\Gamma) \subseteq Wh(\Gamma')$.*

If $Wh(\Gamma')$ is disconnected or has a cut vertex, then so is $Wh(\Gamma)$ - this follows from the fact that they have the same vertex sets.

To prove Theorem **??**, it is therefore enough to show that the cycle graph of a primitive element maps (by a graph map respecting labellings) into a connected graph Δ whose Whitehead graph is disconnected or admits a cut vertex (in other words, that w is readable in a graph with a cut vertex).

Definition 12.6: *A graph Δ endowed with orientation and labelling of the edges by $a_1, \ldots, a_n$ is called an almost-rose if its fundamental group has rank n, it has no valence 1 vertices, and it folds onto the rose R_n by a single fold.*

Lemma 12.7: *If w is primitive and cyclically reduced then it is readable in an almost-rose.*

Proof. Let $w = w_1, \ldots, w_n$ be a basis for $\mathbb{F}_n$ and let $\hat{R}_n$ be the subdivided rose whose petals are labelled by the reduced words $w_1, \ldots, w_n$. There is a graph map $f : \hat{R}_n \to R_n$, whose corresponding morphism f_* is the identity. We can factor f as $j \circ f_r \circ \ldots \circ f_1$ as usual, with $f_i : \Gamma_i \to \Gamma_{i+1}$ a folding map and j an isomorphism. Note that since w is cyclically reduced every vertex in $\hat{R}_n$ is adjacent to at least two edges with distinct labels, hence this is true of all the graphs appearing in the sequence of foldings and in particular they have no valence 1 vertex.

The graph Γ_r is therefore an almost rose. The word w is readable in $\hat{R}_n$, therefore it is readable in Γ_r and we are done. □

Lemma 12.8: *The Whitehead graphs of almost roses have cut vertices.*

Proof. We will in fact describe explicitly all possible almost roses: let Δ be an almost rose, and let e, e' be the edges folded to get a rose, and assume wlog that e, e' are both labelled by a_1. Since all edges of a rose are loops, at least one of e, e' is a loop. Since the rank of the fundamental group does not decrease, e and e' do not have the same endpoint. Thus we may assume e is a loop based at a vertex u and e' joins u to another vertex v. Since the rose has a unique vertex, u and v are the only vertices of Δ, and since Δ has no valence 1 vertex, there is at least one other edge adjacent to v.

Now it is easy to see that Δ has $n+1$ edges, $e_1, \ldots, e_n, e'$ where e_i is labelled by a_i, and there are indices k, l such that $1 \leq k \leq l \leq n$, $k < n$ and

- edges $e_1, \ldots, e_k$ are loops at u;
- edges $e_{k+1}, \ldots, e_l$ and e' join u and v (orientations may vary, wlog we can assume it starts from u and ends in v);
- edges $e_{l+1}, \ldots, e_n$ are loops at v.

Now we claim that v_1 is a cut vertex of the Whitehead graph of Δ. Indeed, the induced subgraphs on $\{v_1^{\pm}, \ldots, v_k^{pm}, v_{k+1}^-, \ldots, v_l^-\}$ and $\{v_1^+, v_{k+1}^+, \ldots, v_l^+, v_{l+1}^{\pm}, \ldots, v_n^{pm}\}$ are complete graphs, and their union is $Wh(\Gamma)$. Since they have only the vertex v_1^+ in common, it is a cut vertex of Δ. □

Remark 12.9: *We could have proved the more general result that if a set S with at least one cyclically reduced element can be extended to a basis then the Whitehead graph of the rose subdivided with petals labelled by the words in S has a cut vertex.*