FREE PRODUCT FACTORIZATION

PETER NELSON

Abstract. In this paper proves a result about free products of groups, namely, that if a group is finitely generated, it can be factored uniquely into the free product of "indecomposable" groups. This proof uses the theory of fundamental groups and covering spaces, so some basics of algebraic topology are introduced along the way

Contents

1. Introduction

The goal of this paper is to prove a result on free products of groups, using topological methods. We assume that the reader has some knowledge of basic group theory, in particular, the notions of subgroup, homomorphism, normal subgroup, conjugation, and quotient groups, as well as group presentations. Some comfort with point set topology may be helpful, especially the product and quotient topologies. Basic definitions and theorems from algebraic topology (but not proved) are given and used often, but familiarity with them is not necessary.

2. Free Products

Our goal in this paper is to study properties of a certain kind of product of groups, called the free product, so there seems to be no place better to start than with the definition:

Definition 2.1. The *free product* of two groups G_1 and G_2, denoted $G_1 * G_2$ is defined as the unique group H such that there exist injective homomorphisms $i_1 : G_1 \to H$ and $i_2 : G_2 \to H$ and for any homomorphisms $f : G_1 \to X$, $g : G_2 \to X$ (where X is any group), there exists a unique homomorphism $\varphi : H \to X$ such that the following diagram commutes.

Date: August 20, 2008.

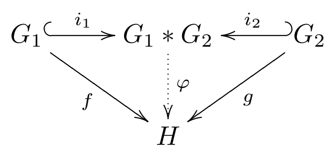

For those familiar with categorical language, this is just the coproduct in the category of groups. For those not so familiar, a more concrete description will likely be more useful. The free product of two groups G and H has as an underlying set the "reduced" words on nonidentity elements of the groups, alternating between the groups. (Here, "reduced" means if two elements of the same group are next to each other, we replace them with their product in the same group. That is if $g_1, g_2, \ldots \in G$ and $h_1, h_2, \ldots \in H$, then $g_1h_1g_2$ and $h_2g_2h_1g_2h_3$ are "legal" words in the free product, for example, but $g_1g_2h_2$ is not, since it could be reduced by substituting the *element* $g_1g_2 \in G$. If g_1 and g_2 are inverses, we get the identity, which we then simply remove. The binary operation in the free product is given by "concatenate and reduce," that is, stick one word after the other and reduce until the result is a reduced word. Then the identity is the empty word, and inverses are given by reversing the order of letters and then replacing letters by their inverses. One can check that this is indeed a group, and satisfies the property in the definition.

A couple more things deserve mentioning. Free products are never abelian unless one group is trivial. Also, one can take the free product of an arbitrary number of groups, by requiring one injective homomorphism for each factor in the definition, and then given one homomorphism from each factor to X, there is a unique homomorphism from the free product making the larger diagram commute.

One special case is very important:

Definition 2.2. A *free group* is the free product of some number (possibly infinitely many) copies of $\mathbb{Z}$. The number of copies is called the *rank* of the free group. Also, the trivial group is the free group of rank 0 (being the free product of 0 copies of X).

One can also think of this as reduced words on some number of letters (one for each copy of $\mathbb{Z}$), again under the operation "concatenate and reduce." We will also need a generalization of free product a few times.

Definition 2.3. Let A, B, C be groups, and f, g be homomorphisms, $f : C \to A$, $g : C \to B$. Then the *free product of A and B amalgamated over f and g* (or *the free product amalgamated over C* if the maps are injective and understood, which they usually will be) is the unique group D and pair of homomorphisms $\varphi : A \to D$ and $\psi : B \to D$ such that given any other pair of homomorphisms $\varphi' : A \to X, \psi' : B \to X$, there is a unique homomorphism h making the following diagram commute:

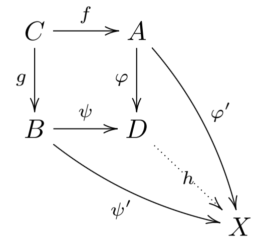

This is written as $A *_C B$.

Again, this may not be so enlightening for those unfamiliar with categorical language (for those who are familiar, its just a pushout). A more concrete description is that it is the free product of the group, but identifying the images of C (we quotient by the normal subgroup generated by things of the form $f(c)g(c^{-1})$ or $g(c)f(c^{-1})$. A more intuitive way of thinking about this comes if f, g are injections. Then C can be thought of as a subgroup of A and of B, and then the free product of A and B amalgamated over C is the free product of A and B, but with the subgroups corresponding to C thought of as the same subgroup. One should also note that if C is the trivial group, this is just the free product.

