Groups Geom. Dyn. 18 (2024), 235–264
DOI 10.4171/GGD/746

A fine property of Whitehead's algorithm

Dario Ascari

Abstract. We develop a refinement of Whitehead's algorithm for primitive words in a free group. We generalize to subgroups, establishing a strengthened version of Whitehead's algorithm for free factors. These refinements allow us to prove new results about primitive elements and free factors in a free group, including a relative version of Whitehead's algorithm and a criterion that tests whether a subgroup is a free factor just by looking at its primitive elements. We develop an algorithm to determine whether or not two vertices in the free factor complex have distance d for $d = 1, 2, 3$, as well as $d = 4$ in a special case.

1. Introduction

An algorithm to determine whether an element of a free group is primitive or not was first found by Whitehead in 1936; it is based on the following theorem.

Theorem A (Whitehead). *Let w be a cyclically reduced word, which is primitive but not a single letter. Then there is a Whitehead automorphism φ such that the cyclic length of $\varphi(w)$ is strictly smaller than the cyclic length of w.*

For a proof of the Theorem A, the reader can refer to Whitehead's paper [11]. Theorem A has been proved and studied in several different ways over the years. Whitehead's original proof involves working with three-manifolds, performing surgery on embedded paths and surfaces (see [11]). Of particular importance are the peak-reduction techniques introduced by Rapaport in [9]; see also [6] for a simplified version of the argument. Another surprisingly short proof, based on Stallings' folding operations, appeared recently in [5].

Let $x_1, \dots, x_n$ be a fixed basis for F_n. We recall that a Whitehead automorphism φ is an automorphism such that, for some $a \in \{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\}$, we have $\varphi(a) = a$ and $\varphi(x_j) \in \{x_j, ax_j, x_j\overline{a}, ax_j\overline{a}\}$ for each other generator $x_j \neq a, \overline{a}$. A generic element $w \in F_n$ consists of a (reduced) sequence of symbols in $\{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\}$, and in order to obtain the image $\varphi(w)$, we can just apply φ letter by letter to the sequence of symbols (and then reduce the resulting word). In the present paper, we shall build on the following refinement of Whitehead's Theorem A.

2020 Mathematics Subject Classification. Primary 20E05; Secondary 20F65.
Keywords. Free groups, Whitehead's algorithm, free factor.

Theorem B (Theorem 3.7). *The automorphism in Theorem A can be chosen in such a way that every letter a or $\overline{a}$ that is added when we apply φ to w letter by letter, immediately cancels (in the cyclic reduction process).*

Theorem B can be deduced fairly directly from Whitehead's original argument [11]; however, we were not able to find this statement in the literature. It can also be derived with the techniques of [5], as will be shown in the body of the present paper. It is difficult to imagine how Theorem B might be proved with the peak-reduction techniques of [9].

We make use of Theorems A and B to prove the following theorem about free factors in a free group.

Theorem C (Theorem 4.1). *Let $H \leq F_n$ be a finitely generated subgroup. Suppose that every element of H which is primitive in H is also primitive in F_n. Then H is a free factor.*

We point out that the additional property of Whitehead's algorithm really plays a key role in the proof of Theorem C.

We also generalize the fine property of Theorem B to subgroups. Generalizations of Whitehead's algorithm for free factors already exist, as shown in [4], but the known proofs are based on the peak-reduction techniques, which, as we have noted, do not seem appropriate for our refinement. We introduce the concept of Whitehead graph for a subgroup, and we generalize the ideas of [5] in order to get a refined statement of Whitehead's algorithm for subgroups. The standard statement is the following Theorem D, and we add the fine property into Theorem E.

Theorem D (See Theorems 5.5 and 5.6). *Let $H \leq F_n$ be a free factor, and suppose* $\mathrm{core}(H)$ *has more than one vertex. Then there is a Whitehead automorphism φ such that* $\mathrm{core}(\varphi(H))$ *has strictly fewer vertices and strictly fewer edges than* $\mathrm{core}(H)$.

Theorem E (Theorem 5.7). *The automorphism in Theorem D can be chosen in such a way that* $\mathrm{core}(\varphi(H))$ *can be obtained from* $\mathrm{core}(H)$ *by means of a quotient that collapses some of the edges to points whilst preserving the labels and orientations on the other edges.*

This additional property turns out to have several interesting features, and in particular, it behaves well with respect to subgroups (see Lemmas 5.17 and 5.19 in the body of the paper). It also allows us to deduce a relative version of Whitehead's algorithm. Let $x_1, \dots, x_n$ be a fixed basis for F_n.

Theorem F (Theorem 5.21). *Let $w \in F_n$ be a primitive element that is not a single letter. Suppose there is an automorphism $\theta\colon F_n \to F_n$ such that $\theta(\langle x_1, \dots, x_k \rangle) = \langle x_1, \dots, x_k \rangle$ and $\theta(w) = x_{k+1}$. Then there is a Whitehead automorphism φ such that*

(i) $\varphi(x_i) = x_i$ *for* $i = 1, \dots, k$.

(ii) *The length of $\varphi(w)$ is strictly smaller than the length of w.*

(iii) *Every letter, which is added to w when applying φ to w letter by letter, immediately cancels (in the free reduction process).*

We then use the techniques developed in order to investigate the structure of the free factor complex FF_n. The free factor complex of a free group is analogous to the curve complex of a surface. A famous rigidity theorem of Ivanov states that the isometries of the curve complex are essentially the mapping class group of the surface; in the same way, there is a rigidity theorem due to Bestvina and Bridson stating that the isometries of FF_n are essentially the outer automorphisms of F_n, see [1]. Just as the curve complex turned out to be hyperbolic, FF_n is hyperbolic too: in [2], Bestvina and Feighn study the large-scale geometry of FF_n and give a detailed description of lines which are geodesics up to a reparametrization and up to distance C. Unfortunately, the constant C is quite large, so those techniques do not give much information about the local geometry of FF_n. In the present paper, we investigate existence of an algorithm for computation of the exact distance in FF_n. There are algorithms to compute the distance between any two points in the curve complex of a surface, but the same question for FF_n remains open. We are able to determine whether two vertices are at distance d for $d = 1, 2, 3$; we are also able to do that for $d = 4$ in the particular case when one of the two vertices represents a conjugacy class of free factors of rank $n - 1$.

In order to do this, we make use of the tools developed earlier in this paper, together with an idea which appeared in [3], which is the following. Given a subgroup $H \leq F_n$ and an element $w \in H$, even if we have a bound on the size (i.e., the number of edges) of $\mathrm{core}(H)$, we cannot really control the cyclic length of w; we instead look at the map $i_*\colon \mathrm{core}(\langle w \rangle) \to \mathrm{core}(H)$ induced by the inclusion, and we look at the image $\mathrm{im}(i_*)$. We have that $\mathrm{im}(i_*)$ is a subgraph of $\mathrm{core}(H)$, and we are able to bound the size of $\mathrm{core}(H)$: it follows that $\mathrm{im}(i_*)$ can only take a finite number of values, and is thus much easier to control. We notice that the same observation remains true when we replace the word w with a subgroup $K \leq H$, and we make use of this observation to produce algorithms to recognize whether two vertices of the free factor complex are at distance d for $d = 1, 2, 3$, and $d = 4$ in a special case.

2. Preliminaries and notations

We work inside a finitely generated free group F_n of rank n, generated by $x_1, \dots, x_n$. We write $\overline{x}_i = x_i^{-1}$. We denote by R_n the standard n-rose, i.e., the graph with one vertex $*$ and n oriented edges labeled $x_1, \dots, x_n$. The fundamental group $\pi_1(R_n, *)$ will be identified with F_n: the path going along the edge labeled x_i (with the right orientation) corresponds to the element $x_i \in F_n$.

Definition 2.1. An *F_n-labeled graph* is a graph G together with a map $f\colon G \to R_n$ sending each vertex of G to the unique vertex of R_n, and each open edge of G homeomorphically to one edge of R_n.

This means that every edge of G is equipped with a label in $\{x_1, \dots, x_n\}$ and an orientation, according to which edge of R_n it is mapped to; the map $f\colon G \to R_n$ is called *labeling map* for G.

Definition 2.2. Let G_0 and G_1 be F_n-labeled graphs. A map $h: G_0 \to G_1$ is called *label-preserving* if it sends each vertex to a vertex and each edge to an edge with the same label and orientation.

We notice that if $h: G_0 \to G_1$ is label-preserving and $f_0: G_0 \to R_n$ and $f_1: G_1 \to R_n$ are the labeling maps, then $f_1 \circ h = f_0$.

Core graph of a subgroup

Definition 2.3. Let G be a graph which is not a tree. Define its *core graph* $\mathrm{core}(G)$ as the subgraph given by the union of its non-degenerate loops, i.e., all the images $f(S^1)$ for a continuous locally injective map $f: S^1 \to G$ from the circle.

Notice that $\mathrm{core}(G)$ is connected, and every vertex has valence at least 2. We say that a graph G is *core* if $\mathrm{core}(G) = G$.

In the following, we will often consider graphs with a basepoint. We always mean that the basepoint is a vertex of the graph.

Definition 2.4. Let $(G, *)$ be a pointed graph which is not a tree. Define its *pointed core graph* $\mathrm{core}_*(G)$ as the subgraph given by the union of all the images $f([0,1])$ for a continuous locally injective map $f: [0,1] \to G$ with $f(0) = f(1) = *$.

We will often abbreviate $(G, *)$ to G. For a pointed graph G, there is a unique shortest path σ (either trivial or embedded) connecting the basepoint to $\mathrm{core}(G)$; the graph $\mathrm{core}_*(G)$ consists exactly of the union $\mathrm{core}(G) \cup \mathrm{im}(\sigma)$.

Given a nontrivial subgroup $H \leq F_n$, we can build the corresponding pointed covering space $p: (\mathrm{cov}(H), *) \to (R_n, *)$; this means that $\mathrm{cov}(H)$ is the unique covering space for R_n such that $p_*(\pi_1(\mathrm{cov}(H), *)) = H$ as subgroups of $\pi_1(R_n, *) = F_n$. Define the core graph $\mathrm{core}(H)$ and the pointed core graph $\mathrm{core}_*(H)$ to be the core and the pointed core of $(\mathrm{cov}(H), *)$, respectively. Notice that the labeling map $f: \mathrm{core}_*(H) \to R_n$ induces an injective map $\pi_1(f): \pi_1(\mathrm{core}_*(H), *) \to F_n$, and the image of such map is exactly the subgroup H. We observe that conjugate subgroups have the same core graph, but distinct pointed core graphs. We also observe that H is finitely generated if and only if $\mathrm{core}(H)$ is finite (and if and only if $\mathrm{core}_*(H)$ is finite).

Stallings' folding

We will assume that the reader has some confidence in the classical Stallings' folding operation, for which we refer to [10]. Let us briefly recall the main properties that we are going to use.

Let G be a finite connected F_n-labeled graph and suppose there are two distinct edges e_1, e_2 with endpoints v, v_1 and v, v_2, respectively. Suppose that e_1 and e_2 have the same label and orientation. We can identify v_1 with v_2, and e_1 with e_2: we then get a quotient map of graphs $q: G \to G'$.

Definition 2.5. The quotient map $q: G \to G'$ is called *Stallings' folding*.

We notice that q is label preserving. Fix a basepoint $* \in G$, which induces a basepoint $* \in G'$; then the map $\pi_1(f)\colon \pi_1(G) \to \pi_1(R_n)$ and the map $\pi_1(f')\colon \pi_1(G') \to \pi_1(R_n)$ give the same subgroup $\pi_1(f)(\pi_1(G)) = \pi_1(f')(\pi_1(G')) \le \pi_1(R_n) = F_n$. The map $\pi_1(q)\colon \pi_1(G) \to \pi_1(G')$ is surjective; however, it is not injective in general.

Definition 2.6. A Stallings' folding $q\colon G \to G'$ is called *rank-preserving* if the induced map $\pi_1(q)$ is an isomorphism.

Being rank-preserving is equivalent to the requirement $v_1 \neq v_2$ (i.e., that we are identifying two distinct vertices) (see also Figure 1). In fact, the rank of the fundamental group of a finite connected graph is $E - V + 1$, where E is the number of edges and V is the number of vertices; during a folding operation, the number of edges always decreases by exactly one; if we identify two distinct vertices, then the number of vertices decreases by one too, and thus the rank is preserved; if the vertices v_1, v_2 coincide, then the number of vertices remains the same, and thus the rank decreases by one. Notice that being rank-preserving does not depend on the choice of the basepoint.

Figure 1. Examples of configurations where a folding operation is possible. The two examples on the left produce rank-preserving folding operations. The two examples on the right produce non-rank-preserving folding operations.

