BACHELOR'S DEGREE THESIS

Degree in Mathematics

Title: Algorithmic problems about subgroups of free groups

Author: Guillermo Santamaría

Advisor: Enric Ventura

Department: MA III

Academic year: 2012-2013

Facultat de Matemàtiques i Estadística

UNIVERSITAT POLITÈCNICA DE CATALUNYA

Algorithmic problems about subgroups of free groups

Guillermo Santamaría

June 20, 2013

The study of the lattice of subgroups of F_k changed completely when J.Stallings published the paper [1] in 1983: in that paper Stallings developed some special graphs in order to solve several algorithmic problems about subgroups like the membership or the intersection problem. The eficient and nice solution using this machinery contrasted with the complex methods developed before. Stallings constructed his theory using an algebraic topology approach, and some papers like [2] use automata theory. On the contrary, in this paper we are going to use purely graph theory: this kind of approach needs its own definitions and sometimes could be rambling but it's really useful when talking about algorithms. So, in the first chapter we develop Stallings theory from the beginning and in the second one we explain several algorithmic applications and present some examples.

1 Introduction: the basic theory

The goal of this first chapter is to introduce the key concepts concerning free groups and graphs and to prove a particular case for graphs of a topology classical theorem. It says that, if X is a topological space with some connectivity properties then, for any subgroup S of its fundamental group based at v, $S \leq \pi_1(X, v)$, there is, up to isomorphism, one and only one connected covering space with a distinguished point w, $p : (X_h, w) \to (X, v)$ such that $p(\pi_1(X_h, w)) = H$. There is a topological proof in [3]. This theorem holds for graphs and we will prove it, in fact, using a graph theory approach. It's the basic result used to develop the following chapters.

1.1 Fundamental group of a graph

We are going to consider combinatorial non-oriented multigraphs with distinguished orientations, i.e.:

A combinatorial *graph* Γ consists of two sets E and V, and two maps $^* : E \to E$ and $\iota, \tau : E \to V$ such that $(e^*)^* = e$ and $e^* \neq e$. On the other hand $\iota(e) = \tau(e^*)$ and viceversa.

If $v = \iota(e)$ or $v = \tau(e)$ we will say that the e is *incident* to v. The number of incident edges to a vertex v is called the *degree* of v. If two edges are incident to the same vertex v, we will say that those edges are *adjacent*.

An *orientation* θ of the edges of Γ is a choice of exactly one element in each pair $\{e, e^*\}$. If the chosen element is e, we will denote it by $(\iota(e), \tau(e))$ and we will say that it is oriented *forwards* with respect to the vertex $\iota(e)$, otherwise, if the choosen element is e^*, we will denote it by $(\tau(e), \iota(e))$ and will say that it is oriented *backwards* with respect to the vertex $\iota(e)$.

A *map of graphs* $f : \Gamma \to \Delta$ consists of a pair of functions, f_e from edges to edges, and f_v from vertices to vertices, preserving the structure, i.e.: $f_e(e^*) = f_e(e)^*$ and $f_v(\iota(e)) = \iota(f_e(e))$.

The *rose of k petals* R_k is the graph with one vertex and k edges with an orientation, each one labeled by a letter of $\{a_1, \ldots a_n\}$, a given alphabet. Then, if Δ is finite and have an orientation, consider a map q_k from Δ to R_k, with $k = |E(\Delta)|$, that sends the edges of Δ to the edges of R_k injectivelly and preserves orientations. Then, for a map of graphs $f : \Gamma \to \Delta$ composing q_k with f we get a map from Γ to R_k, that induces an orientation on the edges of Γ that is preserved by $q_k \circ f(e)$, and assigns the label of $q_k \circ f(e)$ to e. We will usually think on a map f as those assigned labels in Γ by $q_k \circ f(e)$, with the induced orientation.

A *path* p in Γ, of *lenght* $n = |p|$ with *initial vertex* u and *terminal vertex* v, is an n-tuple of edges of Γ, $p = e_1 e_2 \cdots e_n$, such that $\tau(e_i) = \iota(e_{i+1})$ for $i = 1 \ldots n-1$, and $u = \iota(e_1)$ and $v = \tau(e_n)$. We can also think a path as a map of graphs from the line graph with $n+1$ vertices to Γ, $p : L_n \to \Gamma$.

In the set of paths of Γ, $P(\Gamma)$, we can define an operation, called *concatenation*, as follows: for paths $p = e_1 e_2 \ldots e_n$ and $q = e'_1 e'_2 \ldots e'_m$ with $\tau(e_n) = \iota(e'_1)$ we can form the new path $pq = e_1 e_2 \ldots e_n e'_1 \ldots e_m$ of lenght $n + m$. On the other hand, a map of graphs induces a lenght-preserving homomorphism from $P(\Gamma)$ to $P(\Delta)$.

A *round-trip* is a path of the form ee^*. If a path contains two adjacent edges forming a round-trip, we can delete it and get a path p' of lenght

$|p|-2$. Although the round-trip can be in the two first edges of the path or in the two last, p' has the same initial and terminal vertex as p. We call this operation *elementary reduction*, and write $p \sim p'$ if p' is obtained from p by successive elementary reductions or viceversa. Then, it's clear that $\sim$ is an equivalence relation, called *homotopy*. Furthermore, concatenacion of paths is compatible with homotopy since two homotopical paths have the same initial and terminal vertices, and then $p \sim p'$, $q \sim q'$ implies $pq \sim p'q'$.

From now, consider the set of $\sim -$classes of paths in Γ *based* in a vertex v, i.e. with v as initial and terminal vertex, we denote this set as $P(\Gamma, v)/\sim$. We will choose as a representant of each class the *reduced path*, i.e. the only path it contains without round-trips.

By definition of the concatenation operation and its compatibleness with $\sim$, we have that this operation is well defined and closed in $P(\Gamma, v)/\sim$. Furthermore, each element has an inverse, taking the path of lenght zero as the identity element: let $[p]$ be an homotopy class, with p the reduced path. If $|p| = 0$, $p = v$, the path of lenght 0 and $[p]$ is the class of the identity. If $|p| > 0$, then $p = re$, where e is the last edge of p, and define recursively $p^* = e^*r^*$, and then $[p]^{-1} = [p^*]$, $[p][p]^{-1} = [pp^*] = [v]$ and $[p]^{-1}[p] = [p^*p] = [v]$. In other words, if $p = e_1e_2 \dots e_n$ take $p^* = e_n^* \dots e_2^* e_1^*$. Then $P(\Gamma, v)/\sim$ with the operation of concatenation is a group, $\pi_1(\Gamma, v)$ called the *fundamental group* of Γ based at v.

Given a map of graphs $f : \Gamma \to \Delta$, it induces an homomorphism in the fundamental groups, denoted by the same symbol:

$$f : \pi_1(\Gamma, v) \to \pi_1(\Delta, f(v))$$

1.2 Computing the fundamental group of a graph

Here we list some classical properties of graphs in order to compute very easily the elements of their fundamental groups. The graphs are thought as non-oriented ones: e and e^* are identified in one edge that can be reached forwards or backwards. First we need some definitions about graphs:

A graph is *connected* if for every pair of vertices there exists a path from one to the other. A path with the same initial and terminal vertex is called a *circuit* and a graph is a forest if the only reduced circuits have lenght 0. A *tree* is a connected forest.

Proposition 1.1.

(a) Every connected graph contains a maximal tree.

(b) Every maximal tree in a connected graph Γ contains all the vertices of Γ.

(c) Let v be a vertex of a connected graph Γ, and $T \subseteq G$ a maximal tree, that exists by (a). Let θ be an orientation of Γ. Then $\pi_1(\Gamma, v)$ is free with basis.

$$B = \{P_e \,|\, e \in \theta \backslash E(T)\}$$

Where $P_e = T[v, \iota(e)]\, e\, T[\tau(e), v]$, with $T[u, v]$ the unique reduced path contained in the tree from u to v.

Proof.
The result follows from Zorn's lemma. For a connected graph, consider the set of all its subtrees with the partial order induced by the inclusion of trees as sets, i.e. $T_1 \leq T_2$ if and only if $T_1 \subseteq T_2$. For a chain of inclusions of trees, $T_1 \subseteq T_2 \subseteq T_3 \subseteq \ldots$, we have that $\bigcup_{i=1}^{\infty} T_i$ is a tree, as follows: suppose by contradiction that contains a non-trivial circuit, then, since that circuit has finite lenght by definition, it must lie in a tree of the chain, a contradiction to the definition of tree. That $\bigcup_{i=1}^{\infty} T_i$ is connected follows from the same argument: any pair of vertices must be contained in a tree of the chain, so, taking a path between them in that tree, we have that that path is containted in $\bigcup_{i=1}^{\infty} T_i$ also. So, $\bigcup_{i=1}^{\infty} T_i$ is a tree containing each T_i thus we can apply Zorn's lemma and get that there exists a maximal tree in the graph.

To prove (b), let T be a maximal tree of Γ and suppose by contradiction that exists a vertex v that is not in T. For an arbitrary vertex w of the tree, let $P = e_1 \ldots e_n$ a path in Γ with w and v as initial an terminal vertices, respectively. This path exists because Γ is connected. Let $j = min\{i = 1 \ \ldots n \,|\, \tau(e_i) \in V(\Gamma)\backslash V(T)\}$, then $T' = T \cup e_j$ is a tree, otherwise $\tau(e_i)$ would be already in T. This contradicts de maximality of T.