3. Some Topology

To do what we want to do with free products, we also need some machinery from topology. We will not prove that any of it works and take it as black box. For those who want proofs, see Hatcher's book [1].

In topology, if one space can be "continuously deformed" into another, we don't really distinguish them (as in the classic doughnut and coffee mug example). In more technical language, two such spaces are "homotopy equivalent." We now make this intuition more rigorous.

Definitions 3.1. A *homotopy* between two continuous maps $f, g : X \to Y$ is a continuous map $h : X \times [0,1] \to Y$ (in the product topology), such that for all $x \in X, h(x,0) = f(x)$ and $h(x,1) = g(x)$. Often, $h(x,t)$ is written as $h_t(x)$. When such a map exists, we call f and g *homotopic* and write $f \simeq g$.

Definition 3.2. Two topological spaces X, Y are called *homotopy equivalent* if there exists continuous maps $f : X \to Y$ and $g : Y \to X$ which are "inverses up to homotopy," that is, $f \circ g \simeq id_Y$ and $g \circ f \simeq id_X$ (where id_X is the identity map on X and similarly for Y). In this case, we call the maps f and g *homotopy equivalences* and write $X \simeq Y$.

Definition 3.3. A topological space X is called *contractible* if it is homotopy equivalent to a point.

Definition 3.4. A *loop* in a space X with basepoint x is a continuous map $\gamma : [0,1] \to X$ with $\gamma(0) = \gamma(1) = x$. Alternatively, a loop is a map $\gamma : S^1 \to X$ with the x in the image of γ.

In this paper, we care about a certain way to associate groups to spaces, because it will allow us to translate questions about groups to ones about spaces. We call this the fundamental group of a space.

Definition 3.5. The *fundamental group* of a topological space X at the point x is the set of homotopy equivalences of loops with basepoint x, with the group operation "do one, then do the other." That is, if γ, γ' are representatives of some homotopy classes then $[\gamma][\gamma'] = [\gamma'']$, where γ'' is the loop defined as

$$\gamma''(t) = \begin{cases} \gamma(2t) & \text{if } 0 \leq t \leq 1/2 \\ \gamma'(2t-1) & \text{if } 1/2 \leq t \leq 1 \end{cases}$$

One can verify that all this is well defined and forms a group with the identity the constant loop at x and inverses given by "going around backwards."

Often, the fundamental group is actually independent of the choice of basepoint.

Definitions 3.6. A *path* between two points $x, y \in X$ is a continuous map $\gamma : [0,1] \to X$ such that $\gamma(0) = x$ and $\gamma(1) = y$. A space X is called path connected if for any two points $x, y \in X$, there is a path between x and y.

Proposition 3.7. *If X is path connected, and $x, y \in X$, then $\pi_1(X, x) \cong \pi_1(X, y)$.*

If this is the case, we drop the basepoint from the notation, and just write $\pi_1(X)$.

Proposition 3.8. *If X is homotopy equivalent to Y, with basepoints x, y respectively, then $\pi_1(X, x) \cong \pi_1(Y, y)$*

Definition 3.9. A space X is called *simply connected* if it is path connected and its fundamental group is trivial.

We now give a very important example of a nontrivial fundamental group.

Proposition 3.10. $\pi_1 S^1 \cong \mathbb{Z}$

Fundamental groups of spaces are fairly easy to explicitly calculate due to the following theorem:

Theorem 3.11 (Van Kampen)**.** *If X is path connected and if $X = A \cup B$ (where A and B are open subspaces of X) and if $A \cap B$ is simply connected (that is $\pi_1(A \cap B)$ is trivial), then $\pi_1(X) = \pi_1(A) * \pi_1(B)$. More generally, if $X = A \cup B$ and A, B are open in X, then and $C = \pi_1(A \cap B)$ then $\pi_1(X) = \pi_1(A) *_C \pi_1(B)$.*

Definition 3.12. The *wedge sum* of two spaces X, Y with distinguish basepoints x, y is the disjoint union of X and Y, with the points x and y identified.

For example, the wedge of two circles with any point the basepoint is a "figure-8" space.

In particular, by Theorem 3.11 and Proposition 3.10, the fundamental group of a wedge of α (not necessarily finite!) circles is a free group of rank α.

We now describe, without being very formal, a very nice class of spaces to work with, CW complexes.

