Dan Lidral-Porter

FREE PRODUCTS AND BRITTON'S LEMMA

1. Free Products

I determined that the best jumping off point was to start with free products. Free products are an extension of the notion of free groups. Recall that a free group F of some set X is a group containing X where every function $f : X \to G$ (where G is a group) can be extended to a homomorphism $\varphi : F \to G$. The free group is thus the "largest" group that can be generated from X, and similarly the free product is the "largest" group generated from a collection of groups in such a way that the groups retain their structure.

Definition. *The **free product** P of some collection of groups $\{A_\alpha\}$ is a group such that:*

(i) *For every A_α, there exists a subgroup of P that is isomorphic to A_α. That is, there is a homomorphism $\iota_\alpha : A_\alpha \to P$ that is injective.*

(ii) *For every group G with a collection of homomorphisms $f_\alpha : A_\alpha \to G$, there exists a unique homomorphism $\varphi : P \to G$ such that for each α, $\varphi\iota_\alpha = f_\alpha$.*

The first condition is analogous to the requirement that the free group F contains X, and the second is almost identical to the salient feature of a free group. Indeed, the commutative diagrams are almost identical.

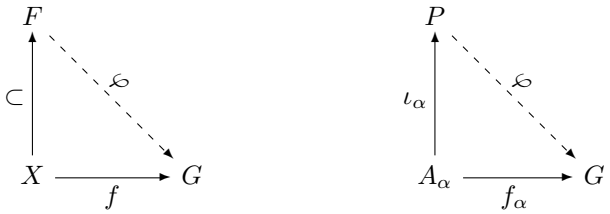

The only differences is that the maps across the left and bottom in the first diagram are set maps, while the ones in the second diagram are group homomorphisms. Between the two, a little is lost, and a little is gained—while the maps that φ must extend are now well-behaved, there are a plurality of maps, instead of just one. The construction of a free product of some given groups is very similar to the construction of a free group of a given set. However, before the free product can be constructed, some notation which will be used throughout this writeup must be introduced.

First call the set $\cup A_\alpha$ the *alphabet*. Note that it is assumed that the A_α are pairwise disjoint; this can always be achieved by replacing some of the A_α with isomorphic copies. The elements of the alphabet are called *letters*. Call the set of all sequences of letters S. A word w in $\cup A_\alpha$ is a sequence

$$w = \{a_1, a_2, a_3, ..., a_{k-1}, 1, 1, 1, ...\} \in S.$$

That is, a sequence in S where there is some integer k such that $n \geq k \Rightarrow a_n = 1$. Note that 1 may appear in the word before a_k. The sequence where $k = 1$ is called

the *empty word* and is denoted as 1. A word in $\cup A_\alpha$ is called a *reduced word* if it satisfies the following:

(i) a_k is the first 1 to appear in the word.

(ii) a_i and a_{i+1} are never members of the same A_α.

Words will be denoted by writing their finitely many non-1 terms. Note that spellings of reduced words are unique by the definition of sequence equality.

Existence of the Free Product. *Let $\{A_\alpha\}$ be a collection of groups, then there is free product of those groups.*

Proof. The free product is constructed out of the reduced words of $\cup A_\alpha$, with the operation being juxtaposition of the words. Let W be the set of all reduced words in $\cup A_\alpha$. For every $a \in \cup A_\alpha$, one can define an operation $|a|$ on W, where

$$|a|(a_1, a_2, ..., a_n) = \begin{cases} a, a_1, a_2, ..., a_n & \text{if } a \text{ and } a_1 \text{ are not in the same } A_\alpha \\ aa_1, a_2, ..., a_n & \text{otherwise.} \end{cases}$$

Note that since both $|a||a^{-1}|$ and $|a^{-1}||a|$ are the identity on W, $|a|$ and $|a^{-1}|$ must be isomorphisms that take given permutations of W to different permutations. Denote the set of permutations of W as S_W. Let P be the subgroup of S_W generated by $\{|a| : a \in \cup A_\alpha\}$. The claim is that P is the free product of the A_α.

