1. Group presentations.

A group presentation gives a means of specifying a group up to isomorphism. It is the basis of the now "classical" combinatorial theory of groups. In Section 2, we will give a more geometrical intepretation of these constructions.

1.1. Notation.

Throughout the course, we will use the following fairly standard notation relating to groups.

$G \subseteq \Gamma$: G is a subset of Γ.

$G \leq \Gamma$: G is a subgroup of Γ.

$G \triangleleft \Gamma$: G is a normal subgroup of Γ.

$G \cong \Gamma$: G is isomorphic to Γ.

$1 \in \Gamma$ is the identity element of Γ.

$[\Gamma : G]$ is the index of G in Γ.

We write $|A|$ for the cardinality of a set A. In other words, $|A| = |B|$ means that there is a bijection between A and B. (This should not be confused with the fairly standard notation for "realisations" of complexes, used briefly in Section 2.)

We use **N**, **Z**, **Q**, **R**, **C** respectively for the natural numbers (including 0), the intergers, and the rational, real and complex numbers.

We shall generally view $\mathbf{Z}^n$ and $\mathbf{R}^n$ from different perspectives. We shall normally think of $\mathbf{Z}^n$ as group under addition, and $\mathbf{R}^n$ as a metric space with the euclidean norm.

1.2. Generating sets.

Let Γ be a group and $A \subseteq \Gamma$. Let $\langle A \rangle$ be the intersection of all subgroups of Γ containing the set A. Thus, $\langle A \rangle$ is the unique smallest subgroup of Γ containing the set A. In other words, it is characterised by the following three properties:

(G1) $A \subseteq \langle A \rangle$,

(G2) $\langle A \rangle \leq \Gamma$, and

(G3) if $G \leq \Gamma$ and $A \subseteq G$, then $\langle A \rangle \subseteq G$.

We can give the following explicit description of $\langle A \rangle$:

$$\langle A \rangle = \{a_1^{\epsilon_1} a_2^{\epsilon_2} \cdots a_n^{\epsilon_n} \mid n \in \mathbf{N}, a_i \in A, \epsilon_i = \pm 1\}.$$

(If $n = 0$ this expression is interpreted to be the identity element, 1.) To see this, we verify properties (G1), (G2) and (G3) for the right hand side: note that the set of such elements forms a group containing A, and that any subgroup containing A must also contain every element of this form.

Definition : Γ is *generated* by a subset A if $\Gamma = \langle A \rangle$. In this case, A is called a *generating set* for Γ.

We say that Γ is *finitely generated* if it has a finite generating set.

In other words, $\Gamma = \langle a_1, \ldots, a_n \rangle$ for some $a_1, \ldots, a_n \in \Gamma$. (Note, $\langle a_1, \ldots, a_n \rangle$ is an abbreviation for $\langle \{a_1, \ldots, a_n\} \rangle$.) We will frequently abbreviated "finitely generated" to "f.g."

By definition, Γ is *cyclic* if $\Gamma = \langle a \rangle$ for some $a \in \Gamma$. It is well known that a cyclic group is isomorphic to either $\mathbf{Z}$ or $\mathbf{Z}_n$ for some $n \in \mathbf{N}$. To be consistent, we shall normally use multiplicative notation for such groups. Thus, the infinite cyclic group will be written as $\{a^n | n \in \mathbf{Z}\}$. This is also called the *free abelian group* of *rank* 1.

Similarly, $\mathbf{Z}^2 = \mathbf{Z} \oplus \mathbf{Z}$ is generated by two elements $a = (1, 0)$ and $b = (0, 1)$. It is called the *free abelian group* of *rank* 2. We again use multiplicative notation, and write it as $\{a^m b^n \mid m, n \in \mathbf{Z}\}$. Note that $ab = ba$. This is an example of a "relation" between generators.

More generally, we refer to the (isomorphism class of) the group $\mathbf{Z}^n$ as the *free abelian group* of *rank* n. It is generated by the n elements, $e_1, \ldots, e_n$, of the form $(0, \ldots, 0, 1, 0, \ldots 0)$. The free abelian group of rank 0 is the trivial group.

Exercise: If $\mathbf{Z}^m \cong \mathbf{Z}^n$, then $m = n$.

Warning: The term "free abelian" should be thought of as one word. A free abelian groups is not "free" in the sense shortly to be defined (unless it is trivial or cyclic).