A CW complex is a space that is built up out of disks of various dimensions. We start with some collection of points and call these *0-cells.* Then we attach some number of intervals to these points (at the endpoints of the intervals), and call these *1-cells.* Now we might have loops (which are just images of S^1 in this space), so we can attach 2-dimensional disks to these loops along the boundaries of the disks. We can keep doing this inductively, attaching n-balls to $(n-1)$-spheres, calling them n-cells. What we get from this process is called a CW complex. (The letters CW refer to the somewhat technical topology we put on this space. It doesn't really matter if we're treating it as blackbox.) If a space only has cells of dimension n or less, we call it an *n-dimensional CW complex.* Finally, we sometimes want to consider only cells of a given dimension or less in a given CW complex X. We call this space then *n-skeleton* of X

CW complexes have some nice properties.

Proposition 3.13. *For CW complexes, path-connected is equivalent to connected. Furthermore, a CW complex is path-connected if and only if its 1-skeleton is connected*

Proposition 3.14. *If X is a CW complex, and A is a contractible subcomplex of X, then X/A is homotopy equivalent to X*

A natural question is which groups occur as fundamental groups of spaces. The answer is, even for a restricted set of spaces, all of them, by the following proposition.

Proposition 3.15. *For any group G, there is a 2-dimensional CW complex X_G with $\pi_1(X_G) \cong G$*

We will also need the notion of a covering space, and how it interacts with the fundamental group.

Definition 3.16. A *covering space* of X is a space $\tilde{X}$ and a map $p : \tilde{X} \to X$ with the following properties: There is an open cover $\{U_\alpha\}$ of X such that for each $\alpha, p^{-1}(U_\alpha)$ is a disjoint union of open sets in $\tilde{X}$ which are homeomorphic to U_α, with p the homeomorphism.

Another reason CW complexes are nice is that a covering space of one is also a CW complex.

Proposition 3.17. *Every covering space of an n-dimensional CW complex is an n-dimensional CW complex.*

The main reason why we care about covering spaces, at least for the purposes of this paper, is that they correspond exactly to subgroups of the fundamental group. This allows gives a dictionary between group theory and topology that we will make good use of.

Theorem 3.18 (Galois Correspondence)**.** *If X is nicely behaved, then the subgroups of $\pi_1(X)$ correspond exactly to the connected covering spaces of X. More precisely, if H is a subgroup of $\pi_1(X)$ then there is exactly one covering space (up to homeomorphism) $\tilde{X}$ of X with $\pi_1(\tilde{X}) \cong H$. Also, if $\tilde{X}$ is a connected covering space of X, then $\pi_1(\tilde{X}$ is isomorphic to a subgroup of $\pi_1(X)$.*

Here nicely behaved means path connected and a couple of other technical assumptions. In particular, CW complexes are nicely behaved.

Finally, we want to say some things about graphs (not quite a combinatorialist's graphs, but those are a special case).

Definition 3.19. A *graph* is a 1-dimensional CW complex.

Definition 3.20. A *tree* is a contractible graph.

Proposition 3.21. *If X is a connected graph, then $\pi_1(X)$ is a free group.*

Proof. We first define a *maximal tree* in a graph X to be a tree Y contained in X such that Y is not contained in any other tree in X. The fact that every graph has a maximal tree follows easily from Zorn's Lemma.[1] Then by Propositions 3.14 and 3.8, collapsing a maximal tree to a point gives a space which has the same fundamental group as X, but is a wedge of some number of circles (one for each edge of the graph not in the maximal tree). Then by Proposition 3.21, X has a free group for its fundamental group. □

Proposition 3.22. *Every subgroup of a free group is free.*

Proof. Let G be a free group, and let H be a subgroup of G. Then let X be a graph with fundamental group G (for example, a wedge of circles). By Theorem 3.18, there is a connected covering space $\tilde{X}$ of X with fundamental group H. But by Proposition 3.17, $\tilde{X}$ is a graph. But then by Proposition 3.21, H is a free group. □

4. Main Results

Now we state the primary theorem whose proof is the goal of this paper:

Theorem 4.1. *Any finitely generated group G is the free product of finitely many indecomposable groups $G_1 * G_2 * \ldots * G_n$ and this decomposition is unique, that is, if we also have $G = H_1 * H_2 * \ldots * H_m$ then $n = m$ and each G_i is isomorphic to precisely one H_j.*

The proof of this will involve some theorems about free products that rely on some basic facts about covering spaces that we mentioned before.