Given a finite F_n-labeled graph G, we can successively apply folding operations to G in order to get a sequence $G = G^{(0)} \to G^{(1)} \to \cdots \to G^{(l)}$. Notice that the number of edges decreases by 1 at each step, and thus the length of any such chain is bounded (by the number of the edges of G).

Proposition 2.7. *Let G be a finite connected F_n-labeled graph, and let $G = G^{(0)} \to G^{(1)} \to \cdots \to G^{(l)}$ be a maximal sequence of folding operations. Also, fix a basepoint $* \in G$, inducing a basepoint $* \in G^{(i)}$. Then we have the following:*

(i) *Each such sequence has the same length l and the same final graph $G^{(l)}$.*

(ii) *For each such sequence, and for each label, the sequence has the same number of folding operations involving edges with that label.*

(iii) *Each such sequence has the same number of rank-preserving folding operations.*

(iv) *For $i = 1, \dots, l$, let $f^{(i)}\colon G^{(i)} \to R_n$ be the labeling map. Then the image of the map $\pi_1(f^{(i)})\colon \pi_1(G^{(i)}) \to \pi_1(R_n)$ is the same subgroup $H \le F_n$ for every $i = 1, \dots, l$.*

(v) *For* $i = 1, \dots, l$, *there is a unique label-preserving map* $h^{(i)}\colon G^{(i)} \to \mathrm{cov}(H)$ *preserving the basepoint. The image* $\mathrm{im}(h^{(i)})$ *is the same subgraph of* $\mathrm{cov}(H)$ *for every* $i = 1, \dots, l$.

(vi) *The map* $h^{(l)}$ *is an embedding of* $G^{(l)}$ *as a subgraph of* $\mathrm{cov}(H)$. *Moreover,* $G^{(l)}$ *contains* $\mathrm{core}_*(H)$.

An F_n-labeled graph is called *folded* if no folding operation is possible on G.

Definition 2.8. Let G be a finite connected F_n-labeled graph. Define its *folded graph* $\mathrm{fold}(G)$ to be the F_n-labeled graph $G^{(l)}$ obtained from any maximal sequence of folding operations as in Proposition 2.7.

The following lemma gives us information about vertices of valence one that may appear along a chain of folding operations.

Lemma 2.9. *Let* $(G, *)$ *be a finite connected* F_n*-labeled graph, and let* $G \to G^{(1)} \to \dots \to G^{(l)}$ *be any maximal sequence of folding operations as in Proposition 2.7, with maps* $p_k\colon G \to G^{(k)}$ *given by the composition of folding operations. Suppose that for some vertex* v *of* G*, the vertex* $p_k(v)$ *of* $G^{(k)}$ *has valence one. Then all the edges going out of* v *in* G *have the same label and orientation.*

Proof. Suppose that $p_k(v)$ has valence one in $G^{(k)}$. Let e, e' be edges of G going out of v: we observe that $p_k(e)$, $p_k(e')$ are both edges of $G^{(k)}$ going out of $p_k(v)$. But then they must coincide, since there is only one edge of $G^{(k)}$ going out of $p_k(v)$. In particular, $p_k(e)$, $p_k(e')$ have the same label and orientation, and thus e, e' have the same label and orientation too, since p_k is label-preserving. ■

Corollary 2.10. *Let* $(G, *)$ *be a finite connected* F_n*-labeled graph, and let* $G \to G^{(1)} \to \dots \to G^{(l)}$ *be any maximal sequence of folding operations as in Proposition 2.7. Suppose for each vertex* v *of* G*, there are two edges going out of* v *with different labels, or with the same label but different orientations. Then there is no valence-*1 *vertex in any graph of the sequence.*

Proof. Suppose some graph $G^{(k)}$ contains a valence-1 vertex u. Let $p_k\colon G \to G^{(k)}$ be the map given by the composition of the folding operations: since p_k is surjective, we can find a vertex v of G with $p_k(v) = u$. But then by Lemma 2.9, all the edges going out of v have the same label and orientation, contradicting the hypothesis. ■

Suppose we are given a finite set of reduced words $w_1, \dots, w_k \in F_n$ of lengths $l_1, \dots, l_k$, and let $H = \langle w_1, \dots, w_k \rangle$. We can construct the graph G given by a basepoint $*$ and pairwise disjoint loops $\gamma_1, \dots, \gamma_k$ starting and ending at the basepoint. The loop γ_i is subdivided into l_i edges, labeled and oriented according to the letters of the word w_i, in such a way that, when going along γ_i, we read exactly the word w_i. We have the following lemma.

Lemma 2.11. *We have that* $\mathrm{core}_*(H) = \mathrm{fold}(G)$. *Moreover, along the chain of folding operations, we never have a valence-*1 *vertex, except possibly for the basepoint.*

Proof. Proposition 2.7 tells us that $\mathrm{fold}(G)$ is a finite subgraph of $\mathrm{cov}(H)$ and that it contains $\mathrm{core}_*(H)$. Lemma 2.9 tells us that $\mathrm{fold}(G)$ has no valence-1 vertex, except possibly for the basepoint. But a finite subgraph of $\mathrm{cov}(H)$, containing $\mathrm{core}_*(H)$, and with no valence-1 vertex except possibly the basepoint, must be equal to $\mathrm{core}_*(H)$. Thus $\mathrm{fold}(G) = \mathrm{core}_*(H)$, as desired. ■

In particular, $w_1, \dots, w_n$ are a basis for F_n if and only if, with the above construction, the resulting F_n-labeled graph $\mathrm{fold}(G)$ is the standard n-rose R_n. In that case, each folding operation has to be rank-preserving, since the fundamental group of G has the same rank as the fundamental group of $\mathrm{fold}(G) = R_n$.

Primitive elements and free factors

We will be interested in the study of primitive elements and free factors.

Definition 2.12. An element $w \in F$ is called *primitive* if it is part of some basis for the group.

Definition 2.13. A subgroup $H \leq F$ is called a *free factor* if some basis (equivalently, every basis) for H can be extended to a basis for F.

Notice that an element $w \in F$ is primitive if and only if the cyclic subgroup $\langle w \rangle$ is a free factor; in this sense, the notion of free factor is a natural generalization of the notion of primitive element.

The following lemma is immediate.

Lemma 2.14. *Let G be a pointed graph, and let G' be a connected subgraph containing the basepoint. Then the map $\pi_1(G') \to \pi_1(G)$ induced by the inclusion is injective, and $\pi_1(G')$ is a free factor in $\pi_1(G)$.*

The following proposition turns out to be very useful in several situations.

Proposition 2.15. *Let $K \leq F_n$ be a finitely generated subgroup, and let $H \leq F_n$ be a free factor. Then $H \cap K$ is a free factor in K.*

Proof. Suppose H has rank r; without loss of generality, we can assume that $H = \langle x_1, \dots, x_r \rangle \leq F_n$.

Let $G = \mathrm{core}_*(K)$ be the basepointed F_n-labeled graph which represents K. A word w belongs to $H \cap K$ if and only if it can be represented by a path inside G which starts and ends at the basepoint, and which only crosses edges labeled with $x_1, \dots, x_r$. Consider the subgraph $G' \subseteq G$ which is given by the union of the basepoint and of all such paths. Then $\pi_1(G')$ is exactly $H \cap K$.

But since G' is a subgraph of G, we have that $\pi_1(G')$ is a free factor in $\pi_1(G)$, meaning that $H \cap K$ is a free factor in K, as desired. ■

We can easily obtain several corollaries from the above proposition.

Corollary 2.16. *Consider the standard inclusion $F_k = \langle x_1, \dots, x_k \rangle \le \langle x_1, \dots, x_n \rangle = F_n$ with $k < n$. Let $w \in F_k$. Then w is primitive in F_k if and only if w is primitive in F_n.*

Corollary 2.17. *Consider the standard inclusion $F_k = \langle x_1, \dots, x_k \rangle \le \langle x_1, \dots, x_n \rangle = F_n$ with $k < n$. Let $H \le F_k$. Then H is a free factor in F_k if and only if H is a free factor in F_n.*

In particular, when talking about primitive elements and free factors, it often makes sense to omit mention of the ambient group.

Corollary 2.18. *Let $H, H' \le F_n$ be free factors. Then $H \cap H'$ is a free factor.*

In the above statement, we mean that it is a free factor in F_n, and also in both H and H'.

3. A fine property of Whitehead's algorithm

Whitehead automorphisms and Whitehead graph

Definition 3.1. Let $a \in \{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\}$, and let $A \subseteq \{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\} \setminus \{a, \overline{a}\}$. Define the *Whitehead automorphism* $\varphi = (A, a)$ as the automorphism given by $a \mapsto a$ and

$$\begin{cases} x_j \mapsto x_j & \text{if } x_j, \overline{x}_j \notin A, \\ x_j \mapsto a x_j & \text{if } x_j \in A \text{ and } \overline{x}_j \notin A, \\ x_j \mapsto x_j \overline{a} & \text{if } x_j \notin A \text{ and } \overline{x}_j \in A, \\ x_j \mapsto a x_j \overline{a} & \text{if } x_j, \overline{x}_j \in A. \end{cases}$$

The letter a will be called the *acting letter*, and the set A will be the set of letters we *act on*. Notice that our notation for Whitehead automorphisms is slightly different from the one found in Lyndon and Schupp's book [7]: they choose to include the acting letter a inside the set A, while we prefer not to do so.

Let w be a cyclically reduced word, whose reduced form is $w = b_1 \dots b_l$, where we have $b_j \in \{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\}$. Let $\varphi = (A, a)$ be a Whitehead automorphism. We can substitute each b_j with the sequence $\varphi(b_j)$ (which is either b_j or ab_j or $b_j\overline{a}$ or $ab_j\overline{a}$): this produces a new writing $\varphi(b_1) \dots \varphi(b_l)$ which represents the word $\varphi(w)$; this writing will not be cyclically reduced in general. In what follows, when we speak of the free or cyclic reduction, we mean any sequence of moves in which an adjacent pair of letters $x_i \overline{x}_i$ or $\overline{x}_i x_i$ is replaced by the empty word.

Lemma 3.2. *Let $w = b_1 \dots b_l$ be a cyclically reduced word, and let $\varphi = (A, a)$ be a Whitehead automorphism. Then, in the process of cyclic reduction for the sequence $\varphi(b_1) \dots \varphi(b_l)$, no letter different from a gets cancelled.*

Proof. For simplicity of notation, we prove the proposition only for the free reduction process; the proof for the cyclic reduction process is completely analogous.

Fix a process of free reduction for $\varphi(b_1)\dots\varphi(b_l)$. Suppose some cancellation takes place, involving a letter which is not a nor $\overline{a}$. Consider the first such cancellation. Suppose this cancellation involves a letter from the block $\varphi(b_j)$ and one from the block $\varphi(b_k)$, with $j < k$. Then we must have $b_k = \overline{b_j}$, and either all the letters inbetween are a, or all of them are $\overline{a}$. This means that the word w has the form either $\dots b_j a^d \overline{b_j} \dots$ or $\dots b_j \overline{a}^d \overline{b_j} \dots$ for some $d \geq 0$. Also, since the writing $b_1 \dots b_l$ was reduced, we must have $d > 0$.

We assume w has the form $\dots b_j a^d \overline{b_j} \dots$, the other case being completely analogous. If $\overline{b_j} \notin A$, then the sequence $\varphi(b_1)\dots\varphi(b_l)$ has the form $\dots b_j a^d \overline{b_j} \dots$, and at least one a letter survives between b_j and $\overline{b_j}$, and thus b_j is not allowed to cancel with $\overline{b_j}$. If $\overline{b_j} \in A$, then the sequence $\varphi(b_1)\dots\varphi(b_l)$ has the form $\dots b_j \overline{a} a^d a \overline{b_j} \dots$, and again we see that at least one a letter survives between b_j and $\overline{b_j}$, and thus b_j is not allowed to cancel with $\overline{b_j}$. This contradiction completes the proof. ∎

Definition 3.3. Let w be a cyclically reduced word. Define the *Whitehead graph* of w as follows:

(i) We have $2n$ vertices labeled $x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n$.

(ii) For every pair of consecutive letters in w, we draw an (unoriented) arc from the inverse of the first letter to the second. We also draw an arc connecting the inverse of the last letter of w to the first letter of w, as if they were adjacent.

Notice that, w being cyclically reduced, we never have any arc connecting a vertex to itself.

Definition 3.4. Let w be a cyclically reduced word. A vertex a in the Whitehead graph of w is called a *cut vertex* if it is non-isolated and at least one of the following two configurations happens:

(i) The connected component of a does not contain $\overline{a}$.

