Chapter One

Group Theory

It is not surprising, given the syntactical nature of arguments in logic and in combinatorial group theory, that one of the earliest examples of a natural, widely studied, mathematical problem to be shown algorithmically unsolvable lies in group theory. Our unsolvability (and dichotomy) results in topology and geometry will come about by encoding group theory into these subjects, which will be done in chapter 2. The goal of this chapter is to provide the necessary background in group theory.

This subject contains a number of deep theorems. Some of the highlights of this part are:

- the unsolvability of the word problem by Novikov and Boone, and the triviality problem by Adian and Rabin (section 1.2);
- Higman's embedding theorem (section 1.2);
- the recent work of Sapir with Birget and Rips on Dehn functions (section 1.3);
- the results of Borel-Wallach and of Clozel on the cohomology of arithmetic groups (section 1.5);
- the theorems of Baumslag and Dyer with Heller and Miller on the group homology of finitely presented groups (section 1.6).

1.1 PRESENTATIONS OF GROUPS

Let G be a group. G is said to be *finitely generated* if there are finitely many elements $g_1, g_2, \ldots, g_k$ such that every element of G is a product of these elements (many times) and their inverses. All finitely generated groups are countable, but the rational numbers give an example of a nonfinitely generated countable group.

Note that saying that G is finitely generated is exactly the same thing as saying that there is a surjection from a free group $F_k \to G$ for some k. We say that a subgroup H of G is *finitely normally generated* if there is a finite set S such that H is the smallest normal subgroup of G containing S. Alternatively, the elements of H are products of things of the form gsg^{-1}, where s is from S, and g lies in G.

The group G is finitely presented if there is a surjection $F_k \to G$ whose kernel is finitely normally generated. A set R which normally generates the kernel is called a set of relations for G. One uses the notation

$$G = \langle g_1, g_2, \ldots, g_k \mid r_1, r_2, \ldots, r_s \rangle,$$

where the r's are a list of the elements of R. The r's should be thought of as being words in the g-letters. One can easily show that the kernel of one surjection of a free group to G is finitely normally generated iff it is for any other surjection. Using this, one can show, for instance, that the following group is not finitely presented:

$$\langle x, y \mid [x, x^a y x^{-a}] \text{ where } a = 0, 1, 2, 3, \ldots \rangle = \mathbb{Z} \wr \mathbb{Z},$$

where $[g, h] = ghg^{-1}h^{-1}$ is the commutator of g and h, and $\wr$ denotes the wreath product (for those familiar with this concept; we will never have any use for it).

One thinks of the finite presentation as giving the group defined by these generators, subject to the given set of relations (like an axiomatic system). The word problem asks one to give an algorithm for deciding whether a combination of generators represents the trivial element of G; in other words, whether a certain potential relation is a consequence of the relations that are already part of the set R. (Again, this is a property of a group, not of the way the group is presented.) We will discuss the word problem in the next section.

Remark. It is not at all easy to tell if two finite presentations define the same group. The Tietze theorem asserts that two presentations define the same group iff there is a sequence of elementary moves (and their inverses) that relate the two presentations. The first move adds a new generator and a new relation that defines the generator in terms of the old ones, for example,

$$\langle x, y \mid \ \rangle = \langle x, y, z \mid z = xyx^{-215}y^{12}x \rangle.$$

(We use the convention that a relation of the form $r = s$ is just a user friendly way of writing rs^{-1}.) The second move allows one to add a relation that is a "consequence" (i.e., that lies in the normal closure) of the others. So

$$\langle x, y \mid yx = xy \rangle = \langle x, y \mid yx = xy, xyxy = yxyx \rangle.$$

These moves always increase the complexity of presentations, but that means that their inverses can decrease complexity. We will see that relating two "simple presentations" by Tietze moves could involve going through very complicated intermediate presentations. (We will never really need the Tietze theorem, but it will sometimes be useful for illustration purposes.)

Besides their natural logical interest, finitely presented groups occur extremely naturally throughout mathematics. Probably the simplest, general sources of these are fundamental groups of compact manifolds and arithmetic groups (or more general lattices in real Lie groups).

The reason that compact smooth manifolds have finitely presented fundamental groups is because they are all homeomorphic to finite polyhedra (= finite simplicial complex) by a relatively difficult theorem of Cairns and Whitehead (CW), or by Morse theory, which, at least, shows that they are homotopy equivalent to finite CW complexes.

Given a finite CW complex or simplicial complex one can very simply write down a presentation of the fundamental group. The 1-skeleton is always homotopy equivalent to a wedge of circles; the 2-cells of the 2-skeleton then give the relations. The proof of this description goes by way of van Kampen's theorem, which describes the fundamental group of the union of two spaces that intersect "nicely." To describe the answer, we shall need the constructions of "amalgamated free product" and "HNN extension."

Definition. Let A, B, and C be groups, and let C be a subgroup of both A and B. $A *_C B$ is the group defined to have the universal property that A and B both map to $A *_C B$, and if $f\colon A \to D$ and $g\colon B \to D$ are homomorphisms that agree on C, then there is a unique extension of f and g to a map $A *_C B \to D$.

General nonsense implies that $A *_C B$ is unique (up to canonical isomorphism) if it exists. To construct it, we can give a formula. Suppose $A = \langle a_1, a_2, \ldots \mid r_1, r_2, \ldots \rangle$, $B = \langle b_1, b_2, \ldots \mid s_1, s_2, \ldots \rangle$, and C is generated by $c_1, c_2, \ldots$. As C is a subgroup of both A and B, we can write $c_1 = c_1(a\text{'s})$ and $c_1 = c_1(b\text{'s})$, $c_2 = c_2(a\text{'s})$ and $c_2 = c_2(b\text{'s})$, etc. Now

$$\begin{aligned} A *_C B = \langle & a_1, a_2, \ldots, b_1, b_2, \ldots \mid r_1, r_2, \ldots, s_1, s_2, \ldots, c_1(a\text{'s}) \\ &= c_1(b\text{'s}),\ c_2(a\text{'s}) = c_2(b\text{'s}), \ldots \rangle. \end{aligned}$$

(Note that actually one needs only a homomorphism from C into both A and B for the definition and universal property to make sense. The universal property is helpful for a conceptual understanding of why the amalgamated free product is independent of all the choices involved.) With this preparation, we can state van Kampen's theorem.

Van Kampen's theorem

If two spaces with fundamental groups A and B intersect along a connected space with a fundamental group C, then their union has the fundamental group $A *_C B$.

The HNN extension is a natural analogue of the amalgamated free product, and comes up in determining the fundamental group of a union when the intersection is not connected.

Definition. Let A be a group and B_j, $j = 1, 2$, isomorphic subgroups (let ϕ be the isomorphism). Then the HNN extension $A*_B$ is defined by the universal property that, if $f : A \to D$ is a homomorphism which restricts to conjugate maps on the two copies of B, then it extends uniquely to $A*_B$. In terms of generators and relations, the formula is

$$A*_B = \langle a_1, a_2, \ldots, t \mid r_1, r_2, tb_1t^{-1} = \phi(b_1), tb_2t^{-1} = \phi(b_2) \ldots \rangle$$

with the obvious notation.

If one glues a cylinder $Y \times I$ whose fundamental group is B along both of its ends to a space with fundamental group A, then one gets into a situation where the HNN extension is defined. Indeed, the fundamental group of this union is the HNN extension.

Exercise

Show that van Kampen's theorem together with the above addendum for cylinders together suffice to deal with all unions of polyhedra along subpolyhedra.

It is worth noting that these constructions can behave quite oddly when the "subgroups" are really groups with homomorphisms with nontrivial kernels. For instance, if we consider $\mathbb{Z}_2$ and $\mathbb{Z}_3$ amalgamated "along" $\mathbb{Z}$ which maps surjectively to both groups, one obtains the trivial group.

On the other hand, if all of the "inclusions" are really injections, then A automatically injects into $A *_C B$ and into $A*_B$. In fact, there is quite a natural normal form for elements in the amalgamated free product and HNN extension, given a choice of coset representatives for the subgroup. In the next section we will make extensive use of these constructions, and, in particular, this injectivity statement.

The proof of this and several of the other basic theorems of combinatorial group theory can be written down in a completely opaque combinatorial fashion, but are actually quite transparent from a geometric point of view.

Definition. A graph of groups is a graph Γ such that each vertex v is associated a group G_v and each edge is assigned G_e. For each of the two inclusions of endpoints u, v in an edge e, there are given injections $G_e \to G_u$ and $G_e \to G_v$.

