Certainly, any $\Pi g_k r_{j_k}^{\varepsilon_k} g_k^{-1} \in N$ is equal to 1 in G, since each $r_j = 1$. Conversely, suppose a word $w = 1$ in G. We shall show that each insertion or deletion of $r_j^{\pm 1}$ in w can be accomplished by multiplying w by $g_k^{-1} r_j^{\pm 1} g_k$ for some g_k.

Note firstly that deletion of $r_j^{\pm 1}$ can always be accomplished by insertion of $r_j^{\mp 1}$ next to it, followed by cancellation (which is valid in F). Thus it remains to deal with insertions.

Let $w = uv \to ur_j v$ be the insertion of r_j between the factors u, v of w. We can obtain the same result by multiplying w by $v^{-1} r_j v$, since $ur_j v$ is freely equivalent to $uv \cdot v^{-1} r_j v$.

Repetition of this process for each insertion in the sequence required to convert w to 1 gives a word

$$w \Pi g_k^{-1} r_{j_k}^{\varepsilon_k} g_k$$

which is freely equivalent to 1, and therefore

$$w = \Pi g_k r_{j_k}^{-\varepsilon_k} g_k^{-1} \quad \text{in } F$$

so that $w \in N(r_1, r_2, \ldots)$. □

Dyck's 1882 paper is the beginning of combinatorial group theory as a subject, and the first to recognize the fundamental role of free groups. Dyck viewed free groups as the most general groups, since any other group is obtainable by imposing relations on them. The explanation of relations in terms of normal subgroups and quotients suggests a reconstruction of combinatorial group theory in more conventional algebraic terms. This can indeed be done, including the definition of free groups themselves, but it proves to be an object lesson in the impotence of abstract algebra. All substantial theorems in combinatorial group theory still require honest toil with words and relations, and the best labour-saving device turns out to be the topological interpretation of 0.5.1, rather than algebra.

Exercise 0.5.6.1. If G is any group show that the result of adding relations $v_1 = 1$, $v_2 = 1, \ldots,$ to G is $G/N(v_1, v_2, \ldots)$.

0.5.7 The Word Problem and Cayley Diagrams

When a group G arises as a fundamental group, as in 0.5.1, null-homotopic paths correspond to words w which equal 1 in G. Thus the problem of deciding null-homotopy (contractibility to a point) is reduced to deciding whether a given word $w = 1$ in G. Even though we can compute a presentation of G, this problem is not trivial, and its fundamental importance for topology and group theory was first recognized by Dehn 1910, who called it the *word problem*.

Early topologists, such as Poincaré, Tietze, and Reidemeister, frequently commented on the difficulty of group-theoretic problems in topology, on occasion (Reidemeister) saying that the fundamental group seemed merely to translate hard topological problems into hard group-theoretic problems. This pessimism was vindicated when Novikov 1955 proved that the word problem (for specific, finitely presented G) was unsolvable. Novikov's proof is based on the idea of Post 1947 of simulating Turing machines by systems of generators and relations. A word corresponds (roughly) to the tape expression on a Turing machine M, and the relations permit the word to be changed to reflect the atomic acts of M. (The technical difficulty, which is absent in the semigroup case, is the presence of relations $a_i a_i^{-1} = a_i^{-1} a_i = 1$ which do not correspond to acts of M.)

Solution of the word problem for G is equivalent to the construction of a figure $\mathscr{C}_G$ called the *Cayley diagram* of G, introduced for finite groups by Cayley 1878 and for infinite groups by Dehn 1910. If G is generated by $a_1, a_2, \ldots$; then $\mathscr{C}_G$ is a graph with a vertex P_g for each distinct $g \in G$ and an oriented edge labelled a_i from P_g to P_{ga_i} for each generator a_i. It follows that each vertex has exactly one outgoing, and one incoming, edge for each generator. Examples (labelling each vertex g instead of P_g for simplicity) are given in Figure 45. The last example is constructed by noting that there are six distinct elements $g = 1, b, b^2, a, ab, ab^2$, then multiplying each of these by a, b and using the defining relations to reduce each product to one of the six forms already chosen.

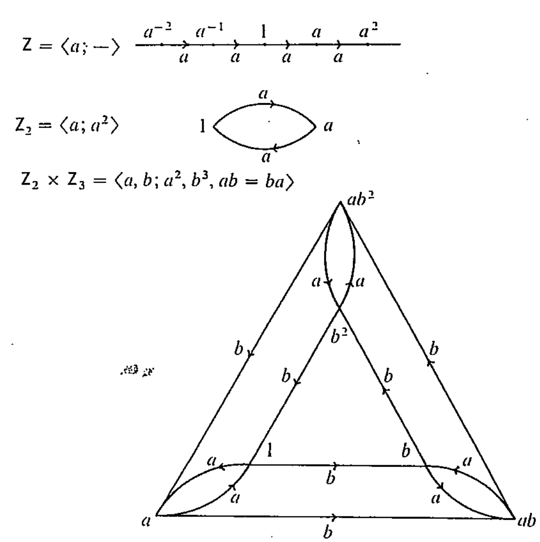

Figure 45

Each word

$$w = a_{i_1}^{\varepsilon_1} \cdots a_{i_k}^{\varepsilon_k}$$

determines a path from P_1 to P_w by following the labels $a_{i_1}, \ldots, a_{i_k}$ in succession, with or against the arrow according as the exponent ε is $+1$ or -1. It follows that $w = 1$ if and only if the path is closed.