First, we will go about proving uniqueness. The following theorem will be helpful:

Theorem 4.2 (Kuroš)**.** *Let $G = G_1 * G_2$ and let H be a subgroup of G. Then $H = F * H_1 * \ldots * H_k$, where F is a free group and $H_1, \ldots, H_k$ are subgroups of conjugates of G_1 and G_2, within $G_1 * G_2$.*

Proof. First, we make a CW complex X with $\pi_1(X) = G$ using a "dumbbell construction". Let X_1', X_2' be CW complexes with fudamental groups G_1, G_2, respectively. Now connect the basepoints of X_1' and X_2' with an interval E, and declare its midpoint, v to be the basepoint of the new space, which we now name X. We also define X_1 to be the component of $X \setminus \{v\}$ which contains X_1', together with the $\{v\}$(and again, similarly for X_2). Then X_1 deformation retracts onto X_1' (send the extra interval to the basepoint in X_1'). Thus, $\pi_1(X_1) \cong G_1$. (And the same for X_2. Now to find the fundamental group of X, we use Van Kampen's Theorem (Theorem 3.11). We have that $X = X_1 \cup X_2$ and also $X_1 \cap X_2 = v$, which has trivial fundamental group, so then $\pi_1(X) = \pi_1(X_1) * \pi_1(X_2) = G_1 * G_2$.

Now let H be a subgroup of G. By the correspondence between connected covering spaces of X and subgroups of $\pi_1(X)$ (Theorem 3.18) there is a connected covering space $\tilde{X}$ of X and with projection map $p : \tilde{X} \to X$. Since $\tilde{X}$ is a covering space of X, $p^{-1}(X_1)$ is a covering space of X_1 (p is still a local homeomorphism) and similarly for X_2. However, $p^{-1}(X_1)$ may not be connected, but each component of $p^{-1}(X_1)$ will be a connected covering space of X_1, and thus will correspond to a subgroup of $\pi_1(X_1)$. Also, $p^{-1}(E)$ is a union of connected covering spaces of E.

[1]And is in fact, equivalent to the Axiom of Choice.

Since E is contractible, there is only one connected covering space of E, that is, E itself (again, by Theorem 3.18). Thus, $p^{-1}(E)$ is a union of copies of E. These copies of E attach to the connected covering spaces of X_1 and X_2 at elements of the preimages of the basepoints. Thus, $\tilde{X}$ is a union of connected covering spaces of X_1 and X_2 connected to each other with intervals.

We want to proceed here by saying that $\pi_1(X)$ is the free product of the fundamental groups of these smaller covering spaces, as well as an extra free group. This would certainly be true if X were a graph with the small covering spaces glued onto the vertices. However, this is not quite the case. The intervals connecting to one of the "components" do not all attach at the same point (if they did, no neighborhood of this point would be homeomorphic to any neighborhood of v). This is not too difficult to deal with though. The idea is to replace the smaller covering spaces with homotopy equivalent spaces, where the connecting intervals all attach at the same point.

Choose a component of $p^{-1}(X_1)$ and call it $\tilde{X}_1$ First, recall that a covering space of a CW complex is another CW complex (Proposition 3.17), and that in this case, connected and path-connected are equivalent. Additionally, if a CW complex is path-connected, then its 1-skeleton is. Thus, we may chose one 0-cell $y \in \tilde{X}_1$ where an interval connects and then connect all other such 0-cells to it via paths in the 1-skeleton. If any of these make a loop around a non-contractible portion of the 1-skeleton (e.g. a circle), we may replace that segment of the path with a constant at its basepoint. Thus, y is connected to all the "connection points" via contractible paths in the 1-skeleton. The union of all of these form a contractible subcomplex A of $\tilde{X}_1$, and thus, $\tilde{X}_1/A$ is homotopy equivalent to $\tilde{X}_1$. (By Proposition 3.14) Doing this procedure to each connected covering space of X_1 or X_2 in X, we get a space which is actually a graph with extra spaces glued on at the vertices. Now by choosing a maximal tree in this graph and collapsing it to a point, we get a space which is the wedge sum of some covering spaces of X_1 and X_2 as well as some more circles, which came from the edges of the graph which were not in the maximal tree. Thus, $\pi_1(X)$ isomorphic to the free product of subgroups of G_1 and G_2 as well as an additional free group.