Notice that an edge can have both endpoints being the same vertex, in which case one has a group with two isomorphic subgroups. Thus, the data for an edge are exactly the data necessary for defining amalgamated free products or HNN extensions. For any connected graph, one can then inductively define, using amalgamated free products or HNN extensions on individual vertices, the fundamental group of the graph of groups. If all the vertex and edge groups are trivial, this is just the fundamental group of the graph.

One could define this notion in a less ad hoc way by an appropriate universal property; we leave this for the reader. One can also see that it is the fundamental group of a union of spaces, which overlap only in pairs, and that "fit together" according to the pattern of the graph (vertices corresponding to spaces, and edges to overlaps) and with the fundamental groups of every piece determined by the labels on the graph.

The following proposition gives a connection between group actions on trees and graphs of groups.

Proposition *If a group G acts simplicially on a tree T without inversions (i.e., invariant edges are fixed), then the quotient T/G is a graph of groups, with fundamental group G, wherewith each vertex or edge is associated its stabilizer. Conversely, every graph of groups comes from an action of its fundamental group.*

The proof is little more than covering space theory. Note that, as a consequence, A injects into $A *_C B$, since the latter group acts on a tree, with A and B as vertex groups and C as the edge group. Similarly, A injects into the HNN extension $A{*}_B$ because the latter also acts on a tree with A as a stabilizer subgroup for a vertex. Here are some other corollaries:

1. Subgroups of free groups are free. (Freeness is equivalent to having a free action on a tree; free actions restricted to subgroups remain free.)
2. Exercise: Show that any subgroup of finite index in a nonabelian free group is a free group of higher rank. Show that the commutator subgroup of a nonabelian free group is a free group of infinite rank. What are its generators?
3. Generalized Kurosh subgroup theorem: Any subgroup of a graph of groups is a graph of groups where the edge and vertex groups are subgroups of the original vertex and edge groups.

The usual Kurosh theorem corresponds to free products, that is, where edge groups are trivial, so for the subgroup all edge groups are also trivial, so that the fundamental group is a free product of subgroups of the free factors and free groups.

Exercise

Supposing that G and H are nontrivial groups, not both $\mathbb{Z}_2$, show that $G * H$ contains a free group of rank 2.

1.2 PROBLEMS ABOUT GROUPS

Now let us return to the theory of group presentations. Recall that the word problem for a finitely presented group is to determine when a word represents the trivial element.

Example 1 *Finite groups have a solvable word problem. (Use their multiplication tables.)*

So do finitely presented residually finite groups, that is, groups G with enough homomorphisms to finite groups to catch any nontrivial element. (In other words, each nontrivial element g of G is mapped nontrivially by some homomorphism of G into a finite group.) The proof of this goes as follows. Start two machines going. The first lists elements of the normal closure of R systematically (i.e., going though the elements of G to get many conjugates of the elements of R, and taking many products of these) and checks to see if the given word occurs. If this happens, the machine yells "w is trivial." The second machine lists all homomorphisms from G into any finite group (why is this algorithmic?) and then checks if the word is mapped trivially. If it is not, this machine yells "w is nontrivial." Clearly only one of these will happen, and if G is residually finite, one of them will.

One could want to know a bound on how long it would take an algorithm to determine if w is trivial. In general, it could be quite bad.

Another example is the free group, where the algorithm is quite simple; one just looks for appearances of symbols like xx^{-1} within the word, and removes these. When one is done, one has a reduced word, and every group element has only one expression as a reduced word. This is an extremely fast algorithm. (It is linear in the length of the word.) There is a large and rich class of groups with a linear time solution to the word problem (subgroups of products of hyperbolic groups in the sense of Gromov; see the references), but I do not know of many general theorems for them.

Nowadays, there are even examples of finitely presented solvable groups with unsolvable word problems! But this is running ahead in our story. The wonderful theorem of Boone and (P. S.) Novikov is simply the following:

Theorem *There are finitely presented groups with an unsolvable word problem.*

We will not even sketch the proof here. The overall idea is to encode a Turing machine into a finite presentation (using a series of amalgamated free products and HNN extensions) so "the normal form theorem" (for elements of an amalgamated free product or HNN extension) implies that the only way that a word will be trivialized is via the appropriate computation of the Turing machine. That, the construction, and the proof that it works, will give you the theorem. A number of refinements and extensions will be of importance to us later. We will get to these.

A rather different approach to these problems comes about via the following landmark result:

Higman Embedding Theorem

A finitely generated group is a subgroup of a finitely presented group iff it has a computably enumerable set of relations.

(For infinitely generated groups with a computable set of generators, the same is correct.) The necessity of the computable enumerability of the relation set is easy. The generators of the subgroup are some specific words, and all their relations are specific relations that hold in the original group. We can always enumrate the relations in a finitely presented group, by taking all products of conjugates of the relators.

By now there are a number of different proofs of Higman's theorem. We will give some references below. The theorem is remarkable in that it relates a basic group theoretic notion, embeddability in a finitely presented group, to a computation theoretic one. Moreover, it very quickly gives rise to a proof of the Novikov-Boone theorem, as follows. Let S be a c.e. set which is not computable. Consider the finitely generated, computably presented group

$$G = \langle a, b, c, d \mid a^k b a^{-k} = c^k d c^{-k} \text{ for } k \in S \rangle.$$

Note that G is a free product with amalgamation of two free groups $\langle a, b \rangle$ and $\langle c, d \rangle$. The relation $a^k b a^{-k} = c^k d c^{-k}$ is true iff k is an element of S, by the normal form theorem for amalgamated free products. Thus we cannot tell in G whether a word represents the trivial element.

Embedding G into a finitely presented group gives a finitely presented group with unsolvable word problem. By the way, Higman's technique shows the existence of a "universal" finitely presented group, which contains all others. Later we will make use of such groups.

Now let us turn to the triviality problem, which was solved by Adian and Rabin in much greater generality.

Definition. A *Markov property* of a group is a property such that (1) if G has this property, so does any subgroup of G, and (2) there is some group H not possessing this property.

Theorem (Adian and Rubin) *There is no algorithm to decide if a finite presentation has any particular Markov property.*

So one cannot tell if a group is trivial, if it is finite, abelian, nilpotent, solvable, free,[30] has a solvable word problem, is torsion-free, or contains infinitely divisible elements.

[30]Exercise: Show that it is impossible to decide whether a group is freely generated by a specific set of elements. Hint: Use HNN extensions.

We shall give the proof of this, since it is quite simple, given what we know about amalgamated free products and because the method is very important.

Proof. Given a finitely presented (f.p.) group G and an element w of G, we shall construct a new f.p. group G_w such that (1) either w was the trivial element, in which case G_w is the trivial group, or (2) w is nontrivial, in which case G_w contains G as a proper subgroup.

Note that this immediately gives the impossibility of deciding triviality, since an algorithm for this would give an algorithm for deciding whether w is trivial. It also implies the theorem in general, by picking G to be a free product $H * K$, where K has an unsolvable word problem, and H does not have the given Markov property. The group $(H * K)_h$ has the Markov property iff h is the trivial element, which cannot be discerned algorithmically.

Let $G = \langle x_1, x_2, \ldots, x_k \mid R\rangle$ and let w be an element of G. We form $G * F_2$, where F_2 is generated by t, s.

The normal form theorem for elements in a free product gives us a free group of rank $k+1$ in this free product, and we will choose one generated by $w[w, t]$, and $x_{a'}^{\alpha}$, where $\alpha_a = s^a t^a$. We use the standard notation that [,] denotes the commutator and "exponentiation indicates conjugation" in groups. The condition on the α_a ensures that the words $x_{a'}^{\alpha}$ have little cancellation possible among them, and therefore generate a free group. (This is a formal matter given the normal form theorem for free products, and we leave an exact construction as a worthwhile exercise for the reader.) One can add on conjugates of s and t by complicated words and still have a free subgroup of $G * F_2$.

Let A be a group containing a free subgroup F on $k+3$ generators, the first one of which normally generates A. (Many of the fundamental groups of knot complements in the 3-sphere have this property. The meridian[31] of the knot is always an element which normally generates the group; the fundamental group of the Seifert surface is always a free group, and if we omit one of them, then the remaining ones together with the meridian will also generate a free group; within this free group one can increase the rank at will.) Set $G_w = (G * F_2) *_F A$.

Clearly, if w is nontrivial, this free product is nontrivial. If w is trivial, then A dies because w trivial kills $w[w, t]$, which has been identified with a normal generator of A. Once A dies, so do all the elements of F, but these each go to conjugates of the remaining generators of $G * F_2$, and thus this whole group is killed as well, completing the verification of the construction.