(ii) The connected component of a becomes disconnected if we remove a.

Whitehead's algorithm

We are now ready to state Whitehead's theorem and our refinement of it (Theorem 3.7).

Theorem 3.5. *Let w be a cyclically reduced word, which is primitive but not a single letter. Then the Whitehead graph of w contains a cut vertex.*

Theorem 3.6. *Let w be a cyclically reduced word, and suppose the Whitehead graph of w contains a cut vertex. Then there is a Whitehead automorphism φ such that the cyclic length of $\varphi(w)$ is strictly smaller than the cyclic length of w.*

Theorem 3.7. *The automorphism in Theorem 3.6 can be chosen in such a way that every a or $\overline{a}$ letter, which is added when we apply φ to w letter by letter, immediately cancels (in the cyclic reduction process).*

Example. Let $w = xyxyx\overline{y}z$ in F_3. Consider the automorphism $\varphi = (\{\overline{x}\}, y)$, meaning that $x \mapsto x\overline{y}$, $y \mapsto y$ and $z \mapsto z$; the word becomes $\varphi(w) = xxx\overline{y}\overline{y}z$. This one is shorter, meaning that the automorphism φ would be suitable for Theorem 3.6 applied to the word w. But φ does not satisfy Theorem 3.7, because a letter $\overline{y}$ appears between the last x and the z, and it does not cancel.

We are now going to prove Theorems 3.5, 3.6 and 3.7. The proof of Theorem 3.5 that we give is essentially contained in [5], but we will make use of variations of the argument in what follows, so we prefer to rewrite it here.

Proof of Theorem 3.5. Let w be cyclically reduced and primitive. We will assume that w contains all the letters $x_1, \dots, x_n$ at least once; otherwise, if w only contains the letters $x_1, \dots, x_k$, then we can just apply the same argument to the free factor $\langle x_1, \dots, x_k\rangle \leq \langle x_1, \dots, x_n\rangle = F_n$ (using Corollary 2.16).

Since w is primitive, we can take a basis $w = w_1, w_2, \dots, w_n$ of reduced words. We can build the graph G given by a basepoint $*$, together with a path p_i for each w_i: the path p_i goes from $*$ to $*$, and contains an edge for each letter appearing in w_i (in such a way that, moving around the path p_i, we read exactly the word w_i). Let $G(w)$ denote the subgraph of G given by the only cycle corresponding to the generator w.

We now apply a sequence of folding operations to the graph G, in order to get a sequence $G \to G' \to \cdots \to G^{(l-1)} \to G^{(l)}$ as in Proposition 2.7: each map $G^{(i)} \to G^{(i+1)}$ consists of a single folding operation, and no further folding operation can be applied to $G^{(l)}$. Since $w_1, \dots, w_n$ is a basis, we have that $G^{(l)}$ is the standard n-rose R_n. Since w_1 is cyclically reduced and $w_2, \dots, w_n$ are reduced, Corollary 2.10 yields that no graph $G^{(i)}$ contains any valence-1 vertex. A folding operation can decrease the rank of the fundamental group, but it cannot increase it; since $\pi_1(G)$ has the same rank as $\pi_1(R_n)$, we must have that each folding operation is rank-preserving.

We now look at the graph $G^{(l-1)}$: it does not contain any valence-1 vertex, and a single rank-preserving folding operation sends it to the standard n-rose. It is quite easy to see that $G^{(l-1)}$ has to be of the form described in Figure 2 below, for some $1 \leq \alpha \leq \beta \leq n$ with $\alpha < n$ (up to permutation of the letters, and up to substitution of some letter with its inverse).

We have a map of graphs $f\colon G(w) \to G^{(l-1)}$ preserving the orientations and the labels on the edges (given by the inclusion $G(w) \to G$ followed by the sequence of foldings). Suppose we have two adjacent letters $w = \dots yz \dots$; this means that $G(w)$ contains two edges labeled y and z with a common endpoint u. We either have

$$f(u) = v \quad \text{or} \quad f(u) = v',$$

meaning that $\overline{y}$ and z are either both in $\{x_1, \overline{x}_1, x_2, \overline{x}_2, \dots, x_\alpha, \overline{x}_\alpha, x_{\alpha+1}, \dots, x_\beta\}$ or both in $\{\overline{x}_1, \overline{x}_{\alpha+1}, \dots, \overline{x}_\beta, x_{\beta+1}, \overline{x}_{\beta+1}, \dots, x_n, \overline{x}_n\}$. This tells us that, if we remove the vertex $\overline{x}_1$ from the Whitehead graph of w, we get the disjoint union $V \sqcup V'$ of two separate graphs: V with vertices $x_1, x_2, \overline{x}_2, \dots, x_\alpha, \overline{x}_\alpha, x_{\alpha+1}, \dots, x_\beta$ and V' with vertices $\overline{x}_{\alpha+1}, \dots, \overline{x}_\beta, x_{\beta+1}, \overline{x}_{\beta+1}, \dots, x_n, \overline{x}_n$.

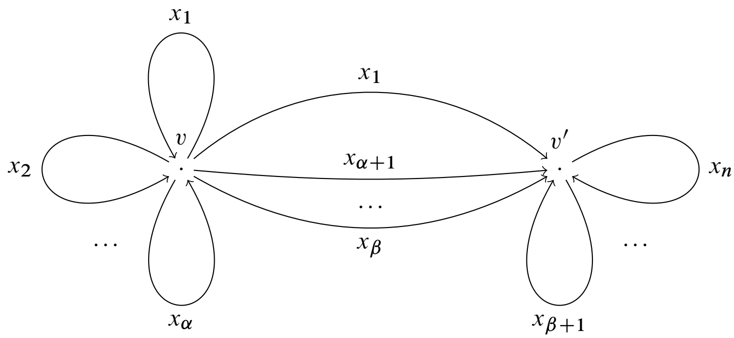

Figure 2. The generic graph $G^{(l-1)}$. This contains exactly one edge with each label, except for the two edges labeled x_1. Those two edges have to be folded in order to obtain the n-rose.

If the image $f(G(w)) \subseteq G^{(l-1)}$ crosses both the edges labeled x_1, then in the Whitehead graph of w we have that $\overline{x}_1$ is connected to at least one vertex in V and to one vertex in V'; this means $\overline{x}_1$ is a cut vertex (because it satisfies condition (ii) of Definition 3.4). If $f(G(w))$ crosses the edge labeled x_1 with distinct endpoints, but not the other, then $\overline{x}_1$ is connected to V' but not to V; and again $\overline{x}_1$ is a cut vertex (because it satisfies condition (i) of Definition 3.4). If $f(G(w))$ does not cross the arc labeled x_1 with distinct endpoints, then we make use of the assumption that w contains every letter at least once; we get that $f(G(w))$ has to contain the edge $x_{\alpha+i}$ for some $1 \le i \le \beta$; this gives that any of $x_{\alpha+i}$, $\overline{x}_{\alpha+i}$ is a cut vertex for the Whitehead graph of w (because they satisfy condition (i) of Definition 3.4). ■

We now prove that Theorem 3.5 implies Theorem 3.6, together with the fine property of Theorem 3.7.

Proof of Theorems 3.6 *and* 3.7. Let w be cyclically reduced, and let a be a cut vertex in its Whitehead graph.

If the connected component of a does not contain $\overline{a}$, then we take the set A to be that connected component (excluding a itself). Otherwise, take the connected component of a and remove a itself: we are left with at least two nonempty connected components, and at least one of these components does not contain $\overline{a}$; take A to be such a component. In both cases we consider the Whitehead automorphism $\varphi = (A, a)$. We look at what happens between two consecutive non-a non-$\overline{a}$ letters in w when we apply the automorphism φ.

Suppose we have two consecutive letters $w = \dots yz \dots$ with $y, z \in \{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\} \setminus \{a, \overline{a}\}$. Then we have an arc from $\overline{y}$ to z in the Whitehead graph, so we are either acting on neither of $\overline{y}$, z or on both of them. If we are not acting on them, then $\varphi(w) = \dots yz \dots$ and the word is not affected between y and z. If we are acting on both of them, then $\varphi(w) = \dots y\overline{a}az \dots = \dots yz \dots$ and every a and $\overline{a}$ which appears there immediately cancels, and again the word is not affected between y and z.

Suppose we have in w a segment of the form $w = \dots y a^k z \dots$ with $y, z \in \{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\} \setminus \{a, \overline{a}\}$ and $k \geq 1$ (the case $k \leq -1$ is analogous). This means that we have an arc from $\overline{a}$ to z, and thus we are not acting on z. If we are not acting on $\overline{y}$, then $\varphi(w) = \dots y a^k z \dots$ and the word is not affected between y and z. If we are acting on $\overline{y}$, then $\varphi(w) = \dots y \overline{a} a^k z \dots = \dots y a^{k-1} z \dots$ and the number of a letters strictly decreases between y and z.

This shows that the property of Theorem 3.7 holds for this automorphism. To conclude, we notice that, since there is at least one arc between a and a vertex of A, at least one cancellation takes place, giving a (strict) decrease in the cyclic length of w, yielding Theorem 3.6. ∎

Remark 3.8. The above also shows that it is possible to count the number of cancellations that take place, by just looking at the Whitehead graph of w and at the Whitehead automorphism $\varphi = (A, a)$; it is also shown in [8, Proposition 2.2]. To be precise, let E be the number of edges of the Whitehead graph of w that connect the vertex a to a vertex of A; then we have $|\varphi(w)|_c = |w|_c - E$, where $|w|_c$ denotes the cyclic length of the word w.

4. On primitive elements in a subgroup of a free group

Let $H \leq F_n$ be any subgroup. If H is a free factor, then, of course, every element which is primitive in H has to be primitive in F_n too. We here deal with a converse: if H is not a free factor, then there is an element which is primitive in H, but not in F_n. This section is completely dedicated to the proof of the existence of such a witness.

Theorem 4.1. *Let $H \leq F_n$ be a finitely generated subgroup. Suppose that every element of H which is primitive in H is also primitive in F_n. Then H is a free factor.*

Proof. The proof proceeds by induction on the rank of the subgroup. For the base step, we notice that for a subgroup of rank 1 the statement is trivially true. For the inductive step, suppose we know the statement to be true for subgroups of rank k, and we want to prove it for subgroups of rank $k + 1$.

Take a subgroup $H = \langle w_1, \dots, w_{k+1} \rangle$ of rank $k + 1$ (meaning that $w_1, \dots, w_{k+1}$ is a basis for H), and suppose that every element v which is primitive in H is also primitive in F. We consider the subgroup $H' = \langle w_1, \dots, w_k \rangle$ and notice that every element v which is primitive in H' is also primitive in H, and thus is primitive in F. Then H' has rank k and satisfies the hypothesis of the theorem, so by inductive hypothesis we get that H' is a free factor. So we can take an automorphism $\theta \colon F \to F$ with $\theta(w_i) = x_i$ for $i = 1, \dots, k$. Instead of proving that H is a free factor in F, we prove that $\theta(H)$ is a free factor in F; but $\theta(H) = \langle x_1, \dots, x_k, \theta(w_{k+1}) \rangle$, so it is enough to prove the statement of Theorem 4.1 for subgroups H of the form $H = \langle w, x_1, \dots, x_k \rangle$.

The statement in the case of subgroups of the form $H = \langle w, x_1, \dots, x_k \rangle$ will be proved by induction on the length of w. The base step where w has length one is trivial.

We observe that if the first (or the last) letter of w is $y \in \{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$, then we can define the shorter word $w' = \overline{y}w$ (or $w\overline{y}$). But then $H = \langle w, x_1, \dots, x_k \rangle = \langle w', x_1, \dots, x_k \rangle$, and so we are done by inductive hypothesis. Thus in the following, we assume this is not the case.

We consider the word $v = x_1 w x_1 \widetilde{w}$, where $\widetilde{w}$ is any word in the letters $\{x_1, \dots, x_k\}$ with the following properties:

(i) The Whitehead graph of $\widetilde{w}$ contains at least one edge joining each pair of distinct vertices in $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$.

(ii) When we write $x_1 w x_1 \widetilde{w}$, we get a cyclically reduced word, without any cancellation needed.

For example, we may take

$$\widetilde{w} = (x_1 x_1)(x_2 x_2) \dots (x_k x_k) \prod_{1 \le i < j \le k} x_1 (x_i x_j)(x_i \overline{x}_j),$$

where the factors in the product are ordered lexicographically (but any ordering works).

The word v is primitive in the subgroup H, so by the hypothesis, it has to be primitive in F, and, in particular, we can take an automorphism φ satisfying Theorems 3.6 and 3.7. We now look at what happens to the letters $x_1, \dots, x_k$ and to the word w.