For (c), first of all observe that given two vertices of the graph, there exists only one reduced path contained in T from one to the other, otherwise T would contain a non-trivial circuit. This implies that the paths P_e are well defined. In order to prove that they generate $\pi_1(\Gamma, v)$, take $P = e_1 e_2 \ldots e_n$ a reduced v-based path. Let $e_{i_1} \ldots e_{i_k}$ be the edges of P contained in $E(G)\backslash E(T)$, and then P can be expressed as $P = p_1 e_{i_1} \dots e_{i_k} p_{k+1}$, where p_i are subpaths inside T. Then we claim that the path $(P_{e_{i_1}})^{\pm} \dots (P_{e_{i_k}})^{\pm}$, where the exponent of $(P_{e_{i_j}})^{\pm}$ is positive if $e_{i_j} \in \theta$ and negative otherwise, is homotopic to the path P. This follows easily from the uniqueness of the paths between two vertices in T: since P and $(P_{e_{i_1}})^{\pm} \dots (P_{e_{i_k}})^{\pm}$ reach the same edges outside T in the same order, and there is only one reduced path in T from e_{i_j} to $e_{i_{j+1}}$, for $1 \leq j \leq k-1$, we have that they are homotopic.

To see that it is a free set, consider a path $(P_{e_{i_1}})^{\varepsilon_1} \dots (P_{e_{i_n}})^{\varepsilon_n}$, with $\varepsilon_j = \pm$, homotopic to the trivial one, with $(P_{e_{i_j}})^{\varepsilon_j} \neq (P_{e_{i_{j+1}}})^{-\varepsilon_{j+1}}$, for $1 \leq j \leq n-1$. Then it can be written as: $p_1\, e_1\, p_2\, e_2 \dots p_{n-1}\, e_{n-1}\, p_n$ with $p_1, \dots p_n$ reduced paths in T and with $p_1 = T[v, \iota(e_1)]$ and $p_n = T[\tau(e_n), v]$. Consequently, any further reduction can only occur at a subpath $e_{j-1} p_j e_j$, with p_j the trivial path and $e_{j-1} = e_j^*$. But then $(P_{e_{i_j}})^{\varepsilon_j} = (P_{e_{i_{j+1}}})^{-\varepsilon_{j+1}}$, a contradiction.

$\square$

A basic corollary of **c)** is that the number of generators of $\pi_1(\Gamma, v)$, called the rank of $\pi_1(\Gamma, v)$, is given by the number of edges of Γ that are not contained in T. If Γ is finite we have that $rk(\pi_1(\Gamma, v)) = |E(\Gamma)| - |E(T)|$, and using that for any finite tree $|E(T)| = |V(T)| - 1$, then

$$rk(\pi_1(\Gamma, v)) = |E(\Gamma)| - |V(T)| + 1 = |E(\Gamma)| - |V(\Gamma)| + 1.$$

1.3 Coverings

We have already talked about R_k, the rose of k petals, a graph with one vertex and k edges, labeled by $\{a_1, a_2, \dots, a_k\}$, with an orientation. It is easy to see that the fundamental group of R_k is the free group of rank k with generators $\{a_1, a_2, \dots, a_k\}$, called F_k: every reduced path in R_k can be read as a word with the labels of its edges as letters. Then, for a graph Λ, assigning labels $\{a_1, a_2, \dots, a_k\}$ to its edges and choosing an orientation, we can construct a map $p : \Lambda \to R_k$ sending each vertex of Λ to the vertex of R_k and each edge of Λ to the edge in R_k with the same label, preserving orientations.

Suppose Λ has exactly $2k$ edges meeting at each vertex, then we will say that $p : \Lambda \to R_k$ is a *covering* if for each vertex k incident edges have label $a_1, a_2, \dots, a_k$ respectively and are oriented forwards with respect to that vertex, and the other k edges have label $a_1, a_2, \dots, a_k$ respectively and are oriented backwards. In other words, the local picture near each vertex of

Λ is the same as the picture of R_k.

From now, we will say that $p(\pi_1(\Lambda, v))$, a subgroup of F_k, is the *fundamental group induced by the covering* Λ or just the *fundamental group of the covering* Λ. To simplify, we will say also that the pair $(\Lambda, f : \Lambda \to R_k)$ is a *labeled graph* and we will usually denote it just by Λ. The definition of a map of graphs can be extended to labeled graphs, as follows:

DEFINITION 1.1. *Given* $(\Gamma, h_1 : \Gamma \to R_k)$*,* $(\Delta, h_2 : \Delta \to R_k)$ *two labeled graphs, a map of labeled graphs* $f : \Gamma \to \Delta$ *is a map of graphs such that* $h_2 \circ f = h_1$*.*

We claim now a lemma about coverings that will be very useful in the following theorems:

Lemma 1.1. *Let* $p : \Lambda \to R_k$ *be a covering, then given a path* $P = e_1 \dots e_n$ *in* R_k *and* u *a vertex of* Λ*, there exists one and only one path* $Q = e'_1 \dots e'_n$ *in* Λ*, with* $p(Q) = P$ *and* u *as initial vertex.*

Proof. Since each $e_j \in \{a_1, a_2, \dots, a_k\}^{\pm}$, if the exponent of e_1 is positive/negative take the unique edge incident to u and oriented forwards/backwards with the label of e_1, called e'_1. Inductively, given the path $Q = e'_1 \dots e'_{n-1}$ with u as initial vertex, there exists again one and only one edge $e_n^{'}$ oriented forwards/backwards, depending on the sign of e_n, adjacent to e'_{n-1} and with the label of e_n. $\square$

A covering of R_k is a particular case of an *immersion*: a map from a graph Θ, with an orientation, to R_k, such that, for any vertex of Θ, at most one forward incident edge has the label a_i and at most one backward incident edge has the label a_i, for any a_i. Then, for immersions, the $2k$-regularity is not required. Coverings can be thought as locally bijective maps and immersions as locally injective.

Given an immersion is easy to construct a covering, extending it with new incident edges to the vertices that are not $2k$-regular and labelling those edges with the corresponding letters. The problem is that, once we have done this process, we have new vertices that are not $2k$-regular. We can do it again with the new vertexs and go on. In order to formalize this process we define the following graph:

DEFINITION 1.2 (k-Cayley graph). *For* F_k *with generators* $\{a_1, \dots a_k\}$ *we define the k-Cayley graph as a labeled graph with vertices in one to one correspondence with the elements of* F_k*,* $V = \{v_x \mid x \in F_k\}$ *and set of oriented* *edges* $E = \{(v_x, v_{x'}) \mid x, x' \in F_k,\ \exists a_i \in \{a_1, \dots, a_k\}$ *such that* $xa_i = x'\}$. *Then* a_i *is the label of* $(v_x, v_{x'})$.

Proposition 1.2. *The* k*-Cayley graph is a covering and a tree.*

Proof.
That the k-Cayley graph is a covering follows from the definition of its edges: a vertex v_x has two and only two incident edges labeled with the letter a_j: the edge (v_x, v_{xa_j}) and the edge $(v_{xa_j^{-1}}, v_x)$, the first one oriented forwards and the other one oriented backwards. To check that the k-Cayley graph is connected, take v_x, v_y two vertices. Then, since the k-Cayley graph is a covering and using **Lemma 1.1**, there exists a path with v_x as initial vertex and labeled by $x^{-1}y$. Then, the terminal vertex of that path is v_y. Suppose that there exists another reduced path from v_x to v_y, labeled by z. Then, we have that $xz = y$, thus $z = x^{-1}y$. Then the path from v_x to v_y is unique and this implies that the k-Cayley graph is a tree. □
Now, since the k-Cayley is a tree, the deletion of any edge e from the Cayley produces two different connected components, A_e, B_e, with B_e containing the vertex v_1. Using this, we make the following definiton:

Definition 1.3 ($a_j^{\pm}$-branch of the k-Cayley graph). *Given* $a_j^{\pm}$ *with* a_j *a generator of* F_k*, take* e*, an edge of the* k*-Cayley graph, with* $e = (v_1, v_{a_j})$ *if the exponent of* a_j *is positive and* $e = (v_{a_j}, v_1)$ *if it is negative. Then, the* $a_j^{\pm}$*-branch of the* k*-Cayley graph is defined as* $A_e \cup e$*. The vertex of* $A_e \cup e$ *whose unique incident edge is* e *is called the attaching vertex.*

Then, given an immersion $p : \Theta \to R_k$, if one of the vertices of Θ does not have an incident edge oriented forwards/backwards and labeled by the letter a_j, we just attach the a_j^+/a_j^--branch of the k-Cayley by the attaching vertex. The crucial point is that, since the Cayley graph is a tree, the $a_j^{\pm}$-branch also, and then, using **Proposition 1.1**, the fundamental group of the new graph remains the same and the induced fundamental group as well. The bad news are that the Cayley have as many vertices as the elements of F_k and the branches a $\frac{1}{2k}$ part of the number of elements of F_k, so both are infinite graphs. Then, although we will prove the theorems for coverings, in practice we will work with immersions for obvious reasons. On the other hand, we could wonder whether, among all the immersions of a covering, there exists one containing the same information about the fundamental group as the covering and of minimum size:

Definition 1.4. *Given a map of graphs $f : (\Gamma, v) \to R_k$, let $\Omega = \{\Theta \subseteq \Gamma \,|\, v \in \Theta, f(\pi_1(\Theta, v)) = f(\pi_1(\Gamma, v))\}$. Then the core of Γ, denoted by $C(\Gamma)$, is defined as $C(\Gamma) = \bigcap_{\Theta \in \Omega} \Theta$.*

Lemma 1.2. *Given a map of graphs $f : (\Gamma, v) \to R_k$, then in $C(\Gamma)$ every vertex has degree equal or greater than 2 except, maybe, v.*

Proof.
Suppose that there exists w a vertex of $C(\Gamma)$, $w \neq v$, with degree 1 and let w' be its adjacent vertex. There exists a reduced v-based path P that reaches w, otherwise w' wouldn't be in $C(\Gamma)$. But then, the previous and the following vertex of w in P must be w', so P is not reduced. $\square$

The following theorem is the key of this theory, because it characterizes the subgroups of F_k using the fundamental groups of the coverings of R_k.