What we actually want, though, is for these isomorphisms of subgroups to be given by conjugation. To show this, let $\tilde{v}$ be the basepoint of $\tilde{X}$ and remember that $H = p_*(\pi_1(\tilde{X}, \tilde{v}))$. Also, let C be some component of $p^{-1}(X_1)$, and let γ be a path from some element of $p^{-1}(v)$ which is in C to $\tilde{v}$. The $p \circ \gamma$ is a loop in X_1 based at v, so $p_*(\pi_1(C, \tilde{v}))$ is conjugate to some subgroup of G_1 (via some element of the fundamental group which is represented by $p \circ \gamma$. □

As a consequence of this theorem, we get the following corollary, and another proof of Proposition 3.22:

Corollary 4.3. *Let H be indecomposable (that is, H is not a free product of any two nontrivial groups), and suppose $H \ncong \mathbb{Z}$. If $H \leq G_1 * G_2$, then H is a subgroup of a conjugate of G_1 or G_2*

Proof. Since $H \leq G_1 * G_2$, H is isomorphic to a free product of subgroups of G_1 and G_2 as well as a free group F, by the theorem. But then since H is indecomposable, F must be trivial or $\mathbb{Z}$, otherwise , we could express F as a free product of another nontrivial free group and $\mathbb{Z}$, and then H would be the free product of two nontrivial groups. Now if the $F \cong \mathbb{Z}$, all of the subgroups of G_1 and G_2 composing H must be trivial, or again, H would not be indecomposable. But then $H \cong \mathbb{Z}$, which we assumed it wasn't. Thus, F is trivial. Then H is the free product of subgroups of conjugates of G_1 and G_2, but all but one of these must be trivial, since H is indecomposable. Thus, H is a subgroup of a conjugate of G_1 or G_2. □

Alternate Proof of Proposition 3.22

Proof. Let F be a free group and let $H \leq F$. Then $F \cong \mathbb{Z} * \mathbb{Z} * \mathbb{Z} * \ldots$. Then by the Theorem 4.2, H is a free product of conjugates of subgroups of $\mathbb{Z}$ and a free group. But any nontrivial subgroup of $\mathbb{Z}$ is isomorphic to $\mathbb{Z}$, and conjugates of a subgroup are still isomorphic to that subgroup. Thus, H is the free product of two free groups, which is free, since free groups are the free product of several copies of Z. □

Now we just need to state one more lemma, and we can prove uniqueness of the free product decomposition.

Lemma 4.4. *Let $G = G_1 * G_2$ (assuming G_1, G_2 nontrivial) and $w \in G$. Then if $w^{-1}G_1w \cap G_i$ $(i \in \{1, 2\})$ is nontrivial, then $i = 1, w \in G_1$, and thus, $w^{-1}G_1w = G_1$.*

Proof. Let g be a nontrivial element of G_1 with $w^{-1}gw \in G_i$. Also, we can let $w = \alpha w'$, where $\alpha \in G_1$ and w' is a reduced word in starting in G_2. Then $w^{-1}gw = w'^{-1}\alpha^{-1}g\alpha w' = w'^{-1}g'w'$, where g' is a nontrivial element of G_1. Thus, $w'^{-1}g'w'$ is a reduced word, since w is a reduced word and the last letter of w'^-1 and the first letter of w are both in G_2. But, $w'^{-1}g'w'$ is in G_i, and thus has length 1. Then w' is trivial, and thus, $w = \alpha \in G_1$. So, $w^{-1}G_1w = G_1$. But then $w^{-1}G_1w \cap G_i = G_1 \cap G_i$ is nontrivial, but this can be the case only if $i = 1$, since G is the free product of G_1 and G_2. □

We may now prove uniqueness.

Theorem 4.5. *If $G = G_1 * G_2 * \ldots * G_n$ and $G = H_1 * H_2 * \ldots * H_m$, where each G_i and H_j is indecomposable and nontrivial, then $m = n$ and (possibly by permuting the $H_i's$), $G_i \cong H_i$ for all i.*

Proof. First suppose that, for all $i, G_i \cong \mathbb{Z}$. Then G is free, and then since $H_i \leq G$ for all i, H_i is free by Corollary 3.22 and since H_i is indecomposable and nontrivial, $H \cong \mathbb{Z}$. Then by abelianizing $G_1 * \ldots * G_n$ and $H_1 * \ldots * H_m$, we get $\mathbb{Z}^n \cong \mathbb{Z}^m$, which is true only if $m = n$.