Note that every reduced word $a_1, a_2, ..., a_n$ can be factored into the series of operations $|a_1||a_2|...|a_n|1$. Because spellings of reduced words are unique, this factorization yields distinct permutations of W. At this point the injective homomorphisms ι_α can be constructed. The map takes every $a \in A_\alpha$ to $|a|$. The map is injective due to the uniqueness of the factorizations, and it is a homomorphism because

$$\begin{aligned} \iota_\alpha(aa') &= |aa'|1 \\ &= aa' \\ &= |a||a'|1 \\ &= \iota_\alpha(a)\iota_\alpha(a'). \end{aligned}$$

Let a group G and a collection of homomorphisms $\{f_\alpha : A_\alpha \to G\}$ be given. The desired homomorphism $\varphi : P \to G$ is defined as

$$\varphi(|a_1||a_2|...|a_n|) = f_\alpha(|a_1|)f_{\alpha'}(|a_2|)...f_{\alpha''}(|a_n|)$$

Because factorizations of reduced words are unique, φ is well defined. Now, consider two words β and $\delta \in P$. We need to show that $\varphi(\beta\delta) = \varphi(\beta)\varphi(\delta)$. The only concern is the possibility of combining adjacent elements from the same A_α in one of those terms but not the other. Since φ is defined in terms of the homomorphisms f_α, such combinations will pass right through φ, so φ is homomorphic. It is clear that φ is unique, since it is defined by the actions of the f_α on the generating elements of P. □

Now we know that the free product can be constructed. However, its construction was quite ungainly, and I would be lying if I said that I was confident of my treatment of its construction. It required the introduction of an entirely new concept (words and reduced words), and the argument lacked a sense of coherency.

As contrast, the proof of the uniqueness of the free product highlights the utility of the category-theoretic approach. Whereas in the construction of the free product one becomes bogged down in the elements of the groups in question, the uniqueness

of the free product follows naturally when one considers the natures of the maps between the objects in question.

Uniqueness of the Free Product. *Let $\{A_\alpha\}$ be a family of groups. If G and H are free products of those A_α, then $G \approx H$.*

Proof. Since G and H are both free products of the A_α, they have collections of injective homomorphisms, ι_α and κ_α respectively, which take each A_α to the respective free product. Thus we have the commutative diagram:

$$\begin{array}{ccc} G & & H \\ {\scriptstyle \iota_\alpha}\uparrow & & \uparrow{\scriptstyle \kappa_\alpha}. \\ A_\alpha & \xrightarrow[\mathrm{id}_\alpha]{} & A_\alpha \end{array}$$

Let $f_\alpha = \kappa_\alpha \mathrm{id}_\alpha$. Since f_α is a homomorphism from A_α to H, and since G is a free product of the A_α there is a unique extension of f_α to a homomorphism $\varphi : G \to H$. Now let $g_\alpha = \iota_\alpha \mathrm{id}_\alpha$. By the same reasoning, there exists a unique extension of g_α to a homomorphism $\psi : H \to G$. Their compose $\psi\varphi : G \to G$ which gives the diagram:

$$\begin{array}{ccc} G & \xrightarrow{\psi\varphi} & G \\ {\scriptstyle \iota_\alpha}\uparrow & & \uparrow{\scriptstyle \iota_\alpha} \\ A_\alpha & \xrightarrow[\mathrm{id}_\alpha]{} & A_\alpha \end{array}$$

This diagram does commute. This is due to the properties of φ and ψ, specifically, that $\varphi\iota_\alpha = f_\alpha = \kappa_\alpha$, and $\psi\kappa_\alpha = g_\alpha = \iota_\alpha$. Thus the compose of the left and top sides of the diagram is

$$\begin{aligned} \psi\varphi\iota_\alpha &= \psi\kappa_\alpha \\ &= \iota_\alpha. \end{aligned}$$

The definition of the free product states that the homomorphism which extends the f_α is unique. Note that the identity on G also has the property $I_G\iota_\alpha = \iota_\alpha$, so $\psi\varphi = I_G$. Similarly, $\varphi\psi = I_H$, so the homomorphisms φ and ψ must be isomorphisms. □

Free products with amalgamated subgroups are a slight variation on free products, which figure centrally in Britton's Lemma to the Novikov-Boone Theorem.

Definition. *Let B be a group and let $\{A_\alpha\}$ be a collection of groups. For each A_α, let there be a subgroup B_α of A_α, and an isomorphism $\varphi_\alpha : B \to B_\alpha$. The **free product of the A_α with amalgamated subgroup** B is the group*

$$((*A_\alpha) * B)/N,$$

where N is the normal subgroup of the free product generated by all elements of the form $b\varphi_\alpha(b^{-1})$ where $b \in B$.

