A course on geometric group theory.

Brian H. Bowditch
School of Mathematics, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
http://www.maths.soton.ac.uk/staff/Bowditch/
and
Department of Mathematical and Computing Science
Tokyo Institute of Technology
2-12-1, O-okayama, Meguro-ku, Tokyo, Japan.

[October 2005]

Preface.

These notes are based on a series of lectures I gave at the Tokyo Institute of Technology from April to July 2005. They constituted a course entitled "An introduction to geometric group theory" totalling about 20 hours. The audience consisted of fourth year students, graduate students as well as several staff members. I therefore tried to present a logically coherent introduction to the subject, tailored to the background of the students, as well as including a number of diversions into more sophisticated applications of these ideas. There are many statements left as exercises. I believe that those essential to the logical developments will be fairly routine. Those related to examples or diversions may be more challenging.

The notes assume a basic knowledge of group theory, and metric and topological spaces. We describe some of the fundamental notions of geometric group theory, such as quasi-isometries, and aim for a basic overview of hyperbolic groups. We describe group presentations from first principles. We give an outline description of fundamental groups and covering spaces, sufficient to allow us to illustrate various results with more explicit examples. We also give a crash course on hyperbolic geometry. Again the presentation is rather informal, and aimed at providing a source of examples of hyperbolic groups. This is not logically essential to most of what follows. In principle, the basic theory of hyperbolic groups can be developed with no reference to hyperbolic geometry, but interesting examples would be rather sparse.

In order not to interupt the exposition, I have not given references in the main text. We give sources and background material as notes in the final section.

I am very grateful for the generous support offered by the Tokyo Insititute of Technology, which allowed me to complete these notes, as well as giving me the freedom to pursue my own research interests. I am indebted to Sadayoshi Kojima for his invitation to spend six months there, and for many interesting conversations. I thank Toshiko Higashi for her constant help in making my stay a very comfortable and enjoyable one. My then PhD student Ken Shackleton accompanied me on my visit, and provided some tutorial assistance. Shigeru Mizushima and Hiroshi Ooyama helped with some matters of translatation etc.

0. Introduction.

In 1872, Klein proposed group theory as a means of formulating and understanding geometrical constructions. The resulting programme has been termed the "Erlingen programme". Since that time the two subjects have been closely linked. The subject of geometric group theory might be viewed as Klein's programme in reverse — geometrical ideas are used to give new insights into group theory. Although largely a creation of the last twenty years or so, its anticedents can be traced back to the early 20th century. For example, Dehn used hyperbolic geometry to solve the word problem in a surface group. His ideas were subsequently formalised in terms of "small cancellation theory", in some sense a forerunner of modern geometric group theory (while remaning an active field in itself). The observation, due to Efremovich, Schwarz and Milnor, that a group acting discretely compactly on a proper space resembles, on a large scale, the space on which they act, is key to the development of the subject.

The subject draws on ideas from across mathematics, though one can identify two particular sources of inspiration. One is low-dimensional topology, in particular 3-manifold theory. Another is hyperbolic geometry. The work of Thurston in the late 1970s showed that these two subjects were intimately linked. The resulting flurry of activity might be seen as the birth of geometric group theory as a subject in its own right. The work of Gromov in the 1980s was particularly influential. We note especially his papers on hyperbolic groups and asymptotic invariants.

The subject has now grown into a major field. It would impossible to give even a representative overview in notes such as these. I have directed attention towards giving a basic introduction to hyperbolic groups and spaces. These are of fundamental importance, though of course, many other directions would have been possible. (We say almost nothing, for example, about the vast subject of non-positively curved spaces.) Similarly, many of the diversions are informed by my own intersts, and I offer my apologies if these have been over-represented.

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.

$G \le \Gamma$: G is a subgroup of Γ.

$G \lhd \Gamma$: 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 these 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 $\mathbf{N}$, $\mathbf{Z}$, $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{C}$ respectively for the natural numbers (including 0), the intergers, and the rational, real and complex numbers.

We shall generally view $\mathbf{Z}^2$ 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 $\langle \{a_1, \ldots, a_n\} \rangle$).

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 notion, 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,
$\mathbf{Q}$ (exercise), $\mathbf{R}$, $GL(n, \mathbf{R})$, etc.

Execises: (1) If $N \lhd \Gamma$ and Γ is f.g., then Γ/N is f.g.
(2) If $N \lhd \Gamma$, 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 hypthothesis, 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 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 $[a \mapsto \bar{a}]$ this 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, $w = w_0, w_1, \dots, w_n = w'$, 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 $[a_1 \dots a_n]^{-1} = [\bar{a}_n \dots \bar{a}_1]$.)
(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 $a_1 \dots a_n$ 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 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.

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), wheras 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 $S = \{a^n b a^{-n} | n \in \mathbf{N}\} \subseteq F_2$. 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) $\langle\langle A \rangle\rangle \triangleleft G$,

(3) If $N \triangleleft G$, $A \subseteq N$, 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 $\langle \{x_1, \ldots, x_n\} \mid \{r_1, \ldots, r_m\} \rangle$ 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 $F_2 = \langle a, b | \rangle$ 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 consdering 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.

2. Cayley graphs.

A Cayley graph gives us a means by which a finitely generated group can be viewed as a geometric object. The starting point is a finite generating set. The dependence of the construction on the choice of generating set will be discussed in Section 3.

Many combinatorial constructions can be interpreted geometrically or topologically, and this often results in the most efficient means of proof. This is the main theme of this course. Nevertheless, the earlier combinatorials tools developed by Higman, Neumann etc, still remain a powerful resource, and a rich source of interesting examples.

2.1. Basic terminology and notation.

Let K be a graph. Formally this is thought of as a set, $V(K)$, of vertices together with a set $E(K)$ of edges. It will be convenient to allow multiple edges (edges connecting the same pair of vertices) and loops (edges starting and finishing at the same vertex). We recall some standard terminology from graph theory:

Definition : A (combinatorial) *path* consists of a sequence of edges with consecutive edges adjacent.

An *arc* is an embedded path.

A *cycle* is a closed path.

A *circuit* is an embedded cycle.

A graph is *connected* if every pair of vertices are connected by a path (and hence also by an arc).

The *valence* of a vertex is the number of incident edges (counting muliplicities of multiple edges, and counding each loop twice).

A graph is *locally finite* if each vertex has finite valence.

It is *n-regular* if every vertex has valence n.

We will often write a combinatorial path as a sequence of vertices rather than edges, though if there are loops or multiple edges one also needs to specify the edges connecting them.

2.2. Graphs viewed as metric and topological spaces.

Let K be a graph. We can realise K as a 1-complex. Each edge of K corresponds to a copy of the unit interval with vertices at its endpoints. (A loop gets bent around into a circle.) This gives a description of the 1-complex as a set, denoted $|K|$. We can define a distance on $|K|$ as follows. We first fix a parameterisation for each edge. This defines a length for each interval in an edge, so that the whole interval has length 1. We can generalise the notion of a path allowing it to start and finish in the interiors of edges as well as at a vertex. A path, π, thus has a well defined *length*, $l(\pi) \in [0, \infty)$. Given $x, y \in K$ we defined $d(x, y) \in [0, \infty]$ to be the minimum length of an connecting x to y. We set this equal to ∞ if there is no such path. Otherwise, the minimum is attained, since the set of lengths of such arcs is discrete. If K is connected, then $(|K|, d)$ is a metric space. It thus induces a topology on $|K|$. This topology makes sense even if K is not connected. The space $|K|$ might be called the *realisation* of K. We note (exercise):

$|K|$ is compact $\Leftrightarrow$ K is finite.

$|K|$ is locally compact $\Leftrightarrow$ K is locally finite.

$|K|$ is (topologically) connected $\Leftrightarrow$ K is connected.

We frequently omit the $|.|$, and simply write $|K|$ as K. Thus, a graph K can be viewed in three ways: as a combinatorial object, as a geometric object (metric space) or as a topological object.

Remark : For most purposes here, we will only be interested in locally finite graphs. In this case, the topology described here is the only "sensible" one. In the non locally finite case, however, there are other natural topologies such as the "CW-topology" which is different. We will not need to worry about these issues in this course.

Suppose that a group Γ acts on the graph K. It then acts by isometries on its realisation. We say that the action on K is *free* if it is free on both $V(K)$ and $E(K)$. This means that no element of Γ fixes a vertex or inverts an edge (i.e. swaps the vertices of an edge). In this case, we can form the quotient graph K/Γ.

Simple example : We can view the real line, $\mathbf{R}$, as the (realisation of) a graph with vertices at the integers. The group $\mathbf{Z}$ acts by translation. $\mathbf{R}/\mathbf{Z}$ consists of a single loop. $\mathbf{R}/n\mathbf{Z}$ is a circuit with n vertices. Both are topological circles.

Definition : A *tree* is a connected graph with no circuits.

Thus (exercise) a graph is a tree if and only every pair of vertices are connected by a unique arc. Also (exercise) any two n-regular trees are isomorphic. The 2-regular tree is the real line described above.

2.3. Graphs associated to groups.

Let Γ be a group, and $S \subseteq \Gamma$ a subset not containing the identity. Let $\bar{S}$ be a set of "formal inverses" of elements of S (as in the construction of a free group) and let $A = S \sqcup \bar{S}$. If $S \cap S^{-1} = \emptyset$, then we can identify A as a subset of Γ, by identifying the formal inverse $\bar{a}$ of an element $a \in S$, with its actual inverse, a^{-1}, in Γ. Indeed, we can usually arrange that $S \cap S^{-1} = \emptyset$: if both a and a^{-1} lie in S, then just throw one of them away. This is only a problem if $a = a^{-1}$. In this case, we need to view the formal inverse of a as distinct from a. This is a somewhat technical point, and we will generally write a^{-1} for the inverse whether used formally or as a group element. Only if $a^2 = 1$ do we need to make this distinction.

We construct a graph $\Delta = \Delta(\Gamma; S)$ as follows. Let $V(\Delta) = \Gamma$. We connect vertices $g, h \in \Gamma$ by and edge in Δ if $g^{-1}h \in A$. In other words, for each $g \in \Gamma$ and $a \in A$ we have an edge connected g to ga. We imagine the directed edge from g to ga as being labelled by the element a. The same edge with the opposite orientation is thus labelled by a^{-1}. (If it happens that $a^2 = 1$, then we connect g to ga by a pair of edges, labelled by a and its formal inverse.) Note that Δ is $|A|$-regular. It is locally finite if and only if $|S|$ is finite. (Note $|A| = 2|S|$.)

Now Γ acts on $\Gamma = V(\Delta)$ by left multiplicaton, and this extends to an action on Δ: if $g, h, k \in \Gamma$, with $g^{-1}h \in A$ then $(kg)^{-1}(kh) = g^{-1}h \in A$. This action is free and preserves the labelling. The quotient graph consists of a single vertex and $|S|$ loops. Topologically this is a wedge of $|S|$ circles — a space we will meet again in relation to free groups in Section 4.

Let $W(A)$ be the set of words in A. There is a natural bijection between $W(A)$ and the set of paths in Δ starting from the identity. More precisely, if $w = a_1 a_2 \ldots a_n \in W(A)$, then there is a unique path $\pi(w)$ starting at $1 \in V(\Delta)$ so that the ith directed edge of $\pi(w)$ is labelled by a_i. We see, inductively, that the final vertex of $\pi(w)$ is the group element obtained by multiplying together the a_i in Γ. Thus, by interpreting a word in the generators a group element, we retain only the final destination point in Δ, and forget about how we arrived there. We write $p : W(A) \longrightarrow \Gamma$ for the map obtained in this way. Thus $\pi(w)$ a path from 1 to $p(w)$. We can similarly start from any group element $g \in \Gamma$, and get a path from g to $gp(w)$. It is precisely the image, $g\pi(w)$, of $\pi(w)$ by g in the above group action.

Lemma 2.1 : $\Gamma = \langle S \rangle \Leftrightarrow \Delta(\Gamma; S)$ *is connected.*

This follows from the above construction and the observation that $\langle S \rangle$ is precisely the set of elements expressible as a word in the alphabet A. (Note that, in general, $\langle S \rangle$ is the set of vertices of the connected component of Δ containing 1.)

Definition : If $S \subseteq \Gamma$ is a generating set of Γ, then $\Delta(\Gamma; S)$ is the *Cayley graph* of Γ with respect to S.

Note that, in summary, we have seen that any finitely generated group acts freely on a connected locally finite graph. Converses to this statement will be discussed in Section 3.

Examples.
(0) The Cayley graph of the trivial group with respect to the empty generating set is just a point.
(1a) $\mathbf{Z} = \langle a \rangle$. In this case Δ is the real line (Figure 2a). $\mathbf{Z}$ acts on it by translation with quotient graph a circle: a single vertex and a single loop.

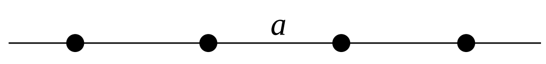

Figure 2a.

(b) If we add a generator $\mathbf{Z} = \langle a, a^2 \rangle$ we get an infinite "ladder" (Figure 2b).

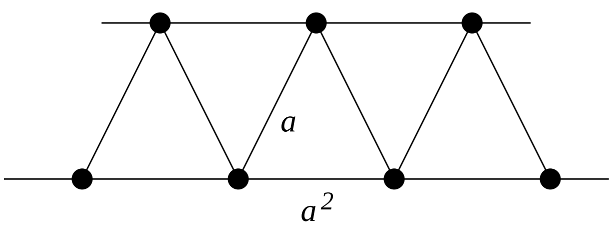

Figure 2b.

(c) $\mathbf{Z} = \langle a^2, a^3 \rangle$, see Figure 2c.

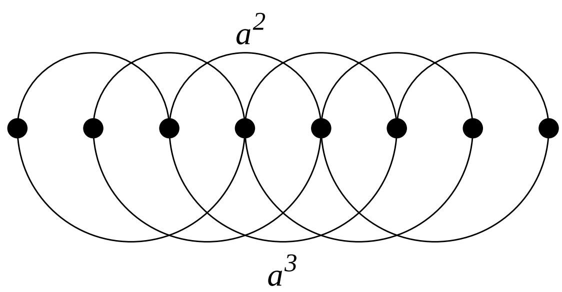

Figure 2c.

In these examples, the graphs are all combinatorially different, but they all "look like" the real line from "far away": in a sense that will be made precise later (see Section 3).

(2) $\mathbf{Z}_n = \langle a \mid a^n = 1 \rangle$. This is a circuit of length n (see Figure 2d, where $n = 5$).

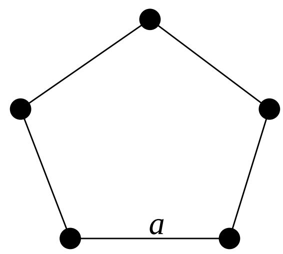

Figure 2d.

(3) $\mathbf{Z} \oplus \mathbf{Z} = \langle a, b \mid ab = ba \rangle$. Here Δ is the 1-skeleton of the square tessellation of the plane, $\mathbf{R}^2$ (Figure 2e). It "looks like" $\mathbf{R}^2$ from "far away".

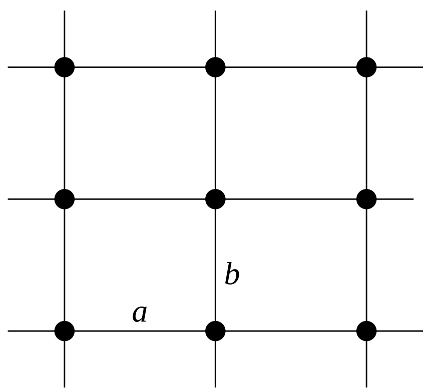

Figure 2e.

(4) The dihedral group $\langle a, b \mid b^2 = baba = a^n = 1 \rangle$
and the infinite dihedral group $\langle a, b \mid b^2 = baba = 1 \rangle$ (exercise).

(5) The icosahedral group $\langle a, b, c \mid a^2 = b^2 = c^2 = (ab)^3 = (ca)^5 = (bc)^2 = 1 \rangle$ (exercise).

Note that in a Cayley graph, a word representing the identity gives a path starting and finishing at 1, in other words a cycle through 1. In particular, the relators give rise to cycles. For example, $aba^{-1}b^{-1}$ gives us the boundary of a square tile in example (3) above. Since a free group has no relations, the following should be no surprise:

Theorem 2.2 : *Suppose F is freely generated by $S \subseteq F$, then $\Delta(F;S)$ is a tree.*

Indeed the converse holds provided we assume, in that case, that $S \cap S^{-1} = \emptyset$.

Before we begin, we make the observation that a word w corresponds to a path $\pi = \pi(w)$, then any subword of the form aa^{-1} would mean that we cross an edge labelled a, and then immediately cross back again, in other words, we *backtrack* along this edge. Cancelling this word corresponds to eliminating this backtracking. A word is therefore reduced if and only if the corresponding path has no backtracking.

Proof of Theorem 2.2 : We can assume that F has the form $F(S)$ as in the construction of free groups. We want to show that Δ has no circuits. We could always translate such a circuit under the action of F so that it passed through 1, and so can be though of a path starting and ending at 1. Now such a circuit, σ, corresponds to a word in $S \cup S^{-1}$ representing the identity element in F. Thus, there is a finite sequence of reductions and inverse reductions that eventually transform this word to the empty word. Reinterpreting this in terms of paths in the Cayley graph, we get a sequence of cycles, $\sigma = \sigma_0, \ldots, \sigma_n$, where σ_n is just the constant path based at 1, and each σ_i is obtained from σ_{i-1} by either eliminating or introducing a backtracking along an edge.

Now, given any cycle, τ, in Δ, let $O(\tau) \subseteq E(\Delta)$ be the set of edges through which τ passes an odd number of times It is clear that $O(\tau)$ remains unchanged after eliminating a backtracking from τ. In particular, $O(\sigma_i)$ remains constant throughout. But $O(\sigma) = E(\sigma)$ since σ is a circuit, and $O(\sigma_n) = \emptyset$. Thus, $E(\sigma) = \emptyset$, and so σ could only have been the constant path at 1. $\diamondsuit$

Remark : We are really observing that the $\mathbf{Z}_2$-homology class of a cycle in $H_1(\Delta; \mathbf{Z}_2)$ remains unchanged under cancellation of backtracking, and that the $\mathbf{Z}_2$-homology class of a non-trivial circuit is non-trivial.

This is the main direction of interest to us. The converse is based on the observation that in a tree, any path with no backtracking is an arc. Moreover there is only one arc connecting 1 to any given element $g \in F$. This tells us that each element, g, of F has a unique representative as a reduced word. If we have a map, ϕ, from S into any group Γ, we can use such a reduced word to define an element, $\hat{\phi}(g)$, in Γ by multiplying together the ϕ-images of the letters in our reduced word. We need to check that $\hat{\phi}$ is a homomorphism from F into Γ, and that there was no choice in its definition. This then shows that S is, by definition, a free generating set. The details of this are left as an execise.

Note that putting together Theorem 2.2 with the above observation on arcs in trees, we obtain the uniqueness part of Proposition 1.3: every element in a free group has a unique representative as a reduced word in the generators.

A picture of the Cayley graph of the free group with two generators, a and b, is given in Figure 2f. To represent it on the page, we have distorted distances. In reality, all edges have length 1.

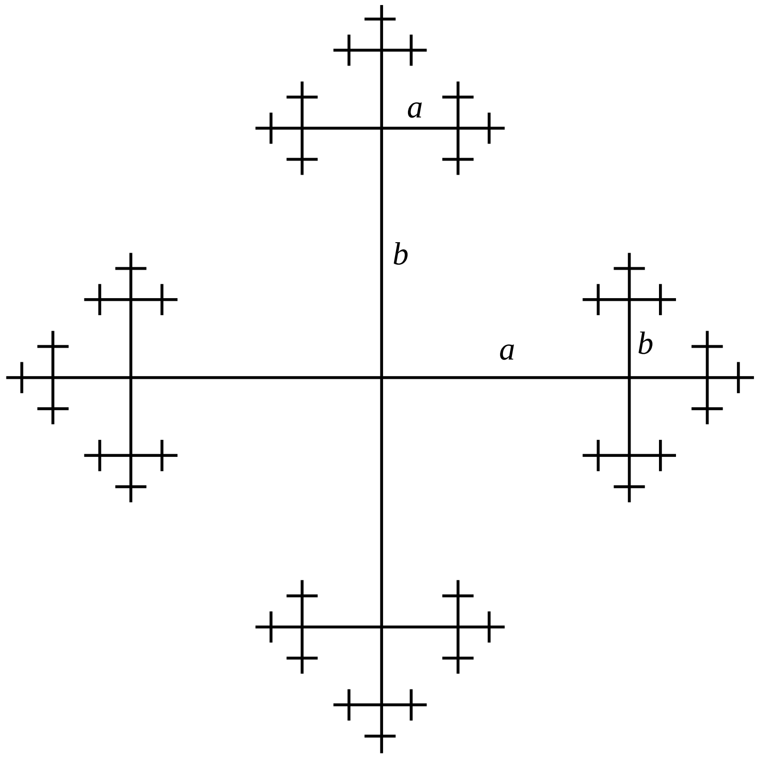

Figure 2f.

3. Quasi-isometries.

In this section we will define the notion of a "quasi-isometry" — one of the fundamental notions in geometric group theory. First, though, we need to describe some more general terminology, and make a few technical observations. Many of the technical details mentioned below can be forgotten about in the main cases of interest us, where the statements will be apparent. However, we might as well state them in general.

3.1. Metric Spaces.

Let (M, d) be a metric space.