Theorem 1.1. *If $p_1 : (\Lambda_1, u) \to R_k$ and $p_2 : (\Lambda_2, w) \to R_k$ are two connected coverings of R_k, then Λ_1 and Λ_2 are isomorphic via a map of labeled graphs, $f : \Lambda_1 \to \Lambda_2$ with $f(u) = w$ if and only if $p_1(\pi_1(\Lambda_1, u)) = p_2(\pi_1(\Lambda_2, w))$.*

Proof.
If there is an isomorphism $f : \Lambda_1 \to \Lambda_2$, with $f(u) = w$, since f is a map of labeled graphs we have $p_1 = p_2 \circ f$ and hence $\pi_1(\Lambda_1, u) \leq \pi_1(\Lambda_2, w)$, and from $p_2 = p_1 \circ f^{-1}$, $\pi_1(\Lambda_2, w) \leq \pi_1(\Lambda_1, u)$.

Conversely, we define $f : \Lambda_1 \to \Lambda_2$, as follows. First $f(u) = w$. For a vertex $r \in V(\Lambda_1)$, $r \neq u$, take a reduced path $P = e_1 \dots e_n$ from u to r and consider $p_1(P)$. Then, using **Lemma 1.1**, there exists a unique path Q in Λ_2 with w as initial vertex, and with $p_1(P) = p_2(Q)$. Let s be the terminal vertex of Q. Then $f(r) = s$. In order to see that this f is well defined take P, P' two reduced different paths from u to r, and Q, Q' two paths from w with $p_1(P) = p_2(Q)$ and $p_1(P') = p_2(Q')$. Suppose that s and s', the terminal vertices of Q and Q' respectively, are distinct. Then, since the path PP'^{-1} is a u-based path and using that $\pi_1(\Lambda_1, v) \leq \pi_1(\Lambda_2, w)$, there exists a w-based path K with $p_2(K) = p_1(PP'^{-1})$. Spliting K in two parts, $K = K_1K_2$ with $p_2(K_1) = p_1(P)$ and $p_2(K_2) = p_1(P'^{-1})$, since K_1 and Q have the same initial vertex and the same image by p_2 and using **Lemma 1.1** again, we get that $K_1 = Q$. Analogously, $K_2 = Q'^{-1}$, and then $s = s'$. On the other hand, to check that f is a map of labeled graphs and is well defined on the edges, take an edge $e = (x, y)$ of Λ_1. Let P be the reduced path from u to x and Q the path in Λ_2 from w with $p_1(P) = p_2(Q)$ and $f(x)$ as terminal vertex. Then, taking e' the unique edge incident to $f(x)$ with the same label and orientation as e, we have that $p_1(Pe) = p_2(Qe')$, and hence the terminal vertex of Qe' is $f(y)$, thus $f(e) = f(e')$. Define $f^{-1} : \Lambda_2 \to \Lambda_1$ in the same fashion, using now that $\pi_1(\Lambda_2, v) \leq \pi_1(\Lambda_1, w)$. To see that $f^{-1} \circ f(v) = v$ just take P the reduced path from u to v and Q the path in Λ_2 from w with $p_1(P) = p_2(Q)$ and $f(v)$ as terminal vertex. Now, $f^{-1} \circ f(v)$ is again the terminal vertex of P and then v. The same for $f \circ f^{-1}(v)$.

□

1.4 Stallings foldings

Our goal now is, for any given subgroup A of F_k, to find an algorithmic way to construct a covering of R_k whose fundamental group is A. Then, using **Theorem 1.1**, it will be easy to construct a bijection between the subgroups of F_k and the coverings of R_k.

Given $A =< w_i(a_1^{\pm}, \ldots a_k^{\pm}) >_{i=1}^{m}$ a subgroup of F_k, where w_i is a reduced word, we can think of w_i as a tuple of elements of $\{a_1 \ldots a_k\}^{\pm}$. Now, construct a graph as the disjoint union of exactly m line graphs, $P_1, \ldots , P_m$, with lenght of P_i equal to the lenght of w_i as a word. Finally indentify the initial vertex and terminal vertex of all line graphs in one vertex, called v. Then our graph consists of m v-based paths, $P_1', \ldots , P_m'$ and it's easy to construct a map from this graph to R_k, sending the fundamental group to A: label the edges of P_i' with the letters of the tuple w_i and choose orientations according to their exponents in w_i. We call this construction the *flower graph of A*. The question now is how we could construct, from this flower graph, an immersion of F_k with fundamental group A (and then, from that inmersion we will be able to construct a covering immediately as we explained before). Basically, the problem is that the local picture near each vertex of the flower graph may not be the same as the picture of R_k. In other words, there could be two edges, meeting at the same vertex, with the same label and both oriented forwards or backwards. These two edges must be incident to v because the words w_i are simplified and then the paths P_i can not contain two such edges. This kind of pair of edges is called *inadmissible*. The solution to this problem is due to Stallings, in the paper [1], and is very simple: identifying each pair of inadmissible edges doesn't change the induced fundamental group:

Theorem 1.2. *Let $p : \Gamma \to R_k$ with a defined orientation in Γ. If (e_1, e_2) is* *an inadmissible pair of edges, i.e. e_1 and e_2 have one vertex v in common, both are oriented forwards (backwards) with respect to v and $p(e_1) = p(e_2)$, then for $p' : \Gamma' = \Gamma/[e_1 = e_2] \to R_k$, the map induced by the identification of e_1 with e_2, we have that $p'(\pi_1(\Gamma, [v]))$ remains the same.*

Proof.

For $a \in p(\pi_1(\Gamma, v))$, let x with $p(x) = a$. If e is an edge of x different from e_1 and e_2, $p'([e]) = p(e)$. Otherwise, if e is either e_1 or e_2, since $p'([e_1]) = p'([e_2]) = p(e_1) = p(e_2)$, $p'([e]) = p(e)$. On the other hand, $[x]$ is a well defined path in Γ' because $[e]$ is incident, at least, to the same edges as e. Then $p'([x]) = a$.

Conversely, suppose, without loss of generality, that $e_1 = (v, u)$ and $e_2 = (v, w)$. For a', an element of $p'(\pi_1(\Gamma', [v]))$, and $[x'] = [x'_1] \ldots [x'_n]$ with $p'([x']) = a'$, define a path $y = r_1 \ldots r_n$ in Γ as the concatenation of the paths r_i. In order to define those paths we distinguish two cases, first, $u \neq w$. Then:

(a) If $[x'_i] \neq [e_1]$ and $x'_i = (x, u)$ for some vertex x. Then, if $x'_j = (u, y)$ for $j \equiv i + 1 \bmod n$, define $r_i = x'_i$ and then $p(r_i) = p'([x'_i])$. And, if $x'_j = (w, y)$ define $r_i = x'_i e_1^{-1} e_2$, and then $p(r_i) = p(x'_i e_1^{-1} e_2) = p(x'_i) p(e_1^{-1}) p(e_2) = p(x'_i) = p'([x'_i])$.

(b) If $[x'_i] \neq [e_1]$ and $x'_i = (x, w)$ for some vertex x, this case is analogous to the previous one.

(c) If $[x'_i] \neq [e_1]$ but x'_i is not in one of the previous cases, define $r_i = x'_i$.

(d) If $[x'_i] = [e_1]$ and $x'_j = [(u, x)]$ for $j \equiv i + 1$ and for some vertex x, then define $r_i = e_1$, and then $p(r_i) = p(e_1) = p(e_2) = p'([e_1])$.

(e) If $[x'_i] = [e_1]$ and $x'_j = [(w, x)]$ for $j \equiv i + 1$ and for some vertex x, then define $r_i = e_2$.

It's easy to check that this path is well defined: if $[x'_i] \neq [e_1]$ but x'_i is incident to u or to w (cases (a) and (b)) we look at the next edge of the path and if it's not adjacent to x'_i in Γ, we reach it running through $e_1^{-1} e_2$ or $e_2^{-1} e_1$. On the other hand, if $[x'_i] = [e_1]$, we look at the next edge of the path to see whether it has been reached from e_1 or from e_2. The second case is $u = w$:

(a) If $[x'_i] \neq [e_1]$ define $r_i = x'_i$. Then $p(r_i) = p'([x'_i])$.

(b) If $[x'_i] = [e_1]$ define $r_i = e_1$. Then $p(r_i) = p'([e_1])$.

The well definition is clear.

Observe that in the first case, after the identification, Γ' has one vertex and one edge less than Γ, so:

$$rk(\pi_1(\Gamma', [v])) = |E(\Gamma')| - |V(\Gamma')| + 1 =$$

$$(|E(\Gamma)| - 1) - (|V(\Gamma)| - 1) + 1 = rk(\pi_1(\Gamma, v)).$$

On the other hand, in the second case Γ' has one edge less but the same number of vertices than Γ, so:

$$rk(\pi_1(\Gamma', [v])) = |E(\Gamma')| - |V(\Gamma')| + 1 = (|E(\Gamma)| - 1) - (|V(\Gamma)|) + 1 =$$

$$rk(\pi_1(\Gamma, v)) - 1.$$

□

Although the previous proof seems very technical, it provides us an algorithm to reconstruct a path in Γ from a path of $\Gamma' = \Gamma/[e_1 = e_2]$, both describing the same element in the free group F_k. We will use this algorithm in the next section.

Theorem 1.3. *If F_k is the free group of rank k with generators $\{a_1, a_2, ..., a_k\}$, then exists a set-bijection $\psi : \mathcal{S}_{F_k} \to \mathcal{C}_{F_k}$ from the set of finitely generated subgroups of F_k to the set of connected coverings of R_k with a distinguished vertex and with finite core, such that $p(\pi_1(\psi(S), v)) = S$, for every $\{p : (\psi(S), v) \to R_k\} \in \mathcal{S}_{F_k}$ and v the distinguished vertex of $\psi(S)$.*

Proof.
For $S \leq F_k$, let $\{s_1, s_2, ..., s_j\}$ elements of F_k generating S. Construct Γ the flower graph of $\{s_1, s_2, ..., s_j\}$ and call v to the central vertex. Then, $\pi_1(\Gamma, v) = S$. Now, by **Theorem 1.2** and because Γ is finite, exists a finite chain of foldings $\Gamma \to \Gamma' \to \cdots \to \Gamma^{(k)}$, with $\pi_1(\Gamma, v) = \pi_1(\Gamma^i, v)$ and with $\Gamma^{(k)}$ immersion of R_k. Attaching $a_j^{\pm}$-branches of the k-Cayley graph we can get a covering Δ. Observe that since $\Gamma^{(k)}$ is finite and $C(\Delta)$ is a subgraph of $\Gamma^{(k)}$, $C(\Delta)$ is finite. Then, define $\psi(S) = (\Delta, v)$. Conversely, for a covering with finite core, $p : (\Lambda, w) \to R_k$ we have that $p(\pi_1(\Lambda, w))$ is finetely generated. Then define $\psi^{-1}(\Lambda, w) = p(\pi_1(\Lambda, w))$. Now, $\psi^{-1}(\psi) = id$ holds automatically, and $\psi(\psi^{-1})$ follows from **Theorem 1.1**: if two coverings have the same fundamental group they are isomorphic as labelled graphs. □

From now, to simplify, if we talk of $p(\pi_1(\psi(S), v))$, v will be the distinguished vertex of $\psi(S)$.

