COURSE NOTES FOR MATH 797: COARSE GEOMETRY AND TEICHMÜLLER THEORY

INSTRUCTOR:
ALEX WRIGHT

MAIN AUTHORS:
SHABARISH CHENAKKOD, JACK HAVILAND, CHINMAYA KAUSIK,
SAYANTAN KHAN, KATIA SHCHETKA, TERESA YU

WITH CONTRIBUTIONS BY:
FRANCISCO ARANA-HERRERA, DYLAN CORDARO, MALAVIKA MUKUNDAN,
YING WANG, CHRIS ZHANG, BRADLEY ZYKOSKI

AND GUEST LECTURES BY:
MARK HAGEN, TIM SUSSE

Contents

Introduction: These are notes on a graduate topics course at the University of Michigan in Winter 2022. Corrections are welcome and may be sent via email to Alex Wright.

Audience and scope: These notes might be useful for students who would like additional preparation before tackling the literature on what has been called the "Masur-Minsky machine", and is now in a broader context called "hierarchical hyperbolicity". We attempt to begin without assuming familiarity with either Teichmüller theory or coarse geometry, and proceed with an example focused approach. The main goal is to get a good idea of what a hierarchically hyperbolic space is and why Teichmüller space is an example. The course concludes with guest lectures by Tim Susse and Mark Hagen on connections between cube complexes and hierarchical hyperbolicity.

Authorship: For each lecture, one course participant was designated as the author, and another as the editor. The notes for each lecture are labelled with the initials of the author followed by the initials of the editor. In addition to the listed main authors and contributors, Paul Apisa and Giuseppe Martone each served as the editor for one or two days each.

Citations: Only a very small number of citations are provided.

Acknowledgements: Alex Wright would like to thank Alessandro Sisto for generously and with great clarity and insight explaining many things related to these notes. He would also like to thank Jacob Russell for numerous very helpful conversations, and Tim Susse and Mark Hagen for sharing their insights as guest lecturers during the last four meetings of the course.

During the period in which this course was taught, there were a number of seminar talks related to hierarchical hyperbolicity at Michigan. These were not part of the course, and are not recorded in these notes, but nonetheless they had a large indirect benefit on the course. Alex Wright would like to thank Daniel Berlyne, Alexandre Martin, Jacob Russell, Alessandro Sisto, and Bin Sun for speaking. Alex Wright would like to especially thank Kasra Rafi and Howard Masur, who both gave expository seminar talks specifically designed to complement the course.

1. Hyperbolicity and quasi-geodesics (01/05, TY, CZ)

Hyperbolic space $\mathbb{H}^n$ is defined by its curvature being a constant -1. While curvature is a local property, this condition on curvature leads to many global features. In contrast to curvature, coarse geometry studies features that are unaffected by changes to the geometry in a small neighborhood. One can think of zooming out so that local features cannot be distinguished. You can only see the forest, not the trees. A Gromov hyperbolic space defined below will be a generalization of hyperbolic space from the point of view of coarse geometry.

Definition 1.1. A **geodesic** is an isometry from an interval in $\mathbb{R}$ to a metric space.

Definition 1.2. A metric space is called **geodesic** if every pair of points is joined by a geodesic.

Example 1.3. The space $\mathbb{R}^2$ minus a ball (with the induced metric from $\mathbb{R}^2$) is not geodesic, as two points on either side of the ball cannot be joined by a geodesic.

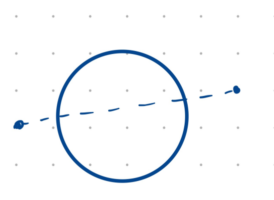

Figure 1

Definition 1.4. A geodesic metric space is called **δ-hyperbolic** if, for any geodesic triangle (a triple of geodesics, each ending where the next begins), each edge is contained in the closed δ-neighborhood of the union of the other two edges.

Definition 1.5. A space is called **(Gromov) hyperbolic** if there exists a $\delta \geqslant 0$ such that it's δ-hyperbolic.

Example 1.6. The space $\mathbb{R}^2$ is not hyperbolic. Taking equilateral triangles of greater and greater side lengths, we also need larger and larger δ in order for one edge to be contained in the δ-neighborhood of the other two edges.

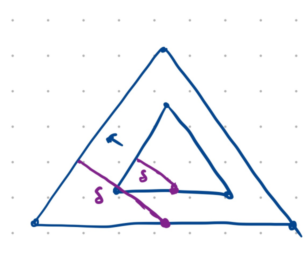

Figure 2

Example 1.7. Trees (collections of vertices and edges with no cycles, and edges have length 1) are 0-hyperbolic. Given a geodesic triangle, every point on one of the geodesics is also on one of the other two geodesics.

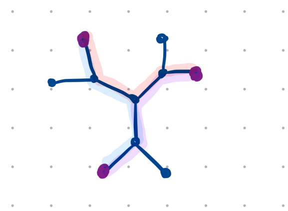

FIGURE 3

Remark 1.8. Trees should be thought of as the best example of hyperbolic spaces since δ can be taken to be 0.

Example 1.9. The real line $\mathbb{R}$ is also 0-hyperbolic, similar to the previous example.

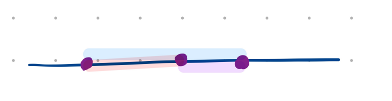

FIGURE 4

Example 1.10. The hyperbolic plane $\mathbb{H}^2$ is hyperbolic.

Example 1.11 (Alexandrov)**.** Let $\kappa < 0$. Any complete, simply connected manifold with curvature $\leqslant \kappa$ is Gromov hyperbolic.

See [BH99, Theorem 1A.6, page 173] for a proof.

Example 1.12. Every space of finite diameter is hyperbolic.

Definition 1.13.
- The map $f : X \to Y$ is a (K, C)**-quasi-isometric embedding** if for all $x, y \in X$,
$$\frac{d(x,y)}{K} - C \leqslant d(f(x), f(y)) \leqslant K d(x,y) + C.$$
- The map f is a **quasi-isometry** if there exists a constant $C \geqslant 0$ such that for all $y \in Y$, there exists $x \in X$ such that $d(f(x), y) \leqslant C$.
- A (K, C)**-quasi-geodesic** in X is a (K, C)-quasi-isometric embedding of an interval in $\mathbb{R}$ into X.

Remark 1.14. The notation of quasi-isometry will be like a notation of isomorphism for coarse geometry. For example, we will discuss later that if a space X is quasi-isometric to a hyperbolic space, then X is hyperbolic.

The following two examples show that quasi-geodesic is not a good notion in $\mathbb{R}^2$ since there are too many quasi-geodesics. We will see later that in a hyperbolic space, a quasi-geodesic will closely follow a geodesic.

Example 1.15. We'll show that the following is an example of a $(2,0)$-quasi-geodesic in $\mathbb{R}^2$:

$$f : \mathbb{R} \to \mathbb{R} : t \mapsto \begin{cases} (t, 0) & \text{if } t \geqslant 0, \\ (0, -t) & \text{if } t \leqslant 0. \end{cases}$$

In words, the real line is wrapped about the origin along the positive y- and x-axes.

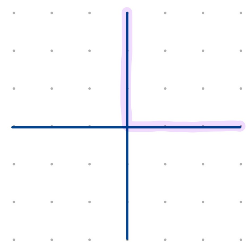

FIGURE 5

Let $s, t \in \mathbb{R}$ with $s < t$. If s, t are both non-negative or both non-positive, then $d(f(s), f(t)) = d(s, t)$, and so

$$\frac{d(s,t)}{2} \leqslant d(f(s), f(t)) = d(s,t) \leqslant 2d(s,t).$$

If $s < 0$ and $t > 0$, then $d(f(s), f(t)) \leqslant d(s,t)$ by the triangle inequality. We also have that $\max(|s|, |t|) \geqslant \frac{d(s,t)}{2}$, and so $\frac{d(s,t)}{2} \leqslant d(f(s), f(t))$. Thus, we also have in this case that

$$\frac{d(s,t)}{2} \leqslant d(f(s), f(t)) \leqslant d(s,t) \leqslant 2d(s,t).$$

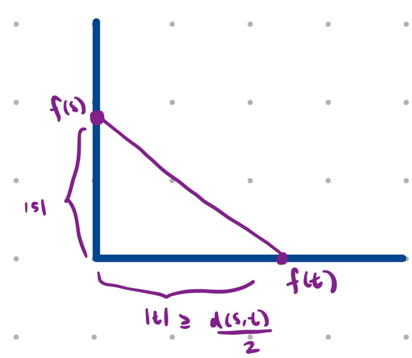

FIGURE 6

Compare to the fact that the L^2 and L^1 metrics are comparable on $\mathbb{R}^2$.

Example 1.16. Consider the embedding f of the interval $[0, \infty)$ in $\mathbb{R}^2$ as a spiral with side lengths given by the sequence $\{\ell_i\}_{i \geqslant 1}$.

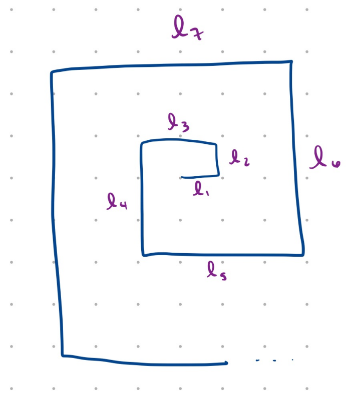

FIGURE 7

For $s, t \in [0, \infty)$ with $s < t$, we always have that $d(f(s), f(t)) \leqslant d(s, t)$. Now, assume that we have the following two conditions on the sequence $\{\ell_i\}$ of side lengths: for each $i > 1$,

(A.) $\ell_i \geqslant \sum_{j=1}^{i-1} \ell_j$, and
(B.) $\ell_i \geqslant 16\ell_{i-1}$.

We'll show that f is a $(8, 0)$-quasi-geodesic. By above, we have the upper bound $d(f(s), f(t)) \leqslant d(s, t) \leqslant 8d(s, t)$, so we'll now show the lower bound, $\frac{d(s,t)}{8} \leqslant d(f(s), f(t))$. Suppose $f(s)$ and $f(t)$ are on the ℓ_j and ℓ_i segments of the spiral respectively, where $i \geqslant j$. We proceed by cases.

(1) If $i = j$, then we have that $d(f(s), f(t)) = d(s, t)$, and so the lower bound holds.
(2) If $i = j + 1$, then as in Example 1.15, we have that $\frac{d(s,t)}{2} \leqslant d(f(s), f(t))$, and so the desired lower bound holds here.
(3) If $i > j + 1$, then let r be such that $f(r)$ is the point between the ℓ_{i-1} and ℓ_{i-2} segments, so $s < r < t$.

By (A.), $\ell_{i-1} \geqslant \sum_{j=1}^{i-2} \ell_i$, and so

$$d(r, t) \geqslant \ell_{i-1} \geqslant \sum_{j-1}^{i-2} \ell_j \geqslant d(s, r).$$

This implies that $d(r, t) \geqslant \frac{d(s,t)}{2}$.

Then, we have that

$$d(f(s), f(t)) \geqslant d(f(r), f(t)) - d(f(s), f(r)) \geqslant \frac{d(r, t)}{2} - 2\ell_{i-2},$$

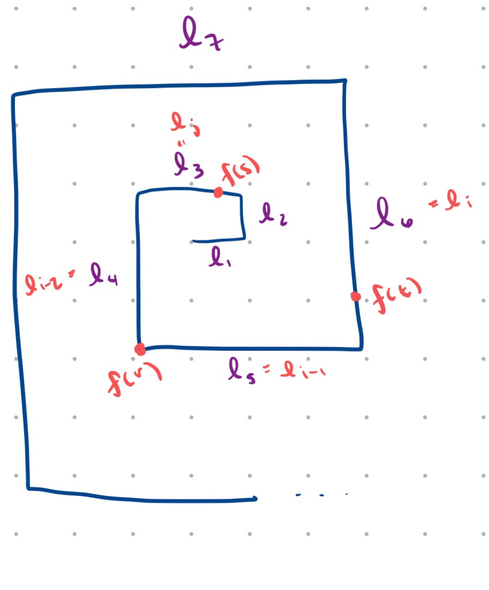

FIGURE 8

where the first inequality comes from the triangle inequality, and the second inequality comes from $d(f(r), f(t)) \geqslant \frac{d(r,t)}{2}$ (similar to Example 1.15) and $d(f(s), f(r)) \leqslant 2\ell_{i-2}$, as $f(s), f(r)$ are contained in a square with side lengths ℓ_{i-2}.

By above, $\frac{d(r,t)}{2} \geqslant \frac{d(s,t)}{4}$, and by (B.), $\frac{d(s,t)}{8} \geqslant \frac{\ell_{i-1}}{8} \geqslant 2\ell_{i-2}$. Thus, we see that

$$d(f(s), f(t)) \geqslant \frac{d(r,t)}{2} - 2\ell_{i-2} \geqslant \frac{d(s,t)}{4} - \frac{d(s,t)}{8} = \frac{d(s,t)}{8},$$

and so the desired lower bound holds in this case.

Optional Exercise 1. Prove that the wedge of two hyperbolic spaces is hyperbolic.

Optional Exercise 2. Prove the easier direction of Manning's Bottleneck Criterion [Man05, Theorem 4.6].

Optional Exercise 3. Prove that the composition of two quasi-isometric embeddings is a quasi-isometric embedding, and that being quasi-isometric is an equivalence relation.

Optional Exercise 4. Prove $\mathbb{R}^2$ and $\mathbb{R}^3$ are not quasi-isometric.

Optional Exercise 5. Formulate and prove a thinness result for geodesic n-gons in a δ-hyperbolic space.

Optional Exercise 6. Consider the wedge of infinitely many copies of the real line. Convince yourself that anything quasi-isometric to this space is hard to compactify. Show that the Gromov boundary is not compact.

2. FELLOW TRAVELLING (01/07, JH, KS)

Last time, we saw that quasi-geodesics are not well-behaved in Euclidean space. In hyperbolic spaces, they are much nicer and naturally arise when using quasi-isometry as a coarse notion of equivalence. Here, "nice" means that quasi-geodesics are not too different from actual geodesics, in the sense that a quasi-geodesic between two points is always close to a geodesic between the points. To formalize this, we need a notion of distance between sets; the usual definition is Hausdorff distance.

Definition 2.1. If $A, B \subseteq X$, the **Hausdorff distance** between them is

$$d_{\text{Haus}}(A,B) = \inf \left\{ \begin{array}{c} R \geqslant 0 \text{ s.t. every point of } A \text{ is distance} \\ \leqslant R \text{ from a point of } B \text{ and vice-versa.} \end{array} \right\}$$

More concisely,

$$d_{\text{Haus}}(A,B) = \inf\{R \geqslant 0 \text{ s.t. } A \subseteq N_R(B) \text{ and } B \subseteq N_R(A)\},$$

where N_R denotes the R-neighborhood of a set.

Now, we can make precise the idea that quasi-geodesics are always close to geodesics in hyperbolic spaces. One reference for this is [Sisb, Proposition 5.4.2].

Proposition 2.2 (Fellow Traveller Property)**.** *Let X be δ-hyperbolic. Then for all K, C there is a D such that any (K,C)-quasi-geodesic α has Hausdorff distance at most D from any geodesic γ joining its endpoints.*

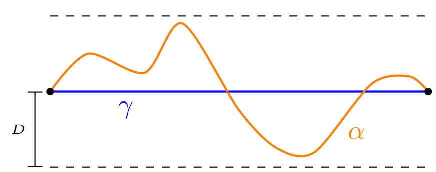

FIGURE 9

Remark 2.3. The quasi-geodesic could be a geodesic, hence "geodesics are coarsely unique in hyperbolic spaces." For this reason, many people refer to "the" geodesic between two points.

Remark 2.4. Proposition 2.2 is very false in $\mathbb{R}^2$, which we can see using Example 1.15, the simpler example from last time. The legs form a quasi-geodesic, even though

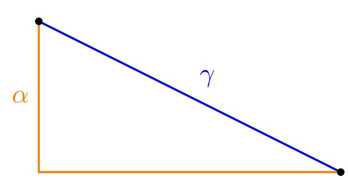

FIGURE 10

the point at their intersection can be arbitrarily far from the actual geodesic (the hypotenuse).

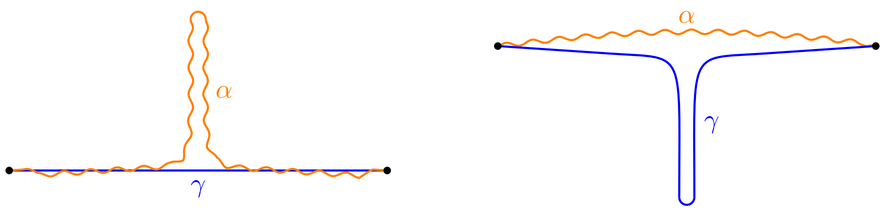

FIGURE 11

Remark 2.5. To prove Proposition 2.2, we need to rule out *both* of the following cases: That is, we need to show both that the quasi-geodesic stays close to the geodesic *and* that the geodesic stays close to the quasi-geodesic.

Example 2.6. In a tree, a (continuous) quasi-geodesic must go through every vertex on the geodesic between two points; this says that the geodesic is never far from the quasi-geodesic. For the other direction, when a quasi-geodesic leaves the geodesic, it must eventually return, and the defining inequalities ensure that this excursion never takes points on the quasi-geodesic too far from the geodesic. More concretely,

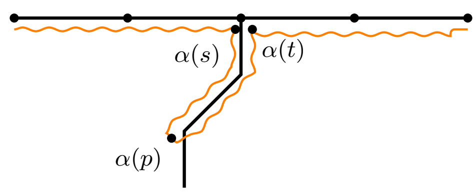

FIGURE 12

$$\frac{|s-t|}{K} - C \leqslant d(\alpha(s), \alpha(t)) = 0 \quad \Longrightarrow \quad |s-t| \leqslant KC,$$

so

$$d(\alpha(s), \alpha(p)) \leqslant K|p-s| + C \leqslant K|s-t| + C \leqslant K^2C + C.$$

That is, we can take $D = K^2C + C$ in Proposition 2.2.

Lemma 2.7. *Let X be δ-hyperbolic . Let $\alpha\colon [a,b] \to X$ be a path with*

$$d(\alpha(s), \alpha(t)) \leqslant K|t-s| + C.$$

Let p be any point on a geodesic from $\alpha(a)$ to $\alpha(b)$, and assume further that $b - a > 1$. Then

$$d(p, \alpha) \leqslant \delta \log_2(b-a) + D,$$

where $D = D(\delta, K, C)$ depends only on δ, K, C.

Remark 2.8. We think of $b - a$ as the length of α. If we want to get from x to y while avoiding a ball centered at the midpoint of radius $d(x,y)/2$ (as shown above), the length required is exponential in $d(x,y)$. The situation is comparable to the fact that balls have circumference $\sim \exp(\text{radius})$ in $\mathbb{H}^2$. This is different than the Euclidean case, where the required length is only linear (a factor of π).

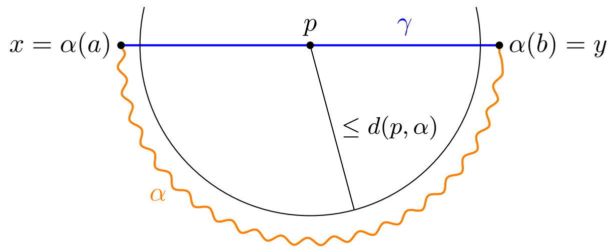

FIGURE 13

For trees, it's actually impossible to avoid large enough balls; this is because there are unavoidable intermediate points when going between any two points in a tree. That being said, small balls (e.g., balls of radius less than $C/2$ for a (K, C)-quasi-geodesic) can be avoided because quasi-geodesics are not required to be continuous, so they could "jump over" the ball since the points on opposite sides would still be within the required distance of one another.

Proof of Lemma 2.7. First, if $b - a \leqslant 2$, then for all $s \in [a, b]$,

$$\begin{aligned} d(p, \alpha(s)) &\leqslant d(p, \alpha(b)) + d(\alpha(b), \alpha(s)) \\ &\leqslant d(\alpha(a), \alpha(b)) + d(\alpha(b), \alpha(s)) \\ &\leqslant K|b-a| + K|b-s| + 2C \leqslant 4K + 2C. \end{aligned}$$

So, take $D = 4K + 2C$.

We induct on n with $2^{n-1} \leqslant b - a < 2^n$. We just took care of the base case $n = 1$, so let $n > 1$ and assume the result holds for $n - 1$. Set $m = (a + b)/2$, the midpoint of a and b, and note that $2^{n-2} \leqslant m - a, b - m < 2^{n-1}$ by our definition of n. Say p

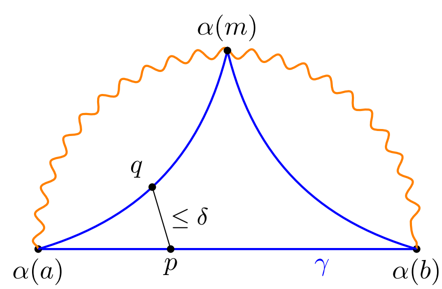

FIGURE 14

is distance at most δ from a point q which lies on the geodesic from $\alpha(a)$ to $\alpha(m)$ (q could also lie on the geodesic from $\alpha(m)$ to $\alpha(b)$, but we may assume WLOG that it does not). By induction,

$$d(q, \alpha|_{[a,m]}) \leqslant \delta \log_2 \left(\frac{b-a}{2} \right) + D = \delta \log_2(b-a) - \delta + D.$$

So,

$$\begin{aligned} d(p,\alpha) &\leqslant d(p,q) + d(q,\alpha|_{[\alpha,m]}) \\ &= \delta + \delta \log_2(b-a) - \delta + D \\ &= \delta \log_2(b-a) + D. \end{aligned}$$

□

Remark 2.9. This proof basically gives us an algorithm to find a point of α close-ish to p.

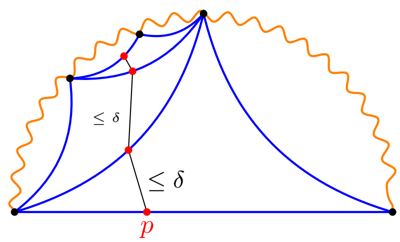

FIGURE 15

We now upgrade Lemma 2.7 for the case when α is a quasi-geodesic (before, we only assumed the upper bound).