Case 1. Suppose the acting letter a is different from $x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k$. Then $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$ is either contained in A or disjoint from A. This is because the vertices $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$ are all pairwise connected in the Whitehead graph of v.

Subcase 1.1. Suppose that $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$ is disjoint from A. This means that $\varphi(x_1) = x_1, \dots, \varphi(x_k) = x_k$ and $\varphi\langle w, x_1, \dots, x_k \rangle = \langle \varphi(w), x_1, \dots, x_k \rangle$. We have $\varphi(v) = x_1 \varphi(w) x_1 \widetilde{w}$ which is cyclically reduced, and thus has to be strictly shorter than $v = x_1 w x_1 \widetilde{w}$. But in $x_1 \varphi(w) x_1 \widetilde{w}$ we can only have cancellations inside $\varphi(w)$, and so we are able to deduce that $\varphi(w)$ is strictly shorter than w, and we are done by inductive hypothesis.

Subcase 1.2. Let $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$ be contained in A. This means that $\varphi(x_1) = \overline{a} x_1 a, \dots, \varphi(x_k) = \overline{a} x_k a$. Moreover, we have $\varphi\langle w, x_1, \dots, x_k \rangle = \overline{a}\langle a\varphi(w)\overline{a}, x_1, \dots, x_k \rangle a$. We have $\varphi(v) = \overline{a} x_1 (a\varphi(w)\overline{a}) x_1 \widetilde{w} a$ which cyclically reduces to $x_1 (a\varphi(w)\overline{a}) x_1 \widetilde{w}$, and thus this has to be strictly shorter than $v = x_1 w x_1 \widetilde{w}$. But in $x_1 (a\varphi(w)\overline{a}) x_1 \widetilde{w}$, we can only have cancellations inside $(a\varphi(w)\overline{a})$, and so we deduce that $a\varphi(w)\overline{a}$ is strictly shorter than w, and we are done by the inductive hypothesis.

Case 2. Suppose the acting letter a is one of $x_2, \dots, x_k, \overline{x}_2, \dots, \overline{x}_k$. Then $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$ is disjoint from A, because all the vertices $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\} \setminus \{a, \overline{a}\}$ are connected to $\overline{a}$ in the Whitehead graph of v. Now we have $\varphi(x_1) = x_1, \dots, \varphi(x_k) = x_k$ and $\varphi\langle w, x_1, \dots, x_k \rangle = \langle \varphi(w), x_1, \dots, x_k \rangle$. We proceed exactly as in subcase 1.1. The key point is that, since $a \neq x_1, \overline{x}_1$, when we write $x_1 \varphi(w) x_1 \widetilde{w}$, the two x_1 letters cannot cancel against $\varphi(w)$. We get that $\varphi(w)$ is strictly shorter than w, and we are again done by the inductive hypothesis.

Case 3. Suppose $a = x_1$ (the case $a = \overline{x}_1$ is completely analogous). As in case 2, we get that $\{x_1, \dots, x_k, \overline{x}_1, \dots, \overline{x}_k\}$ is disjoint from A, and thus $\varphi(x_1) = x_1, \dots, \varphi(x_k) = x_k$ and $\varphi\langle w, x_1, \dots, x_k\rangle = \langle \varphi(w), x_1, \dots, x_k\rangle$. We have $\varphi(v) = (x_1\varphi(w))x_1\widetilde{w}$ which is cyclically reduced, and has thus to be strictly shorter than $v = x_1 w x_1 \widetilde{w}$. In $(x_1\varphi(w))x_1\widetilde{w}$, the only cancellations can happen inside $(x_1\varphi(w))$, so we are able to deduce that $x_1\varphi(w)$ has to be strictly shorter than $x_1 w$. Now some care is needed. If x_1 does not cancel against $\varphi(w)$, then we get that $\varphi(w)$ is strictly shorter than w, and we are done by the inductive hypothesis. Otherwise, let t be the first letter of w, and we must have $t \in A$. We look at the arcs between $a = x_1$ and A in the Whitehead graph of v: we certainly have at least one arc from x_1 to t.

Subcase 3.1. If we have more than one arc from x_1 to t, or if we have any other arc from x_1 to A, then all of those arcs give cancellations inside $\varphi(w)$. We already knew that $x_1\varphi(w)$ was strictly shorter than $x_1 w$, but now we got at least one additional cancellation inside $\varphi(w)$, so we are able to deduce that $x_1\varphi(w)$ is strictly shorter than w. Now we observe that $\varphi\langle w, x_1, \dots, x_k\rangle = \langle x_1\varphi(w), x_1, \dots, x_k\rangle$, and we are done by inductive hypothesis.

Subcase 3.2. Suppose we only have one arc from x_1 to t, and no other arc from x_1 to A. If the vertex t has degree at least 2, then we use t instead of x_1 as cut vertex for the Whitehead graph of v, and we end up in case 1, and we are done. If the vertex t has degree 1, then the letter t appears exactly once inside w; in this case, we have the automorphism $\eta\colon F \to F$ which keeps all the letters fixed, except for $t \mapsto w$; this gives $\eta\langle t, x_1, \dots, x_k\rangle = \langle w, x_1, \dots, x_k\rangle$, showing that the subgroup is a free factor, as desired. ∎

Remark 4.2. We proved that Theorem 4.1 holds for a subgroup H which is finitely generated, but that hypothesis can easily be waived. For a subgroup of rank greater than the rank of F (and, as a consequence, for a subgroup of infinite rank), the hypothesis of Theorem 4.1 cannot hold.

5. Whitehead's algorithm for free factors

Whitehead graph for subgroups

Definition 5.1. Let G be a F_n-labeled graph, and let $v \in G$ be a vertex. Define the *letters at* v to be the subset $L(v) \subseteq \{x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n\}$ of the labels of the edges coming out of v. More precisely, we have $x_i \in L(v)$ if and only if G contains an edge labeled x_i coming out of v, and $\overline{x}_i \in L(v)$ if and only if G contains an edge labeled x_i going into v.

Definition 5.2. Let G be a F_n-labeled graph. Define the *Whitehead graph* of G as follows:

(i) We have $2n$ vertices labeled $x_1, \dots, x_n, \overline{x}_1, \dots, \overline{x}_n$.

(ii) For every vertex $v \in G$ and for every pair $y, z \in L(v)$ of distinct letters at v, we draw an (unoriented) arc from y to z in the Whitehead graph.

This means that the Whitehead graph of G contains an edge for every (legal) turn in G. Notice that the Whitehead graph contains a complete subgraph with vertex set $L(v)$ for every vertex $v \in G$; moreover, the Whitehead graph is exactly the union of these complete subgraphs. Notice that, when we apply a folding operation to a F_n-labeled graph G, the Whitehead graph of G can gain new edges, but it does not lose any.

We define the Whitehead graph of a nontrivial finitely generated subgroup $H \leq F_n$ to be the Whitehead graph of $\mathrm{core}(H)$. When the subgroup H is generated by a single word $H = \langle w \rangle$, the Whitehead graph of H coincides with the Whitehead graph of the cyclic reduction of w; in this sense, this notion of Whitehead graph is a generalization of the previous Definition 3.3. We can also define the notion of cut vertex for the Whitehead graph of a subgroup exactly as in Definition 3.4.

Whitehead automorphisms and subdivision of graphs

We are now interested in how the core graph of a subgroup changes when we apply a Whitehead automorphism. We are thus going to describe an operation which we call *subdivision*, being performed on an F_n-labeled graph. In what follows, we work with a fixed Whitehead automorphism $\varphi = (A, a)$, and we assume that $a \in \{x_1, \ldots, x_n\}$. The case where $a \in \{\overline{x}_1, \ldots, \overline{x}_n\}$ is completely analogous: whenever we would have an edge labeled a with a certain orientation, we consider instead an edge labeled $\overline{a}$ and with opposite orientation.

Let G be an F_n-labeled graph. Choose an edge $e \in G$, oriented and labeled with a letter $y \in \{x_1, \ldots, x_n\}$. If $y, \overline{y} \notin A$, we do not change the edge e. If $y \in A$ and $\overline{y} \notin A$, then we subdivide e into two edges, and to the first, we give the label a and the orientation of e, and to the second, we give the label y and the same orientation. If $y \notin A$ and $\overline{y} \in A$, then we subdivide e into two edges, and to the first, we give the label y and the same orientation of e, and to the second, we give the label a and the opposite orientation. If $y, \overline{y} \in A$, then we perform both transformations on e, as in Figure 3.

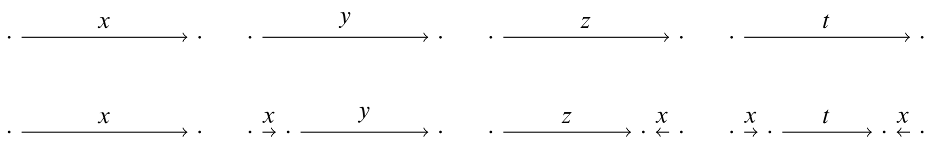

Figure 3. The effect of the subdivision operation upon the single edges. Here $F_4 = \langle x, y, z, t \rangle$ and $\varphi = (\{y, \overline{z}, t, \overline{t}\}, x)$, meaning that $\varphi(x) = x$ and $\varphi(y) = xy$ and $\varphi(z) = z\overline{x}$ and $\varphi(t) = xt\overline{x}$. Above, we see the edges before the subdivision, while below we see them after the subdivision.

We apply the subdivision operation to each edge of G, in order to obtain another F_n-labeled graph, which will be called *φ-subdivided graph*, to be denoted by $\mathrm{subd}_\varphi(G)$.

Notice that for every vertex $v \in G$, we also have a vertex $v \in \mathrm{subd}_\varphi(G)$: we say that $v \in \mathrm{subd}_\varphi(G)$ is an *old vertex*. For every edge $e \in G$, the subdivision on e gives at most three edges, and exactly one of them has the same label as e: we say that the edge is an *old* *edge*. If we have a vertex $v \in G$ and a letter $b \in L(v) \cap A$, then the subdivision operation upon the corresponding edge will create a new vertex u together with an edge labeled a going from v to u: we say that the vertex u is a *new vertex near* v, and that the edge from v to u is a *new edge near* v. For every vertex of $\mathrm{subd}_\varphi(G)$, it is either an old vertex or a new vertex near a unique $v \in G$. For every edge of $\mathrm{subd}_\varphi(G)$, it is either an old edge, or a new edge near a unique $v \in G$.

To an F_n-labeled graph G, we associate the subgroup $H \leq F_n$ given by the image $H = \pi_1(f)(\pi_1(G))$, where $f: G \to R_n$ is the labeling map. It is immediate to see that, if G is associated with the subgroup H, then $\mathrm{subd}_\varphi(G)$ is associated with the subgroup $\varphi(H)$. This observation yields the following (which is essentially the same as [3, Lemma 2]).

Proposition 5.3. *Let $H \leq F_n$ be a non-trivial finitely generated subgroup, and let $\varphi = (A, a)$ be a Whitehead automorphism. Then* $\mathrm{core}(\varphi(H)) = \mathrm{core}(\mathrm{fold}(\mathrm{subd}_\varphi(\mathrm{core}(H))))$.

Proof. For the F_n-labeled graph $\mathrm{core}(H)$ with labeling map $f: \mathrm{core}(H) \to R_n$, we have that $f_*(\pi_1(\mathrm{core}(H))) = H$. After the subdivision operation, if we call

$$g\colon \mathrm{subd}_\varphi(\mathrm{core}(H)) \to R_n$$

the labeling map, we have that $g_*(\pi_1(\mathrm{subd}_\varphi(\mathrm{core}(H)))) = \varphi(H)$. In particular, Proposition 2.7 tells us that $\mathrm{fold}(\mathrm{subd}_\varphi(\mathrm{core}(H)))$ can be embedded in $\mathrm{cov}(\varphi(H))$ as a subgraph containing $\mathrm{core}(\varphi(H))$. It follows that $\mathrm{core}(\mathrm{fold}(\mathrm{subd}_\varphi(\mathrm{core}(H)))) = \mathrm{core}(\varphi(H))$, as desired. ∎

We are also interested in having a detailed description of how the folding operations take place. Some partial description is already given in [3, Lemma 3], but we provide the more precise technical Lemma 5.4, which is a generalization of Lemma 3.2.