Notation. Given $x \in r$ and $r \geq 0$, write $N(x, r) = \{y \in M \mid d(x, y) \leq r\}$ for the closed *r-neighbourhood* of x in M. If $Q \subseteq M$, write $N(Q, r) = \bigcup_{x \in Q} N(x, r)$. We say that Q is *r-dense* in M if $M = N(Q, r)$. We say that Q is *cobounded* if it is r-dense for some $r \geq 0$. Write $\mathrm{diam}(Q) = \sup\{d(x, y) \mid x, y \in Q\}$ for the *diameter* of Q. We say that Q is *bounded* if $\mathrm{diam}(Q) < \infty$. Note that any compact set is bounded.

Definition : Let $I \subseteq \mathbf{R}$ be an interval. A *(unit speed) geodesic* is a path $\gamma : I \longrightarrow M$ such that $d(\gamma(t), \gamma(u)) = |t - u|$ for all $t, u \in M$.

(Sometimes, we may talk about about a *constant speed* geodesic, where $d(\gamma(t), \gamma(u)) = \lambda|t - u|$ for some constant "speed" $\lambda \geq 0$.)

Note that a geodesic is an arc, i.e. injective (unless it has zero speed).

Warning: This terminology differs slightly from that commonly used in riemannian geometry. There a "geodesic" is a path satisfying the geodesic equation. This is equivalent to being *locally* geodesic of consant speed in our sense.

Suppose $\gamma : [a, b] \longrightarrow M$ is any path. We can define its *length* as

$$\sup \left\{ \sum_{i=1}^{n} d(t_{i-1}, t_i) \mid a = t_0 < t_1 < \cdots < t_n = b \right\}.$$

If $-\infty < a \leq b < \infty$, we say that γ is *rectifiable* if its length is finite. In general we say that a path is rectifiable if its restriction to any finite subinterval is rectifiable. There certainly exist non-rectifiable paths e.g. the "snowflake" curve. However, all the paths we deal with in this course will be sufficiently nice that this will not be an issue.

A (slightly technical) exercise shows that $\text{length}(\gamma) = d(\gamma(a), \gamma(b))$ if and only if $d(\gamma(a), \gamma(b)) = d(\gamma(a), \gamma(t)) + d(\gamma(t), \gamma(b))$ for all $t \in [a, b]$. If γ is also injective (i.e. an arc), then we can reparametrise as γ as follows. Define $s : [a, b] \longrightarrow [0, d(a, b)]$ by $s(t) = d(\gamma(a), \gamma(t))$. Thus s is a homeomorphism. Now, $\gamma' = \gamma \circ s^{-1} : [0, d(a, b)] \longrightarrow M$ is a geodesic. Indeed this gives us another description of a geodesic up to parameterisation, namely as an arc whose length equals the distance between its endpoints.

Exercise: If $\gamma : I \longrightarrow M$ is any rectifiable path then we can find a paramerisation so that γ has unit speed, i.e. for all $t < u \in I$, the length of the subpath $\gamma|[t, u]$ between t and u has length $u - t$. (This time our map s might not be injective if γ stops for a while.)

In any case, the above observation should be clear in cases where we actually need it.

We will sometimes abuse notation and write $\gamma \subseteq M$ for the image of γ in M — even if γ is not injective.

Definition : A metric space (M, d) is a *geodesic space* (sometimes called a *length space*) if every pair of points are connected by a geodesic.

Such a geodesic need not in general be unique.

Examples.

(1) Graphs with unit edge lengths. (Essentially from the definition of the metric.)

(2) $\mathbf{R}^n$ with the euclidean metric: $d(\underline{x}, \underline{y}) = \sqrt{\sum_{i=1}^{n} |x_i - y_i|^2}$, and

(3) Any convex subset of $\mathbf{R}^n$.

(4) Hyperbolic space, $\mathbf{H}^n$, and any convex subset thereof (see later).

(5) In fact, any complete riemannian manifold (from the definition of the metric and the Hopf-Rinow theorem). We won't be needing this formally in this course.

Non examples.

(1) Any non-connected space.

(2) $\mathbf{R}^2 \setminus \{(0,0)\}$: there is no geodesic connecting $\underline{x}$ to $-\underline{x}$.

Indeed any non-convex subset of $\mathbf{R}^n$ with the euclidean metric.

(3) Define a distance on the real line, $\mathbf{R}$, by setting $d(x,y) = |x-y|^p$ for some constant p. This is a metric if $0 < p \leq 1$, but $(\mathbf{R}, d)$ is a geodesic space only if $p = 1$ (exercise).

(4) If we allow different edge lengths on a locally infinite graph, the result might not be a geodesic space. For example, connect two vertices x, y by infinitely many edges (e_n) where n varies over $\mathbf{N}$, and assign e_n a length $1 + \frac{1}{n}$. Thus $d(x,y) = 1$, but there is no geodesic connecting x to y.

Definition : A metric space (M, d) is *proper* if it is complete and locally compact.

Proposition 3.1 : *If (M,d) is a proper geodesic space then $N(x,r)$ is compact for all $r \geq 0$.*

One way to see this is to fix x, and consider the set

$$A = \{r \in [0, \infty) \mid N(x,r) \text{ is compact}\}.$$

If $A \neq [0, \infty)$ one can derive a contradiction by considering $\sup(A)$. We leave the details for the reader.

In practice, the conclusion of Proposition 3.1 will be clear in all the cases of interest to us here: euclidean space, locally finite graphs etc., so the technical details need not worry us.

Suppose that M is proper and that $Q \subseteq M$ is closed. Given $x, y \in Q$, let $d_Q(x,y)$ be the miminum of the lengths of rectifiable paths connecting x to y. This is ∞ if there is no such path. Another technical exercise, using Proposition 3.1, shows that the minimum is attained. Again, this is apparent in the cases of interest to us.

If d_Q is always finite, then (Q, d_Q) is a geodesic metric space. We refer to d_Q as the *induced path metric*. Clearly $d(x,y) \leq d_Q(x,y)$. In cases of interest (Q, d_Q) will have the same topology as (Q, d), though one can concoct examples where its topology is strictly finer.

3.2. Isometries.

Let (X, d) and (X', d') be metric spaces. A map $\phi : X \longrightarrow X'$ is an *isometric embedding* if $d'(\phi(x), \phi(y)) = d(x,y)$ for all $x, y \in X$. It is an *isometry* if it is also surjective. Two spaces are *isometric* if there is an isometry between them.

The set of self-isometries of a metric space, X, forms a group under composition — the *isometry group* of X, denoted $\mathrm{Isom}(X)$.

For example, the isometries of euclidean space $\mathbf{R}^n$ are precisely the maps of the form $[\underline{x} \mapsto A\underline{x} + \underline{b}]$, where $A \in O(n)$ and $\underline{b} \in \mathbf{R}^n$.

Let X be a proper length space. Suppose that Γ acts on X by isometry. Given $x \in X$, we write $\Gamma x = \{gx \mid g \in \Gamma\}$ for the *orbit* of x under Γ, and $\mathrm{stab}(x) = \{g \in \Gamma \mid gx = x\}$ for its *stabiliser*. The action is *free* if the stabiliser of every point is trivial. (This has nothing to with "free" in the sense of "free groups" or of "free abelian groups"!)

Definition : We say that the action on X is *properly discontinuous* if for all $r \geq 0$ and all $x \in X$, the set $\{g \in \Gamma \mid d(x, gx)\}$ is finite.

Using Proposition 3.1, we can express this without explicit mention of the metric: it is equivalent to the statement that $\{g \in \Gamma \mid gK \cap K \neq \emptyset\}$ is finite for all compact $K \subseteq M$.

If the action is properly discontinuous, then the quotient, X/Γ, is hausdorff (and complete and locally compact). Indeed we can define a metric, d' on X/Γ by setting $d'(\Gamma x, \Gamma y) = \min\{d(p, q) \mid p \in \Gamma x, q \in \Gamma y\} = \min\{d(x, gy) \mid g \in \Gamma\}$.
Exercise: this is a metric, and it induces the quotient topology on X/Γ.

Definition : A properly discontinuous action is *cocompact* if X/Γ is compact.

Exercise: The following are equivalent:
(1) The action is cocompact,
(2) Some orbit is cobounded,
(3) Every orbit is cobounded.

We will frequently abbreviate "properly discontinuous" to p.d., and "properly discontinuous and cocompact" to p.d.c.

Examples.
(1) The standard action of $\mathbf{Z}$ on $\mathbf{R}$ by translation ($n.x = n + x$) is p.d.c. The quotient, $\mathbf{R}/\mathbf{Z}$, is a circle.
(2) The action of $\mathbf{Z}$ on $\mathbf{R}^2$ by horizontal translation ($n.(x, y) = (n + x, y)$) is p.d. but not cocompact. The quotient is a bi-infinite cylinder.
(3) The standard action of $\mathbf{Z}^2$ on $\mathbf{R}^2$ (namely $(m, n).(x, y) = (m + x, n + y)$) is p.d.c. The quotient is a torus.
(4) If S is a finite generating set of a group Γ, then the action of Γ on the Cayley graph $\Delta(\Gamma, S)$ is p.d.c. (Note that example (1) is a special case of this.)

In fact, all the actions described above are free.

3.3. Definition of quasi-isometries.

Let X, X' be metric spaces. We will normally assume them to be geodesic spaces, though this is not formally required for the following definitions.

Definition : A map $\phi : X \longrightarrow X'$ is *quasi-isometric* if there are constants, $k_1 > 0, k_2, k_3, k_4 \geq 0$ such that for all $x, y \in X$,

$$k_1 d(x,y) - k_2 \leq d'(\phi(x), \phi(y)) \leq k_3 d(x,y) + k_4.$$

A quasi-isometric map, ϕ, is a *quasi-isometry* if, in addition, there is a constant, $k_5 \geq 0$, such that

$$(\forall y \in X')(\exists x \in X)(d(y, \phi(x)) \leq k_5.$$

Thus, a quasi-isometry preserves distances to within fixed linear bounds and its image is cobounded.

Notes: (1) We do not assume that ϕ is continuous. We are trying to capture one the "large scale" geometry of our spaces. They cannot be expected to respect small scale structure such as topology. Indeed, certain basic observations below would fail if we were to impose such a constraint.

(2) A fairly simple observation is that if two maps ϕ, ψ agree up to bounded distance (i.e. there is a constant $k \geq 0$ such that $d'(\phi(x), \psi(x)) \leq k$ for all $x \in X$) then ϕ is a quasi-isometry if and only if ψ is. This is an example of a more general principle. In coarse geometry we are usually only interested in things up to bounded distance. Indeed, we will frequenly only specify maps up to a bounded distance.

(3) We will be giving various constructions that construct new quasi-isometries from old. (Moving points a bounded distance, as above, might be considered one example.) Usually, in such cases, the new constants of quasi-isometry (the k_i) will depend only on the old ones and any other constants involved in the construction. In principle, one can keep track of this dependence through various arguments, though we not usually bother to do this explicitly.

From now on we will assume that our spaces are length spaces. The following are the basic properties of quasi-isometries. We leave the proofs as an exercise.

Proposition 3.2 : *(1) If $\phi : X \longrightarrow Y$ and $\psi : Y \longrightarrow Z$ are quasi-isometries, then so is $\psi \circ \phi : X \longrightarrow Z$.*

(2) If $\phi : X \longrightarrow Y$ is a quasi-isometry, then there is a quasi-isometry $\psi : Y \longrightarrow X$ with $\psi \circ \phi$ and $\phi \circ \psi$ a bounded distance from the identity maps.

For (2), given $y \in Y$, choose any $x \in X$ with $\phi(x)$ a bounded distance from y and set $\psi(y) = x$. We refer to such a map ψ as a *quasi-inverse* of ϕ. Note that a quasi-inverse cannot necessarily be made continuous, even if ϕ happens to be continuous. A quasi-inverse is unique up to bounded distance — which is the best one can hope for in this context. Note that a quasi-isometric map is a quasi-isometry if and only if it has a quasi-inverse.

Definition : Two length spaces, X and Y, are said to be *quasi-isometric* if there is a quasi-isometry between them.
In this case, we write $X \sim Y$.

Note that, by Proposition 3.2, $X \sim X$ $X \sim Y \Rightarrow Y \sim X$ and $X \sim Y \sim Z \Rightarrow X \sim Z$.

Examples:

(1) Any non-empty bounded space is quasi-isometric to a point.

(2) $\mathbf{R} \times [0,1] \sim \mathbf{R}$: Projection to the first coordinate is a quasi-isometry.

(3) By a similar construction, the Cayley graph of $\mathbf{Z}$ with resepct to $\{a, a^2\}$ is quasi-isometric to $\mathbf{R}$. Recall that is $\mathbf{R}$ is also the Cayley graph of $\mathbf{Z}$ with respect to $\{a\}$.

(4) Similarly $\Delta(\mathbf{Z}; \{a^2, a^3\}) \sim \mathbf{R}$.

(5) The Cayley graph of $\mathbf{Z}^2$ with respect to the standard generators $\{a, b\}$ is quasi-isometric to the plane, $\mathbf{R}^2$. Recall that we can identify this Cayley graph with the 1-skeleton of a square tessellation of the plane, and its inclusion into $\mathbf{R}^2$ is a quasi-isometry.
Any quasi-inverse quasi-isometry will be discontinuous. For example, puncture each square tile at the centre and retract by radial projection to the boundary (Figure 3a). We can send the centre to any boundary point of the tile.

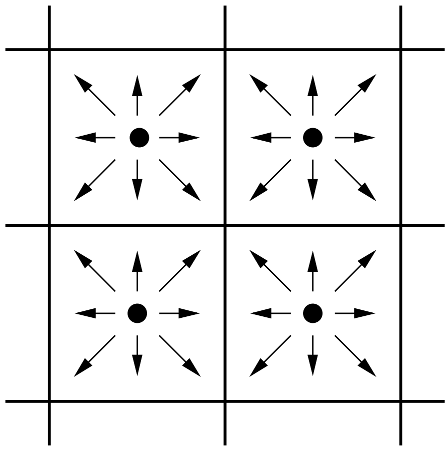

Figure 3a.

(6) By a similar argument, the Cayley graph of $\mathbf{Z}^n$ with the standard generators is quasi-isometric to $\mathbf{R}^n$ — it is the 1-skeleton of a regular tessellation of $\mathbf{R}^n$ by unit cubes.

(7) Let T_n be the n-regular tree. We claim that $T_3 \sim T_4$. To see this, colour the edges of T_3 with three colours so that all three colours meet at each vertex. Now collapse each edge of one colour to a point so as to obtain the tree T_4. (In Figure 3b, the edges of one colour a highlighted in bold.) The quotent map from T_3 to T_4 is a quasi-isometry: clearly

it is distance non-increasing, and arc of length at most $2k+1$ in T_3 can get mapped to an arc of length k in T_4.

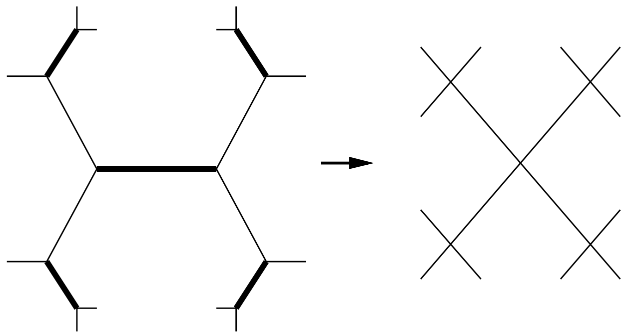

Figure 3b.

Exercise: For all $m, n \geq 3$, $T_m \sim T_n$. Indeed if T is any tree such that the valence of each vertex is at least 3 and at most some constant k, then $T \sim T_3$.

On the other hand, finding quasi-isometry invariants to show that spaces are not quasi-isometric can be more tricky. A significant part of geometric group theory centres around the search for such invariants. Here are a few relatively simple cases.

Non-examples.

(0) The empty set is quasi-isometric only to itself.

(1) Boundedness is a quasi-isometry invariant. Thus, for example, $\mathbf{R} \not\sim [0,1]$.

(2) $[0,\infty) \not\sim \mathbf{R}$. To see this, one can argue as follows. Suppose that $\phi : \mathbf{R} \longrightarrow [0,\infty)$ were a quasi-isometry. Then as $t \to \infty$, $\phi(t) \to \infty$ and $\phi(-t) \to \infty$. Also, $|\phi(n) - \phi(n+1)|$ is bounded. Choose some a much larger that $\phi(0)$, as described shortly. Now the sequence $(\phi(n))_{n \in \mathbf{N}}$ must eventually pass within a bounded distance of a. In other words, there is some $p \in \mathbf{N}$ with $|a - \phi(p)|$ bounded. (Let $p = \max\{n \in \mathbf{N} \mid \phi(n) < a\}$.) Similarly, we can find some $q < 0$ with $|a - \phi(q)|$ bounded. Thus, $|\phi(p) - \phi(q)|$ is bounded, and so $p - q$ is bounded. Thus, $p \leq p - q$ is also bounded. But $|\phi(0) - \phi(p)|$ agrees with $a - \phi(0)$ up to an additive constant. We are free to choose a as large as we want without affecting any of these constants, and so if we take it large enough, we get a contradiction.

Remark: we have only used the fact that ϕ distorts distances by a linearly bounded amount. Thus, in fact, there is no quasi-isometric map from $\mathbf{R}$ to $[0,\infty)$.

We have used a discrete version of the Intermediate Value Theorem. We could not apply this theorem directly since our map, ϕ, was not assumed continuous.

Exercise: write out the above argument more formally with explicit reference to the quasi-isometric constants, k_1, k_2, k_3, k_4.

(3) $\mathbf{R}^2 \not\sim \mathbf{R}$. We sketch a proof using the theorem that any continuous map of the circle to the real line must identify some pair of antipodal points. (This theorem can be deduced from the Intermediate Value Theorem.) By taking a sufficiently large circle we get a contradiction. Since quasi-isometies are not assumed continuous, we will need some kind of approximation argument. One way to formulate this is as follows.

Let $||.||$ denote the euclidean norm on $\mathbf{R}^2$. Suppose $\phi : \mathbf{R}^2 \longrightarrow \mathbf{R}$ were a quasi-isometry. Choose $n \in \mathbf{N}$ sufficienly large (as below) and let $x_0, x_1, \ldots, x_{2n} = x_0$ be $2n$ equally spaced points around the circle S of radius n centred at the origin. Thus $x_{i+n} = -x_i$ and $||x_i - x_{i+1}|| \le \pi$. It follows that $|\phi(x_i) - \phi(x_{i+1})|$ is bounded. We can now define a continuous map $f : S \longrightarrow \mathbf{R}$ by setting $f(x_i) = \phi(x_i)$ and sending the arc of S between x_i and x_{i+1} onto the interval between $\phi(x_i)$ and $\phi(x_{i+1})$ in $\mathbf{R}$. As observed above, this interval has bounded length. From the above theorem, there is some $x \in S$ with $f(x) = f(-x)$. Choose some x_i nearest x in S. Thus $||x - x_i|| = ||(-x) - (-x_i)|| < \pi$ and so $|f(x_i) - f(x)|$ and $|f(-x_i) - f(-x)|$ are both bounded, and so $|\phi(x_i) - \phi(-x_i)| = |f(x_i) - f(-x_i)|$ is bounded. Thus $||x_i - (-x_i)|| = 2||x_i|| = 2a$ is bounded. But we could have chosen a arbitrarily large giving a contradiction.

Indeed we have shown that there is no quasi-isometric map from $\mathbf{R}^2$ into $\mathbf{R}$. We therefore see also that $\mathbf{R}^2 \not\sim [0, \infty)$.

Remark: The Borsuk-Ulam Theorem says that any continuous map from the n-sphere S^n to $\mathbf{R}^n$ must identify some pair of antipodal points. Using this one can deduce that if there is a quasi-isometric map from $\mathbf{R}^m$ to $\mathbf{R}^n$, then $m \le n$. One then sees that if $\mathbf{R}^m \sim \mathbf{R}^n$ then $m = n$. Thus the question of quasi-isometric equivalence is completely resolved for euclidean spaces.

(4) The 3-regular tree, T_3 is not quasi-isometric to $\mathbf{R}$. (Exercise). We thus also have a complete classification for regular trees.

Exercise. Suppose that $f : \mathbf{R}^n \longrightarrow \mathbf{R}^n$ is a proper continuous map. ("Proper" means that $f^{-1}K$ is compact for all compact K.) Suppose there is some $k \ge 0$ such that for all $x \in \mathbf{R}^n$, $\mathrm{diam}(f^{-1}(x)) \le k$. Then f is surjective.

The idea of the proof to extend f to a continuous map between the one-point compactifications $f : \mathbf{R}^n \cup \{\infty\} \longrightarrow \mathbf{R}^n \cup \{\infty\}$, and using appropriate identifications of $\mathbf{R}^n \cup \{\infty\}$ with the sphere, S^n, we can apply the Borsuk-Ulam theorem to get a contradiction.

As a corollary one can get the following.

Any quasi-isometric map from $\mathbf{R}^n$ to itself is a quasi-isometry.

This of course calls for some approximation construction, as in the examples above.

Further quasi-isometry invariants arise from the notion of (Gromov) hyperbolicity that we will encounter later. Indeed some of the above examples can be seen in these terms.

3.4. Cayley graphs again.

Let S, S' be finite generating sets for some group, Γ, and let $\Delta = \Delta(\Gamma; S)$ and $\Delta' = \Delta(\Gamma; S')$ be the corresponding Cayley graphs. We write d, d' for the geodesic metrics on these graphs. Now $V(\Delta) = V(\Delta') = \Gamma$, and we can extend the identity map, $V(\Delta) \longrightarrow$

$V(\Delta')$ to a map $\phi : \Delta \longrightarrow \Delta'$ by sending an edge of Δ linearly to a geodesic in Δ' with the same endpoints. By choosing these geodesics appropriately, we can arrange that the map ϕ is equivariant, that is, $g\phi(x) = \phi(gx)$ for all $x \in \Delta$ and all $g \in \Gamma$. Let $r = \max\{d'(1, a) \mid a \in S\}$. Then each edge of Δ gets mapped to a path of length at most r in Δ'. We see that $d'(\phi(x), \phi(y)) \leq rd(x, y)$ for all $x, y \in \Delta$.