If not all G_i are isomorphic to $\mathbb{Z}$, we may reorder the G_i's so that $G_1, \ldots, G_k \ncong \mathbb{Z}$ and $G_{k+1}, \ldots, G_n \cong \mathbb{Z}$. Now $G_1 \leq H_1 * \ldots * H_m$, and it is also indecomposable, so by Corollary 4.3 G is a subgroup of a conjugate of some H_i, but we may reorder and make it H_1. Then for some $u \in G, u^{-1}G_1u \subseteq H_1$. So then $H_1 \ncong \mathbb{Z}$ and we may apply Corollary 4.3 again, this time to $H_1 \leq G_1 * \ldots * G_n$, so for some $v \in G$ and some i, $v^{-1}H_1v \subseteq G_i$. But then $(uv)^{-1}G_1uv \subseteq G_i$, and by Lemma 4.4, $uv \in G_1$ and $i = 1$. Now we have

$$G_1 = (uv)^{-1}G_1uv \subseteq v^{-1}H_1v \subseteq G_1.$$

Thus H_1 is conjugate to G_1 in G, and so $G_1 \cong H_1$.

From this same argument we get that $G_i \cong H_1$ for $1 \leq i \leq k$. For any H_j, there is only one G_i conjugate to it, since otherwise two different G_i's would be conjugate, and that can't happen by the Lemma. So then $G_1 * \ldots * G_k \cong H_1 * \ldots * H_k$.

Now let G' be the normal subgroup generated by G containing $G_1 * \ldots * G_k$ and let H' be the smallest normal subgroup of G containing $H_1 * \ldots * H_k$. Since G_i is conjugate to H_i for $1 \leq i \leq k$, $G' = H'$, so

$$G_{k+1} * \ldots * G_n \cong G/G' = G/H' \cong H_{k+1} * \ldots * H_m.$$

But $G_{k+1} * \ldots * G_n$ is free, since for $i > k, G_i \cong \mathbb{Z}$. But then $H_{k+1} * \ldots * H_m$ is free, since its isomorphic to $G_{k+1} * \ldots * G_n$, so for each $i > k, H_i \cong \mathbb{Z}$ and by abelianizing, we see that $m - k = n - k$. Thus, $m = n$, and $H_i \cong G_1$ for $1 \leq i \leq n$. □

Remark 4.6. Notice that this proof works for any group that can be written as the free product of indecomposable groups, not just finitely generated ones. The hypothesis that G be finitely generated is only necessary to guarantee the existence of such a decomposition. There are groups which do not admit such decompositions, and we will now give an example of one.

First, remember that for two elements a, b of a group G, the notation $[a, b]$ means the $aba^{-1}b^{-1}$, which is called the commutator of a and b, and is the identity iff $ab = ba$.

Example 4.7. The group $G = \langle a_0, a_1, a_2, \ldots, b_1, b_2, \ldots | [a_n, b_n]a_{n-1}^{-1}, n \geq 1 \rangle$ cannot be written as *any* free product of indecomposable groups.

Proof. First, note that $G = \langle b_1 \rangle * \langle a_1, a_2, \ldots, b_2, b_3, \ldots | [a_n, b_n]a_{n-1}^{-1}, n \geq 2 \rangle \cong \mathbb{Z} * G$. (We dropped the a_0 since it is identified with $[a_1, b_1]$.) Thus, we may write G as a free product of any finite number of groups. Because of this, G cannot be the free product of any finite number of indecomposable groups, because if we could write it as the free product of n indecomposables, our uniqeness result would apply and it could not be a free product of any $n + 1$ groups, which is a contradiction. So G is not the free product of any finite number of indecomposable groups.

To rule out the infinite case, we need a lemma, which will be stated without proof.

Lemma 4.8. *If $G = A * B$, A, B nontrivial, and if $g = [g_1, g_2]$ is a nontrivial element of A, then $g_1, g_2 \in A$.*

Also, note that $G = \langle a_1, b_1 \rangle *_{C_1} \langle a_2, b_2 \rangle *_{C_2} \ldots$, where $C_i \cong \mathbb{Z}$ for all i, and where the inclusion maps are given by sending a generator of C_i to a_i (for $C_i \hookrightarrow \langle a_i, b_i \rangle$) and the same generator to $[a_{i+1}, b_{i+1}$ (for $C_i \hookrightarrow \langle a_{i+1}, b_i + 1 \rangle$).