Lemma 2.10. *If α is a (K,C)-quasi-geodesic from x to y and γ is a geodesic from x to y, then for all $p \in \gamma$, $d(p,\alpha) \leqslant E$ for some constant $E = E(K,C,\delta)$.*

Remark 2.11. This lemma is an upgrade because it gives a constant bound that does not depend on the length of the parametrizing interval instead of the logarithmic bound from Lemma 2.7.

Proof. Pick a point p on γ as far from α as possible and set $E = d(p,\alpha)$ (such a p exists because γ is compact and distance to α is continuous). Our goal is to find a bound on E that does not depend on p.

Pick points as follows (the picture is more helpful than the descriptions):

- y' by traveling along γ from p for a distance of $2E$, or $y' = y$ if that is not possible.
- y'' on α such that $d(y',y'') \leqslant E$, which is possible since E is the largest distance from a point on γ to a point on α. If $y' = y$, choose $y'' = y' = y$.
- x' by traveling backwards along γ from p for a distance of $2E$, or $x' = x$ if that is not possible.
- x'' on α such that $d(x',x'') \leqslant E$. If $x' = x$, choose $x'' = x' = x$.

Define β to be the concatenation of the geodesic from x' to x'', then α until y'', then the geodesic from y'' to y'. Note the following:

(1) $d(x'',y'') \leqslant 6E$ as there is a path of length $6E$ joining x'' and y''.
(2) Say $\alpha(s) = x''$, $\alpha(t) = y''$. Since α is a quasi-geodesic,

$$\frac{|t-s|}{K} - C \leqslant d(x'',y'') \leqslant 6E,$$

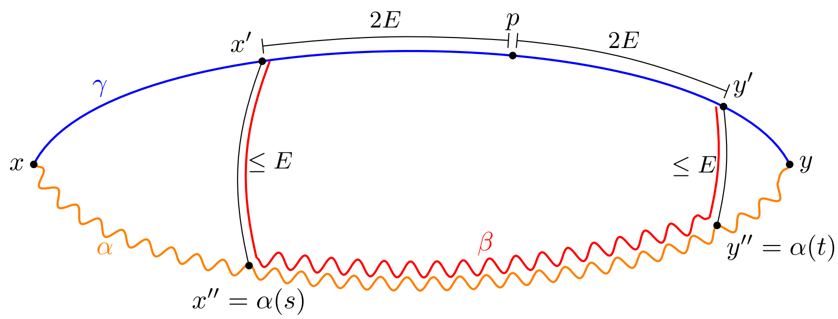

FIGURE 16

hence $|t-s| \leqslant 6EK + CK$.

(3) With its natural parametrization, β is parametrized by an interval of length at most $6EK + CK + 2E$.

(4) It follows from the triangle inequality and α being a quasi-geodesic that β satisfies

$$d(\beta(c), \beta(d)) \leqslant K|d-c| + C.$$

Claim 2.12. $d(p, \beta) \geqslant E$.

Because $d(p, \alpha) = E$ (by definition), we expect equality to hold.

Proof of Claim. If q is a point on the α-part of β, then $d(p, q) \geqslant E$ by the definition of E and choice of p as the farthest point on the geodesic from α. If q is a point on the geodesic from x' to x'' (or y' to y''), then

$$d(p, q) \geqslant d(p, x') - d(x', q) = 2E - E = E,$$

unless $x' = x''$, which happens if $d(p, x) \leqslant 2E$. But in this case, the geodesic from x' to x'' is simply the point x', so $q = x'$ and q lies on α. As before, this means $d(p, q) \geqslant E$ by definition of E and p. ⊗

Finally, applying Lemma 2.7 to β gives

$$d(p, \beta) \leqslant \delta \log_2(6EK + CK + 2E) + D,$$

so Claim 2.12 gives

$$E \leqslant \delta \log_2(6EK + CK + 2E) + D. \tag{1}$$

This implies an upper bound on E because the left hand side grows faster than the right hand side, so if E can be arbitrarily large then the inequality would break. This bound depends only on (1); since $D = D(\delta, K, C)$, this means bound on E depends only on δ, K, and C, as desired. □

3. Fellow travelling (01/10, DC, CK)

Proposition 3.1. *(This is Prop 2.2, the Fellow Traveller Property) Let X be δ-hyperbolic. Then $\forall K, C$, there is a $D = D(\delta, K, C)$ such that any (K, C)-quasi-geodesic has a Hausdorff distance $\leqslant D$ from any geodesic joining its endpoints.*

Last class, we did the harder direction as encapsulated by lemma 2.10, which says that geodesics are uniformly close to quasi-geodesics. The lemma below is the opposite direction, which is easier.

Lemma 3.2. *Let X be a δ-hyperbolic space. If $\alpha : [a, b] \to X$ is a (K, C) quasi-geodesic and γ is a geodesic joining its endpoints, then every point on α is a uniformly bounded distance from some point on γ.*

Proof. The moral goal of this lemma is to exclude the case where the quasi-geodesic has a protruding segment that goes far away from the geodesic (the straight line), as shown below. Before we launch into the proof, the idea is as follows: given a point q (thought of as a point on this long segment), we break α into the parts before and after q and find a point $p \in \gamma$ close to both parts. Using the closeness of p to both sides of q, we show that such a long protruding segment can't be too long and so q can't be too far away from γ.

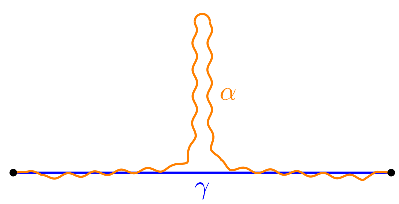

FIGURE 17

Pick a point $q \in \alpha, q = \alpha(c)$. By Lemma 2.10, there exists an $E = E(K, C, \delta)$ such that every point p on γ satisfies $d(p, \alpha) \leqslant E$.

So, each $p \in \gamma$ is distance $\leqslant E$ close to a point on either $\alpha([a, c]) =: \alpha_1$ or $\alpha([b, c]) =: \alpha_2$, or both. The union of the closed neighbourhoods $\bar{N}_E(\alpha_1) \cup \bar{N}_E(\alpha_2)$ thus contains γ, where $N_r(S)$ is the (open) r-neighbourhood of the set S.

Since both the closed neighbourhoods intersect γ (one at each endpoint), $\bar{N}_E(\alpha_1) \cap \gamma$ and $\bar{N}_E(\alpha_1) \cap \gamma$ are both non-empty closed subsets of γ that exhaust it. By the connectedness of γ, $\bar{N}_E(\alpha_1) \cap \bar{N}_E(\alpha_2) \cap \gamma$ is non-empty. So, there exists a point $p \in \gamma$ that satisfies $d(p, \alpha_1) \leqslant E$ and $d(p, \alpha_2) \leqslant E$. See Figure 18.

This implies that there exist points on the segments α_1, α_2 at most E away from p: $\exists s, t$ such that $a \leqslant s \leqslant c \leqslant t \leqslant b$ with

$$d(\alpha(t), p) \leqslant E, \qquad d(\alpha(s), p) \leqslant E$$

By the triangle inequality, these 2 inequalities imply that

$$d(\alpha(s), \alpha(t)) \leqslant 2E.$$

Since α is a (K, C)-quasi-geodesic, we see that

$$\frac{t - s}{K} - C \leqslant 2E \implies t - s \leqslant K(2E + C).$$

This last inequality implies the result because we can show the distance from q to the geodesic is uniformly bounded:

$$\begin{aligned} d(q,p) = d(\alpha(c),p) \leqslant d(\alpha(c),\alpha(t)) + d(\alpha(t),p) &\leqslant (K d_{\mathbb{R}}(c,t) + C) + E \\ &\leqslant K d_{\mathbb{R}}(s,t) + C + E \leqslant K[K(2E+C)] + E + C \end{aligned}$$

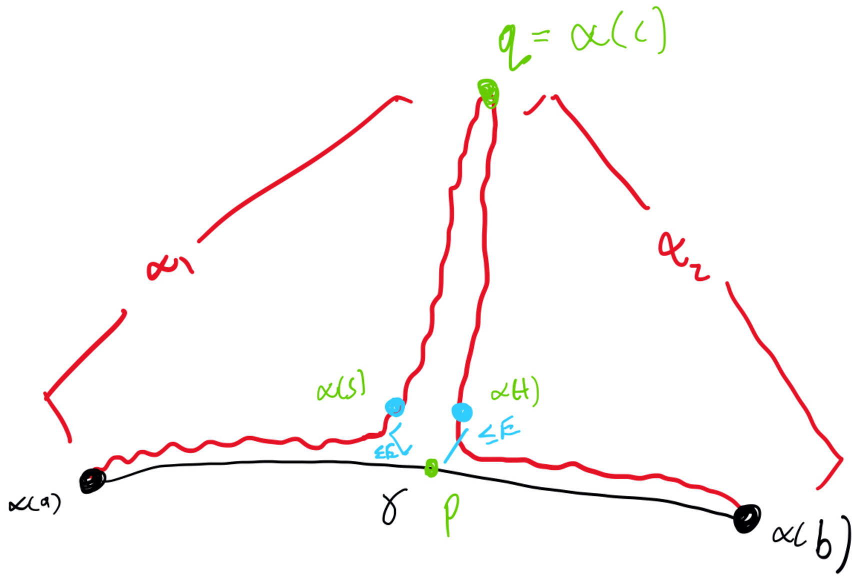

FIGURE 18. An image of the first case in Lemma 3.2, when there is a point $p \in \gamma$ close to both parts of α.

□

Remark 3.3. Suppose $f : X \to Y$ is a (K, C)-quasi-isometric embedding.
If $\gamma : I \to X$ is a geodesic, then $f \circ \gamma : I \to Y$ is a quasi-geodesic.

Proof. Use the definition of a quasi-isometric embedding and that $d_X(\gamma(s), \gamma(t)) = |t - s|$, as we have defined geodesics to be isometries. □

Corollary 3.4. *$\forall \delta, K, C \, \exists \delta'$ such that if Y is δ-hyperbolic and $f : X \to Y$ is a (K, C)-quasi-isometric embedding, then X is δ'-hyperbolic.*

Proof. Let $\gamma_1, \gamma_2, \gamma_3$ form a geodesic triangle T in X, which we aim to show is δ'-thin. We know from Remark 3.3 that $f(\gamma_1), f(\gamma_2), f(\gamma_3)$ is a quasi-geodesic triangle in Y which we will call $f(T)$. Form a geodesic triangle T' in Y from the 3 vertices of $f(T)$. The key is to use the fact that by the Fellow Traveller Property, there is a uniform bound on the change in the thinness of a triangle when we tighten the quasi-geodesics to geodesics. See Figure 19 below.

Since Y is δ-hyperbolic, we know from Proposition 2.2 that the Hausdorff distance between an edge on the geodesic triangle T' and a corresponding edge on the quasi-geodesic triangle $f(T)$ in Y is bounded by a uniform constant $D = D(\delta, K, C)$. Let

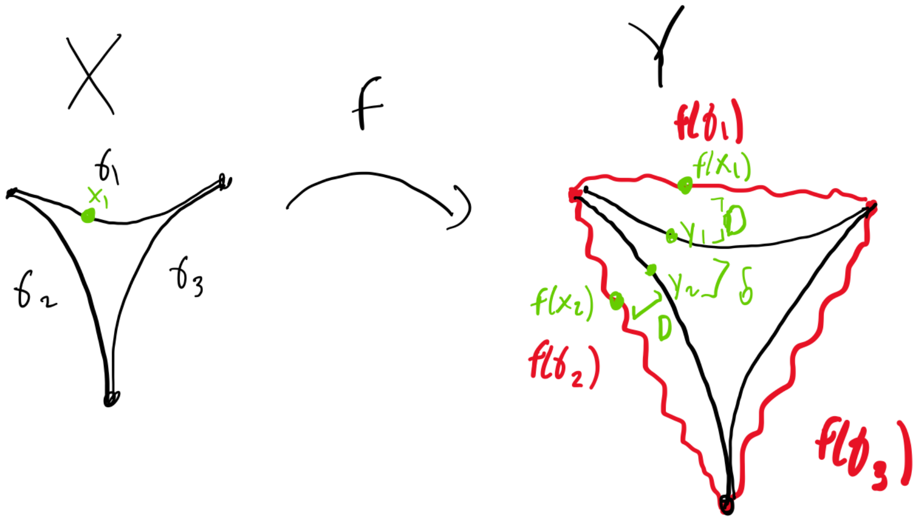

FIGURE 19. For Cor 3.4. The left geodesic triangle is T, which maps to $f(T)$ on the right and is tightened to a geodesic triangle T'.

$x_1 \in \gamma_1$ in T. Since $f(x_1) \in f(\gamma_1)$, there exists a point y_1 on the geodesic triangle $T' \subset Y$ such that

$$d(f(x_1), y_1) \leqslant D.$$

Because Y is $\delta-$hyperbolic, there exists a point y_2 on another edge of the geodesic triangle T' (WLOG the edge corresponding to the quasi-geodesic $f(\gamma_2)$) such that

$$d(y_1, y_2) \leqslant \delta.$$

Finally, we use Proposition 2.2 again to conclude that there exists some $x_2 \in \gamma_2$ such that

$$d(y_2, f(x_2)) \leqslant D.$$

The triangle inequality shows that for the 2 points x_1, x_2 in γ_1, γ_2 respectively

$$d(f(x_1), f(x_2)) \leqslant d(f(x_1), y_1) + d(y_1, y_2) + d(y_2, f(x_2)) \leqslant 2D + \delta.$$

Now, if we use the fact that f is also a quasi-isometric embedding, we see that we have found $x_2 \in \gamma_2$ so that

$$d(x_1, x_2) \leqslant Kd(f(x_1), f(x_2)) + C \leqslant K(2D + \delta) + C =: \delta'.$$

Since x_1 was arbitrary, we have shown that X is δ'-hyperbolic.

□

We will use the techniques used to prove the Fellow Traveller Property (Proposition 2.2) to show the geodesic guessing lemma in the next class.

Optional Exercise 7. Consider a non-Gromov-hyperbolic metric space. Prove that there are geodesics and quasi-geodesics joining the same pair of points that don't stay close together. So the fellow traveling result is actually an "if and only if" characterization of hyperbolicity.

Optional Exercise 8. Suppose X is a geodesic metric space, $C > 0$, and S is a subset of X that is C-dense. Suppose that all geodesic triangles in X with vertices in S are δ-slim. Show that X is $\delta' = \delta'(\delta, C)$ hyperbolic.

Optional Exercise 9. Prove that a geodesic metric space X is quasi-isometric to a proper metric space if there exists a real number $R > 0$ such that any set of points that have pairwise distance in $[2R, 4R]$ is finite.

Conversely, suppose that a metric space X is (K, C)-quasi-isometric to a proper metric space. Show that any set of points that have pairwise distance in $[R, R']$ is finite, for any $R' > R > KC$.

4. Geodesic guessing (01/12, SK, SC)

We will now prove the *Geodesic Guessing Lemma* (which will be labelled as a Proposition rather than a lemma because of its relative importance). Before we state the result though, we define what a family of *geodesic guesses* is, and what it means for the family to be thin.

Definition 4.1 (Geodesic guesses)**.** Let X be a geodesic metric space, and let $D > 0$ be a constant. A family of geodesic guesses is a set of paths $\eta(x, y)$ for all x and y in X satisfying the following conditions.

(1) If $d(x, y) \leqslant 1$, then the diameter of $\eta(x, y)$ is less than D.
(2) For any x' and y' in the path $\eta(x, y)$, any segment of $\eta(x, y)$ starting at x' and ending at y' is within D Hausdorff distance of $\eta(x', y')$[1].

The family is called thin if for all triples $\{x, y, z\}$ in X, $\eta(x, z)$ is contained in a D-neighbourhood of $\eta(x, y) \cup \eta(y, z)$.

Proposition 4.2 (Geodesic Guessing Lemma)**.** *Let X be a geodesic metric space, and suppose it admits a thin family η of geodesic guesses. Then*

- *For all x and y in X, the geodesic γ from x to y is at most Hausdorff distance K from $\eta(x, y)$, where K is some constant only depending on D.*
- *X is a $(2K + D)$-hyperbolic space.*

One reference for this is [Sisa]. Before we prove Proposition 4.2, we need the following lemma.

Lemma 4.3. *Suppose $\gamma : [a, b] \to X$ is a path from x to y satisfying the following inequality for some constants K and C and for all $c, d \in [a, b]$.*

$$d(\gamma(c), \gamma(d)) \leqslant K|c - d| + C$$

Then for any point p on $\eta(x, y)$, we have the following estimate on $d(p, \gamma)$, where A is constant only depending on K, C, and D.

$$d(p, \gamma) \leqslant D \log_2(\max(1, |b - a|)) + A$$

Proof. The proof of this lemma is almost identical to the proof of Lemma 2.7, with all the geodesic segments replaced by the corresponding geodesic guesses, so we skip the proof. $\square$

[1] Note that there may be multiple segments in $\eta(x, y)$ starting and ending at x' and y'.

Remark 4.4. The proof of Proposition 4.2 is quite similar to the proof of Proposition 2.2, with the geodesic guesses in this setting playing the role geodesics played in the proof of Proposition 2.2 and the actual geodesics playing the role quasi-geodesics played in the proof of Proposition 2.2.

Proof of Proposition 4.2. It will suffice to prove the first statement of the proposition, since the second statement follows from the first statement, and the fact that the family of geodesic guesses is D-thin.

It will suffice to prove the following two claims:

(1) Every point on $\eta(x,y)$ is within K distance of a point on a geodesic from x to y.
(2) Every point on a geodesic from x to y is within K distance of a point on $\eta(x,y)$.

Proof of Claim (1). Let γ be a geodesic from x to y, and let $p \in \eta(x,y)$ be a point that is as far as possible from γ. Let $E = d(p,\gamma)$, and pick points x' and y' on $\eta(x,y)$ such that a segment of $\eta(x,y)$ from x' to y' contains p and x' and y' are distance $2E$ from p. If such points don't exist on $\eta(x,y)$, then set $x = x'$ or $y = y'$ (or both, if necessary). Let x'' and y'' be points on γ closest to x' and y'. Let β be the segment obtained by concatenating a geodesic from x' to x'', followed by the segment of γ from x'' to y'', and then a geodesic from y'' to y'. Finally, consider the geodesic guess $\eta(x',y')$ (see Figure 20 for all the points and segments depicted).

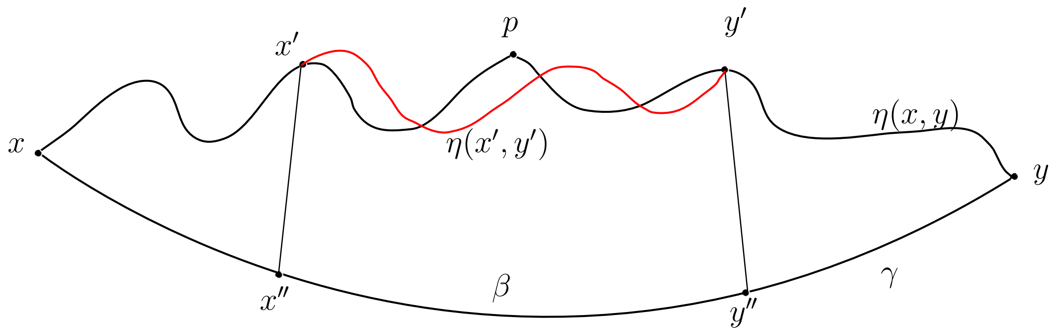

FIGURE 20. The geodesic guesses $\eta(x,y)$ and $\eta(x',y')$ and the geodesic segment γ from x to y.

First of all, note the geodesic segment from x'' to y'' has length at most $6E$, since there's a path from x'' to x' to p to y' to y'' with length at most $6E$. This means that the unit speed parameterization of β has length at most $8E$. Also note that β with the unit speed parameterization satisfies the hypothesis required from the path in Lemma 4.3 with $K = 1$ and $C = 0$. Using Lemma 4.3, we have the following bound for any point $p' \in \eta(x',y')$.

$$d(p',\beta) \leqslant D\log_2(8E) + A$$

We pick p' to be the point on $\eta(x',y')$ closest to p, we use property (2) in the definition of geodesic guesses to deduce that $d(p,p') \leqslant D$, giving us the following inequality.

$$d(p,\beta) \leqslant D\log_2(8E) + A + D$$

We now claim that any point on β that minimizes distance to p lies on γ, and not on the geodesic segments from x' to x'' or y' to y''. The proof of this claim is the same as the proof of Claim 2.12. The fact that a point minimizing the distance between p and β lies on γ gives us a lower bound on $d(p, \beta)$, namely E which is the minimal distance between p and γ. We therefore have the following inequalities.

$$E \leqslant d(p, \beta) \leqslant D \log_2(8E) + A + D$$

From this, we conclude that E is bounded above by a fixed constant J, depending only on D.

Proof of Claim (2). Let γ be a geodesic from x to y, and let q be a point on γ that is as far as possible from $\eta(x, y)$. We know from Claim (1) that every point on $\eta(x, y)$ is within distance J of γ. We color each point in $\eta(x, y)$ red if it is within distance J of the geodesic segment from x to q (we call this the first half of the geodesic) and blue if it is within distance J of the geodesic segment from q to y (we call this the second half of the geodesic). There are two possible cases that can now arise. We describe what happens in the first case, and rule out the second case.

Case 1: There exist points on $\eta(x, y)$ which are red and points which are blue: Since a path is the image of an interval under a coarsely continuous map, there must be arbitrarily close points of the interval whose images p and p' are coloured red and blue respectively. Note that p and p' can be at most D distance apart. Let a and b be points on the geodesic segment from x to q and q to y which are within J distance of p and p' respectively. We then have that $d(a, b) \leqslant 2J + D$, which means $d(p, q) \leqslant 2J + D$, since q lies on the geodesic segment between a and b. (see Figure 21)

Case 2: All of $\eta(x, y)$ is either red or blue: Note that x is always colored red, and y is always colored blue. We therefore rule out this case.

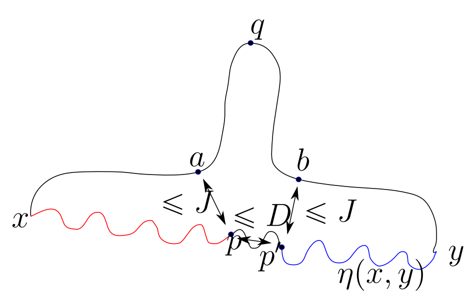

FIGURE 21. Proving that the geodesic lies in a bounded neighbourhood of the geodesic guess.