Finally, armed with this construction, one easily builds groups $H * G_w$ which have a given Markov property if and only if $w \neq e$.

Appendix: Some Refinements and Extensions

The study of the algorithmic problems about groups did not end with the unsolvability of the word problem; indeed, that was just a beginning.

[31] This is the small circle that bounds a tiny 2-disk which intersects the knot once.

To understand the first extension, we need the notion of Turing reducibility. Let S and T be sets of natural numbers. We will say that $S \leq T$ if one can compute S from an oracle that decides membership in T.

Notice that, even if T is c.e., S need not be; for instance, the complement of T is always computable from T. However, we shall restrict our attention to c.e. sets. We shall say that S and T *have the same (c.e.) degree*[32] if $S \leq T$ and $T \leq S$; in other words, if each can be computed in terms of the other. Note that we shall identify c.e. sets with the Turing machines that define them (say, as being the set of inputs on which the machine halts).

Now, the set of words representing the trivial element is a c.e. set, and therefore represents a c.e. degree. After learning that it is possible for this set to be noncomputable, it becomes natural to ask whether there are any restrictions on the c.e. degrees. The following theorem completely answers this question:

Theorem (Fridman, Clapham, and Boone) *If D is a c.e. degree, then there is a finitely presented group with degree exactly D. More precisely, there is a uniform construction starting from a Turing machine T producing a group $G(T)$, whose word problem is of the same degree as (the halting set of) T.*

One way to prove this is to give a precise version of Higman's embedding theorem which preserves the c.e. degree of the word problem. Then the construction described above certainly would yield this theorem.

There are other natural problems about elements of a f.p. group that these theorems imply are algorithmically undecidable, and of arbitrary degree, as one varies the group.

Theorem *There is no algorithm to decide, in general, whether a group element*

1. *is of finite order,*
2. *is of the form* $[x, y]$*,*
3. *lies in the center,*
4. *commutes with another given element, or*
5. *is conjugate to another given element.*

The proof for (5) is obvious; conjugacy is a more general problem than triviality. (3) and (4) can be proven simultaneously by considering a group $\mathbb{Z} * G$, for G a group with an unsolvable word problem. The only element that commutes with tg (t in $\mathbb{Z}$, g nontrivial in G) is the identity, which is the whole center, so the word problem for G reduces to either of these problems. To prove (1) it would suffice to observe that one can produce groups as above which are torsion-free, which is true. (They can be built up from the trivial

[32] These are sometimes referred to as Turing degrees, or as degrees of unsolvability.

group by HNN extensions and amalgamated free products; as the reader can check, this implies that the group is torsion free.[33])

Another approach would be to produce a recursively presented group and then Higman embed. So let

$$G = \langle a_1, a_2, \dots \mid a_k^{f(k)} \rangle,$$

where the exponent $f(k)$ is 0 if the kth Turing machine does not halt on input k and is the number of steps that machine takes to halt, if it does. This G clearly has unsolvable "torsionality," and it can be embedded into a f.p. group.

Only (2) requires (as far as I can see) more trickery. Rather than prove it, let me just point out that it is a close cousin to the word problem and the conjugacy problem (5). If one considers a space with fundamental group G, the word problem asks one to decide whether a curve bounds a disk; the conjugacy problem asks whether a pair of curves bound an annulus. (2) asks whether a curve is the boundary of a punctured torus, whereas (4) asks whether two curves could be homotoped to lie on torus.

In this context it is worth observing that the question of whether a curve is the boundary of some compact surface is solvable. This is the same as asking whether it is a product of some number of elementary commutators, i.e., whether the element is trivial in the abelianization of the group.

Remark. The conjugacy problem is clearly "harder" than the word problem. In fact, Clapham showed that for an arbitrary pair of c.e. degrees $D \le E$, there is a finitely presented group with word problem of degree D and conjugacy problem of degree E.

While we have mentioned only briefly (in the Introduction) that there are degrees besides computable and K, the degree of the halting problem, actually, the structure of the set of degrees, is extremely rich and complicated. In particular, it is a dense partial ordering, and there are also many noncomparable degrees, and so on.

We will close this section with the comment that for arithmetic groups, the word problem and conjugacy problems are in fact solvable (although it would be nice to get good bounds). However, the following generalized word problem is not solvable:

Theorem (Mikhailova) *There is no algorithm in $F_2 \times F_2$ to decide whether a given element lies in the group generated by a given finite set of elements.*

Since $F_2 \times F_2$ is a subgroup of $\mathrm{SL}_n(\mathbb{Z})$ for $n > 3$, we obtain the unsolvability of the generalized word problem for these arithmetic groups.

[33]Notice that the groups produced by HNN and amalgamated free products starting from the trivial group all have finite-dimensional Eilenberg-MacLane complexes, which gives a "geometric proof" of the nonexistence of nontrivial torsion.

The beautiful construction is irresistible. Let H be a two-generator group with an unsolvable word problem. (Such exist; even without this fact, one could modify the construction slightly, as the reader can readily see.) Write $H = \langle x, y \mid R \rangle$. Now consider the subgroup S of $F_2 \times F_2$ generated by (x, x), (y, y), and $(1, r)$ for r in R. It is trivial to check that the pair (u, v) lies in S iff u and v represent the same element of H.

1.3 DEHN FUNCTIONS

The Dehn function of a presentation of a group is a concrete measure of how hard it is to solve the word problem. More precisely, we ask, for all words of length $\leq n$ that actually represent the trivial element, what is the largest number of relations it is necessary to use (when a relation is used twice, it is counted twice) in order to prove this.

This function does depend on the presentation. However, its order of growth does not. More precisely, the Dehn functions of two different presentations satisfy a relation of the sort

$$g(n/B) - Cn - D < f(n) < g(Bn) + Cn + D \qquad (1)$$

for some constants B, C, D. (The linear term is there for fairly trivial reasons: see example 1 below.)

The famous Gromov hyperbolic groups are those whose Dehn functions grow linearly. Remarkably, if the Dehn function is subquadratic, it is automatically linear. Thus, there is a nontrivial subject of the possible Dehn functions of f.p. groups.

In this section, I would like to explain a little bit about their theory and describe some remarkable work that gives an almost complete solution to the problem of characterizing Dehn functions by M. Sapir with J. C. Birget and E. Rips.

It is important to be clear that Dehn functions measure the difficulty of a particular method of trying to solve the word problem; there could be other solutions that are much more rapid. They are the analogues of the stopping times of particular Turing machines defining a given c.e. set. However, unlike the Turing machine situation, for a given group, the Dehn function has a much stronger well-definedness property (1).

Although "eastern philosophy" as we introduced it in the Introduction is phrased in terms of arbitrary algorithms, in chapters 3 and 4, we will give versions that can make use of Dehn functions, rather than arbitrary stopping times for arbitrary solutions of the word problem. Let us compute some examples.

Example 1 *The Dehn function of $\mathbb{Z}$. Let $\mathbb{Z} = \langle x \mid \rangle$. Then the Dehn function is trivial. However, for the presentation $\langle x, y \mid y \rangle$, the Dehn function is linear. One removes all y's one at a time, and the number of y's present can be linear in the word length.*

Exercise

Verify the linearity of the Dehn function for a free group.

