

Part IV — Topics in Geometric Group Theory

Theorems with proof

Based on lectures by H. Wilton

Notes taken by Dexter Chua

Michaelmas 2017

These notes are not endorsed by the lecturers, and I have modified them (often significantly) after lectures. They are nowhere near accurate representations of what was actually lectured, and in particular, all errors are almost surely mine.

The subject of geometric group theory is founded on the observation that the algebraic and algorithmic properties of a discrete group are closely related to the geometric features of the spaces on which the group acts. This graduate course will provide an introduction to the basic ideas of the subject.

Suppose Γ is a discrete group of isometries of a metric space X . We focus on the theorems we can prove about Γ by imposing geometric conditions on X . These conditions are motivated by curvature conditions in differential geometry, but apply to general metric spaces and are much easier to state. First we study the case when X is *Gromov-hyperbolic*, which corresponds to negative curvature. Then we study the case when X is *CAT(0)*, which corresponds to non-positive curvature. In order for this theory to be useful, we need a rich supply of negatively and non-positively curved spaces. We develop the theory of *non-positively curved cube complexes*, which provide many examples of CAT(0) spaces and have been the source of some dramatic developments in low-dimensional topology over the last twenty years.

- Part 1. We will introduce the basic notions of geometric group theory: Cayley graphs, quasiisometries, the Schwarz–Milnor Lemma, and the connection with algebraic topology via presentation complexes. We will discuss the word problem, which is quantified using the Dehn functions of a group.
- Part 2. We will cover the basic theory of word-hyperbolic groups, including the Morse lemma, local characterization of quasigeodesics, linear isoperimetric inequality, finitely presentedness, quasiconvex subgroups etc.
- Part 3. We will cover the basic theory of CAT(0) spaces, working up to the Cartan–Hadamard theorem and Gromov’s Link Condition. These two results together enable us to check whether the universal cover of a complex admits a CAT(1) metric.
- Part 4. We will introduce cube complexes, in which Gromov’s link condition becomes purely combinatorial. If there is time, we will discuss Haglund–Wise’s *special* cube complexes, which combine the good geometric properties of CAT(0) spaces with some strong algebraic and topological properties.

Pre-requisites

Part IB Geometry and Part II Algebraic topology are required.

Contents

| | | |
|----------|---|-----------|
| 1 | Cayley graphs and the word metric | 4 |
| 1.1 | The word metric | 4 |
| 1.2 | Free groups | 5 |
| 1.3 | Finitely-presented groups | 5 |
| 1.4 | The word problem | 5 |
| 2 | Van Kampen diagrams | 8 |
| 3 | Bass–Serre theory | 11 |
| 3.1 | Graphs of spaces | 11 |
| 3.2 | The Bass–Serre tree | 11 |
| 4 | Hyperbolic groups | 13 |
| 4.1 | Definitions and examples | 13 |
| 4.2 | Quasi-geodesics and hyperbolicity | 13 |
| 4.3 | Dehn functions of hyperbolic groups | 17 |
| 5 | CAT(0) spaces and groups | 19 |
| 5.1 | Some basic motivations | 19 |
| 5.2 | CAT(κ) spaces | 19 |
| 5.3 | Length metrics | 20 |
| 5.4 | Alexandrov’s lemma | 20 |
| 5.5 | Cartan–Hadamard theorem | 21 |
| 5.6 | Gromov’s link condition | 22 |
| 5.7 | Cube complexes | 22 |
| 5.8 | Special cube complexes | 22 |

1 Cayley graphs and the word metric

1.1 The word metric

Theorem. For any two finite generating sets S, S' of a group Γ , the identity map $(\Gamma, d_S) \rightarrow (\Gamma, d_{S'})$ is a quasi-isometry.

Proof. Pick

$$\lambda = \max_{s \in S} \ell_{S'}(s), \quad \lambda' = \max_{s \in S'} \ell_S(s),$$

We then see

$$\ell_S(\gamma) \leq \lambda' \ell_{S'}(\gamma), \quad \ell_{S'}(\gamma) \leq \lambda \ell_S(\gamma).$$

for all $\gamma \in \Gamma$. Then the claim follows. \square

Lemma (Schwarz–Milnor lemma). Let X be a proper geodesic metric space, and let Γ act properly discontinuously, cocompactly on X by isometries. Then

- (i) Γ is finitely-generated.
- (ii) For any $x_0 \in X$, the orbit map

$$\begin{aligned} \Gamma &\rightarrow X \\ \gamma &\mapsto \gamma x_0 \end{aligned}$$

is a quasi-isometry $(\Gamma, d_s) \underset{qi}{\simeq} (X, d)$.

Proof of Schwarz–Milnor lemma. Let $\bar{B} = \bar{B}(x, R)$ be such that $\Gamma \bar{B} = X$. This is possible since the quotient is compact.

Let $S = \{\gamma \in \Gamma : \gamma \bar{B} \cap \bar{B} \neq \emptyset\}$. By proper discontinuity, this set is finite.

We let

$$r = \inf_{\gamma' \notin S} d(\bar{B}, \gamma' \bar{B}).$$

If we think about it, we see that in fact r is the *minimum* of this set, and in particular $r > 0$.

Finally, let

$$\lambda = \max_{s \in S} d(x_0, s x_0).$$

We will show that S generates Γ , and use the word metric given by S to show that Γ is quasi-isometric to X .

We first show that $\Gamma = \langle S \rangle$. We let $\gamma \in \Gamma$ be arbitrary. Let $[x_0, \gamma x_0]$ be a geodesic from x_0 to γx_0 . Let ℓ be such that

$$(\ell - 1)r \leq d(x_0, \gamma x_0) < \ell r.$$

Then we can divide the geodesic into ℓ pieces of length about r . We can choose $x_1, \dots, x_{\ell-1}, x_\ell = \gamma x_0$ such that $d(x_{i-1}, x_i) < r$.

By assumption, we can pick $\gamma_i \in \Gamma$ such that $x_i \in \gamma_i \bar{B}$, and further we pick $\gamma_\ell = \gamma, \gamma_0 = e$. Now for each i , we know

$$d(\bar{B}, \gamma_{i-1}^{-1} \gamma_i \bar{B}) = d(\gamma_{i-1} \bar{B}, \gamma_i \bar{B}) \leq d(x_{i-1}, x_i) < r.$$

So it follows that $\gamma_{i-1}^{-1} \gamma_i \in S$. So we have

$$\gamma = \gamma_\ell = (\gamma_0^{-1} \gamma_1)(\gamma_1^{-1} \gamma_2) \cdots (\gamma_{\ell-1}^{-1} \gamma_\ell) \in \langle S \rangle.$$

This proves $\Gamma = \langle S \rangle$.