Thus if $\mathscr{C}_G$ can be effectively constructed we have a solution of the word problem for G.

Conversely, if the word problem for G can be solved, we can construct $\mathscr{C}_G$.

Effectively list the words of G as $w_1, w_2, \ldots$ and as each w_j appears, use the solution of the word problem to decide whether $w_j =$ any w_i earlier on the list (see if $w_i w_j^{-1} = 1$). If not, put w_j on a *second list*. The second list is then an effective enumeration of the distinct elements of G, which we use as labels for the vertices of $\mathscr{C}_G$.

As each vertex P_{w_j} is constructed, we again use the solution of the word problem to find which of the words $w_j a_i$ is equivalent to a w_k already on the second list (if an equivalent is not found, one will be found later by repeated checking as the second list grows). For each such word we construct an oriented edge labelled a_i from P_{w_j} to $P_{w_j a_i} = P_{w_k}$. This is an effective process which eventually gives each vertex and edge in $\mathscr{C}_G$. □

Since G has many different presentations, $\mathscr{C}_G$ is not unique. However, if there is a solution to the word problem for one finite presentation of G there is a solution for any other finite presentation of G, hence the effective constructibility of $\mathscr{C}_G$ does not depend on the presentation chosen.

Exercise 0.5.7.1. Prove the last remark.

Exercise 0.5.7.2. Show that $\{w: w = 1 \text{ in } G\}$ is r.e. when G is finitely presented.

Exercise 0.5.7.3. Sketch the Cayley diagram of the free group $F_2 = \langle a, b; -\rangle$.

Exercise 0.5.7.4. Describe the Cayley diagrams of the *free abelian groups* $\mathbb{Z} \times \mathbb{Z} \times \ldots \times \mathbb{Z} = \langle a_1, \ldots, a_m; a_i a_j = a_j a_i \ (i, j \leq m)\rangle$ as figures in $\mathbb{R}^m$.

Exercise 0.5.7.5. Figure 46 shows the Cayley diagram of a group. Why is this group nonabelian?

Show that the group is the group of symmetries of an equilateral triangle.

0.5.8 Tietze Transformations

Tietze transformations are simply the obvious ways of transforming a finite presentation $\langle a_1, \ldots, a_m; r_1, \ldots, r_n\rangle$.

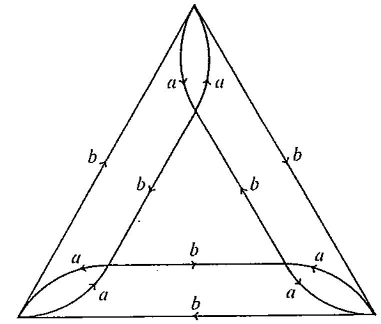

Figure 46

T_1: Add a relation $r_{n+1}(=1)$ which is a consequence of $r_1, \ldots, r_n$. (That is, r_{n+1} is equivalent to 1 with respect to the relations $r_1 = \cdots = r_n = 1$. We write this $r_1, \ldots, r_n \Rightarrow r_{n+1}$.)

T_2: Add a generator a_{m+1} together with a relation

$$a_{m+1} = w(a_1, \ldots, a_m)$$

which defines it as a word in the old generators.

The inverse transformations, which we denote by T_1^{-1}, T_2^{-1}, can also be applied when meaningful.

Tietze's Theorem. *Any two finite presentations of a group G are convertible into each other by a finite sequence of Tietze transformations.*

Suppose G has presentations $\langle a_1, \ldots, a_m; r_1, \ldots, r_n\rangle$ and $\langle a'_1, \ldots, a'_{m'}; r'_1, \ldots, r'_{n'}\rangle$, which we abbreviate to $\langle a_i; r_j(a_i)\rangle$ and $\langle a'_i; r'_j(a'_i)\rangle$. We use the notation $w(x_i)$ to express the fact that w is a word in the letters x_i, and denote the result of substituting a word χ_i for x_i in $w(x_i)$ by $w(\chi_i)$.

Since both presentations denote the same group, there are words α'_i in $a_1, a_2, \ldots$ representing the a'_i and hence satisfying the relations $r'_j(\alpha'_i)$. Then the $r_j(a_i) \Rightarrow$ the $r'_j(\alpha'_i)$ since *all* relations in the a_i are consequences of the $r_j(a_i)$. Similarly there are words α_i in $a'_1, a'_2, \ldots$, representing the a_i, and the $r'_j(a'_i) \Rightarrow$ the $r_j(\alpha_i)$.