Example 2 *The Dehn function of* $\mathbb{Z}^2$ *is quadratic. Consider* $\mathbb{Z}^2 = \langle x, y \mid [x, y] \rangle$. *Now, we know that in* $\mathbb{Z}^2$ *we have* $[x^n, y^n] = e$. *However the "proof" of this is quite large:* $[x^n, y^n] \sim [x, y]^{n^2}$ *(see figure 8, where "$\sim$" means that on the right-hand side we have a product of that number of conjugates of* $[x, y]$*). It is not very hard to see, by hand, that there is no smaller product of (conjugates of) relations that gives* $[x^n, y^n]$*, so one has that* $D(4n) \geq n^2$. *It is also easy enough to see that a quadratic number of uses of the relations suffices to "reduce" any word to normal form.*

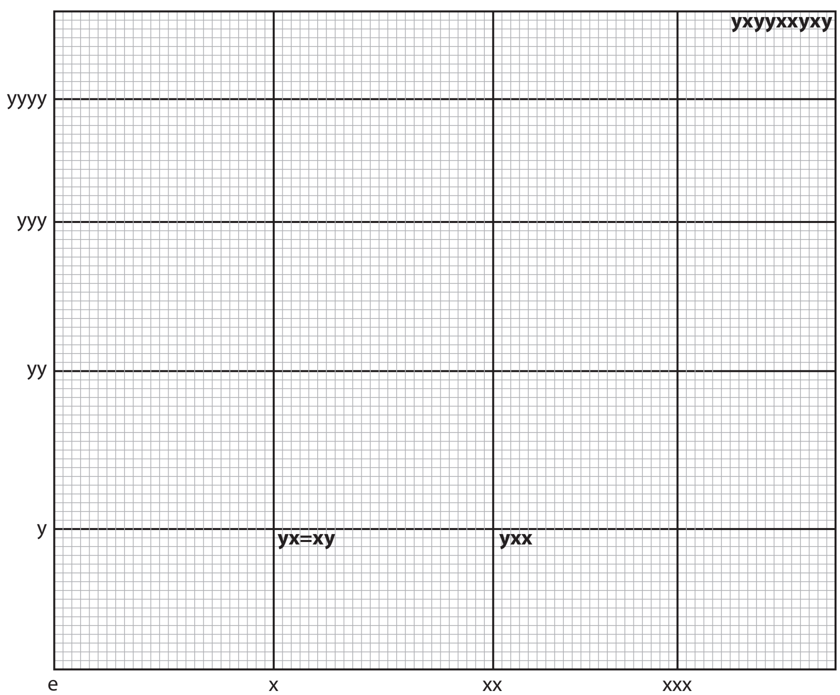

Figure 8. The Cayley graph of Z^2. (Vertices are labeled by group elements; we have labeled them using words that describe specific paths from e to the given node.) Note that the "curve" bounded by a large square is the boundary of quite a lot of small squares. The small unit squares represent the basic relator.

Note, by the way, that using the two homomorphisms $\mathbb{Z}^2 \to \mathbb{Z}$ (killing x and y, respectively), we can get a linear solution to the word problem for $\mathbb{Z}^2$. I am not sure whether there are any reasonable results about the class of groups that have a linear time solution to their word problems. One can often very substantially compress Dehn functions using "nonsyntactic" algorithms. Before going further, it is worth pointing out a limit to this:

Theorem *The Dehn function of a f.g. group is computable iff the word problem is solvable.*

If the Dehn function is computable, then one knows how many relators to try to combine to trivialize a word; conversely, given a solution to the word problem, one can then search out products of relations for all the words of length n that are, in fact, trivial. When this has been accomplished, then look at the length of the longest one you found. This is a computable upper bound for the Dehn function. Then check all products with a smaller number of relations to make sure you did not miss the shortest product realizing these words, to actually compute the Dehn function.

Example 3 *The Heisenberg group H of 3×3 upper diagonal unipotent integer matrices has a presentation with two generators and two relations:*

$$H = \langle x, y \mid [x, [x, y]], [y, [x, y]] \rangle,$$

and has a cubic Dehn function. We have $[x^n, [x^n, y^n]] = e$ in H, but it is a product of a cubic number of relators.

The way to see this is to think somewhat geometrically about the meaning of Dehn functions. Let Z be the universal cover of a space (say a manifold, to be specific) with fundamental group π. Words that represent the trivial element are essentially the same thing as closed nullhomotopic loops in the base, and, by covering space theory, these lift to closed loops in Z. Since they are nullhomotopic, these loops bound 2-disks. The Dehn function is equivalent to seeking the smallest area disk in Z that can be found to bound an arbitrary closed curve of length L.

In examples 1 and 2, we were discussing the Euclidean line and plane, respectively, where the answers are linear and quadratic. (By the way, note that as the dimension of the Euclidean space increases, the Dehn function remains quadratic.) To do example 3, consider the map $f : H(\mathbb{R}) \to \mathbb{R}^2$ (of real 3×3 unipotent upper diagonal matrices) to $\mathbb{R}^2$ (giving the entries immediately above the diagonal) by killing the commutator. Observe that $f^*(dx \wedge dy) \leq dA$ (the area 2-form). Since $dx \wedge dy$ is a closed 2-form, one can integrate it over any disk bounding a given closed curve, like the one represented by $[x^n, [x^n, y^n]]$, and (1) the integral is independent of the bounding disk (even surface) and (2) by the area inequality, it gives a lower bound on the area of such a bounding disk. Playing with these gives the cubic nature of the Heisenberg group.

Remark. This cubic nature is due to the identity $[x^n, y^n] = [x, y]^{n^2}$, so we have n^2 commutators to commute with x^n. This identity asserts that the copy of the integers generated by $[x, y]$ in H is *quadratically distorted*. Using this fact, one can also see that the number of distinct elements of the group represented by words of length n (the *volume growth* of the group) is quartic.

Remark. A finitely generated subgroup of $F_2 \times F_2$ has nonrecursive distortion exactly when the generalized word problem associated with that subgroup is unsolvable.

Example 4 *Let us now give some exponential and superexponential examples; with a bit of trickery, one can promote these higher up the Ackermann hierarchy.*[34]

The solvable Baumslag-Solitar group $G = \langle x, y \mid xyx^{-1} = y^2\rangle$ *is actually a linear group. Then* $[y, x^n yx^{-n}]$ *is a linear length word that requires an exponential number of relations to kill. Note that* $x^n yx^{-n} = y^{2^n}$. *Now, let us add a new generator as follows:*

$$\langle x, y, t \mid xyx^{-1} = y^2, tyt^{-1} = y, txt^{-1} = x^2\rangle.$$

Then $[y, [t^n xt^{-n} y, t^{-n} x^{-1} t^n]]$ *is a linear length word that now needs* 2^{2^n} *relations to kill, and so on.*

We leave the verification to the reader. A useful method for doing these types of calculations is to try to consider "van Kampen diagrams," which are directed planar graphs, where every edge is labeled by a generator of the group, and every face is labeled by a relator of the group. A word w *is trivial iff there is a van Kampen diagram over the given presentation, whose boundary is that word.*

In all of our examples, the relations are designed for easy examination of possible van Kampen diagrams. We shall not pursue this direction any further,[35] but now move on to some very general theorems.

Theorem *(Birget-Sapir) If* $D(n)$ *is the Dehn function of a finitely presented group* π*, then it is the stopping time of some nondeterministic Turing machine that solves the word problem for* π.

Here is the idea of this restriction. Suppose we knew that (a function equivalent to) n^α were a bound on the Dehn function of G. Then I would check 2^{n^α} products of relations to see which words were trivialized, and this would give me $D(n)$ exactly. Thus, if the Dehn function were actually (up to equivalence) of the form n^β, we would be able to compute β quite quickly. It turns out that, when you unravel this, it implies that the first k digits of β can be computed in

[34]The Ackerman hierarchy measures how much induction is used in the definition of a function. nm is adding m to itself n times; iterating that, we have m^n, which is multiplying m by itself n times; then one can consider $m^{[n]}$, which would exponentiate itself n times. After that, it becomes more awkward: $m^{[[n]]}$ would be doing $m^{[m]}$ to itself n times, and so on. The Ackermann function is the diagonalization of this $A(2) = 2^2$, $A(3) = 3^{[3]}, \ldots$.

[35]It seems quite possible that the extension of some of our geometric results to dimension four will depend precisely on the continuation of these explicit methods.

time $2^{2^{2^k}}$. In other words, if the Dehn function grows quite slowly, it cannot take too long to compute it. The best result known currently is almost the converse of this observation.

Theorem *If $T^4(n)$ is superadditive ($T^4(m+n) \geq T^4(m) + T^4(n)$), and T is the stopping function for a (perhaps nondeterministic[36]) Turing machine, then $T^4(n)$ is equivalent to the Dehn function of a finitely presented group π.*

The group is built out of the machine whose stopping time is T. More precisely, if M is this machine, there is an injective map from the input words of M to the words of p such that word size does not get distorted under this map, and the word is accepted by M iff its image is trivial in π.

The moral is that, besides the extra fourth power and the assumed superadditivity, which seem like technical conditions, one gets a remarkably close connection between stopping times of general Turing machines and Dehn functions.

1.4 GROUP HOMOLOGY

Group homology and cohomology is a very basic tool that developed simultaneously with homological algebra. My impression is that early researchers were very highly motivated by Hopf's result that the cokernel of the two-dimensional Hurewicz homomorphism depends only on the fundamental group of the space. Nowadays, we write this as the Hopf exact sequence

$$\pi_2(X) \to H_2(X) \to H_2(\pi_1(X)) \to 0.$$

(Of course the first homology of a space depends only on its fundamental group; it is the abelianization.)