2 Applications of the covering theory

Here we list and prove some applications of the machinery developed in the previous chapter. Problems like the membership or the subgroups conjugation have a really simple and elegant solution using this techinque, that constrasts with the complex methods used before the construction was developed.

2.1 The membership problem

Given a finitely-generated subgroup of a free group, $S \leq F_k$, we wonder whether, given an element $a \in F_k$, that element belongs to S also. In other words, we wonder whether operating the generators of S and their inverses in some way, we can get a. The solution of this problem is maybe the best example of how Stallings construction can be really powerful. The first step to solve it is to find a basis of S:

Theorem 2.1 (Algorithm to find a basis)**.** *Given S a finitely-generated subgroup of a free group F_k, let $p : (\Theta, v) \to R_k$ be an immersion with $p(\pi_1(\Theta, v)) = S$. Let $\{b_1, \dots, b_j\}$ be a basis of $\pi_1(\Theta, v)$, then $\{p(b_1), \dots, p(b_j)\}$ is a basis of S.*

Proof.
Since $p(\pi_1(\Theta, v)) = S$ and the elements $\{b_1, \dots, b_j\}$ generate $\pi_1(\Theta, v)$ we get that $\{p(b_1), \dots, p(b_j)\}$ generate S. To see that they are free generators, take $P = e_1 \dots e_n$ a non trivial reduced path in Θ and suppose that $p(P)$ is the trivial path. But then, for some i, $p(e_i) = p(e_{i+1})^{-1}$, so e_i and e_{i+1} are two adjacent edges with the same label and both oriented forwards o backwards with respect to their common vertex, that is imposible because Θ is an immersion. So, for any non-trivial concatenation of $\{b_1, \dots, b_j\}$, the image of this path by p is non trivial, thus it's a basis. $\square$
Algorithmically, if S is generated by $\{s_1, \dots s_m\}$, just construct the flower graph Γ of $\{s_1, s_2, ..., s_m\}$. Then, apply successively Stallings foldings to get the chain $\Gamma \to \Gamma' \to \cdots \to \Gamma^{(k)}$ where $\Gamma^{(k)}$ is an immersion of R_k with induced fundamental group A. Find a basis of $\Gamma^{(k)}$ using **Theorem 1.1** and using the previous theorem, the image of this basis by $p^{(k)}$ is a basis of S.

A first important fact that we can easily get from this machinery is that every subgroup of a free group is free, as follows. From **Theorem 1.3**, given a subgroup S of F_k, exists one and only one covering of R_k such that

$p(\pi_1(\psi(S), w)) = S$. But then, since every covering is an immersion and using the previous theorem, since $\pi_1(\psi(S), w)$ is free, S is free also.

Now, given a subgroup S and an element a of F_k, to check if a belongs to S becomes a really easy problem: take $\Gamma^{(k)}$, the last graph of the chain of foldings and then an immersion, and try to find P, a v-based path in $\Gamma^{(k)}$, labeled by a. Furthermore: by the definition of immersion, if that path exists, it's unique. If a is an element of S, we can go one step farther and know how the generators of S have to be operated in order to get a. In order to do this we use the algorithm defined in the proof of **Theorem 1.2.**: take a reduced path in $\Gamma^{(k)}$ such that $p^{(k)}(P^{(k)}) = a$. Then, using the algorithm, construct a reduced path $P^{(k-1)}$ in $\Gamma^{(k-1)}$ such that $p^{(k-1)}(P^{(k-1)}) = a$. Interatively, we can get a path in P in Γ with $p(P) = a$. Once simplified, this path is a concatenation of v-based subpaths, $r_1 \dots r_m$, with $p(r_i)$ one of the generators of S.
Here we have an example:

Example 2.1. *Let S be a subgroup generated by $\{a^3, ab, a^2ba\}$. Find free generators of S, check if the element $a^{-2}ba^{-1}ba$ is containted in S and how the generators of S must be operated in order to get that element.*
First we construct the flower graph of those generators:

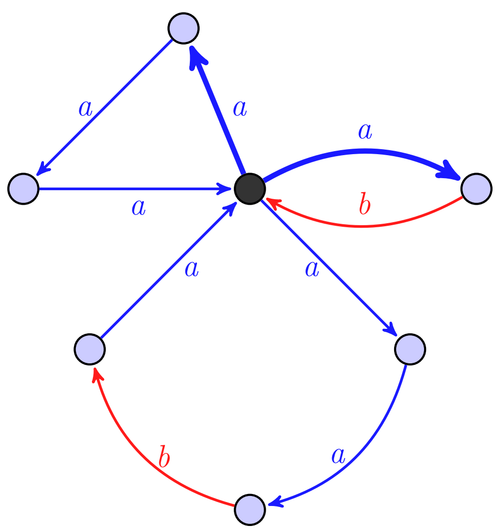

And now we apply successive Stallings foldings in order to get an immersion with induced fundamental group S. The edges that will be identified in the next step have a thicker line

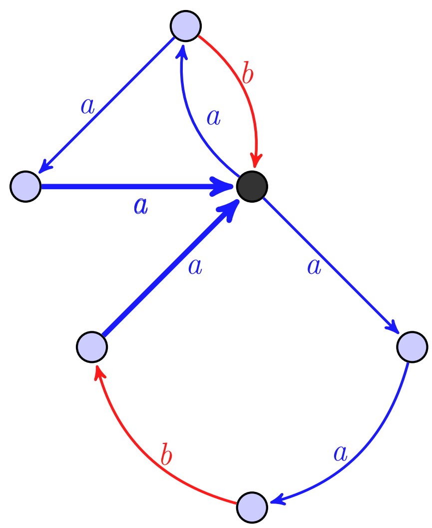

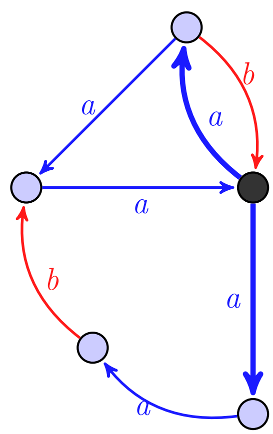

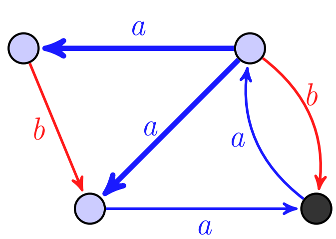

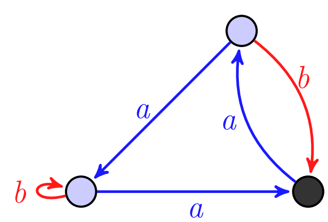

Observe that $rk(S) = |E(\Gamma)| - |V(\Gamma)| + 1 = 5 - 3 - 1 = 3$*. So* $\{a^3, ab, a^2ba\}$ *was already a free basis of* S*. There is a path in the immersion labeled by* $a^{-2}ba^{-1}ba$*, so it's an element of* S*. Using the algorithm of* ***Theorem 1.2*** *this path can be lifted to the flower graph, as follows:*

$$a^{-2}ba^{-1}ba \Rightarrow a^{-2}ba^{-1}(a^{-1}a)ba \Rightarrow a^{-2}ba^{-1}a^{-1}(a^{-1}a)aba \Rightarrow$$

$a^{-2}ba^{-1}a^{-1}(a^{-1}a)aba \Rightarrow a^{-2}(a^{-1}a)ba^{-1}a^{-1}a^{-1}aaba \Rightarrow (a^3)^{-1}(ab)(a^3)^{-1}(a^2ba)$

So we know how the generators of S must be operated in order to get $a^{-2}ba^{-1}ba$, that would be really hard to know without this construction.

2.2 The conjugacy problem

Given S, L subgroups of F_k we want to guess if S and L are conjugate, i.e. if exists an element z of F_k such that $S = zLz^{-1}$. Obviously, if S and L are conjugate they have the same rank.