Lemma 5.4. *Let $H \leq F_n$ be a non-trivial finitely generated subgroup, and let $\varphi = (A, a)$ be a Whitehead automorphism. Consider the graph* $\mathrm{subd}_\varphi(\mathrm{core}(H))$. *For every vertex $v \in \mathrm{core}(H)$, fold together all the edges of* $\mathrm{subd}_\varphi(\mathrm{core}(H))$ *going out of v and labeled with a. Then, after these folding operations, no further folding operation is possible.*

In particular, for every folding sequence starting from $\mathrm{subd}_\varphi(\mathrm{core}(H))$, *the sequence contains only rank-preserving folding operations involving edges labeled a.*

Proof. Let $G = \mathrm{core}(H)$ and denote by $L(v)$ the set of letters at the vertex v in G. For every vertex $v \in G$, take all the edges in $\mathrm{subd}_\varphi(G)$ labeled a and going out of v, and fold them all together; call the resulting graph G'. Our aim is to show that no folding operation is possible on G'. This is equivalent to showing that, for every vertex $u \in G'$ and for every letter, there is at most one edge with that label going out of u.

Suppose we have a vertex $v \in G$ with $L(v) \cap A = \emptyset$. Then no new vertex is created near v in the subdivision operation.

Suppose we have a vertex $v \in G$ with $L(v) \cap A \neq \emptyset$ and $a \in L(v)$; this means that G contains an edge labeled a going from v to u. Then every new vertex, which is created near v with the subdivision operation, gets identified with u in G'.

Suppose we have a vertex $v \in G$ with $L(v) \cap A \neq \emptyset$ and $a \notin L(v)$. Then all the new vertices, which are created near v with the subdivision operation, are identified together into a vertex which we call $v_1 \in G'$.

Thus G' contains only two types of vertices: old vertices u obtained from a vertex $u \in G$, and new vertices u_1 near a vertex $u \in G$ with $L(u) \cap A \neq \emptyset$ and $a \notin L(u)$.

Case 1. Suppose we have a vertex $u \in G$ with $\overline{a} \notin L(u)$. Since $\overline{a} \notin L(u)$, we have that the vertex $u \in \mathrm{subd}_\varphi(G)$ does not get identified with any other vertex during the folding operations that produce G'. The edges going out of $u \in G$ with label in $L(u) \cap A^c$ give edges going out of $u \in G'$ with the same label. The edges going out of $u \in G$ with label in $L(u) \cap A$ give edges labeled a going out of $u \in \mathrm{subd}_\varphi(G)$; all the edges labeled a and going out of $u \in \mathrm{subd}_\varphi(G)$ are folded together in G', meaning that there is at most one edge labeled a going out of $u \in G'$. Thus in this case, the vertex $u \in G'$ has at most one edge with each label going out of it.

Case 2. Suppose we have a vertex $u \in G$ with $\overline{a} \in L(u)$. This means that G contains an edge labeled a going from v to u (see also Figure 4). It is possible that the subdivision operation creates new vertices near v; the folding operation identifies all such vertices with u. Thus, for every letter $l \in L(v) \cap A$, the vertex $u \in G'$ has one edge labeled l going out of it. The edges going out of u with label in $L(u) \cap A^c$ give edges going out of $u \in G'$ with the same label. Notice that $L(v) \cap A$ and $L(u) \cap A^c$ are disjoint, so we cannot get two vertices with the same label in this way. As in case 1, the edges going out of $u \in G$ with label in $L(u) \cap A$ give edges labeled a going out of $u \in \mathrm{subd}_\varphi(G)$; all the edges labeled a going out of $u \in \mathrm{subd}_\varphi(G)$ are folded together in G', meaning that there is at most one edge labeled a going out of $u \in G'$. Thus in this case, the vertex $u \in G'$ has at most one edge with each label going out of it.

Case 3. Suppose we have a vertex $u \in G$ with $L(u) \cap A \neq \emptyset$ and $a \notin L(u)$. This means that all the vertices which are created near u fold together into a single vertex $u_1 \in G'$. Each edge going out of $u \in G$ with label in $L(u) \cap A$ gives one edge going out of $u_1 \in G'$ with the same label. The vertex u_1 also has one edge labeled $\overline{a}$ going out of it, and notice that $\overline{a} \notin L(u) \cap A$. It follows that the vertex $u_1 \in G'$ has at most one edge with each label going out of it.

Since we examined each vertex of G', we conclude that no folding operation is possible on G'. According to Proposition 2.7, we have that each folding sequence starting from $\mathrm{subd}_\varphi(G)$ can only contain rank-preserving folding operations involving edges labeled a, as desired. ∎

Whitehead's algorithm for subgroups

We are now ready to state the analogs of Theorems 3.5, 3.6 and 3.7 for free factors.

Theorem 5.5. *Let $H \leq F_n$ be a free factor, and suppose* $\mathrm{core}(H)$ *has more than one vertex. Then the Whitehead graph of H contains a cut vertex.*

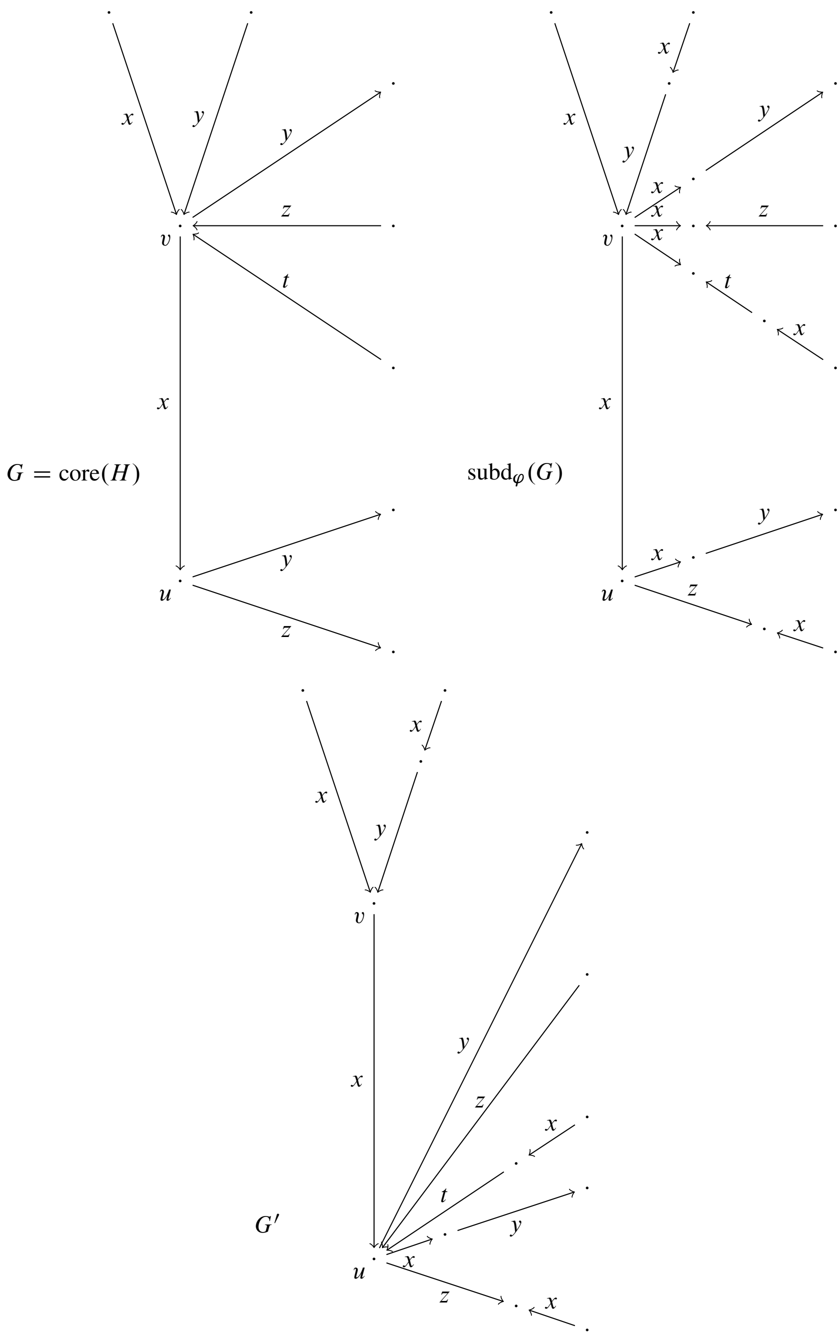

Figure 4. A local picture of the graph G and how it changes during the proof of Lemma 3.2 (with a focus on case 2). Here $F_4 = \langle x, y, z, t \rangle$ and $\varphi = (\{y, \bar{z}, t, \bar{t}\}, x)$ is the same as in Figure 3. We see a portion of the starting graph G, its subdivision, and the corresponding portion of G'.

Theorem 5.6. *Let $H \leq F_n$ be a free factor, and suppose the Whitehead graph of H contains a cut vertex. Then there is a Whitehead automorphism φ such that* $\operatorname{core}(\varphi(H))$ *has strictly fewer vertices and strictly fewer edges than* $\operatorname{core}(H)$.

In the following theorem, we write $L(v)$ as introduced in Definition 5.1.

Theorem 5.7. *The automorphism $\varphi = (A, a)$ in Theorem 5.6 can be chosen in such a way that, at each vertex v of* $\operatorname{core}(H)$, *exactly one of the following configurations takes place:*

(i) $L(v) \cap A = \emptyset$.

(ii) $L(v) \subseteq A$.

(iii) $a \in L(v)$ *and* $L(v) \subseteq A \cup \{a\}$.

Remark 5.8. Case (i) means that we do not act on any of the letters at v. Case (ii) means that we act on all the letters at v. Case (iii) means that we act on all the letters at v except for a.

Remark 5.9. Notice that if $\overline{a} \in L(v)$, then v necessarily falls into case (i).

The proof of Theorem 5.5 is analogous to the proof of Theorem 3.5.

Proof of Theorem 5.5. Let H be a free factor such that $\operatorname{core}(H)$ has more than one vertex. Up to conjugation, we can assume that the basepoint belongs to $\operatorname{core}(H)$. We also assume that $\operatorname{core}(H)$ contains each letter $x_1, \dots, x_n$ at least once; otherwise, if $\operatorname{core}(H)$ only contains the letters $x_1, \dots, x_k$, then we can just apply the same argument in the free factor $\langle x_1, \dots, x_k \rangle \leq \langle x_1, \dots, x_n \rangle = F_n$ (using Corollary 2.17).

Since H is a free factor, we can take a basis for H and add reduced words $w_1, \dots, w_r$ in order to make it a basis for F_n. Take the graph $\operatorname{core}(H)$ and add r paths from the basepoint to itself, corresponding to the words $w_1, \dots, w_r$, in order to get a graph G. Then, apply a sequence of folding operations $G \to G' \to \cdots \to G^{(l)}$ until no further folding operation is possible, as in Proposition 2.7. Since $\langle H, w_1, \dots, w_r \rangle = F_n$, we must have that $G^{(l)} = R_n$ is the standard n-rose. Using Corollary 2.10, we can see that no graph in the sequence contains any valence-1 vertex. Also, since $\pi_1(G)$ has the same rank as $\pi_1(R_n)$, we must have that each folding operation is rank-preserving.

Thus $G^{(l-1)}$ has no valence-1 vertex, and produces the standard n-rose with just one rank-preserving folding operation. It is easy to see that $G^{(l-1)}$ has to be of the form described in Figure 2, for some $1 \leq \alpha \leq \beta \leq n$ with $\alpha < n$ (up to permutation of the letters, and up to substitution of some letter with its inverse) (and the two edges labeled x_1 are the ones to be folded in order to obtain the n-rose).

We have a map of graphs $f\colon \operatorname{core}(H) \to G^{(l-1)}$ which preserves orientations and labels of edges. The image of $f(\operatorname{core}(H)) \subseteq G^{(l-1)}$ contains each letter at least once, meaning that it has to cross at least one of the edges connecting v to v' (see Figure 2). If it crosses the edge labeled x_1, then $\overline{x}_1$ is a cut vertex for the Whitehead graph of H. If it does not cross the edge labeled x_1, then it has to cross the edge $x_{\alpha+i}$ (for some $1 \leq i \leq \beta$), and thus any of $x_{\alpha+i}, \overline{x}_{\alpha+i}$ is a cut vertex for the Whitehead graph of H. ∎

Theorems 5.6 and 5.7 are a consequence of Theorem 5.5.

Proof of Theorems 5.6 and 5.7. Let a be a cut vertex in the Whitehead graph of H.

If the connected component of a does not contain $\overline{a}$, then we take the set A to be that connected component (excluding a itself). Otherwise, take the connected component of a and remove a itself: we remain with at least two nonempty connected components, and at least one of these components does not contain $\overline{a}$; take A to be such a component. We consider the Whitehead automorphism $\varphi = (A, a)$.