Letting $K = 2J + D$ proves statement 1 of the proposition, and therefore the entire proposition. □

5. The horoball construction (01/14, YW, SK)

We will show how to construct a hyperbolic space from any geodesic metric space using a construction called the *horoball construction*. Given a group G acting upon a metric space X, the horoball $H(X)$ is a hyperbolic space upon which the group G

continues to act via isometries. This is quite useful, since the group G need not be hyperbolic, giving us plenty of examples of non-hyperbolic groups acting on hyperbolic spaces. References include [GM08, Definition 3.1], [MS20, Definition 2.1].

Definition 5.1 (Horoball construction)**.** Let (X, d) be a geodesic metric space. A horoball $H(X)$ on X is defined as

$$\left(\bigcup_{k=0}^{\infty} X \times \{k\}, 2^{-k} d \right),$$

with an edge of length 1 added from (x, k) to $(x, k+1)$ for any $x \in X, k \geqslant 0$.

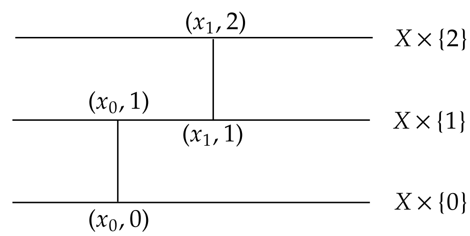

Figure 22

There is a natural path metric on $H(X)$ given by

$$d_{H(X)}(x, y) = \inf \left\{ \sum_{i=0}^{n+1} d(x_i, x_{i+1}) \;\middle|\; \begin{array}{l} x_0 = x, x_{n+1} = y, \text{ for each } i, \, x_i \text{ and } x_{i+1} \\ \text{are either on the same added edge or on} \\ X \times \{k\} \text{ for the same } k. \end{array} \right\}$$

We now give an alternative definition of horoball.

Definition 5.2 (Combinatorial horoball)**.** Let Γ be a graph. Define the combinatorial horoball $H(\Gamma)$ on Γ to be the following metric graph:

(1) vertices are $V(\Gamma) \times \{0, 1, 2, \cdots\}$
(2) If $v, w \in V(\Gamma)$ are joined by an edge, then for any $i \geqslant 0$, join (v, i) and (w, i) by an edge of length 2^{-i}.
(3) For any $v \in V(\Gamma), i \geqslant 0$, join (v, i) and $(v, i+1)$ by an edge of length 1.

Remark 5.3. Every geodesic metric space X is quasi-isometric to a graph Γ by the following construction. Let $V(\Gamma) = \{v_x : x \in X\}$, and join v_x and v_y by an edge of length 1 if $d(x, y) \leqslant 1$. Then the horoball $H(X)$ is quasi-isometric to the combinatorial horoball $H(\Gamma)$.

A natural question to ask is what are the geodesics of $H(X)$. We first prove a lemma that shows we do not need to consider complicated paths in the definition of $d_{H(X)}$ in Definition 5.1:

Lemma 5.4. *Define an up-geodesic-down path to be a path $\gamma = \gamma_1 \cdot \gamma_2 \cdot \gamma_3$ where γ_1 is a path from (x_0, i) to (x_0, j) for $j > i$ via the added edges, γ_2 is a geodesic from (x_0, j) to (y_0, j) in $X \times \{j\}$, and γ_3 is a path from (y_0, j) to (y_0, k) for $k < j$ via the added edge. Similarly we define geodesic-down paths and so on.*

For any path γ in $H(X)$ that equals to a concatenation of finitely many paths in $X \times \{k\}$ for some k or some added edges, there is a up-geodesic-down path joining the same endpoints of at most the same length.

Proof. Let γ be as given. Suppose γ has a segment γ_{ab} from $a = (x, i)$ to $b = (y, j)$ as in the graph. Then γ_{ab} has length $j - i + d$ for $d = d_{X\times\{i\}}((x, i), (y, i))$. But we can replace the segment γ_{ab} by γ'_{ab}, which has a smaller length $j - i + 2^{-j+i}d$. Then replace the horizontal movements by a geodesic in the metric space $(X, 2^{-j}d_X)$. Similarly, we can replace the down-geodesic segments by the geodesic-down segments, and cancel the up-down segments. □

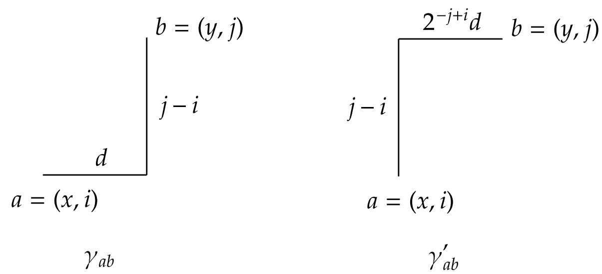

FIGURE 23

Remark 5.5. Suppose $d(x, y) = s$. Then

$$d_{H(X)}((x, i), (y, j)) = \inf\{2h - i - j + 2^{-h}s \mid h \geqslant i, j\}.$$

In the above distance formula, as h becomes very large, the length of the path starts increasing, once $2^{-h}s$ is less than 1. That means the inf of the above quantity is actually achieved in a finite range of values for h, and therefore can be replaced by a min. This shows that $H(X)$ is a geodesic metric space, (recall Definition 1.1).

We now realize the infimum in 5.5 using a concrete formula. Suppose $d_{X\times\{i\}}((x, i), (y, i)) = d$, and we want to see if going up one level higher will end up with a shorter path.

Lemma 5.6. *Suppose we have a path γ from (x, i) to (y, i) in $X \times \{i\}$. We say that γ' is γ going up one level if $\gamma' = \gamma_1 \cdot \gamma_2 \cdot \gamma_3$ where γ_1 is the path from (x, i) to $(x, i+1)$ via the added edge, γ_2 is the image of γ via the translation map $X \times \{i\} \to X \times \{i + 1\}$, and γ_3 is the path from $(y, i + 1)$ to (y, i) via the added edge.*

For any $i \geqslant 0$, going up one level gives up a shorter path if

$$d_{X\times\{i\}}((x, i), (y, i)) = d > 4,$$

and if going up by more than one level reduces length, then going up by exactly one level also reduces length.

Combined with 5.4, we conclude that an up-geodesic-down path is a geodesic if and only if the height of the horizontal geodesic segment is $h = \max(i, j, \lceil \log_2 d \rceil - 2)$.

Proof. The length of the path that goes up one level is $2 + d/2$. Solving for $2 + d/2 < d$ gives us $d > 4$. So the top horizontal segment is smaller than 4 if γ is a geodesic. Now suppose going up by more than one level reduces the length of the path. This means for any $n \geqslant 1$, we have $2n + d/2^n < d$. Solving for this inequality gives $\frac{2^{n+2}n}{2^n-1} < d$. But $\frac{2^{n+2}n}{2^n-1} > \frac{2^{n+2}n}{2^n} = 4n \geqslant 4$. That means going up by one level would have also reduced the length of the path.

In order for a path from (x, i) to (y, j) be a geodesic, the highest copy $X \times \{h\}$ of X we can achieve is when $d_{(X,h)}((x, h), (y, h)) \leqslant 4$. This is, $2^{-h} d_\Gamma(x, y) \leqslant 4$, or $h \geqslant \lceil \log_2 d_\Gamma(x, y) \rceil - 2$. ☐

Corollary 5.7. *If* $d_\Gamma(x, y) > 4$, *then* $d_{H(X)}((x, 0), (y, 0)) \approx 2 \log_2 d$.

Proof. Since $d_\Gamma(x, y) = d > 4$, by 5.6, we have $h = \lceil \log_2 d \rceil - 2$. Then

$$d_{H(X)}((x, 0), (y, 0)) = 2h + 2^{-h} d \approx 2 \log_2 d - 2 + 2^{-\log_2 d} 2^2 d + 2 \approx 2 \log_2 d + 2.$$

☐

Example 5.8. We provide some examples of how geodesics are achieved between $(x, i), (y, j)$. In each of the example, the horizontal segment has length no more than 4.

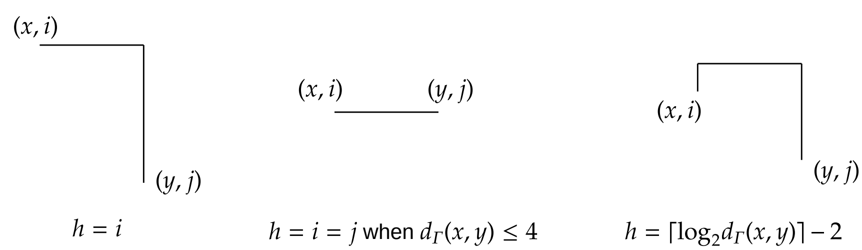

FIGURE 24

Proposition 5.9. *For any geodesic metric space* X, *all triangles in* $H(X)$ *with vertices at integer height are* 5*-slim.*

It follows from this that $H(X)$ is hyperbolic.

Remark 5.10. As a warm up, you might want to verify yourself that if x_1, x_2, x_3 are the vertices of an equilateral triangle in X, which may very well be far from slim, that $(x_1, 0), (x_2, 0), (x_3, 0)$ are the vertices of a slim triangle in $H(X)$. Speaking extremely vaguely, one might say that the horoball construction is "pulling" the midpoints of the edges up to a height at which the metric in the X direction is so contracted that they become close.

Since the geodesics in $H(X)$ always have horizontal distance no more than 4 by 5.6, the vertical movement is the major contribution to the distance between two points. So we need to understand how vertical movements between any three vertices are related to each other, which is studied by the next lemma:

Lemma 5.11. *Let $(x_i, n_i) \in H(X)$ for $i = 1, 2, 3$. For $a \neq b \in \{1, 2, 3\}$, let h_{ab} be the height achieved in 5.6 for paths from (x_a, n_a) to (x_b, n_b). Then $h_{ab} \leqslant \max(h_{ac}, h_{bc}) + 1$.*

Proof. We prove two cases for this lemma and leave the third case to the next class. Recall $h_{ab} = \max(n_a, n_b, \lceil \log_2 d_\Gamma(x_a, x_b)\rceil - 2)$. If $h_{ab} = n_a$, then since $h_{ac} = \max(n_a, n_c, \lceil \log_2 d_\Gamma(x_a, x_c)\rceil - 2)$, we have $h_{ab} = n_a \leqslant h_{ac}$. Similarly $h_{ab} = n_b$ implies $h_{ab} \leqslant h_{bc}$. □

6. The horoball construction (01/19, CK, KS)

Recall that our aim is to show Proposition 5.9, restated below.

Proposition 6.1. *For any geodesic metric space X, all triangles in $H(X)$ with vertices at integer height are 5-slim.*

We were in the process of proving lemma 5.11, restated below.

Lemma 6.2. *Let $(v_a, i_a) \in H(X)$ for $a = 1, 2, 3$. For $a \neq b \in \{1, 2, 3\}$, let h_{ab} be the height achieved in 5.6 for paths from (v_a, i_a) to (v_b, i_b). Then $h_{ab} \leqslant \max(h_{ac}, h_{bc}) + 1$.*

Proof. Recall $h_{ab} = \max(i_a, i_b, \lceil \log_2 d_\Gamma(x_a, x_b)\rceil - 2)$. If $h_{ab} = i_a$, then since $h_{ac} = \max(i_a, i_c, \lceil \log_2 d_\Gamma(x_a, x_c)\rceil - 2)$, we have $h_{ab} = i_a \leqslant h_{ac}$. Similarly $h_{ab} = i_b$ implies $h_{ab} \leqslant h_{bc}$.

If $h_{ab} = \lceil \log_2 d(v_a, v_b)\rceil - 2$, then we note that the triangle inequality implies that

$$d(v_a, v_b) \leqslant d(v_b, v_c) + d(v_a, v_c) \leqslant 2\max(d(v_b, v_c), d(v_a, v_c))$$

Applying $\lceil \log_2(\cdot)\rceil - 2$ to both sides, we have

$$h_{ab} \leqslant 1 + \max(\lceil \log_2 d(v_b, v_c)\rceil - 2, \lceil \log_2 d(v_a, v_c)\rceil - 2) \leqslant \max(h_{bc}, h_{ac}) + 1$$

□

Proof of Proposition 6.1. See Figure 25. The picture is only representative and has some loss of generality, but the argument below holds in general.

Pick p on the geodesic triangle T formed by $(v_a, i_a) \in H(X)$ for $a = 1, 2, 3$. WLOG, p lies on the 12 side. If p lies on the **horizontal part** of 12, then it is at a distance of 4 from either endpoint (see the topmost point p in Figure 25). WLOG, $h_{23} = \max(h_{13}, h_{23})$. So $h_{12} \leqslant h_{23} + 1$ and the endpoints of the horizontal part of 12 are either on the vertical part of 23 or one edge above it. Hence, p is at most a distance $4+1 = 5$ away from 23.

When p is on a vertical edge, WLOG $p = (v_1, k); i_1 \leqslant k \leqslant h_{12}$. We have two cases:

- **Case 1:** $k \leqslant h_{13} + 1$. In that case, p is at most one edge away from the 13 side. So it is at most at a distance of 1 away from the 13 side.

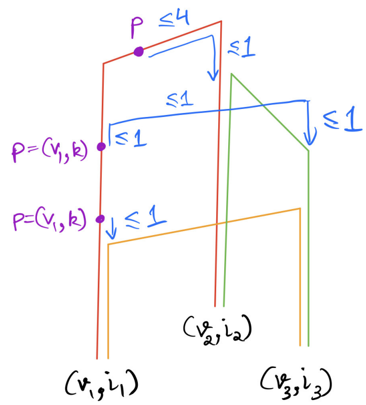

FIGURE 25. The red, yellow and green sides represent sides of the geodesic triangle in the horoball. The proof addresses the three possibilities for p: case 1, case 2 and horizontal edge respectively, bottom to top.

- **Case 2:** $k > h_{13} + 1$. So $\lceil k \rceil \geqslant h_{13} + 2$ and since $d((v_1, h_{13}), (v_3, h_{13})) \leqslant 4$, $d((v_1, \lceil k \rceil), (v_2, \lceil k \rceil)) \leqslant 4 * 2^{-2} = 1$. Note also that Lemma 6.2 implies that $k \leqslant h_{23} + 1$. So, as illustrated by the second path in the figure, p can travel up to $(v_1, \lceil k \rceil)$ by a distance of at most 1, from $(v_1, \lceil k \rceil)$ to $(v_3, \lceil k \rceil)$ by a distance of at most 1 to end up at most 1 edge above the 23 side. So it is at most at a distance of $1 + 1 + 1 = 3$ away from the 23 side.

$\square$

Remark 6.3. Given a space X and a "nice" collection of subspaces X_i, X is called hyperbolic relative to X_i if it is hyperbolic after gluing $H(X_i)$ to $X_i \subset X$. (We might not discuss relative hyperbolicity much in this course, so we'll leave this a bit vague for now.)

We now define a "downward extension" $E(X)$ of the horoball $H(X)$. $E(X)$ can be thought of as related to the horoball the way the hyperbolic plane is to the horodisk at ∞.

Definition 6.4. If (X, d) is a geodesic metric space, we define the *extended horoball* on X as

$$E(X) = \left(\bigcup_{k=-\infty}^{\infty} (X \times \{k\}, 2^{-k} d) \right),$$

with an edge of length 1 added from (x, k) to $(x, k+1)$ for any $x \in X, k \in \mathbb{Z}$. See Figure 26.

It is actually a fact that $E(\mathbb{R})$ is quasi-isometric to $\mathbb{H}^2$ and $H(\mathbb{R})$ is quasi-isometric to the horodisk at ∞ in $\mathbb{H}^2$.

Definition 6.5. For $k \in \mathbb{Z}$ define

$$H_k(X) = \left(\bigcup_{r=k}^{\infty} (X \times \{r\}, 2^{-r} d) \right),$$

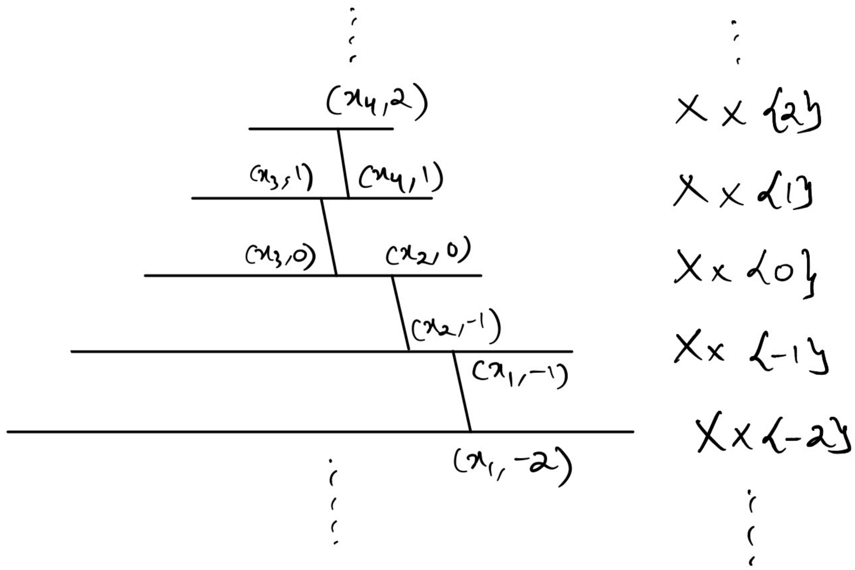

FIGURE 26. The extended horoball.

with an edge of length 1 added from (x, r) to $(x, r+1)$ for any $x \in X, r \geqslant k$.

Remark 6.6. Note that $H_k(X) \cong H((X, 2^{-k}d))$ under the map $(x, r) \mapsto (x, r-k)$.

The remark above shows that each $H_k(X)$ is 5-hyperbolic, since by Proposition 6.1, any horoball is 5-hyperbolic.

Remark 6.7. Under the map $(x, r) \mapsto (x, r)$

$$H_k(X) \subset H_{k-1}(X) \subset H_{k-2}(X) \subset \ldots \text{ and } \bigcup_{k=0}^{\infty} H_{-k}(X) = E(X)$$

By repeating our argument from the previous class, geodesics in $E(X)$ must take the up-geodesic-down path as well. Then $H_k(X)$ is convex in $E(X)$, since any up-geodesic-down path does not go lower in height than its endpoints.

Lemma 6.8. *$E(X)$ is hyperbolic.*

Proof. Pick any three points $x, y, z \in E(X)$. Each lies in some $H_{-k}(X)$. Pick the highest $k = j$, say. By our inclusions in the remark above, all three points lie in $H_{-j}(X)$. Since this set is convex, any geodesic triangle joining the three points also lies in $H_{-j}(X)$. Since $H_{-j}(X)$ is 5-hyperbolic, this triangle is 5-slim. So, $E(X)$ is 5-hyperbolic. □

We can also collapse a horoball in the extended horoball to get an important construction, the hyperbolic cone.

Definition 6.9. If (X, d) is a geodesic metric space, define the *hyperbolic cone* on X to be the quotient $C(X) = E(X)/H_1(X)$, equivalently defined as

$$\{0\} \bigcup \left(\bigcup_{k=0}^{\infty} (X \times \{k\}, 2^{-k}d) \right),$$

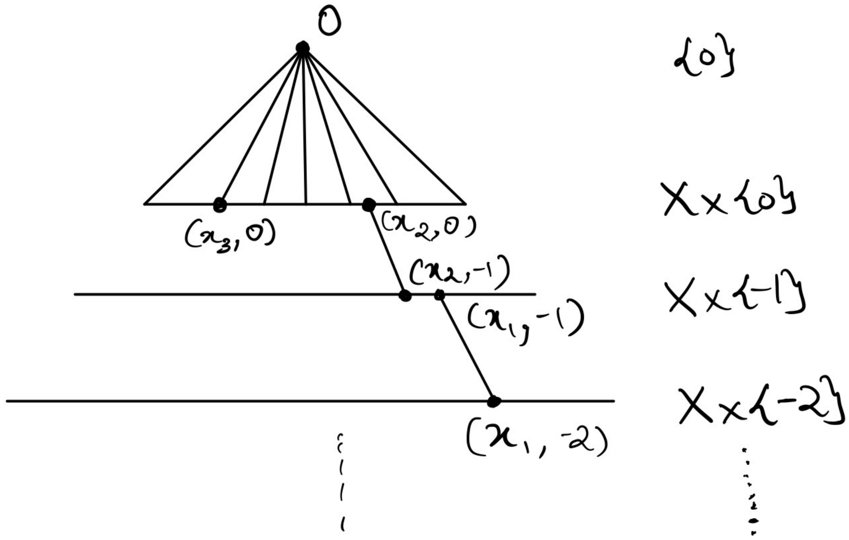

FIGURE 27. The hyperbolic cone.

with an edge of length 1 added from (x, k) to $(x, k+1)$ for any $x \in X, k \geqslant 0$ as well as an edge of length 1 from 0 to $(x, 0)$ for any $x \in X$.

We will show that $C(X)$ is hyperbolic in the next lecture.

Optional Exercise 10. Show that $H(\mathbb{R})$ is quasi-isometric to a the subset of the upper half plane with $\mathrm{Im}(z) > 1$, with the hyperbolic metric.

Optional Exercise 11. Show that for any X, the Gromov boundary of $H(X)$ is a point.

7. Electrification (01/21, SC, KH)

We will now exhibit an operation that allows us to get new hyperbolic spaces from existing ones.

Definition 7.1 (Electrification/Coning Off)**.** Let X be a metric space and let $\{X_j\}_{j\in J}$ be a collection of subspaces. We define $Cone_{\{X_j\}_{j\in J}}(X)$ to be the disjoint union of X with a collection of points $\{c_j\}_{j\in J}$ along with an edge of length 1 from each point of X_j to c_j, for every j. This procedure is called coning off X along $\{X_j\}_{j\in J}$. The resulting space is also called the electrification of X along $\{X_j\}_{j\in J}$.

