Definability of Orbits, Algebraic Closures and Elementary Subgroups in the Free Group

Jonathan Fruchter

Preliminaries

The proof of the homogeneity of the free group consists mainly of construcing a formula $\psi_g(x)$ stating that there exists an automorphism of the free group mapping g to x; in other words, that g and x lie in the same orbit under the action of the automorphism group. We show (in a slightly more general setting) that this formula indeed defines the orbit of g, and use the fact that such orbits are definable in order to give an explicit description of the algebraic closure of a subset of the free group. Finally, this description of algebraic closures in the free group will help us conclude the following: every elementary subgroup of the free group is a free factor.

We begin by reminding a few previous results and stating a few facts to be used throughout this document. We denote by $\mathbb{F}$ a non-abelian free group and by $\mathrm{Aut}_A(\mathbb{F})$ the group of automorphisms of $\mathbb{F}$ fixing the subset $A \subset \mathbb{F}$. Unless stated otherwise, we assume that $\mathbb{F}$ is freely indecomposable relative to A and denote by Λ a JSJ decomposition of $\mathbb{F}$ relative to A.

Theorem 0.1. *Suppose that $\langle A \rangle$ is non-abelian. Then one can assume that Λ satisfies the following properties:*

1. *The graph of groups underlying Λ is finite, and every vertex group in this graph is finitely generated.*
2. *A is contained in a unique rigid vertex group H_A, that is H_A is not of surface-type.*
3. *Every edge group is maximal abelian in its endpoints' vertex groups.*

Remark 0.2. We can assume that property 1. above holds also in the case where $\langle A \rangle$ is abelian.

Theorem 0.3. *The group of modular automorphisms of $\mathbb{F}$ fixing A, $Mod_A(\mathbb{F})$, has finite index in $Aut_A(\mathbb{F})$. Furthermore, if $\phi \in Mod_A(\mathbb{F})$ then ϕ restricts to conjugation on non-abelian rigid vertices of Λ and maps surface-type vertex groups isomorphically onto conjugates of themselves.*

Definition 0.4. We say that two homomorphisms $\phi, \varphi : \mathbb{F} \longrightarrow \mathbb{F}$ are $\Lambda-$*related* if the following holds:

1. For each rigid vertex group H of Λ there exists $u_H \in \mathbb{F}$ such that $\phi|_H = \mathrm{Conj}(u_H) \circ \varphi|_H$.
2. For every surface-type vertex group S of Λ, if $\phi(S)$ is non-abelian then $\varphi(S)$ is non-abelian.

Remark 0.5. By theorem 0.3, if $\phi = \sigma \circ \varphi$ for $\sigma \in \mathrm{Mod}_A(\mathbb{F})$ then ϕ and φ are $\Lambda-$related. Furthermore, if a homomorphism $h : \mathbb{F} \longrightarrow \mathbb{F}$ is $\Lambda-$related to the identity then h is an automorphism.

Theorem 0.6. *Fix a homomorphism $h_A : \langle A \rangle \longrightarrow \mathbb{F}$. Then there exist finitely many proper quotients of $\mathbb{F}$, $\{\eta_i : \mathbb{F} \longrightarrow Q_i\}_{i \leq n}$, such that for any non-injective homomorphism $h : \mathbb{F} \longrightarrow \mathbb{F}$ extending h_A there exists a modular automorphism $\sigma \in Mod_A(\mathbb{F})$ such that $h \circ \sigma$ factors via η_i for some $i \leq n$.*

We conclude by mentioning another theorem of a model-theoretic character, opposed to the algebraic theorems appearing above:

Theorem 0.7. *If G is a free factor of $\mathbb{F}$ then G is an elementary subgroup of $\mathbb{F}$.*

1 Definability of Orbits in the Free Group

Theorem 1.1. *Suppose $A \subset \mathbb{F}$ is not contained in any proper free factor of $\mathbb{F}$ and let $\bar{g}$ be a tuple in $\mathbb{F}$. Then the orbit of $\bar{g}$ under $Aut_A(\mathbb{F})$ is definable (over A).*