Now we can apply the above construction in the reverse direction to give us an equivariant map $\psi : \Delta' \longrightarrow \Delta$. One can now easily check that these are quasi-inverse quasi-isometric maps, and hence quasi-isometries. We have shown:

Theorem 3.3 : *Suppose that S and S' are finite generating sets for a group Γ. Then there is an equivariant quasi-isometry from $\Delta(\Gamma; S)$ to $\Delta(\Gamma; S')$.* $\diamondsuit$

In particular, the Cayley graph of a finitely generated group is well-defined up to quasi-isometry. If we are only interested in its quasi-isometry class, we can simply denote it by $\Delta(\Gamma)$ without specifying a generating set. This leads us to the following definitions.

Definition : If Γ and Γ' are f.g. groups, we say that Γ is *quasi-isometric* to Γ' if $\Delta(\Gamma) \sim \Delta(\Gamma')$.
We write $\Gamma \sim \Gamma'$.

Examples.
(1) All finite groups are q.i. to each other — their Cayley graphs are bounded.
(2) If $p, q \geq 2$, then $F_p \sim F_q$: Note that with respect to free generating sets, the Cayley graphs are the regular trees T_{2p} and T_{2q}.
(3) If $p \geq 2$, then $F_p \not\sim \mathbf{Z}$.
(4) $\mathbf{Z} \sim \mathbf{Z} \times \mathbf{Z}_2$: Exercise: constuct a Cayley graph for $\mathbf{Z} \times \mathbf{Z}_2$.

Definition : A finitely generated group, Γ is *quasi-isometric* to a geodesic space, X, if $\Delta(\Gamma) \sim X$.
We write $\Gamma \sim X$.

Examples.
(1) $\mathbf{Z} \sim \mathbf{R}$.
(2) $\mathbf{Z}^2 \sim \mathbf{R}^2$.

Note that from the above it follows that $\mathbf{Z} \not\sim \mathbf{Z}^2$. Indeed, from the earlier remark, we know that $\mathbf{Z}^m \sim \mathbf{Z}^n \Rightarrow m = n$.

We thus have complete q.i. classifications of both f.g. free groups and f.g. free abelian groups. Indeed it will follow from results later in the course that if $F_m \sim \mathbf{Z}^n$ then $m = n = 1$, so we can, in fact, classify the union of these two classes by q.i. type. (See Section 6.)

3.5. A useful construction.

Suppose a group Γ acts p.d.c. on a proper geodesic space, X. Fix any $a \in X$. Thus Γa is r-dense in X for some $r \geq 0$. Let $k = 2r + 1$. Construct a graph, Δ, with vertex set $V(\Delta) = \Gamma$ by connecting $g, h \in \Gamma$ by and edge if $d(ga, ha) \leq k$. Since the action is p.d., Δ is locally finite. Also:

Lemma 3.4 : *Δ is connected.*

Proof : Given any $g, h \in \Gamma$, let $\alpha \subseteq X$ be a geodesic connecting ga to ha. Choose a sequence of points, $ga = x_0, x_1, \ldots, x_n = ha$ along α, such that $d(x_i, x_{i+1}) \leq 1$ for all i. For each i choose $g_i \in \Gamma$ so that $d(x_i, g_i a) \leq r$. We can take $g_0 = g$ and $g_n = h$ (Figure 3c). Note that $d(g_i a, g_{i+1} a) \leq k$ for all i, and so g_i is adjacent to g_{i+1} in Δ. Thus the path $g_0 g_1 \cdots g_n$ connects g to h in Δ. $\diamondsuit$

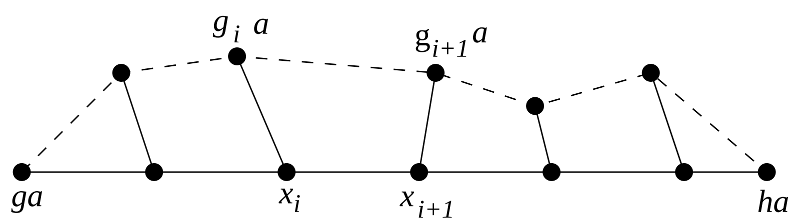

Figure 3c.

Now let $A = \{g \in \Gamma \setminus \{1\} \mid d(a, ga) \leq k\}$. Thus A is finite and symmetric, and $g, h \in \Delta$ if and only if $g^{-1}h \in A$. We see that Δ is, in fact, the Cayley graph of Γ for the generating set A (at least after identifying any double edge corresponding to an order-2 element). From the discussion in Section 1, we see that A generates Γ. Thus:

Theorem 3.5 : *If Γ acts p.d.c. on a proper geodesic space, X, then Γ is finitely generated.* $\diamondsuit$

In fact, we can refine our useful construction to get more information. First note that we have a map $f : \Delta \longrightarrow X$ obtained by setting $f(g) = ga$ and sending the edge between two adjacent $g, h \in \Gamma$ linearly to a geodesic from ga to ha in X. (Again by taking suitable geodesics, we can arrange that f is equivariant.)

Suppose $g, h \in \Gamma$. We can choose the points x_i, as in the proof of Lemma 3.4, evenly spaced so that $n \leq d(ga, ha) + 1 = d(f(g), f(h)) + 1$, where n is the length of the path constructed from g to h in Δ. Conversely, if $d_\Delta(g, h) \leq n$, then $d(f(g), f(h)) \leq rn$. Now $\Gamma = V(\Delta)$ is cobounded in Δ, and $f(V(\Delta)) = \Gamma a$ is cobounded in X. It now follows easily that f is a quasi-isometry from Δ to X. Since Δ is a Cayley graph for Γ we see:

Theorem 3.6 : *If Γ acts properly discontinuously cocompactly on a proper length space X, then $\Gamma \sim X$.* $\diamondsuit$

This tells us once more some things we already knew, for example that $\mathbf{Z}^n \sim \mathbf{R}^n$. We can also get new information.

Proposition 3.7 : *Suppose that Γ is finitely generated and $G \leq \Gamma$ is finite index. Then Γ is finitely generated and $G \sim \Gamma$.*

Proof : Let Δ be any Cayley graph of Γ. We restrict the action of Γ on Δ to an action of G. This is also p.d.c. We now apply Theorems 3.5 and 3.6. $\diamondsuit$

Note that the vertices of Γ/G correspond to the cosets of G in Γ.

3.6. Quasi-isometry and commensurability.

Definition : Two groups Γ and Γ' are *commensurable* if there exist finite index subgroups $G \leq \Gamma$ and $G' \subseteq \Gamma'$ with $G \cong G'$. We write $\Gamma \approx \Gamma'$.

Note that from Theorem 3.5 and an earlier exercise in Section 1, we see that if $\Gamma \approx \Gamma'$ then Γ is finitely generated if and only if Γ' is. Also:

Exercise: The relation $\approx$ is transitive.

We can thus talk about commensurability classes of (f.g.) groups.
Applying Propostion 3.7, we see:

Proposition 3.8 : *If Γ and Γ' are f.g., then $\Gamma \approx \Gamma' \Rightarrow \Gamma \sim \Gamma'$.* $\diamondsuit$

Definition : A group Γ is *torsion-free* if for any $g \in \Gamma$ and $n \in \mathbf{N}$, then $g^n = 1$ implies $g = 1$.

Definition : If "P" is any property, we say that a group is *virtually* P if it has a finite index subgroup that is P.

For example, we have "virually abelian", "virtually free", "virtually torsion-free" etc. Note that all finite groups are "virtually trivial".

Theorem 3.9 : *Suppose that Γ is a f.g. group quasi-isometric to $\mathbf{Z}$. The Γ is virtually $\mathbf{Z}$.*

Proof : This is quite subtle, and we only give the outline.

First, let us suppose that we have found an infinite order element $g \in \Gamma$. Let $G = \langle g \rangle \equiv \mathbf{Z}$. We claim that $[\Gamma : G] < \infty$. To see this, let Δ be any Cayley graph of Γ, so that $V(\Delta) \cong \Gamma$. Note that $d(g^m, g^n) = d(1, g^{m-n})$ and that $d(1, g^n) \to \infty$ as $n \to \pm\infty$. Let

$\phi : \Delta \longrightarrow \mathbf{R}$ be a quasi-isometry. Define a map, $f : \mathbf{Z} \longrightarrow \mathbf{R}$ by $f(n) = \phi(g^n)$. From the above we see that

(1) For all n, $|(f(n) - f(n+1)|$ is bounded, and

(2) $(\forall r \geq 0)(\exists p \in \mathbf{N})$ such that if $|f(n) - f(m)| \leq r$ then $|m - n| \leq p$.

Now it is an exercise to show that the image of any map from $\mathbf{Z}$ to $\mathbf{R}$ satisfying the above is cobounded in $\mathbf{R}$. It then follows that $G \subseteq V(\Delta)$ is cobounded in Δ. Thus, Δ/G is a finite graph, G has finite index in Γ as claimed. (Note that the vertices of Δ/G correspond to cosets of G in Γ.) We that thus proven the theorem in this case.

We still need to find an infinite order element, g. The idea is fairly simple, but the details take a while to write out. We just give the general idea. We shall find some $g \in \Gamma$ and some subset $A \subseteq \Gamma = V(\Delta)$, such that gA is properly contained in A. It then follows easily that g must have infinite order.

Suppose $A = V(\Delta) \cap \phi^{-1}[0, \infty)$ where $\phi : \Delta \longrightarrow \mathbf{R}$ is our quasi-isometry. Now any $g \in \Gamma$ acts by isometry on Δ, and so determines (via ϕ) a quasi-isometry ψ from $\mathbf{R}$ to itself. It is an exercise to show that $\psi([0, \infty))$ a bounded distance from $[\phi(0), \infty)$ (i.e., each point of one set is a bounded distance from some point of the other) or else is a bounded distance from $(-\infty, \psi(0)]$. Now if the former is the case, and if $\psi(0)$ is much greater than 0, it then follows that gA is properly contained in A, and so we are done.

To find such a g, we take two elements $h, k \in \Gamma$, so that $1, g, k$ are all very far apart in Δ. Thus, $\phi(1), \phi(h), \phi(k)$ are all far apart in $\mathbf{R}$. One can now apply the argument of the previous paragraph, considering the images of $[0, \infty)$. At least two of the sets A, hA, kA are nested (one properly contained in the other). We can then take g to be one of the elements h, k, hk^{-1} or $h^{-1}k$. $\diamondsuit$

We remark that this result is a weak version of a result of Hopf from around 1940 that a f.g. group with "two ends" is virtually $\mathbf{Z}$.

Prompted by Theorem 3.9, one can ask the following:

General question: When does $\Gamma \sim \Gamma'$ imply $\Gamma \approx \Gamma'$?

In general this is very difficult to answer.

Some positive examples.

(1) This is true if one of the groups is finite: then they are both finite.

(2) True if one of the groups is (virtually) $\mathbf{Z}$, by Proposition 3.9.

(3) True if both groups are virtually abelian. We can argue as follows. Let G be a finite index subgroup of Γ. Then by Proposition 3.7, G is finitely generated. In fact, we can assume that G is also torsion free, since we could write $G \cong G' \times T$, were G' is torsion free, and T is finite, and then replace G by G'. Now any finitely generated torsion-free abelian group is isomorphic to $\mathbf{Z}^n$ for some n (from the classification of f.g. abelian groups). In other words, Γ is virtually $\mathbf{Z}^n$. Similarly Γ' is virtually $\mathbf{Z}^m$ for some m. Thus $\mathbf{Z}^n \sim \Gamma \sim \Gamma' \sim \mathbf{Z}^m$ and so $m = n$. Thus $\Gamma \approx \Gamma'$.

(4) In fact this remains true if we only assume that one of these groups is virtually abelian. In other words any f.g. group q.i. to a virtually abelian group (or equivalently a euclidean plane) is itself virtually abelian. This is however a much deeper theorem. The first proof of this used the result of Gromov: "Any group of polynomial growth is virtually nilpotent" (in turn using another deep result of Montgomery and Zippin characterising Lie groups) and then using some q.i. invariants of nilpotent groups. A more direct, though still difficult, proof has since been given by Shalom, by very different arguments.

(5) If both groups are (virtually) free, the statement is true. Given that $F_n \sim \mathbf{Z}$ only if $n = 1$, by the above results, it remains to show that if $m, n \geq 2$, then F_m and F_n have an isomorphic finite index subgroup. We will discuss this again later (Section 4).

(6) In fact, the above holds if only one group is assumed virtually free: any group q.i. to a (virtually) free group is virtually free. This again uses some sophisticated machinery. It follows from a result of Dunwoody on "accessibility" together with work of Stallings on group splittings.

(7) Other examples relating to surfaces will be disussed in the context of hyperbolic geometry in Section 5.

Some negative results.

There certainly (many) examples where the statement fails: non-commensurable groups that are q.i. However, I don't know of a simple example that one can easily verify. A standard example comes from 3-manifold theory. There are compact hyperbolic 3-manifolds, M and N, which do not have any common finite cover. Then $\pi_1(M) \not\approx \pi_1(N)$ but both groups are q.i. to hyperbolic 3-space. We discuss this again in Section 5.

3.7. Quasi-isometry invariants.

Properties of groups that are invariant under quasi-isometry are often termed "geometric". There are many geometric invariants which we won't have time to look at seriously. Here are a few examples:

(1) Finite presentability: If $\Gamma \sim \Gamma'$ then Γ is f.p. if and only if Γ' is f.p.

(2) The word problem. Suppose Γ is f.p. A word in the generators and their inverses represents some element of the group. Is there an algorithm to decide if this is the identity element? If so then the group is said to have solvable word problem. For finitely presented groups this turns out to be a geometric property, and follows from work of Alonso and Shapiro (see the discussion at the end of Section 6). For f.g. groups it appears to open whether having solvable word problem is geometric.

(3) As alluded to earlier, by the result of Gromov, the property of being virtually nilpotent is geometric.

There are many other results and open problems. One issue arises from recognising torsion from the geometry of a group. We saw, for example, that it was somewhat complicated to show that a group q.i. to $\mathbf{Z}$ contained an infinite order element. To further illustrate this, it appears to be open as to whether a torsion free group can be q.i. to a torsion group, i.e. a group in which every element has finite order.

4. Fundamental groups

In this section we give a review of some background material relating to fundamental groups and covering spaces. Since these serve primarily as illustrative examples, our presentation will be informal, and we omit proofs. We will describe some consequences for free groups.

4.1. Definition of fundamental groups.

The fundamental group is an invariant of a topological space. We are not interested here in "pathological" examples. We will generally assume that our spaces are reasonably "nice". For example, manifolds and simplicial complexes are all "nice", and in practice that is all we care about.

Let X be a topological space. Fix a "basepoint" $p \in X$. A *loop* based at p is a path $\alpha : [0,1] \longrightarrow X$ with $\alpha(0) = \alpha(1) = p$. Two such loops, α, β, are *homotopic* if one can be deformed to the other through other loops; more precisely, if there is a map $F : [0,1]^2 \longrightarrow X$ with $F(t,0) = \alpha(t)$, $F(t,1) = \beta(t)$ and $F(0,u) = F(1,u) = p$ for all $t, u \in [0,1]$. This defines an equivalence relation on the set of paths. Write $[\alpha]$ for the homotopy class of α.

Given loops α, β, write $\alpha * \beta$ for the path that goes around α (twice as fast) then around β (i.e. $\alpha * \beta(t)$ is $\alpha(2t)$ for $t \leq 1/2$ and $\beta(2t-1)$ for $t \geq 1/2$) (Figure 4a). Write $[\alpha][\beta] = [\alpha * \beta]$.

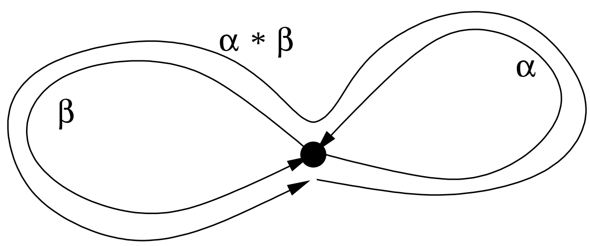

Figure 4a.

Exercise: This is well-defined and gives the set of homotopy classes of loops based at p the structure of a group.

Definition : We call this group the *fundamental group* of X (based at p).

It is denoted by $\pi_1(X, p)$.
The following is not hard to see:

Fact: If the points $p, q \in X$ are connected by a path, then $\pi_1(X, p) \cong \pi_1(X, q)$.

To define the isomorphism, given a loop based at p, we can obtain a loop based at q by following the path from q to p, then going around the loop, and then following the same path back to q. We leave the details as an exercise.

Definition : A space X is *path connected* if any two points are connected by a path.

Note that any path-connected space is connected, and any "nice" connected space will also be path-connected. (There are counterexamples to the latter statement, but these are not nice.)

Thus if X is path-connected, then the fundamental group is well-defined up to isomorphism. It is denoted $\pi_1(X)$. Clearly homeomorphic spaces have isomorphic fundamental groups.

Examples.

(1) $\pi_1(\text{point}) = \{1\}$.

(2) $\pi_1(S^1) = \mathbf{Z}$.

(3) $\pi_1(S^1 \times [0, 1]) = \mathbf{Z}$.

(4) $\pi_1(S^1 \times \mathbf{R}) = \mathbf{Z}$.

(5) $\pi_1(\text{torus}) = \mathbf{Z}^2$. This is generated by two (homotopy classes of) loops, a, b, on the torus that cross just once at p (Figure 4b). It is easily seen that $aba^{-1}b^{-1} = 1$. This gives us a presentation of $\mathbf{Z}$. Of course, a lot more work would be needed to prove directly that this is actually the fundamental group. This will follow from a result stated below.

(6) $\pi_1(\text{figure of eight}) = F_2$. The two generators correspond to the two loops. This time there is no relation.

(7) More generally the fundamental group of a wedge of circles is free. (A "wedge" is constructed by gluing together is a collection of spaces at a single point.) For a finite wedge of n circles, we get F_n. A wedge of 5 circles is illustrated in Figure 4c. A "figure of eight" is a wedge of 2 circles.

Definition : A space X is *simply connected* if it is path-connected and $\pi_1(X)$ is trivial.

This means that every closed path bounds a disc. More precisely, if D is the unit disc in $\mathbf{R}^2$ then any map $f : \partial D \longrightarrow X$ extends to a map $f : D \longrightarrow X$.

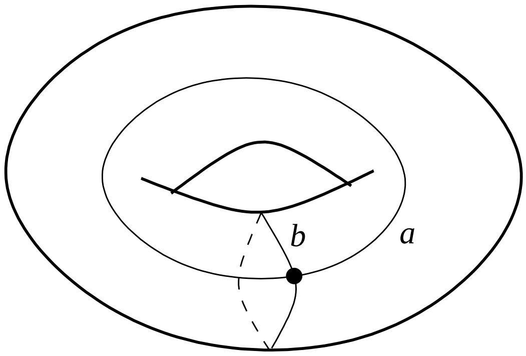

Figure 4b.

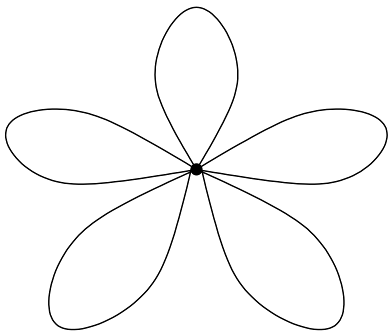

Figure 4c.

Examples: A point is simply connected! So is $\mathbf{R}^n$ for any n. So is any tree. So is the n-sphere for any $n \geq 2$.

Suppose X is "nice" and $Y \subseteq X$ is a "nice" closed simply connected subset. Then it turns out that the fundamenal group is unchanged by collapsing Y to a point.

For example, if X is a graph, and Y is a subtree (a subgraph that is a tree) then we can collapse Y to a single vertex and get another graph. If we take Y to be a maximal tree, we get a wedge of circles (Figure 4d).

Since a maximal tree always exists (if you believe the Axiom of Choice in the case of an infinite graph) we can deduce the following facts:

Facts:

(1) The fundamental group of a graph is free.

(2) The fundamental group of a finite graph is F_n for some n.

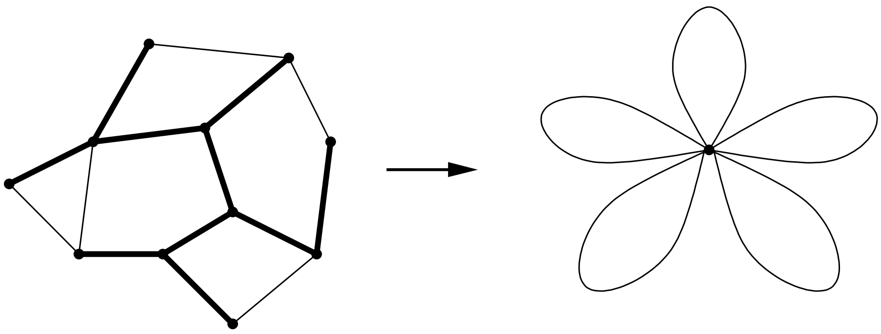

Figure 4d.

4.2. Covering spaces.

Covering spaces give an alternative viewpoint on fundamental groups for nice spaces. Indeed the discussion here could serve to give an equivalent definition of fundamental group, suited to our purposes. We shall not approach the subject systematically here. We just say enough to give the general ideas needed for subsequent applications.