We also have the path metric defined similarly as in the case of the horoball -

$$d_{Cone(X)}(x, y) = \inf \left\{ \sum_{i=0}^{n+1} d(x_i, x_{i+1}) \;\middle|\; \begin{array}{l} x_0 = x, x_{n+1} = y, \text{ for each } i, x_i \text{ and } x_{i+1} \\ \text{are either on the same added edge or on } X \end{array} \right\}$$

We see that each X_j has diameter less than or equal to 2 under the path metric. We now see some examples,

Example 7.2. Consider the coning off of X along a single subspace X_0. The distance between any two points x and y changes only if there is a path passing through c_0 with

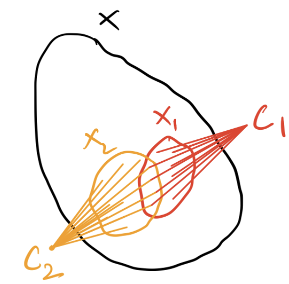

FIGURE 28. Coning off X along X_1 and X_2

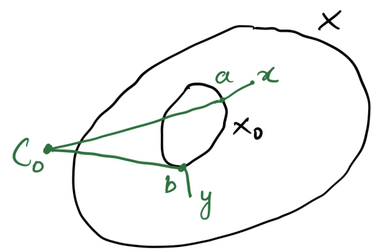

FIGURE 29. The path realising the distance between x and y in the coned off space

smaller length than $d(x, y)$. In this case, we may assume that the path visits c_0 only once. In case X is proper and X_0 is closed, we have points a and b in X_0 that are closest to x and y respectively, and the distance in the coned off space then realised along a path from x to a, a to c_0, c_0 to b and b to y (Figure 29). We thus have the expression,

$$d_{Cone_{X_0}(X)}(x, y) = \min(d_X(x, y), 2 + d_X(x, a) + d_X(y, b))$$

Since the distance function in the space resulting from collapsing X_0 to a point has the same form as above with 2 replaced by 0, we see that the coning off of X along X_0 is $(1, 2)$ quasi-isometric to the space resulting from collapsing X_0 to a point.

This is not true in general. If the X_j's overlap, then collapsing and coning off could be very different, as can be seen with the help of the following example-

Example 7.3. Let X be $\mathbb{R}$ and $X_j = [j, j + 10], j \in \mathbb{Z}$. In this case, collapsing all the X_j's gives us only a point, but coning off X_j's gives a space that is quasi-isometric to $\mathbb{R}$ with the quasi-isometry coarsely reducing distances by a factor of 5 (Figure 30).

Example 7.4. Recall the space $C(X)$ from the previous lecture obtained by collapsing a horoball in the extended horoball. The preceding discussion tells us that $C(X)$ is quasi-isometric to $Cone_{H_1(X)}E(X)$ (Figure 31).

We will show that coning off a hyperbolic space along quasiconvex subspaces gives us a hyperbolic space, for which we first define what quasiconvex subsets are.

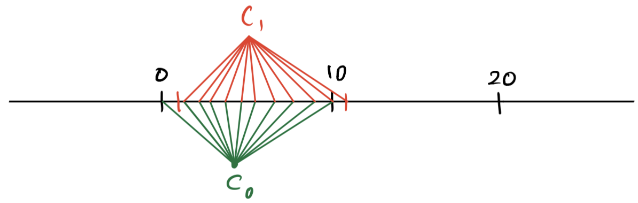

FIGURE 30. Coning off $\mathbb{R}$ along subspaces of the form $[j, j+10]$

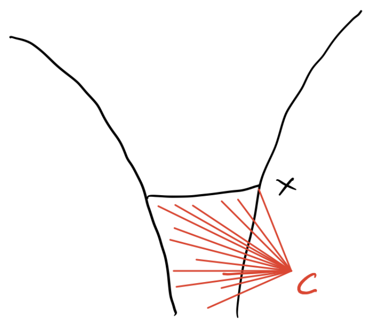

FIGURE 31. The space $Cone_{H_1(X)}E(X)$

Definition 7.5 (Quasiconvexity)**.** A subset Y of a geodesic metric space X is said to be C-quasiconvex if any geodesic segment of X with end points in Y stays in $N_C(Y)$.

For example, the graph of $\sin(x)$ in $\mathbb{R}^2$ is 2-quasiconvex, since the geodesic between any two points of the graph stays in $\mathbb{R} \times [-1, 1]$, and the 2-neighbourhood of the graph of $\sin(x)$ contains this set. On the other hand the graph of $|x|$ is not quasiconvex, since no D-neighbourhood of the graph contains the convex hull of the graph. We also note that with this definition, a 0-quasiconvex set is simply a convex set. We can now state the proposition,

Proposition 7.6. *Let X be δ-hyperbolic and let $X_j \subset X, j \in J$ all be C-quasiconvex, then $Cone_{\{X_j\}_{j\in J}}(X)$ is $\delta' = \delta'(\delta, C)$ hyperbolic, if it is a geodesic metric space.*

Proposition 7.6 and Example 7.4 together give us the corollary,

Corollary 7.7. *For a geodesic metric space X, $C(X)$ is hyperbolic.*

The heavy lifting in proving Proposition 7.6 will be done by the following lemma from [KR14, Proposition 2.5].

Lemma 7.8. *Let X be δ-hyperbolic and Y be a geodesic metric space. Let $f : X \to Y$ be a map satisfying the following conditions -*

(1) $d_Y(f(x), f(y)) \leqslant L d_X(x, y)$ for all $x, y \in X$, i.e., f is L-Lipschitz, for some $L > 0$.
(2) $N_C(f(X)) = Y$, i.e., f is C-surjective, for some $C > 0$.

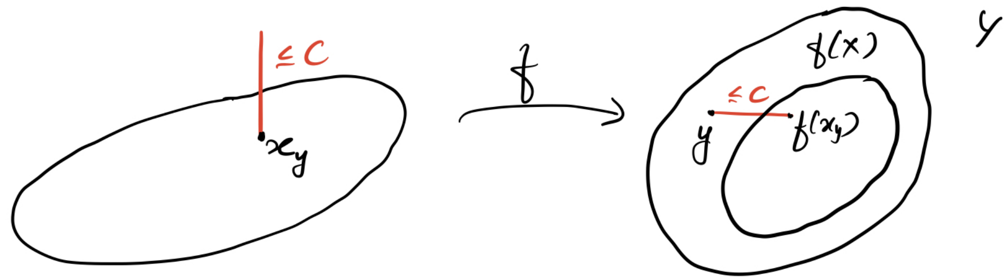

FIGURE 32. Wedging lines with the aim of extending f to a surjective map

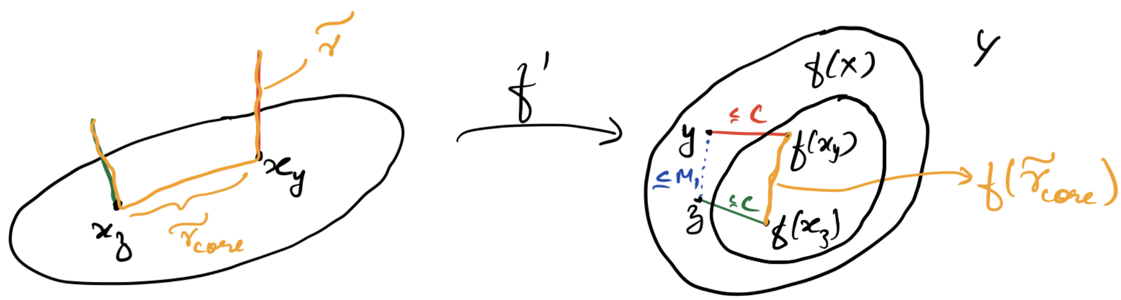

FIGURE 33. Showing that f' satisfies condition 3 of Lemma 7.8

(3) There exist $M_1, M_2 > 0$ such that if γ is a geodesic from x to y and $d(f(x), f(y)) \leqslant M_1 + 2C$, then $\operatorname{diam}(f(\gamma)) \leqslant M_2$.

Then, Y is hyperbolic and there exists D such that, for each geodesic segment γ in X, we have a corresponding geodesic segment γ' in Y with $d_{Haus}(f(\gamma), \gamma') \leqslant D$.

Remark 7.9. Restriction (3) is absolutely crucial. Indeed, every finitely generated group is a quotient of a free group via a 1-Lipschitz, surjective map; but of course not every finitely generated group is hyperbolic!

Proof. We first show that, without loss of generality, we may assume f to be surjective. To see this, we construct a new space X' from X by wedging a line segment on X for each $y \in Y \backslash f(X)$. For every such y, by C-surjectivity, there exists an $x_y \in X$ with $d_Y(f(x_y), x) \leqslant C$. The line segment in X' associated to y has length $d_Y(f(x_y), y)$ and is wedged to X along the point x_y (Figure 32). We then have a natural extension f' to X' of the map f on X, given by mapping each new line segment to a geodesic segment between $f(x)$ and y. It can be shown that the map f' is $(\max(1, L))$-Lipschitz.

It remains to show that f' satisfies condition 3 of the lemma. To see this, consider $y_1, y_2 \in Y$ with $d(y_1, y_2) \leqslant M_1$. Let $\tilde{\gamma}$ be a geodesic from a point in $(f')^{-1}(\{y_1\})$ to a point in $(f')^{-1}(\{y_2\})$. Then, there is a segment $\tilde{\gamma}_{core}$ of $\tilde{\gamma}$ lying entirely in X (Figure 33). There are at most two segments of $\tilde{\gamma}$ other than $\tilde{\gamma}_{core}$ at each end of $\tilde{\gamma}$, and each of these segments has length at most C. f' maps these segments isometrically onto their images.

An application of the triangle inequality then shows that the end points of $f(\tilde{\gamma}_{core})$ are not more than $2C + M_1$ apart, which then gives us that $f(\tilde{\gamma}_{core})$ has diameter less than or equal to M_2. Consequently, $f(\tilde{\gamma})$ has diameter less than or equal to $M_2 + 2C$. This shows that f' satisfies condition 3 of the lemma with new constants $M_1' = M_1 - 2C$ and $M_2' = M_2 + 2C$. Thus, we henceforth assume that f is surjective.

Next, for all $y_1, y_2 \in Y$, we pick a geodesic $\gamma(y_1, y_2)$ joining a point in $(f')^{-1}(\{y_1\})$ to a point in $(f')^{-1}(\{y_2\})$, and define $\eta(y_1, y_2)$ to be $f(\gamma(y_1, y_2))$. We shall use these η's as geodesic guesses in the geodesic guessing lemma and continue the proof in the next lecture. □

8. Electrification (01/24, KS, TY)

Our goal is to prove Proposition 7.6. Let us restate it here.

Proposition 8.1. *Let X be δ-hyperbolic and let $X_j \subset X, j \in J$ all be C-quasiconvex, then $Cone_{\{X_j\}_{j\in J}}(X)$ is $\delta' = \delta'(\delta, C)$ hyperbolic, if it is a geodesic metric space.*

To prove the proposition we need Lemma 7.8, which is also restated below.

Lemma 8.2. *Let X be δ-hyperbolic and Y be a geodesic metric space. Let $f : X \to Y$ be a map satisfying the following conditions*

(1) $d_Y(f(x), f(y)) \leqslant L d_X(x, y)$ for all $x, y \in X$, i.e., f is L-Lipschitz, for some $L > 0$.
(2) $N_C(f(X)) = Y$, i.e., f is C-coarsely surjective, for some $C > 0$.
(3) There exist $M_1, M_2 > 0$ such that if γ is a geodesic from x to y and $d(f(x), f(y)) \leqslant M_1 + 2C$, then $\operatorname{diam}(f(\gamma)) \leqslant M_2$.

Then, Y is hyperbolic and there exists D such that, for each geodesic segment γ in X, we have a corresponding geodesic segment γ' in Y with $d_{Haus}(f(\gamma), \gamma') \leqslant D$.

Proof. In the previous lecture we proved that WLOG $C = 0$, i.e., f is surjective.

For all $y_1, y_2 \in Y$ pick a geodesic $\gamma(y_1, y_2)$ in X from a point of $f^{-1}(y_1)$ to a point of $f^{-1}(y_2)$, and set $\eta(y_1, y_2) = f(\gamma(y_1, y_2))$. We will check conditions in the Geodesic Guessing Lemma 4.2.

The first condition (if $d(y_1, y_2) \leqslant 1$, then the diameter of $\eta(y_1, y_2)$ is bounded by a constant) is satisfied by assumption (3).

We'll now show the second condition: for any z_1 and z_2 in the path $\eta(y_1, y_2)$, the segment of $\eta(y_1, y_2)$ starting at z_1 and ending at z_2 is within bounded Hausdorff distance of $\eta(z_1, z_2)$ (see Figure 34).

It suffices to check that if $\gamma, \gamma' : [0, 1] \to X$ are constant speed geodesics and $f(\gamma(0)) = f(\gamma'(0))$ and $f(\gamma(1)) = f(\gamma'(1))$, then $f(\gamma)$ and $f(\gamma')$ are Hausdorff close (in Figure 34, γ corresponds to the green segment and γ' to the red segment).

WLOG, $p \in \gamma$ (the argument is symmetric with respect to γ and γ'), and so p is δ-close to a point $p' \in \alpha_1 \cup \alpha_2$ by the thinness of geodesic triangles in X (see Figure 35).

FIGURE 34. Showing that segment of $\eta(y_1, y_2)$ (in green) is Hausdorff close to $\eta(z_1, z_2)$

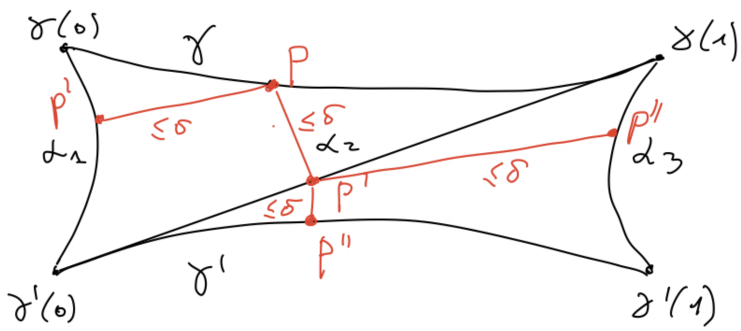

FIGURE 35. Showing that $f(\gamma)$ is Hausdorff close to $f(\gamma')$

There are two cases:

- Suppose $p' \in \alpha_1$. Note that the endpoints of α_1 map to the same point of Y, so diam $f(\alpha_1) \leqslant M_2$ by assumption. So $f(p')$ is close to $f(\gamma'(0))$. Since f is L-Lipschitz, $f(p')$ is at most δL from $f(p)$, so $f(p)$ is close to $f(\gamma')$.
- Suppose $p' \in \alpha_2$. Then p' is δ-close to a point $p'' \in \gamma' \cup \alpha_3$. If $p'' \in \alpha_3$, then it is the same situation as in the previous case. If $p'' \in \gamma'$, then $f(p)$ is $2\delta L$-close to $f(p'')$, so we are done.

Finally, we check the third condition of the Geodesic Guessing Lemma: that for all triples $\{y_1, y_2, y_3\}$, $\eta(y_1, y_3)$ is contained in a D-neighbourhood of $\eta(y_1, y_2) \cup \eta(y_2, y_3)$. Consider $y_1, y_2, y_3 \in Y$. By the work we just did when checking the second condition, it suffices to show that, for $i = 1, 2, 3$, if $x_i \in f^{-1}(y_i)$ and γ_{ij} is a geodesic from x_i to x_j, then the $f(\gamma_{ij})$ form a slim triangle. This follows from $\gamma_1 \cup \gamma_2 \cup \gamma_3$ being slim and f being L-Lipschitz. □

Now let us prove the proposition.

Proof of Proposition 8.1. Let $f : X \hookrightarrow Cone_{\{X_j\}}(X)$ be the inclusion. It is 1-Lipschitz and 1-coarsely surjective. Let γ be a geodesic from x to y in X, and assume $d(f(x), f(y)) <$

2.1. So, the geodesic γ' in $Cone_{\{X_j\}}(X)$ can go through at most one cone point. If γ' doesn't go through a cone point then $d_X(x,y) = d(f(x), f(y)) < 2.1$, so diam $\gamma < 2.1$ and diam $f(\gamma) < 2.1$.

Now assume γ' does go through a cone point: Say it goes from x to $x' \in X_j$ along a geodesic in X, then takes the edge to the cone point c_j associated to X_j, then takes the edge from c_j to $y' \in X_j$, and then goes to y along a geodesic in X. See Figure 36. We need to use this information to show γ has bounded diameter in $Cone_{\{X_j\}}(X)$, which we will do by showing γ stays close to X_j and hence c_j.

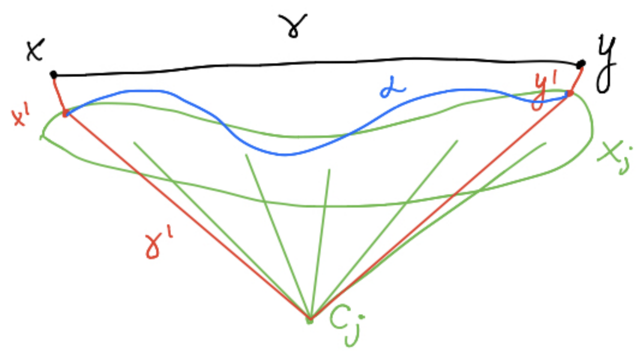

FIGURE 36. Showing that geodesic γ is close to X_j

Since $2.1 > d(x,y) = 1 + 1 + d_X(x,x') + d_X(y,y')$ we have $d_X(x,x') < 0.1$ and $d_X(y,y') < 0.1$. Let α be a geodesic from x' to y' in X. By quasiconvexity of X_j, for any $p \in \alpha$, one has $d(p, X_j) \leqslant C$. Then every point on γ is at most distance $2\delta + 1$ from α. Then $d(p, X_j) \leqslant 2\delta + C + 1$, so diam $f(\gamma) \leqslant 2(2 + 2\delta + C + 1)$. In particular, we can then apply Lemma 8.2 to see that $Cone_{\{X_j\}}(X)$ is hyperbolic. $\square$

Example 8.3. The following non-example shows that the quasiconvexity condition cannot be dropped for electrification to be hyperbolic.

Let $X = [0, \infty)$ and $X_0 = \{2^n, n \geqslant 0\}$. Then $Cone_{X_0}(X)$ is not hyperbolic. It looks like a wedge of an infinte number of bigger and bigger circles, and in particular it doesn't have coarsely unique geodesics.

Example 8.4. Let X be hyperbolic and X_j for $j \in J$ be a geodesic segment. The following is an analogy for the electrification of a space along a collection of geodesic segments. There are a bunch of train lines, which all go straight (but can go over and under each other). It is a magical train: there is always a train waiting for you at any point on any train line; it leaves immediately when you get on; and it takes exactly 2 minutes to get to where you want to get off (on the same line), no matter how far away it is. Despite the magic, even if your start and endpoints are connected by the train system, it might take a long time to get between them, simply because you might need to transfer lines a lot of times.

Example 8.5. Fix δ and C, and suppose we are considering the electrification of a δ-hyperbolic space X along a collection of C-quasi-convex subsets X_j. We actually know that every geodesic in X stays D close to a geodesic in the coned off space, where D depends only on, and can be explicitly computed from, δ and C.

If we happen to know that any two X_j are more than D apart in X, it is much easier to understand the coned off space. For example, suppose γ is a geodesic in X, and γ' is a geodesic joining the same endpoints in the coned off space. Suppose γ' goes through the cone point c_j associated to X_j, coming along the edge from $x' \in X_j$ and exiting along the edge to $y' \in X_j$. So, x' and y' are within distance D of γ in the coned off space. Since the different X_i are more than D apart, we get that x' and y' are within distance D of γ in X. So, γ' can't go through c_j unless γ comes close to X_j. This isn't a complete answer to the question of which cone points γ' goes through, but it does make that question comparatively easy.

This isn't true in general: especially if D is large and the X_j overlap, it can be quite challenging to figure out which cone points a geodesic γ' will go through, and it might go through cone points of sets X_j that aren't close to the original geodesic $\gamma \subset X$.

Optional Exercise 12. Let T be a tree, and let S be a connected subset of T. (a) Show that each point in T has a unique closest point in S. (b) Let π_S denote the closest point projection. Show that if U is a connected subset of T disjoint from S, then $\pi_S(U)$ is a point.

Optional Exercise 13. Let X be a tree, and let $X_j, j \in J$ be an arbitrary collection of connected subtrees. Prove that $Cone_{\{X_j\}}(X)$ is a quasi-tree, using Manning's bottleneck condition (from Exercise 2). Is there a more explicit way to see this space is qi to a tree?

Optional Exercise 14. Determine the Gromov boundary of $E(X)$ and $C(X)$ for arbitrary X.

Optional Exercise 15. Show that every bi-infinite geodesic in X gives rise to a qi-embedding of the hyperbolic plane in $E(X)$.

Optional Exercise 16. Let F_n be the free group, and let X be its Caley graph, i.e. the $2n$-regular tree. (a) Let H_S be the subgroup generated by a subset S of the generators. Show that H_S, as well as all its cosets, are convex in X. (b) Let R denote a set of proper subsets S of the set of generators of F_n. Consider the space Y defined by coning off all cosets of all $H_S, S \in R$. Show R is hyperbolic. (c) Show Y is infinite diameter.

Optional Exercise 17. Let X be a hyperbolic space, and let $r : [0, \infty) \to X$ be a geodesic ray. Let R be the image of r. Relate the Gromov boundary of X to that of $Cone_R(X)$. (See [DT17, Theorem 3.2] for the ultimate generalization of this.)

Optional Exercise 18. Let X be a hyperbolic space, and let $p \in X$. Fix numbers

$$0 = r_0 < r_1 < r_2 < \ldots < r_k.$$

For any $q \in X$, let i be the largest number with $d(p,q) > r_i$, and pick a geodesic segment γ_q from q to a point that is distance r_i from p. Let Y be X with all $\gamma_q, q \in X$ coned off. Show Y has bounded diameter.

A natural special case is when X is the hyperbolic plane, and γ_q goes straight towards p. This special case can be interesting is because a geodesic in Y between two points of $X \subset Y$ might go through cone points of some γ_q that are far in X from the geodesic joining the two points. Why doesn't this contradict "geodesics in Y are Hausdorff close to geodesics in X"?

9. Closest point projections (01/26, JH, YW)

We now develop several results about closest point projections in δ-hyperbolic spaces. First, we establish notation for the distance from a point to a subspace.

Definition 9.1. For any metric space X, if $x \in X$ and $S \subseteq X$, set

$$d(x,S) = \inf_{s \in S} d(x,s) = d_{\text{Haus}}(\{x\}, S).$$

The first thing we want to show about closest point projections is that they are (coarsely) well-defined.