Proof. Recall that Λ denotes a JSJ decomposition of $\mathbb{F}$ relative to A and fix $q_i \in Q_i$ where $Q_1, ..., Q_n$ are as in theorem 0.6 above (taking h_A to be the inclusion map). We aim to formulate the statement "there exists a homomorphism $h : \mathbb{F} \longrightarrow \mathbb{F}$ fixing A such that $h(\bar{g}) = \bar{x}$, and for every h' which is $\Lambda-$related to h, $h'(q_i)$ is nontrivial for every $i \leq n$".

Note that if this statement holds for some $\bar{g}'$ then by theorem 0.6 there exists an injective $h : \mathbb{F} \longrightarrow \mathbb{F}$ mapping $\bar{g}$ to $\bar{g}'$ and fixing A; by the relative co-Hopf property, h is an automorphism, which implies that $\bar{g}'$ lies in the orbit of $\bar{g}$ under $\mathrm{Aut}_A(\mathbb{F})$. On the other hand, suppose that $\bar{g}'$ lies in the orbit of $\bar{g}$ under $\mathrm{Aut}_A(\mathbb{F})$ so $\bar{g}' = h(\bar{g})$ for some automorphism h of $\mathbb{F}$ which fixes A. If h' is any homomorphism which is $\Lambda-$related to h then $h' \circ h^{-1}$ is $\Lambda-$related to the identity, and by remark 0.5 $h' \circ h^{-1}$ is an automorphism. Hence h' is an automorphism so $h'(q_i)$ is nontrivial for every $i \leq n$ and the statement holds for $\bar{g}'$. Thus, if $\psi_{\bar{g}}(x)$ is a formulation of the statement above then $\psi_{\bar{g}}(\mathbb{F})$ is exactly the orbit of $\bar{g}$ under $\mathrm{Aut}_A(\mathbb{F})$.

Fix a generating set $\bar{s} = (s_1, ..., s_n)$ for $\mathbb{F}$ and write $g_i = w_i(\bar{s})$ where $\bar{g} = (g_1, ..., g_k)$. We divide the construction of the formula $\psi_{\bar{g}}(x)$ into three parts:

1. A formula $\Psi(t_1, ..., t_n, \bar{x})$ stating that the homomorphism defined via $s_i \mapsto t_i$ maps $\bar{g}$ to $\bar{x}$ and fixes A. For every $a \in A$ fix a word w_a such that $w_a(\bar{s}) = a$; consider the system of equations

$$\Sigma_A(\bar{s}, \bar{t}) = \left\{ w_a(\bar{s}) \left(w_a(\bar{t}) \right)^{-1} = 1 : \ a \in A \right\}$$

(where $\bar{s}$ and $\bar{t}$ are tuples of variables). Since free groups are equationally noetherian, $\Sigma_A(\bar{s}, \bar{t})$ is equivalent to a finite subsystem consisted of equations corresponding to some elements $a_1, ..., a_m \in A$. This enables us to write

$$\Psi(\bar{t}, \bar{x}) = \left(\bigwedge_{i=1}^{k} w_i(\bar{t}) = x_i \right) \ \wedge \ \left(\bigwedge_{i=1}^{m} w_{a_i}(\bar{t}) = a_i \right)$$

2. A formula $\Phi(t_1, ..., t_n, r_1, ..., r_n)$ stating that the homomorphism h_t defined via $s_i \mapsto t_i$ is $\Lambda-$related to the homomorphism h_r defined via $s_i \mapsto r_i$. Suppose that the rigid vertices of