Take a vertex v in $\mathrm{core}(H)$, and notice that the letters in $L(v)$ are vertices of a complete subgraph of the Whitehead graph of H. Thus $L(v)$ has to be contained either in $A \cup \{a\}$ or in A^c. This yields the trichotomy of Theorem 5.7.

We now examine in more detail what happens in each of the three cases. The folding takes place according to Lemma 5.4. For each vertex v of $\mathrm{core}(H)$, we look at the vertices which are created near v in $\mathrm{subd}_\varphi(\mathrm{core}(H))$.

Case (i): $L(v) \subseteq A^c$. This means no new vertex is created near v. The total number of vertices remains unchanged.

Case (ii): This means that, for every edge with endpoint v, a new vertex is created near v. All these new vertices are then folded together into a vertex v_1. The vertex v becomes a valence-1 vertex, and can thus be removed from the graph. Thus we lose the vertex v and we gain the vertex v_1 in the core graph: the total number of vertices is unchanged.

Case (iii): $a \in L(v)$ and $L(v) \subseteq A \cup \{a\}$. This means that $\mathrm{core}(H)$ contains an edge e labeled a going from v to u. For every other edge with endpoint v, a new vertex is created near v. All these new vertices are then folded together with the vertex u. The vertex v becomes a valence-1 vertex, and can thus be removed from the graph. The total number of vertices decreases by 1.

In each of the cases (i), (ii) and (iii), the number of vertices and edges of the core graph does not increase. Also, since the Whitehead graph contains at least an edge between a and A, we have that case (iii) happens at least once, giving a strict decrease in the number of vertices and edges. This yields Theorem 5.6. ∎

Remark 5.10. We notice that, if $\mathrm{core}(H)$ has rank r, the number of edges of $\mathrm{core}(H)$ is the number of vertices plus $r - 1$. The same holds for $\mathrm{core}(\varphi(H))$, which has rank r too. Thus the decrease in the number of vertices is the same as the decrease in the number of edges.

Remark 5.11. One may try to look for generalization of Remark 3.8: we would like to compute the decrease in the number of edges of $\mathrm{core}(H)$ by just looking at the Whitehead graph of H; this is unfortunately not easy, because the valence of certain vertices of H comes to play a role. Let S be the set of vertices of $\mathrm{core}(H)$ that fall into case (iii) of the trichotomy of Theorem 5.7: then we have that $|\mathrm{core}(\varphi(H))|_e = |\mathrm{core}(H)|_e - |S|$, as we will show later in the paper (see Lemma 5.16, and refer also to Figure 5), where $|\mathrm{core}(H)|_e$

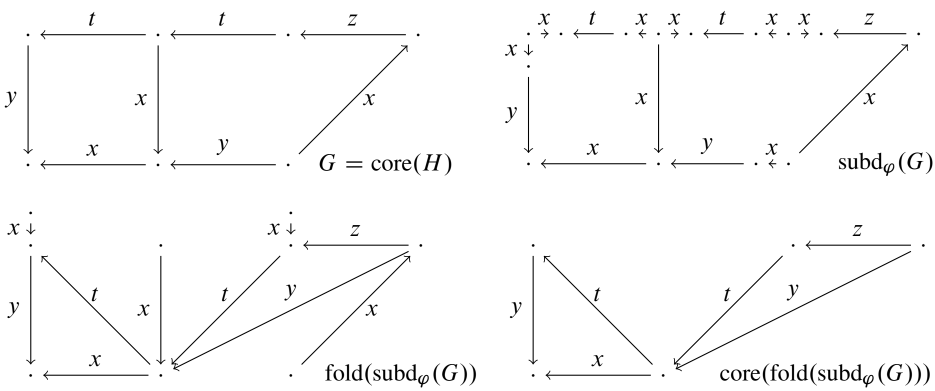

Figure 5. Here $F_4 = \langle x, y, z, t\rangle$ and we consider the free factor $H = \langle ty\bar{x}^2, x\bar{y}xzt\rangle$. The Whitehead transformation $\varphi = (\{y, \bar{z}, t, \bar{t}\}, x)$ satisfies the trichotomy of Theorem 5.7 for the graph core(H). In the figure we start with core(H), we subdivide it, we fold the result and we remove the valence-1 vertices: the result is core$(\varphi(H))$. We observe that core$(\varphi(H))$ can be obtained from core(H) in the following way: take the vertices of core(H) which fall into case (iii) of the trichotomy, and collapse to a point the x-edges at those vertices; Theorem 5.14 makes this formal.

denotes the number of edges of the graph core(H). Let E be the number of edges of the Whitehead graph of H that connect the vertex a to a vertex of A: the number E can in general be different from $|S|$; in fact, E and S are related by $E = \sum_{v\in S}(d(v) - 1)$, where $d(v)$ denotes the degree of a vertex v of core(H).

The quotient map

In the following, we use the notation $L(v)$ as introduced in Definition 5.1.

Definition 5.12. Let $H \le F_n$ be a finitely generated non-trivial subgroup, and let $\varphi = (A, a)$ be a Whitehead automorphism. We say that the action of φ on H is *fine* if for each vertex $v \in$ core(H), exactly one of the following configurations takes place:

(i) $L(v) \cap A = \emptyset$.

(ii) $L(v) \subseteq A$.

(iii) $a \in L(v)$ and $L(v) \subseteq A \cup \{a\}$.

Remark 5.13. This is exactly the property given by the trichotomy of Theorem 5.7.

Suppose now the action of φ on H is fine. Let v be a vertex of core(H) that falls into case (iii) of the trichotomy: since $a \in L(v)$, there is a unique edge labeled a going out of v. For each vertex v of core(H) that falls into case (iii), collapse that a-edge to a single point. We obtain a quotient F_n-labeled graph Q together with a quotient map $q\colon$ core$(H) \to Q$.

Theorem 5.14. *There is an isomorphism $\theta\colon Q \to \mathrm{core}(\varphi(H))$ of graphs sending each edge to an edge with the same label and orientation.*

Before proving the above theorem, we need to introduce another map first. According to Proposition 5.3, we have $\mathrm{core}(\varphi(H)) = \mathrm{core}(\mathrm{fold}(\mathrm{subd}_\varphi(\mathrm{core}(H))))$. Consider the map $r_1\colon \mathrm{core}(H) \to \mathrm{subd}_\varphi(\mathrm{core}(H))$ which sends each edge e of $\mathrm{core}(H)$ to the edge-path $\mathrm{subd}_\varphi(e)$. Consider also the map $r_2\colon \mathrm{subd}_\varphi(\mathrm{core}(H)) \to \mathrm{fold}(\mathrm{subd}_\varphi(\mathrm{core}(H)))$ which is given by the quotient map induced by the folding operations. Consider finally the map $r_3\colon \mathrm{fold}(\mathrm{subd}_\varphi(\mathrm{core}(H))) \to \mathrm{core}(\mathrm{fold}(\mathrm{subd}_\varphi(\mathrm{core}(H))))$ given by the retraction which collapses each edge with a valence-1 endpoint to the other endpoint. The composition of these three maps gives a map $r = r_3 \circ r_2 \circ r_1\colon \mathrm{core}(H) \to \mathrm{core}(\varphi(H))$.

Theorem 5.15. *The isomorphism $\theta\colon Q \to \mathrm{core}(\varphi(H))$ in Theorem 5.14 can be chosen in such a way that*

(i) *For each vertex v of $\mathrm{core}(H)$, we have $r(v) = (\theta \circ q)(v)$.*

(ii) *For each edge e of $\mathrm{core}(H)$, the map $r\big|_e$ is a (weakly monotone) reparametrization of $\theta \circ q\big|_e$.*

In particular, r and $\theta \circ q$ are homotopic relative to the 0-skeleton of $\mathrm{core}(H)$.

See also Figure 6.

$$\begin{tikzcd}
& \mathrm{core}(H) \arrow[dl, "q"'] \arrow[dr, "r"] & \\
Q \arrow[rr, "\theta"] & & \mathrm{core}(\varphi(H))
\end{tikzcd}$$

Figure 6. The maps q, r and θ.

Proof of Theorems 5.14 and 5.15. We examine cases (i), (ii), (iii) of the trichotomy of Definition 5.12.

Let v in $\mathrm{core}(H)$ be a vertex which falls into case (i). Then we have that no new vertex is created near v, and the core graph remains unchanged.

Let v in $\mathrm{core}(H)$ be a vertex which falls into case (ii). Then we have that, for each edge at v, a new vertex is created near v. All these new vertices fold together into a new vertex v_1, and v becomes a valence-1 vertex and is thus removed from the core graph. The vertex v_1 takes the place of the vertex v, and the core graph does not change.

Let v in $\mathrm{core}(H)$ be a vertex which falls into case (iii). We consider the unique edge e labeled a and going from v to another vertex u. For every other edge at v, we have that a new vertex is created near v. All these new vertices are then folded together and with u, and the vertex v becomes a valence-1 vertex, and is thus removed from the core graph. The effect on the core graph is exactly the same as collapsing the edge e to a single point.

This shows that the quotient graph Q is isomorphic to $\mathrm{core}(\varphi(H))$, yielding Theorem 5.14.

Take now an edge e of $\mathrm{core}(H)$ with endpoints u, v. Notice that if e gets collapsed by the quotient map q, then it is collapsed by the map r too, and the thesis holds; so assume this is not the case. If both u and v fall into case (i) of the trichotomy, then the quotient map q and the map r send e homeomorphically onto the same edge $q(e) = r(e)$ of $\mathrm{core}(\varphi(H))$. If u falls into case (i) but v falls into case (ii) or (iii), then q sends e homeomorphically onto an edge $q(e)$ of $\mathrm{core}(\varphi(H))$. The map r_1 maps e to an edge path containing two edges, and the map r_3 collapses one of those two edges to a point, and sends the other homeomorphically onto $q(e)$. This yields the conclusion for the edge e. The case where u falls into case (ii) or (iii) is completely analogous. ■

For brevity, in the following we write $\overline{q} = \theta \circ q\colon \mathrm{core}(H) \to \mathrm{core}(\varphi(H))$. The above Theorems 5.14 and 5.15 have several interesting consequences.

Lemma 5.16. *Let $H \leq F_n$ be a non-trivial finitely generated subgroup, and let $\varphi = (A, a)$ be a Whitehead automorphism such that the action of φ on H is fine. If case* (iii) *takes place for exactly $p \geq 1$ vertices $v \in \mathrm{core}(H)$, then $\mathrm{core}(\varphi(H))$ has exactly p fewer vertices and p fewer edges than $\mathrm{core}(H)$. If case* (iii) *never happens, then $\mathrm{core}(\varphi(H)) = \mathrm{core}(H)$ and the restriction of φ to H is conjugation by some element $u \in F_n$.*

Proof. The map $\overline{q}$ collapses exactly one edge for each vertex of $\mathrm{core}(H)$ that falls into case (iii). This proves the first part of the lemma.

For the second part, suppose that case (iii) never happens for a vertex of $\mathrm{core}(H)$. This means that the map $\overline{q}\colon \mathrm{core}(H) \to \mathrm{core}(\varphi(H))$ is an isomorphism of F_n-labeled graphs.

Consider the pointed core graph $\mathrm{core}_*(H)$, and let σ be the (possibly trivial) shortest path from the basepoint to a vertex v of $\mathrm{core}(H)$; moreover, call t the word that you read while going along σ, from the basepoint to v.

Take an element $w \in H$ and think of the corresponding (reduced) path α in $\mathrm{core}_*(H)$ from the basepoint to itself. The path α consists of σ followed by β followed by the reverse of σ, for some (reduced) path β in $\mathrm{core}(H)$ from v to itself. This gives a decomposition $w = tw'\overline{t}$, where w' is the word that we read while going along β.

Notice that $\overline{q}$ sends β isomorphically onto $\overline{q}(\beta)$, preserving labels and orientation on the edges. If v falls into case (i) of the trichotomy, then this gives $\varphi(w') = w'$, meaning that φ acts on H as the conjugation by $\varphi(t)\overline{t}$. If v falls into case (ii) or (iii) of the trichotomy, then this gives $\varphi(w') = aw'\overline{a}$, meaning that φ acts on H as the conjugation by $\varphi(t)a\overline{t}$. ■

We now show that the trichotomy of Definition 5.12 has a nice behavior when we consider subgroups.