Example 9.2. As a motivating example, suppose $S \subseteq \mathbb{R}^2$ is convex and closed. Then any $x \in \mathbb{R}^2$ has a unique closest point in S. Indeed, if there were two distinct closest points $s, s' \in S$, then their midpoint would be closer (see Figure 37).

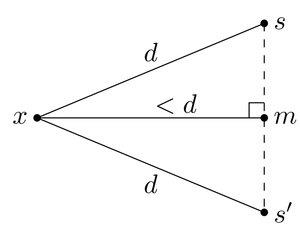

Figure 37. Taking the midpoint of two projections yields a closer projection in $\mathbb{R}^2$.

If $m = (s+s')/2$ is the midpoint of s and s' and $d = d(x,s) = d(x,s')$, then

$$d(x,m) = \sqrt{d(x,s)^2 - d(m,s)^2} < d.$$

More simply, $d(x,m) < d$ because the legs of a right triangle are shorter than its hypotenuse. Since S is convex, $m \in S$, so this contradicts the fact that s was a closest point in S to x.

More generally, the same argument goes through in CAT(0) spaces.

In our coarse setting, hyperbolicity yields a similar result, although as usual strict uniqueness is lost. In this case, rather than a single closest point, we have a bounded diameter set of closest points. The loss of uniqueness is ultimately inconsequential because from the coarse perspective, a bounded-diameter set is essentially the same as a point.

Lemma 9.3. *Suppose X is δ-hyperbolic and $S \subseteq X$ is C-quasi-convex. For $x \in X$, set*

$$\Pi_S(x) = \{s \in S : d(x,s) \leqslant d(x,S) + 1\}.$$

Then $\operatorname{diam} \Pi_S(x) \leqslant D = D(\delta, C)$.

We think of $\Pi_S(x)$ as being the closest points in S to x. We take the points with distance at most $d(x,S) + 1$, rather than just $d(x,S)$, so that we don't need to worry about whether the infimum in $d(x,S)$ is realized. If $x \in S$, then the infimum is always realized (by x), but because of the $+1$, $\Pi_S(x)$ is the closed ball of radius 1 around x, intersected with S.

Remark 9.4. We could arbitrarily pick a point $\pi_S(x) \in \Pi_S(x)$ to get a coarse projection map $\pi_S \colon X \to S$. For this reason, some authors/papers treat Π_S as an actual closest-point projection map, even though it is only coarsely well-defined. In this case, π_S is not strictly a projection map, since it may not restrict to the identity on S; but it will only move points in S distance at most 1 and it would be easy to modify π_S to restrict to the identity.

Proof of Lemma 9.3. Take $y, z \in \Pi_S(x)$ and let m be the midpoint of a geodesic from y to z (see Figure 38). Setting $D = d(y,z)$, we want to find a uniform bound on D depending only on δ and C.

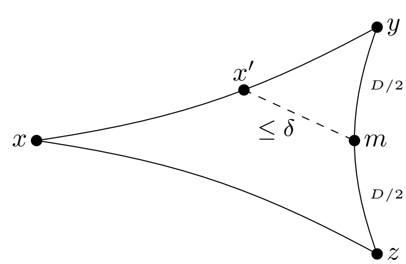

FIGURE 38. Showing two projections of the same point are bounded distance apart.

As S is C-quasi-convex, $d(m,S) \leqslant C$. Since the geodesic x-y-z triangle is thin, there is a point x' on the geodesic from x to (without loss of generality) y such that $d(m,x') \leqslant \delta$. Now,

$$\begin{aligned} d(x,x') + d(x',y) &= d(x,y) \\ &\leqslant d(x,S) + 1 \\ &\leqslant d(x,x') + d(x',m) + d(m,S) + 1 \\ &\leqslant d(x,x') + \delta + C + 1. \end{aligned}$$

This implies $d(x',y) \leqslant \delta + C + 1$, so finally,

$$\frac{D}{2} = d(m,y) \leqslant d(m,x') + d(x',y) \leqslant \delta + \delta + C + 1.$$

Hence, $D \leqslant 4\delta + 2C + 2$, giving the desired bound. $\square$

Remark 9.5. The 2 in this bound comes from the fact that we allowed a maximum error of 1 in the definition of $\Pi_S(x)$, and the $2C$ comes from quasi-convexity. Because we could modify the definition of $\Pi_S(x)$ to use some smaller amount of error (as long as we still use a positive number), and C disappears if we assume S is actually convex, in some sense the essential part of the bound is 4δ. The case $\delta = 0$ occurs in trees, so in trees we do have a well-defined closest point projection onto closed convex sets.

Next, we consider projections of subsets, especially geodesics.

Definition 9.6. If $\gamma \subseteq X$, define

$$\Pi_S(\gamma) = \bigcup_{x\in\gamma} \Pi_S(x) = \operatorname{image}(\Pi_S|_\gamma).$$

Note that if we replace Π_S by a projection map π_S, then we could instead define $\pi_S(\gamma) = \{\pi_S(x) : x \in \gamma\}$.

Example 9.7. As a warm-up for the next lemma, let X be a tree, $S \subseteq X$ convex (equivalently, connected) and closed, and $\gamma \subseteq X$ convex and closed with $\gamma \cap S = \varnothing$. In this case, Π_S can be an honest map (the strict closest point projection) because our subspaces are closed, and $\Pi_S(\gamma)$ is a point. The idea is that any two connected subsets of a tree can be separated by removing a single point of the tree, so any geodesic from one subset to the other has to go through that bottleneck.

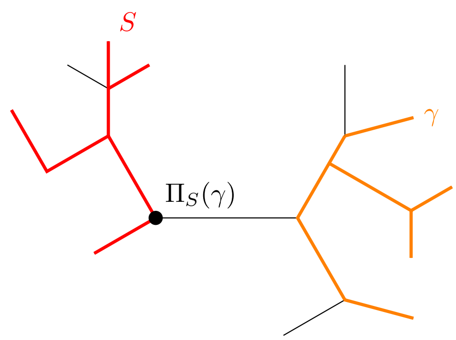

FIGURE 39. Projecting a convex subset of a tree onto another yields a single point.

Remark 9.8. Example 9.7 fails in $\mathbb{R}^2$. For example, when S and γ are two distinct parallel lines, $\Pi_S(\gamma)$ is all of S.

Even though the result does not hold in $\mathbb{R}^2$, we know that hyperbolic spaces are more like trees than $\mathbb{R}^2$, so we expect a version of Example 9.7 to hold.

Lemma 9.9. *Let X be δ-hyperbolic and $S \subseteq X$ a C-quasi-convex subspace. Then there exists a $B = B(\delta, C) > 0$ such that if γ is a geodesic segment with*

$$\gamma \cap N_{C+2\delta+1}(S) = \varnothing,$$

then $\operatorname{diam} \Pi_S(\gamma) \leqslant B$.

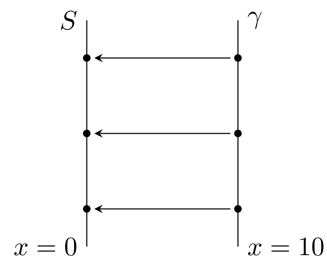

FIGURE 40. Projecting a line onto a parallel line.

This says that if γ is a geodesic and stays far enough away from S, then its projection onto S has bounded diameter. The proof will be very similar to the proof of Lemma 9.3.

Proof. Say γ joins x to y. Pick $p_x \in \Pi_S(x)$, $p_y \in \Pi_S(y)$, and say $B = d(p_x, p_y)$. We want to find a uniform bound on B depending only on δ and C. Let m be the midpoint of a geodesic from p_x to p_y, and let α_x and α_y be geodesics from x to p_x and y to p_y (see Figure 41).

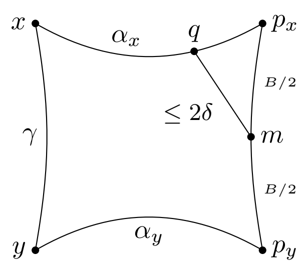

FIGURE 41. Showing the projections of points in γ have bounded distance apart.

Note that $d(m, \gamma) \geqslant 2\delta + 1$ because γ doesn't come within distance $C + 2\delta + 1$ of S and $d(m, S) \leqslant C$ by quasi-convexity. Recall that quadrilaterals in a δ-hyperbolic space are 2δ-thin (this is easily proved by drawing a diagonal and applying thinness of triangles and the triangle inequality). Therefore, there is a point q in $\alpha_x \cup \alpha_y \cup \gamma$ with $d(m, q) \leqslant 2\delta$. Since $d(m, \gamma) > 2\delta$, q cannot lie on γ, so without loss of generality assume q lies on α_x.

Now,

$$\begin{aligned} d(x, q) + d(q, p_x) &= d(x, p_x) \\ &\leqslant d(x, S) + 1 \\ &\leqslant d(x, q) + d(q, m) + d(m, S) + 1 \\ &\leqslant d(x, q) + 2\delta + C + 1. \end{aligned}$$

This implies that $d(q, p_x) \leqslant 2\delta + C + 1$, so

$$\frac{B}{2} = d(m, p_x) \leqslant d(m, q) + d(q, p_x) \leqslant 2\delta + 2\delta + C + 1.$$

Hence, $B \leqslant 8\delta + 2C + 2$, giving the uniform bound we wanted. □

Remark 9.10. There is not really a weaker condition than γ being a geodesic (or at least quasi-geodesic) that we could use for this lemma. We need the thinness condition, and the geodesic guessing lemma tells us that any reasonable family of paths with the thinness property stays Hausdorff close to geodesics, so there is no difference from the coarse perspective.

Example 9.11. The condition $\gamma \cap N_{C+2\delta+1}(S) = \varnothing$ cannot be replaced with $\gamma \cap S = \varnothing$. For example, let T be an arbitrary hyperbolic space. Then $T \times [0,1]$ is also hyperbolic, and $T \times \{0\}$ and $T \times \{1\}$ are disjoint and convex, but the projection of one onto the other is surjective. The point is that "disjoint" is not very robust notion in coarse geometry.

Lemma 9.9 has an application that sets up an important perspective. Suppose that we have a convex (or quasi-convex) set S and two points x and y such that $\Pi_S(x)$ and $\Pi_S(y)$ are "far apart" (see Figure 42). Lemma 9.9 tells us that if a geodesic from x to y does not get close to S, then its projection onto S has bounded diameter. That is, moving along the geodesic cannot significantly change the projection to S. Therefore, if two points *do* have significantly different projections onto S, any geodesic joining them must get close to S.

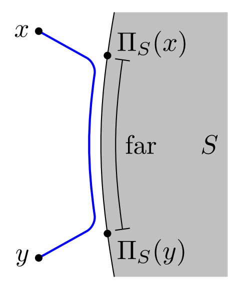

FIGURE 42. Understanding geodesics by projecting them onto a convex set.

This means knowing the projection of two points x, y onto S, and especially knowing the projections are far apart, gives us partial information about the points. It tells us that a geodesic γ joining them must travel to S, then stay close to S, and then travel back away from S (see Figure 42). Moreover, while γ is traveling to S from x and to y from S, it does not make any appreciable progress in S, in the sense that the projections of those parts of γ have bounded diameter.

This idea is related to the fact that it is exponentially expensive to avoid going through a ball (see Remark 2.8) in hyperbolic space; in fact it is really a more general version. It turns out that understanding geodesics between points by looking at their projections onto a convex subspace is a very useful way to get information. It will be an important perspective when we study Teichmüller space — we will talk a lot about projections as a way to build intuition for things that would be very mysterious otherwise.

10. Intersection number (01/28, MM, SC)

Let $S = S_g$ be a surface of genus $g \geqslant 2$.

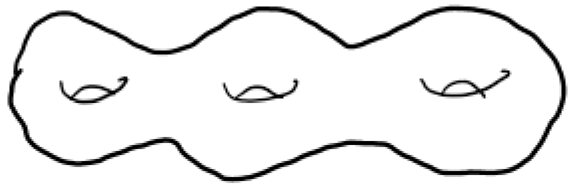

Figure 43. Genus 3 surface

A curve α on S is called *essential* if it is not null-homotopic, and *simple* if it has no self-intersections. (A simple curve is null-homotopic if and only if it bounds a disc.) Throughout this section, we will refer to simple, essential closed curves as simple closed curves (scc's).

The *type* of α is given by the homeomorphism class of $S \backslash \alpha$. In particular,

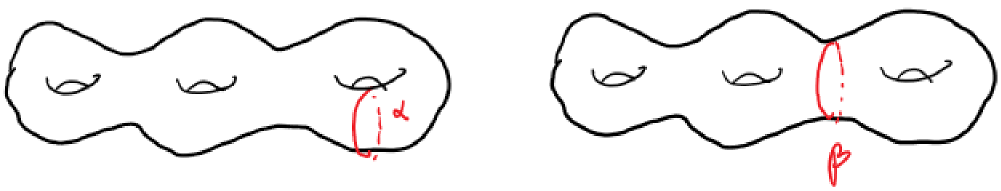

Figure 44. α is non-separating; β is separating

- α is called *non- separating* if $S \backslash \alpha$ is a genus $g - 1$ surface with two boundary circles
- α is called *separating* if $S \backslash \alpha$ is the union of a surface of genus g_1 with one boundary circle and a surface of g_2 with one boundary circle, where $g_1 + g_2 = g$

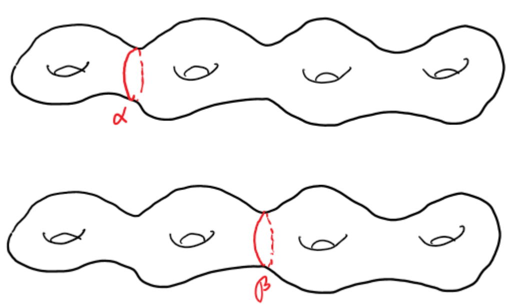

Figure 45. Types of separating curves of a genus 4 surface

As the genus increases, we get more and more types of separating curves. For a genus g surface, there are $\lfloor \frac{g}{2} \rfloor$ types of separating scc's.

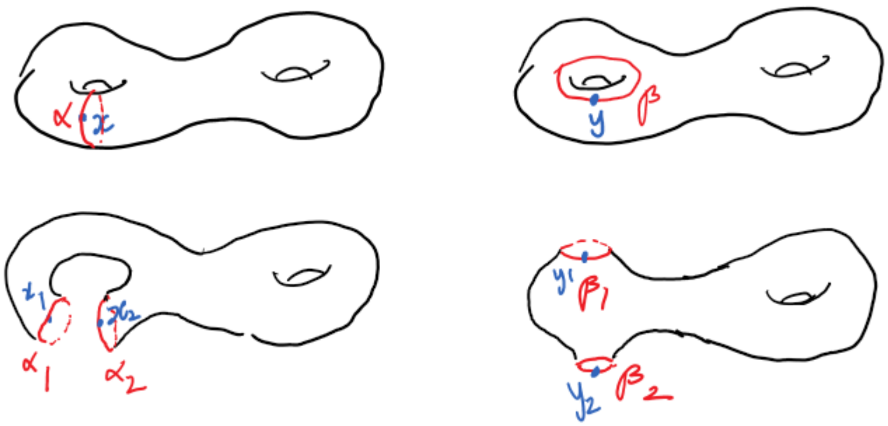

FIGURE 46. Separating curves of same type

Lemma 10.1. *α, β are scc's of same type if and only if there exists a homeomorphism $e : S \longrightarrow S$ such that $e(\alpha) = \beta$*

Proof. $\Longleftarrow$ is clear.

For $\Longrightarrow$:

If α, β have same type, then there exists a homeomorphism $\tilde{e} : S\backslash\alpha \longrightarrow S\backslash\beta$. Let $\alpha_1, \alpha_2, \beta_1, \beta_2$ be the copies of the cut curves α, β respectively as shown in Figure 46. Any $x \in \alpha$ has copies $x_i \in \alpha_i$, and any $y \in \beta$ has copies $y_i \in \beta_i$.

The only problem in gluing up via $\tilde{e}$ is that it is possible that for some $x \in \alpha$, $\tilde{e}(x_1), \tilde{e}(x_2)$ are not copies of the same point $y \in \beta$.

But this problem can be easily fixed by slightly altering $\tilde{e}$ in a neighborhood of the boundary of $S\backslash\alpha$. That is, $\exists e$ isotopic to $\tilde{e}$ such that $\forall x \in \alpha$, $e(x_1) = y_1, e(x_2) = y_2$ for a unique $y \in \beta$. This e glues up into a homeomorphism $e : S \longrightarrow S$ with $e(\alpha) = \beta$. ☐

The above lemma is proved in Chapter 1 of [FM12]. Its philosophy is analogous to the fact that all bases of a linear vector space are equivalent.

Definition 10.2. Given scc's α, β, we define the intersection number $i(\alpha, \beta)$ as the minimum number of intersections between scc's α', β' in the homotpy classes of α, β respectively. That is,

$$i(\alpha, \beta) = \min\{\#(\alpha' \cap \beta') | \alpha' \text{ is homotopic to } \alpha, \beta' \text{ is homotopic to } \beta\}$$

Example:

- $i(\alpha, \alpha) = 0$ (orientability of S is necessary here. If α is the central curve of a Möbius strip for instance, $i(\alpha, \alpha) = 1$).
- See Figure 47.

Definition 10.3. Given scc's α, β, a *bigon* is a disc bounded by a segment of α and a segment of β.

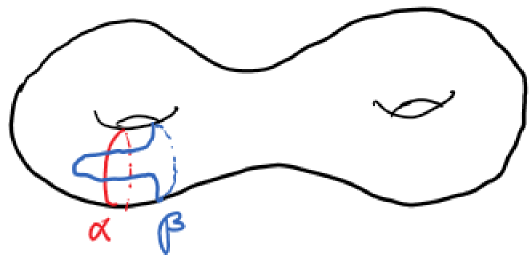

FIGURE 47. β can be homotoped into a curve that is disjoint from α, and so, $i(\alpha, \beta) = 0$

Example:

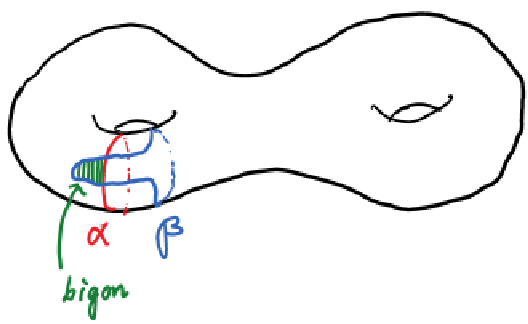

FIGURE 48. Bigon

Lemma 10.4. *(The Bigon Criterion)*
If simple closed curves α and β are transverse and do not make bigons, then

$$i(\alpha, \beta) = \#(\alpha \cap \beta)$$

Remark 10.5. If α, β make a bigon, then we expect to be able to slide β over the bigon to produce a curve β' homotopic to β such that $\#(\alpha \cap \beta') < \#(\alpha \cap \beta)$. See Exercise 19 for how to pick which bigon to start with to make this totally clear. This partially proves the lemma. But we need to check that in the absence of bigons, there exist no homotopies that reduce the intersection number.

Corollary 10.6. *If $\#(\alpha \cap \beta) = 1$, then $i(\alpha, \beta) = 1$.*

The above is clear since two curves intersecting at a single point cannot form a bigon.

Lemma 10.7. *Any collection of scc's can be homotoped so that no pair of them forms a bigon.*

Remark 10.8. An isotopy between two scc's is a homotopy $h : S \times [0, 1] \longrightarrow S$ between them such that $\forall t \in [0, 1]$, $h(., t)$ traces out a simple closed curve. Two scc's are homotopic if and only if they are isotopic. We will denote the isotopy class of α by $[\alpha]$.

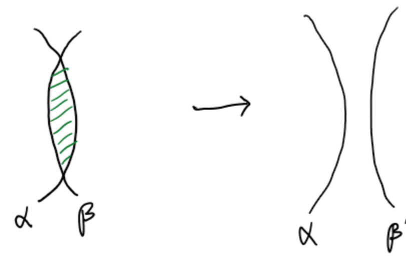

FIGURE 49. Sliding over a bigon

We will now introduce the *curve graph* $\mathscr{C}S$. It has connections with the geometry of Teichmüller space. We will later see that it is connected and hyperbolic, and is not locally compact.

Definition 10.9. The curve graph $\mathscr{C}S$ is defined in the following way:

- its vertex set is $\{[\alpha]\}$ where $[\alpha]$ ranges over all isotopy classes of scc's on S
- there is an edge joining $[\alpha]$ and $[\beta]$ if and only if $i(\alpha, \beta) = 0$

This graph has a natural metric - the distance between any two vertices is the length of the shortest length path connecting them. (You might say two vertices with no path connecting them are an infinite distance apart, but we'll see immediately that this doesn't happen.) One reference on the curve complex is [Sch].

If $i(\alpha, \beta) = 0$, then $dist([\alpha], [\beta]) = 1$. We also have the following.

Lemma 10.10. *If* $i(\alpha, \beta) \neq 0$*, then*

$$dist([\alpha], [\beta]) \leqslant 2 \log_2 i(\alpha, \beta) + 2$$

A reference is [Hem01, Lemma 2.1].

Corollary 10.11. *$\mathscr{C}S$ is connected.*

We'll start with some warm up discussion related the lemma, which in particular will provide the base case when we do an inductive proof of the lemma next class.