Λ are $v_1, ..., v_r$ and that the surface-type vertices of Λ are $u_1, ..., u_s$; denote by H_v the vertex group of the vertex $v \in V(\Lambda)$. Using theorem 0.1, fix a finite generating set $\bar{a}^v$ for H_v and write $\bar{a}^v = \bar{w}^v(\bar{s})$ for every $v \in V(\Lambda)$. Recall that h_t and h_r are $\Lambda-$related if and only if the following is true: for every $i \leq r$ there exists $z_i \in \mathbb{F}$ such that $h_t|_{H_{v_i}} = \text{Conj}(z_i) \circ h_r|_{H_{v_i}}$, and for every $i \leq s$ if $h_t(H_{u_i})$ is non-abelian then so is $h_r(H_{u_i})$. Denote by $\alpha(\bar{x})$ the formula stating that every two elements taken from the tuple $\bar{x}$ commute (fix such formula for every arity); hence $\alpha(\bar{x})$ holds if and only if the tuple $\bar{x}$ generates an abelian group. Therefore one can write

$$\Phi(\bar{t}, \bar{r}) = \left(\exists z_1 \cdots \exists z_r \bigwedge_{i=1}^{r} \left(\bar{w}^{v_i}(\bar{t}) = z_i \left(\bar{w}^{v_i}(\bar{r}) \right) z_i^{-1} \right) \right) \wedge \left(\bigwedge_{i=1}^{s} \left(\neg\alpha(\bar{w}^{u_i}(\bar{t})) \longrightarrow \neg\alpha(\bar{w}^{u_i}(\bar{r})) \right) \right)$$

3. A formula $\Upsilon(r_1, ..., r_n)$ stating that the homomorphism defined via $s_i \mapsto r_i$ does not kill $q_1, ..., q_n$. The construction of such a formula is rather simple: write $q_i = w_{q_i}(\bar{s})$ and set

$$\Upsilon(\bar{r}) = \left(\bigwedge_{i=1}^{n} w_{q_i}(\bar{r}) \neq 1 \right)$$

At last, the three formulae above give us the desired formula:

$$\psi_g(\bar{x}) = \exists \bar{t} \ \left(\Psi(\bar{t}, \bar{x}) \ \wedge \ \forall \bar{r} \ \left(\Phi(\bar{t}, \bar{r}) \longrightarrow \Upsilon(\bar{r}) \right) \right)$$

□

2 A Description of Algebraic Closures

The goal of this section is to describe the algebraic closure of a subset A of $\mathbb{F}$, in the case where the group generated by A is non-abelian. Recall that the algebraic closure of A in $\mathbb{F}$, denoted by $\text{acl}(A)$, concludes of all elements $g \in \mathbb{F}$ which are contained in a finite set which is definable over A. We remark that in the case where $\langle A \rangle$ is abelian (and hence cyclic), a slight modification of the arguments presented in this section can be used to derive the following result:

Theorem 2.1. *If $A \subset \mathbb{F}$ generates an abelian group then $acl(A)$ is the envelope of $\langle A \rangle$, that is*

$$acl(A) = \{ g \in \mathbb{F} : \ \exists n \ g^n \in \langle A \rangle \}$$

In the case where A contains non-commuting elements, we have the following:

Theorem 2.2. *Suppose $A \subset \mathbb{F}$ generates a non-abelian group. Denote by F_A the smallest free factor of $\mathbb{F}$ containing A and let Λ be a JSJ decomposition of F_A relative to A. Then $acl(A)$ is the vertex group of Λ containing A.*

Remark 2.3. Bring to mind that theorem 0.1 implies that there exists a *unique* vertex group in Λ containing A.

Before getting hands on with the proof, observe the following: if the orbit of an element $g \in \mathbb{F}$ is infinite, then $g \notin \text{acl}(A)$. This follows from the fact that any element g' in the orbit of g shares the same type as g (over A), and thus the orbit of g is contained in every definable set $X \subset \mathbb{F}$

containing g. Therefore, our strategy of choice will be to show that whenever g is not in the vertex group of Λ containing A then the orbit of g under $\mathrm{Aut}_A(\mathbb{F})$ is infinite, hence $g \notin \mathrm{acl}(A)$. Denote this vertex group by H_A, and the two lemmas to follow will imply that $\mathrm{acl}(A) \subset H_A$.