An important observation is that generating sets are not unique. For example:
$\mathbf{Z} = \langle a^2, a^3 \rangle$ (Note that $a = (a^3)(a^2)^{-1}$.)
$\mathbf{Z}^2 = \langle ab, a^2b^3 \rangle$ (Note that $a = (ab)^3(a^2b^3)^{-1}$ and $b = (a^2b^3)(ab)^{-2}$.)

It is sometimes convenient to use "symmetric" generating sets in the following sense.

Given $A \subseteq \Gamma$, write $A^{-1} = \{a^{-1} \mid a \in A\}$.

Definition : A is *symmetric* if $A = A^{-1}$.

Note that for any set, A, $A \cup A^{-1}$ is symmetric. Thus, a finitely generated group always has a finite symmetric generating set.

Note that in this case, each element of Γ can be written in the form $a_1 a_2 \cdots a_n$, where $a_i \in A$. Such an expression is called a "word" of "length" n in the elements of A. (We give a more formal definition shortly). A word of length 0 represents the identity, 1.

There are very many naturally arising finitely generated groups. A few examples include:
(1) Any finite group (just take $\Gamma = A$),
(2) $\mathbf{Z}^n$ for $n \in \mathbf{N}$,
(3) Finitely generated free groups (defined later in this section).
(4) Many matrix groups, for example $GL(n, \mathbf{Z})$, $SL(n, \mathbf{Z})$, $PSL(n, \mathbf{Z})$, etc.
(5) In particular, the (discrete) Heisenberg group has some interesting geometrical properties. It can be defined as:

$$\left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \mid x, y, z \in \mathbf{Z} \right\}.$$

(6) The fundamental groups of compact manifolds (see Section 4).
(7) The fundamental groups of finite simplicial complexes (see Section 4).
(8) Hyperbolic groups in the sense of Gromov (see Section 6).

(9) Mapping class groups, braid groups, Artin groups. etc.
(10) Many others.

On the other hand, many groups are not f.g. For example, **Q** (exercise), **R**, $GL(n, \mathbf{R})$, etc.

Exercises: (1) If $N \triangleleft \Gamma$ and Γ is f.g., then Γ/N is f.g.
(2) If $N \triangleleft \Gamma$, N is f.g. and Γ/N is f.g., then Γ is f.g.
(3) If $G \leq \Gamma$ and $[\Gamma : G] < \infty$ ("finite index") then Γ is f.g. if and only if G is f.g. (The "only if" part might be a bit tricky. It will follow from a more general discussion later, see section 3.)

Remark: There are examples of f.g. groups Γ which contain subgroups which are not finitely generated. (Indeed "most" f.g. groups will have such subgroups.) We will see examples later (Section 4).

We remark however, that any subgroup of a f.g. abelian group will be finitely generated.

1.3. Free groups.

Note that in any group, we will always have relations obtained by cancelling elements with their inverses, that is $aa^{-1} = 1$ and $a^{-1}a = 1$. The idea behind a free group is that there should be no other relations. Informally, a group F is "freely generated" by a subset $S \subseteq F$ if the only relations arise out of cancelling pairs aa^{-1} and $a^{-1}a$ for $a \in A$. Of course the meaning of "arising out of" has not yet been defined. Here is a more formal definition which, as we will eventually see, captures this idea.

Definition : A group F is *freely generated* by a subset $S \subseteq F$ if, for any group Γ and any map $\phi : S \longrightarrow \Gamma$, there is a unique homomorphism $\hat{\phi} : F \longrightarrow \Gamma$ extending ϕ, i.e. $\hat{\phi}(x) = \phi(x)$ for all $x \in S$.

Note that we have not said that S is finite. In most examples of interest to us it will be, though we don't need to assume that for the moment.

Lemma 1.1 : *If F is freely generated by S, then it is generated by S (i.e. $F = \langle S \rangle$).*

Proof : Let $\Gamma = \langle S \rangle$. The inclusion of S into Γ extends to a (unique) homomorphism, $\theta : F \longrightarrow \Gamma$. If we compose this with the inclusion of Γ into F, we get a homomorphism from F to F, also denoted θ. But this must be the identity map on F, since both θ and the identity map are homomorphisms extending the inclusion of S into F, and such an extention is, by hypothesis, unique. It now follows that $\Gamma = F$ as required. $\diamondsuit$

Lemma 1.2 : *Suppose that F is freely generated by $S \subseteq F$, that F' is freely generated by $S' \subseteq F'$, and that $|S| = |S'|$. Then $F \cong F'$.*