Suppose that a group Γ acts properly discontinuously on a proper space X. For our purposes here, we can take X to be a proper geodesic space and assume that Γ acts by isometry. We shall also assume the action is *free*. This is yet another usage of "free". Here it means there are no fixed points, i.e. if $gx = x$ for some $x \in X$, then $g = 1$. We can form the quotient space X/Γ. The space X, together with the quotient map to X/Γ is a (particular) example of a "covering space" as we disuss below.

Of particular interst is the case where X is simply connected. In this case, we have the following:

Fact: In the above situation, $\pi_1(X/\Gamma) \cong \Gamma$.

The isomorphism can be seen by fixing some $p \in X$. Given $g \in \Gamma$, connect p to gp by some path α in X. (Since X is simply connected, it doesn't really matter which one.) This projects to a loop in X/Γ with a fixed basepoint, and hence determines an element of $\pi_1(X/\Gamma)$. This gives a homomorphism from Γ to $\pi_1(X/\Gamma)$, which turns out to be an isomorphism.

Conversely, given a nice space, Y, one can construct a simply connected space X, and free p.d. action of $\Gamma = \pi_1(Y)$ on X such that $Y = X/\Gamma$. Then, X is called the *universal cover* of Y. (It is well-defined up to homeomorphism.) It is often denoted $\tilde{Y}$.

Examples.

(1) $\Gamma = \mathbf{Z}$. $X = \mathbf{R}$, $Y = S^1$.

(2) $\Gamma = \mathbf{Z}$, $X = \mathbf{R} \times [0,1]$, $Y = S^1 \times [0,1]$.

(3) $\Gamma = \mathbf{Z}$, $X = \mathbf{R}^2$, $Y = S^1 \times \mathbf{R}$.
(4) $\Gamma = \mathbf{Z}^2$, $X = \mathbf{R}^2$, $Y = S^1 \times S^1$ is the torus.
(5) $\Gamma = F_n$, $X = T_{2n}$, Y is a wedge of n circles. We are taking the action of F_n on its Cayley graph, T_{2n}, which here happens to be the universal cover.

Note that if $G \leq \Gamma$ is a subgroup, we also get a natural map from $Z = X/G$ to $Y = X/\Gamma$. This is a more general example of a "covering space". Formally we say that a map $p : Z \longrightarrow Y$ is a *covering map* if every point $y \in Y$ has a neighbourhood U such that if we restrict p to any connected component of $p^{-1}U$ we get a homeomorphism of this component to U. In this situation, Z is called a *covering space*. We will not worry too much about this formal definition here. The following examples illustrate the essential points:

(1) Consider the action of $\mathbf{Z}$ on $\mathbf{R}$, and the subgroup $n\mathbf{Z} \leq \mathbf{Z}$. We get a covering $\mathbf{R}/n\mathbf{Z} \longrightarrow \mathbf{R}/\mathbf{Z}$. This is a map from the circle to itself wrapping around n times.
(2) We get a covering map of the cylinder to the torus, $\mathbf{R}^2/\mathbf{Z} \longrightarrow \mathbf{R}^2/\mathbf{Z}^2$.

The main point to note is that the fundamental group of a cover is a subgroup of the fundamental group of the quotient. If both spaces are compact, then the subgroup will have finite index. These statements can be thought of in terms of the actions on the universal cover.

Exercise. If G happens to be normal in Γ, then there is a natural action of the group Γ/G on Z, and Y can be naturally identified as the quotient of Z by this action. The covering space $Z \longrightarrow Y$ is then the quotient map.

If G is not normal, then the cover will not arise from a group action; so the notion of a covering space is more general that of a free p.d. group action. An example described at the end of this section illustrates this.

4.3. Applications to free groups.

As well as being useful to illustrate later results, these constructions have implications for free groups.

Theorem 4.1 *Any subgroup of a free group is free.*

Proof : Suppose F is free, and $G \leq F$. Its universal cover, X, is a tree. (It will only be a proper space if F is finitely generated, but that doesn't really matter here.) Now G acts of T and T/G is a graph. By the earlier discussion, $G \cong \pi_1(T/G)$ is free. $\diamondsuit$

This is a good example of a result that is relatively easy by topological/geometric means, but quite hard to prove by direct combinatorial means.

We note that G need not be f.g. even if F is. As an example consider the $F_2 = \langle a, b\rangle$ and let $G = \langle \{b^n a b^{-n} \mid n \in \mathbf{Z}\}\rangle$. In this case, the covering space, $K = T_4/G$, is the real

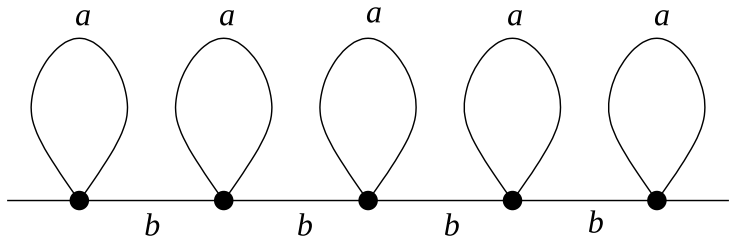

Figure 4e.

line with a loop attached to each integer point (Figure 4e).

Collapsing the real line to a point we get an infinite wedge of circles, and so F is free on an infinite set, and so cannot be finitely generated (this was an exercise in Section 1). In fact, $\{b^n a b^{-n} \mid n \in \mathbf{Z}\}$ is a free generating set. This is one of the simplest examples of an infinitely generated subgroup of a finitely generated group.

Exercise. $G = \langle\langle a \rangle\rangle$. In particular, $G \triangleleft F$. In fact, G corresponds to the set of words in a, a^{-1}, b, b^{-1} with the same number of b's and b^{-1}'s. Writing $J = F/G$, we have $J \cong \mathbf{Z}$. Now J acts by translation on the graph K, and the quotient graph, K/J, is a "figure of eight", which is naturally identified with T_4/F.

The map $K \longrightarrow K/J$ is another example of a covering space.

The subgroup $H = \langle a, bab^{-1}, b^2ab^{-2}, \ldots \rangle$ is also infinitely generated (but not normal). In this case, the covering space T_4/H is a bit more complicated, and the covering map to the figure of eight does not arise from a group action.

Theorem 4.2 *If $p, q \geq 2$, then $F_p \approx F_q$.*

Proof : Let K_n be the graph obtained by taking the circle, $\mathbf{R}/n\mathbf{Z}$, and attaching a loop at each point of $\mathbf{Z}/n\mathbf{Z}$ — that is n additional circles (see Figure 4f, where $n = 5$).

We can collapse down a maximal subtree of K_n to give us a wedge of $n+1$ circles. Thus $\pi_1(K_n) = F_{n+1}$. (The universal cover of K_n is T_4.) We also note that if for any $m \in \mathbf{N}$, K_{mn} is a cover of K_n.

Now given $p, q \geq 2$, set $r = pq - p - q + 2 = (p-1)(q-1) + 1$, and note that K_r covers both K_p and K_q. Since these are all compact, we see that F_r is a finite index subgroup of both F_p and F_q. $\diamondsuit$

This proves something we commented on earlier, namely that two f.g. free groups are q.i. if and only if they are commensurable. There are three classes: $F_0 = \{1\}$, $F_1 = \mathbf{Z}$, and F_n for $n \geq 2$.

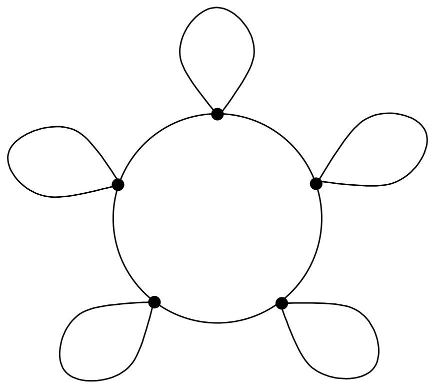

Figure 4f.

5. Hyperbolic geometry.

Here we give a brief outline of some of the main features of hyperbolic geometry. Again, this will serve mainly as a source of examples and motivation, and we will not give detailed proofs. One-dimensional hyperbolic space is just the real line, so we begin in dimension 2. The main ideas generalise to higher dimensions.

5.1. The hyperbolic plane.

We describe the "Poincaré model" for the hyperbolic plane. For this it is convenient to use complex coordinates. Let $D = \{z \in \mathbf{C} \mid |z| < 1\}$. Suppose $\alpha : I \longrightarrow D$ is a smooth path. We write $\alpha'(t) \in \mathbf{C}$ for the complex derivative at t. Thus, $|\alpha'(t)|$ is the "speed" at time t. The euclidean length of α is thus given by the formula $l_E(\alpha) = \int_I |\alpha'(t)| dt$. This is equal to the "rectifiable" length as defined in Section 3.

We now modify this by the introducion a scaling factor, $\lambda : D \longrightarrow (0, \infty)$. The appropriate formula is: $\lambda(z) = 2/(1 - |z|^2)$. The hyperbolic length of α is thus given by $l_H(\alpha) = \int_I \lambda(\alpha(t))|\alpha'(t)| dt$.

Note that as z approaches ∂D in the euclidean sense, then $\lambda(z) \to \infty$. Thus close to ∂D, things big in hyperbolic space may look very small to us in euclidean space. Indeed, since $\int_0^\infty \frac{2}{1-x^2} dx = \infty$, one needs to travel an infinite hyperbolic distance to approach ∂D. For this reason, ∂D, is often referred to as the *ideal* boundary — we never actually get there.

Given $x, y \in D$, write $\rho(x, y) = \inf\{l_H(\alpha)\}$ as α varies over all smooth paths from x to y. In fact, the minimum is attained — there is always a smooth geodesic from x to y. The remark about the ideal boundary in the previous paragraph boils down to saying that this metric is complete. Moreover, if we want to get between two points x and y as quickly as possible, it would seem a good idea to move a little towards the centre of the disc, in the euclidean sense. Thus we would expect hyperbolic geodesics approach the middle of the disc relative to their euclidean counterparts.

For a more precise analysis, we need the notion of a *Möbius transformation*. This is a map $f : \mathbf{C} \cup \{\infty\} \longrightarrow \mathbf{C} \cup \{\infty\}$ of the form $f(z) = \frac{az+b}{cz+d}$ for constants $a, b, c, d \in \mathbf{C}$ with $ad - bc \neq 0$. We set $f(\infty) = c/d$ and $f(-c/d) = \infty$ (though we don't really have to worry about ∞ here). It is usual to normalise so that $ad - bc = 1$. A Möbius transformation is bijective, and one sees easily that the set of such transformations forms a group under composition. Since it is complex analytic, any Möbius tranformation is conformal, i.e. it preserves angles.

Here are some useful observations about Möbius transformations, which we can leave as exercises:

Exercises

(1) A Möbius transformation sends euclidean circles to euclidean circles, where we allow a straight line union ∞ to be a "circle". (Warning: it need not preserve centres of circles.)

(2) If $d = \bar{a}$ and $c = \bar{b}$ (the complex conjugates) and $|a|^2 - |b|^2 > 0$ then $f(D) = D$. (We shall normalise so that $|a|^2 - |b|^2 = 1$.) In fact, any Möbius tranformation preserving D must have this form.

(3) Such an f (as in (2)) is an isometry of (D, ρ). For this, one needs to check that if α is a smooth path, then $l_H(f \circ \alpha) = l_H(\alpha)$. This follows from the formula, $\lambda(f(z))|f'(z)| = \lambda(z)$, which can be verified by direct calculation.

(4) If $z, w \in D$, then there is some such f sending z to w. (Without loss of generality, $w = 0$.)

(5) If $p, q, r \in \partial D$ are distinct, and $p', q', r' \in \partial D$ are distinct, and the orientation of p, q, r is the same as that of p', q', r', then there is some such f with $f(p) = p'$, $f(q) = q'$ and $f(r) = r'$. Here is one way to see this. First show there is a (unique) Möbius transformation taking any three distinct points of $\mathbf{C} \cup \{\infty\}$ to any other three distinct points. (Since they form a group, we could take one set to be $\{0, 1, \infty\}$.) By (1), if all six points lie in ∂D then the Möbius transformation must preserve ∂D, since three points determine a euclidean circle. The condition about orientation is needed so that f sends the interior of D to the interior, rather than the exterior.

Putting together (3) and (4), we see that (D, ρ) is *homogeneous* that is, there is an isometry taking any point to any other point. It thus looks the same everywhere. In fact, any rotation about the origin is clearly an isometry (and a Möbius tranformation). Thus (D, ρ) is also *isotropic* — it looks the same in all directions. It thus shares these properties with the euclidean plane.

Now by symmetry it is easily seen that any euculidean diameter of D, (for example, the interval $(-1, 1)$) is a bi-infinite geodesic with respect to the metric ρ. Indeed it is the unique geodesic between any pair of points on it. Under isometries of the above type it is mapped onto arcs of euclidean circles othogonal to ∂D. Since any pair of points of D lie on such a circle, we see that all geodesics must be of this type, and so we conclude (Figure 5a):

Proposition 5.1 : *Bi-infinite geodesics in the Poincaré disc are arcs of euclidean circles othogonal to the ∂D (including diameters of the disc).* $\diamondsuit$

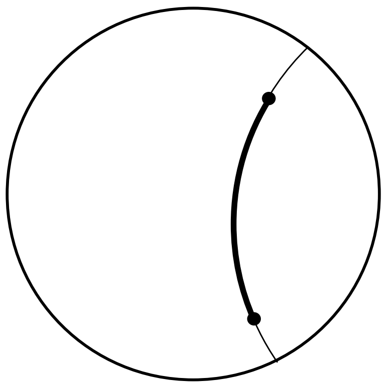

Figure 5a.

We remark that, in fact, all orientation preserving isometries of (D, ρ) are Möbius transformations of the above type. This is not very hard to deduce given our description of geodesics, but we shall not formally be needing it.

The isometry type of space we have just constructed is generally referred to as the *hyperbolic plane*, and denoted $\mathbf{H}^2$. It has many descriptions. The one we have given is called the *Poincaré model*

Hyperbolic geometry has its roots in attempts to understand the "parallel postulate" as formulated by Euclid about 2300 years ago. Its discovery, due to Bolyai, Lobachevsky and Gauß in the 1830s is one of the great landmarks in mathematics. Varous explicit models of hyperbolic geometry were subsequently discovered by Beltrami, Poincaré, Klein etc.

5.2. Some properties of the hyperbolic plane.

(1) In general angles in hyperbolic geometry are "smaller" than in the corresponding situation in Euclidean geometry. For example, if T is a triangle with angles p, q, r, then $p + q + r < \pi$. In fact, one can show that the area of T is $\pi - (p + q + r)$. One can allow for one or more of the vertices to lie in the ideal boundary, ∂D, in which case the corresponding angle is deemed to be 0. An *ideal triange* is one where all three vertices are ideal. Its area is π.

(2) Triangles are "thin". One way of expressing this is to say that there is some fixed constant $k > 0$, so that if T is any triangle there is some point, $x \in \mathbf{H}^2$, whose distance from all three sides is at most k. (One can, in fact, take x in the interior of T.) To verify this, one can first deal with case of an ideal triangle. By excercise (5) above, we see that there is a hyperbolic isometry carrying any ideal triangle to any other. In particular, we

can suppose that its vertices are three equally spaced points in ∂D. A calculation now shows that the centre, 0, of this triangle is a distance $\frac{1}{2}\log 3$ from each of the three sides. In fact, the largest disc in the interior of an ideal triangle is the one of radius $\frac{1}{2}\log 3$ about the centre. It touches all three sides. Now, for a general triangle, take a disc in the interior of maximal radius, r, say. This must touch all three sides, otherwise we could make it bigger. The centre of the disc is thus a distance r from all three sides. Now by pushing the three vertices of the triangle towards the ideal boundary, we can place our triangle inside an ideal triangle. It therefore follows that $r \leq \frac{1}{2}\log 3$, and so the same constant, $k = \frac{1}{2}\log 3$, works for all triangles (Figure 5b).

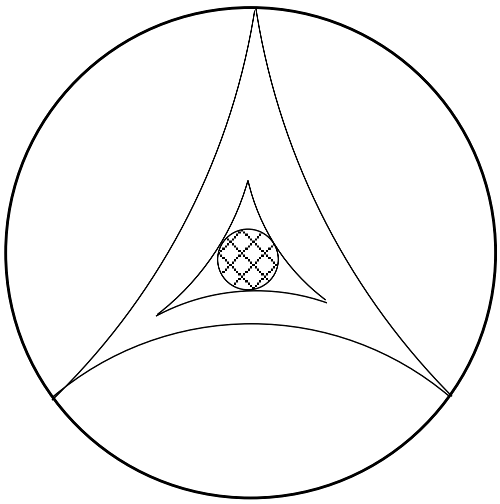

Figure 5b.

(3) A (round) circle of radius r in $\mathbf{H}^2$ has length $2\pi\sinh r$. A round disc $B(r)$, of radius r has area $2\pi(\cosh r - 1)$. (This is an exercise in integration: note that a hyperbolic circle about the origin, 0, of the Poincaré model is also a euclidean circle.) We note in particular, that $\mathrm{area}(B(r)) \leq \mathrm{length}(\partial B(r))$. In fact, if B is any topological disc in $\mathbf{H}^2$, one can show with a little bit of work that $\mathrm{area}(B) \leq \mathrm{length}(\partial B)$. (It turns out that the round disc is the "worst case", but this is much harder to show.) This is in contrast to the euclidean plane, where such bounds are quadratic. Inequalities of this sort are called "isoperimetric inequalities". We will briefly mention this again in Section 6.

5.3. Tessellations of $\mathbf{H}^2$.

Suppose that $n \in \mathbf{N}$, $n \geq 3$. The regular euclidean n-gon has all angles equal to $(1 - \frac{2}{n})\pi$. If $0 < \theta < (1 - \frac{2}{n})\pi$, then one can construct a regular hyperbolic n-gon with all angles equal to θ. This can be seen simply by using a continuity argument: start with a very small regular n-gon centred at the origin of the Poincaré disc. Now push all the vertices out to the ideal bounday at a uniform rate. The angles must all tend to 0. Thus at some intermediate point, they will all equal θ. (In fact the angles decrease monotonically, so the polygon is unique.) Now it θ has the form $2\pi/m$ for some $m \in \mathbf{N}$ $n \geq 3$, we get a regular tessellation of the hyperbolic plane by repeatedly reflecting the polygon in its edges. (This requires some sort of formal argument, for example, applying "Poincaré's Theorem", but there are lots of pretty computer images to demonstrate that it works.) Note that the condition $\frac{2\pi}{m} < (1 - \frac{2}{n})\pi$ reduces to $\frac{1}{m} + \frac{1}{n} < \frac{1}{2}$. Thus we get:

Proposition 5.2 : *If $m, n \in \mathbf{N}$ with $\frac{1}{m} + \frac{1}{n} < \frac{1}{2}$, then there is a regular tessellation of the hyperbolic plane by regular n-gons so that m such n-gons meet at every vertex.* $\diamondsuit$

We remark that in the euclidean situation, the corresponding condition is $\frac{1}{m} + \frac{1}{n} = \frac{1}{2}$. In this case, we just get the three familiar tilings where $(m, n) = (3, 6), (4, 4), (6, 3)$.

5.4. Surfaces.

This has consequences for the geometry and topology of surfaces. For simplicity, we consider here only closed orientable surfaces. These are classified by their "genus", which is a non-negative integer.

The sphere (genus 0) clearly has "spherical geometry" — as the unit sphere in $\mathbf{R}^3$. It is simply connected. In other words, its fundamental group is trivial. Moreover, it is compact, and therefore quasi-isometric to a point. At this point, geometric group theorists start to lose interest in spherical geometry.

The torus, T, (genus 1) is a bit more interesting. We can think of the torus topologically as obtained by gluing together the opposite edges of the unit square, $[0, 1]^2$. Note that at the vertex we get a total angle of $4(\pi/2) = 2\pi$, so there is no singularity there. We get a metric on the torus which is locally euclidean (i.e. every point has a neighbourhood isometric to an open subset of the euclidean plane). Such a metric is often referred to simply as a "euclidean structure". The universal cover of the torus is then, in a natural way, identified with the euclidean plane, $\mathbf{R}^2$, with the fundamental group acting by translations. (Given the remarks on universal covers in Section 4, this in fact gives a proof that $\pi_1(T) \cong \mathbf{Z}^2$.) The square we used to construct T lifts to a regular square tiling of the plane. Note that the edges of the square project to loops representing generators, a, b of $\pi_1(T)$. Reading around the boundary of the square we see that $[a, b] = aba^{-1}b^{-1} = 1$. The 1-skeleton of the square tessellation of the plane can be identified with the Cayley graph of $\pi_1(T)$ with respect to these generators. In particular, we see that $\pi_1(T)$ is quasi-isometric to the euclidean plane.

We have made most of the above observations, in some form, before. We now move into new territory.

Let S be the closed surface of genus 2. We can construct S by taking a (regular) octagon and gluing together its edges. We do this according to the cyclic labelling $ABA^{-1}B^{-1}CDC^{-1}D^{-1}$, so that the first edge gets mapped to the third with opposite orientation etc. (Figure 5c).