Note that

$$dist([\alpha], [\beta]) = \min\{n | \exists \text{ scc's } \gamma_0, \gamma_1, ..., \gamma_n \text{ with } \gamma_0 = \alpha, \gamma_n = \beta, \gamma_i \cap \gamma_{i+1} = \varnothing \forall i\}$$

- If $i(\alpha, \beta) = 1$, we let γ be the boundary of the ϵ neighborhood $N_\epsilon(\alpha \cup \beta)$ of the two curves. For ϵ small enough, $\overline{N_\epsilon(\alpha \cup \beta)}$ is a torus with one boundary component (see Figure 50). We note that

  $$S \backslash \gamma = N_\epsilon(\alpha \cup \beta) \sqcup S'$$

  In particular, if $g \geqslant 2$, S' cannot be a disk, and so, γ is not null-homotopic. Clearly,

  $$i(\alpha, \gamma) = i(\gamma, \beta) = 0$$

  Therefore,

  $$dist([\alpha], [\beta]) = 2$$

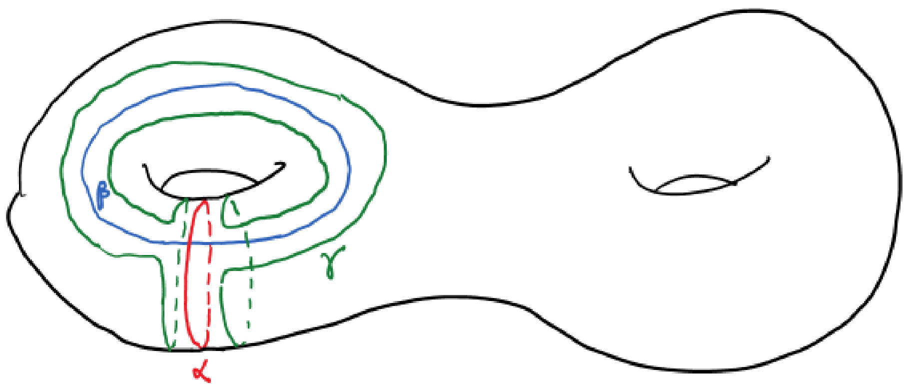

FIGURE 50. α, β have geometric intersection number 1

- Suppose $dist([\alpha],[\beta]) > 2$, then for any scc γ, $i(\alpha,\gamma) > 0$ or $i(\beta,\gamma) > 0$ or both. In this case, we say that α and β *fill* the surface. Now cut along α, β. This gives a collection of polygons, each of which have an even number of edges (see Figure 51). These constitute a cell-decomposition of S, with V vertices, E

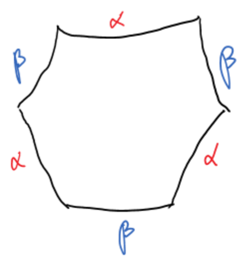

FIGURE 51. A possible connected component of $S \backslash (\alpha \cup \beta)$

edges end F faces. We note that

$$\begin{aligned} V &= i(\alpha,\beta) \\ E &= 2i(\alpha,\beta) \\ F &\geqslant 1 \end{aligned}$$

Thus,

$$\begin{aligned} 2g-2 &= -\chi(S) \\ &= -V+E-F \\ &\leqslant -i(\alpha,\beta)+2i(\alpha,\beta)-1 \\ &\implies i(\alpha,\beta) \geqslant 2g-1 \end{aligned}$$

Optional Exercise 19. Suppose α and β are simple closed curves on S. Suppose there is disc on the surface bounded by a segment of α and a segment of β. Prove there is a component of $S-(\alpha\cup\beta)$ that is a disc and is bounded by a segment of α and

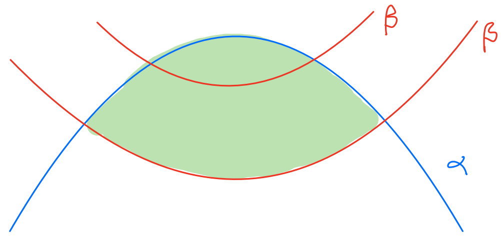

FIGURE 52. The green bigon is not disjoint from β

a segment of β. That is, show that if α and β have a bigon, then they have a bigon where the interior of the bigon is disjoint from α and β. (Sometimes this disjointness is included in the definition. See Figure 52.)

Optional Exercise 20. Think about intersection number for non-simple curves: Suppose α and β are primitive, which means they aren't obtained by traversing a smaller loop multiple times. Suppose neither has immersed monogons, which means they have simple lifts in the universal cover. Prove that the number of intersections is equal to the intersection number if and only if there are no *immersed* bigons. Or equivalently, if and only if any lift of α to the universal cover does not form any bigons with any lift of β. (Hint: It seems helpful to think about geodesic representatives for a hyperbolic metric.)

11. Connectivity of the curve complex (01/31, TY, CK)

Proof of Lemma 10.10. Recall that α, β are assumed to have no bigons. We proceed by induction on $i(\alpha, \beta)$.

First, if $i(\alpha, \beta) = 1$, then we saw last class that $dist(\alpha, \beta) = 2$. Thus, we see that

$$dist(\alpha, \beta) = 2 \leqslant 2\log_2 i(\alpha, \beta) + 2 = 2\log_2 1 + 2 = 2,$$

as desired.

Now, assume $i(\alpha, \beta) \geqslant 2$, and that the lemma holds for all pairs of curves with smaller intersection number. Orient α. Then at least one of the following cases occurs (they are not mutually exclusive).

- **Case 1:** Suppose that there exists a pair of intersections between α and β that are consecutive on β such that the orientations of α along each intersection are in the same direction, as in the left of Figure 53.

  If α starts at the top intersection (end 2) and returns to the top intersection (end 1) before reaching the bottom intersection (end 3), it will never reach the bottom intersection, since α is the image of S^1. So, it must go to the bottom intersection after the top one ($2 \to 3 \to 4 \to 1$). An example of this is in the left of Figure 54.

  Let α' be a curve that starts to the right of α at the top intersection (end 2), follows α to the right until reaching just before the bottom intersection (end 3), and then crosses over α and β vertically to close the path. Similarly, let α'' be a curve that starts to the right of α after the bottom intersection (end 4), follows

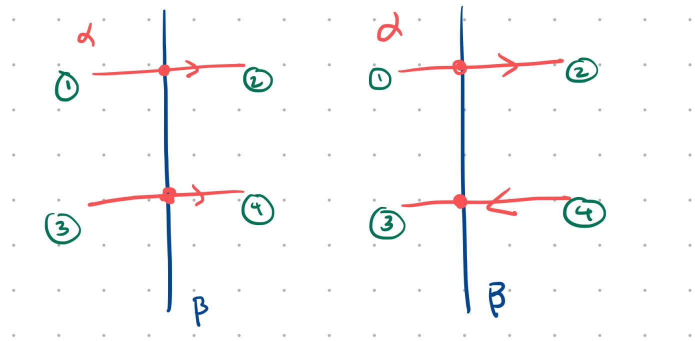

FIGURE 53. Possible cases for orientations of α around consecutive intersections on β.

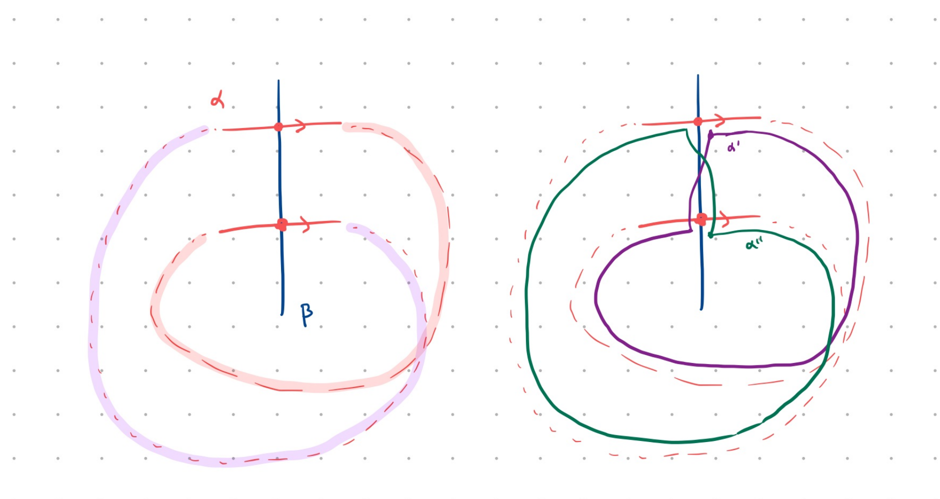

FIGURE 54. Paths α' and α''.

α to the right until reaching just before the top intersection (end 1), and then crosses over β and then α to close the path vertically. See the right of Figure 54.

Notice that α' and α'' basically divide α into two halves. Every intersection between α and β that doesn't occur in the vertical segment leads to either an intersection of β with α', or an intersection of β with α''. With the middle segment, there are two intersections between α and β and two intersections between $\alpha' \cup \alpha''$ and β. Hence, there is no increase in the number of intersections with β upon going from α to α' and α''. See Figure 55.

To rephrase,

$$i(\alpha', \beta) + i(\alpha'', \beta) \leqslant i(\alpha, \beta)$$

Here, we have an inequality (rather than an equality) since the no-bigon condition holds for α, β by assumption, but might not hold for β with α' or α''.

In addition, we've constructed α', α'' so that

$$i(\alpha, \alpha') = 1 = i(\alpha, \alpha'').$$

Note that this also implies that α', α'' are essential; if not, they wouldn't have intersection numbers 1 with something else on the closed surface.

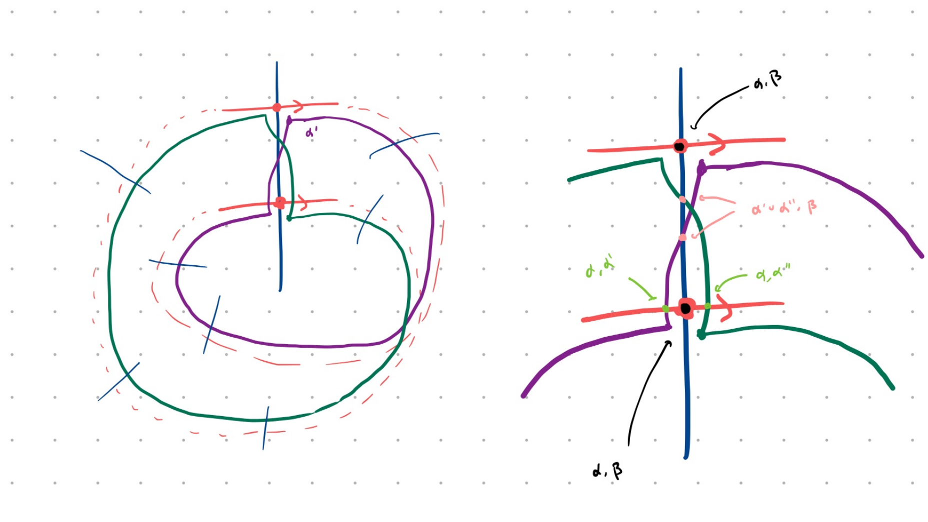

FIGURE 55. Outside of the middle segment, intersections between α and β occur as an intersection between β with either α' or α''. The right is a zoom-in on intersections in the middle segment.

Now, from the above inequality, we have WLOG that $i(\alpha', \beta) \leqslant \frac{i(\alpha,\beta)}{2}$, and so our claim holds for α', β by the induction hypothesis. By the triangle inequality,

$$dist(\alpha, \beta) \leqslant dist(\alpha, \alpha') + dist(\alpha', \beta).$$

Also, by the base case, $dist(\alpha, \alpha') = 2$ since $i(\alpha, \alpha') = 1$. Combining all this with the induction hypothesis, we see that

$$\begin{aligned} dist(\alpha, \beta) &\leqslant 2 + (2\log_2 i(\alpha', \beta) + 2) \\ &\leqslant 2 + \left(2 + \log_2 \frac{i(\alpha, \beta)}{2} + 2\right) \\ &= 2 + 2\log_2 i(\alpha, \beta), \end{aligned}$$

as desired.

- **Case 2:** Now suppose that there exists a pair of intersections between α and β that are consecutive on β such that the orientations of α along each intersection are in opposite directions.

  Again, α must go from the top intersection (end 2) to the bottom intersection (end 4) before returning to the top (end 1), and vice versa.

  As in the previous case, let α' be a curve starting to the right of α after the top intersection, follows to the right of α until right before the bottom intersection, and then travels up along β to close the path. Similarly define α''. See Figure 56.

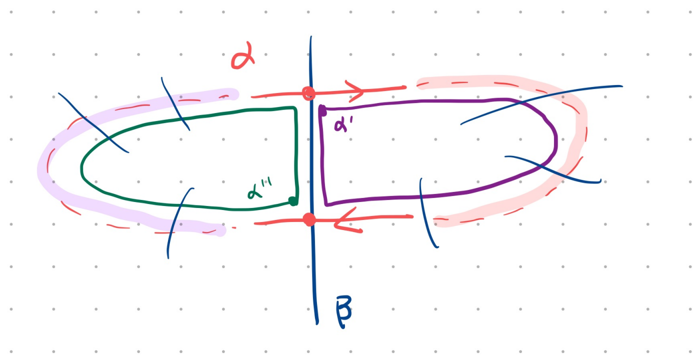

FIGURE 56. Paths α' and α'', along with their intersections with β.

By construction, we have that

$$i(\alpha, \alpha') = 0 = i(\alpha, \alpha''),$$

as α does not intersect either α' or α''; see Figure 56. In addition, we see that

$$i(\alpha', \beta) + i(\alpha'', \beta) + 2 \leqslant i(\alpha, \beta).$$

This inequality is shown in a way similar to the previous case: every intersection of α with β outside the middle segment occurs exactly once as an intersection between either α' or α'' with β, and $\alpha' \cup \alpha''$ does not intersect β along the middle segment, whereas α, β intersect twice there.

We have that α', α'' are simple, as they couldn't self-intersect on the portions where they follow α, nor on the portions where they follow β, as α and β are simple and β doesn't intersect α along the middle segment.

We also have to show that α', α'' are essential. This is true as they are bicorns of α, β (see Definition 11.1 and Remark 11.2 below).

Now, the desired inequality follows as in the first case since we have the weaker inequalities $dist(\alpha, \alpha') = 1 = dist(\alpha, \alpha'') \leqslant 2$ and $i(\alpha', \beta) + i(\alpha'', \beta) \leqslant i(\alpha, \beta)$.

□

Definition 11.1. Given scc's α, β with no bigons, a **bicorn** curve γ is a simple closed curve that consists of an arc of α and an arc of β.

Remark 11.2. Since α, β have no bigons, a bicorn γ cannot bound a disk. Thus, γ is essential.

Next, we'll show that the curve complex is hyperbolic by making paths from scc's α to β of bicorn curves of that are geodesic curves. To do so, we'll interpolate from one curve to another via bicorns: we'll first follow all of α, then start interpolating by following more and more of β and less and less of α, until we get to following all of just β.

To get this interpolation, we'll prove the following lemma next class.

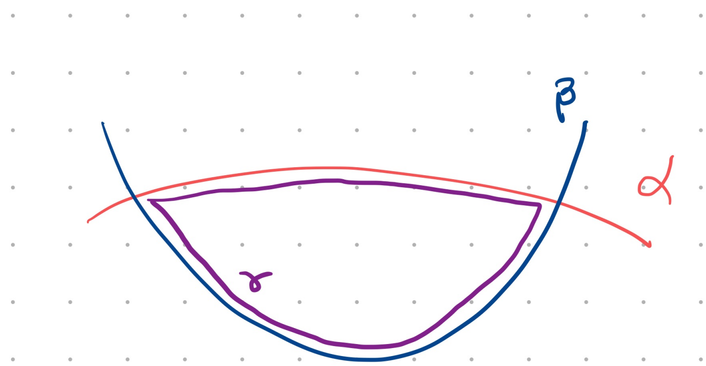

FIGURE 57. Bicorn γ of α and β. This is only representative, since α and β don't have any bigons.

Lemma 11.3. *Let γ be a bicorn of α and β. Then if $i(\gamma, \beta) > 2$, then there exists a bicorn γ' of α, β whose β segment strictly contains the β segment of γ, and $i(\gamma, \gamma') \leqslant 1$.*

Note that the path of bicorns we're constructing could be very inefficient, but this does not affect our application of the Geodesic Guessing Lemma (Proposition 4.2), since its conditions don't refer to the parametrization of the path, only to its image.

Optional Exercise 21. Suppose $d_{\mathcal{CS}}(\alpha, \beta) = D$. Prove that there is a connected degree D cover S' of S where α and β have disjoint lifts (meaning there are disjoint scc α', β' on S' such that α' maps to (a power of) α, and similarly for β').

Optional Exercise 22. Use the last exercise to show that if $\mathcal{CS}$ had finite diameter, there would be a finite cover $S' \to S$ such that any two scc on S have disjoint lifts on S'.

12. Improved geodesic guessing (02/02, CK, JH)

We will start proving a generalization of the Geodesic Guessing Lemma, namely one that removes the need for the second coherence condition for geodesic guesses. References are [MS13, Theorem 3.15] and [Bow14a, Proposition 3.1].

Proposition 12.1. *For any given h, if G is a connected graph such that $\forall x, y \in V(G)$, we have a chosen connected subgraphs $\eta(x, y)$ so that*

- $d(x, y) \leqslant 1 \implies \operatorname{diam}(\eta(x, y)) \leqslant h$ *and*
- $\eta(x, y) \subset N_h(\eta(x, z) \cup \eta(z, y))\ \forall x, y, z \in V(G)$,

then G is hyperbolic. Moreover, there is a $D = D(h)$ so that for any geodesic γ joining x and y, $d_{Haus}(\gamma, \eta(x, y)) \leqslant D$.

Remark 12.2. We state this for graphs, but any geodesic metric space X is quasi-isometric to a graph described as follows: Let $V(G) = X$ as a set and join any two vertices of distance $\leqslant 1$ in X by an edge. This is a connected graph by our assumption on X.

If we have path-connected geodesic guesses in X satisfying these conditions, then one can check that the subgraphs generated by their images in $V(G)$ satisfy the same conditions with a slightly different h.

Proof. The breakdown of the proof is similar to the original version: we get a logarithmic bound on distances of points on geodesic guesses to paths, then show that distances from guesses to geodesics are bounded and vice versa. We first prove the logarithmic bound. We only prove it for paths that are concatenations of geodesics since length is troublesome to define otherwise and more generality is not needed in the proof.

Lemma 12.3. *Given a concatenation of geodesics $\gamma : [a, b] \to G$, define its length to be $l(\gamma) = b - a$. Then for any point $p \in \eta(x, y)$,*

$$d(p, \gamma) \leqslant h \log_2(l(\gamma)) + 2h = h \log_2(b - a) + 2h$$

if $b - a \geqslant 1$.

Proof. This is identical to the proof of a similar lemma used in the original Geodesic Guessing Lemma. Consider the midpoint $m = \gamma(\frac{b+a}{2})$. By the second condition

$$\eta(x, y) \subset N_h(\eta(x, m) \cup \eta(m, y)).$$

So, p is within h of $\eta(x, m)$ or $\eta(m, y)$. WLOG, p is within h of $\eta(x, m)$. Pick $p' \in \eta(x, m)$ so that $d(p, p') \leqslant h$. Let γ' be the part of γ from x to m. See Figure 58.

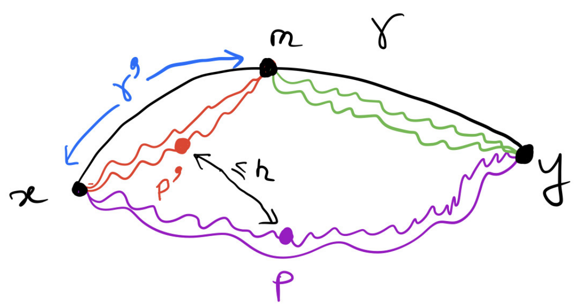

FIGURE 58. The squiggly sets are the geodesic guesses and the black curve is the geodesic γ.

Now if $b - a \leqslant 2$, then $l(\gamma') = \frac{b-a}{2} \leqslant 1$. So, $d(x, m) \leqslant 1$. By the first condition, $\operatorname{diam}(\eta(x, m)) \leqslant h$ and so

$$d(p', \gamma) \leqslant d(p', \gamma') \leqslant \operatorname{diam}(\eta(x, m)) \leqslant h.$$

Hence, $d(p, \gamma) \leqslant d(p, p') + d(p', \gamma) \leqslant 2h \leqslant h \log_2(b - a) + 2h$.

Assume that this holds for paths of length $\leqslant n$. Then if $b-a \leqslant n+1$, $l(\gamma') = \frac{b-a}{2} \leqslant n$. By the induction hypothesis, $d(p', \gamma) \leqslant d(p', \gamma') \leqslant h \log_2(\frac{b-a}{2}) + 2h$. Then,

$$d(p, \gamma) \leqslant d(p, p') + d(p', \gamma') \leqslant h + h \log_2\left(\frac{b-a}{2}\right) + 2h = h \log_2(b - a) + 2h.$$

By induction, the lemma holds. $\square$

Note that for $b-a \geqslant 2, h\log_2(b-a)+2h \leqslant 3h\log_2(b-a)$. The rest of the lemma is proved like last time. We first show that any point on a geodesic guess is at most $c = c(h)$ from a corresponding geodesic and then vice versa. We will continue next class. □

13. Improved geodesic guessing (02/04, SK, KS)

Proving the key lemma. In this section, we finish proving Proposition 12.1, which is a version of the geodesic guessing lemma where the guesses are connected subgraphs.

Proposition 12.1. *For any given h, if G is a connected graph such that $\forall x, y \in V(G)$, we have a chosen connected subgraphs $\eta(x,y)$ so that*

- *$d(x,y) \leqslant 1 \implies \operatorname{diam}(\eta(x,y)) \leqslant h$ and*
- *$\eta(x,y) \subset N_h(\eta(x,z) \cup \eta(z,y))\ \forall x,y,z \in V(G)$,*

then G is hyperbolic. Moreover, there is a $D = D(h)$ so that for any geodesic γ joining x and y, $d_{Haus}(\gamma, \eta(x,y)) \leqslant D$.

Remark 13.1. Note that this version of the geodesic guessing lemma differs from the previous version in two ways: instead of $\eta(x,y)$ being coarsely continuous paths, they are now connected subgraphs. This lets us drop the requirement that $\eta(x',y')$ be bounded distance away from the subpath of $\eta(x,y)$ between x' and y'.

The proof is mostly identical to the proof of Proposition 4.2; the only part of the proof that differs is the proof of the following lemma.

Lemma 13.2. *There exists a constant $c > 0$ such that for any geodesic γ from x to y, and any $p \in \eta(x,y)$, $d(p,\gamma) \leqslant c$.*

Proof. In Lemma 12.3, we proved the following bound on $d(p,\gamma)$, where $\ell(\gamma)$ is the length of the geodesic γ from x to y, and $\ell(\gamma) \geqslant 1$.

$$d(p,\gamma) \leqslant h\log_2(\ell(\gamma)) + 2h$$

Let $E_\gamma = \max_{p\in\eta(x,y)} d(p,\gamma)$, i.e. the furthest $\eta(x,y)$ gets from γ: it will suffice to get an upper bound on E_γ, independent of γ.