Theorem 2.2. *For S and L different subgroups of F_k and $p_1 : (\Lambda_1, u) \to R_k$, $p_2 : (\Lambda_2, w) \to R_k$ two connected coverings with induced fundamental group S and L respectively, we have that S and L are conjugate subgroups if and only if Λ_1 is isomorphic to Λ_2 as a labeled graphs but with the image of u not w. In that case, the conjugator element z is the image of a reduced path from u to w.*

Proof.
If S and L are conjugate with $S = zLz^{-1}$, let Q be the path in Λ_2 with terminal vertex w and, such that, $p_2(Q) = z$. The existence of this path follows from **Lemma 1.1**. Let v be the initial vertex of Q. Then, for any w-based path P, the concatenation QPQ^{-1} is a v-based path, and then we have that $zp_2(\pi_1(\Lambda_2, w))z^{-1} \leq p_2(\pi_1(\Lambda_2, v))$. On the other hand, for any v-based path P', the concatenation $Q^{-1}P'Q$ is a w-based path, and then $z^{-1}p_2(\pi_1(\Lambda_2, v))z \leq p_2(\pi_1(\Lambda_2, w))$. Thus $p_2(\pi_1(\Lambda_2, v)) = zp_2(\pi_1(\Lambda_2, w))z^{-1} = zLz^{-1} = S$. And, since (Λ_1, u) and (Λ_2, v) have the same induced fundamental group, from **Theorem 1.1** we get that (Λ_1, u) is isomorph to (Λ_2, v) as a labeled graph. Conversely consider the path Q from u to w in $\Lambda_1 \cong \Lambda_2$ and let $z = p_1(Q)$. For each P, w-based path, we have that QPQ^{-1} is a u-based path and then $zp_2(\pi_1(\Lambda_2, w))z^{-1} \leq p_1(\pi_1((\Lambda_2, u)))$. On the other hand, for each P', u-based path, $Q^{-1}P'Q$ is a w-based path and then $z^{-1}p_1(\pi_1(\Lambda_1, u))z \leq p_2(\pi_1(\Lambda_2, w))$. Thus $zp_2(\pi_1(\Lambda_2, w))z^{-1} = p_1(\pi_1(\Lambda_2, u))$ and then S and L are conjugate subgroups with z as the conjugator element.
□

Unfortunately, we can't define an algorithm using coverings because they could be infinite. The problem now is that, if we work with the immersion obtained from the flower graph of L after applying Stallings foldings, the previous theorem can be false. In fact, for a conjugate subgroup of L, zLz^{-1}, there could not exists a path whose image is z in that immersion. We have to go inside the a_j-branches to find it. The next example is quite illustrative:

Example 2.2. *Consider S, a subgroup of F_2 generated by $\{b^2, bab^{-1}\}$. First of all we compute an immersion with fundamental group S. The flower graph of S is:*

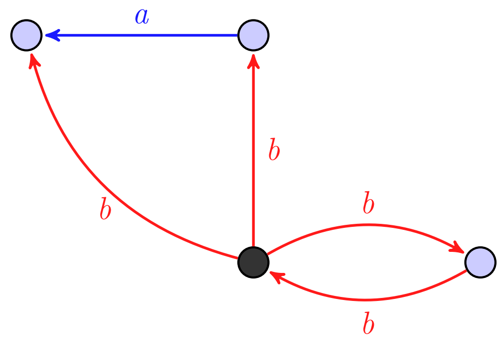

We only have to apply two foldings to get an immersion:

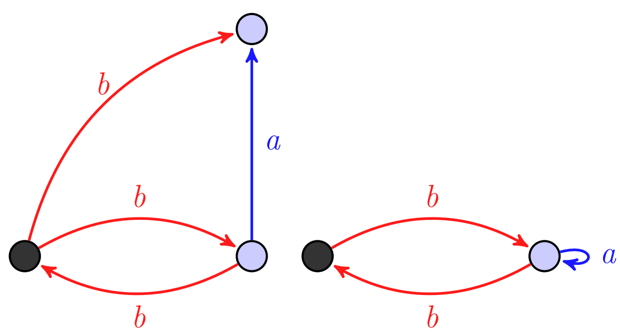

Recall that, in order to get a covering from this immersion we have to attach a_j-branches of the 2-Cayley graph. So, attaching just some edges of those a_j-branches we will get a new immersion:

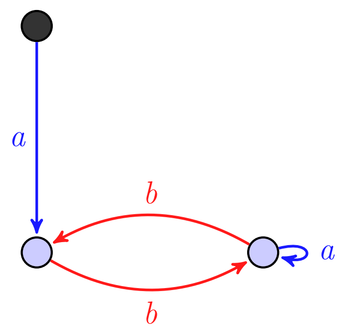

The fundamental group based at that new vertex in black is generated by $\{ab^2a^{-1}, abab^{-1}a^{-1}\}$ *so it's a conjugate of the first one, but the immersions are not isomorphic as a graphs.*

Observe that in the previous example both immersions are cores with fundamental group S and aSa^{-1} respectively, so **Theorem 2.2** doesn't work with cores either. The problem here is that the base vertex of the second immersion is inside the a_j-branch and it could be as deep as we want to. So, the solution in some way is to forget all the graphs whose base vertex is in the same a_j-branch of the core, and then we will get the isomorphism of **Theorem 2.2**. In order to do this, we already know that in a core the only vertex that can have degree 1 is the base vertex. If this happens, it means that the base vertex is in a a_j-branch, so let's try to move it outside. Let R be a reduced path with the base vertex v as initial vertex and such that the only vertex of degree equal or greater than 3 is its terminal vertex, called $r(v)$. This path always exists if the fundamental group of the core is not the trivial one and is unique because the degree of the basic vertex is 1. In some sense, we are moving the base vertex v to $r(v)$. Then the core without R has no vertices of degree 1. We formalize this in the following definition:

Definition 2.1 (Reduced core). *Given a map of graphs* $f : (\Gamma, v) \to R_k$*, let* $C(\Gamma)$ *be its core. If* v *has degree equal or greater than 2 in* $C(\Gamma)$ *then, define the reduced core of* Γ*, denoted by* $C_r(\Gamma)$ *as* $C_r(\Gamma) = C(\Gamma)$*. Otherwise, let* R *be the path defined above and* $r(v)$ *its terminal vertex. Then define* $C_r(\Gamma) = C(\Gamma)\backslash(R\backslash r(v))$ *and take* $r(v)$ *as a distinguished vertex.*

Then observe also that the reduced core have only vertices of degree equal or greater than 2. So, once we have moved the basic vertex of the core outside the a_j-branch, **Theorem 2.2** holds:

Theorem 2.3. *For S and L different subgroups of F_k and $p_1 : (\Lambda_1, u) \to R_k$, $p_2 : (\Lambda_2, w) \to R_k$ two connected coverings with induced fundamental group S and L respectively, we have that S and L are conjugate subgroups if and only if $C_r(\Lambda_1)$ is isomorph to $C_r(\Lambda_2)$ as a labeled graphs. In that case, the conjugator element z is $p_1(R)f_1(Q)p_2((R')^{-1})$ where R and R' are the paths of the Definition 2.1 and Q is a path from $r(u)$ to $r(w)$.*

Proof.
If S and L are conjugate subgroups, using **Theorem 2.2**, Λ_1 and Λ_2 are isomorphic as a labeled graphs. Then, since the reduced cores are in one to one correspondence with the connected coverings, $C_r(\Lambda_1)$ and $C_r(\Lambda_2)$ are isomorphic as a labeled graphs as well. Conversely consider the path Q from $r(u)$ to $r(w)$ in $C_r(\Lambda_1) \cong C_r(\Lambda_2)$ and let $z = p_1(Q)$. Using the same argument as in the proof of **Theorem 2.2** we get that:

$$zp_2(\pi_1(C_r(\Lambda_2), r(w)))z^{-1} = p_1(\pi_1(C_r(\Lambda_1), r(u))).$$

Now, let R and R' be the paths of the **Definition 2.1**, then

$$\pi_1(C_r(\Lambda_1, r(u)) = R^{-1}\pi_1(C(\Lambda_1), u)R, \quad \pi_1(C_r(\Lambda_2), r(w)) = (R')^{-1}\pi_1(C(\Lambda_2), w)R'.$$

And then:

$$zp_2(R'^{-1})p_2(\pi_1(C(\Lambda_2), w))p_2(R')z^{-1} = p_1(R^{-1})p_1(\pi_1(C(\Lambda_2), u)p_1(R) \Leftrightarrow$$

$$(p_1(R)zp_2(R'^{-1}))p_2(\pi_1(C(\Lambda_2), w))(p_2(R')z^{-1}p_1(R^{-1})) = p_1(\pi_1(C(\Lambda_2), u) \Leftrightarrow$$

$$(p_1(R)zp_2(R'^{-1}))L(p_2(R')z^{-1}p_1(R^{-1})) = S.$$

□

2.3 The normal subgroup problem

Recall that a subgroup $N \leq F_k$ is called normal if it is invariant under conjugation by every element of F_k. In other words, for any z, element of F_k, $zNz^{-1} = N$. This means that, before the Stallings construction was developed, the normality of a subgroup was a property really hard to check.

Using graphs, on the contrary, it becomes astonishingly easy. Using the previous section we know that to conjugate a subgroup by an element a means to change the base point of its covering from v to v', the terminal vertex of a as a path. Then, if the subgroup is normal, the induced fundamental group using this new base point remains the same. In other words, the v'-based labeled paths are in one to one correspondence with the v-based labeled paths, and thus, using **Theorem 1.1**, we will get that there exists a labeled graph automorphism, sending v to v'.

Theorem 2.4. *Given N, a subgroup of F_k, we have that N is normal if and only if for any w, vertex of $(\psi(N), v)$, there exists an isomorphism of labeled graps, $f : \psi(N) \to \psi(N)$, with $f(v) = f(w)$.*

Proof.
Let $p : (\psi(N), v) \to R_k$ be the connected covering of N and w a vertex of $\psi(N)$. Then, since $\psi(N)$ is connected, there exists a reduced path P from v to w. Let $a \in F_k$ with $a = p(P)$. Then, using **Theorem 2.2**, we have that aNa^{-1} has covering $(\psi(N), w)$. Thus, using that N is normal we have know that $N = \pi_1(\psi(N), v) = \pi_1(\psi(N), w)$. So, using **Theorem 1.1** there exists an isomorphism of labeled graphs $f\psi(N) \to \psi(N)$ with $f(v) = f(w)$. Conversely, for any a element of F_k, there exists a path P in $\psi(N)$ with $p(P) = a$ and v as initial vertex. Then, let w be the terminal vertex of P. Using that $p : (\psi(N), w) \to R_k$ is the covering of aNa^{-1}, then, by **Theorem 1.1** again, $N = aNa^{-1}$. □

2.4 The intersection problem

Given S, L subgroups of F_k, given an element $a \in F_k$, we wonder again whether that element belongs to $S \cap L$. Or, more generally, we want to find generators of $S \cap L$. The procedure is similar to the previous one: construct the covering of the subgroup $S \cap L$ and find generators of the induced fundamental group. In order to do this we need some definitions:

DEFINITION 2.2 (Product of a map of graphs). *Given $f_1 : (\Gamma_1, u) \to R_k$ and $f_2 : (\Gamma_2, w) \to R_k$ two maps of graphs with distinguished vertex, we define the product graph, $\Gamma_1 \times \Gamma_2$, as the graph with set of vertices $V_{\Gamma_1 \times \Gamma_2} = V_{\Gamma_1} \times V_{\Gamma_2}$ and set of edges $E = \{((u_1, u_2), (w_1, w_2)) \,|\, (u_1, w_1) \in E_{\Gamma_1}, (u_2, w_2) \in E_{\Gamma_2}$ and $f_1(u_1, w_1) = f_2(u_2, w_2)\}$. Then, the product of the maps f_1 and f_2,*

$f_1 \times f_2 : \Gamma_1 \times \Gamma_2 \to R_k$, *is the map sending each edge* $((u_1, u_2), (w_1, w_2))$ *to* $f_1(u_1, w_1)$.