We shall avoid any discussion of homological algebra and work purely topologically. Obstruction theory very quickly leads to the following fact:

Theorem *For any countable group π, there is a space $B\pi$ whose fundamental group is p and whose universal cover is contractible. This space is unique up to pointed homotopy type. Moreover, for any homomorphism $\pi \to \pi'$ between groups, there is a unique pointed homotopy class of maps between spaces $B\pi \to B\pi'$ whose induced map on fundamental groups is the given homomorphism.*

This theorem embeds group theory within homotopy theory. As a result, any homotopy functor gives a functor on groups: so we can have homology and

[36]Essentially, the difference between a deterministic and nondeterministic Turing machine is not one of calculability, but one of speed; it measures the difference between discovering and verifying membership. Note that, in a group with an unsolvable word problem, not only is it hard to discover that a word is trivial, but it is also hard to verify that it is—the "certificate of triviality," for example, the product of relations one must use, can be noncomputably long.

cohomology, and more exotic things like K-theory and stable homotopy theory, whatever is useful.

Thus, we will write $H_i(\pi)$ for $H_i(B\pi)$ (and similarly for cohomology and coefficients, etc.). Some of the low groups have special interpretations.

Proposition *If A is abelian, then $H^1(\pi, A) = Hom(\pi, A)$. Also $H^2(\pi, A)$ is in a $1 : 1$ correspondence with central extensions of π by A, i.e., exact sequences*

$$0 \to A \to E \to \pi \to e.$$

(We denote the trivial group by e and the trivial abelian group by 0.)

The condition that the extension be central means that A lies in the center of E. Thus, for instance, $\mathbb{Z}_2$ has two central extensions by $\mathbb{Z}$: $\mathbb{Z} \times \mathbb{Z}_2$ and $\mathbb{Z}$ (mapping onto $\mathbb{Z}_2$); the infinite dihedral group $\mathbb{Z}_2 * \mathbb{Z}_2$ mapping onto $\mathbb{Z}_2$ is a noncentral extension by $\mathbb{Z}$.

If π is *perfect*, that is, $H_1(\pi) = 0$, meaning that π is its own commutator subgroup, then π has a universal central extension, that is, one that all others map through. For this extension, the center is $H_2(\pi)$, and it corresponds to the tautologous element of $H^2(\pi; H_2(\pi))$ that defines the Kronecker pairing from cohomology to the dual of homology.

Unraveling this a bit, under the assumption of perfection, we note that the universal coefficient theorem identifies $H^2(\pi; H_2(\pi)) = \mathrm{Hom}(H_2(\pi); H_2(\pi))$; our element is the one that corresponds to the identity. Let us give a few simple examples of group homology.

Example 1 *The trivial group. Here $B\pi$ is a point, and homology vanishes above dimension zero.*

Example 2 *Free groups. If π is F_k, then $B\pi$ is a wedge of k circles, with $H_i = 0$ for $i > 1$, $H_0 = \mathbb{Z}$, and $H_1 = \mathbb{Z}^k$. It is worth mentioning in this example that when $B\pi$ is a finite complex, it makes sense to discuss its Euler characteristic, which is a well-defined integer (as the Euler characteristic is a homotopy invariant of finite complexes), and computable from the homology with coefficients in any field.*

Example 3 *Free abelian groups. If $\pi = \mathbb{Z}^k$, then $B\pi$ is T^k, the k-torus. Now, the Künneth formula can be applied to show that the homology of $B\pi$ is torsion-free of rank the binomial coefficient $k!/a!(k-a)!$.*

Example 4 *Cyclic groups. For $\pi = \mathbb{Z}_k$, it is a little harder to see what $B\pi$ looks like. One approach is to consider S^{2n-1} as the unit sphere of $\mathbb{C}^n$ on which $\mathbb{Z}_k$ acts freely, just by multiplying each coordinate by a primitive root of unity. As n gets large these quotient spaces resemble $B\pi$ more and more closely, and we can use their (co)homology to compute the group (co)homology.*

In this case, it is not hard to find a cell complex for the quotients $S^{2n-1}/\mathbb{Z}_k$ by induction on n. For $n = 1$, the quotient is a circle. For larger n, one gets a cell decomposition with one cell in every dimension until $2n - 1$. Explicit calculation then leads to

$$\begin{aligned} &H_a(\mathbb{Z}_k) = \mathbb{Z}_k \quad H^a(\mathbb{Z}_k) = 0 \quad \text{with } a \text{ odd,} \\ &H_a(\mathbb{Z}_k) = 0 \quad\ \ H^a(\mathbb{Z}_k) = \mathbb{Z}_k \quad \text{with } a \text{ even} > 0. \end{aligned}$$

The Künneth formula can then be applied to give the calculation for abelian groups.

Remark. In general, one can build a $B\pi$ using a construction of Milnor. Consider the infinite join $\pi * \pi * \cdots$. It is contractible and has a free π action. The quotient is our desired space. (Recall that $X * Y$ is the space made up of lines connecting a point of X to a point of Y. The join of k discrete spaces is a wedge of $2k+1$ spheres, up to homotopy type.) The chain complex obtained in this way for the computation of group homology is called the "bar resolution." (The n-chains are the free abelian group of n-tuples of group elements whose product is the identity.)

Before proceeding to methods of calculation, it seems worth mentioning a couple of applications. Throughout the rest of the book we will be giving many more geometric, but more involved, examples.

The first is to groups that act freely and properly discontinuously on Euclidean space. If π is such a group, then so is any subgroup. Note that if π so acts, it has a $B\pi$ which is finite dimensional, so its homology vanishes in all sufficiently large dimensions. In particular, by example 4, π cannot have any nontrivial finite cyclic subgroups, that is, π is torsion-free.

In fact, similar reasoning shows that if π acts freely on $\mathbb{R}^n$ or even a contractible manifold, the cohomology with arbitrary coefficients vanishes in dimension greater than n. If the quotient is noncompact, then even in dimension n one would get vanishing. If the quotient were compact, then the integral cohomology would be $\mathbb{Z}$ in dimension n; in fact, one could see that $H^n(\pi; \mathbb{Z}\pi) = \mathbb{Z}$ and vanishes otherwise. This last condition turns out to imply that $B\pi$ (if a finite complex) satisfies Poincaré duality.

It is an important conjecture that the converse might hold, namely, that given such a π there exists (a unique) free cocompact action on some contractible X iff $H^n(\pi; \mathbb{Z}\pi) = \mathbb{Z}$ and vanishes otherwise. (In the appendix to section 2.3, we will see, following M. Davis, that X cannot always be taken to be Euclidean space.)

Here is another rather different application. Suppose we are interested in the existence of short exact sequences of the form

$$1 \to F_k \to E \to \mathbb{Z}_r \to 0,$$

where E is torsion-free. Note that the kernel of the composite surjection $F_s \to \mathbb{Z} \to \mathbb{Z}_r$ is isomorphic to $F_{r(s-1)+1}$. So a sufficient condition is that $r \mid (k-1)$.

Using homological ideas, one can show that the converse holds. The idea is that if E existed, we could consider BE. Let us suppose first that this is a finite complex. Then by covering space theory its r-fold cover would be homotopy equivalent to BF_k, a wedge of k circles. Since, by definition, the Euler characteristic is multiplicative in coverings, one would obtain that $1 - k = \chi(BF_k) = r\chi(BE)$, giving necessity. An actual proof goes like this. One shows that under these conditions, the chain complex of BE is still (chain equivalent to) a finite projective (over $\mathbb{Z}E$) chain complex. This allows one to use the Lefschetz fixed-point theorem for the $\mathbb{Z}_r$ action on BF. Then examination of the possible rational representations on H_1 forces the latter module to be a sum of one trivial summand and a number of copies of $\mathbb{Q}[e^{2\pi i/r}]$. This implies that $k \equiv 1 \bmod r$.

For our purposes, probably the most important calculational tool is the Mayer-Vietoris sequence of the amalgamated free product and HNN extension.

Theorem *If B is a subgroup of A and C (i.e., we are in the injective cases of the amalgamated free product and HNN extension) then*

$$B(A*_B C) = BA \cup_{BB} BC,$$
$$B(A*_B) = BA \cup_{BB\times\{0,1\}} BB \times [0,1].$$

Consequently, one obtains exact sequences

$$\cdots \to H_k(B) \to H_k(A) \times H_k(C) \to H_k(A *_B C) \to H_{k-1}(B) \to \cdots,$$

$$\cdots H_k(B) \to H_k(A) \to H_k(A*_B) \to H_{k-1}(B) \to \cdots,$$

where the map $H_k(B) \to H_k(A)$ in the second exact sequence is the difference of the two inclusions.