This is in rough terms the "largest" group where all the subgroups B_α of the A_α remain identified through the isomorphisms φ_α. Note that if the subgroup B is trivial, the free product with amalgamated subgroup reduces to simply being the free product.

2. Turing Machines, the Novikov-Boone Theorem, and Britton's Lemma

The Novikov-Boone Theorem is more or less the group-theory analogue of the proof of the endecidability of the halting problem. It states that there are groups that have an unsolvable word problem, which means that there are simple questions about these groups that no decidable algorithm can answer.

This writeup will not address the Novikov-Boone Theorem in particular, but it will lay down some of the groundwork, in particular, Britton's Lemma.

In this section groups will often be referred to by their presentations, that is, their sets of generators and relations, denoted

$$G = (X|\Delta),$$

where X is the set of generators and Δ the set of relations. The groups we will be concerned with are those that have a finite presentation.

Definition. *A group has a* ***finite presentation*** *if its presentation contains a finite number of generators and a finite number of relations.*

It has been proven that there exist finitely generated groups that have an infinite number of relations. Some classes of groups are easily shown to be finitely presented:

- Every finite group is finitely presented, since if the group has rank n, there are at most n^2 relations (write out the multiplication table).
- Every free group of a finite set is finitely presented, since free groups have no relations.
- Every free product of finitely presented groups is also finitely presented, since the free product does not introduce any further relations into the group.

Suppose we have some group G with finite presentation

$$G = (x_1, x_2, ..., x_n|\Delta).$$

The **word problem** for G is **solvable** if there exists an algorithm to decide all questions of the form "is the word w in the x_i the identity element of G?" An algorithm is said to **decide** a question if it can be guaranteed to return an answer in a finite number of steps.

Note that these words are not necessarily reduced. The **length** of a word w in some x_i is the number of x_i and x_i^{-1} which appear in the spelling of the word. Thus, the empty word has length 0, while the word $x_1x_1^{-1}$ has length 2 (even though these words are equivalent after cancellation).

As an example of a group with a solvable word problem, we will show that every free group's word problem is solvable. A free group has presentation

$$F = (x_1, x_2, ..., x_n|\phi).$$

Let a word w in the generators of F be given. The following is an algorithm that decides F's word problem:

(1) if w has length 0 or 1, go to step 4.
(2) Underline the first consecutive pair $x_i x_i^{-1}$ which appears in the word. If no such pair appears in the word, go to step 4.
(3) Remove the two underlined letters from the word w. Go to step 1.
(4) If w has length 0, it is the identity element of F. Otherwise, it's not.

Since any word in the x_i has finite length, this algorithm can be guaranteed to finish in a finite number of steps (because steps 2 and 3 either reduce the length of w or halt).

To bring the question of decidability into the algebraic sphere, an algebraic description of Turing Machines is defined. This writeup assumes the reader is familiar with the operations of a Turing Machine.

Let $s_1, s_2, ...$ and $q_1, q_2, ...$ be infinite lists which will be used for the tape alphabet and states of the tape head, respectively. A Turing Machine is then a collection of 4-tuples over these two lists and two additional symbols L and R, where each 4-tuple is of one of the following forms:

1. $q_i s_j s_k q_l$
2. $q_i s_j R q_l$
3. $q_i s_j L q_l$,

and no two 4-tuples have the same q_i and s_j. The first type of 4-tuple is interpreted as replacing the current tape symbol, the second as moving the tape head right, and the third as moving the tape head left.

Note that the state of a Turing Machine can be described as a word on the s_i and q_j with exactly one q_j (which is not at the end of the word). This means that, at that moment, the tape consists of the s_i in the order that they appear in the word, the tape head's state is that of the q_j which appears in the word, and it is currently scanning the s_i which appears after the q_j in the word. Such a word is called an **instantaneous description**.

Now suppose we have a Turing Machine T, and two instantaneous descriptions α and β. If there are tuples $\in T$ that allow the Turing Machine to start in the description α, and then (in one move) end in the description β. If this is the case, then we say that $\alpha \to \beta$.