Lemma 5.17. *Let $K \le H \le F_n$ be non-trivial finitely generated subgroups, and let $\varphi = (A, a)$ be a Whitehead automorphism. If the action of φ on H is fine, then the action of φ on K is fine.*

Proof. It is enough to notice that for every vertex $u \in \operatorname{core}(K)$, we have that its image $i_*(u) = v \in \operatorname{core}(H)$ satisfies $L(u) \subseteq L(v)$. Since the vertex v satisfies the trichotomy of Definition 5.12, so does u. ∎

Remark 5.18. In the hypothesis of Lemma 5.17, we have that, for each subgroup $K \le H$, the automorphism φ either strictly decreases the size of $\operatorname{core}(K)$, or it acts on K as a conjugation by an element of F_n. We can actually be more precise than just that. Consider the map $\overline{q}\colon \operatorname{core}(H) \to \operatorname{core}(\varphi(H))$, and let $G \subseteq \operatorname{core}(H)$ be the subgraph given by the union of all the edges which are not collapsed by q. Observe that the inclusion $i\colon K \to H$ induces a locally injective label-preserving map of graphs $i_*\colon \operatorname{core}(K) \to \operatorname{core}(H)$. Then, φ acts on K as a conjugation automorphism if and only if $i_*(\operatorname{core}(K)) \subseteq G$.

We conclude this section with a technical lemma which will be useful to us later. Let again $K \le H \le F_n$ be finitely generated non-trivial subgroups. Let $i\colon K \to H$ be the inclusion, consider the map of graphs $i_*\colon \operatorname{core}(H) \to \operatorname{core}(K)$ and consider the subgraph $i_*(\operatorname{core}(K)) \subseteq \operatorname{core}(H)$. For an automorphism $\varphi\colon F_n \to F_n$, let $j\colon \varphi(K) \to \varphi(H)$ be the inclusion, let

$$j_*\colon \operatorname{core}(\varphi(H)) \to \operatorname{core}(\varphi(K))$$

be the corresponding map of graphs, and consider the subgraph

$$j_*(\operatorname{core}(\varphi(K))) \subseteq \operatorname{core}(\varphi(H)).$$

Lemma 5.19. *Let $K \le H \le F_n$ be non-trivial finitely generated subgroups. Let $\varphi = (A, a)$ be a Whitehead automorphism such that the action of φ on K is fine. Then, with the above notation, $j_*(\operatorname{core}(\varphi(K)))$ has at most as many edges as $i_*(\operatorname{core}(K))$.*

Proof. Let $\overline{q} = \theta \circ q\colon \operatorname{core}(K) \to \operatorname{core}(\varphi(K))$ be as in Theorem 5.14. Suppose we have two edges e, e' in $\operatorname{core}(K)$ such that their image is the same edge $i_*(e) = i_*(e')$ of $\operatorname{core}(H)$, and suppose $\overline{q}$ does not collapse either of e, e'. Then the two edges $\overline{q}(e)$, $\overline{q}(e')$ of $\operatorname{core}(\varphi(K))$ are sent to the same edge

$$j_*(\overline{q}(e)) = j_*(\overline{q}(e'))$$

of $\operatorname{core}(\varphi(H))$. We now divide the edges of K into equivalence classes $E_1, \dots, E_\alpha$, where each equivalence class is the set of edges with a given image in $\operatorname{core}(H)$ (and, in particular, the image $i_*(\operatorname{core}(K))$ has exactly α edges); similarly, we divide the edges of $\operatorname{core}(\varphi(K))$ into equivalence classes $F_1, \dots, F_\beta$, based on their image in $\operatorname{core}(\varphi(H))$. Then each F_j is a union of some E_is, implying that $\beta \le \alpha$, as desired. ∎

Remark 5.20. Lemma 5.19 becomes false if we try to count the number of vertices, instead of counting the number of edges.

A relative version of Whitehead's algorithm

Let $F_n = \langle x_1, \ldots, x_n \rangle$ and consider the free factor $\langle x_1, \ldots, x_k \rangle$ for $1 \le k \le n-1$.

Theorem 5.21. *Let $w \in F_n$ be primitive and not a single letter. Suppose there is an automorphism $\theta \colon F_n \to F_n$ such that $\theta(\langle x_1, \ldots, x_k \rangle) = \langle x_1, \ldots, x_k \rangle$ and $\theta(w) = x_{k+1}$. Then there is a Whitehead automorphism $\varphi = (A, a)$ such that*

(i) $\varphi(x_i) = x_i$ *for* $i = 1, \ldots, k$.

(ii) *The length of $\varphi(w)$ is strictly smaller than the length of w.*

(iii) *Every letter a, which is added to w when applying φ to w letter-by-letter, immediately cancels (in the free reduction process).*

Remark 5.22. Notice that the word w is not required to be cyclically reduced. In (ii), we mean the length and not the cyclic length. In (iii), we consider the free reduction process and not the cyclic reduction process.

For the proof, we need the following straightforward lemma.

Lemma 5.23. *Consider the inner automorphism $\gamma_a(w) = aw\bar{a}$ of F_n. Then for every Whitehead automorphism (A, a), the identity $(A, a) = \gamma_a \circ (A^c \setminus \{a, \bar{a}\}, \bar{a})$ holds.*

Proof of Theorem 5.21. Consider the free factor $H = \langle x_1, \ldots, x_k \rangle * \langle w \rangle$. Notice that $\mathrm{core}(H)$ consists of $\mathrm{core}_*(w)$ together with k edges from the basepoint to itself, labeled with the letters $x_1, \ldots, x_k$ (and here it is important that $k \ge 1$). We apply Theorems 5.6 and 5.7 to H in order to get a Whitehead automorphism $\varphi = (A, a)$. If the basepoint of $\mathrm{core}(H)$ would fall into case (ii) or (iii) of the trichotomy of Theorem 5.7, then we apply Lemma 5.23 and consider the Whitehead automorphism

$$\varphi = (A^c \setminus \{a, \bar{a}\}, \bar{a})$$

instead. Then φ satisfies all of the desired properties. ∎

We observe that Theorem 5.21 can also be generalized to subgroups (and the proof is the same, so will be omitted).

Theorem 5.24. *Let $H \le F_n$ be a free factor of rank $r \ge 1$, and suppose that $\mathrm{core}_*(H)$ has at least two vertices. Suppose there is an automorphism $\theta \colon F_n \to F_n$ such that*

$$\theta(\langle x_1, \ldots, x_k \rangle) = \langle x_1, \ldots, x_k \rangle \quad \textit{and} \quad \theta(H) = \langle x_{k+1}, \ldots, x_{k+r} \rangle.$$

Then there is a Whitehead automorphism $\varphi = (A, a)$ such that

(i) $\varphi(x_i) = x_i$ *for* $i = 1, \ldots, k$.

(ii) *The graph $\mathrm{core}_*(\varphi(H))$ has strictly fewer vertices and edges than $\mathrm{core}_*(H)$.*

(iii) *The trichotomy of Theorem 5.7 holds at each vertex $v \in \mathrm{core}_*(H)$. Moreover, the basepoint always falls into case* (i) *of the trichotomy.*

6. About computation of distances in the complex of free factors

For an element $w \in F_n$, we denote by $[w]$ the conjugacy class of that element. For a subgroup $H \leq F_n$, we denote by $[H]$ the conjugacy class of that subgroup.

We are now going to define a simplicial complex FF_n, starting with its 0-skeleton and its 1-skeleton. The 0-skeleton FF_n^0 has a point $[H]$ for every conjugacy class of free factors $H \leq F_n$. The 1-skeleton FF_n^1 is defined as follows: add a 1-simplex with vertices $[H_0]$, $[H_1]$ if and only if $[H_0] \neq [H_1]$ and there are representatives $H_0' \in [H_0]$ and $H_1' \in [H_1]$ and a permutation $\sigma\colon \{0,1\} \to \{0,1\}$ such that $H'_{\sigma(0)} \leq H'_{\sigma(1)}$. Define FF_n as the flag complex over the 1-skeleton FF_n^1: we have a k-simplex with endpoints $[H_0], \dots, [H_k]$ if and only if $[H_0], \dots, [H_k]$ are pairwise connected by 1-simplices in FF_n^1. Equivalently, we have a k-simplex with endpoints $[H_0], \dots, [H_k]$ if and only if $[H_0], \dots, [H_k]$ are pairwise distinct and there are representatives $H_0' \in [H_0], \dots, H_k' \in [H_k]$ and a permutation $\sigma\colon \{0,\dots,k\} \to \{0,\dots,k\}$ such that $H'_{\sigma(0)} \leq \dots \leq H'_{\sigma(k)}$.

Definition 6.1. The simplicial complex FF_n defined above is called *complex of free factors*.

It is shown in [2] that FF_n is connected. We would like to determine whether there is an algorithm that, given two vertices of FF_n, gives as output their distance in a finite time, where distance is the combinatorial distance in the 1-skeleton FF_n^1. We here furnish algorithms for distances 1, 2, 3, and also an algorithm for distance 4 when one of the free factors has rank $n-1$.

Distance one

It is easy to check whether two conjugacy classes of free factors $[H]$, $[K]$ are at distance 1 or not. Assume $\operatorname{rank}(H) \geq \operatorname{rank}(K)$. We look for representatives $H' \in [H]$ and $K' \in [K]$ with an inclusion $K' \leq H'$. This is equivalent to looking for a locally injective map of graphs $\operatorname{core}(K) \to \operatorname{core}(H)$. Each such map, if it exists, is uniquely determined by the image of a given vertex; thus we only have to deal with a finite number of tries.

Distance two

We will rely on the following proposition.

Proposition 6.2. *Let H and K be non-trivial free factors, and suppose that* $\operatorname{core}(H) \sqcup \operatorname{core}(K)$ *contains at least one edge with each label. Suppose there are free factors $H' \in [H]$ and $K' \in [K]$ and $J \neq F_n$ such that $H', K' \leq J$. Then there is a Whitehead automorphism $\varphi = (A, a)$ such that*

$$\operatorname{core}(\varphi(H)) \sqcup \operatorname{core}(\varphi(K))$$

has strictly fewer vertices and strictly fewer edges than $\operatorname{core}(H) \sqcup \operatorname{core}(K)$.

Remark 6.3. With $\mathrm{core}(H) \sqcup \mathrm{core}(K)$ we mean the (not connected) F_n-labeled graph defined as the disjoint union of $\mathrm{core}(H)$ and $\mathrm{core}(K)$. Similarly, with $\mathrm{core}(\varphi(H)) \sqcup \mathrm{core}(\varphi(K))$ we mean the F_n-labeled graph defined as the disjoint union of $\mathrm{core}(\varphi(H))$ and $\mathrm{core}(\varphi(K))$.

Proof. Since J is a free factor, by a recursive application of Theorems 5.6 and 5.7, we obtain a chain of Whitehead automorphisms $\varphi_1, \ldots, \varphi_l$ such that $\mathrm{core}(\varphi_l \circ \cdots \circ \varphi_1(J))$ is a rose with labels only in $\{x_1, \ldots, x_{n-1}\}$. By Lemmas 5.17 and 5.16, we have that either $\mathrm{core}(\varphi_1(H)) \sqcup \mathrm{core}(\varphi_1(K)) = \mathrm{core}(H) \sqcup \mathrm{core}(K)$ or $\mathrm{core}(\varphi_1(H)) \sqcup \mathrm{core}(\varphi_1(K))$ has strictly fewer vertices and strictly fewer edges than $\mathrm{core}(H) \sqcup \mathrm{core}(K)$. If $\mathrm{core}(\varphi_1(H)) \sqcup \mathrm{core}(\varphi_1(K)) = \mathrm{core}(H) \sqcup \mathrm{core}(K)$, then we repeat the reasoning with φ_2 instead of φ_1; and so on. If $\mathrm{core}(\varphi_l \circ \cdots \circ \varphi_1(H)) \sqcup \mathrm{core}(\varphi_l \circ \cdots \circ \varphi_1(K)) = \mathrm{core}(H) \sqcup \mathrm{core}(K)$, then we have a contradiction, since $\mathrm{core}(\varphi_l \circ \cdots \circ \varphi_1(H)) \sqcup \mathrm{core}(\varphi_l \circ \cdots \circ \varphi_1(K))$ only contains edges with the labels $\{x_1, \ldots, x_{n-1}\}$, while $\mathrm{core}(H) \sqcup \mathrm{core}(K)$ contains edges with all possible labels by hypothesis. So we can take the smallest m such that

$$\mathrm{core}(\varphi_m(H)) \sqcup \mathrm{core}(\varphi_m(K)) \neq \mathrm{core}(H) \sqcup \mathrm{core}(K),$$

and the Whitehead automorphism φ_m satisfies the thesis. ■

Let $[H]$, $[K]$ be conjugacy classes of non-trivial free factors. We want to check (i) whether or not there are representatives with non-trivial intersection and (ii) whether or not there are representatives contained in a common proper free factor.