The proof of the theorem goes as follows. First, the right-hand sides have the right fundamental groups by van Kampen's theorem. One just wants to analyze the universal covers and see that they are contractible. This follows from the tree picture. The universal covers are made of copies of the universal covers of the BA's and BC's (which are themselves contractible) glued along universal covers of the BB's which are contractible. This guarantees the contractibility of the universal covers of $BA \cup_{BB} BC$ and of $BA \cup_{BB\times[0,1]} BB \times [0,1]$.

Example 5 *Consider the group $G = \langle a, b, c, d \mid [a, b] = [c, d] \rangle$. It is clearly of the form $F_2 *_Z F_2$ where the free groups are generated by $\langle a, b\rangle$ and $\langle c, d\rangle$. The relation is the amalgamation of a $\mathbb{Z}$ which is generated by the commutators. In this case the maps from the subgroup into the groups are trivial, so one gets the calculation that*

$$H_1(G) = \mathbb{Z}^4 \text{ and } H_2(G) = \mathbb{Z}.$$

The perspicacious reader probably noticed that this group is just the fundamental group of a surface of genus two, and we have computed the group *homology just by noticing that the surface is a BG! Therefore, it is worth noting that we would get the exact same calculation for group homology if we used* $[a, b]^2 = [c^2, d^3]$. *Indeed, for any words u, v in the commutator subgroups of* $\langle a, b\rangle$ *and* $\langle c, d\rangle$, *respectively, one would obtain for* $\langle a, b, c, d \mid u = v\rangle$ *the same homology. (However, these groups are not Poincaré duality groups, because they do not satisfy duality with respect to arbitrary coefficient modules.)*

A straightforward argument shows the following, which opens the way to applying spectral sequence techniques:

Proposition *If one has a short exact sequence of groups* $1 \to K \to E \to L \to 1$, *then there is a fibration* $BK \to BE \to BL$.

As a special case (where the fibration is a circle bundle, and the spectral sequence becomes the Gysin sequence), one has for a central $\mathbb{Z}$-extension $1 \to \mathbb{Z} \to E \to L \to 1$ the sequence

$$\cdots \to H^{k-2}(L) \to H^k(L) \to H^k(E) \to H^{k-1}(L) \to \cdots .$$

We have chosen to write this sequence in cohomology because there one can interpret the map $H^{k-2}(L) \to H^k(L)$ concretely as cup product with the Euler class of the circle bundle.

Example 6 *Let us compute the homology of the Heisenberg group H of* 3×3 *upper triangular unipotent matrices. We have an exact sequence extension* $1 \to \mathbb{Z} \to H \to \mathbb{Z}^2 \to 1$. *The Euler class is the generator of* $H^2(\mathbb{Z}^2)$. *One thus obtains (via the Gysin sequence) that*

$$H_1(H) = \mathbb{Z}^2,\ H_2(H) = \mathbb{Z}^2,\ \text{ and } H_3(H) = \mathbb{Z}.$$

Exercise

Write H as an HNN extension with $A = B = \mathbb{Z}^2$ and use a Mayer-Vietoris sequence to do the same calculation.

1.5 ARITHMETIC GROUPS

Arithmetic groups are groups that are defined similarly to $\mathrm{SL}_n(\mathbb{Z})$, the group of invertible matrices with determinant one. They arise naturally all over mathematics, and they have been studied from many points of view.

In this section, we will review a few special theorems regarding the homology of arithmetic groups that we will need in chapter 4.

Consider a subgroup G of GL_n defined by polynomial relations with coefficients in the rational numbers $\mathbb{Q}$. In other words, we shall assume that there is a set of polynomials in the entries of the matrices and $\det^{-1}$ that define the group

G. It makes sense to discuss the $\mathbb{F}$ points of G, denoted $G(\mathbb{F})$, for any field $\mathbb{F}$ of characteristic 0. By $G_{\mathbb{Z}}$, we mean $G(\mathbb{Q}) \cap \mathrm{GL}_n(\mathbb{Z})$. A discrete subgroup Γ of $G(\mathbb{Q})$ is called *arithmetic* if it is commensurable with $G_{\mathbb{Z}}$. Such subgroups are often *lattices*, that is, the natural metric on G/Γ has finite volume (see the theorem of Borel and Harish-Chandra below). For various reasons, it usually makes more sense to look at $K\backslash G/\Gamma$; for instance, it is an Eilenberg-MacLane space when Γ is torsion-free.

Note that given the real Lie group $G(\mathbb{R})$, there are many "$\mathbb{Q}$-forms," and these will give rise to different commensurability classes of arithmetic groups. For instance, we can define $O(n, 1)$ using any quadratic form over the rationals that has signature $(n, 1)$, and they will all give rise to arithmetic groups, but, unless these quadratic forms are homothetic over $\mathbb{Q}$, it is quite unlikely that these lattices will be commensurable. (A little below we will give a somewhat more general way of generating arithmetic lattices.)

We have already met a number of arithmetic groups: all finite groups, finitely generated abelian groups, finitely generated torsion free nilpotent groups (theorem of Malcev) such as the Heisenberg group, free groups (lie in $\mathrm{SL}_2(\mathbb{Z})$), surface groups. Given a quadratic form f in n variables over the rationals, then $\mathrm{SO}(n, f)$ defines a most interesting arithmetic group.

Remark. There is nothing sacred about $\mathbb{Q}$ in these definitions; using E, a finite extension of $\mathbb{Q}$, in its stead can be useful in defining more examples; in theory, this does not change the class of arithmetic groups, because if E is degree d over $\mathbb{Q}$, one can view $\mathrm{GL}_n(E)$ as a subgroup of $\mathrm{GL}_{nd}(\mathbb{Q})$. However, it is quite a bit simpler (and provides more insight) to write formulas using the general E rather than forcing them to be subgroups of $\mathrm{GL}_{nd}(\mathbb{Q})$.

Here is an important example that, among other things, shows the need for a slight modification of the definition of arithmetic. Let us consider an orthogonal group of the quadratic form

$$x_1^2 + x_2^2 + \cdots + x_n^2 - \sqrt{2}x_{n+1}^2$$

for $\mathbb{Q}[\sqrt{2}]$. There are two embeddings of $\mathbb{Q}[\sqrt{2}]$ in $\mathbb{R}$. Thus $O(x_1^2 + x_2^2 + \cdots + x_n^2 - \sqrt{2}x_{n+1}^2, \mathbb{Z}[\sqrt{2}])$ is a lattice in $O(n, 1) \times O(n+1)$(in the usual positive embedding, where $\sqrt{2}$ is positive, this quadratic form is of type $(n, 1)$; in the embedding where $\sqrt{2}$ is negative, the form is positive definite).

The image of this lattice in $O(n, 1)$ is a lattice there as well, because all that we are doing is projecting $O(n, 1) \times O(n+1) \to O(n, 1)$, which has a compact kernel. (The discreteness of the lattice means that we kill at most a finite normal subgroup of it when projecting.) The theorems we will explain presently show that the lattices just described are cocompact hyperbolic lattices. Torsion-free subgroups of finite index provide compact hyperbolic manifolds.

It is not at all obvious, but it is true, that arithmetic groups are finitely presented; they have a solvable word problem, are virtually torsion-free, and residually finite (this is a general fact about linear groups called Selberg's lemma).

Even the conjugacy problem is solvable in these groups (although not in all residually finite groups). However, as we saw above, the generalized word problem is usually not solvable in these groups.

Another remarkable property of these groups is that they are (Bieri-Eckmann) duality groups. This means that $H^k(\Gamma; \mathbb{Z}\Gamma)$ is nonzero for only one value of k. (This then implies a twisted Poincaré duality for Γ, where the twist is by that module.) This follows from the theory of Borel and Serre, which shows that the open manifold $K\backslash G/\Gamma$ can be compactified to a manifold with boundary, and their analysis of what the universal cover of its boundary looks like.

Recall that the *radical* of an algebraic group is its maximal connected algebraic solvable normal subgroup. The *(E-)rank* of $G(E)$ is the dimension of the maximal split torus (i.e., products of GL_1) defined over E that can be embedded in $G(E)$. G is *semisimple* if its radical is trivial.