If we try the above constuction with a euclidean octagon we would end up with an angle of $8(3\pi/4) = 6\pi > 2\pi$ at the vertex, so our euclidean structure would be singular. It

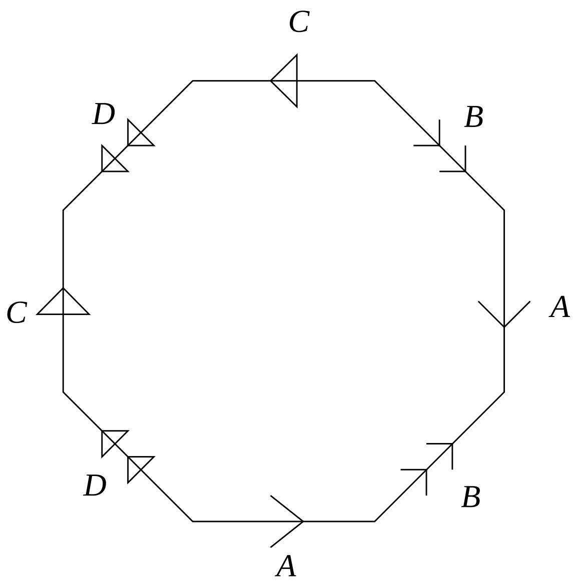

Figure 5c.

is therefore very natural to take instead the regular hyperbolic octagon all of whose angles are $\pi/4$. In this way we get a metric on S that is locally hyperbolic, generally termed a *hyperbolic structure* in S. The universal cover is $\mathbf{H}^2$ and our octagon lifts to a tessellation of the type $(m,n) = (8,8)$ described above. The edges of the octagon project to loops a, b, c, d, and reading around the boundary, we see that $[a,b][c,d] = aba^{-1}b^{-1}cdc^{-1}d^{-1} = 1$ (Figure 5d).

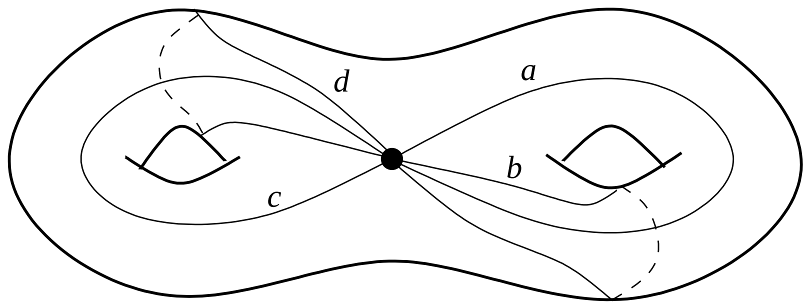

Figure 5d.

In fact, it turns out that $\langle a, b, c, d \mid [a,b][c,d] = 1\rangle$ is a presentation for $\pi_1(S)$. Its Cayley graph can be identified with the 1-skeleton of our $(8,8)$ tessellation of $\mathbf{H}^2$. We see that $\pi_1(S) \sim \mathbf{H}^2$. Indeed, for this, it is enough, by Theorem 3.6, to note that S is the quotient of a p.d.c. isometric action on $\mathbf{H}^2$.

More generally, if S is a closed surface of genus, $g \geq 2$, we get a similar story taking a regular $4g$-gon with cone angles $\pi/2g$. We get

$$\pi_1(S) \cong \langle a_1, b_1, a_2, b_2, \ldots, a_g, b_g \mid [a_1, b_1][a_2, b_2] \cdots [a_g, b_g] = 1 \rangle.$$

The Cayley graph is the 1-skeleton of a $(4g, 4g)$ tesselation of $\mathbf{H}^2$.

In summary we see:

Proposition 5.3 : *If S is a closed surface of genus at least 2, then $\pi_1(S) \sim \mathbf{H}^2$.* $\diamondsuit$

Remark: We just showed that such a surface admits some hyperbolic structure. However, there are lots of variations on this construction, and in fact, there are many different hyperbolic structures one could put on such a surface. Indeed there is a whole "Teichmüller space" of them, and Teichmüller theory is a vast subject in itself.

We also note:

Proposition 5.4 : *If S and S' are closed orientable surfaces of genus at least 2, then $\pi_1(S) \approx \pi_1(S')$.*

Proof : We can do this by a similar argument to Theorem 4.2. Imagine embedding the graph K_n in $\mathbf{R}^3$, and thickenning it up to a 3-dimensional object (called a "handlebody") whose boundary is a surface of genus $n + 1$. (See Figure 5e, where $n = 5$.) Now we do essentially the same construction, to see that a surface of genus $p \geq 2$ and a surface of genus $q \geq 2$ are both covered by a surface of genus $pq - p - q + 2$. $\diamondsuit$

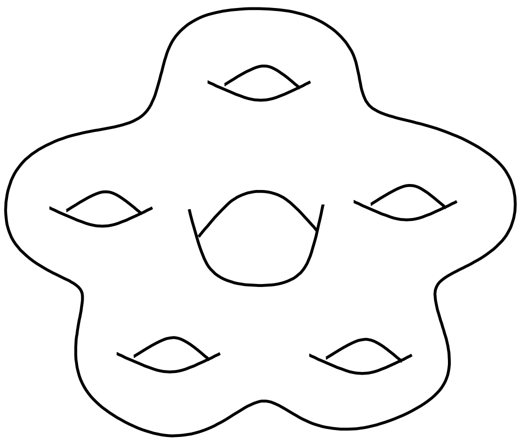

Figure 5e.

(Note that the above was a topological not a geometric construction. The covers need not respect any given hyperbolic structures.)

Theorem 5.5 : *Suppose S and S' are closed orientable sufaces. If $\pi_1(S) \sim \pi_1(S')$ then $\pi_1(S) \approx \pi_1(S')$.*

If we believe that the euclidean plane is not quasi-isometric to the hyperbolic plane, then there are exactly three quasi-isometry classes — one each for the sphere, the torus, and all higher genus surfaces. The result then follows by Propostion 5.4. The fact that $\mathbf{R}^2 \not\sim \mathbf{H}^2$ will follow from our discussion of hyperbolicity in Section 6.

Alternatively you believe that a higher genus surface group is not (virtually) $\mathbf{Z}^2$, then the fact that $\mathbf{R}^2 \not\sim \mathbf{H}^2$ also follows from the q.i. invariance of virually abelian groups, cited but not proven, in Section 3. This is, however, much more work than is necessary for this result.

Fundamental groups of closed surface, other than the 2-sphere, are generally just referred to as *surface groups*. (Any non-orientable surface is double covered by an orientable surface, and so non-orientable surfaces can easily be brought into the above discussion.)

Fact: Any f.g. group quasi-isometric to a surface group is a virtual surface group.

The case of the torus was already discussed in Section 3. The hyperbolic case (genus at least 2) is a difficult result of Tukia, Gabai and Casson and Jungreis.

In fact any group quasi-isometric to a complete riemannian plane is a virtual surface group. This was shown by Mess (modulo the completion of the above theorem which came later).

5.5. 3-dimensional hyperbolic geometry.

Our construction of the Poincaré model makes sense in any dimension n, except that we don't have such convenient complex coordinates. In this case, we take the disc $D^n = \{\underline{x} \in \mathbf{R}^n \mid ||\underline{x}|| < 1\}$. We scale the metric by the same factor $\lambda(\underline{x}) = \frac{2}{1-||\underline{x}||^2}$. We get a complete geodesic metric, ρ, and the isometry class of (D, ρ) is referred to as "hyperbolic n-space", $\mathbf{H}^n$. It is homogeneous and isotropic. Its ideal boundary, ∂D, is an $(n-1)$-sphere. Once again, bi-infinite geodesics are arcs of euclidean circles (or diameters) orthogonal to ∂D. More generally any euclidean sphere of any dimension, meeting ∂D othogonally intersects D in a hyperbolic subspace isometric to $\mathbf{H}^m$ for some $m < n$ — there is an isometry of (D, ρ) that maps it to a euclidean subspace through the origin, thereby giving us a lower dimensional Poincaré model.

There has been an enormous amount of work on 3-dimensional hyperbolic geometry, going back over a hundred years. One can construct lots of examples of polyhedra and tessellations, though the situation becomes more complicated. One can also use these to construct examples of compact hyperbolic 3-manifolds. The "Seifert-Weber space" is a nice example made out of a dodecahedron. In the 1970s people, such as Riley, started to notice that "many" 3-manifolds admitted hyperbolic structures. In the late 1970s Thurston revolutionised the subject by making the conjecture that every compact 3-manifold can be cut into pieces in a natural way so that each piece has a geometric structure. There are eight geometries in dimension 3, but by far the richest source of examples is hyperbolic geometry. In 2003, Perelman claimed a proof of Thurston's conjecture, building on earlier work of Hamilton.

This is a vast subject, we won't have time to look into here. We'll just mention a few facts relevant to group theory.

Let M be a closed hyperbolic 3-manifold, and $\Gamma = \pi_1(M)$. Then Γ is finitely generated (in fact, finitely presented), and $\Gamma \sim \mathbf{H}^3$. Thus any two such groups are q.i. There are however examples of closed hyperbolic 3-manifolds which do not have any common finite cover. (In contrast to the 2-dimenional case, the hyperbolic structure on a closed 3-manifold is unique, and it follows that covers are forced to respect hyperbolic metrics.) The verification of this involves algebraic number theory, so we won't describe it here. The point is that a closed hyperbolic 3-manifold has associated to it "stable trace field", a finite extention of the rationals, which one can compute. If these are different, then the groups are incommensurable. (It is likely that, in some sense, one would expect a "random" pair of hyperbolic 3-manifolds to be incommensurable.) This is therefore again in contrast to the 2-dimensional case, where any two surface groups of genus at least 2 are commensurable.

6. Hyperbolic groups.

In this section we will explore some of the basic properties of hyperbolic groups. The notion of a hyperbolic group was introduced by Gromov around 1985. They arise in many different contexts, and there is a sense in which a "generic" finitely presented group is hyperbolic.

For much of the discussion we will just deal with geodesic spaces. One can get quite a long way with just elementary metric space theory as we shall see.

6.1. Definition of a hyperbolic space.

Let (X, d) be a geodesic metric space.

Definition : A *(geodesic) triangle*, T, in X consists of three geodesics segments, (α, β, γ) cyclically connecting three point (called the *vertices* of T). We refer to the geodesics segments as the *sides* of T.

Definition : If $k \geq 0$, a point, $p \in X$ is said to be a *k-centre* for the triangle T if $\max\{d(p, \alpha), d(p, \beta), d(p, \gamma)\} \leq k$.

See Figure 6a. (In the figures in this section, geodesics are often depicted curved inwards, rather than as euclidean straight lines. This is meant to evoke the Poincaré model of the

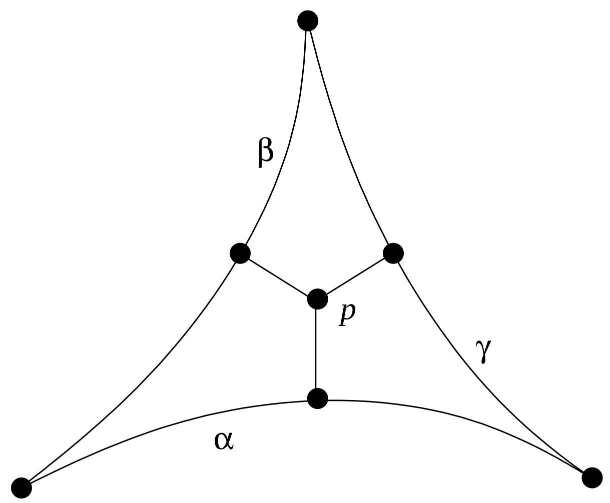

Figure 6a.

hyperbolic plane, to which hypebolic spaces have a more natural resemblence.)

Definition : We say that X is *k-hyperbolic* if every triangle has a k-centre.

Definition : We say that X is *hyperbolic* if it is k-hyperbolic for some $k \geq 0$. We refer to such a k as a *hyperbolicity constant* for X.

Examples.

(1) Any space of finite diameter, k, is k-hyperbolic.

(2) Any tree is 0-hyperbolic

(3) The hyperbolic plane $\mathbf{H}^2$ is $(\frac{1}{2}\log 3)$-hyperbolic.

(4) In fact, hyperbolic space $\mathbf{H}^n$ of any dimension is $(\frac{1}{2}\log 3)$-hyperbolic: any triangle in $\mathbf{H}^n$ lies in some 2-dimensional plane.

(5) Indeed, any complete simply connected riemanian manifold with curvatures bounded above by some negative constant $-\kappa^2 < 0$ is $(\frac{1}{2\kappa}\log 3)$-hyperbolic. For example, complex and quaternionic hyperbolic spaces are $(\frac{1}{2}\log 3)$-hyperbolic.

In (2) we can consider more general trees than those considered in Section 2. In particular, we can allow any positive length assigned to an edge (rather than just unit length). The result will always be a 0-hyperbolic geodesic space.

Remark: In fact, any 0-hyperbolic geodesic space is a more general sort of tree known as an "$\mathbf{R}$-tree". Here one can allow branch points (i.e. valence ≥ 3 points) to accumulate, so such a tree need not be a graph. (Indeed there are examples where every point is a branch point.) The theory of $\mathbf{R}$-trees was introduced by Morgan and Shalen and developed by Rips and many others, and it is now an important tool in geometric group theory.

Non-examples.

(1) Euclidean space, $\mathbf{R}^n$ for $n \geq 2$ is not hyperbolic.

(2) The 1-skeleton of the regular square tessellation of the plane is not hyperbolic. In fact, this example illustrates a slightly subtle point. It turns out that any three points of this graph can be connected by three geodesics so the triangle formed has a 1-centre (excercise). However not every triangle has this property. In fact, in this graph, we can have two geodesics between the same pair of points which go an arbitrarily long way apart before coming back together again.

6.2. Basic properties.

Before studing properties of a geodesic space, we make a couple of observations that hold in any geodesic metric space.

Let (X, d) be a metric space. We will often abbreviate $d(x, y)$ to xy. Given $x, y, z \in X$, write

$$\langle x, y \rangle_z = \frac{1}{2}(xz + yz - xy).$$

This is sometimes called the "Gromov product". The triangle inequality tells us this is non-negative. One way to think of it is as follows. Set $r = \langle y, z \rangle_x$, $s = \langle z, x \rangle_y$ and $r = \langle x, y \rangle_z$. Then

$$\begin{aligned} xy &= r + s \\ yz &= s + t \\ zx &= t + r. \end{aligned}$$

We can construct a "tripod" consisting of three edges meeting at a vertex of valence three and place the points x, y, z at the other endpoints of these edges (Figure 6b). If we assign the edge lengths r, s, t to these edges, we see that distances between x, y, z in X agree with those in the tripod. (This tripod might not be isometrically embedable in X.)

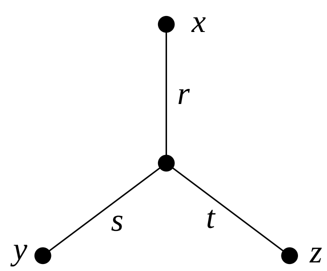

Figure 6b.

Another point to note is:

Lemma 6.1 : *If α is any geodesic from x to y, then $d(z,\alpha) \geq \langle x,y\rangle_z$.*

Proof : If $a \in \alpha$, then we have

$$\begin{aligned} xy &= xa + ax \\ xz &\leq xa + az \\ yz &\leq ya + az, \end{aligned}$$

and so $az \geq \langle x,y\rangle_z$. ♢

Note also that if z lies on any geodesic from x to y, then $\langle x,y\rangle_z = 0$.

Suppose now that (X,d) is k-hyperbolic. We prove a series of lemmas involving various constants. We aim to provide arguments that are fairly simple, rather that ones that will optimise the constants involved. With more careful arguments, one can probably do better in this regard.

Suppose that $T = (\alpha,\beta,\gamma)$ is a geodesic triangle. If p is any k-*centre*, we can find some $a \in \alpha$ with $ap \leq k$. Such a point a, is then a $2k$-centre for T.

Lemma 6.2 : *Suppose $x,y,z \in X$, and α is any geodesic connecting x to y. Then $d(z,\alpha) \leq \langle x,y\rangle_z + 4k$.*

Proof : Let $t = \langle x,y\rangle_z$. Let β,γ be geodesics from z to x and y respetively (Figure 6c). Let $a \in \alpha$ be a $2k$-centre for the triangle (α,β,γ). Thus

$$\begin{aligned} xa + ay &\leq xz + 4k \\ ya + az &\leq yz + 4k \\ xa + ay &= xy. \end{aligned}$$

Adding the first two of these and subtracting the third, we get $2az \leq 2t + 8k$, and so $az \leq t + 4k$ as required. ♢

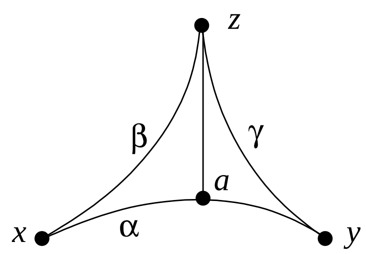

Figure 6c.

Corollary 6.3 : *If α and β are two geodesics connecting the same pair of points, then $\alpha \subseteq N(\beta, 4k)$ and $\beta \subseteq N(\alpha, 4k)$.*

Proof : Let the common endpoints be x and y, and suppose $x \in \beta$. Then $\langle x, y \rangle_z = 0$ and so by Lemma 6.2, $d(z, \alpha) \leq 4k$. This proves the first inclusion, and the other follows by symmetry. $\diamondsuit$

Thus in a hyperbolic space, any two geodesics with the same endpoints remain a bounded distance apart. We we also see that, up to an additive constant, we can think of the Gromov product, $\langle x, y \rangle_z$ as the distance between z and any geodesic from x to y.

Notation: Given any path α and points a, b on α, we write $\alpha[a, b]$ for the subpath of α between a and b.

The following terminology is not standard, but will be useful for our purposes.

Definition : A path α is t-taut if length$(\alpha) \leq xy + t$, where x, y are the endpoints of α.

Thus a geodesic is a 0-taut path. Also (exercise) any subpath of a t-taut path is t-taut.

We have the following generalisation of Lemma 6.3:

Lemma 6.4 : *Suppose α is a geodesic and β is a t-taut path with the same endpoints. Then:*
(1) $\beta \subseteq N(\alpha, \frac{1}{2}t + 4k)$, and
(2) $\alpha \subseteq N(\beta, t + 8k)$.

Proof : Let x, y be the endpoints of α.
(1) If $z \in \beta$, then $\langle x, y \rangle_z \leq t/2$, and so by Lemma 6.2, $d(z, \alpha) \leq \frac{1}{2}t + 4k$.
(2) Suppose $w \in \alpha$. By a connectedness argument using part (1), we can find some $z \in \beta$ a distance at most $\frac{1}{2}t + 4k$ from points a and b in α, on different sides of w. (Consider the closed subsets, $\beta \cap N(\alpha[x, w], \frac{1}{2}t + 4k)$ and $\beta \cap N(\alpha[y, w], \frac{1}{2}t + 4k)$. By (1) these cover β and so must intersect.) Thus $ab \leq t + 8k$ and $w \in \alpha[a, b]$, so w is a distance at most $\frac{1}{2}t + 4k$ from one of the points a or b (Figure 6d).

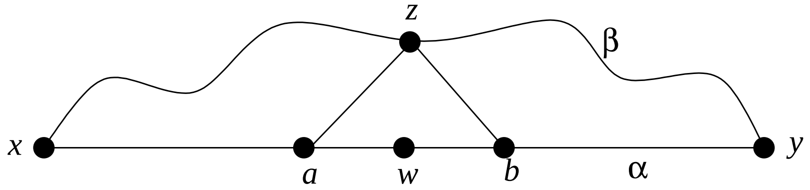

Figure 6d.

It follows that $wz \leq t + 8k$ as required. $\diamondsuit$

Lemma 6.5 : *If (α, β, γ) is a geodesic triangle, then $\alpha \subseteq N(\beta \cup \gamma, 6k)$.*

Proof : Let $a \in \alpha$ be a 2k-centre of (α, β, γ). This cuts α into two segments $\alpha[a, x]$ and $\alpha[a, y]$. Let δ be any geodesic rom z to a. Since $d(a, \beta) \leq 2k$, the path $\delta \cup \alpha[a, x]$ is $4k$-taut and so by Lemma 6.4(1), $\alpha[a, x] \subseteq N(\beta, 6k)$. Similarly, $\alpha[a, y] \subseteq N(\gamma, 6k)$. $\diamondsuit$

Remark: The concusion of Lemma 6.5 gives us an alternative way of defining hyperbolicity. Suppose (α, β, γ) is a geodesic triangle with $\alpha \subseteq N(\beta \cup \gamma, k')$ for some $k' \geq 0$, then by a connectedness argument (similiar to that for proving Lemma 6.4(2)), we can find some point $a \in \alpha$ a distance at most k' for both β and γ. This a will be a k'-centre from (α, β, γ). Thus we can define a space to be hyperbolic if for every geodesic triangle, each edge is a bounded distance from the union of the other two. This definition is equivalent to the one we have given, though the hyperbolicity constants involved my differ by a some bounded multiple.

6.3. Projections.

Suppose $x, y, z \in X$ and α is a geodesic connecting x to y. We describe a few different, but essentially equivalent ways of thinking of the notion of a "projection" of z to α.

(P1) One way, we have already seen, is to take geodesics β, γ from z to x and y respectively, and let $a \in \alpha$ be a $2k$-centre for the triangle (α, β, γ). A-priori, this might depend on the choice of β and γ. Here are another two constructions.

(P2) Let $b \in \alpha$ be the unique point so that $xb = \langle y, z \rangle_x$. It follows that $yb = \langle x, z \rangle_y$.

(P3) Choose some $c \in \alpha$ so as to minimise zc. This "neasest point" construction is the closest to what one normally thinks of as projection.

We want to show that these three constructions agree up to bounded distance.

To see this, first note that

$$\begin{aligned} xz &\leq xa + az \leq xz + 4k \\ yz &\leq ya + az \leq yz + 4k \\ xy &= xa + ay, \end{aligned}$$

and so we get $xa - 2k \leq \langle y, z \rangle_x \leq xa + 2k$. It follows that $ab \leq 2k$.