Without loss of generality, we can replace γ with the smallest $\tilde{\gamma}$ such that $E_{\tilde{\gamma}} \geqslant E_\gamma$. This lets us assume that if δ is a geodesic shorter than γ, then $E_\delta \leqslant E_\gamma$. Consider now points q, x', and y' on γ: q is a point on γ that is distance E_γ from p, and x' and y' are points in $\eta(x,y)$ which are distance at least $2h+1+2E_\gamma$ from q and closer to x and y respectively than y and x (see Figure 59).

Let z be any point on the open geodesic segment from x to x' or y to y'. In case $x = x'$ and $y = y'$, we have the logarithmic bound on E_γ from the previous section, which gives us the following inequality.

$$E_\gamma \leqslant h\log(4h + 2 + 4E_\gamma) + 2h$$

This provides an upper bound for E_γ.

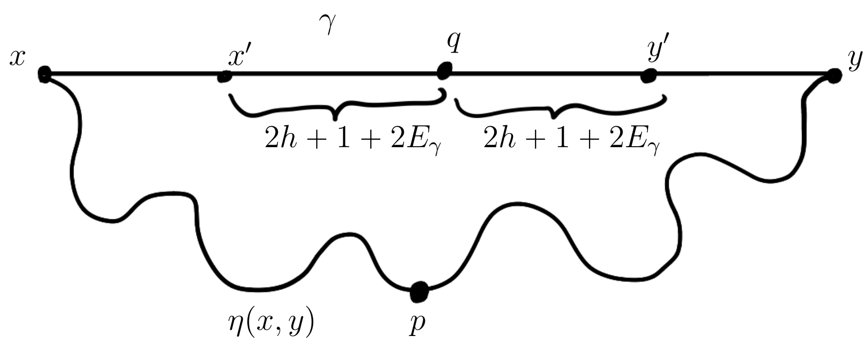

FIGURE 59. The guess $\eta(x, y)$ and the geodesic γ.

When at least one of $x \neq x'$ and $y \neq y'$ holds, we have the following lower bound on $d(z, p)$ from the triangle inequality.

$$\begin{aligned} d(z,p) &\geqslant d(z,q) - d(p,q) \\ &\geqslant 2h + 1 + 2E_\gamma - E_\gamma \\ &= 2h + 1 + E_\gamma \end{aligned}$$

Without loss of generality, we can assume that $x \neq x'$ and z lies on the geodesic γ' from x to x'. We must have that $E_{\gamma'} \leqslant E_\gamma$ since $\ell(\gamma') < \ell(\gamma)$. In particular, this means that for any $p' \in \eta(x, x')$, $d(p', \gamma') \leqslant E_\gamma$. The triangle inequality then lets us conclude the following.

$$d(p, \eta(x, x')) \geqslant 2h + 1$$

Consider now the $2h$-slim quadrilateral formed by $\eta(x, x')$, $\eta(x', y')$, $\eta(y', y)$, and $\eta(x, y)$. Since both $\eta(x, x')$ and $\eta(y, y')$ are distance at least $2h + 1$ from $\eta(x, y)$ by the previous argument, there must exist some point q' on $\eta(x', y')$ such that $d(p, q') \leqslant 2h$.

We also have that $d(q', \gamma'') \leqslant k \log(4h + 2 + 4E_\gamma)$, where γ'' is the geodesic segment from x' to y'. Finally, since q is contained in γ'', $d(p, \gamma) = d(p, \gamma'')$. See Figure 60 for all new points and labels added to the picture.

Combining all these inequalities, we get the following chain of inequalities.

$$\begin{aligned} E_\gamma &= d(p, \gamma) \\ &= d(p, \gamma'') \\ &\leqslant d(p, q') + d(q', \gamma'') \\ &\leqslant 4h + h \log(4h + 2 + 2E_\gamma) \end{aligned}$$

This inequality shows that E_γ is bounded uniformly, independently of γ, proving the lemma. □

The rest of the proof of the Proposition 12.1 is identical to the proof of Proposition 4.2, so we leave the verification to the reader.

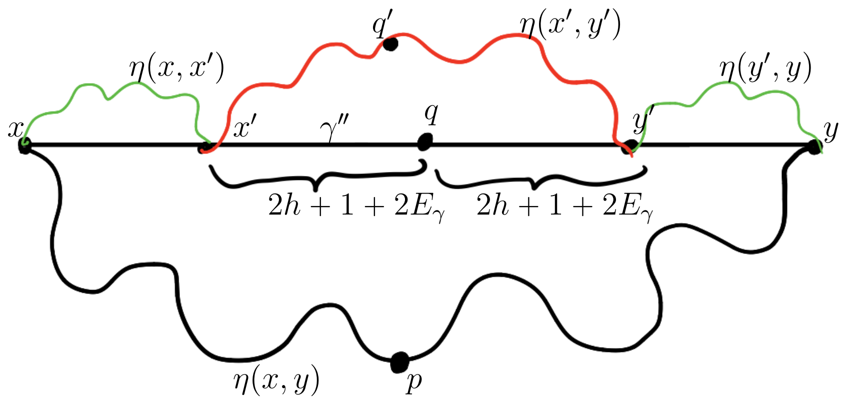

FIGURE 60. The geodesic quadrilateral and the point q'.

Constructing the family of geodesic guesses for $\mathcal{C}(S)$. To use Proposition 12.1 to prove that $\mathcal{C}(S)$ is hyperbolic, we first need to construct connected paths $\eta(x, y)$ between points x and y in $\mathcal{C}(S)$. To do so, we will need to prove the following key lemma.

Lemma 13.3 (Constructing coarsely connected paths)**.** *Let γ be a bicorn between α and β. If $i(\gamma, \beta) > 0$, then there exists another bicorn γ' whose β segment strictly contains the β segment of γ and $i(\gamma, \gamma') \leqslant 1$.*

Proof. Let α' and β' be the segments of α and β that form the bicorn γ. We extend the segment β' along β starting at one of the corners until the curve intersects α again in the interior of α' then turn and follow α until we reach the other corner of the bicorn. This is guaranteed to happen since $i(\gamma, \beta) > 0$. There are 2 cases to consider depending on whether the segment turns left or right to reach the other corner. We illustrate the 2 cases in Figure 61.

It is clear from the picture that the resulting curve γ' in both the cases is a bicorn between α and β and contains a larger segment of β.[2] This proves the result. □

For any α and β in the curve complex $\mathcal{C}(S)$, we define the geodesic guess $\eta(\alpha, \beta)$ to be $\{\alpha, \beta\} \cup \{$bicorns between α and β and the interpolating length 2 paths$\}$. By the above lemma, this is a connected subgraph of the curve complex.

14. HYPERBOLICITY OF THE CURVE COMPLEX (02/07, SC, MM)

In today's class we prove the following theorem. A reference is [PS17, Theorem 2.1].

Theorem 14.1. *$\mathcal{C}(S)$ is hyperbolic.*

[2]Note that the extension of β' may intersect α many times before it first intersects α'. We only stop extending β' the first time it intersects α'.

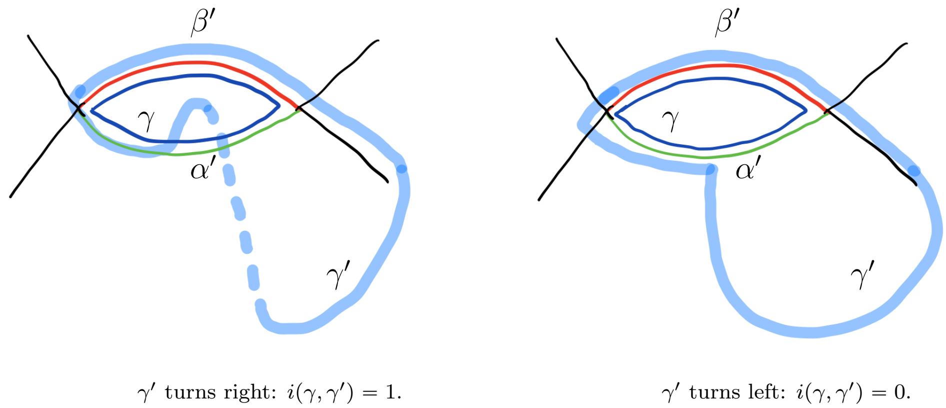

FIGURE 61. The 2 cases to consider for bicorn interpolation.

This will follow once we show that the geodesic guesses $\eta(a, b)$ defined earlier satisfy the conditions required by the upgraded geodesic guessing lemma, Proposition 12.1. Recall that each geodesic guess $\eta(a, b)$ consists of a, b and all bicorns between a and b along with all the interpolating paths between these curves of length at most two.

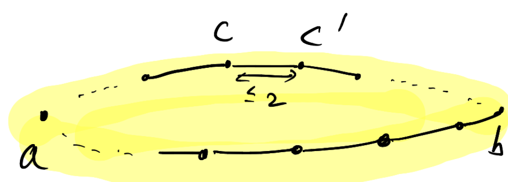

FIGURE 62. The geodesic guess $\eta(a, b)$. c and c' are bicorns between a and b. There may be multiple paths from a to b each given by series of bicorns between a and b.

Proof. We just need to check the two conditions in the upgraded geodesic guessing lemma, Proposition 12.1. To check the first condition, we consider curves a and b with $d(a, b) \leqslant 1$. This means either $a = b$, or a and b are disjoint. In the former case, $\eta(a, b)$ consists of just a point, and in the latter case, $\eta(a, b)$ consists of a, b and the edge between a and b (since disjoint curves cannot form bicorns). In either case, $\eta(a, b)$ has diameter at most 1.

Next, we show that for curves a, b and d in $\mathcal{C}(S)$, $\eta(a, b) \subset N_3\big(\eta(a, d) \cup \eta(d, b)\big)$. This will follow from the following lemma,

Lemma 14.2. *Let a, b and d be curves on S with no bigons between any pair of them, and c be a bicorn of a and b. Then, there exists a curve $c' \in \eta(a,d) \cup \eta(d,b)$ that intersects c at most twice.*

To see that this lemma implies $\eta(a,b) \subset N_3\big(\eta(a,d) \cup \eta(d,b)\big)$, we observe that if c' intersects c at most twice, $d(c,c') \leqslant 2$. This is because if $d(c,c') > 2$, then c and c' fill, and we saw on January 28 that filling curves intersect at least $2g-1 \geqslant 3$ times. Then, if $\tilde{c} \in \eta(a,b)$, it is either a bicorn between a and b or it lies on an interpolating path of length two between bicorns of a and b. In either case, we have a bicorn c (where c could possibly be $\tilde{c}$) at distance at most 1 from $\tilde{c}$ to which we apply the lemma (Figure 63).

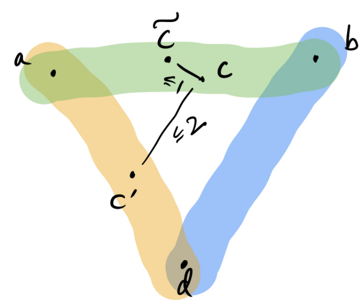

FIGURE 63. Showing that $\eta(a,b) \subset N_3\big(\eta(a,d) \cup \eta(d,b)\big)$. The shaded region between a and b is $\eta(a,b)$.

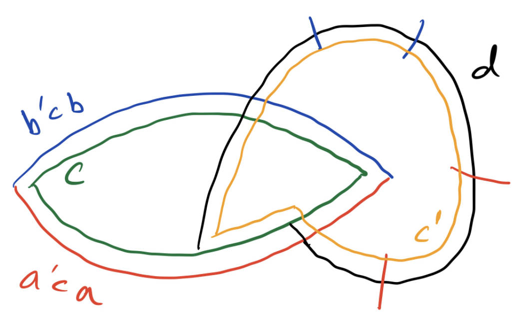

FIGURE 64. The bicorn c' when d' intersects b' once and the two ends of d' are on different sides of a'.

Proof of lemma. If $i(c,d) \leqslant 2$, we let $c' = d$. Else, let a' and b' be the segments from a and b respectively in the bicorn c as shown in Figure 64. Let d' be a minimal segment of d with both end points on a' or both endpoints on b'. WLOG, assume that d' has both end points on a'. Then, d' does not intersect a' except at its end points and it intersects b' at most once. For example, in Figure 64, d' interects b' once, and in Figure 65 d' does not intersect b'.

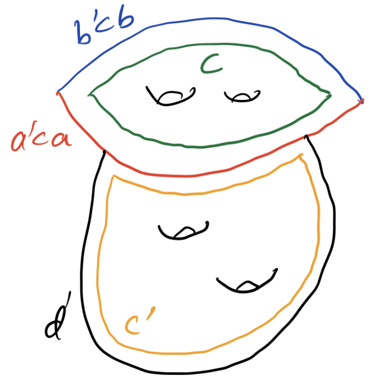

FIGURE 65. The bicorn c' when d' does not intersect b' and the two ends of d' are on the same side of a'.

We define c' to be the bicorn of d and a formed by closing up d' along a' as shown in Figure 64. Since d' does not intersect a' except at its end points, d' clearly does not interesect the segment of a' that it forms the bigon c' with. This tells us that the curve c' is simple.

The intersections between c' and c can come either from intersections between c and d' or intersections between c and the segment of c' joining the two ends of d'. There can be at most one intersection between c and d', which is a possible intersection between d' and b'. The segment of c' joining the ends of d' along a' intersects c at most once when the two end points of d' are on different sides of a' (Figure 64). When the two end points of d' are on the same side of a', then the segment of c' joining the ends of d' does not intersect c (Figure 65). □

□

Later on, we will show that $\mathcal{C}(S)$ is quasi-isometric to an electrification of $\mathcal{T}_g$.

Remark 14.3. A bicorn is somewhat related to a "closest return" of an IET. For example, suppose α is the core curve of a horizontal cylinder on an Abelian differential, and β is a core curve of a vertical cylinder. Think of α as length 1 and β as huge. α is like the interval (the transversal) on which the IET is defined. If you start at a point of intersection, you can follow β upwards, and it will return to α many times. You can't always close up such a vertical segment to get a simple curve, but you can if the return is as close as the segment has gotten so far to the start point. Such a closing up is a bicorn.

This maybe could provide a bit of intuitions for bicorn paths having something to do with Teichmüller geodesics, and hence being reasonable guesses.

15. Hyperbolic surfaces (02/09, KS, SC)

In this lecture we discuss hyperbolic surfaces and their properties. Hyperbolic surfaces are surfaces with a Riemannian metric of constant curvature equal to -1. They are locally isometric to $\mathbb{H} = \{x + iy : y > 0\}$ with the metric $\frac{dx^2+dy^2}{y^2}$. Hyperbolic surfaces have a lot of nice geometric properties which we state as key facts below.

- **Key fact 1.** For any $l_1, l_2, l_3 > 0$ there exists a unique hyperbolic metric on $S_{0,3}$ (a.k.a. "pants", see Figure 66) such that boundary circles are geodesics of length l_1, l_2, l_3.

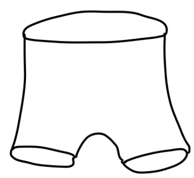

FIGURE 66. "Pants" $S_{0,3}$, i.e. sphere with three boundary components

"Pants" can be glued to each other to construct closed geodesic surfaces of different genus (see Figure 67). Note that we may have twists.

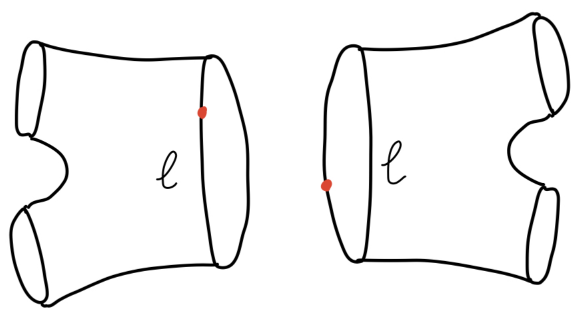

FIGURE 67. "Pants" $S_{0,3}$, i.e., sphere with three boundary components, two boundaries with the same length l are glued together

- **Key fact 2.** There do not exist hyperbolic geodesic bigons. More generally, that's a feature of spaces of non-positive curvature (see Figure 68).
- **Key fact 3.** Every closed curve is homotopic to a unique geodesic.

Let us describe how "pants" can look like. By the Gauss-Bonnet formula they always have an area of 2π. For example, if $l_i, i = 1, 2, 3$ are small, boundaries have to be far from each other in order to keep area equal to 2π. In the limit $l_i \to 0, i = 1, 2, 3$ we get a sphere with three cusps (see Figure 69). For other possible cases, see Figure 70.

Now let us see how to get a hyperbolic surface by gluing "pants" (see Figure 71) with an example of a genus 2 surface. Given a hyperbolic surface of genus g, we can get its topological $2g - 2$ "pants" decomposition, then take $3g - 3$ geodesic representatives as "pants" boundaries. They are disjoint so we don't have any bigons.

Now let us discuss properties of simple closed geodesics on hyperbolic surfaces.

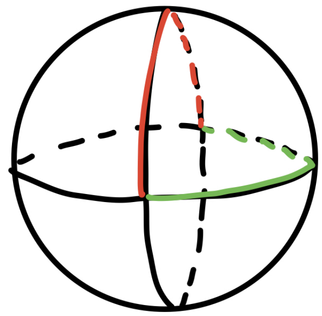

FIGURE 68. Geodesic bigon on a sphere.

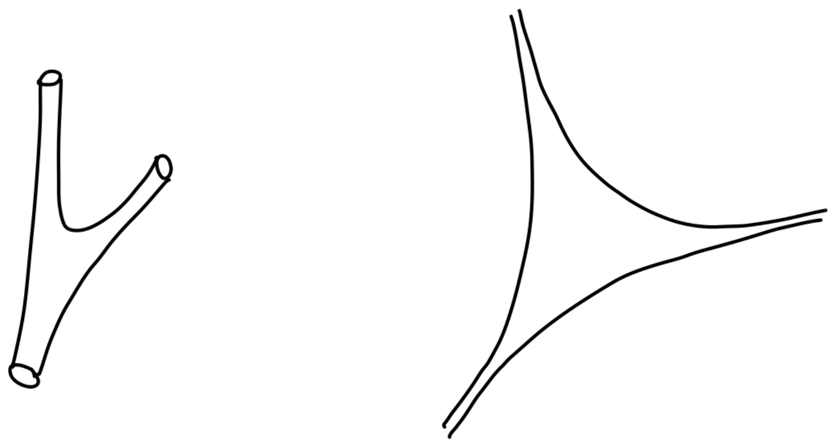

FIGURE 69. On the left: "pants" with three small boundaries. On the right: limit as $l_i \to 0, i = 1, 2, 3$, i.e., sphere with three cusps.

- **"Bad" news.** For any $c > 0$, there exists a hyperbolic surface with genus large depending on c so that every simple closed geodesic has length greater than c.
- **"Good" news (part 1).** For any g, there exists $C_g > 0$ such that all hyperbolic surfaces of genus g have a simple closed geodesic of length at most C_g. One can show $C_g \sim \log g$.
- **"Good" news (part 2).** For any g there exists $B_g > 0$ such that all hyperbolic surfaces of genus g have "pants" decomposition with all geodesics of length not greater than B_g. Notice that B_g is a lot bigger than C_g.
- **Final remark.** There exists $\delta > 0$ such that for all hyperbolic surfaces of any genus, any two geodesics of length less than δ are disjoint and simple.

16. Systole and Teichmüller space (02/11, JH, MM)

We cover one more topic on hyperbolic surfaces before moving on to some basic Teichmüller theory.

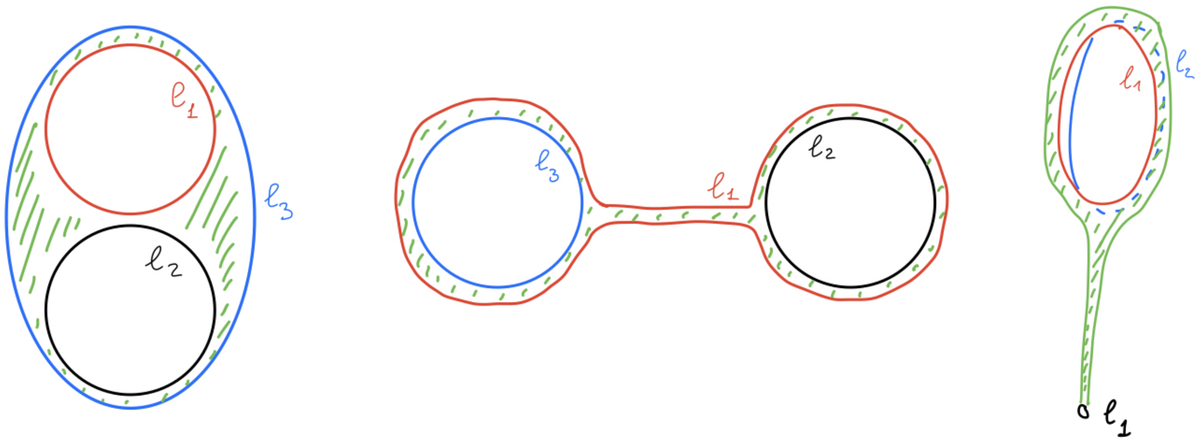

Figure 70. On the left: $l_1, l_2, l_3 \gg 1$, in the middle: $l_1 \gg l_2, l_2$, on the right: $l_1 \ll l_2, l_3$.

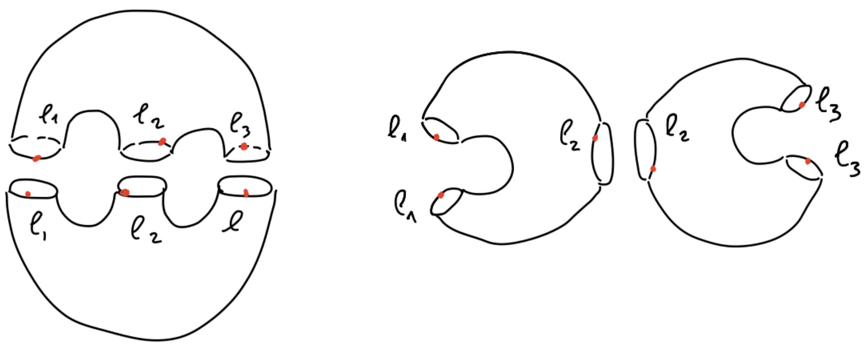

Figure 71. "Pants" decomposition for genus 2 surface.