To prove the second part, we simply note that

$$r\ell - r \leq d(x_0\gamma x_0).$$

We also saw that ℓ is an upper bound for the word length of γ under S . So we have

$$r\ell_s(\gamma) - r \leq d(x_0, \gamma x_0).$$

On the other hand, by definition of λ , we have

$$d(x_0, \gamma x_0) \leq \lambda\ell_s(\gamma).$$

So the orbit map is an orbit-embedding, and quasi-surjectivity follows from cocompactness. \square

1.2 Free groups

Corollary. $\pi_1(X_r)$ has the universal property of $F(S)$, with $S = \{a_1, \dots, a_r\}$. So $\pi_1(X_r) \cong F_r$.

Proof. Given any map $f : S \rightarrow G$, define $\tilde{f} : F(S) \rightarrow G$ by

$$\tilde{f}(a_{i_1}^{\pm 1} \dots a_{i_n}^{\pm 1}) = \tilde{f}(a_{i_1})^{\pm 1} \dots \tilde{f}(a_{i_n})^{\pm 1}$$

for any reduced word $a_{i_1}^{\pm 1} \dots a_{i_n}^{\pm 1}$. This is easily seen to be well-defined and is the unique map making the diagram commute. \square

1.3 Finitely-presented groups

1.4 The word problem

Lemma. Let $\Gamma = \langle S \mid R \rangle$. Then the elements of $\ker(F(S) \rightarrow \Gamma)$ are precisely those of the form

$$\prod_{i=1}^d g_i r_i^{\pm 1} g_i^{-1},$$

where $g_i \in F(S)$ and $r_i \in R$.

Proof. We know that $\ker(F(S) \rightarrow \Gamma) = \langle\langle R \rangle\rangle$, and the set described above is exactly $\langle\langle R \rangle\rangle$, noting that $gxyg^{-1} = (gxg^{-1})(yyg^{-1})$. \square

Proposition. The word problem for \mathcal{P} is solvable iff $\delta_{\mathcal{P}}$ is a computable function.

Proposition. If P and Q are two finite presentations for Γ , then $\delta_{\mathcal{P}} \approx \delta_{\mathcal{Q}}$.

Lemma. If $R' \subseteq \langle\langle R \rangle\rangle$ is a finite set, and

$$\mathcal{P} = \langle S \mid R \rangle, \quad \mathcal{Q} = \langle S \mid R \cup R' \rangle,$$

then $\delta_{\mathcal{P}} \simeq \delta_{\mathcal{Q}}$.

Proof. Clearly, $\delta_{\mathcal{Q}} \leq \delta_{\mathcal{P}}$. Let

$$m = \max_{r' \in R'} \text{Area}_{\mathcal{P}}(r').$$

It is then easy to see that

$$\text{Area}_{\mathcal{P}}(w) \leq m \text{Area}_{\mathcal{Q}}(w). \quad \square$$

Lemma. Let $\mathcal{P} = \langle S \mid R \rangle$, and let

$$\mathcal{Q} = \langle S \amalg T \mid R \amalg R' \rangle,$$

where

$$R' = \{tw_t^{-1} : t \in T, w_t \in F(S)\}.$$

Then $\delta_{\mathcal{P}} \approx \delta_{\mathcal{Q}}$.

Proof. We first show $\delta_{\mathcal{P}} \preceq \delta_{\mathcal{Q}}$. Define

$$\begin{aligned} \rho : F(S \amalg T) &\rightarrow F(S) \\ s &\mapsto s \\ t &\mapsto w_t. \end{aligned}$$

In particular, $\rho(r) = r$ for all $r \in R$ and $\rho(r') = 1$ for all $r' \in R'$.

Given $w \in F(S)$, we need to show that

$$\text{Area}_{\mathcal{P}}(w) \leq \text{Area}_{\mathcal{Q}}(w).$$

Let $d = \text{Area}_{\mathcal{Q}}(w)$. Then

$$w = \prod_{i=1}^d g_i r_i^{\pm 1} g_i^{-1},$$

where $g_i \in F(S \amalg T)$, $r_i \in R \cup R'$. We now apply ρ . Since ρ is a retraction, $\rho(w) = w$. Thus,

$$w = \prod_{i=1}^d \rho(g_i) \rho(r_i)^{\pm 1} \rho(g_i)^{-1}.$$

Now $\rho(r_i)$ is either r_i or 1. In the first case, nothing happens, and in the second case, we can just forget the i th term. So we get

$$w = \prod_{i=1, r_i \in R}^d \rho(g_i) r_i^{\pm 1} \rho(g_i)^{-1}.$$

Since this is a valid proof in \mathcal{P} that $w = 1$, we know that

$$\text{Area}_{\mathcal{P}}(w) \leq d = \text{Area}_{\mathcal{Q}}(w).$$

We next prove that $\delta_{\mathcal{Q}} \preceq \delta_{\mathcal{P}}$. It is unsurprising that some constants will appear this time, instead of just inequality on the nose. We let

$$C = \max_{t \in T} |w_t|_S.$$

Consider a null-homotopic word $w \in F(S \amalg T)$. This word looks like

$$w = s_{i_1}^{\pm} s_{i_2}^{\pm} \cdots t_{j_1} \cdots s \cdots t \cdots \in F(S \amalg T).$$

We want to turn these t 's into s 's, and we need to use the relators to do so.

We apply relations from R' to write this as

$$w' = s_{i_1}^{\pm} s_{i_2}^{\pm} \cdots w_{t_{j_1}} \cdots s \cdots w_t \cdots \in F(S).$$

We certainly have $|w'|_S \leq C|w|_{S \amalg T}$. With a bit of thought, we see that

$$\begin{aligned} \text{Area}_{\mathcal{Q}}(w) &\leq \text{Area}_{\mathcal{P}}(w') + |w|_{S \amalg T} \\ &\leq \delta_{\mathcal{P}}(C|w|_{S \amalg T}) + |w|_{S \amalg T}. \end{aligned}$$