Theorem (Borel and Harish-Chandra) *Let Γ be an arithmetic subgroup of G.*

1. *$G(\mathbb{R})/\Gamma$ has finite G-invariant volume iff there are no $\mathbb{Q}$-homomorphisms from the identity component of G to GL_1.*
2. *$G(\mathbb{R})/\Gamma$ is compact iff G has no subgroup isomorphic to GL_1 which is iff it has finite volume and every unipotent element of $G(\mathbb{Q})$ lies in its radical.*

For example, $\mathrm{GL}_m(\mathbb{Z})$ is not a lattice in $\mathrm{GL}_n(\mathbb{R})$ because of the homomorphism to $\mathbb{R}^*$ ($= \mathrm{GL}_1$) given by the determinant. Recall that unipotents are elements differing from the identity by a nilpotent. Condition 2 is equivalent to saying that the $\mathbb{Q}$-rank of G is 0.

We can use this theorem to check that the hyperbolic lattices produced above are actually cocompact. Since we are using embeddings of $\mathbb{Q}[\sqrt{2}]$ in $\mathbb{R}$, an element of the lattice is unipotent iff it is under either embedding. However, the embedding where $\sqrt{2}$ is negative gave rise to the orthogonal group, which has no nontrivial unipotents; a fortiori, neither does the lattice.

Example *For quadratic forms, (1) holds unless $n = 2$ and f represents 0 (i.e., there are nontrivial nullvectors in E), and (2) holds whenever the form is anisotropic (i.e., has no nullvectors).*

We shall confine the rest of our discussion to the semisimple case, and, as indicated above, we shall extend the definition of an arithmetic subgroup of a real Lie group H to be the image of an arithmetic group in a group G defined over $\mathbb{Q}$ under a Lie homomorphism from G onto an open subgroup of H which has a compact kernel.

Every semisimple group has a $\mathbb{Q}$-form that gives it arithmetic lattices. In fact, G contains both uniform (= cocompact) and nonuniform lattices.

Theorem *(Margulis's arithmeticity.) If $\mathbb{R}$-rank$(H) > 1$, then all irreducible lattices in H are arithmetic.*

In rank one, the existence of nonarithmetic lattices depends strongly on the Lie group and has been the object of intensive study. For instance, all the SO(n, 1)'s have nonarithmetic lattices, but Sp(n, 1) does not. It is unknown whether $U(n, 1)$ has such lattices when n is large.

Later, I will be interested in the cohomology of certain arithmetic lattices. While I cannot go into the details here, it seems worth just *mentioning* some of the ideas that have been brought to bear on this problem. Everything is a lot simpler in the uniform (i.e., cocompact) case, although, with more work, similar results can often be obtained for the nonuniform case.

The first key point is that for any semisimple group the coset space $K\backslash G$ with its right invariant metric has nonpositive curvature.[37] Consequently, it is a Euclidean space, and the manifold[38] $K\backslash G/\Gamma$ is a $B\Gamma$. Thus, the group homology is the study of the homology of this manifold. The following discussion is more straightforward if we assume that Γ is a uniform lattice, that is, that $K\backslash G/\Gamma$ is compact.

For convenience, we switch to cohomology and make use of an essentially tautological isomorphism:

$$H^*(\Gamma; \mathbb{C}) = H^*(K\backslash G/\Gamma; \mathbb{C}) = H^*(\mathfrak{g}, \mathfrak{k}; C^\infty(G/\Gamma))$$

where the last term is Lie algebra cohomology; the isomorphism is a consequence of a cochain complex isomorphism between the deRham model of the cohomology of $K\backslash G/\Gamma$ and the defining complex of relative Lie algebra cohomology with coefficients.

Now, one can show that the decomposition of $L^2(G/\Gamma)$ as a sum of irreducible representations,

$$L^2(G/\Gamma) \approx \bigoplus m(\pi, \Gamma) H_\pi$$

(where the m's are multiplicities and the H_π are the irreducible Hilbert space representations of G that are summands of the regular representation of G), gives one of $H^*(\mathfrak{g}, \mathfrak{k}; C^\infty(G/\Gamma))$ as well. That is,

$$H^*(\Gamma; \mathbb{C}) \approx \bigoplus m(\pi, \Gamma) H^*(\mathfrak{g}, \mathfrak{k}; H_\pi). \tag{2}$$

(This is essentially some kind of smoothing theorem, analogous to the fact that smooth singular homology and the continuous version are isomorphic.)

This result is the *Matsushima formula*. It gives a lot of useful information, including useful vanishing and nonvanishing theorems. A very useful result is that for lattices in semisimple groups, through some range linear in the $\mathbb{R}$-rank, the terms not coming from the trivial representation give vanishing contributions. This means that the cohomology (in some range) is independent of the lattice! (The cohomology associated with the trivial representation is isomorphic to that of the *compact dual* of $K\backslash G$.)

[37] We will discuss the elementary geometry necessary to follow this discussion in chapter 3. In any case, the trusting reader can just skip a sentence or so.

[38] Orbifold, if Γ has torsion.

I should emphasize that this is true only for coefficients with characteristic 0. For finite coefficients, one can see that the opposite is almost always true. (Hint: Think about lattices corresponding to p-Sylow subgroups of finite quotients of a given lattice.) It also fails strongly around the rank, as one can see for surfaces.

Remarks

Although our discussion assumed cocompactness, there are versions of the Matsushima formula and the vanishing theorems that are true for general lattices. Moreover, one can also generalize a great deal of the theory of arithmetic groups (including hard things like arithmeticity and cohomology calculations, although not the soft parts like finite generation!) to "S-arithmetic groups." These are lattices in products of real and p-adic fields, that is, groups like $\mathrm{SL}_n(\mathbb{Z}[1/k])$ for an integer k.

One of the most striking results of this development is that for any number field E (finite extension of $\mathbb{Q}$)

$$H^*(G(E)) = H^*_{\text{cont}}(G_\infty) \tag{3}$$

where the subscript "cont" means continuous cohomology and G_∞ refers to the copies of G at the infinite places. In other words, it looks as if equation (2) holds, but with no contributions of any of the other representations besides the trivial one!

The formula (3) is based on the fact that, by considering all of the E-points, one has essentially arranged for the rank of the "lattice" to be infinite.

A second useful method is L^2-cohomology. (The applications of this considerably transcend the study of lattices.) While the results are not quite precise, they give conclusions such as that, if $K\backslash G$ is odd dimensional, the Betti numbers of regular covers of a given lattice grow as o(volume) (i.e., sublinearly in the index of the cover), and in even dimensions, all but the middle cohomology groups do the same. On the other hand, the middle-dimensional groups do have ranks that are asymptotic to a multiple of the volume. (Some people even believe that this behavior is typical of residually finite groups that are fundamental groups of aspherical manifolds.) The drawback of this method is that it is hard to go from an L^2 calculation back to an ordinary calculation.

Later on, we will need sharp information about vanishing and nonvanishing of cohomology for negatively curved manifolds. This necessitates a look at lattices of $\mathbb{R}$-rank one. A very deep theorem of Clozel that gives sharp vanishing and nonvanishing results for a class of arithmetic lattices in $U(n, 1)$.

Theorem *For every n, there are complex hyperbolic n-manifolds*[39] *whose cohomology is nonzero in exactly the following dimensions:*

1. *there is a rank-one piece in every even dimension $0 \le d \le 2n$;*

[39] These are of complex dimension n, and thus of real dimension $2n$.

2. *for any divisor a of $n+1$ (less than $n+1$) there are elements in every second dimension between $n-a+1$ and $n+a-1$.*

It is worth making a couple of comments about this theorem. First, the relative Lie algebra cohomology is nonzero in the vanishing range here; the theorem is an arithmetic phenomenon, and, indeed, it is known that it fails for other lattices in $U(n,1)$. (Of course, what is happening is that the multiplicities occurring in the Matsushima formula are zero.)

Second, there are a large number of contributions to the proof of this theorem coming from deep number theory, à la Langlands' program. While I cannot say anything that really elucidates what is going on, I should probably mention that the "baby example" of these ideas is Deligne's proof of the Weil conjectures. These conjectures give an arithmetic method for computing the cohomology of smooth projective varieties. According to this work, the number of points on the variety over the various finite fields contains exactly the same information as the rational cohomology. The cohomology of $\mathbb{CP}^n$ "corresponds" to the number of points in $\mathbb{P}^n(\mathbb{F}_q)$ being $1+q+q^2+\cdots+q^n$.