Proof : The statement that $|S| = |S'|$ means that there is a bijection, ϕ, between S and S'. Let $\theta = \phi^{-1} : S' \longrightarrow S$ be the inverse bijection. These extend to homomorphisms $\hat{\phi} : F \longrightarrow F'$ and $\hat{\theta} : F' \longrightarrow F$. As with Lemma 1.1, we see that the composition $\hat{\theta} \circ \hat{\phi} : F \longrightarrow F$ must be the identity map of F. Thus, both ϕ and θ must be isomorphisms. $\diamondsuit$

If $|S| = n < \infty$, we denote F by F_n.

Definition : The group F_n is the *free group* of *rank* n.

By Lemma 1.2, it is well defined up to isomorphism.

Fact : If $F_m \cong F_n$, then $m = n$. (We will sketch a proof of this at the end of this section.)

Exercise: A free group is torsion-free. (i.e. if $x^n = 1$ then $x = 1$.)

We have given a characterisation of free groups, and shown a uniqueness property. However, we have not yet said anything about their existence. For this, we need the following construction.

1.4. Construction of free groups.

First, we introduce some terminology. Let A be any set, which we shall call our *alphabet*.

Definition : A *word* is a finite sequence of elements of A.

More formally, it is a map $\{1, \ldots, n\} \longrightarrow A$. We denote the image of i by $a_i \in A$, and write the word as $a_1 a_2 \ldots a_n$. We refer to n as its *length*. The a_i are the *letters* in this word. If $n = 0$, we refer to this as the *empty word*. Note that we can *concatenate* a word $a_1 \ldots a_m$ of length m with a word $b_1 \ldots b_n$ of length n, in the alphabet A, to give us a word $a_1 \ldots a_m b_1 \ldots b_n$ of length $m + n$. We will write $W(A)$ for the set of all words in the alphabet A. We identify A as the subset of $W(A)$ of words of length 1. A *subword*, w', of a given word w is a word consisting of a sequence of consecutive letters of w. That is, w can be written as a concatenation of (possibly empty) words $w = uw'v$, with $u, v \in W(A)$.

Warning: If A happens to be a subset of a group, Γ, then this notation is ambiguous, since $a_1 \ldots a_n$ might be interpreted either as a (formal) word in the above sense, or else as the product of the elements, a_i, in Γ. We would therefore need to be clear in what sense we are using this notation. A word will determine an element of the group, but a given element might be represented by many different words, for example ab and ba are different words, but represent the same element in an abelian group.

Now suppose that B is any set. Let $\bar{B}$ be another, disjoint set, with a bijection to B. We shall denote our bijection $B \longrightarrow \bar{B}$ by $[a \mapsto \bar{a}]$ as a ranges over the elements of B. We use the same notation for the inverse bijection. In other words $\bar{\bar{a}} = a$. Let $A = B \sqcup \bar{B}$. The map $[a \mapsto \bar{a}]$ gives an involution on A. (The idea behind this construction is that B will give us a free generating set, and $\bar{a}$ will give us the inverse element to a.) We consider A to be our alphabet, and let $W = W(A)$ be the set words.

Definition : Suppose $w, w' \in W$. We say that w' is a *reduction* of w if it is obtained from w by removing subword of the form $a\bar{a}$ for some $a \in A$ (or equivalently either $a\bar{a}$ or $\bar{a}a$ for some $a \in B$).

(More formally this means there exist $u, v \in W$ and $a \in A$ so that $w' = uv$ and $w = ua\bar{a}v$.)

Let $\sim$ be the equivalence relation on W generated by reduction. That is, $w \sim w'$ if there is a finite sequence of words, $w = w_0, w_1, \ldots, w_n = w'$, such that each w_{i+1} is obtained from w_i by a reduction or an inverse reduction.

Let $F(B) = W(A)/\sim$. We denote the equivalence class of a word, w, by $[w]$. We define a multiplication on $F(B)$ by writing $[w][w'] = [ww']$.

Exercise:

(1) This is well-defined.

(2) $F(B)$ is a group. (Note that $[a_1 \ldots a_n]^{-1} = [\bar{a}_n \ldots \bar{a}_1]$.)

(3) $F(B)$ is freely generated by the subset $S = S(B) = \{[a] \mid a \in B\}$.