So it follows that

$$\delta_{\mathcal{Q}}(n) \leq \delta_{\mathcal{P}}(Cn) + n. \quad \square$$

Proof of proposition. Let $\mathcal{P} = \langle S \mid R \rangle$ and $\mathcal{Q} = \langle S' \mid R' \rangle$. If \mathcal{P}, \mathcal{Q} present isomorphic groups, then we can write

$$s = u_s \in F(S') \text{ for all } s \in S$$

Similarly, we have

$$s' = v_{s'} \in F(S) \text{ for all } s' \in S'$$

We let

$$\begin{aligned} T &= \{su_s^{-1} \mid s \in S\} \\ T' &= \{s'v_{s'}^{-1} \mid s' \in S'\} \end{aligned}$$

We then let

$$\mathcal{M} = \langle S \amalg S' \mid R \cup R' \cup T \cup T' \rangle.$$

We can then apply our lemmas several times to show that $\delta_{\mathcal{P}} \approx \delta_{\mathcal{M}} \approx \delta_{\mathcal{Q}}$. \square

2 Van Kampen diagrams

Lemma (van Kampen’s lemma). Let $\mathcal{P} = \langle S \mid R \rangle$ be a presentation and $w \in S^*$. Then the following are equivalent:

- (i) $w = 1$ in Γ presented by \mathcal{P} (i.e. w is null-homotopic)
- (ii) There is a van Kampen diagram for w given \mathcal{P} .


If so, then

$$\text{Area}_a(w) = \min\{\text{Area}_g(D) : D \text{ is a van Kampen diagram for } w \text{ over } \mathcal{P}\}.$$


Proof. In one direction, given

$$w = \prod_{i=1}^n g_i r_i^\pm g_i^{-1} \in F(S)$$

such that $w = 1 \in \Gamma$, we start by writing down a “lollipop diagram”



This defines a diagram for the word $\prod_{i=1}^n g_i r_i^\pm g_i^{-1}$ which is equal in $F(S)$ to w , but is not exactly w . However, note that performing elementary reductions (or their inverses) correspond to operations on the van Kampen diagram that does not increase the area. We will be careful and not say that the area doesn’t change, since we may have to collapse paths that look like



In the other direction, given a diagram D for w , we need to produce an expression

$$w = \prod_{i=1}^d g_i r_i^\pm g_i^{-1}$$

such that $d \leq \text{Area}(D)$.

We let e be the first 2-cell that the boundary curve arrives at, reading anticlockwise around ∂D . Let g be the path from p to e , and let r be the relator read anti-clockwise around e . Let

$$D' = D - e,$$

and let $w' = \partial D'$. Note that

$$w = g r^{\pm 1} g^{-1} w'.$$

Since $\text{Area}(D') = \text{Area}(D) - 1$, we see by induction that

$$w = \prod_{i=1}^d g_i r_i^{\pm 1} g_i^{-1}$$

with $d = \text{Area}(D)$. □

Corollary. If w is null-homotopic, then we can write w in the form

$$w = \prod_{i=1}^d g_i r_i^{\pm 1} g_i^{-1},$$

where

$$|g_i|_S \leq \text{Diam } D$$

with D a minimal van Kampen diagram for w . We can further bound this by

$$(\max |r_i|_S) \text{Area}(D) + |w|_S \leq \text{constant} \cdot \delta_{\mathcal{P}}(|w|_S) + |w|_S.$$

Proposition. The word problem for a presentation \mathcal{P} is solvable iff $\delta_{\mathcal{P}}$ is computable.

Proof.

- (\Leftarrow) By the corollary, the maximum length of a conjugator g_i that we need to consider is computable. Therefore we know how long the partial algorithm needs to run for.
- (\Rightarrow) To compute $\delta_{\mathcal{P}}(n)$, we use the word problem solving ability to find all null-homotopic words in $F(S)$ of length $\leq n$. Then for each d , go out and look for expressions

$$w = \prod_{i=1}^d g_i r_i^{\pm 1} g_i^{-1}.$$

A naive search would find the smallest area expression, and this gives us the Dehn function. □

Theorem (Novikov–Boone theorem). There exists a finitely-presented group with an unsolvable word problem.

Corollary. $\delta_{\mathcal{P}}$ is sometimes non-computable.

Theorem. Let $n \geq 4$ and $\Gamma = \langle S \mid R \rangle$ be a finitely-presented group. Then we can construct a closed, smooth, orientable manifold M^n such that $\pi_1 M \cong \Gamma$.

Proof. Let $S = \{a_1, \dots, a_m\}$ and $R = \{r_1, \dots, r_n\}$. We start with

$$M_0 = \#_{i=0}^m (S^1 \times S^{n-1}).$$

Note that when we perform this construction, as $n \geq 3$, we have

$$\pi_1 M_0 \cong F_m$$

by Seifert–van Kampen theorem. We now construct M_k from M_{k-1} such that

$$\pi_1 M_k \cong \langle a_1, \dots, a_m \mid r_1, \dots, r_k \rangle.$$

We realize r_k as a loop in M_{k-1} . Because $n \geq 3$, we may assume (after a small homotopy) that this is represented by a smooth embedded map $r_k : S^1 \rightarrow M_{k-1}$.

We take N_k to be a smooth tubular neighbourhood of r_k . Then $N_k \cong S^1 \times D^{n-1} \subseteq M_{k-1}$. Note that $\partial N_k \cong S^1 \times S^{n-2}$.

Let $U_k = D^2 \times S^{n-2}$. Notice that $\partial U_k \cong \partial N_k$. Since $n \geq 4$, we know U_k is simply connected. So we let

$$M'_{k-1} = M_k \setminus \overset{\circ}{N}_k,$$

a manifold with boundary $S^1 \times S^{n-2}$. Choose an orientation-reversing diffeomorphism $\varphi_k : \partial U_k \rightarrow \partial M'_{k-1}$. Let

$$M_k = M'_{k-1} \cup_{\varphi_k} U_k.$$

Then by applying Seifert van Kampen repeatedly, we see that

$$\pi_1 M_k = \pi_1 M_{k-1} / \langle\langle r_k \rangle\rangle,$$

as desired. □

Theorem (Douglas, Radú, Murray). If γ is embedded, then there is a least-area filling disc.

Theorem (Filling theorem). let M be a closed Riemannian manifold. Then $FiU^M \simeq \delta_{\pi_1 M}$.

3 Bass–Serre theory

3.1 Graphs of spaces

Proposition. For all groups G there exists an aspherical space $BG = K(G, 1)$ such that $\pi_1(K(G, 1)) \cong G$. Moreover, for any two choices of $K(G, 1)$ and $K(H, 1)$, and for every homomorphism $f : G \rightarrow H$, there is a unique map (up to homotopy) $\bar{f} : K(G, 1) \rightarrow K(H, 1)$ that induces this homomorphism on π_1 . In particular, $K(G, 1)$ is well-defined up to homotopy equivalence.