DEFINITION 2.3. *The projection maps* $x : E_{\Gamma_1 \times \Gamma_2} \to E_{\Gamma_1}$ *and* $y : E_{\Gamma_1 \times \Gamma_2} \to E_{\Gamma_1}$ *are defined as* $x((u_1, u_2), (w_1, w_2)) = (u_1, w_1)$ *and* $y((u_1, u_2), (w_1, w_2)) = (u_2, w_2)$.

Theorem 2.5. *Given* S, L *subgroups of* F_k, *if* $p_1 : (\Lambda_1, u) \to R_k$ *and* $p_2 : (\Lambda_2, w) \to R_k$ *are two coverings with induced fundamental group* S *and* L *respectively, and if* $C_{(u,w)} \subseteq \Lambda_1 \times \Lambda_2$ *is the connected component containing the vertex* (u, w), *then* $(p_1 \times p_2)_{\|C_{(u,w)}} : C_{(u,w)} \to R_k$ *is a covering with induced fundamental group* $S \cap L$.

Proof.
That $C_{(u,w)}$ is a covering follows inmmediately from the definition of the product of maps and from the fact that Λ_1 and Λ_2 are coverings. To see that the induced fundamental group is $S \cap L$, let $p_1 \times p_2(a)$ be an element of $p_1 \times p_2(\pi_1(C_{(u,w)}, u \times w))$, then projecting every edge of a by the maps x and y, we get a u-based path P_1 in Λ_1 and a w-based path P_2 in Λ_2 such that $p_1(P_1) = p_2(P_2) = p_1 \times p_2(a)$ by the definition of the product of maps. Then $p_1 \times p_2(a) \in p_1(\pi_1(\Lambda_1, u)) \cap p_2(\pi_1(\Lambda_2, w)) = A \cap B$. Conversely, if P_1 is a u-based path in Λ_1 and P_2 a w-based path in Λ_2 with $p_1(P_1) = p_2(P_2)$, then the graph product of p_1 and p_2 contains a (u, w)-based path in $\Lambda_1 \times \Lambda_2$ that induces the same element in R_k as P_1 and P_2. $\square$

Then, once we have an immersion with induced fundamental group $S \cap L$, we can find easily a basis. It's easy to check that the product of two immersions is also an immersion, so, algorithmically, if S and L are finite generated, $C(\Lambda_1)$ and $C(\Lambda_2)$ are finite cores, and then the corresponding connected component of $C(\Lambda_1) \times C(\Lambda_2)$ will be a finite inmersion of $S \cap L$, where we can get a basis. On the other hand, the graph $\Lambda_1 \times \Lambda_2$ could have more than one connected component. Althought in the previous theorem we only need the connected component containing (u, w), we can wonder what the other connected components mean. We can choose properly a vertex u' of Λ_1 so that (u', w) can lie in any of the other connected components. And the same can be done for a vertex w' of Λ_2. Then using the previous section, we know that (Λ_1, u) and (Λ_1, u') induce conjugate subgroups, so the components of $\Lambda_1 \times \Lambda_2$ mean intersections of conjugates subgroups of Λ_1 and Λ_2. Here we have one interesting example:

Example 2.3. *Given $S =< b, a^3, a^2bab^{-1}a >$ and $L =< a^3, ab, a^2ba >$, compute $S \cap L$. L is the same subgroup of the Example 2.1 and we have already get the immersion with S as fundamental group, so the product of both covering becomes:*

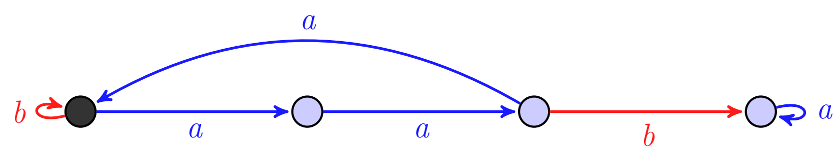

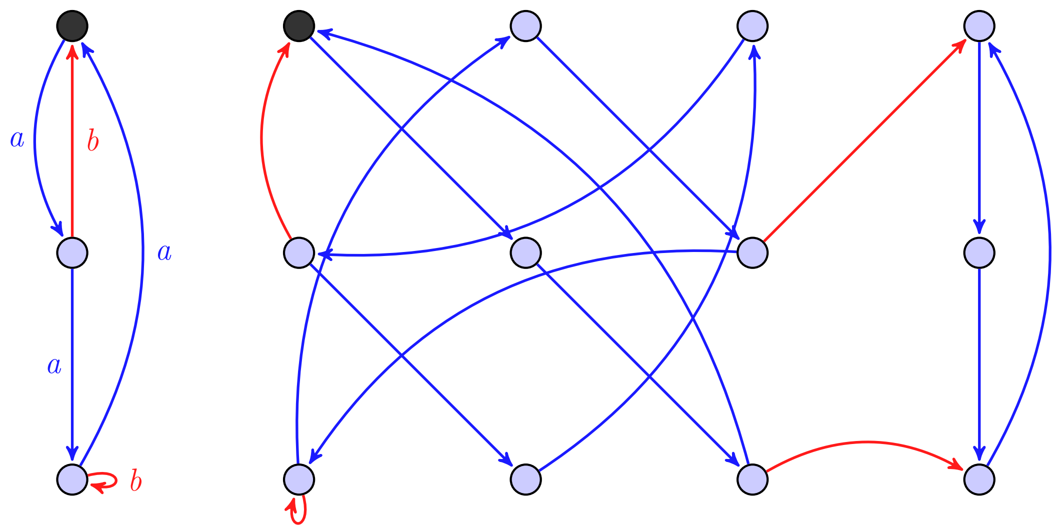

So, the $rk(S \cap L) = |E(\Gamma_1 \times \Gamma_1)| - |V(\Gamma_1 \times \Gamma_1)| + 1 = 16 - 12 + 1 = 5$. Taking a maximal tree and using Proposition 1.1 we get the following basis of $S \cap L$:

$$\{b^{-1}a^3b, a^3, a^2ba^3b^{-1}a^{-2}, a^2ba^{-2}b^{-1}a^3ba^2b^{-1}a^{-2}, a^2ba^{-2}b^{-1}aba^{-1}ba^2b^{-1}a^{-2}\}$$

So, oberserve that the rank of the subgroup intersection is greater than the ranks of S and L.

This example proves another surprising fact about the free group: the rank of $S \cap L$ can be larger than the rank of S and the rank of L. Nevertheless, we know that if the rank of S and the rank of L are finite, then the rank of the intersection is finite: since $C(\Lambda_1)$ and $C(\Lambda_2)$ are finite graphs, $C(\Lambda_1)\times C(\Delta_2)$ as well, and then the rank of $S\cap L$ is finite also. But, can it be arbitrarily large? The answer is no, and, in fact, it can be bounded in terms of the product of rank of S and rank of L. Before this theorem we need a definition:

Definition 2.4. *(Reduced rank) If S is a free group with $rk(S) = k$, we define $\widetilde{rk}(S) = min\{0, k-1\}$*

And, with this definition:

Theorem 2.6 (Hanna-Neumann)**.** *Given S, L subgroups of F_k, then*

$$\widetilde{rk}(S\cap L) \leq 2\widetilde{rk}(S)\widetilde{rk}(L) \tag{1}$$

In order to prove that result we need a pair of lemmas:

Lemma 2.1. *Given Γ a connected finite graph, then:*
(a) If Γ is not a tree, then $\sum_{v\in V(\Gamma)}(d(v)-2) = 2\widetilde{rk}(\Gamma)$.
(b) If Γ is a tree, then $\sum_{v\in V(\Gamma)}(d(v)-2) = -2$.

Proof.
We already know that $rk(\Gamma) = |E(\Gamma)| - |V(\Gamma)| + 1$, then if Γ is not a tree:

$$\widetilde{rk}(\Gamma) = |E(\Gamma)| - |V(\Gamma)| = \frac{1}{2}\sum_{v\in V(\Gamma)} d(v) - |V(\Gamma)| = \frac{1}{2}\sum_{v\in V(\Gamma)}(d(v)-2).$$

And if Γ is a tree:

$$-1 = \frac{1}{2}\sum_{v\in V(\Gamma)}(d(v)-2), \quad \sum_{v\in V(\Gamma)}(d(v)-2) = -2.$$

□
The following lemma bounds the degree of (u, v) in terms of the degree of u and v.