Now suppose $G = G_1 * G_2 * \ldots$, where each G_i is indecomposable and nontrivial. (G is at most a countable free product, since it is countable generated) Consider the element $a_0 \in G$. Then for some n, a_0 must lie in $G_1 * \ldots * G_n$. (This is actually true for any element of G, since every element is a finite word.) Now set $A = G_1 * \ldots * G_n$ and $B = G_{n+1} * G_{n+2} * \ldots$. Then $G = A * B$, where A and B are nontrivial, and $a_0 \in A$. The decomposition of G into an infinite amalgamated free product given above shows that each a_i is nontrivial in G. Then by Lemma 4.8, since $a_0 \in A, a_1, b_1 \in A$. But then for all i, a_i and b_i are in A, so $G \subseteq A$, and B is trivial, which is a contradiction. Thus, G cannot be written as the free product of any number of indecomposable groups. □

Finally, we go about proving the existence of a free product decomposition for finitely generated groups. This will follow as a corollary to to the following theorem.

Theorem 4.9 (Gruško)**.** *Let F be a free group, $G = G_1 * G_2$ and let $\varphi : F \to G$ be a surjective homomorphism. Then there are $F_1, F_2 \leq F$ with $F = F_1 * F_2$ and $\varphi(f_i) = G_i$.*

Remark 4.10. Note that this theorem is really talking about the generators of G. It says that if G is the free product of two groups, and can be generated be n elements, then there is a set of n generators of G, each of which is an element of G_1 or G_2, and not a longer word.

Proof. Again we use the "dumbell" construction. Let X_1', X_2' be CW complexes with $\pi_1(X_i') \cong G_i$, and connect their basepoints with and interval E and call this space X. Declare the basepoint v to be the midpoint of E. Finally, let X_1 be the closure of the component of $X \setminus \{v\}$ which contains X_1' (so that X_1 also contains v) and define X_2 similarly. Also, X_1 and has fundamental group G_1, since X_1 deformation retracts to X_1', and similarly for X_2 and G_2.

Now we make a quick definition.

Definition 4.11. Let K be an arbitrary pointed topological space, and let $f : K \to X$ be a basepoint-preserving continuous map. We say that f *represents* φ if there is an isomorphism ι between $\pi_1(K)$ and f with the diagram below commuting

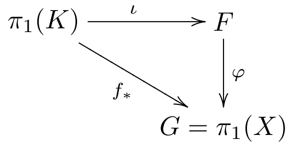

Actually, for the rest of the proof, we only care about the case where K is a 2-dimensional CW complex and f is a cellular map.

Maps representing φ exist, since we can take $K = \vee_{i=1}^n S^1$ (n being the rank of F). In this case f_* is basically equal to φ since $\pi_1(K) \cong F$. However, this may not give the "right" map, so we modify it until it suits our purposes. "Suits our purposes" in this case means that we want $f^{-1}(v)$ to be a tree, because then the theorem follows almost immediately.

So, let K_0 be the wedge of n circles (with n the rank of F), and pick a cellular map f_0 representing φ so that $f^{-1}(v)$is a finite number of 0-cells in K_0. Particularly, $f_0^{-1}(v)$ is a forest (a disjoint union of trees). If it is connected, it's just a single point, but in particular, its a tree. If it's not a tree, we want to find another space K and another map f which represents φ which is a tree. The following lemma allows us to do that.

Lemma 4.12. *Let K be a pointed CW complex and $f : K \to X$ be a map representing φ such that $f^{-1}(v)$ is a forest with $n \geq 2$ components. Then there exists a pointed CW complex K' and a map $f' : K' \to X$ representing φ and $f'^{-1}(v)$ a forest with $n - 1$ components.*

Proof. Let ℓ be a path in K with endpoints in two distinct component of $f^{-1}(v)$ (that is, ℓ connects two components of $f^{-1}(v)$). Now attach to K an interval e which connects the endpoints of ℓ, and also a 2-cell D to the loop defined by ℓ and e. Call this new CW complex K'. Now we are done if we can extend the map f

to a map $f' : K' \to X$ with the properties $f(e) = v$ and $f'^{-1}(v) \cap B^{\circ} = \emptyset$ (that is, nowhere in the interior of B maps to v, or the image of the interior of B is contained entirely in one of $X_1 or X_2$). This is because $f'^{-1}(v) = f^{-1}(v) \cup e$, which is a tree with one fewer component than $f^{-1}(v)$, and f' represents φ, since K' deformation retracts to K (collapse all of B to the path ℓ), so $\pi_1(K') \cong \pi_1(K)$.