Moreover, we can choose $K(G, 1)$ functorially, namely there are choices of $K(G, 1)$ for each G and choices of \bar{f} such that $\overline{f_1 \circ f_2} = \bar{f}_1 \circ \bar{f}_2$ and $\overline{\text{id}_G} = \text{id}_{K(G, 1)}$ for all f, G, H .

3.2 The Bass–Serre tree

Lemma. If \mathcal{X} is a graph of spaces and $\hat{X} \rightarrow X$ is a covering map, then \hat{X} naturally has the structure of a graph of spaces $\hat{\mathcal{X}}$, and p respects that structure.

Proof sketch. Consider

$$\bigcup_{v \in V(\Xi)} X_v \subseteq X.$$

Let

$$p^{-1} \left(\bigcup_{v \in V(\Xi)} X_v \right) = \coprod_{\hat{v} \in V(\hat{\Xi})} \hat{W}_{\hat{v}}.$$

This defines the vertices of $\hat{\Xi}$, the underlying graph of $\hat{\mathcal{X}}$. The path components $\hat{X}_{\hat{v}}$ are going to be the vertex spaces of $\hat{\mathcal{X}}$. Note that for each \hat{v} , there exists a unique $v \in V(\Xi)$ such that $p : \hat{X}_{\hat{v}} \rightarrow X_v$ is a covering map.

Likewise, the path components of

$$p^{-1} \left(\bigcup_{e \in E(\Xi)} X_e \times \{0\} \right)$$

form the edge spaces $\coprod_{e \in E(\Xi)} \hat{X}_{\hat{e}}$ of $\hat{\mathcal{X}}$, which again are covering spaces of the edge space of X .

Now let's define the edge maps $\partial_{\hat{e}}^{\pm}$ for each $\hat{e} \in E(\hat{\Xi}) \mapsto e \in E(\Xi)$. To do so, we consider the diagram

$$\begin{array}{ccccc} \hat{X}_{\hat{e}} & \xrightarrow{\sim} & \hat{X}_{\hat{e}} \times [-1, 1] & \dashrightarrow & \hat{X} \\ \downarrow & & \downarrow & & \downarrow p \\ X_e & \xrightarrow{\sim} & X_e \times [-1, 1] & \longrightarrow & X \end{array}$$

By the lifting criterion, for the dashed map to exist, there is a necessary and sufficient condition on $(\hat{X}_{\hat{e}} \times [-1, 1] \rightarrow X_e \times [-1, 1] \rightarrow X)_*$. But since this condition is homotopy invariant, we can check it on the composition $(\hat{X}_{\hat{e}} \rightarrow X_e \rightarrow X)_*$ instead, and we know it must be satisfied because a lift exists in this case.

The attaching maps $\partial_{\hat{e}}^{\pm} : \hat{X}_e \rightarrow \hat{X}$ are precisely the restriction to $\hat{X}_e \times \{\pm 1\} \rightarrow \hat{X}$.

Finally, check using covering space theory that the maps $\hat{X}_{\hat{e}} \times [-1, 1] \rightarrow \hat{X}$ can be injective on the interior of the cylinder, and verify that the appropriate maps are π_1 -injective. \square

Lemma (Britton’s lemma). For any vertex Ξ , the natural map $G_v \rightarrow G$ is injective.

Proof sketch. Observe that the universal cover \tilde{X} can be produced by first building universal covers of the vertex space, which are then glued together in a way that doesn’t kill the fundamental groups. \square

Theorem (Normal form theorem). Every element can be represented by a reduced loop, and the only reduced loop representing the identity is the trivial loop.

Proof idea. It all boils down to the fact that the Bass–Serre tree is a tree. Connectedness gives the existence, and the simply-connectedness gives the “uniqueness”. \square

4 Hyperbolic groups

4.1 Definitions and examples

4.2 Quasi-geodesics and hyperbolicity

Theorem (Morse lemma). For all $\delta \geq 0, \lambda \geq 1$ there is $R(\delta, \lambda, \varepsilon)$ such that the following holds:

If X is a δ -hyperbolic metric space, and $c : [a, b] \rightarrow X$ is a (λ, ε) -quasigeodesic from p to q , and $[p, q]$ is a choice of geodesic from p to q , then

$$d_{\text{Haus}}([p, q], \text{im}(c)) \leq R(\delta, \lambda, \varepsilon),$$

where

$$d_{\text{Haus}}(A, B) = \inf\{\varepsilon > 0 \mid A \subseteq N_\varepsilon(B) \text{ and } B \subseteq N_\varepsilon(A)\}$$

is the *Hausdorff distance*.

Corollary. There is an $M(\delta, \lambda, \varepsilon)$ such that a geodesic metric space X is δ -hyperbolic iff any (λ, ε) -quasigeodesic triangle is M -slim.

Corollary. Suppose X, X' are geodesic metric spaces, and $f : X \rightarrow X'$ is a quasi-isometric embedding. If X' is hyperbolic, then so is X .

In particular, hyperbolicity is a quasi-isometrically invariant property, when restricted to geodesic metric spaces.

Lemma. Let X be a geodesic space. For any (λ, ε) -quasigeodesic $c : [a, b] \rightarrow X$, there exists a continuous, rectifiable (λ, ε') -quasigeodesic $c' : [a, b] \rightarrow X$ with $\varepsilon' = 2(\lambda + \varepsilon)$ such that

(i) $c'(a) = c(a), c'(b) = c(b)$.

(ii) For all $a \leq t < t' \leq b$, we have

$$\ell(c'|_{[t, t']}) \leq k_1 d(c'(t), c'(t')) + k_2$$

where $k_1 = \lambda(\lambda + \varepsilon)$ and $k_2 = (\lambda\varepsilon' + 3)(\lambda + 3)$.

(iii) $d_{\text{Haus}}(\text{im } c, \text{im } c') \leq \lambda + \varepsilon$.