We can therefore make the following modifications of the group presentation by Tietze transformations:

$$\langle a_i; r_j(a_i)\rangle$$
$$\to \langle a_i; r_j(a_i), r'_j(\alpha'_i)\rangle \quad \text{by} \quad T_1 \text{ since the } r_j(a_i) \Rightarrow \text{the } r'_j(\alpha'_i)$$
$$\to \langle a_i, a'_i; r_j(a_i), r'_j(\alpha'_i), a'_i = \alpha'_i\rangle \quad \text{by } T_2$$
$$\to \langle a_i, a'_i; r_j(a_i), r'_j(a'_i), r'_j(\alpha'_i), a'_i = \alpha'_i\rangle \quad \text{by } T_1$$
$$\to \langle a_i, a'_i; r_j(a_i), r'_j(a'_i), a'_i = \alpha'_i\rangle \quad \text{by } T_1^{-1}$$
$$\to \langle a_i, a'_i; r_j(a_i), r'_j(a'_i), a'_i = \alpha'_i, a_i = \alpha_i\rangle \quad \text{by } T_1 \qquad (*)$$

since the relations $a_i = \alpha_i$ are true in the group and hence consequences of the relations already present. But (*) is symmetric with respect to primed and unprimed symbols, so it could equally well be obtained from $\langle a_i'; r_j'(a_i')\rangle$. By reversing the latter derivation we obtain

$$\langle a_i; r_j(a_i)\rangle \to (*) \to \langle a_i'; r_j'(a_i')\rangle. \qquad \square$$

Since we can effectively enumerate all consequences of a given finite set of relations, and hence all possible sequences of Tietze transformations which can be applied to a given presentation, Tietze's theorem shows that we can effectively enumerate all finite presentations of a given group. Thus the problem of deciding when two presentations are the same, the *isomorphism problem* of Tietze 1908, is similar to the word problem—in both cases we can effectively enumerate the pairs of equal objects, and the difficulty is to find the pairs of unequal objects. It actually follows from basic results of recursive function theory (see Rogers 1967) that the two problems are of the same degree of unsolvability, that is, a solution of one would effectively yield a solution of the other. (In particular, the isomorphism problem is unsolvable.) In individual cases, however, the isomorphism problem is usually harder to solve than the word problem.

On the positive side, the Tietze theorem is often a slick way to prove existence of algorithms or semidecision procedures. For example, if G has a property that can be recognized from one of its presentations we can eventually verify this property by enumerating all the presentations of G. Examples of such properties are:

(i) being abelian (all generators commute)
(ii) being finite (all relations of the form $a_i a_j = a_k$)
(iii) being a specific finite group (relations given by multiplication table)
(iv) being free (no relations).

Exercise 0.5.8.1. Show that $\langle a, b; abab^{-1}\rangle = \langle c, d; c^2d^2\rangle$.

Exercise 0.5.8.2. Suppose that *infinitely many* consequence relations or new generators can be added in a transformation of type T_1 or T_2 respectively. Deduce that any two presentations of the same group are then convertible to each other by a finite sequence of Tietze transformations.

Exercise 0.5.8.3. Give an algorithm for finding α_i and α_i' from two presentations $\langle a_i; r_j(a_i)\rangle$ and $\langle a_i'; r_j'(a_i')\rangle$ of the same group. (This gives a "uniform" solution to Exercise 0.5.7.1.)

Exercise 0.5.8.4. If G has a finite presentation, show that in any presentation

$$G = \langle a_1, \ldots, a_n; r_1, r_2, \ldots\rangle$$

all but a finite number of relations are superfluous.

0.5.9 Coset Enumeration

As a final example of the way finiteness can be discovered by systematically enumerating words, consider the case of a subgroup H of a finitely presented group G. If the set of cosets Hg for $g \in G$ is finite, H is said to be of *finite index* in G. In this case there is a finite set $\{g_1, \ldots, g_k\}$ of *coset representatives* such that

$$G = Hg_1 \cup \cdots \cup Hg_k.$$

We now show how to find such a set, if one exists.

$G = Hg_1 \cup \cdots \cup Hg_k$ if and only if the set $\{Hg_1, \ldots, Hg_k\}$ is closed under right multiplication by the generators of G and their inverses. That is

$$Hg_j a_i = \text{some } Hg_{j'}, \; Hg_j a_i^{-1} = \text{some } Hg_{j''}$$

for each generator $a_1, \ldots, a_m$ of G. Now assuming H is effectively enumerable, we can verify the equality of two cosets by enumerating their members, along with an enumeration of equal words in G, until we find a common element.

It therefore suffices to enumerate all the finite sets $\{g_1, \ldots, g_k\}$ in G, and for each one try to verify that $\{Hg_1, \ldots, Hg_k\}$ is closed under right multiplication by looking for equal pairs $Hg_j a_i, Hg_{j'}$ and $Hg_j a_i^{-1}, Hg_{j''}$. Eventually such a verification will succeed. $\square$

A more practical version of the above idea is known as the *Todd–Coxeter coset enumeration method* (Todd, Coxeter 1936).