Now note that $zc = d(z, \alpha) = d(z, \alpha[c, x])$. Applying Lemma 6.2 (with α replaced by $\alpha[c, x]$), we see that $zc \leq \langle z, c \rangle_x + 4k$, and so $2zc \leq (zc + zx - cx) + 8k$ giving $zc + cx \leq zx + 8k$. Thus $\langle x, z \rangle_c \leq 4k$ and so by Lemma 6.2 again (with z replaced by c and α replaced by β) we get $d(c, \beta) \leq 4k + 4k \leq 8k$. Similarly, $d(c, \gamma) \leq 8k$. In other words, c is an $8k$-centre for (α, β, γ). We can now apply the argument of the previous paragraph again. The constants have got a bit bigger, and this time we get $bc \leq 8k$.

This shows the above three definitions of projections agree up to bounded distance, depending only on k. It is also worth noting that there is some flexibility in the definitions.

For example, if we took a to be any t-centre, or chose $c \in \alpha$ to be any point with $d(z,c) \leq d(z,\alpha) + t$, then we get similar bounds depending only on t and k.

One consequence of this construction is the following:

Lemma 6.6 : *Suppose that $x, y, z \in X$. Then a, b are t-centres of triangles with vertices x, y, z, then ab is bounded in terms of t and k.*

Proof : By Corollary 6.3, a t-centre of one triangle will be a $t + 4k$ centre of any other with the same vertices. We can therefore assume that a and b are centres of the same triangle. We can also assume that they lie on some edge, say α, of this triangle (replacing t by $2t$). The situation is therefore covered by the above discussion. $\diamondsuit$

Of course, one can explicitly calculate the bound in terms of t and k, though such calculations eventually become tedious, and for most purposes it is enough to observe that some formula exists.

We also note that a centre, a, for x, y, z, is decribed up to bounded distance by saying that $ax \leq \langle y, z \rangle_x + t$, $ay \leq \langle z, x \rangle_y + t$ and $ax \leq \langle x, y \rangle_z + t$, for some constant t.

Here is another consequence worth noting. Given $x, y, z \in X$, let a be a centre for x, y, z. Let δ, ϵ, ζ be geodesics connecting a to x, y and z respectively. Let τ be the "tripod" $\delta \cup \epsilon \cup \zeta$. This is a tree in X, with extreme points (valence 1 vertices) x, y, z. Note that distances in τ agree with distances in X up to a bounded constant. This is an instance of a much more general result about the "treelike" nature of hyperbolic spaces.

Notation. Given $x, y \in X$ we shall write $[x, y]$ for some choice of geodesic between x and y. If $z, w \in [x, y]$, we will assume that $[z, w] \subseteq [x, y]$.

Of course this involves making a choice, but since any two such geodesics remain a bounded distance apart, in practice this will not matter much. This is just for notational convenience. Formally we can always rephrase any statement to refer to a particular geodesic.

6.4. Trees in hyperbolic spaces.

The following expresses the "treelike" nature of a hyperbolic space:

Proposition 6.7 : *There is a function $h : \mathbf{N} \longrightarrow [0, \infty)$ such that if $F \subseteq X$ with $|F| = n$, then there is a tree, τ, embedded in X, such that for all $x, y \in F$, $d_\tau(x, y) \leq xy + kh(n)$.*

Here d_τ is distance measured in the tree τ. Note that we can assume that all the edges of τ are geodesic segments. We can also assume that every extreme (i.e. valence 1) point of τ lies in F. In this case, τ will be $kh(n)$-taut, in the following sense:

Definition : A tree $\tau \subseteq X$ is t-taut if every arc in τ is t-taut.

We will refer to a such a tree, τ, as a "spanning tree" for F.
To prove Propostion 6.7, we will need the following lemma:

Lemma 6.8 : *Suppose $x, y, z \in X$. Suppose that β is a t-taut path from x to y and that y is the nearest point on β to z. The $\beta \cup [y, z]$ is $(3t + 24k)$-taut.*

Proof : Let α be any geodesic from x to y (Figure 6e). By Lemma 6.5(2), $\alpha \subseteq N(\beta, t+8k)$. By hypothesis, $d(z, \alpha) = yz$, and so $d(z, \alpha) \geq yz - (t + 8k)$. Thus, by Lemma 6.2,

$$\begin{aligned} \langle x, y \rangle_z &\geq d(z, \alpha) - 4k \\ &\geq yz - t - 12k. \end{aligned}$$

That is, $xz + yz - xy \geq 2yz - 2t - 24k$, and so $xy + yz \leq xz + 2t + 24k$. It follows that

$$\begin{aligned} \text{length}(\beta \cup [y, z]) &\leq (xy + t) + yz \\ &\leq xz + 3t + 24k. \end{aligned}$$

$\diamondsuit$

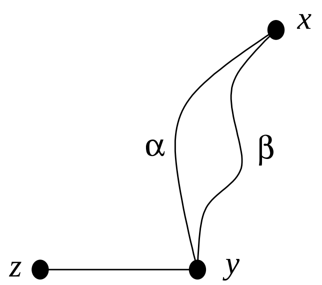

Figure 6e.

Corollary 6.9 : *Suppose that τ is t-taut tree and $z \in X$. Let $y \in \tau$ be a nearest point to z. Then $\tau \cup [y, z]$ is $(3t + 24k)$-taut.* $\diamondsuit$

Proof of Lemma 6.7: Let $F = \{x_1, x_2, \ldots, x_n\}$. Construct τ inductively. Set $\tau_2 = [x_1, x_2]$, and define $\tau_i = \tau_{i-1} \cup [y, x_i]$, where y is a nearest point to x_i in τ_{i-1} (Figure 6f). We now apply Corollary 6.9 inductively, and set $\tau = \tau_n$. $\diamondsuit$

We remark that this argument gives $h(n)$ exponential in n. In fact, one can show that the same construction gives $h(n)$ linear in n, but this is more subtle. One cannot do better than linear for an arbitrary ordering of the points of F (exercise). A different construction can be used to give a tree with $h(n) = O(\log(n))$, which is the best possible:

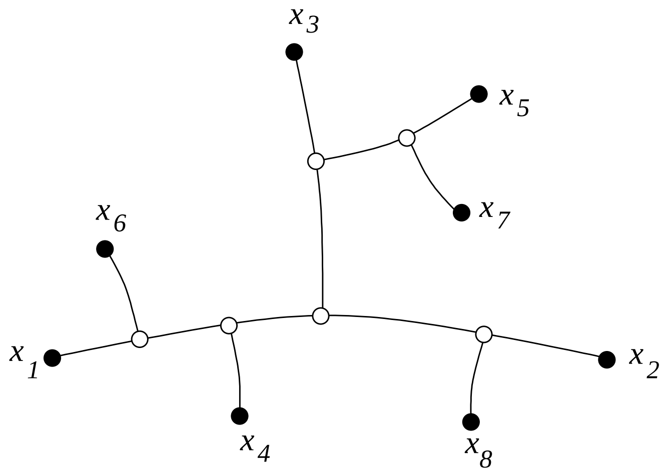

Figure 6f.

Exercise: Let F be a set of n equally spaced points around a circle of radius $r \geq \log(n)$ in $\mathbf{H}^2$. Then no spanning tree can be better than t-taut, where $t = O(r) = O(\log n)$. (Use the fact that the length of a circle of radius r is $2\pi \sinh r$.)

As far as I know, the following question remains open (even for $\mathbf{H}^2$):

Question: In the construction of the tree in Proposition 6.7, can one choose the order of the points $(x_i)_i$ so as always to give a tree with $h(n) = O(\log n)$?

Proposition 6.7 is very useful. We are frequently in a situation where we are dealing just with a bounded number of points. If we are only interested in estimating something up to an additive constant (depending on k), then we can assume we are working in a tree.

For many applications, it is enough to embed our set F is some tree τ, and do not need to know that τ is actually embedded in X. It is possible to construct such a tree by a more direct argument, though we won't describe the construction here.

6.5. The four-point condition.

Let us suppose that τ is a tree containing four points $x, y, z, w \in \tau$. One can see that, measuring distances in τ, we have

$$xy + zw \leq \max\{xz + yw, xw + yz\}.$$

Suppose, for example, that the arcs from x to y and from z to w meet in at most one point (Figure 6g).

In this case, we write $(xy|zw)$, and this situation we see that $xy + zw \leq xz + yw = xw + yz$. Whatever the arrangement of the three points, it is easily seen that at least one of $(xy|zw)$, $(xz|yw)$ or $(xw|yz)$ must hold, thereby verifying the above inequality.

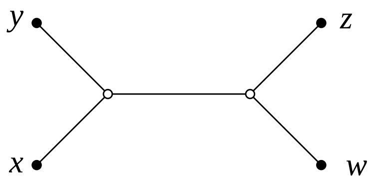

Figure 6g.

From this we can deduce:

Lemma 6.10 : *Given $k \geq 0$, there is some $k' \geq 0$ such that if X is a k-hyperbolic geodesic spaces, and $x, y, z, w \in X$, then*

$$xy + zw \leq \max\{xz + yw, xw + yz\} + k'.$$

Proof : By Lemma 6.7, we can find a tree τ, containing x, y, z, w, so that distances in τ agree with distances in X up to an additive constant $kh(4)$. We can now apply the above observation. $\diamondsuit$

As usual, k' is some particular multiple of k, which we could calculate explicitly (exercise). (In fact, there are more direct routes to this particular result that would probably give better constants.)

It turns out that hyperbolicity is characterised by this property. Let us suppose, for the moment, that (X, d) is any geodesic space and $k' \geq 0$ is some constant. We suppose:

$$(*) \qquad (\forall x, y, z, w \in X)(xy + zw \leq \max\{xz + yw, xw + yz\} + k').$$

.

Given $x, y, z \in X$ and a geodesic α from x to y, let $a \in \alpha$ be the point with $xa = \langle y, z \rangle_x$ (cf. the earlier discussion of projections).

Lemma 6.11 : *$xa + az \leq xz + k'$ and $yz + az \leq yz + k'$.*

Proof : Let $r = \langle y, z \rangle_x$, $s = \langle z, x \rangle_y$ and $t = \langle x, y \rangle_z$. Thus,

$$\begin{aligned} xy &= r + s \\ yz &= s + t \\ zx &= t + r \\ xa &= r \end{aligned}$$

and

$$ya = x.$$

Let
$$za = u.$$
We now apply $(*)$ to $\{x, y, z, a\}$. The three distance sums are
$$\begin{aligned} &r + s + u \\ &r + s + t \\ &r + s + t, \end{aligned}$$
and so $u \le t + k'$. But now $xa + az = r + u \le r + t + k' = xz + k'$ and $ya + az = s + u \le s + t + k' = yz + k'$. $\diamondsuit$

Lemma 6.12 : *In the above situation, let β and γ be geodesics from z to x and from z to y repectively. Then a is a $(3k'/2)$-centre for the triangle (α, β, γ).*

Proof : Let b be the projection of a to β in the above sense. Applying Lemma 6.11 to a and β (in place of z and α) we see that
$$\begin{aligned} ab + bx &\le ax + k' \\ ab + bz &\le az + k'. \end{aligned}$$
Adding we get
$$\begin{aligned} 2ab + (xb + bz) &\le xa + az + k' \\ 2ab + xz &\le xa + az + 2k' \\ &\le xz + 2k' \end{aligned}$$
applying Lemma 6.11 again to z and α. We see that $ab \le 3k'/2$. We have shown that $d(a, \beta) \le 3k'/2$.

Similarly $d(a, \gamma) \le 3k'/2$ as required. $\diamondsuit$

We have shown that under the assumption $(*)$ every triangle has a $(3k'/2)$-centre. Putting this together with Lemma 6.10, we get:

Proposition 6.13 : *For a geodesic metric space, the condition $(*)$ is equivalent to hyperbolicity.* $\diamondsuit$

We remark that $(*)$ makes no reference to geodesics, and so, in principle, makes sense for any metric space. Its main application, however is to geodesic spaces.

Remark: The "four point" condition $(*)$ is frequently given in the following equivalent form:
$$(\forall x, y, z, w \in X)(\langle x, y\rangle_w \ge \min\{\langle x, z\rangle_w, \langle y, z\rangle_w\} - k''$$
(where $k'' = k'/2$). Indeed this was the first definition of hyperbolicity given in Gromov's original paper on the subject.

6.6. Exponential growth of distances.

We observed in Section 5 that the length of a hyperboloc circle grows exponentally in the diameter. The following can be viewed as a more general expression of this phenomenon.

We fix a basepoint $p \in X$. We write $N(x,r) = \{y \in X \mid d(x,y) \leq r\}$, and $S(x,r) = \{y \in X \mid d(x,y) = r\}$. We write $N^0(x,r) = N(x,r) \setminus S(x,r)$.

Proposition 6.14 : *There are constants $\mu > 0$ and $K \geq 0$ such that for all $r \geq 0$, if α is a path in $X \setminus X \setminus N^0(p,r)$ connecting $x, y \in S(x,r)$, then* $\operatorname{length} \alpha \geq e^{\mu d(x,y)} - K$.

See Figure 6h.

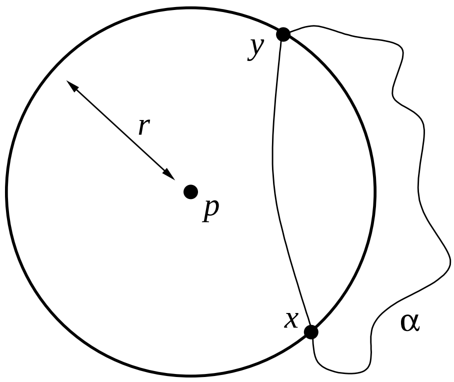

Figure 6h.

The idea behind the proof is that projections from a sphere to a smaller concentric sphere will tend to reduce distances by a a uniform factor less than 1. Because of the additive constants involved in the definition of hyperbolicity, we will need to express the argument as a discrete process.

Our proof will be use the following two related observations.

Lemma 6.15 : *For all sufficiently large h in relation to the hyperbolicity constant, if $x, y \in X$ with $d(x,y) \leq h/2$, then $d(x',y') \leq h$ where $x' \in [p,x]$ and $y' \in [p,y]$ with $d(p,x') = d(p,y')$.*

Lemma 6.16 : *For all sufficiently large h in relation to the hyperbolicity constant, if $x, y \in X$ with $d(p,x) = d(p,y)$ and $d(x,y) \leq 2h$, then $d(x',y') \leq h$, where $x' \in [p,x]$ and $y' \in [p,y]$ with $d(x,x') = d(y,y') = h$.*

We leave the proofs as an excersise — for example, either verify the statements in a spanning tree for the five points p, x, y, x', y', or else by a more direct argument by considering the triangle $[p, x], [p, y], [x, y]$.

Proof of Proposition 6.14 : Let $l = \text{length}\,\alpha$. We can assume that $l \geq 4h$. We can thus find some $n \in \mathbf{N}$ so that $2^{m-1}h \leq l \leq 2^{m-1}h$. Let $x = x_0, x_1, \ldots, x_{2^m} = y$ be a sequence of $2^m + 1$ points along α so that $d(x_i, x_{i+1}) \leq h/2$ for all i. Let $y_i \in [p, x_i]$ with $d(p, y_i) = r$. Thus $y_0^0 = x$, $y_{2^m}^0 = y$ and $y_i^0 \in S(r)$ for all i. By Lemma 6.15, $d(y_i^0, y_{i+1}^0) \leq h$ for all i. (Figure 6i.)

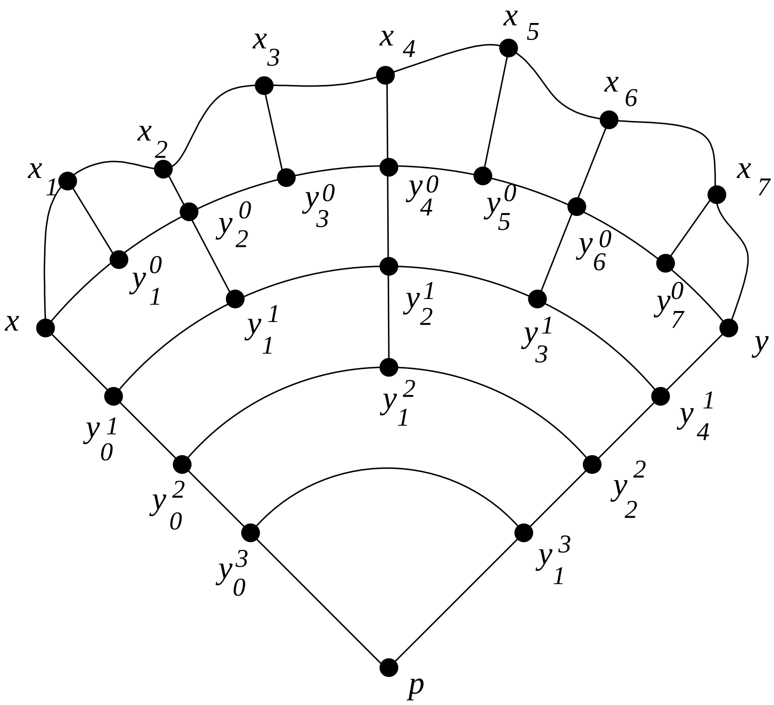

Figure 6i.

Now define a sequence $y_0^1, \ldots, y_{2^{m-1}}^1$ as follows. If $r \leq h$ set $y_i^1 = p$ for all i. If $r \geq h$, let $y_i^1 \in [p, y_{2i}^0]$ be the point with $d(y_i^1, y_{2i}^0) = h$. Now $d(y_{2i}^0, y_{2i+2}^0) \leq 2h$, and so by Lemma 6.16, we have $d(y_i^1, y_{i+1}^1) \leq h$. We now proceed inductively, each time eliminating half the points, and moving the others a distance h towards p (or possibly setting them all equal to p). For each $j = 1, 2, \ldots, m$, we get a sequence $y_0^j, \ldots, y_{2^{m-j}}^j$, with $d(y_i^j, y_{i+1}^j) \leq h$ for all i. We end up with a 2-point sequence, y_0^m, y_1^m. Note that $d(y_0^m, y_1^m) \leq h$. Now for all j, $d(y_0^j, y_0^{j+1}) \leq h$ and so $d(y_0^0, y_0^m) \leq mh$. Similarly, $d(y_{2^m}^0, y_1^m) \leq mh$. But $x = y_0^0$ and $y = y_{2^m}^0$, and so $d(x, y) \leq 2mh + h = (2m+1)h$, and so $m \geq (d(x, y) - h)/2h$.

We see that $l \geq 2^m m - 2h \geq 2^{(d(x,y)-h)/2h}h/4$. This is under the initial assumption that $l \geq 4h$. Thus, in general, we always get an inequality of the form $l \geq e^{\mu d(x,y)} - K$, where μ and K depend only on h, and hence only on the hyperbolicity constant k, as required. $\diamondsuit$

Remark : It turns out that the exponential growth of distances gives another formulation of hyperbolicity — essentially taking the conclusion of Proposition 6.14 as a hypothesis. We will not give a precise formulation of this here.

6.7. Quasigeodesics.

The notion of a quasigeodsesic path is another fundamental notion in geometric group theory. The following definition will make sense in any metric space, though it is mainly of interest in geodesic spaces. In what follows we shall abuse notion slightly and identify a path in X with its image as a subset of X, (even if the path is not injective). Given two point, x, y in a path α, we shall write $\alpha[x, y]$ for the segment of α beween x and y.

Definition : A path, β, is a (λ, h)-*quasigeodesic*, with respect to constants $\lambda \geq 1$ and $h \geq 0$, if for all $x, y \in \beta$, $\mathrm{length}(\beta[x, y]) \leq \lambda d(x, y) + h$. A *quasigeodesic* is a path that is (λ, h)-quasigeodesic for some λ and h.

In other words, it takes the shortest route to within certain linear bounds. Note that a $(1, h)$-quasigeodesic is the same as an h-taut path.

We now suppose that (X, d) is k-hyperbolic again. A basic fact about quasigeodesiscs is that they remain a bounded distance apart (cf. Lemma 6.4).

Proposition 6.17 : *Suppose that α is a geodesic, and β is a (λ, h)-quasigeodesic with the same endpoints. Then $\beta \subseteq N(\alpha, r)$ and $\alpha \subseteq N(\beta, r)$ where r depends only on λ, h, and the hyperbolicity constant k.*

Proof : We first show that α lies a bounded distance from β. (In other words, we proceed in the opposite order from Lemma 6.4.) Let a, b be the endpoints of α.

Choose $p \in \alpha$ so as to maximise $d(p, \beta) = t$, say. Let $a_0, a_1 \in [a, p]$ be points with $d(p, a_0) = t$ and $d(p, a_1) = 2t$. The point a_0 certainly exists, since $d(p, a) \geq t$. If $d(p, a) < 2t$, we set $a_1 = a$ instead. Now $d(a_1, \beta) \leq t$, and so there is some point $a_2 \in \beta$ with $d(a_1, a_2) \leq t$. If $a_1 = a$, we set $a_2 = a$. We similarly define points b_0, b_1, b_2 (Figure 6j).