Proof sketch. Let $\Sigma = \{a, b\} \cup ((a, b) \cap \mathbb{Z})$. For $t \in \Sigma$, we let $c'(t) = c(t)$, and define c' to be geodesic between the points of Σ . Then claims (i) and (iii) are clear, and to prove quasigeodesicity and (ii), let $\sigma : [a, b] \rightarrow \Sigma$ be a choice of closest point in Σ , and then estimate $d(c'(t), c'(t'))$ in terms of $d(c(\sigma(t)), c(\sigma(t')))$. \square


Lemma. let X be δ -hyperbolic, and $c : [a, b] \rightarrow X$ a continuous, rectifiable path in X joining p to q . For $[p, q]$ a geodesic, for any $x \in [p, q]$, we have

$$d(x, \text{im } c) \leq \delta |\log_2 \ell(c)| + 1.$$

Proof. We may assume $c : [0, 1] \rightarrow X$ is parametrized proportional to arc length. Suppose

$$\frac{\ell(c)}{2^N} < 1 \leq \frac{\ell(c)}{2^{N-1}}.$$

Let $x_0 = x$. Pick a geodesic triangle between $p, q, c(\frac{1}{2})$. By δ -hyperbolicity, there exists a point x_1 lying on the other two edges such that $d(x_0, x_1) \leq \delta$. We wlog assume $x_1 \in [p, c(\frac{1}{2})]$. We can repeat the argument with $c|_{[0, \frac{1}{2}]}$.



Formally, we proceed by induction on N . If $N = 0$ so that $\ell(c) < 1$, then we are done by taking desired point on $\text{im } c$ to be p (or q). Otherwise, there is some $x_1 \in [p, c(\frac{1}{2})]$ such that $d(x_0, x_1) \leq \delta$. Then

$$\frac{\ell(c|_{[0, \frac{1}{2}]})}{2^{N-1}} < 1 \leq \frac{\ell(c|_{[0, \frac{1}{2}]})}{2^{N-2}}.$$

So by the induction hypothesis,

$$\begin{aligned} d(x_1, \text{im } c) &\leq \delta |\log_2 \ell(c|_{[0, \frac{1}{2}]})| + 1 \\ &= \delta \left| \frac{1}{2} \log_2 \ell(c) \right| + 1 \\ &= \delta (|\log_2 \ell(c)| - 1) + 1. \end{aligned}$$

Note that we used the fact that $\ell(c) > 1$, so that $\log_2 \ell(c) > 0$.

Then we are done since

$$d(x, \text{im } c) \leq d(x, x_1) + d(x_1, \text{im } c). \quad \square$$


Proof of Morse lemma. By the first lemma, we may assume that c is continuous and rectifiable, and satisfies the properties as in the lemma.

Let p, q be the end points of c , and $[p, q]$ a geodesic. First we show that $[p, q]$ is contained in a bounded neighbourhood of $\text{im } c$. Let

$$D = \sup_{x \in [p, q]} d(x, \text{im } c).$$

By compactness of the interval, let $x_0 \in [p, q]$ where the supremum is attained. Then by assumption, $\text{im } c$ lies outside $B(x_0, D)$. Choose $y, z \in [p, q]$ be such that $d(x_0, y) = d(x_0, z) = 2D$ and y, x_0, z appear on the geodesic in this order (take $y = p$, or $z = q$ if that is not possible).

Let $y' = c(s) \in \text{im } c$ be such that $d(y', y) \leq D$, and similarly let $z' = c(t) \in \text{im } c$ be such that $d(z', z) \leq D$.



Let $\gamma = [y, y'] \cdot c|_{[s,t]} \cdot [z', z]$. Then

$$\ell(\gamma) = d(y, y') + d(z, z') + \ell(c|_{[s,t]}) \leq D + D + k_1 d(y', z') + k_2,$$

by assumption. Also, we know that $d(y', z') \leq 6D$. So we have

$$\ell(\gamma) \leq 6k_1 D + 2D + k_2.$$

But we know that

$$d(x_0, \text{im } \gamma) = D.$$

So the second lemma tells us

$$D \leq \delta |\log_2(6k_1 D + 2D + k_2)| + 1.$$

The left hand side is linear in D , and the right hand side is a logarithm in D . So it must be the case that D is bounded. Hence $[p, q] \subseteq N_{D_0}(\text{im } c)$, where D_0 is some constant.

It remains to find a bound M such that $\text{im } c \subseteq N_M([p, q])$. Let $[a', b']$ be a maximal subinterval of $[a, b]$ such that $c[a', b']$ lies entirely outside $\bar{N}_{D_0}([p, q])$. Since $\bar{N}_D(c[a, a'])$ and $\bar{N}_D(c[b', b])$ are both closed, and they collectively cover the connected set $[p, q]$, there exists

$$w \in [p, q] \cap \bar{N}_{D_0}(c[a, a']) \cap \bar{N}_{D_0}(c[b', b]).$$

Therefore there exists $t \in [a, a']$ and $t' \in [b', b]$ such that $d(w, c(t)) \leq D_0$ and $d(w, c(t')) \leq D_0$. In particular, $d(c(t), c(t')) \leq 2D_0$.

By the first lemma, we know

$$\ell(c|_{[t,t']}) \leq 2k_1 D_0 + k_2.$$

So we know that for $s \in [a', b']$, we have

$$\begin{aligned} d(c(s), [p, q]) &\leq d(c(s), w) \\ &\leq d(c(s), c(t)) + d(c(t), w) \\ &\leq \ell(c|_{[t,t']}) + D_0 \\ &\leq D_0 + 2k_1 D_0 + k_2. \end{aligned} \quad \square$$

Theorem (Gromov). A random group is infinite and hyperbolic.

Theorem. Let X be δ -hyperbolic and $c : [a, b] \rightarrow X$ be a k -local geodesic where $k > 8\delta$. Then c is a (λ, ε) -quasigeodesic for some $\lambda = \lambda(\delta, k)$ and $\varepsilon = \varepsilon(\delta, k)$.


Lemma. Let X be δ -hyperbolic and $k > 8\delta$. If $c : [a, b] \rightarrow X$ is a k -local geodesic, then $\text{im } c$ is contained in the 2δ -neighbourhood of $[c(a), c(b)]$.

Proof. Let $x = c(t)$ maximize $d(x, [c(a), c(b)])$. Let

$$y = c\left(t - \frac{k}{2}\right), \quad z = c\left(t + \frac{k}{2}\right).$$

If $t - \frac{k}{2} < a$, we set $y = c(a)$ instead, and similarly for z .

Let $y' \in [c(a), c(b)]$ minimize $d(y, y')$, and likewise let $z' \in [c(a), c(b)]$ minimize $d(z, z')$.



Fix geodesics $[y, y']$ and $[z, z']$. Then we have a geodesic rectangle with vertices y, y', z, z' . By δ -hyperbolicity, there exists w on the rectangle *not* on c , such that $d(x, w) = 2\delta$.