Lemma 2.4. *If $g \notin F_A$ then its orbit under $Aut_A(\mathbb{F})$ is infinite.*

Proof. Write $\mathbb{F} = F_A * H$ and $g = f_1 h_1 \cdots f_n h_n$ where $f_i \in F_A$ and $h_i \in H$. Since $g \notin F_A$, at least one nontrivial h_i appears in this product. Since F_A is a non-abelian free group, there exists $u \in F_A$ which does not commute with any (nontrivial) f_i. Let $\tau_n \in \mathrm{Aut}_A(\mathbb{F})$ be given by the identity on F_A and conjugation by u^n on H (τ is indeed an automorphism since the homomorphism defined by the identity on F_A and conjugation by u^{-n} on H is its inverse).

It is enough to show that $\tau_n(g) \neq \tau_m(g)$ for every $n \neq m$. And indeed, note that

$$\tau_n(g) = (f_1 u^n)\, h_1 \left(u^{-n} f_2 u^n\right) \cdots \left(u^{-n} f_r u^n\right) h_r \left(u^{-n}\right)$$

and by the uniqueness of the normal form of an element in the free product, since $u^n \neq u^m$ one has that $\tau_n(g) \neq \tau_m(g)$. Clearly any τ_n can be extended to an automorphism of $\mathbb{F}$ fixing A which completes the proof. □

Lemma 2.5. *If $g \in F_A$ is not contained in H_A then its orbit under $Aut_A(\mathbb{F})$ is infinite.*

Proof. Recall that by theorem 0.1 we may assume that each edge group is maximal abelian in its endpoints vertex groups. Fix a maximal subtree $\mathcal{T}$ of Λ, and denote by Λ_g the minimal subgraph of Λ such that g is contained in $\pi_1(\Lambda_g)$ (computed with respect to the maximal subtree $\mathcal{T} \cap \Lambda_g$). If Λ_g contains some edge e with corresponding edge group $\langle c \rangle$, collapse all edges of Λ but e to obtain a splitting of F_A. Distinguish between the two following cases:

1. The splitting obtained is an HNN extension $U *_t$ with $H_A \leq U$ and $g \notin U$.

   Consider the Dehn twist τ_{c^n} which is the identity on U and maps t to tc^n. Since τ_{c^n} can be extended to an automorphism of $\mathbb{F}$ fixing A it is enough to show that $\tau_{c^n}(g) \neq \tau_{c^m}(g)$ for every $n \neq m$; this comes down to a standard calculation with normal forms.
   Denote the two images of $\langle c \rangle$ in H_A by C_1 and C_2, identify c with its image in C_1 and denote by $\bar{c}$ the image of c in C_2; since $g \notin U$ and C_1, C_2 are maximal abelian in H_A, g can be written in normal form $g = g_0 t^{\epsilon_1} g_1 \cdots t^{\epsilon_k} g_k$ where $k > 0$, if $\epsilon_i = 1$ and $\epsilon_{i+1} = -1$ then $g_i \notin C_1$ (and thus g_i does not commute with c) and if $\epsilon_i = -1$ and $\epsilon_{i+1} = 1$ then $g_i \notin C_2$ (and in this case, g_i does not commute with $\bar{c}$). Striving for a contradiction, assume that $\tau_{c^n}(g)\left(\tau_{c^m}(g)\right)^{-1} = 1$ for some $n \neq m$. Note that if ℓ is the maximal index for which g_ℓ is nontrivial, then $\epsilon_{\ell+1}, ..., \epsilon_k$ are all equal. Assume without loss of generality that $\epsilon_{\ell+1} = -1$ (and thus $\epsilon_{\ell=1}$) and note that (in the case where $\epsilon_{\ell+1} = 1$ the proof is carried out in a similar manner using the fact that $tc = \bar{c}t$)