Lemma 2.2. *Given Λ_1, Λ_2, two coverings with finite generated subgroup. Consider $C_r(\Lambda_1)\times C_r(\Lambda_2)$, then, for every $(v, w) \in C_r(\Lambda_1)\times C_r(\Lambda_2)$, $d(v, w) - 2 \leq (d(v)-2)(d(w)-2)$.*

Proof.
Observe that, since $C_r(\Lambda_1)$ and $C_r(\Lambda_2)$ are immersions, for every edge incident to v there exists at most one edge incident to w with the same label. Then, since the number of edges incident to (v, w) is the number of pairs of edges incident to v and to w with the same label, we have that $d((v, w)) \leq min\{d(v), d(w)\}$. Now, suppose without loss of generality that $d(v) \leq d(w)$. Recall that, by the definition of reduced core the degree of each vertex of $C_r(\Lambda_1)$ and $C_r(\Lambda_2)$ is equal or greater than 2. Then, if $d(w) = 2$ the equation holds automatically. If $d(w) \geq 3$, we have that $d((v, w)) - 2 \leq min\{d(v), d(w)\} - 2 \leq d(v) - 2 \leq (d(v) - 2)(d(w) - 2)$ since $d(w) - 2 \geq 1$. $\square$

With this lemmas we can already proof the theorem:
Proof.
Let $p_1 : (\Lambda_1, u) \to F_k$ and $p_2 : (\Lambda_2, w) \to F_k$ be two coverings with induced fundamental group S and L respectively. Let $X = C_{(u,w)}$ be the connected component of $\Lambda_1 \times \Lambda_2$ containing the vertex (u, w), then, since X is a covering, we have that $p_1 \times p_2(\pi_1(X, (u, w))$ and $p_1 \times p_2(\pi_1(C(X), (u, w)))$ have the same reduced rank, by definition of core. Now, consider $C_r(X)$. Recall that $C_r(X) = C(x)\backslash(P\backslash r(u, w))$, so $|E(C_r(X))| - |V(C_r(X))| = (|E(C(X)| + n) - (|V(C(X))| + (n + 1) - 1)$. So $p_1 \times p_2(\pi_1(C(X), (u, w)))$ and $p_1 \times p_2(\pi_1(C_r(X), r(u, w)))$ have the same reduced rank, equal to the reduced rank of $S \cap L$.

Then, if $C_r(X)$ is a tree $0 \leq 2\tilde{rk}(S)\tilde{rk}(L)$. Otherwise, summing up the degrees of (u_i, w_j), the vertices of $C_r(X)$, we get:

$$\sum_{(u_i,wj)\in C_r(X)} (d(u_i, wj) - 2) \underset{\text{lemma 2.1}}{=} 2\tilde{rk}(S \cap L)$$

And then, using **Lemma 2.2**

$$\sum_{(u_i,wj)\in C_r(X)} (d(u_i, wj) - 2) \leq \sum_{(u_i,wj)\in C_r(X)} (d(u_i) - 2)(d(w_j - 2)) \leq$$

$$\sum_{(u_i,wj)\in C(X)} (d(u_i) - 2)(d(w_j - 2)) \leq \sum_{u_i} (d(u_i) - 2) \sum_{w_j} (d(w_j - 2)).$$

Finally, using **Lemma 2.1** again:

$$2\widetilde{rk}(S \cap L) \leq \sum_{u_i}(d(u_i) - 2) \sum_{w_j}(d(w_j - 2)) = 2\widetilde{rk}(S)2\widetilde{rk}(L)$$

Thus:

$$\widetilde{rk}(S \cap L) \leq 2\widetilde{rk}(S)\widetilde{rk}(L)$$

□

This bound was improved by

$$\widetilde{rk}(S \cap L) \leq \widetilde{rk}(S)\widetilde{rk}(L)$$

few years ago. An algebraic proof can be found in [4]. Unfortunately, a proof using only graph theory is not known yet.

2.5 The cosets problem

Theorem 2.7. *Given S, a subgroup of F_k, then the cosets of S, zS for $z \in F_k$ are in one to one correspondence with the vertices of $\psi(S)$.*

Proof.
Let $p : (\psi(S), v) \to R_k$ be the connected covering of S. Let $z \in F_k$ be a representant of the coset zS. Define $f_{\psi(S)}$, our bijection from the cosets of S to the vertices of $\psi(S)$, as follows. Take the path P in $\psi(S)$ with $p(P) = z$ and v as initial vertex. Let w be the terminal vertex of P. Then define $f_{\psi(S)}(zS) = w$. To see that this construction does not depen on the representant of the class, take z' another element of the coset zS, and P' a path in $\psi(S)$ with $p(P') = z'$ and w' as terminal vertex. Then, since $zz'^{-1} = p(PP'^{-1}) \in S$, PP'^{-1} must be a closed path and then the terminal vertex of P must be the inital vertex of P'^{-1} thus $w = w'$. To see that $f_{\psi(S)}$ is injective, suppose that exists a coset yS different from zS with $f(zS) = f(yS) = w$, but then there exists a path Q with $p(Q) = y$, v as initial vertex and w as terminal vertex, thus $y \in zS$ and then $zN = yN$ using that the cosets are disjoint. On the other hand, $f_{\psi(S)}$ is surjective because $\psi(S)$ is a connected graph. □

A corollary of this theorem is that if S is a subgroup of finite index, i.e. the number of cosets of F_k/S is finite, then the number of vertices of $\psi(S)$

is also finite and vice versa. If the number of vertices of $\psi(S)$ is finite then $\psi(S)$ does not contain any a_j-branches of the k-Cayley graph, and then, once we get an immersion from the flower graph of S applying foldings, that immersion must be a covering. Otherwise we should attach a_j-branches in order to get a covering, a contradiction to the finiteness of $\psi(S)$.

We can also wonder whether the intersection of cosets is again a coset and, in this case, how we can obtain an element of that intersection. The answer of the first question is afirmative:

Lemma 2.3. *Given S, L, subgroups of F_k, and $a, b \in F_k$ then $aS \cap bL$ is empty or the coset $z(S \cap L)$ for any $z \in aS \cap bL$.*

Proof.
Suppose that $aS \cap bL$ is not empty, and let $z \in aS \cap bL$. Then exists $s \in S$, $l \in L$ with $z = as = bl$. Let $n \in S \cap L$, then $zn = asn \in aS$ and $zn = bln \in L$ and then $z(S \cap L) \leq aS \cap bL$. On the other hand, for $m \in aS \cap bL$ with $m = as' = bl'$ we have that $m = zs^{-1}s'$ and $m = zl^{-1}l'$, so, since $s^{-1}s' \in S$ and $l^{-1}l' \in L$, $s^{-1}s' = l^{-1}l' \in S \cap L$, thus $aS \cap bL \leq z(S \cap L)$. Then $z(S \cap L) = aS \cap bL$. $\square$

Now, for $p_1 : (\Lambda_1, u) \to F_k$ and $p_2 : (\Lambda_2, w) \to F_k$ two connected coverings with fundamental group S and L respectively, we consider $f_{\Lambda_1}(aS)$, $f_{\Lambda_2}(bL)$ with the f defined in **Theorem 2.7**. So $f_{\Lambda_1}(aS) \in V(\Lambda_1)$ and $f_{\Lambda_2}(bL) \in V(\Lambda_2)$. Then, consider the product graph $\Lambda_1 \times \Lambda_2$. Since the connected component containing (u, w), $C_{(u,w)} \subseteq \Lambda_1 \times \Lambda_2$, is a covering with induced fundamental group $S \cap L$, if the vertex $(f_{\Lambda_1}(aS), f_{\Lambda_2}(bL))$ lies in a different connected component, using that the vertices of $C_{(u,w)}$ are in one to one correspondance with the cosets of $S \cap L$, we get that $aS \cap bL$ is empty. If $(f_{\Lambda_1}(aS), f_{\Lambda_2}(bL))$ lies in $C_{(u,w)}$, $f^{-1}_{\Lambda_1 \times \Lambda_2}(f(aS), f(bL))$ is a coset $z(S \cap L)$ with z the image by $p_1 \times p_2$ of a path with (u, w) as initial vertex and $(f_{\Lambda_1}(aS), f_{\Lambda_2}(bL))$ as terminal vertex.

2.6 Residually finite groups

In this section we prove that F_k is residually finite:

Definition 2.5. *A group G is residually finite if for any g element of G different from 1, there exists a normal subgroup of G of finite index not containing g.*

First of all, for $g \in F_k$ let's construct a finite covering whose induced fundamental group does not contain g. Let $\{a_1, \dots, a_k\}$ be the generators of F_k, then we can think of g as a tuple of elements of $\{a_1, \dots, a_k\}^{\pm}$, $g = a_{i_1}^{\pm} \dots a_{i_n}^{\pm}$. Now, construct the covering as follows. First, take two line graphs P, P' both of lenght n, with edges $e_1 \dots e_n$ and $e'_1 \dots e'_n$ respectively and initial vertices v, v' and terminal vertices w, w'. Label the edges e_j and e'_j, for $1 \leq j \leq n$, with the letter a_{i_j}. Orientate the edge e_j forwards with respect to v if the exponent of a_{i_j} in g is positive, and backwards otherwise. Orientate the edge e'_j, for $1 \leq j \leq n$, in the opposite way: backwards with respect to v' if the exponent of a_{i_j} in g is positive, and forwards otherwise. Then we have a map $p : P \cup P' \to R_k$ that is an inmersion because the element g was a simplified word. Now identify v with v' and w with w' and call this new graph Q. Observe that $q : Q \to R_k$ is still an immersion: the edge e_1 and e'_1 have the same label but opposite orientations with respect to v. The same for e_n and e'_n. Observe also that every vertex of Q has degree 2. Finally, we are going to add edges to Q in order to get a covering. We do this in two steps:
(a) For every x vertex of Q let a_r and a_s be the labels of the two edges incident to x. a_r and a_s could be the same label. Then, for every letter a_j from the set $\{a_1, \dots, a_k\} \backslash \{a_r, a_s\}$, add a x-loop, i.e. and edge (x, x), to Q with label a_j
(b) Consider $x = \iota(e_i) = \tau(e_{i-1})$ and $y = \iota(e'_i) = \tau(e'_{i-1})$ for $2 \leq i \leq n$, the i-th vertex of P and P' respectively. Observe that e_i have the same label as $e_{i'}$ and opposite orientations and the same for e_{i-1} and e'_{i-1}. Then, if the label of e_i is the same as the label of e_{i-1}, Q is, after step (a), a covering in x and y. Otherwise, let a_r, a_s be the labels of e_{i-1} and e_i respectively. Then, add two edges, e, f, to Q, both incident to x and y, with label a_r and a_s, respectively. If e_{i-1} is oriented forwards/backwards with respect to x then e'_{i-1} will be oriented backwards/forwards with respect to y, so orientating the edge e backwards/forwards with respect to $\iota(e_i)$ we will get that it has the opposite orientation with respect to $\iota(e'_i)$. The same for the edge f and so Q will be an immersion in x and y.