If $w \in [y', z']$, then we win. Otherwise, we may wlog assume $w \in [y, y']$. Note that in the case $y = c(a)$, we must have $y = y'$, and so this would imply $w = y \in [c(a), c(b)]$. So we are only worried about the case $y = c(t - \frac{k}{2})$. So $d(y, x) = \frac{k}{2} > 4\delta$. But then by the triangle inequality, we must have $d(y, w) > 2\delta = d(x, w)$.

However,

$$d(x, y') \leq d(x, w) + d(w, y') < d(y, w) + d(w, y') = d(y, y').$$

So it follows that

$$d(x, [c(a), c(b)]) < d(y, y') = d(y, [c(a), c(b)]).$$

This contradicts our choice of x . □

Proof of theorem. Let $c : [a, b] \rightarrow X$ be a k -local geodesic, and $t \leq t' \in [a, b]$. Choose $t_0 = t < t_1 < \dots < t_n < t'$ such that $t_i = t_{i-1} + k$ for all i and $t' - t_n < k$.

Then by definition, we have

$$d(c(t_{i-1}), c(t_i)) = k, \quad d(c(t_n), c(t')) = |t_n - t'|.$$

for all i . So by the triangle inequality, we have

$$d(c(t), c(t')) \leq \sum_{i=1}^n d(c(t_{i-1}), c(t_i)) + d(t_n, t') = |t - t'|.$$

We now have to establish a coarse *lower* bound on $d(c(t), c(t'))$.

We may wlog assume $t = a$ and $t' = b$. We need to show that

$$d(c(a), c(b)) \geq \frac{1}{\lambda} |b - a| - \varepsilon.$$

We divide c up into regular subintervals $[x_i, x_{i+1}]$, and choose x'_i close to x_i . The goal is then to prove that the x'_i appear in order along $[c(a), c(b)]$.

Let

$$k' = \frac{k}{2} + 2\delta > 6\delta.$$

Let $b - a = Mk' + \eta$ for $0 \leq \eta < k'$ and $M \in \mathbb{N}$. Put $x_i = c(ik')$ for $i = 1, \dots, M$, and let x'_i be a closest point on $[c(a), c(b)]$ to x_i . By the lemma, we know $d(x_i, x'_i) \leq 2\delta$.

Claim. x'_1, \dots, x'_m appear in the correct order along $[c(a), c(b)]$.

Let's finish the proof assuming the claim. If this holds, then note that

$$d(x'_i, x'_{i+1}) \geq k' - 4\delta > 2\delta$$

because we know $d(x_i, x_{i+1}) = 6\delta$, and also $d(x_m, c(b)) \geq \eta - 2\delta$. Therefore

$$d(c(a), c(b)) = \sum_{i=1}^M d(x_i, x_{i-1}) + d(x_m, c(b)) \geq 2\delta M + \eta - 2\delta \geq 2\delta(M - 1).$$

On the other hand, we have

$$M = \frac{|b - a| - \eta}{k'} \geq \frac{|b - a|}{k'} - 1.$$

Thus, we find that

$$d(c(a), c(b)) \geq \frac{2\delta}{k'} |b - a| - 4\delta.$$

To prove the claim, let $x_i = c(t_i)$ for all i . We let

$$\begin{aligned} y &= c(t_{i-1} + 2\delta) \\ z &= c(t_{i+1} - 2\delta). \end{aligned}$$

Define

$$\begin{aligned} \Delta_- &= \Delta(x_{i-1}, x'_{i-1}, y) \\ \Delta_+ &= \Delta(x_{i+1}, x'_{i+1}, z). \end{aligned}$$

Both Δ_- and Δ_+ are disjoint from $B(x_i, 3\delta)$. Indeed, if $w \in \Delta_-$ with $d(x_i, w) \leq 3\delta$, then by δ -slimness of Δ_- , we know $d(w, x_{i-1}) \leq 3\delta$, and so $d(x_i, x_{i-1}) \leq 6\delta$, which is not possible.

Therefore, since the rectangle y, z, x'_{i+1}, x'_{i-1} is 2δ -slim, and x_i is more than 2δ away from the sides yx'_{i-1} and zx'_{i+1} . So there must be some $x''_i \in [x'_{i-1}, x'_{i+1}]$ with $d(x_i, x''_i) \leq 2\delta$.

Now consider $\Delta = \Delta(x_i, x'_i, x''_i)$. We know $x_i x'_i$ and $x_i x''_i$ are both of length $\leq 2\delta$. Note that every point in this triangle is within 3δ of x_i by δ -slimness. So $\Delta \subseteq B(x_i, 3\delta)$, and this implies Δ is disjoint from $B(x_{i-1}, 3\delta)$ and $B(x_{i+1}, 3\delta)$ as before.

But $x'_{i-1} \in B(x_{i-1}, 3\delta)$ and $x'_{i+1} \in B(x_{i+1}, 3\delta)$. Moreover, Δ contains the segment of $[c(a), c(b)]$ joining x'_i and x''_i . Therefore, it must be the case that $x'_i \in [x'_{i-1}, x'_{i+1}]$. \square

4.3 Dehn functions of hyperbolic groups

Corollary. Let X be δ -hyperbolic. Then there exists a constant $C = C(\delta)$ such that any non-constant loop in X is *not* C -locally geodesic.

Proof. Take $k = 8\delta + 1$, and let

$$C = \max\{\lambda\varepsilon, k\}$$

where λ, ε are as in the theorem.

Let $\gamma : [a, b] \rightarrow X$ be a closed loop. If γ were C -locally geodesic, then it would be (λ, ε) -quasigeodesic. So

$$0 = d(\gamma(a), \gamma(b)) \geq \frac{|b - a|}{\lambda} - \varepsilon.$$

So

$$|b - a| \leq \lambda\varepsilon < C.$$

But γ is a C -local geodesic. This implies γ is a constant loop. □

Lemma. If Γ has a Dehn presentation, then δ_Γ is linear.

Proof. Exercise. □

Theorem. Every hyperbolic group Γ is finitely-presented and admits a Dehn presentation.

In particular, the Dehn function is linear, and the word problem is solvable.

Proof. Let S be a finite generating set for Γ , and δ a constant of hyperbolicity for $\text{Cay}_S(\Gamma)$.

Let $C = C(\delta)$ be such that every non-trivial loop is *not* C -locally geodesic.

Take $\{u_i\}$ to be the set of all words in $F(S)$ representing geodesics $[1, u_i]$ in $\text{Cay}_S(\Gamma)$ with $|u_i| < C$. Let $\{v_j\} \subseteq F(S)$ be the set of all *non*-geodesic words of length $\leq C$ in $\text{Cay}_S(\Gamma)$. Let $R = \{u_i^{-1}v_j \in F(S) : u_i =_\Gamma v_j\}$.