Note that (3) shows that free groups exist with free generating sets of any cardinality. In fact, putting (3) together with Lemma 1.2, we see that, up to isomorphism, every free group must have the form $F(B)$ for some set B. (This gives us another proof of Lemma 1.1, since it is easy to see explicitly that $F(B)$ is generated by $S(B)$.) The observation that $[a]^{-1} = [\bar{a}]$ justifies the earlier remark that $\bar{a}$ is designed to give us an inverse of a. It can be thought of as a "formal inverse".

There is a natural map from $W(A)$ to $F(B)$ sending w to $[w]$. The restriction to A is injective. It is common to identify B with its image, $S = S(B)$, in $F(B)$, and to omit the brackets [.] when writing an element of $F(B)$. Thus, the formal inverse $\bar{a}$ gets identified with the actual inverse, a^{-1}, in F. As mentioned above, we need to specify when writing $a_1 \ldots a_n$ whether we mean a (formal) word in the generators and their (formal) inverse, or the group element it represents in F.

Exercise: If $F(B)$ is finitely generated, then B is finite.

As a consequence, any finitely generated free group is (isomorphic to) F_n for some $n \in \mathbf{N}$.

Definition : A word $w \in W(A)$ is *reduced* if it admits no reduction.

This means that it contains no subword of the form aa^{-1} or $a^{-1}a$ (adopting the above convention that $\bar{a} = a^{-1}$).

Proposition 1.3 : *If $w \in W(A)$ then there is a unique reduced $w' \in W(A)$ with $w' \sim w$.*

Put another way, every element in a free group group has a unique representative as a reduced word in the generators and their inverses.

The existence of such a word is clear: just take any equivalent reduced word of minimal length. Its uniqueness is more subtle. One can give a direct combinatorial argument. However, we will postpone the proof for the moment, and give a more geometrical argument in Section 2.3.

Remarks

(1) Any subgroup of a free group is free. This is a good example of something that can be seen fairly easily by topological methods (see later), whereas direct combinatorial arguments tend to be difficult. (We shall explain this in Section 4.)

(2) Let a, b be free generators for F_2. Let $S = \{a^n ba^{-n} | n \in \mathbf{N}\} \subseteq F_2$. Then $\langle S \rangle$ is freely generated by S (exercise). Thus, by an earlier exercise, $\langle S \rangle$ is not finitely generated. (Again, this is something best viewed topologically — see Section 4.)

(3) Free generating sets are not unique. If $\mathbf{Z} = \langle a \rangle$, then both $\{a\}$ and $\{a^{-1}\}$ are free generating sets. Less trivially, $\{a, ab\}$ freely generates F_2 (exercise). Indeed F_2 has infinitely many free generating sets. However, all free generating sets of F_n have cardinality n (see the end of this section).

(4) Put another way, the automorphism group of F_2 is infinite. (For

example the map $[a \mapsto a, b \mapsto ab]$ extends to automorphism of F_2). The automorphism groups (and outer automorphism groups) of free groups, are themselves subject to intensive study in geometric group theory.

1.5. Group presentations.

The following definition makes sense for any group, G.

Definition : The *normal closure* of a subset of $A \subseteq G$ is the smallest normal subgroup of G containing A. It is denoted $\langle\langle A \rangle\rangle$.

In other words, $\langle\langle A \rangle\rangle$ is characterised by the following three properties:
(1) $A \subseteq \langle\langle A \rangle\rangle$,
(2) $\langle\langle A \rangle\rangle \lhd G$,
(3) If $N \lhd G$, $A \subseteq N$, then $\langle\langle A \rangle\rangle \subseteq N$.

Exercise: $\langle\langle A \rangle\rangle$ is generated by the set of all conjugates of elements of A, i.e.

$$\langle\langle A \rangle\rangle = \langle \{gag^{-1} \mid a \in A, g \in G\} \rangle.$$

We are mainly interested in this construction when G is a free group. If S is a set, and R any subset of the free group, $F(S)$, we write

$$\langle S \mid R \rangle = F(S)/\langle\langle R \rangle\rangle.$$

Note that there is a natural map from S to $\langle S \mid R \rangle$, and $\langle S \mid R \rangle$ is generated by its image.

Definition : A *presentation* of a group, Γ, is an isomorphism of Γ with a group of the form $\langle S \mid R \rangle$.
Such a presentation is *finite* if both S and R are finite.
A group is *finitely presented* if is admits a finite presentation.