This can be done if the loop formed by $f \circ \ell$ is entirely contained in one of X_1 or X_2, and if this loop is contractible. ($f \circ \ell$ is a loop based at v, since ℓ connects two points in $f^{-1}(v)$. If we have both of these, then $f \circ \ell$ is homotopic to v and so there is a continuous map defined by this homotopy from a closed disc (in this case, B) to X, and we define f' to be f where f is defined, $f'(e) = v$, and define f' by this homotopy on B. Now we just need to guarantee that such an ℓ exists.

Let A, B be components of $f^{-1}(v)$ and suppose L is a path in K joining A to B. Now φ is onto, so f_* is onto, and then there is a loop γ in K based at $L(0)$ so that $f \circ \gamma$ is homotopic to $f \circ L$. Now let $\ell = \gamma^{-1} \circ L$. Then $f \circ \ell$ is a contractible loop in X (since we defined γ so that $f \circ \gamma$ was homotopic to a point). So now we just need ℓ to "not cross through v."

We may assume that ℓ is a cellular map, since we can subdivide the domain of ℓ (which is an interval) and choose L so that this is the case. Then we may express ℓ as a union of smaller paths $\ell_1 \dots \ell_k$ so that the endpoints of each ℓ_i are in $f^{-1}(v)$ and so that each $f \circ \ell_i$ lies entirely in X_1 or X_2 (then this is a loop in X_1 or X_2, alternating between the two. Set $g_i = [f \circ \ell_i] \in \pi_1(X, v)$. Then if for some i, ℓ_i has that g_i is trivial and the endpoints of ℓ_i are in the same component of $f^{-1}(v)$, we can remove it, and replace it by a path connecting the same endpoints, but lying entirely in their component.

Now we have that $1 = g_1 g_2 \dots g_k$ in $\pi_1(X)$, since $\ell = \ell_1 \dots \ell_k$ and $f \circ \ell$ is null homotopic. Also, $\pi_1(X) = G_1 * G_2$, and the g_i's alternate between G_1 and G_2, since the $f \circ \ell_i$'s alternate between X_1 and X_2. Then some g_i must be trivial. Then for the same i, ℓ_i connects two distinct components of $f^{-1}(v)$, and $f \circ \ell_i$ is contractible, so we ℓ_i has all the necessary properties and we are done. □

Now that we have this result, it is easy to prove the theorem.

Take K and $f : K \to X$ be such that $f^{-1}(v)$ is a tree. Then define $Y_i = f^{-1}(X_i), i \in 1, 2$ and $F_i = \pi_1(Y_i)$. Then clearly $Y_1 \cup Y_2 = K$ since $X_1 \cup X_2 = X$. Also $Y_1 \cap Y_2 = f^{-1}(v)$ is a tree, so in particular, its simply connected. Then we apply Van Kampen's Theorem and see that $\pi_1(K) = F_1 * f_2$. By definition $f(Y_i) \subseteq X$, so $f_*(F_i) \subseteq G_i$. Now since G_i injects into G, we have that $f_*(F_i) = G_i$, since φ is surjective. □

Now if we define $\mu(H)$ to be the smallest number of generators of a group H, we now have that

Corollary 4.13. *If* $G = G_1 * G_2, \mu(G) = \mu(G_1) + \mu(G_2)$

Proof. By Theorem 4.9 $\mu(G) \geq \mu(G_1) + \mu(G_2)$. But also, $\mu(G) \leq \mu(G_1) + \mu(G_2)$, since the generators of G_1 together with those of G_2 form a generating set for G, so we actually have $\mu(G) = \mu(G_1) + \mu(G_2)$ □

Only the trivial group has $\mu(H) = 0$, so then, $\mu(G_i) < \mu(G)$ if G_1 and G_2 are nontrivial and G is finitely generated. Now finally we have

Corollary 4.14. *If G is finitely generated, then $G = G_1 * G_2 * \ldots * G_n$ for some n, and each G_i is indecomposable.*

Proof. If G itself is indecomposable, we are done. If not we write $G = G_1 * G_2$, and each of G_1 and G_2 have fewer generators than G. Now we just repeat this process on G_1 and G_2 until we have all the factors indecomposable. This is guaranteed to terminate, because each step reduces the number of generators and G is finitely generated. □

References

[1] Allen Hatcher. Algebraic Topology. Cambridge University Press. 2001.
[2] Peter Scott and Terry Wall. Topological Methods in Group Theory. London Mathematical Society Lecture Notes 36. 1979 pp.137-203