We now just observe that this gives the desired Dehn presentation! Indeed, any non-trivial loop must contain one of the v_j 's, since $\text{Cay}_S(\Gamma)$ is not C -locally geodesic, and we can replace it with u_i ! □

Theorem (Gromov, Bowditch, etc). If Γ is a finitely-presented group and $\delta_\Gamma \preceq n^2$, then Γ is hyperbolic.

Theorem. If Γ is finitely-generated, then the following are equivalent:

- (i) Γ is hyperbolic.
- (ii) Γ has a Dehn presentation.
- (iii) Γ satisfies a linear isoperimetric inequality.
- (iv) Γ has a subquadratic isoperimetric inequality.

Lemma (Ping-pong lemma). Let Γ be hyperbolic and torsion-free (for convenience of statement). If $\gamma_1, \gamma_2 \in \Gamma$ do not commute, then for large enough n , the subgroup $\langle \gamma_1^n, \gamma_2^n \rangle \cong F_2$ and is quasi-convex.

Proposition. Let Γ be hyperbolic, and $\gamma \in \Gamma$. Then $C(\gamma)$ is quasiconvex. In particular, it is hyperbolic.

Corollary. Γ does not contain a copy of \mathbb{Z}^2 .

Theorem (Casson–Jungreis, Gabai). If Γ is hyperbolic and $\partial_\infty \Gamma \cong S^1$, then Γ is virtually $\pi_1 \Sigma$ for some closed hyperbolic Σ .

5 CAT(0) spaces and groups

5.1 Some basic motivations

Theorem (Novikov–Boone theorem). There exists a finitely-presented group with an unsolvable word problem.

Theorem (Gordon). There exists a sequence of finitely generated groups Γ_n such that $H_2(\Gamma_n)$ is not computable.

Theorem (Cartan–Hadamard theorem). Let M be a non-positively curved compact manifold. Then \tilde{M} is diffeomorphic to \mathbb{R}^n . In particular, it is contractible. Thus, $M = K(\pi_1 M, 1)$.

Theorem (Poincaré duality). Let M be an orientable compact n -manifold. Then


$$H_k(M; \mathbb{R}) \cong H_{n-k}(M; \mathbb{R}).$$

5.2 CAT(κ) spaces

Lemma (Convexity of the metric). Let X be a CAT(0) space, and $\gamma, \delta : [0, 1] \rightarrow X$ be geodesics (reparameterized). Then for all $t \in [0, 1]$, we have

$$d(\gamma(t), \delta(t)) \leq (1-t)d(\gamma(0), \delta(0)) + td(\gamma(1), \delta(1)).$$

Proof. Consider the rectangle



Let $\alpha : [0, 1] \rightarrow X$ be a geodesic from $\gamma(0)$ to $\delta(1)$. Applying the CAT(0) estimate to $\Delta(\gamma(0), \gamma(1), \delta(1))$, we get

$$d(\gamma(t), \alpha(t)) \leq d(\overline{\gamma(t)}, \overline{\alpha(t)}) = td(\overline{\gamma(1)}, \overline{\alpha(1)}) = td(\gamma(1), \alpha(1)) = td(\gamma(1), \delta(1)),$$

using what we know in plane Euclidean geometry. The same argument shows that

$$d(\delta(t), \alpha(t)) \leq (1-t)d(\delta(0), \gamma(0)).$$

So we know that

$$d(\gamma(t), \delta(t)) \leq d(\gamma(t), \alpha(t)) + d(\alpha(t), \delta(t)) \leq (1-t)d(\gamma(0), \delta(0)) + td(\gamma(1), \delta(1)). \quad \square$$

Lemma. If X is CAT(0), then X is *uniquely geodesic*, i.e. each pair of points is joined by a unique geodesic.

Proof. Suppose $x_0, x_1 \in X$ and $\gamma(0) = \delta(0) = x_0$ and $\gamma(1) = \delta(1) = x_1$. Then by the convexity of the metric, we have $d(\gamma(t), \delta(t)) \leq 0$. So $\gamma(t) = \delta(t)$ for all t . □

Lemma. Let X be a proper, uniquely geodesic metric space. Then geodesics in X vary continuously with their end points in the compact-open topology (which is the same as the uniform convergence topology).

Proof. This is an easy application of the Arzelá–Ascoli theorem. □

Proposition. Any proper $CAT(0)$ space X is contractible.

Proof. Pick a point $x_0 \in X$. Then the map $X \rightarrow \text{Maps}([0, 1], X)$ sending x to the unique geodesic from x_0 to x is continuous. The adjoint map $X \times [0, 1] \rightarrow X$ is then a homotopy from the constant map at x_0 to the identity map. □


Proposition. Any $CAT(0)$ group Γ satisfies a quadratic isoperimetric inequality, that is $\delta_\Gamma \simeq n$ or $\sim n^2$.

5.3 Length metrics

Theorem (Hopf–Rinow theorem). If a length space X is complete and locally compact, then X is proper and geodesic.


5.4 Alexandrov’s lemma

Lemma (Alexandrov’s lemma). Suppose the triangles $\Delta_1 = \Delta(x, y, z_1)$ and $\Delta_2 = \Delta(x, y, z_2)$ in a metric space satisfy the $CAT(0)$ condition, and $y \in [z_1, z_2]$.



Then $\Delta = \Delta(x, z_1, z_2)$ also satisfies the $CAT(0)$ condition.

Proof. Consider $\bar{\Delta}_1$ and $\bar{\Delta}_2$, which together form a Euclidean quadrilateral \bar{Q} with vertices $\bar{x}, \bar{z}_1, \bar{z}_2, \bar{y}$. We claim that then the interior angle at \bar{y} is $\geq 180^\circ$. Suppose not, and it looked like this:




If not, there exists $\bar{p}_i \in [\bar{y}, \bar{z}_i]$ such that $[\bar{p}_1, \bar{p}_2] \cap [\bar{x}, \bar{y}] = \{\bar{q}\}$ and $\bar{q} \neq \bar{y}$. Now

$$\begin{aligned} d(p_1, p_2) &= d(p_2, y) + d(y, p_2) \\ &= d(\bar{p}_1, \bar{y}) + d(\bar{y}, \bar{p}_1) \\ &= d(\bar{p}, \bar{y}) + d(\bar{y}, \bar{p}_2) \\ &> d(\bar{p}_1, \bar{q}) + d(\bar{q}, \bar{p}_2) \\ &\geq d(p_1, q) + d(q, p_2) \\ &\geq d(p_1, p_2), \end{aligned}$$

which is a contradiction.