We shall abbreviate the presentation $\langle \{x_1, \ldots, x_n\} \mid \{r_1, \ldots, r_m\} \rangle$ to $\langle x_1, \ldots, x_n \mid r_1, \ldots, r_m \rangle$.

It is common to identify x_i with the corresponding element of Γ. In this way, Γ is generated by $\{x_1, \ldots, x_n\}$. An element of R can be written as a word in the x_i and their inverses, and is called a *relator*. Note that such a word corresponds to the identity element in Γ. (Note it is possible for two of the generators x_i, x_j to be equal in Γ, for example if $x_i x_j^{-1}$ is a relator.) We can manipulate elements in a presentation much as we would in a free group, allowing ourselves in addition to eliminate subwords that are relators or to insert relators as subwords, wherever we wish.

Examples

(1) If $R = \emptyset$, then $\langle\langle R \rangle\rangle = \{1\} \subseteq F(S)$, so $\langle S \mid \emptyset \rangle$ is isomorphic to $F(S)$. Thus, $\langle a \mid \emptyset \rangle$ is a presentation of $\mathbf{Z}$, and $\langle a, b \mid \emptyset \rangle$ is a presentation of F_2 etc. Thus a free group is a group with no relators.

(2) $\langle a \mid a^n \rangle$ is a presentation for $\mathbf{Z}_n$.

(3) We claim that $\langle a, b \mid aba^{-1}b^{-1} \rangle$ is a presentation of $\mathbf{Z} \oplus \mathbf{Z}$.

To see this, write $\mathbf{Z} \oplus \mathbf{Z}$ as $\{c^m d^n \mid m, n \in \mathbf{Z}\}$. There is a homomorphism from $F_2 = \langle a, b | \emptyset \rangle$ to $\mathbf{Z} \oplus \mathbf{Z}$ sending a to c and b to d. Let K be its kernel. Thus $\mathbf{Z} \oplus \mathbf{Z} \cong F_2/K$. By definition, $\langle a, b \mid aba^{-1}b^{-1} \rangle \cong F_2/N$, where N is the normal closure of $aba^{-1}b^{-1}$. We therefore want to show that $K = N$. It is clear that $N \subseteq K$. Now F_2/N is abelian. (It is generated by Na and Nb which commute.) Thus a typical element has the form $Na^m b^n$. This gets sent to $c^m d^n$ under the natural map to $\mathbf{Z} \oplus \mathbf{Z} = F_2/K$. If this is the identity, then $m = n = 0$. This shows that $N = K$ as claimed.

The assertion that $aba^{-1}b^{-1} = 1$ is equivalent to saying $ab = ba$. The latter expression is termed a *relation*. This presentation is sometimes written in the notation: $\langle a, b \mid ab = ba \rangle$.

(4) Similarly, $\langle e_1, \ldots, e_n \mid \{e_i e_j e_i^{-1} e_j^{-1} \mid 1 \leq i < j \leq n\} \rangle$ is a presentation of $\mathbf{Z}^n$.

In general, it can be very difficult to recognise the isomorphism type of a group from a presentation. Indeed, there is no general algorithm to recognise if a given presentation gives the trivial group. There are many "exotic" presentations of the trivial group. One well-known example (we won't verify here) is $\langle a, b \mid aba^{-1}b^{-2}, a^{-2}b^{-1}ab\rangle$. The celebrated "Andrews-Curtis conjecture" states that any presentation of the trivial group with the same number of generators as relators can be reduced to a trivial presentation by a sequence of simple moves. This conjecture remains open.

1.6. Abelianisations.

We finish this section with a remark about abelianisations. Let G be a group. The *commutator* of $x, y \in G$ is the element $[x, y] = xyx^{-1}y^{-1}$. Note that $g[x, y]g^{-1} = [gxg^{-1}, gyg^{-1}]$. It follows that the group generated by the set of all commutators is normal. It is usually denoted $[G, G]$. The quotient group $G/[G, G]$ is abelian (exercise), and is called the *abelianisation* of G.

Exercise $F_n/[F_n, F_n] \cong \mathbf{Z}^n$.

This is related to the presentation of $\mathbf{Z}^n$ given above. However it does not make reference to any particular generating set for F_n. (We are considering all commutators, not just those in a particular generating set, though the end result is the same.) Together with an earlier exercise this proves the assertion that $F_m \cong F_n$ implies $m = n$. In particular, all free generating sets of a given finitely generated free group have the same cardinality.