$$\begin{aligned} 1 &= \tau_{c^n}(g)\left(\tau_{c^m}(g)\right)^{-1} \\ &= \left(g_0 t^{\epsilon_1} c^{\epsilon_1 \cdot n} g_1 \cdots g_{\ell-1} (\bar{c})^{\epsilon_\ell \cdot n}\right) t^{\epsilon_\ell} g_\ell c^{\epsilon \cdot (n-m)} (g_\ell)^{-1} t^{-\epsilon_\ell} \left((\bar{c})^{-\epsilon_\ell \cdot m} \cdots g_0^{-1}\right) \end{aligned}$$

   where $\epsilon = \epsilon_{\ell+1} + \cdots + \epsilon_k \neq 0$. By Britton's lemma, since the word above is equal to the trivial element, $g_\ell c^{\epsilon \cdot (n-m)} (g_\ell)^{-1}$ is contained in C_1 (recall we assume that $\epsilon_\ell = 1$). Since g_ℓ

and c do not commute, c^{n-m} must be trivial, a contradiction. Hence, since each τ_{c^n} can be extended to an automorphism of $\mathbb{F}$ fixing A, the orbit of g under $\mathrm{Aut}_A(\mathbb{F})$ is infinite.

2. The splitting obtained is an amalgamated product $U *_{\langle c \rangle} V$ with $H_A \leq U$ and $g \notin U$.

   Again, consider the Dehn twist τ_{c^n} which is the identity on U and conjugation by c^n on V. In this case, similar to the previous one, we claim that $\tau_{c^n}(g) \neq \tau_{c^m}(g)$ for every $n \neq m$. As before, this comes down to a calculation with normal forms which we present briefly: since $g \notin U$, one can write $g = zu_1v_1 \cdots u_kv_k$ where $z \in \langle c \rangle$, at least one v_i is nontrivial, each u_i is taken from a fixed set of coset representatives of $\langle c \rangle$ in U and each v_i is taken from a fixed set of coset representatives of $\langle c \rangle$ in V. Suppose by contradiction that $\tau_{c^n}(g)\,(\tau_{c^m}(g))^{-1} = 1$, and reduce this product to its normal form; this normal form must coincide with the trivial element by uniqueness. During this process, one obtains that v_ℓ, where ℓ is the maximal index for which v_ℓ is nontrivial, must commute with c^{n-m}. Since $\langle c \rangle$ is maximal abelian in the vertex groups of the endpoints of e, v_ℓ must lie in $\langle c \rangle$, a contradiction. It follows that the orbit of g under $\mathrm{Aut}_A(\mathbb{F})$ is infinite.

It is left to verify that the orbit of g under $\mathrm{Aut}_A(\mathbb{F})$ is infinite also in the case where Λ_g does not contain an edge, that is when there is some vertex group H containing g. Note that F_A is obtained from $\pi_1(\mathcal{T})$ by taking finitely many HNN extensions, and thus each automorphism of $\pi_1(\mathcal{T})$ can be extended to an automorphism of F_A (which can be extended to an automorphism of $\mathbb{F}$). Thus it suffices to find infinitely many automorphisms of $\pi_1(\mathcal{T})$ which map g to distinct elements.

Consider the unique path connecting H_A to H in $\mathcal{T}$, and let e an edge appearing along this path. Then again, collapsing all edges of Λ but e, one obtains a splitting of F_A as an amalgamated product $U *_{\langle c \rangle} V$ with $H_A \leq U$ and $g \notin U$. Thus by applying the process described in case 2. above, the fact that the orbit of g under $\mathrm{Aut}_A(\mathbb{F})$ is infinite follows. □

Proof. (of theorem 2.2). With these two lemmas in hand, as observed earlier we obtain that $\mathrm{acl}(A) \subset H_A$. Before showing the inclusion in the other direction, we point out that by theorem 0.3 any automorphism $\phi \in \mathrm{Mod}_A(F_A)$ restricts to conjugation by some u on H_A; since ϕ also fixes the non-abelian group $\langle A \rangle$, u must commute with every $a \in A$ which can only happen in the case where $u = 1$. Thus any modular automorphism of F_A fixing A must fix the entire vertex group H_A.