Thus, we know the right picture looks like this:



To obtain $\bar{\Delta}$, we have to “push” \bar{y} out so that the edge $\bar{z}_1\bar{z}_2$ is straight, while keeping the lengths fixed. There is a natural map $\pi : \bar{\Delta} \rightarrow \bar{Q}$, and the lemma follows by checking that for any $a, b \in \bar{\Delta}$, we have

$$d(\pi(a), \pi(b)) \leq d(a, b).$$

This is an easy case analysis (or is obvious). □

Proposition. If X_1, X_2 are both locally compact, complete $CAT(0)$ spaces and Y is isometric to closed, subspaces of both X_1 and X_2 . Then $X_1 \cup_Y X_2$, equipped with the induced length metric, is $CAT(0)$.

5.5 Cartan–Hadamard theorem

Theorem (Cartan–Hadamard theorem). If X is a complete, connected length space of non-positive curvature, then the universal cover \tilde{X} , equipped with the induced length metric, is $CAT(0)$.

Corollary. A (torsion free) group Γ is $CAT(0)$ iff it is the π_1 of a complete, connected space X of non-positive curvature.

Lemma. If X is proper, non-positively curved and uniquely geodesic, then X is $CAT(0)$.

Proof idea. The idea is that given a triangle, we cut it up into a lot of small triangles, and since X is locally $CAT(0)$, we can use Alexandrov’s lemma to conclude that the large triangle is $CAT(0)$.

Recall that geodesics vary continuously with their endpoints. Consider a triangle $\Delta = \Delta(x, y, z) \subseteq \bar{B} \subseteq X$, where \bar{B} is a compact ball. By compactness, there is an ε such that for every $x \in \bar{B}$, the ball $B_x(\varepsilon)$ is $CAT(0)$.

We let β_t be the geodesic from x to $\alpha(t)$. Using continuity, we can choose $0 < t_1 < \dots < t_N = 1$ such that

$$d(\beta_{t_i}(s), \beta_{t_{i+1}}(s)) < \varepsilon$$

for all $s \in [0, 1]$.

Now divide Δ up into a “patchwork” of triangles, each contained in an ε ball, so each satisfies the $CAT(0)$ condition, and apply induction and Alexandrov’s lemma to conclude. □

Theorem. Let X be a proper length space of non-positive curvature, and $p, q \in X$. Then each homotopy class of paths from p to q contains a *unique* (local) geodesic representative.

5.6 Gromov's link condition

Theorem (Gromov's link criterion). A Euclidean complex X is non-positively curved iff for every vertex v of X , $Lk(v)$ is CAT(1).

Corollary. If X is a 2-dimensional Euclidean complex, then for all vertices v , $Lk(v)$ is a metric graph, and X is CAT(0) iff $Lk(v)$ has no loop of length $< 2\pi$ for all v .

Theorem (Mal'cev). Every finitely generated linear subgroup (i.e. a subgroup of $GL_n(\mathbb{C})$) is residually finite.

Proof sketch. If the group is in fact a subgroup of $GL_n(\mathbb{Z})$, then we just reduce mod p for $p \gg 0$. To make it work over $GL_n(\mathbb{C})$, we need a suitable version of the Nullstellensatz. \square

Theorem (Mal'cev). Every finitely generated residually finite group is Hopfian.

Proof. Finding a proof is a fun exercise! \square

Lemma (Scott's criterion). Let X be a cell complex, and $G = \pi_1 X$. Then G is residually finite if and only if the following holds:

Let $p : \tilde{X} \rightarrow X$ be the universal cover. For all compact subcomplexes $K \subseteq \tilde{X}$, there is a finite-sheeted cover $X' \rightarrow X$ such that the natural covering map $p' : \tilde{X} \rightarrow X'$ is injective on K .

5.7 Cube complexes

Theorem (Gromov). A cube complex is non-positively curved iff every link is flag.

Lemma. For any (simplicial) graph N , the link of the unique vertex of \mathcal{S}_N is $D(\tilde{N})$. In particular, \mathcal{S}_N is non-positively curved.

Theorem. Right-angled Artin groups embed into $GL_n \mathbb{Z}$ (where n depends on N).

5.8 Special cube complexes

Theorem (Haglund–Wise). If X is a compact special cube complex, then there exists a graph N and a local isometry of cube complexes

$$\varphi_X : X \looparrowright \mathcal{S}_N.$$

Corollary. $\pi_1 X \hookrightarrow A_N$.


Proof of corollary. If $g \in \pi_1 X$, then g is uniquely represented by a local geodesic $\gamma : I \rightarrow X$. Then $\varphi \circ \gamma$ is a local geodesic in \mathcal{S}_N . Since homotopy classes of loops are represented by unique local geodesics, this implies $\varphi_X \circ \gamma$ is not null-homotopic. So the map $(\varphi_X)_*$ is injective. \square

Corollary. If X is a special cube complex, then $\pi_1 X$ is linear, residually finite, Hopfian, etc.

Sketch proof of Haglund–Wise. We have to first come up with an N . We set the vertices of N to be the hyperplanes of X , and we join two vertices iff the hyperplanes cross in X . This gives \mathcal{S}_N . We choose a transverse orientation on each hyperplane of X .

Now we define $\varphi_X : X \looparrowright \mathcal{S}_N$ cell by cell.

- Vertices: There is only one vertex in \mathcal{S}_N .
- Edges: Let e be an edge of X . Then e crosses a unique hyperplane H . Then H is a vertex of N . This corresponds to a generator in A_N , hence a corresponding edge $e(H)$ of \mathcal{S}_N . Send e to $e(H)$. The choice of transverse orientation tells us which way round to do it
- Squares: given hyperplanes



Note that we already mapped e_1, e_2 to $e(H)$, and f_1, f_2 to $e(H')$. Since H and H' cross in X , we know $e(H)$ and $e(H')$ bound a square in \mathcal{S}_N . Send this square in X to that square in \mathcal{S}_N .

- There is nothing to do for the higher-dimensional cubes, since by definition of \mathcal{S}_N , they have all the higher-dimensional cubes we can hope for.

We haven't used a lot of the nice properties of special cube complexes. They are needed to show that the map is a local isometric embedding. What we do is to use the hypothesis to show that the induced map on links is an isometric embedding, which implies φ_X is a local isometry. \square