Then, after those new edges are added to Q, we get a covering that doesn't contain the word g because of the path P'

Theorem 2.8. *F_k is residually finite.*

Proof.

For $a \in F_k$, construct $p : (\Gamma_{\bar{a}}, v) \to F_k$ the covering defined before, with $a \notin S = \pi_1(\Gamma_{\bar{a}}, v)$. Let $\{v_1, \ldots v_n\}$, with $v_1 = v$, be the vertices of Γ. Then, define recursively a sequence of coverings $p_j : (\Gamma_j, w_j) \to F_k$, as follows:

$$(\Gamma_1, w_1) = (\Gamma_{\bar{a}}, v),\ (\Gamma_i, w_i) = C_{(w_{i-1}, v_i)}((\Gamma_{i-1}, w_{i-1}) \times (\Gamma_1, v_i))$$

and

$$p_1 := p,\ p_i := p_{i-1} \times p_{1\|C_{(w_{i-1}, v_i)}}$$

where w_{i-1} is the distinguished vertex of Γ_{i-1} and $C_{(w_{i-1}, v_i)}$ is the connected component containing (w_{i-1}, v_i). Then, we have that

$$p_1(\pi_1(\Gamma_1, w_1)) = S,\ p_i(\pi_1(\Gamma_i, w_i)) = p_{i-1}(\pi_1(\Gamma_{i-1}, w_{i-1})) \cap p_1(\pi_1(\Gamma_1, v_i)) =$$

$$p_1(\pi_1(\Gamma_{i-1}, w_{i-1})) \cap P_i S P_i^{-1}$$

where P_i is the reduced path from v to v_i. And finally,

$$p_n(\pi_1(\Gamma_n, w_n)) = S \cap P_2 S P_2^{-1} \cap \ldots \cap P_{n-1} S P_{n-1}^{-1} \cap P_n S P_n^{-1}.$$

We have already seen that every connected component of the product of two coverings is a covering as well, so (Γ_n, w_n) is a covering. Thus the fundamental group induced by (Γ_n, w_n) is the intersection of all the conjugates of S: we have as many conjugates as cosets of S, and using the previous section S has a finite number of cosets since Γ has finitely many vertices. Then S is normal. On the other hand, using previous section again, $p_n(\pi_1(\Gamma_n, w_n))$ is of finite index because Γ_n has a finite number of vertices. Finally, since $a \notin p_1(\pi_1(\Gamma_1, w_1))$, $a \notin p_n(\pi_1(\Gamma_n, w_n))$, and then $p_n(\pi_1(\Gamma_n, w_n))$ is a normal subgroup of finite index and not containing a. $\square$

2.7 The basis problem

One could wonder whether, like in linear algebra, given m generators of a subgroup of rank k, with $m \geq k$, there exists a subset of this generators of cardinality k, generating the same subgroup. In other words, if for any finite A set of elements of F_n we could find a free subset of A generating the same subgroup as A. That is in general not possible, and, in fact, we will show an example where an arbitrary large number of elements generates a subgroup of rank 2 but, if we delete just one of the elements, the remaining elements don't generate the same subgroup.

Example 2.4. *Consider the labeled graph:*

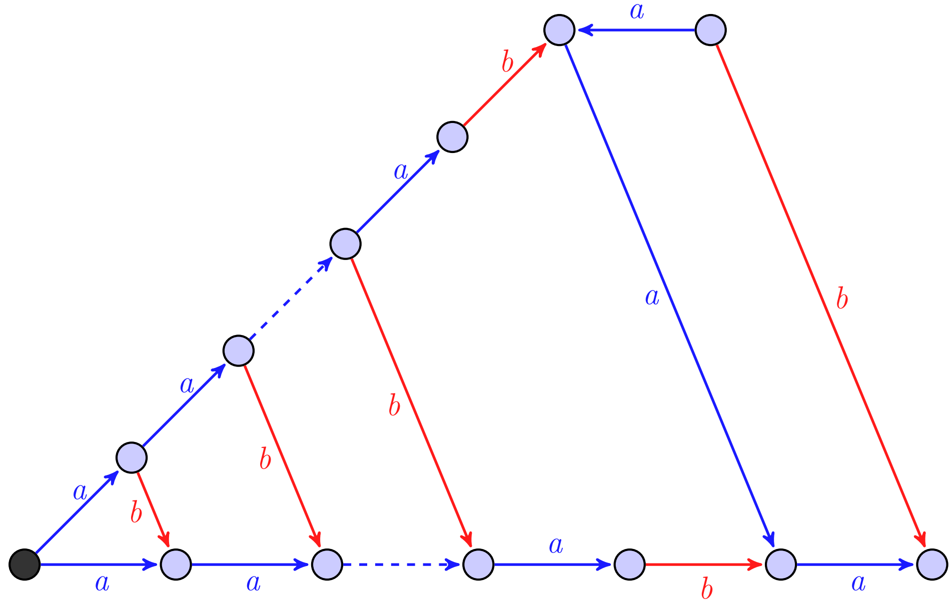

For $n \geq 3$ we consider the previous graph, with n the number of edges in the base of the left triangle. To compute generators of the induced fundamental group of this graph, consider the following maximal tree:

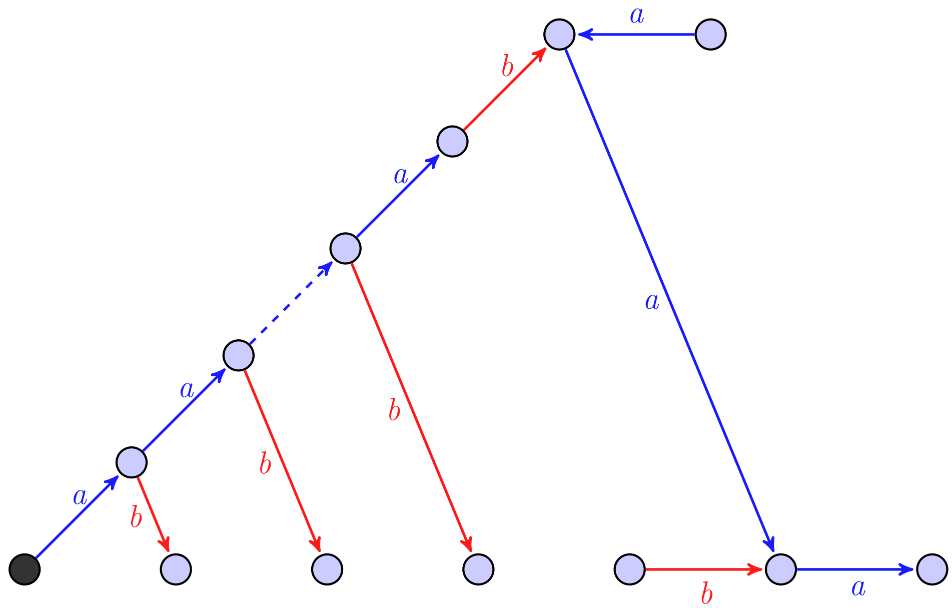

With this tree the generators are:

$$\{aba^{-1}\} \cup \{a^i ba^{-1}b^{-1}a^{-(i-1)}\}_{i=2}^{n-2} \cup \{a^{n-1}bab^{-1}a^{-1}b^{-1}a^{-(n-2)}\} \cup$$

$$\{a^{n-1}ba^{-1}ba^{-2}b^{-1}a^{-(n-1)}\}$$

Thus, it has n generators. But now, observe that the two edges incident to the base vertex are inadmissible. Once we have identified them we will get a new pair of inadmissible edges and finally the triangle will collapse to a line:

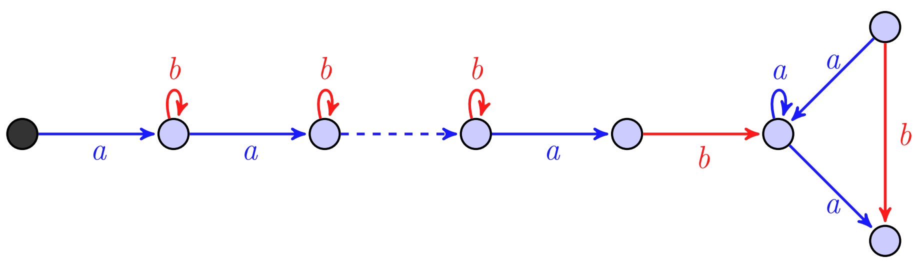

The loop labaled by a on the right and the adjacent edges labeled by a are inadmissible edges also, so we can apply more foldings:

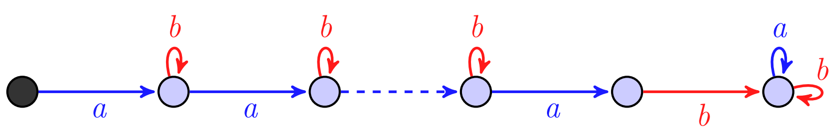

Finally, since we have a double loop in the same vertex, all the edges of the graph collapse in the rose graph R_2, so the elements above generate F_2. Now, suppose we erase one of those elements. This is equivalent to erasing one of the edges of the base of the triangle or the egde labeled by b on the right. But then, the chain of foldings would stop just where the edge has been erased: the triangle woudn't collapse in the line graph and then we wound't get R_k. Thus, we have an arbitrarily large set of elements that generates F_2 but the set minus any element doesn't generate F_2.

References

[1] John R. Stallings, *Topology of Finite Graphs*. Inventione mathematicae 1983.

[2] L.Bartholdi and P.V.Silva *Rational subsets of groups*. Inventione mathematicae 2010.

[3] Allan Hatcher, *Algebraic Topology*. Cambridge University Press 2002.

[4] Warren Dicks, *Simplified Mineyev*. http://mat.uab.es/ dicks/pub.html 2011.