Let $g \in H_A$ and as a consequence of the observation stated above the orbit of g under $\mathrm{Mod}_A(F_A)$ is simply $\{g\}$. By theorem 0.3, $[\mathrm{Aut}_A(F_A) : \mathrm{Mod}_A(F_A)] < \infty$ which implies that the orbit of g under $\mathrm{Aut}_A(F_A)$ is finite. Theorem 1.1 tells us that this orbit is a definable set (over A) in F_A, and thus $g \in \mathrm{acl}^{F_A}(A)$ (where $\mathrm{acl}^{F_A}(A)$ denotes the algebraic closure of A in the group F_A). To finish, recall that F_A is a free factor of $\mathbb{F}$ so by theorem 0.7 $F_A \prec \mathbb{F}$. Hence $\mathrm{acl}^{F_A}(A) = \mathrm{acl}(A)$ which finishes the proof. □

3 Elementary Subgroups of the Free Group

Composing the two results established in the previous sections, one can obtain the converse to theorem 0.7. This gives a complete correspondence between elementary subgroups of $\mathbb{F}$ and free factors of $\mathbb{F}$.

Theorem 3.1. *If H is an elementary subgroup of $\mathbb{F}$ then H is a free factor.*

Proof. Let F be the smallest free factor containing H. Since $H \subset F$, $H \preceq \mathbb{F}$ and $F \preceq \mathbb{F}$, H is an elementary subgroup of F. Assume to derive a contradiction that $H \neq F$, and let Λ be the JSJ decomposition of F relative to H. Consider the following cases: □

1. Λ is nontrivial. Thus Λ contains a vertex whose corresponding group H' does not contain H. Let $g \in H'$ and by lemma 2.5 its orbit under $\mathrm{Aut}_H(F)$ is infinite; the proof of theorem 2.2 implies that the orbit of g does not meet H. By theorem 1.1 there exists a formula $\psi_g(x)$ (with parameters from H) defining the orbit of g in F under $\mathrm{Aut}_H(F)$. Note that $F \models \exists x \psi_g(x)$ and hence the same must hold in H. This yields an element $h \in H$ such that $H \models \psi_g(h)$. Since $H \preceq F$ it follows that $F \models \psi_g(h)$, contradicting the fact that the orbit of g does not meet H.

2. Λ is trivial, that is Λ contains the unique vertex (with corresponding group) F. Let $g \in F \backslash H$ and by theorem 2.2 its orbit under $\mathrm{Aut}_H(F)$ is a finite set, say of size n. Let $\psi_g(x)$ be a formula (with parameters from H) defining this orbit. Define

$$\Psi = \exists x_1 \cdots \exists x_n \left(\bigwedge_{i,j \leq n} x_i \neq x_j \right) \wedge \left(\bigwedge_{i=1}^{n} \psi_g(x_i) \right) \wedge \left(\forall y \left(\bigwedge_{i=1}^{n} y \neq x_i \right) \longrightarrow \neg \psi_g(y) \right)$$

that is Ψ is a sentence (over H) stating that there are exactly n elements in the set $\mathcal{O}$ defined by $\psi_g(x)$. Since $H \preceq F$, the set defined by $\psi_g(x)$ in H is $\mathcal{O} \cap H$. Furthermore, $H \models \Psi$ which implies that $\mathcal{O} \cap H$ contains as many elements as $\mathcal{O}$, that is $\mathcal{O} \cap H = \mathcal{O}$. This contradicts the fact that $g \in \mathcal{O} \backslash H$.