Note that $d(a_2, b_2) \leq 6t$. Let $\delta = \beta[a_2, b_2]$, and let $\gamma = [a_0, a_1] \cup [a_1, a_2] \cup \delta \cup [b_2, b_1] \cup [b_1, b_0]$. Note that $\gamma \cap N^0(p, t) = \emptyset$. Since β is quasigeodesic,

$$\begin{aligned}\mathrm{length}\,\delta &\leq \lambda d(a_2, b_2) + h \\ &\leq 6\lambda t + h,\end{aligned}$$

and so

$$\begin{aligned}\mathrm{length}\,\gamma &\leq 4t + \mathrm{length}\,\delta \\ &\leq (6\lambda t + 4)t + h.\end{aligned}$$

On the other hand, $d(a_0, b_0) = 2t$, and γ does not meet $N^0(p, t)$. Thus applying Proposition 6.17, we get

$$\mathrm{length}\,\gamma \geq e^{\mu(2t)} - K.$$

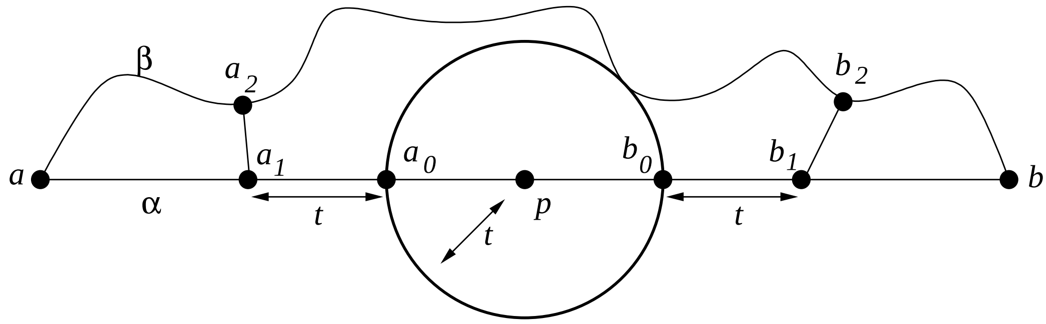

Figure 6j.

Putting these together we get

$$e^{2\mu t} \leq (6\lambda + 4)t + h + K$$

which places an upper bound of t in terms of λ, h, μ, K, and hence in terms of λ, h and k.

To show that β lies in a bounded neighbourhood of α, one can now use a connectedness argument similar to that use in Lemma 6.4 (with the roles of α and β interchanged). $\diamondsuit$

Of course (after doubling the constant r) Propositon 6.17 applies equally well to two quasigeodesics, α and β with the same endpoints.

Using Proposition 6.17, we see that we can formulate hyperbolicity equally well using quasigeodesic triangles, that is where α, β, γ are assumed quasigeodesic with fixed constants. In particular, we note:

Lemma 6.18 : *Any (λ, k)-quasigeodesic triangle (α, β, γ) has a t-centre, where t depends only on λ, h and k.*

Proof : Let $(\alpha', \beta', \gamma')$ be a geodesic triangle with the same vertices. Applying Lemma 6.17, we see that any k-centre of $(\alpha', \beta', \gamma')$ will be a $(k+r)$-centre for (α, β, γ). $\diamondsuit$

6.8. Hausdorff distances.

Before continuing we make the following useful definition.

Definition : Suppose P, Q are subsets of a metric space (X, d). We define the *Hausdorff distance* between P and Q as the infimum of those $r \in [0, \infty]$ for which $P \subseteq N(Q, r)$ and $Q \subseteq N(P, r)$.

Exercise: This is a pseudometric on the set of all bounded subsets of X. (It is only a pseudometric, since the Hausdorff distance between a set and its closure is 0.) Restricted to the set of closed subsets of X, this is a metric.

Note that Proposition 6.17 implies that the Hausdorff distance between two quasi-geodesics with the same endpoints is bounded in terms of the quasigeodesic and hyperbolicity constants.

6.9. Quasi-isometry invariance of hyperbolicity.

This is the key fact that makes the theory of hyperbolic groups work.

Suppose that (X, d) and (X', d') are geodesic spaces and that $\phi : X \longrightarrow X'$ is a quasi-isometry. We would like to say that the image of a geodesic is a quasi-geodesic, but this is complicated by the fact that geodesics are not assumed continuous. The following technical discussion is designed to get around that point.

Fix some $h > 0$. Suppose that α is a geodesic in X from x to y. Choose points $x = x_0, x_1, \ldots, x_n = y$ along α so that $d(x_i, x_{i+1}) \leq h$ and $n \leq l/h \leq n + 1$. Let $y_i = \phi(x_i) \in X'$. Let $\bar{\alpha} = [y_0, y_1] \cup [y_1, y_2] \cup \cdots \cup [y_{n-1}, y_n]$.

Exercise: If α is a geodesic in X and $\bar{\alpha}$ constructed as above, then $\bar{\alpha}$ is quasigeodesic, and the Hausdorff distance between $\bar{\alpha}$ and $\phi(\alpha)$ is bounded. As usual, the statement is *uniform* in the sense that the constants of the conclusion depend only on those of the hypotheses and our choice of h.

We are free to choose h however we wish, though it may be natural to choose it in relation to the other constants of a given argument, such as the constant of hyperbolicy.

Theorem 6.19 : *Suppose that X and X' are geodesic spaces with $X \sim X'$, then X is hyperbolic if and only if X' is.*

Proof : Let $\phi : (X, d) \longrightarrow (X', d')$ be a quasi-isometry and suppose that X' is k-hyperbolic. Let (α, β, γ) be a geodesic triangle in X. Let $\bar{\alpha}, \bar{\beta}, \bar{\gamma}$ be the quasigeodesics a bounded distance from $\phi(\alpha), \phi(\beta), \phi(\gamma)$ as constructed above. By Lemma 6.18, $(\bar{\alpha}, \bar{\beta}, \bar{\gamma})$ has a t-centre, q, where t depends only on k and the quasi-geodesics constants. Since $\phi(X)$ is cobounded, there is some $p \in X$ with $\phi(p)$ a bounded distance from q. Now $\phi(p)$ is a bounded distance from each of $\phi(\alpha)$, $\phi(\beta)$ and $\phi(\gamma)$. It now follows that p is a bounded distance from each of α, β, γ. In other words p is a centre for the triangle (α, β, γ). $\diamondsuit$

In fact, we see that the hyperbolicity constant of X depends only on that of X' and the quasi-isometry constants. (In the construction of $\bar{\alpha}, \bar{\beta}, \bar{\gamma}$ it is natural to take $h = k$. In this way, we get linear bounds between the hyperbolicity contants.)

Theorem 6.19 has some immediate consequences. For example we see:

(1) If $m, n \geq 2$, then $\mathbf{R}^m \not\sim \mathbf{H}^n$.

(2) If $n \geq 2$, then $\mathbf{R}^n$ is not quasi-isometric to any tree.

In particular, we get another proof that $\mathbf{R}^2 \not\sim \mathbf{R}$ and that $\mathbf{R}^2 \not\sim [0, \infty)$.

6.10. Hyperbolic groups.

We are finally ready for the following fundamental notion:

Definition : A group Γ is *hyperbolic* if it is finitely generated and its Cayley graph $\Delta(\Gamma)$ is hyperbolic.

By Theorem 3.3 and Theorem 6.19 this is well defined — it doesn't matter which finite generating set we take to construct the Cayley graph.

We note:

Lemma 6.20 : *Suppose that Γ acts properly discontinously cocompactly on a proper hyperbolic (geodesic) space, then Γ is hyperbolic.*

Proof : By Theorem 3.5, Theorem 3.6 and Theorem 6.19. $\diamondsuit$

Examples:
(1) Any finite group.
(2) Any virtually free group
(3) The fundamental group of any compact hyperbolic manifold. Note that if $\Gamma = \pi_1(M)$, where M is compact hyperbolic, then Γ acts properly discontinously cocompactly on $\mathbf{H}^n$.
(4) In particular, if Σ is any compact (orientable) surface of genus at least 2, then $\pi_1(\Sigma)$ is hyperbolic.

Non-examples:
(1) $\mathbf{Z}^n$ for any $n \geq 2$.
(2) It turns out that a hyperbolic group cannot contain any $\mathbf{Z}^2$ subgroup, so this fact provides many more non-examples. For example, many matrix groups $SL(n, \mathbf{Z})$ etc., knot groups (fundamental groups of knot complements), mapping class groups, braid groups etc. This is not the only obstruction, however.

6.11. Some properties of hyperbolic groups.

This is all we have time to deal with systematically in this course. We shall finish off by listing a few interesting directions currently being pursued. This list is by no means complete.

6.11.1. Subgroups.

(S1) Suppose g is an infinite order element of a hyperbolic group, Γ, so that $\langle g \rangle \cong \mathbf{Z}$. Then $\langle g \rangle$ is a "quasiconvex subgroup". Here this means that if $x \in \Delta(\Gamma)$, then the bi-infinite path $\alpha = \bigcup_{n \in \mathbf{Z}} [g^n x, g^{n+1} x]$ is quasi-geodesic. This is, in fact, the same as saying that the "stable length" $||g|| = \lim_{n\to\infty} \frac{1}{n} d(x, g^n x)$ is positive. (There are many examples of non-hyperbolic groups where this fails.)

(S2) Suppose that, in (S1), $h \in \Gamma$ is another element commuting with g. The bi-infinite path $h\alpha = \bigcup_{n\in\mathbf{Z}}[g^n hx, g^{n+1}hx]$ is a finite Hausdorff distance from α. In fact, using Proposition 6.17, one can show that the Hausdorff distance between α and $h\alpha$ is uniformly bounded, that is, it depends only on the quasi-geodesic constants of α, and not on h. But since $\Delta(\Gamma)$ is locally finite, there are only finitely many possibilities for $h\alpha$. As a result, one can show:

A hyperbolic group cannot contain any subgroup isomorphic to $\mathbf{Z}^2$.

(S3) A "Baumslag-Solitar group" is a group of the form $B(m,n) = \langle g, h \mid g^m h = hg^n \rangle$, where $m, n \geq 1$. (Note that $B(1,1) = \mathbf{Z}^2$.) By a similar argument to (S2), one can show, in fact, that:

A hyperbolic group cannot contain any Baumslag-Solitar subgroup.

(S4) Any hyperbolic group Γ that is not finite or virtually cyclic contains a free subgroup of rank 2, and hence free subgroups of any countable rank. (In particular, Γ has "exponential growth".) The usual way to construct such subgroups uses the so called "small cancellation theory", which long predates the invention of hyperbolic groups (see (F5) below).

(S5) It is a non-trivial question as to which groups can be embedded in hyperbolic groups. In (S4) we saw examples of non–f.g. subgroups of hyperbolic groups: free groups of infinite rank. One can construct hyperbolic groups that contain f.g. subgroups that are not f.p. One can also construct f.p. subgroups which are not hyperbolic, though these constructions become increasingly complicated.

(S6) A hyperbolic group contains only finitely many conjugacy classes of finite subgroups. To see this one can argue as follows. Suppose G is a finite subgroup of the hyperbolic group, Γ. Let a be any vertex of the Cayley graph $\Delta(\Gamma)$. Let $r \in \mathbf{N}$ be minimal such that the orbit, Ga, is contained in $N(b,r)$ for some vertex b of $\Delta(\Gamma)$. Now let B be the set of all vertices b with $Ga \subseteq N(b,r)$. Thus B is G-invariant, and an exercise in hyperbolicity shows that it has diameter bounded in terms of the hyperbolicity constant. There are thus only finitely many possiblities for B up to the Γ-action on $\Delta(\Gamma)$, and it follows that there are only finitely many possibilities for Γ up to conjugacy.

(S7) It is an open question as to whether every hyperbolic group is virtually torsion-free.

6.11.2. Finiteness and computablility properties.

(F1) By hypothesis, a hyperbolic group is f.g. One can show that any hyperbolic group is f.p.

(F2) If Γ is any hyperbolic group, then one can construct a locally finite contractible simplicial complex, K, (the "Rips complex") such that Γ acts properly discontinuously cocompactly on K. This is a strong finiteness condition. For example, a group acts p.d.c. on a locally finite connected complex if and only if it is f.g. (Here we could take K to be a graph. The case of free p.d.c. actions was discussed in Section 2.) Moreover, a group acts freely p.d.c. on a locally finite simply-connnected complex if and only if it is f.p. Here we are assuming that K is contractible, which is equivalent to saying that the homototpy groups, $\pi_n(K)$, are trivial for all $n \in \mathbf{N}$.

(F3) Suppose that Γ is a group which acts p.d.c. on a finite dimensional locally-finite contractible complex, and that Γ has no Baumslag-Solitar subgroups. Is Γ hyperbolic? This seems to be an open question (due to Bestvina) In other words are the conditions (S3) and (F2) together the only obstructions to hyperbolicity?

(F4) The fact that Γ is finitely presented can be strengthenned in another direction to a "linear isoperimetric" inequality. One can construct a p.d.c. action of Γ on a locally finite 2-dimensional simplicial complex K with the property that if α is a curve in the 1-skeleton of length n, then α bounds a disc in K (not nescessarily embedded) meeting at most $f(n)$ 2-simplices, where f is a linear function. One can give an equivalent algebraic statement. Fix any finite presentation of Γ. Suppose that w is a word in the generators an their inverses representing the identity in Γ. Then we can reduce w to the trivial word (of length 0) by repeatedly applying the relations. A linear isoperimetric inequality says that we only need to do this at most $f(n)$ times, where n is the length of w and $f : \mathbf{N} \longrightarrow \mathbf{N}$ is linear. It turns out that:

A group is hyperbolic if and only if it has a linear isoperimetric function.

(In fact, a subquadratic isoperimetric inequality is sufficient.)

(F5) We can make the following additional remarks. It turns out that isoperimetic inequalities of this sort are q.i. invariants (thereby giving a different proof that hyperbolicity is q.i. invariant.) The group $\mathbf{Z}^2$ has a quadratic, but not a linear isoperimetric inquality. The Heisenberg group has a cubic inequality. Other groups have exponential inqualities (or worse).

One can show that, for a f.p. group, a (sub)computable isoperimetric inequality is equivalent to solvability of the word problem. (It puts a computable bound on the work we need to do to check whether or not a word can be reduced to the trival word.) This shows that a solvable word problem is q.i. invariant. It also shows that the word problem in a hyperbolic group is solvable. In fact (though this is not an immediate consequence) it can be solved in linear time. The linear time algorithm is the "Dehn algorithm". Back in the 1920s Dehn used ideas of hyperbolic geometry to show that the word problem in a surface group is solvable. This was the beginning of "small cancellation theory" referred to in (S4). The same basic idea applies to general hyperbolic groups.

(F6) A lot more can be said in relation to computablity. For example, it turns out that a hyperbolic group is "automatic" in the sense of Thurston. This was essentially shown by Cannon, before the either of the notions "hyperbolic" or "automatic" were formally defined. Automaticity is a formal criterion which implies that many calculations can be carried out very efficiently. In particular, the word problem is solvable.

6.11.3. Boundaries.

(B1) Let X be a proper hyperbolic space. A "ray" in X is a semi-infinite geodesic. We say that two rays are "parallel" if the Hausdorff distance between them is finite. This is an equivalence relation, and we write ∂X for the set of equivalence classes. (We can equivalently use quasi-geodesic rays.) One can put a topology on ∂X — informally two rays are close in this topology if they remain close over a long distance in X. It turns out that ∂X is compact and metrisable. We refer to ∂X as the "(Gromov) boundary" of X. Any q.i. $\phi : X \longrightarrow Y$ induces a homeomorphism $\partial X \longrightarrow \partial Y$. Thus the homeomorphism type of ∂X is a quasi-isometry invariant. In particular, it makes sense to talk about the boundary, $\partial\Gamma$, of a hyperbolic group Γ.

(B2) The boundary of a compact space is empty. Thus $\partial(\text{finite group}) = \emptyset$.

(B3) The real line has two boundary points — one for each end. The boundary of the group $\mathbf{Z}$, or any virtually cyclic group, is thus the two-point space.

(B4) If $p \geq 3$, then ∂T_p is a cantor set. Thus ∂F_n is a cantor set for all $n \geq 2$.

(B5) ∂H^n can be identified with the ideal boundary we have already defined — the boundary of the Poincaré disc. It is thus homeomorphic to S^{n-1}. It follows that:

If $\mathbf{H}^m \sim \mathbf{H}^n$ then $m = n$.

Note is particular that the fundamental group of a compact surface cannot be q.i. to the fundamental group of a hyperbolic 3-manifold.

(B6) The work of Tukia, Gabai, Casson and Jungreis referred to earlier shows that if the boundary of a hyperbolic group is homeomorphic to a circle, then the group is a virtual surface group. Cannon asked whether a hyperbolic group with boundary a 2-sphere is a virtual hyperbolic 3-manifold group. This question remains open. The analogous assertion certainly fails in higher dimensions. For example a the fundamental group of a cocompact complex hyperbolic 4-manifold is hyperbolic and has boundary a 3-sphere, but it does not admit any p.d.c. action on (real) hyperbolic 4-space.

(B7) A hyperbolic group Γ acts by homeomorphism on the boundary $\partial\Gamma$. In fact this has a particular dynamical property: it is a "uniform convergence group". This means that induced action on the space of distinct triples (the configuration space of 3-element subsets of the boundary) is p.d.c. The notion of a convergence group was introduced by Gehring and Martin and explored by a number of people, such as Tukia. It turns out (Bowditch) that one can characterise hyperbolic groups in these terms:

If a group acts as a uniform convergence group on a compact metrisable space with no isolated points, then the group is hyperbolic, and the space is equivariantly homeomoprhic to the boundary.

(B8) It turns out that any compact metrisable topological space is homeomorphic to the boundary of some proper hyperbolic space. However, there are constraints on what kinds of spaces can arise as boundaries of hyperbolic groups. In some sense, the "generic boundary" is a Menger curve, but there are many other examples. One can get a lot of information about the algebraic structure of a hyperbolic group from the topology of its boundary.

6.11.4. Other directions.

There are many other directions in the study of hyperbolic groups which I have not had time to mention. The "JSJ splitting" theory of Sela, for example, which has inspired similar results for much wider classes of groups. The Markov property of the boundary, the "geodesic flow" of a hyperbolic group, bounded co-homology, the Novikov conjecture etc.

The subject of "relatively hyperbolic" groups is very fashionable at the moment. These include fundamental groups of finite-volume hyperbolic manifolds, amalgamated free products of groups over finite groups, the "limit groups" defined by Sela in relation to the Tarski problem etc.

There are many naturally arising spaces that are hyperbolic that do not stem directly from groups. For example the Harvey curve complex associated to a surface was shown to be hyperbolic by Masur and Minsky. This has many implications, for example to the study of the mapping class groups.

7. Notes.

General.

At present there seems to be no systematic formal introduction to geometric group theory as such, though the key ideas with many interesting examples are described in de la Harpe's book [Harp]. The fundamental notions are also described in [BriH]. Earlier, more traditional treatments of combinatorial group theory are [MagKS] and [LS]. A list of open problems in geometric group theory can be found in [Bes3].

The theory of non-positively curved spaces grew out of work of Aleksandrov, Toponogov, Busemann and more recently, Gromov. The standard introductory text is [BriH].

7.0. Section 0 :

An overview of basic "small cancellation" theory can be found in Strebel's appendix to [GhH]. The paper [ScW] was influential in introduction topologial ideas into group theory. The works of Stallings and Dunwoody were influential in introducing methods from 3-manifold theory — see the notes on Section 3. Thurston gave an outline of his geometrisation programme in [Th].

7.1. Section 1 :

The basic material of most of this section is standard. More detailed accounts can be found in [MagKS] and [LS]. One important result that illustrates the power of combinatorial methods is Higman's embedding theorem [Hi] (see also [LS]).

A geometrical picture of the Heisenberg group is described in [GhH].

A brief survey of the Andews-Curtis conjecture can be found in [BuM].

7.2. Section 2 :

Again, much of this material can be found in [MagKS], [LS] and other introductory texts.

7.3. Section 3 :

The significance of quasi-isometries, in particular, a version of the the key result, Theorem 3.6, was known to Efremovich, Schwarz amd Milnor in the 1960s. The subject was developed through work of Gromov, see for example [Gro2] and [Gro3]. A more general version of Proposition 3.1 is proven in [Gro4]. A proof of Theorem 3.6 (sometimes called the "Schwarz-Milnor Lemma") along the lines given here can be found in [BriH].

The statement that a group quasi-isometric to $\mathbf{Z}$ is virtually $\mathbf{Z}$ is an immediate consequence of the result of Hopf [Ho] that a two-ended finitely generated group is virtually $\mathbf{Z}$ — at least given the relatively simple fact that the property of being "two-ended" is quasi-isometry invariant. Here "two-ended" can be taken to mean that the complement of any sufficiently large finite subgraph of the Cayley graph has exactly two unbounded components.

A proof of the Borsuk-Ulam theorem can be found in many texts on topology, for example [Ar].

Gromov's theorem on groups of polynomial growth is given in [Gro1]. It uses the solution to Hilbert's fifth problem by Montgomery and Zippin [MonZ]. The "Gromov-Hausdorff" limit argument used in [Gro1] can be conveniently expressed in terms of asymptotic cones [Gro3], which have, in themselves become an important tool in geometric group theory, see for example [Dr]. While there are variations on the theme, this seems to be essentially the only proof known. A quite different approach to deal specifically with virtually abelian groups has been given by Shalom [Sha].

The results of Stallings [St] and Dunwoody [Du] relating to group splittings are good examples the adaptations of ideas from 3-manifolds.

Alonso gives an account of the quasi-isometry invariance of isoperimetric inqualities, and Shapiro's observation conserning the invarians of the word problem, in [Alo].

7.4. Section 4 :

Just about any introductory text on topology will have an account of fundamental groups, covering spaces etc. Our "nice" spaces can all be given the stucture of a simplicial complex. A deeper systematic treatment of such complexes is given in [Sp].

A combinatorial proof of Theorem 4.1 can be found in [LS].

7.5. Section 5 :

Introductory texts on hyperbolic geometry include [Iv] and [An], and a general introduction is also included in [Bea]. The book [W] gives an overview of this subject in connection with Thurston's programme. See also [CanFKP]. The foundational principles of non-euclidean geometry with many historical references can be found in [Gre].

An account of Poincaré's theorem for tessellations in dimension 2 is given in [Bea].

Perelman's account of geometrisation in given in [Pe1,Pe2]. A commentary can be found in [KlL], and a survey in [Mor].