For (i), it is possible to use the technique explained in [10]. There are representatives $H' \in [H]$ and $K' \in [K]$ with non-trivial intersection if and only if the pullback of the two graphs $\mathrm{core}(H)$ and $\mathrm{core}(K)$ contains a non-trivial cycle.

For (ii), we apply Proposition 6.2 repeatedly. If $\mathrm{core}(H) \sqcup \mathrm{core}(K)$ contains only edges with labels from $\{x_1, \ldots, x_{n-1}\}$ (or from any other proper subset of $\{x_1, \ldots, x_n\}$), then they are at distance two; otherwise we look for a Whitehead transformation which strictly reduces the number of vertices of $\mathrm{core}(H) \sqcup \mathrm{core}(K)$: if we do not find it, then they are not at distance two, if we do then we apply it and reiterate the reasoning.

Distance three

Given two conjugacy classes of free factors $[H]$, $[K]$, we want to check whether there are representatives $H' \in [H]$ and $K' \in [K]$ and non-trivial free factors I, J such that $H', J \leq I$ and $J \leq K'$ (there is also a symmetric check to do, but it is completely analogous).

Define the finite oriented graph Θ as follows. The graph Θ has one vertex corresponding to each pair of core folded F_n-labeled graphs (A, B) such that A has at most as many edges as $\mathrm{core}(H)$ and B has at most as many edges as $\mathrm{core}(K)$. There is an oriented edge from (A, B) to (C, D) if and only if there is a Whitehead automorphism φ such that $C = \mathrm{core}(\mathrm{fold}(\mathrm{subd}_\varphi(A)))$ and D is isomorphic to a subgraph of $\mathrm{core}(\mathrm{fold}(\mathrm{subd}_\varphi(B)))$; in that case we label that edge of Θ with the automorphism φ.

Suppose now there are non-trivial free factors I, J such that $H', J \leq I$ and $J \leq K'$. The inclusion $j\colon J \to K'$ gives a locally injective map $j_*\colon \mathrm{core}(J) \to \mathrm{core}(K)$, and thus a subgraph $j_*(\mathrm{core}(J)) \subseteq \mathrm{core}(K)$: the pair $(\mathrm{core}(H), j_*(\mathrm{core}(J)))$ is a vertex of Θ. If $\mathrm{core}(I)$ contains only edges with labels from $\{x_1, \dots, x_{n-1}\}$, then we see that the pair of graphs $(\mathrm{core}(H), j_*(\mathrm{core}(J)))$ only contains edges with those labels too. If $\mathrm{core}(I)$ contains at least one edge with each label, then by Theorems 5.6 and 5.7 there is a Whitehead automorphism φ such that $\mathrm{core}(\varphi(I))$ has strictly fewer edges than $\mathrm{core}(I)$, and such that for each vertex in $\mathrm{core}(I)$ the trichotomy of Theorem 5.7 holds. In particular, by Lemmas 5.16, 5.17 and 5.19, we have that the number of edges of $\mathrm{core}(H)$ and of $j_*(\mathrm{core}(J))$ does not increase either. This means that the pair $(\mathrm{core}(\varphi(H)), j_*(\mathrm{core}(\varphi(J))))$ is a vertex of Θ, and that Θ contains an edge labeled φ and going from $(\mathrm{core}(H), j_*(\mathrm{core}(J)))$ to $(\mathrm{core}(\varphi(H)), j_*(\mathrm{core}(\varphi(J))))$ (here we are using Proposition 5.3).

We now reiterate the same reasoning. By Theorems 5.6 and 5.7, we can take a finite sequence of Whitehead automorphisms $\varphi_1, \dots, \varphi_l$ such that φ_i strictly reduces the number of edges of $\mathrm{core}(\varphi_{i-1} \circ \cdots \circ \varphi_1(I))$, such that for each vertex of $\mathrm{core}(\varphi_{i-1} \circ \cdots \circ \varphi_1(I))$ the trichotomy of Theorem 5.7 holds, and such that $\varphi_l \circ \cdots \circ \varphi_1(I)$ only contains edges with labels from $\{x_1, \dots, x_{n-1}\}$. Then this produces a path in Θ with vertices $(\mathrm{core}(\varphi_i \circ \cdots \circ \varphi_1(H)), j_*(\mathrm{core}(\varphi_i \circ \cdots \circ \varphi_1(J))))$ and which goes from the pair $(\mathrm{core}(H), j_*(\mathrm{core}(J)))$ to a pair containing only edges with labels in $\{x_1, \dots, x_{n-1}\}$. Since the graph Θ is finite, there is an algorithm that tells us whether such a path in Θ exists or not.

Conversely, given two conjugacy classes of free factors $[H]$, $[K]$, suppose there is a path in Θ with vertices $(A_1, B_1), \dots, (A_l, B_l)$ and with an edge labeled φ_i going from (A_i, B_i) to (A_{i+1}, B_{i+1}), such that $A_1 = \mathrm{core}(H)$ and B_1 is a subgraph of $\mathrm{core}(K)$, and such that A_l, B_l only contain edges with labels in $\{x_1, \dots, x_{n-1}\}$. Then we fix basepoints in A_l and B_l and we set $\psi = \varphi_1^{-1} \circ \cdots \circ \varphi_l^{-1}$: we get a segment of length three in FF_n^1 connecting $[H]$ and $[K]$, with vertices $[H] = [\psi(\pi_1(A_l))]$ and $[I] = [\psi(\langle x_1, \dots, x_{n-1}\rangle)]$ and $[J] = [\psi(\pi_1(B_l))]$ and $[K]$.

Thus, given conjugacy classes of free factors $[H]$ and $[K]$, the existence of non-trivial free factors I, J such that $H', J \leq I$ and $J \leq K'$ is equivalent to the existence of a path in Θ from a vertex of the form $(\mathrm{core}(H), B_1)$, with $B_1 \subseteq \mathrm{core}(K)$, to a vertex of the form (A_l, B_l), where $A_l \sqcup B_l$ does not use all the labels in $\{x_1, \dots, x_n\}$. This yields an algorithm to check whether two vertices of FF_n are at distance three or not.

About distance four

We would like to check whether two conjugacy classes of free factors $[H]$, $[K]$ are at distance at most four in FF_n. In order to achieve this, we need to check two conditions:

(1) Whether or not there are representatives $H' \in [H]$ and $K' \in [K]$ and non-trivial free factors J_1, J_2, J_3 such that $J_1 \leq H'$ and $J_1, J_3 \leq J_2$ and $J_3 \leq K'$.

(2) Whether or not there are representatives $H' \in [H]$ and $K' \in [K]$ and non-trivial free factors J_1, J_2, J_3 such that $H', J_2 \leq J_1$ and $J_2, K' \leq J_3$.

We here furnish an algorithm to check condition (1).

Remark 6.4. In the particular case when $\operatorname{rank}(H) = n - 1$, condition (2) reduces to checking distance three. In particular, when one of the free factors has rank $n - 1$, we have an algorithm to check whether they are at distance four or not.

The technique is the same as for distance three. Consider the oriented graph Ω defined as follows. We have one vertex for each pair of core folded F_n-labeled graphs (A, B) such that A has at most as many edges as $\operatorname{core}(H)$ and B has at most as many edges as $\operatorname{core}(K)$. There is an oriented edge from (A, B) to (C, D) if and only if there is a Whitehead automorphism φ such that C is isomorphic to a subgraph of $\operatorname{core}(\operatorname{fold}(\operatorname{subd}_\varphi(A)))$ and D is isomorphic to a subgraph of $\operatorname{core}(\operatorname{fold}(\operatorname{subd}_\varphi(B)))$; in that case, we label that edge of Ω with the automorphism φ.

Suppose there are representatives $H' \in [H]$ and $K' \in [K]$ and non-trivial free factors J_1, J_2, J_3 such that $J_1 \leq H'$ and $J_1, J_3 \leq J_2$ and $J_3 \leq K'$. By means of Theorems 5.6 and 5.7, we take a chain of Whitehead automorphisms $\varphi_1, \dots, \varphi_l$ such that φ_{i+1} strictly reduces the number of edges of $\operatorname{core}(\varphi_i \circ \dots \circ \varphi_1(J_2))$, and such that the trichotomy of Theorem 5.7 holds too. By Lemmas 5.16, 5.17 and 5.19, we have that this produces a path (A_i, B_i) in Ω, where A_i is the image the map

$$\operatorname{core}(\varphi_i \circ \dots \circ \varphi_1(J_1)) \to \operatorname{core}(\varphi_i \circ \dots \circ \varphi_1(H))$$

induced by the inclusion $J_1 \leq K'$, and B_i is the image of the map

$$\operatorname{core}(\varphi_i \circ \dots \circ \varphi_1(J_3)) \to \operatorname{core}(\varphi_i \circ \dots \circ \varphi_1(K))$$

induced by the inclusion $J_3 \leq K'$. The starting point (A_1, B_1) of the path is given by two subgraphs of $\operatorname{core}(H)$ and $\operatorname{core}(K)$ respectively, and the endpoint (A_l, B_l) has the property that $A_l \sqcup B_l$ only contains edges with labels from a proper subset of $\{x_1, \dots, x_n\}$.

Conversely, suppose there is a path $(A_1, B_1), \dots, (A_l, B_l)$ in Ω with an edge from (A_i, B_i) to (A_{i+1}, B_{i+1}) labeled φ_i, and such that A_1, B_1 are subgraphs of $\operatorname{core}(H)$, $\operatorname{core}(K)$, respectively, and $A_l \sqcup B_l$ contains only edges with labels in $\{x_1, \dots, x_{n-1}\}$. Then we fix basepoints in A_l and B_l, we set $\psi = \varphi_1^{-1} \circ \dots \circ \varphi_l^{-1}$, and we produce the free factors $J_1 = \psi(\pi_1(A_l))$ and $J_2 = \psi(\langle x_1, \dots, x_{n-1} \rangle)$ and $J_3 = \psi(\pi_1(B_l))$. For these free factors, there are representatives $H' \in [H]$ and $K' \in [K]$ such that $J_1 \leq H'$ and $J_1, J_3 \leq J_2$ and $J_3 \leq K'$, as desired.

Since the graph Ω is finite, we obtain an algorithm to check condition (1).

Acknowledgments. I would like to thank my supervisor Professor Martin R. Bridson for having given me the interesting Theorem C to think about, and for all the suggestions he gave me when I was stuck while trying to solve it. I also thank him for all the help he gave me while I was thinking about the free factor complex.

Funding. The author was funded by the Engineering and Physical Sciences Research Council.

References

[1] M. Bestvina and M. R. Bridson, Rigidity of the complex of free factors. 2023, arXiv:2306.05941

[2] M. Bestvina and M. Feighn, Hyperbolicity of the complex of free factors. *Adv. Math.* **256** (2014), 104–155 Zbl 1348.20028 MR 3177291

[3] A. Clifford and R. Z. Goldstein, Subgroups of free groups and primitive elements. *J. Group Theory* **13** (2010), no. 4, 601–611 Zbl 1206.20025 MR 2661660

[4] S. M. Gersten, On Whitehead's algorithm. *Bull. Amer. Math. Soc. (N.S.)* **10** (1984), no. 2, 281–284 Zbl 0537.20015 MR 733696

[5] M. Heusener and R. Weidmann, A remark on Whitehead's cut-vertex lemma. *J. Group Theory* **22** (2019), no. 1, 15–21 Zbl 1441.20016 MR 3895635

[6] P. J. Higgins and R. C. Lyndon, Equivalence of elements under automorphisms of a free group. *J. Lond. Math. Soc. (2)* **8** (1974), 254–258 Zbl 0288.20027 MR 340420

[7] R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*. Class. Math., Springer, Berlin, 2001 Zbl 0997.20037 MR 1812024

[8] D. Puder and C. Wu, Growth of primitive elements in free groups. *J. Lond. Math. Soc. (2)* **90** (2014), no. 1, 89–104 Zbl 1308.20022 MR 3245137

[9] E. S. Rapaport, On free groups and their automorphisms. *Acta Math.* **99** (1958), 139–163 Zbl 0080.01705 MR 131452

[10] J. R. Stallings, Topology of finite graphs. *Invent. Math.* **71** (1983), no. 3, 551–565 Zbl 0521.20013 MR 695906

[11] J. H. C. Whitehead, On certain sets of elements in a free group. *Proc. Lond. Math. Soc. (2)* **41** (1936), no. 1, 48–56 Zbl 0013.24801 MR 1575455

Received 10 November 2021.

Dario Ascari
Mathematical Institute, University of Oxford, Andrew Wiles Building, Radcliffe Observatory Quarter, Woodstock Road, OX2 6GG Oxford, UK; ascari.maths@gmail.com