In fact, $\mathbb{CP}^n$ is the compact dual of complex hyperbolic n-space, $U(n+1,1)/U(n)\times U(1)$. The classes accounted for in (1) are the classes coming from the trivial representation in the Matsushima formula, that is, the classes from the compact dual. Geometrically, the dual homology classes can be thought of as intersections of the complex hyperbolic manifold, of a smooth projective variety, with linear subspaces of $\mathbb{CP}^n$. The other classes are harder to account for, although their general placement symmetrically around the middle is Poincaré duality, their upward growth toward the middle is the Lefschetz theorem, and the nonvanishing in the middle can be seen, in even complex dimension, using the Hirzebruch signature theorem, and by the L^2-method.

1.6 REALIZATION OF SEQUENCES OF GROUPS AS GROUP HOMOLOGY

While we will not need the full depths of the following theorems, they are very interesting, and the special cases that we will need are not substantively simpler than the general case.

The basic issue we are interested in is the appearance of the sequence of homology groups of a finitely presented group. Given the Higman embedding theorem, it is perhaps not surprising that there is a strong logical component to this problem. On the other hand, the reader might be surprised to find that, for instance, there is a group G such that

1. for each a, $H_a(G)=0$ or $\mathbb{Z}$, and
2. $\{a \mid H_a(G)=0\}$ is neither c.e. nor the complement of a c.e. set.

We will also see that there is a finitely presented group G such that $H_a(G) = \mathbb{Z}_a$ or $H_a(G)$ is a sum of copies of $\mathbb{Z}$, where the number of copies is the ath digit of $\pi + e^2$.

In fact, there is an almost complete solution to this problem, but as of yet there does not seem to be one to the natural question of what cohomology algebra structures can exist. One has to be careful about the exact formulation of this question, because it is not yet known even for spaces. In fact, the question should be formulated in a way that explicitly compares what happens for groups to what happens for spaces.

There is a sense in which groups are no more special than general spaces:

Theorem *For any connected simplicial complex X, there is a group π and a map $f: B\pi \to X$, which is an isomorphism on homology. In fact, for any covering space of X, the map from the induced cover of $B\pi$ is also an isomorphism on homology.*[40] *Moreover, if X is a finite complex, $B\pi$ can also be taken to be a finite complex. If X is a countable complex, π can be taken countable. Moreover, if X is a c.e. space, then π is a c.e. group.*

We will return to the precise meaning of the c.e. group and c.e. space. Let us concentrate on the proof of the other parts of the theorem.

The construction of $B\pi \to X$ has two steps. The first is the construction of "n-simplices of groups." The second is merely the assembly of these according to the same data that one uses to assemble standard simplices to construct X.

To begin, one needs a nontrivial acyclic group A (that is, a group whose reduced homology vanishes in all dimensions). One can do so using an injection $F_4 \to F_2$ that looks like the projection on homology, and then producing an amalgamated free product $F_2 *_{F_4} F_2$, where the two injections of $F_4 \to F_2$ are such maps, just arranged to be projections to different factors.

Using A, we can easily build a 1-simplex of acyclic groups, using A for each of the two vertices and $A \times A$ for the group associated with the 1-simplex. (Note that here we map the group associated with a vertex into the group associated with an edge, the opposite of what we did with the graph of a group.) Using amalgamated free products and HNN extensions, one can assemble these groups to build a π, such that $B\pi \to G$ for any connected graph G, and by the Mayer-Vietoris exact sequences in section 1.4, this map is an isomorphism in homology (and, by the exact same argument, the same holds for covers). This construction proves the theorem for graphs.

Now, to do two-dimensional complexes, one needs to construct a 2-simplex. That is, we need an acyclic group B that contains the result of applying the construction to a circle, thought of as a triangle, that is, as the boundary of a 2-simplex (and similarly in higher dimensions.)

[40] This notion can be most succinctly described in terms of the "plus construction," which will be explained in chapter 2.

We shall not give a construction of these simplices (all such contructions that I know about are somwhat tricky), and shall instead, rely on the paper [BDH] mentioned in the notes section. In any case, I hope the idea is clear.

Now, let us move on to the notions of a c.e. group and c.e. space. A c.e. group is a group with generators $x_1, x_2, x_3, \ldots$ and a set of relations that is a c.e. set. In other words, there is a Turing machine that constructs the relations. Note that it is entirely equivalent to ask that the set of relations that hold be c.e. or to give a c.e. generating set for these relations. To define the notion of a c.e. space, we will think of the vertices as being the integers (or a finite set of them). We can think of the simplices as being $(n + 1)$tuples of vertices, which can be encoded by natural numbers. So a simplicial complex is just some set of tuples of natural numbers (with the additional property of being closed under inclusion). We shall suppose that our complexes are effectively connected.

Note that the homology groups of a c.e. space are actually c.e. abelian groups. (Hint: First check that a c.e. abelian group, up to computable isomorphism, is equivalent to a c.e. sequence of finitely presented abelian groups with (c.e.) homomorphisms from one group to the next.)

We leave it to the reader to check that the above constructions produce c.e. groups from c.e. spaces.

The space $B\pi$ for π finitely presented (or even c.e.) groups is actually a c.e. simplicial complex, as one can see by going carefully through the Milnor construction. This suggests the following:

Conjecture *For X an effective simplicial complex, whose 2-skeleton is finite (up to homotopy), there is a finitely presented π, and a map $B\pi \to X$, which is an isomorphism on homology.*

This would lead to a characterization of the sequences of groups that can be the homology groups of a finitely presented group. They would be the c.e. sequences of c.e. abelian groups whose first two groups are finitely generated. The following theorem implies something that is quite close.

Theorem *If X is any c.e. simplicial complex, then there is a finitely presented π, and a map $f : B\pi \to \Sigma^2 X$ to the second suspension of X, which is an isomorphism on homology.*

First we find a c.e. group π that resembles X. Then we can embed π in a universal acyclic group U,[41] and form $\pi' = U *_\pi U$, which resembles the suspension of X. This is now a finitely generated c.e. group, which we will denote π'. π' also embeds in U. $U *_{\pi'} U$ resembles the second suspension, and is also finitely presented, proving the theorem.

Corollary *For $n > 3$, for any c.e. abelian group A, there is a f.p. group π such that $H_a(\pi) = A$ for $a = n$, and is 0 otherwise.*

[41] Recall that a universal group is a finitely presented group that contains all others. Baumslag, Dyer, and Miller constructed acyclic universal groups.

Proof. A is a limit of a c.e. sequence of finitely presented abelian groups. Thus, one can form the c.e. sequence of Moore spaces and maps to produce a c.e. space $M(A, n-2)$. Setting $X =$ the limit of these spaces and applying the theorem gives the result.

Remark. To put some more flesh onto the above proof, we should make a few simple remarks. Recall that a Moore space of type (A, k) is a simply connected space all of whose homology groups vanish except for H_k, and $H_k = A$. They exist for $k > 1$, and are unique up to homotopy equivalence. For any homomorphism between A and B, there is a map from $M(A, k) \to M(B, k)$ inducing this homomorphism. (The homotopy class of maps is not unique, however.)

Note also that one can interpolate between two triangulations of a polyhedron P by a triangulation of $P \times [0, 1]$. Arbitrary homotopy classes can also be realized by simplicial maps, which allows one to build a c.e. space from the c.e. sequence of homotopy types and (the not quite well defined) sequence of homotopy classes of maps.

Corollary *If A_i is any c.e. sequence of c.e. abelian groups, such that A_1 and A_2 are finitely generated and A_3 is "untangled" in the sense of [BDM] (see the notes), then there is a finitely presented group with the A's as its homology sequence.*

An abelian group is "untangled" if it has a presentation with a c.e. basis for its relations. This condition is not necessary, and thwarts a complete characterization of the homology sequence of f.p. groups.

Proof. To realize the A_k, $k > 3$, one simply uses as X the wedge of the Moore spaces discussed in the previous proof. We can then take the free product with a group realizing the first three groups from [BDM] (and 0 above dimension three), to obtain our desired π.

NOTES

The elementary topology of fundamental groups, covering spaces, and van Kampen's theorem is all very nicely explained in

W. Massey. *Algebraic Topology: An Introduction*. Reprint of the 1967 edition. Graduate Texts in Mathematics 56. Springer-Verlag, New York, 1977.

A good introduction to basic combinatorial group theory and to the theory of group actions on trees can be found in

P. Scott and C.T.C. Wall. *Topological methods in group theory. In Homological Group Theory*. Proceedings of the Symposium (Durham, 1977), 137–203. London Mathathematical Society Lecture Note Series 36. Cambridge University Press, Cambridge, 1979.