If we consider the symbols s_i and q_j which appear in the tuples of T to be generators, and all moves $\alpha \to \beta$ of T to be relations $\alpha = \beta$, then T gives the presentation of a semigroup $\pi(T)$.

If the reader has some familiarity with Computer Science, then they may suspect that there must exist a semigroup with an unsolvable word problem, since if one wishes to determine whether a particular word in a semigroup is the identity element, one may not perform cancellations (since there are no inverses), but only substitutions according to the relations. Since these relations can be interpreted as moves of a Turing Machine, and there exist Turing Machines with undecidable halting problems, one would expect that this means there are semigroups with undecidable word problems. Indeed, this is was proven by Post [1].

The basic idea of Boone's proof of the Novikov-Boone theorem is to generate a group G from s_i, q_j, t, and k, where the s_i and q_j are the tape and state symbols of a Turing Machine T. He then uses a lemma which states that if $\Sigma \in G$ is an instantaneous description, then $(\Sigma^{-1} t \Sigma)k = k(\Sigma^{-1} t \Sigma)$ in G if and only if $\Sigma = q_0$ in $\pi(T)$. This means that if there were a decidable algorithm to determine whether any

two words in G are equal, then there would be a decidable algorithm to determine whether a Turing Machine halted in the state Σ (that is, if the semigroup $\pi(T)$ had a solvable word problem). Boone proved the lemma with a combinatorial argument; subsequently J.L. Britton developed a proof which uses free product with amalgamated subgroups to argue that certain kinds of elementary operations (that is, canceling inverses, inserting inverses in place of 1, and substituting) can be performed inside a semigroup (where one can only substitute). This writeup presents a partial writeup of Britton's lemma.

First, a few more definitions.

Definition. *A word of the free product of $\{A_\alpha\}$ with amalgamated subgroup B is said to be a **normal form** if it has the form*

$$a_1, a_2, ..., a_n, b$$

where $b \in B, n \geq 0$, and every adjacent a_i lies in a distinct A_α.

Definition. *Let X and Y be words which may not be reduced. We say that $X \equiv Y$ if X and Y have exactly the same spelling (without performing insertions, cancellations, or substitutions).*

Remember that $X = Y$ in the case where X and Y determine the same element of a group.

Definition. *Let W be a word on some generators $x_1, ..., x_n$. Y is a **subword** of W if $W \equiv XYZ$. W **involves** x_i if x_i is a subword of W.*

Definition. *Let $H = (S|D)$ and $H^* = (S^*|D^*)$ be group presentations. $H \leq H^*$ if $S \subset S^*, D \subset D^*$, and for every word W on S, $W = 1$ in H if and only if $W = 1$ in H^*.*

Lemma. *Let $H = (S, D)$ and $H^* = (S, t|D, t^{-1}X_i t = X_i \ \forall i \in I)$, where the X_i are words on the S alone. Let W be a word on the generators of H^* involving t. If $W = 1$ in H^*, then W contains a subword of the form $t^{-1}Ct$ or tCt^{-1}, where*

(i) *C is a word on the S alone.*
(ii) *The element of H determined by C lies in the subgroup of H generated by the X_i.*

Proof. Let X be the subgroup of H generated by the X_i of H^* (this is valid because the X_i of H^* are words on the generators of H). Let X' be isomorphic to X through the isomorphism $\psi : X \to X'$. Note that there is a presentation of X'

$$X' = \ x_i \ \forall i \in I | r_j \ \forall j \in J),$$

where each $x_i = \psi(X_i)$, and the r_j are words on the x_i alone. This is because X' is isomorphic to X, so its generators must be the image of the generators of X. We don't know anything specific about r_j except that they are determined by the set of relations D of H.

Now let $[t]$ be the infinite cyclic group of powers of t. Set $Y = X' \times [t]$. Then take the free product of Y and H in which the subgroups X and X' are amalgamated through ψ. Because free products with amalgamated subgroups do not add any new generators or relations, while preserving the isomorphism of the amalgamated subgroups, one presentation of A must be

$$A = (S, t, x_i \ \forall i \in I | D, r_j = 1 \ \forall j \in J, t^{-1}x_i t = x_i, x_i = X_i \ \forall i \in I).$$

However, due to the isomorphism between X and X', another presentation is

$$A = (S, t, x_i \ \forall i \in I | D, R_j = 1 \ \forall j \in J, t^{-1} X_i t = X_i \ \forall i \in I),$$

where $\psi(R_j) = r_j$ for each j. Now the X_i are words on X, which is a subgroup of H, therefore each $R_j = 1$ in H. Now, because the generators of H are a subset of the generators of A, and the relations of H are a subset of the relations of A, there is a homorphism φ which takes any word in H to "itself" in A. Since all the relations in H are also relations of A, and each $R_j = 1$ in H, each $R_j = 1$ in A.