With regards to the characterisation of virtual surface groups, a seminal piece of work was Mess's paper on the Seifert conjecture [Me]. This used earlier work of Tukia [Tu1], but left open the particular and difficult case of a virtual triangle group. This was resolved in subsequent and independent work of Gabai [Gaba] and Casson and Jungreis [CasJ]. To take care of the "euclidean" case, Mess relies on the theorem of Varopoulos [V] that a group with a recurrent random walk is virtually $\mathbf{Z}^n$ for $n = 0, 1, 2$. This in turn relies on Gromov's result of polynomial growth (see notes on Section 3). An argument that bypasses this, and gives some other characterisations of virtual surface groups, can be found in [Bow5]. Some of the results therein have been generalised by Kleiner.

For introductions to Teichmüller theory see [Ab] or [ImT].

The Mostow rigity theorem [Mo] tells us that any finite-volume hyperbolic structure on a 3-manifold is unique. The "stable trace field" of such a manifold $M = \mathbf{H}^3/\Gamma$ is the field generated by the squares of traces of elements of $\Gamma \subseteq PSL(2, \mathbf{C})$. A consequence of Mostow rigidity and a little algebraic geometry is that such a field is a finite extention of the rationals. It turns out to be a commensurability invariant [Re]. It is not hard to find explicit examples of closed hyperbolic 3-manifolds with different stable trace fields. There is a major project of enumerating small volume hyperbolic 3-manifolds and computing their stable trace fields and other invariants, see [CouGHN].

7.6. Section 6 :

Gromov introduced the notion of a hyperbolic group in [Gro2]. Several expositions of various aspects of this work appeared in the few years that followed: see [GhH], [CooDP], [Sho] or [Bow1]. Since then, the subject has developed in many different directions, though there seems to have been no new systematic general introduction to the subject. Some aspects of hyperbolic groups are discussed in [BriH] and in [Harp].

An introduction to complex hyperbolic geometry can be found [Go].

The notion of an $\mathbf{R}$-tree was introduced by Morgan and Shalen [MorS], in order to prove certain compactness results that formed part of Thurston's work on hyperbolic 3-manifolds. A more geometric approach to their construction was described by Bestvina [Be1]. The subject was then developed by Rips, and elaborations and generalisations of that work can be found in [GaboLP] and [BeF]. $\mathbf{R}$-trees have now become a central tool in geometric group theory. Surveys can be found in [Pau,Bes2] and a general introduction in [Chi]. A key point is that the asymptotic cone (see notes on Section 3) of a hyperbolic space is an $\mathbf{R}$-tree (see for example [Dr]). A typical $\mathbf{R}$-tree can be a quite complicated object. For example, we note that for any cardinal $c \geq 2$ there is a unique complete $\mathbf{R}$-tree with every point of valence c [DyP].

A discussion of spanning trees that approxiamate disances in a hyperbolic space is given in [Gro2], and some elaborations are described in [Bow1], including a construction of logarithmic spanning trees.

Proposition 6.17 is a generalisation of the corresponding lemma for hyperbolic space which is a standard ingredient for the argument proving Mostow rigidity (see notes on Section 5). The proof we present here is based on that given in [Sho].

7.6.1. Subgroups of hyperbolic groups.

The fact that the stable length of an infinite-order element is positive can be found, for example, in [GhH]. In fact, it turns out that they are uniformly rational [Gro2,Del]. For a general group, the stable length of an infinite order element might be 0: consider for example the centre of the Heisenberg group.

An example of a hyperbolic group with a finitely generated subgroup that is not finitely presented is given in [BowM], and is based on a related example in [Kap-mP].

7.6.2. Finiteness conditions.

An account of the Rips complex can be found in [GhH].

Gromov outlined a proof that a subquadratic isoperimetric inequality implies hyperbolicity in [Gro2], and this argument was elaborated upon in [CooDP]. Other arguments are given in [O,Pap,Bow2]. Examples of non-hyperbolic groups with more exotic isoperimetric inequalities (or "Dehn functions") are given in [Bri,BraB] and [SaBR]. This subject has expanded in many directions since.

Automatic structures provide a link with the theory of formal languages. The standard introductory text is [ECHLPT]. Cannon's argument, which can now be interpreted as a proof that a hyperbolic group is automatic, appeared in [Can].

7.6.3. Boundaries.

A general survey of boundaries of hyperbolic groups is [Kap-iB]. A seminal article on the subject was [BesM].

Convergence groups were introduced in the context of Kleinian group by Gehring and Martin [GeM]. A general discussion, applicable to the boundaries of hyperbolic groups is given in [Tu2]. The topological characterisation of a hyperbolic group as a uniform convergence group is given in [Bow4].

The fact that any compact metrisable (topological) space can be realised as the boundary of a proper hyperbolic space can be seen as follows. First embed the space in the unit sphere of a separable Hilbert space. We can view this sphere as the boundary of a Klein model for an infinite dimensional hyperbolic geometry. We take the euclidean convex hull of our set, which is also the hyperbolic convex hull. We use the fact that the convex hull of a compact subset of a Banach space is compact. Thus the convex hull gives us a proper hyperbolic space compactified by our original set. This set is then also the Gromov boundary. Delails are left to the reader familiar with the Klein model of hyperbolic space.

A discussion of "generic" properties of hyperbolic groups is given in [Cha]. This uses the notion of a "geodesic flow" (see below).

7.6.4. Other directions.

The JSJ splitting was introduced by Sela [Se1]. It is another example of a construction inspired by 3-manifold theory, in particular, work of Waldhausen, Johanson [Jo] Jaco and Shalen [JaS]. An account for hyperbolic groups, via boundaries, is given in [Bow3]. There are a number of generalisations, for example, [RiS,DuS,FuP].

An account of the geodesic flow on a hyperbolic group is given in [Mine], and connections with the Baum-Connes conjecture in [MineY]. For earlier work on the Novikov conjecture for hyperbolic groups, see [ConM] and [KasS]. (The former requires a version of the geodesic flow.)

Sela's work on the Tarski problem appears in a series of articles, starting with [Se2]. It makes much use of the JSJ splittling in constructing "Makanin-Razborov" trees. He characterises groups with the same first order theory as free groups as "limit groups". Such groups were shown to be relatively hyperbolic in [Da] (see also [Ali]).

Accounts of relatvely hyperbolic groups can be found in [Fa], [Sz] and [Bow7]. A topological characterisation in terms of convergence groups is given in [Y]. Many results about hyperbolic groups have now been generalised to relatively hyperbolic groups.

Harvey introduced the curve complex in [Harv]. It was shown to be hyperbolic in [MasM] (see also [Bow6,Ham]). This fact was central to the proof of Thurston's ending lamination conjecture [Mins,BroCM], and has wider implications for Teichmüller theory and mapping class groups. This is a particularly active area at the moment. See [Bow8] for a survey of some of this material.

References

[Ab] W.Abikoff, *The real analytic theory of Teichmüller space* : Lecture Notes in Mathematics No. 820, Springer (1980).

[Ali] E.Alibegovic, *A combination theorem for relatively hyperbolic groups* : Bull. London Math. Soc. **37** (2005) 459–466.

[Alo] J.M.Alonso, *Inégalités isopérimetriques et quasi-isométries* : C. R. Acad. Sci. Paris Sér. I Math. **311** (1990) 761–764.

[An] J.W.Anderson. *Hyperbolic geometry* : Springer Undergraduate Mathematics Series (1999).

[Ar] M.A.Armstrong, *Basic topology* : McGraw Hill (1979).

[Bea] A.F.Beardon, *The geometry of discrete groups* : Graduate Texts in Mathematics No. 91, Springer (1983).

[Bes1] M.Bestvina, *Degenerations of the hyperbolic space* : Duke Math. J. **56** (1988) 143–161.

[Bes2] M.Bestvina, **R**-*trees in topology, geometry and group theory* : in "Handbook of geometric topology" (eds. R.J.Daverman, R.B.Sher) North-Hollan (2002) 55-91.

[Bes3] M.Bestvina, *Questions in geometric group theory* : Problem list, Salt Lake City (2000), updated 2004.

[BesF] M.Bestvina, M.Feighn, *Stable actions of groups on real trees* : Invent. Math. **121** (1995) 287–321.

[BesM] M.Bestvina, G.Mess, *The boundary of negatively curved groups* : J. Amer. Math. Soc. **4** (1991) 469–481.

[Bow1] B.H.Bowditch, *Notes on Gromov's hyperbolicity criterion for path-metric spaces* : in "Group theory from a geometrical viewpoint" (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 64–167.

[Bow2] B.H.Bowditch, *A short proof that a subquadratic isoperimetric inequality implies a linear one* : Michigan Math. J. **42** (1995) 103-107.

[Bow3] B.H.Bowditch, *Cut points and canonical splittings of hyperbolic groups* : Acta Math. **180** (1998) 145-186.

[Bow4] B.H.Bowditch, *A topological characterisation of hyperbolic groups* : J. Amer. Math. Soc. **11** (1998) 643–667.

[Bow5] B.H.Bowditch. *Planar groups and the Seifert conjecture* : J. reine angew. Math. **576** (2004) 11-62.

[Bow6] B.H.Bowditch, *Intersection numbers and the hyperbolicity of the curve complex* : to appear in J. reine angew. Math.

[Bow7] B.H.Bowditch, *Relatively hyperbolic groups* : preprint, Southampton (1999).

[Bow8] B.H.Bowditch, *Hyperbolic 3-manifolds and the geometry of the curve complex* : preprint, Southampton (2004).

[BowM] B.H.Bowditch, G.Mess, *A 4-dimensional Kleinian group* : Trans. Amer. Math. Soc. **344** (1994) 391–405.

[BraB] N.Brady, M.R.Bridson, *There is only one gap in the isoperimetric spectrum* : Geom. Funct. Anal. **10** (2000) 1053–1070.

[Bri] M.R.Bridson, *Fractional isoperimetric inequalities and subgroup distorsion* : J. Amer. Math. Soc. **12** (1999) 1103–1118.

[BriH] M.R.Bridson, A.Haefliger, *Metric spaces of non-positive curvature* : Grundlehren der Math. Wiss. No. 319, Springer (1999).

[BroCM] J.F.Brock, R.D.Canary, Y.N.Minsky, *The classification of Kleinian surface groups II: The ending lamination conjecture* : preprint (2004).

[BuM] R.G.Burns, O.Macedonska, *Balanced presentations of the trivial group* : Bull. London Math. Soc. **25** (1993) 513–526.

[Can] J.W.Cannon, *The combinatorial structure of cocompact discrete hyperbolic groups* : Geom. Dedicata **16** (1984) 123–148.

[CanFKP] J.W.Cannon, W.J.Floyd, R.Kenyon, W.R.Parry, *Hyperbolic geometry* : in "Flavors of geometry", (ed. S.Levi), Math. Sci. Res. Inst. Publ. No. 31, Cambridge Univ. Press (1997) 59–115.

[CasJ] A.Casson, D.Jungreis, *Convergence groups and Seifert fibered 3-manifolds* : Invent. Math. **118** (1994) 441–456.

[Cha] C.Champetier, *Propriétés statistiques des groupes de présentation finie* : Adv. Math. **116** (1995) 197–262.

[Chi] I.Chiswell, *Introduction to Λ-trees* : World Scientific (2001).

[ConM] A.Connes, H.Moscovici, *Cyclic cohomology, the Novikov conjecture and hyperbolic groups* : Topology **29** (1990) 345–388.

[CooDP] M.Coornaert, T.Delzant, A.Papadopoulos, *Les groupes hyperboliques de Gromov* : Lecture Notes in Mathematics No. 1441, Springer Verlag (1990).

[CouGHN] D.Coulson, O.A.Goodman, C.D.Hodgson, W.Neumann, *Computing arithmetic invariants of 3-manifolds* : Experimental Math. **9** (2000) 127–152.

[Da] F.Dahmani, *: Combination of convergence groups* : Geom. Topol. **7** (2003) 933-963.

[Dr] C.Druţu *Quasi-isometry invariants and asymptotic cones* : Internat. J. Algebra Comput. **12** (2002) 99-135.

[Du] M.J.Dunwoody, *The accessibility of finitely presented groups* : Invent. Math. **81** (1985) 449–457.

[DuSa] M.J.Dunwoody, M.E.Sageev, *JSJ-splittings for finitely presented groups over slender subgroups* : Invent. Math. **135** (1999) 35–44.

[DyP] A.Dyubina, I.Polterovich, *Explicit constructions of univerisal R-trees and asymptotic geometry of hyoerbolic spaces* : Bull. London. Math. Soc. **33** (2001) 727–734.

[ECHLPT] D.B.A.Epstein, J.W.Cannon, D.F.Holt, S.V.F.Levy, M.S.Paterson, W.P.Thurston, *Word processing in groups* : Jones and Bartlett (1992).

[Fa] B.Farb, *Relatively hyperbolic groups* : Geom. Funct. Anal. **8** (1998) 810-840.

[FuP] K.Fujiwara, P.Papasoglu, *JSJ decompositions of finitely presented groups and complexes of groups* : to appear in Geom. Funct. Anal.

[Gaba] D.Gabai, *Convergence groups are fuchsian groups* : Ann. of Math. **136** (1992) 447–510.

[GaboLP] D.Gaboriau, G.Levitt, F.Paulin, *Pseudogroups of isometries of* $\mathbf{R}$ *and Rips' theorem on free actions on* $\mathbf{R}$*-trees* : Israel J. Math. **87** (1994) 403–426.

[GeM] F.W.Gehring, G.J.Martin, *Discrete quasiconformal groups I* : Proc. London Math. Soc. **55** (1987) 331–358.

[GhH] E.Ghys, P.de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov* : Progress in Maths. No. 83, Birkhäuser (1990).

[Go] W.M.Goldman, *Complex hyperbolic geometry* : Oxford Mathematical Monographs, Oxford University Press (1999).

[Gre] M.J.Greenberg, *Euclidean and non-euclidean geometries: development and history* : W.H.Freeman and Co. (1973).

[Gro1] M.Gromov, *Groups of polynomial growth and expanding maps* : Inst. Hautes Études Sci. Publ. Math. No. 53 (1981) 53–75.

[Gro2] M.Gromov, *Hyperbolic groups* : in "Essays in Group Theory" (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer-Verlag (1987) 75–263.

[Gro3] M.Gromov, *Asymptotic invariants of infinite groups* : in "Geometric group theory, Vol. 2" (eds. G.A.Niblo, M.A.Roller) London Math. Soc. Lecture Notes Series No. 182, Cambridge Univ. Press (1993) 1–295.

[Gro4] M.Gromov (edited by J.LaFontaine and P.Pansu), *Metric structures for Reimannian and non-Riemannian spaces* : Progress in Mathematics No. 152, Birkhäuser (1998).

[Ham] U.Hamenstädt, *Geometry of the complex of curves and of Teichmüller space* : Preprint, Bonn (2005).

[Harp] P.de la Harpe, *Topics in geometric group theory* : Chicago Lectures in Mathematics, University of Chicago Press (2000).

[Harv] W.J.Harvey, *Boundary structure of the modular group* : in "Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference" (ed. I.Kra, B.Maskit), Ann. of Math. Stud. No. 97, Princeton University Press (1981) 245–251.

[Hi] G.Higman, *Subgroups of finitely presented groups* : Proc. Roy. Soc. **262** (1961) 455-475.

[Ho] H.Hopf, *Enden offner Räume und unendliche diskontinuierliche Gruppen* : Comment. Math. Helv. **16** (1943) 81-100.

[ImT] Y.Imayoshi, M.Taniguchi, *An introduction to Teichmüller spaces* : Springer-Verlag (1992).

[Iv] B.Iversen, *Hyperbolic geometry* : London Math. Soc. Student Texts No. 25, Cambridge Univ. Press (1992).

[JaS] W.H.Jaco, P.B.Shalen, *Seifert fibered spaces in 3-manifolds* : Mem. Amer. Math. Soc. Mem. No. 220 (1979).

[Jo] K.Johannson, *Homotopy equivalences of 3-manifolds with boundary* : Springer Lecture Notes in Mathematics, No. 761, Springer Verlag, Berlin (1979)

[Kap-iB] I.Kapovich, N.Benakli, *Boundaries of hyperbolic groups* : in "Combinatorial and geometric group theory" (eds. S.Cleary, R.Gilman, A.G.Myasnikov, V.Shpilrain), Contemp Math. No. 296, Amer. Math. Soc. (2002) 39–93.

[Kap-mP] M.Kapovich, L.Potyagailo, *On absence of Ahlfors' finiteness theorem for Kleinian groups in dimension 3* : Topology and its Applications **40** (1991) 83–91.

[KasS] G.Kasparov, G.Skandalis, *Groups acting properly on "bolic" spaces and the Novikov conjecture* : Ann. of Math. **158** (2003) 165–206.

[KlL] B.Kleiner, J.Lott, *Notes on Perleman's papers* : Notes, Ann Arbor.

[LS] R.C.Lyndon, P.Schupp, *Combinatorial group theory* : Ergebnisse er Mathematik und ihrer Grenzgebiete, No. 89, Springer (1977).

[MagKS] W.Magnus, A.Karrass, D.Solitar, *Combinatorial group theory: presentations of groups in terms of generators and relations* : Interscience (1966).

[MasM] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103-149.

[Me] G.Mess, *The Seifert conjecture and groups which are coarse quasiisometric to planes* : preprint (1988).

[Mine] I.Mineyev, *Flows and joins of metric spaces* : Geom. Topol. **9** (2005) 403–482.

[MineY] I.Mineyev, G.Yu, *The Baum-Connes conjecture for hyperbolic groups : Invent. Math.* **149** *(2002) 97-122.*

[Mins] Y.N.Minsky, *The classification of Kleinian surface groups I : Models and bounds* : preprint, Stony Brook (2002).

[MonZ] D.Montgomery, L.Zippin, *Topological transformation groups* : Interscience (1955).

[Mor] J.W.Morgan, *Recent progress on the Poincaré conjecture and the classification of 3-manifolds* : Bull. Amer. Math. Soc. **42** (2005) 57–78.

[MorS] J.W.Morgan, P.B.Shalen, *Valuations, trees, and degenerations of hyperbolic structures* : Ann. of Math. **120** (1984) 401–476.

[Mos] G.D.Mostow, *Quasi-conformal mappings in n-space an the rigidity of hyperbolic space forms* : Inst. Hautes Études Sci. Publ. Math. No. (1968) 53–104.

[O] A.Yu.Ol'shanskii, *Hyperbolicity of groups with subquadratic isoperimetric inequality* : International J. Alg. and Comp. **1** (1991) 281–289.

[Pap] P.Papasoglu, *On the sub-quadratic isoperimetric inequality* : in Proceedings of the Ohio-State program in geometric group theory (ed. R. Charney and M. Davis), de Gruyter (1995) 149–157.

[Pau] F.Paulin, *Actions de groupes sur les arbres* : Séminaire Bourbaki, Exposé 808, Astérisque No. 241, Société Mathématique de France (1997) 97–138.

[Pe1] G.Perelman, *The entropy formula for Ricci flow and its geometric applications* : preprint, Saint Petersburg (2003).

[Pe2] G.Perelman, *Ricci flow with surgery on 3-manifolds* : preprint, Saint Petersburg (2003).

[Re] A.W.Reid, *A note on trace fields of Kleinian groups* : Bull. London Math. Soc. **22** (1990) 349-352.

[RiS] E.Rips, Z.Sela, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition* : Ann. of Math. **146** (1997) 53–109.

[SaBR] M.V.Sapir, J.-C.Birget, E.Rips, *Isomperimetric and isodiametric functions of groups* : Ann. of Math. **156** (2002) 345–466.

[ScW] P.Scott, T.Wall, *Topological methods in group theory* : in "Homological group theory" (ed. C.T.C.Wall) London Math. Soc. Lecture Note Ser. No. 36, Cambridge Univ. Press (1979) 137–203.

[Se1] Z.Sela, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups II* : Geom. Funct. Anal. **7** (1997) 561–593.

[Se2] Z.Sela, *Diophantine geometry over groups I, Makanin-Razborov diagrams* : Publ. Math. Inst. Hautes Études Sci. No. 90 (2001) 31–105.

[Sha] Y.Shalom, *Harmonic analysis, cohomology, and the large-scale geometry of amenable groups* : Acta Math. **192** (2004) 119–185.

[Sho] H.Short et al., *Notes on word hyperbolic groups* : in "Group theory from a geometrical viewpoint" (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 3–63.

[Sp] E.H.Spanier, *Algebraic topology* : McGraw-Hill (1966).

[St] J.R.Stallings, *Group theory and three-dimensional manifolds* : Yale Math. Monographs No. 4, Yale University Press, New Haven (1971).

[Sz] A.Szczepański, *Relatively hyperbolic groups* : Michigan Math. J. **45** (1998) 611–618.

[Th] W.P.Thurston, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry* : Bull. Amer. Math. Soc. **9** (1982) 357–381.

[Tu1] P.Tukia, *Homeomorphic conjugates of fuchsian groups* : J. reine angew. Math. **391** (1988) 1–54.

[Tu2] P.Tukia, *Convergence groups and Gromov's metric hyperbolic spaces* : New Zealand J. Math. **23** (1994) 157–187.

[V] N.T.Varopoulos, *Théorie du potentiel sur des groupes et des variétés* : C. R. Acad. Sc. Paris, Sér. I **102** (1986) 203–205.

[W] J.R.Weeks, *: The shape of space. How to visualize surfaces and three-dimensional manifolds* : Monographs and textbooks in pure and applied mathematics, No. 96, Marcel Dekker (1985).

[Y] A.Yaman, *A topological characterisation of relatively hyperbolic groups* : J. reine. angew. Math. **566** (2004) 41–89.