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 to 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 $x \in M$ 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 $t, u \in I$.

(Sometimes, we may talk 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

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

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 length(γ) is equal to $d(\gamma(a), \gamma(b))$ 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 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 in Q 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 do 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 $\{g \in \Gamma \mid d(x, gx) \leq r\}$ 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 $K \subseteq X$.

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. In defining a quasi-isometry, we are trying to capture the "large scale" geometry of our spaces. It cannot therefore 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 (in other words, 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 do 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 of a 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.

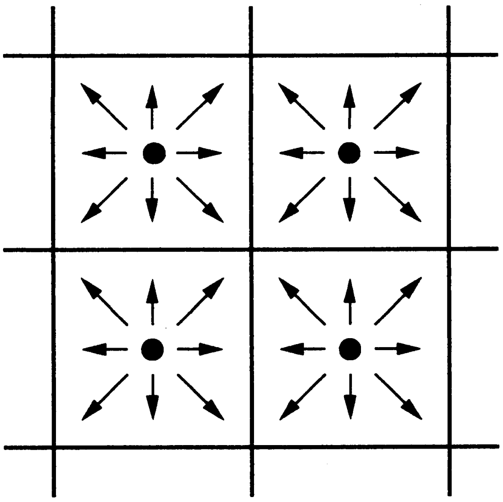

Figure 3a.

(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 are highlighted in bold.) The quotient 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.

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

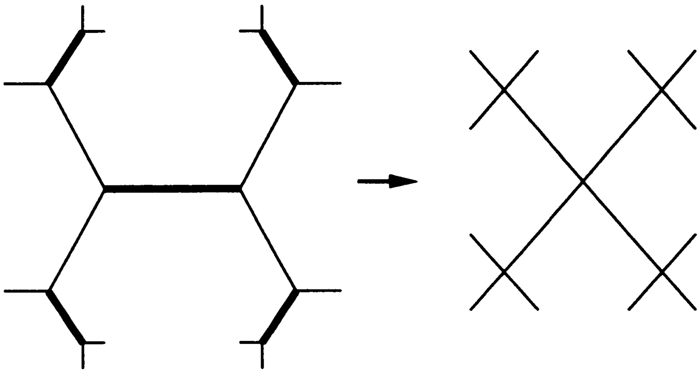

Figure 3b.

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 than $\phi(0)$, 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, $p \le p - q$ 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 $||x_i - x_{i+1}|| \leq \pi$. 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 $m \leq n$. 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 $k \geq 0$ such that for all $x \in \mathbf{R}^n$, $\operatorname{diam}(f^{-1}(x)) \leq k$. Then f is surjective.

The idea of the proof is 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 $d'(\phi(x), \phi(y)) \leq r d(x, y)$ 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. $\diamondsuit$

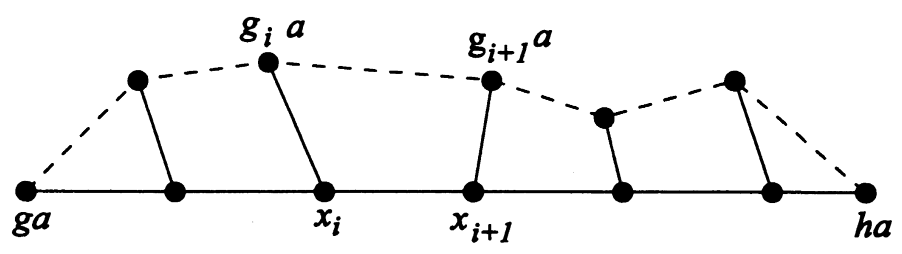

Figure 3c.

Now let $A = \{g \in \Gamma \setminus \{1\} \mid d(a, ga) \leq k\}$. Thus A is finite and symmetric, and $g, h \in \Delta$ are adjacent 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 $\Gamma \approx \Gamma' \Rightarrow \Gamma \sim \Gamma'$.* $\diamond$

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 $V(\Delta) \equiv \Gamma$. 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 have 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 $\psi([0, \infty))$ is a bounded distance from $[\psi(0), \infty)$ (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 $1, h, k$ 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, h^{-1}, k, k^{-1}, hk^{-1} or kh^{-1}. $\diamondsuit$

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 4.3).

(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 5.5.

Some negative results.

There are 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 be 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.