Note: Rotman's treatment claims that this shows that the R_j are thus "superfluous" in the second presentation of A, but I am not entirely sure why this is so. It seems that it must be because the relations D of H supercede the R_j (which makes sense, because the R_j are isomorphic to the r_j, which are relations of an isomorphic copy of a subgroup of H), but in that case why did we have to argue that the R_j are 1 in A, as well as in H? There must be a subtlety I am missing.

Once the superfluity of the R_j in the second presentation of A is granted, we have a third presentation

$$A = (S, t | D, t^{-1} X_i t = X_i \ \forall i \in I),$$

which is identical to the presentation of H^* given at the beginning of the lemma. Therefore, $A \approx H^*$. Because of the properties of the free product with amalgamated subgroup, H^* contains H as a subgroup (up to isomorphism). This means that if a word on the S is 1 in H, it is 1 in H^* as well, and vice versa, so $H \leq H^*$.

Now to prove the lemma. Let W be a word on $\{S, t\}$ involving t such that $W = 1$ in H^*. If W contains tt^{-1} or $t^{-1}t$ as a subword, then we are done. Therefore, assume W takes the form

$$W \equiv W_0 t^{e_1} W_1 ... t^{e_n} W_n,$$

where $n \geq 1$, each e_j is a nonzero integer, each W_j is a word on the S alone, and only W_0 and W_n may be empty. The proof will proceed by induction on n.

If $n = 1$, then $W \equiv W_0 t^{e_1} W_1$. Because $W = 1$ in H^*, $W = 1$ in both H and Y, so $t^{e_1} = W_0^{-1} W_1^{-1}$ in $H \cap Y$. We know from the properties of the amalgamated free product that the intersection of H and Y must be the amalgamated group X, so from this one can conclude that $t^{e_1} \in X$. However, X is generated by the S, so this is a contradiction.

The inductive step uses as justification the normal form theorem, which I did not cover here. It states that every element of the amalgamated free product has a unique representative which is of normal form. If P is the free product of some A_α with amalgamated subgroup B, and the letters of the word $a_1, a_2, ..., a_n$ are elements of P where each a_1 lies in a distinct A_α, then the element determined by that word cannot live in B, because of the normal form theorem (particularly, the word cannot be 1). In this instance, because W is 1 in H^*, one of the W_j where $0 < j < n$ must lie in the amalgamated group X (remember that W_0 and W_n are allowed to be empty). If e_j and e_{j+1} have opposite signs, the lemma is satisfied. Suppose they have the same sign. That means that, in H^*,

$$W \equiv ... t^{e_j} W_j t^{e_{j+1}} W_{j+1} = ... t^{e_j + e_{j+1}} W_j W_{j+1} ...$$

since H^* is the free product of H and Y with amalgamated X', and Y contains the infinite cyclic (and therefore commutative) $[t]$ as a normal subgroup. Note that the second word in the display has $n - 1$ occurrences of a power of t, so it satisfies the

inductive hypotheses. Therefore, the second word contains the desired subword, so W must contain that subword as well.

□

Note: I used Rotman's "The Theory of Groups" [2] most heavily, especially with regards to the word problem. Robinson's [3] "A Course in the Theory of Groups" was also helpful.

References

[1] E. L. Post, "Recursive unsolvability of a problem of thue," *J. Symb. Log.*, vol. 12, no. 1, pp. 1–11, 1947.

[2] J. J. Rotman, *The theory of groups, an introduction [by] Joseph J. Rotman.* Allyn and Bacon, Boston,, 1965.

[3] D. J. Robinson, *A Course in the Theory of Groups*, vol. 80 of *Graduate Texts in Math.* New York, Heidelberg, Berlin: Springer-Verlag, 1982.