Definition 16.1. If X is a (compact) hyperbolic surface, a *systole* of X is a closed geodesic on X of minimal length (excluding single points).

A systole always exists because there are countably many isotopy classes of closed curves on X, and the lengths of the isotopy classes form a discrete set, which must have a minimal element.

Remark 16.2. Every closed geodesic on X is essential, using the same proof that there are no hyperbolic bigons. This is worth remarking because it is specific to hyperbolic surfaces, since they exclude, for example, the situation in Figure 72.

We will develop some basic facts about systoles.

Lemma 16.3. *A systole must be simple.*

Some authors include being simple as part of the definition of being a systole.

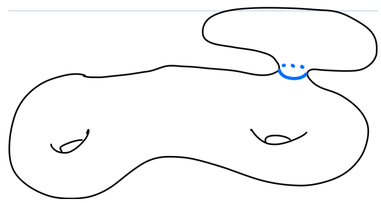

Figure 72. A minimal-length, closed, non-essential geodesic.

Proof. If a systole γ were not simple, take a closed arc α along γ. Then α has no bigons with γ, hence is essential. The geodesic in the homotopy class of α has length at most $l(\alpha) < l(\gamma)$. See Figure 73. □

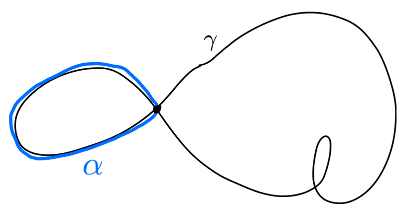

Figure 73. A hypothetical closed arc of a systole.

It would be preferable if systoles were unique, and this is what we expect to happen for "most" spaces, but it is not true in general.

Example 16.4. We can construct a space with two systoles of length ε using two pairs of pants, as shown in Figure 74

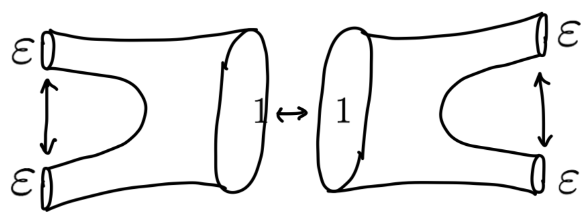

Figure 74. A hyperbolic surface with two systoles.

Even worse than failing uniqueness, systoles don't even have to be disjoint! However, the situation is not too bad.

Lemma 16.5. *If γ_1 and γ_2 are systoles of X, then $i(\gamma_1, \gamma_2) \leqslant 1$.*

This means that in the curve complex, systoles have distance at most 2 from each other, so at least we have coarse uniqueness.

Proof. If $i(\gamma_1, \gamma_2) \geqslant 2$, pick two intersections which divide γ_i into γ_i' and γ_i''. Without loss of generality, assume $l(\gamma_i') \leqslant \frac{1}{2} l(\gamma_i)$. Define $\alpha = \gamma_1' \cup \gamma_2'$, so

$$l(\alpha) \leqslant l(\gamma_1') + l(\gamma_2') \leqslant l(\gamma_1) = l(\gamma_2).$$

See Figure 75. However, since γ_1' and γ_2' are distinct geodesic arcs, they must intersect transversally, so α will have corners where they meet. This means α cannot be a geodesic, so the geodesic representative of α has length strictly less than $l(\alpha) < l(\gamma_1)$. This contradicts minimality of $l(\gamma_1)$, so $i(\gamma_1, \gamma_2) \leqslant 1$. □

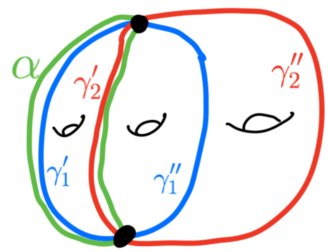

FIGURE 75. Finding a smaller geodesic if two systoles intersect twice.

Example 16.6. We can't do better than bounding the intersection number by 1, since we could have the situation shown in Figure 76. In that case, expanding the two red curves causes the blue curve to shrink, so there is some point where they all have the same length. There are some technical details to worry about with making sure that there are no shorter curves anywhere, but this construction gives the idea for how to obtain two intersecting systoles.

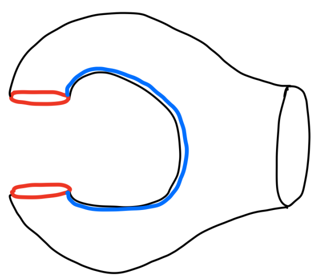

FIGURE 76. Systoles that intersect.

Example 16.7. By making surfaces with lots of symmetry, we can obtain arbitrarily many systoles on a surface, since the image of a systole under an isometry is still a systole. For example, making a genus-2 surface out of a regular octagon gives a surface with four intersecting systoles. See Figure 77.

Remark 16.8. From last time ("good news part 1"), we know that there is a constant $C_g \sim \log_2(g)$, depending only on the genus g, which is an upper bound on the length of a systole of surfaces of genus g.

We want to put all the information we have about hyperbolic surfaces together and look at different surfaces and metrics in a unified way. One way to do this would be to look at the space of all hyperbolic metrics on a fixed surface, but that space is infinite-dimensional and not nice to work with.

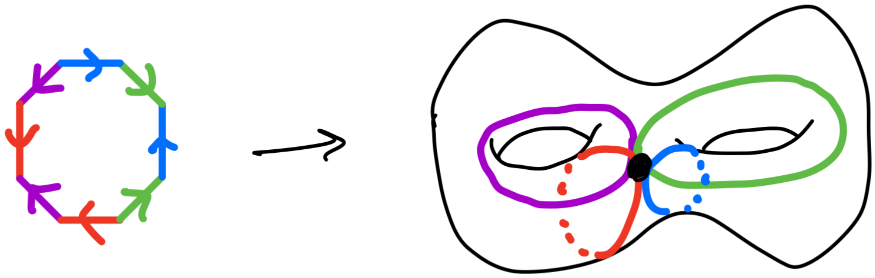

FIGURE 77. A genus-2 surface with 4 systoles.

Definition 16.9. Let S be an oriented surface of genus $g \geqslant 2$. Then the *Teichmüller space* of S, denoted by T_g, is

$$T_g = \frac{\left\{ \varphi\colon S \to X \;\middle|\; \begin{array}{c} X \text{ is an oriented hyperbolic surface;} \\ \varphi \text{ is an orientation-preserving} \\ \text{homeomorphism} \end{array} \right\}}{\sim},$$

where

$$\begin{gathered} \varphi_1\colon S \to X_1 \sim \varphi_2\colon S \to X_2 \\ \Updownarrow \\ \text{there is an orientation preserving isometry } I\colon X_1 \to X_2 \\ \text{such that } I \circ \varphi_1 \text{ is homotopic to } \varphi_2. \end{gathered}$$

We call S the *reference* or *marking* surface.

This is the formal definition of Teichmüller space, but no one really thinks of it this way. A better way to think of points in T_g is as a hyperbolic surface X with some "marking data," in the form of a homotopy class of orientation-preserving homeomorphisms $S \to X$. Since there are countably many such homotopy classes, the second part of the data can be thought of as discrete. For this reason, people (including us, going forward) often write $X \in T_g$, with S, φ, and the equivalence class being implicit.

Remark 16.10 (Key Point). The map $\varphi\colon S \to X$ induces a bijection between homotopy classes of closed curves on S and on X. Concretely, for any closed curve α on X, a marking $[\varphi\colon S \to X]$ determines a homotopy class of curves on X, all homotopic to $\varphi(\alpha)$, with a unique geodesic representative. So we can assume $\varphi(\alpha)$ to be geodesic. So, we can define a map

$$\begin{gathered} \ell_\alpha\colon T_g \to \mathbb{R}_{>0} \\ \ell_\alpha([\varphi\colon X \to S]) = l(\varphi(\alpha)), \end{gathered}$$

mapping $[\varphi\colon X \to S]$ to the length of the image of α.

We can start to see why S is useful to have around. It serves as a universal reference to compare points in T_g without having to designate any individual hyperbolic surface.

There is a natural topology on T_g in which each ℓ_α is continuous, and with this topology, T_g is homeomorphic to $\mathbb{R}^{6g-6}$.

Theorem 16.11 (Fenchel-Nielson Coordinates)**.** *T_g is homeomorphic to $\mathbb{R}_{>0}^{3g-3} \times \mathbb{R}^{3g-3}$.*

Some basic ideas for the proof are as follows: we fix a pants decomposition for S. This gives a pants decomposition of each $X \in T_g$ and induces $6g-6$ parameters.

- The first $3g-3$ are the length coordinates, which are the lengths of the cuffs of the pants. These are positive (hence we use $\mathbb{R}_{>0}$ for these coordinates) and relatively easy to understand.
- The last $3g-3$ coordinates are the twist coordinates, which are more subtle. They tell use which points to identify while gluing each of the circles; i.e., how much to "twist" each of the legs while gluing the pants together. Somewhat surprisingly, going "around the circle" once in these coordinates does not bring us back to the same result, which is why these coordinates use $\mathbb{R}$ and not $(0, 2\pi)$.

We will go into more detail about the twist coordinates next time.

17. Mapping classes and Fenchel-Nielsen coordinates (02/14, CK, TY)

Recall that in the last class, we defined the Teichmuller space T_g of a surface of genus g. Formally, $T_g = \{\phi : S \to X\}/\sim$ under the equivalence relation from last class, with S being our chosen topological surface of genus g, X any hyperbolic surface of genus g and ϕ a homeomorphism. However, we will often suppress the marking ϕ and talk about hyperbolic surfaces $X \in T_g$.

17.1. **Mapping Class Groups.** Consider a pants decomposition of the topological surface S as in Figure 78. Pick the cuff curve γ and consider the red curve α in the figure. Since α and γ don't form any bigons (as they intersect only once), $i(\alpha, \gamma) = 1$. Cut S along γ, twist the top copy to the right by an angle of ϵ in the parametrization, and reglue the curves. See Figure 78.

This gives a new surface S_ϵ with the same Euler characteristic as S and is thus abstractly homeomorphic to S by the classification of surfaces. There is a natural homeomorphism between the cut surfaces before and after twisting. But there is no natural continuous map we can get after the regluing, since we cannot define a continuous extension from $S\backslash\gamma$ to S (there are two choices for the image of each point on γ). However, if you continue twisting and twist by a full circle, there is a natural continuous map that can be defined this way. It is easy to check that this map is a homeomorphism, once one formally defines the twist in a tubular neighborhood of γ, as we will do later.

Intuitively, it make sense that since α twists one extra time around γ, this map is not the identity. We will see that such homeomorphisms can change the marking of a given hyperbolic surface to give a new point in Teichmuller space, assuming that they are not homotopic to the identity. This motivates the following definition.

Definition 17.1. The **mapping class group** $MCG(S)$ of a topological surface S is the group of orientation preserving homeomorphisms quotiented by the normal subgroup of homeomorphisms homotopic to the identity. That is,

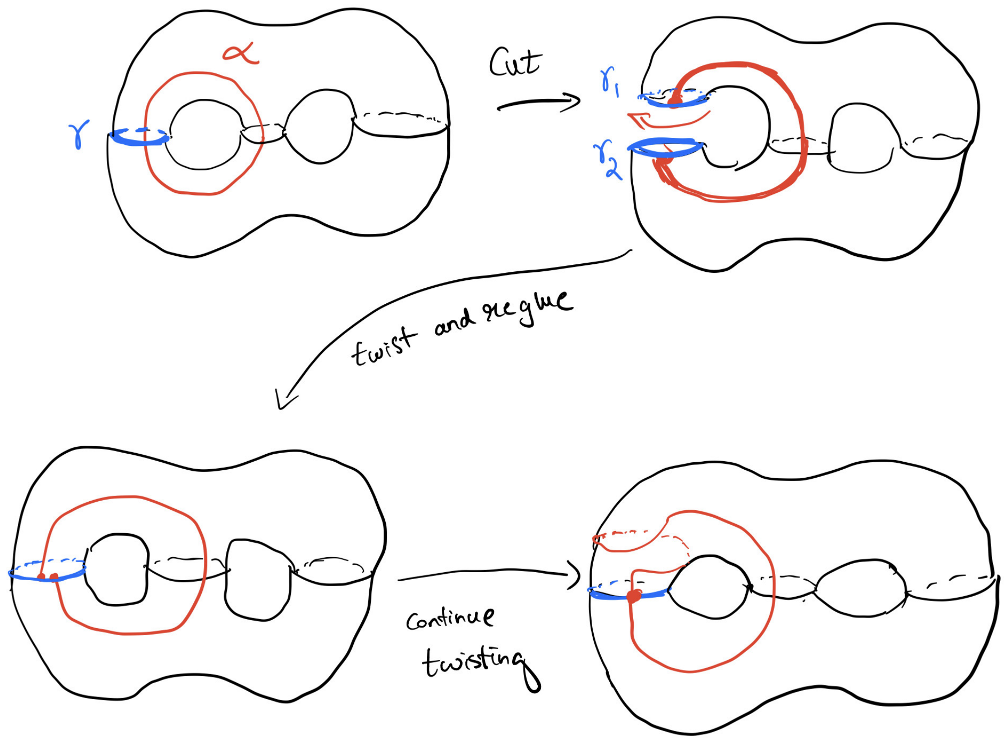

FIGURE 78. Cutting along γ, twisting and regluing produces an abstractly homeomorphic surface. Twisting by a full circle gives us a natural homeomorphism to the new surface.

$$MCG(S) := \frac{\{\phi : S \xrightarrow{\cong \text{ o.p.}} S\}}{\{\phi \sim id\}}$$

Remark 17.2. $MCG(S)$ is countable. In particular, if you extend the group to include all homeomorphisms (those preserving or reversing orientation) modulo homotopy to get $MCG^{\pm}(S)$, then

$$MCG(S) \subset MCG^{\pm}(S) \cong Out(\pi_1(S)) = \frac{Aut(\pi_1(S))}{Inn(\pi_1(S)}$$

Definition 17.3. We define the (left) action of the mapping class group on Teichmuller space by changing the marking as hinted above. That is, if $[f] \in MCG(S)$ (in the sense that f is an orientation preserving homeomorphism) and $[\phi : S \to X] \in T_g$, then

$$[f] \cdot [\phi : S \to X] := [\phi \circ f^{-1} : S \to X]$$

In terms of a commutative diagram,

$$\begin{tikzcd} S \arrow[r, "f^{-1}"] \arrow[dr, "\phi\circ f^{-1}"'] & S \arrow[d, "\phi"] \\ & X \end{tikzcd}$$

Using the tubular neighbourhood theorem, for any s.c.c. γ we can obtain a neighborhood homeomorphic to $S^1 \times [0,1]$. See Figure 79. On this, define a homeomorphism $(\theta, y) \mapsto (\theta + 2\pi y, y)$. Notice that it restricts to the identity on the boundary $S^1 \times \{0,1\}$, so it extends to a homeomorphism f_γ of the surface S.

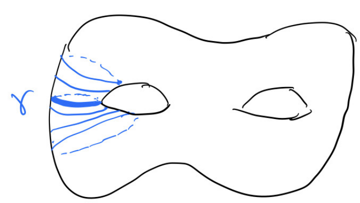

FIGURE 79. The cylindrical tubular neighborhood around an s.c.c γ

Definition 17.4. $[f_\gamma] \in MCG(S)$ as defined above is called the **Dehn twist** about the curve γ. $D_\gamma := [f_\gamma]$

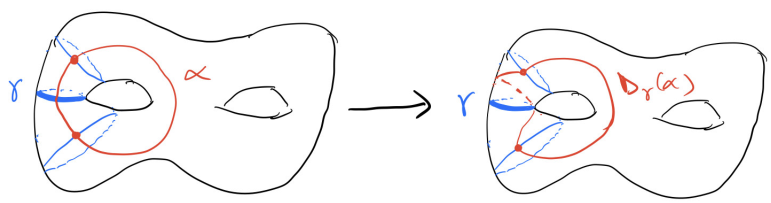

FIGURE 80. The Dehn twist about the curve γ and its effect on a curve α intersecting γ once.

It is easy to see the following.

Claim 17.5. $i(\alpha, D_\gamma^{20}(\alpha)) = 20$

Thus, since $i(\alpha,\alpha) = 0$, $D_\gamma^{20}([\alpha]) \neq [\alpha]$ and so $D_\gamma \neq id$.

Remark 17.6. One can think of changing a twist in a gluing as a "partial Dehn twist."

17.2. **Fenchel-Nielsen Coordinates.** We will now prove a theorem that creates a rough picture of what Teichmuller space looks like.

Theorem 17.7. $T_g \cong \mathbb{R}^{3g-3}_{>0} \times \mathbb{R}_{3g-3}$

Sketch of Proof. We first define both sets of coordinates (namely the length and twist coordinates). Pick a pants decomposition of S given by cuff curves $\gamma_1, ...\gamma_{3g-3}$. See Figure 81.

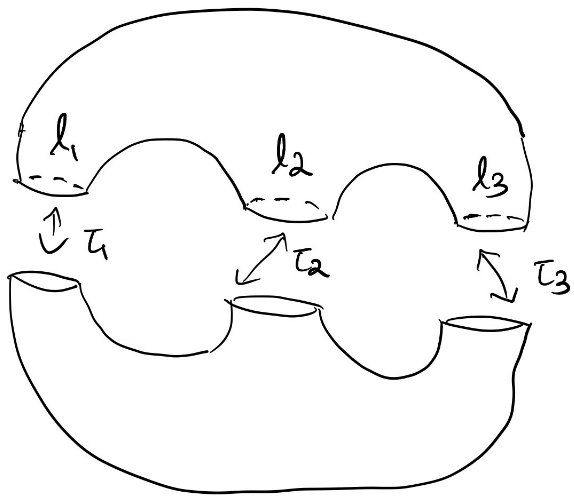

FIGURE 81. Gluing two pairs of pants with length coordinates l_i and twist coordinates τ_i to get a hyperbolic surface of genus 2. There are $3 * 2 - 3$ of each set of coordinates.

- **Length parameters:** Map $X \mapsto l_{\gamma_i}(X)$, so that these $3g - 3$ maps give the $3g - 3$ length coordinates.
- **Twist parameters:** This one is more subtle. Orient all γ_i and for simplicity, assume that they are all non-separating. Let α_i be s.c.c.'s intersecting γ_i exactly once. For example, consider α in Figure 78. It is a fact that there exists a unique curve homotopic to α_i that is made of a geodesic segment orthogonal to γ_i and a segment of γ_i together. We define the twist to be the signed length of this segment.

Note that since twisting by a full circle around any γ_i (that is, by the Dehn twist D_{γ_i}) actually changes the marking and thus the point in T_g, the twist parameter is in fact unbounded.

It is a fact that any 3 length parameters prescribe a unique hyperbolic pant. Also, one can see that the twist parameters describe a unique way to glue these pants together, recovering for us a unique hyperbolic surface up to homotopies that take these cuff curves to non-geodesic-representatives. Thus, these parameters define a unique point in Teichmuller space.

It takes a little bit of work to show that in any of the few definitions of the topology on Teichmuller space, this is actually a homeomorphism. One can also just use this bijection to (non-canonically) define the topology on Teichmuller space and attempt to show that the identity map is a homeomorphism between topologies induced by different pants decompositions. $\square$

Definition 17.8. The **moduli space of hyperbolic metrics on a surface of genus** g is $M_g := T_g/MCG(S)$.

Remark 17.9. Say $[\phi_1 : S \to X], [\phi_2 : S \to X] \in T_g$. Then for $f = \phi_1 \circ \phi_2^{-1}$, $[f] \cdot [\phi_1 : S \to X] = [\phi_2 : S \to X]$. So, any two different markings of a hyperbolic surface are related by a mapping class. The converse is clearly true – when a mapping class acts on a point in T_g, it does not change the hyperbolic surface itself, just the marking. This means that each hyperbolic surface has exactly one representative in M_g.

By the remark above, one can think of M_g as the space of hyperbolic surfaces up to isometry. The natural projection map $T_g \to M_g$ can be represented by the map that takes $[\phi : S \to X] \mapsto X$. That is, it "forgets the marking."

Remark 17.10. It is a fact that the projection $T_g \to M_g$ is an orbifold covering. So, M_g is a "nice" Hausdorff space – namely, an orbifold.

We now want to define an object that leads to the construction of a map from $T_g \to \mathcal{C}(S)$.

Definition 17.11. Let $M_g^{\geqslant \epsilon} = \{\text{hyperbolic surfaces with systolic length} \geqslant \epsilon\}$. This is called the **$\epsilon$-thick part of moduli space**.

In the next class, we will prove the following theorem.

Theorem 17.12 (Mumford's Compactness Criterion)**.** *For every $\epsilon > 0$, the ϵ-thick part $M_g^{\geqslant \epsilon}$ is compact.*

18. The systole map (02/16, BZ, CK)

In this lecture, we will begin the proof that the curve complex is an electrification of Teichmüller space. We start with a characterization of divergent sequences of hyperbolic surfaces.

Let S be an orientable surface of genus g. Recall that

$$\begin{aligned}\mathcal{M}_g &= \mathcal{T}_g/\mathrm{MCG}(S) = \{X \text{ a hyperbolic surface of genus } g\}/\text{isometry},\\ \mathcal{M}_g^{\geqslant \varepsilon} &= \{X \in \mathcal{M}_g \mid \ell_{\mathrm{sys}(X)}(X) \geqslant \varepsilon\}.\end{aligned}$$

We call $\mathcal{M}_g^{\geqslant \varepsilon}$ the *ε-thick part* of the moduli space $\mathcal{M}_g$. The following theorem says that a sequence X_n of hyperbolic surfaces in $\mathcal{M}_g$ "diverges to ∞" if and only if $\mathrm{sys}(X_n) \to 0$.

Theorem 18.1 (Mumford's Compactness Criterion)**.** *For every $\varepsilon > 0$, the ε-thick part $\mathcal{M}_g^{\geqslant \varepsilon}$ is compact.*

The proof of this theorem is clarified by the following fact.

Proposition 18.2. *There is a bijection*

$$\left\{ \begin{array}{c} \textit{Topological types of pants} \\ \textit{decomposition in genus } g \end{array} \right\} \leftrightarrow \left\{ \begin{array}{c} 3\textit{-regular multigraphs on} \\ 2g-2 \textit{ vertices} \end{array} \right\}$$

Proof sketch. Let $\mathcal{P} = \{P_1, \ldots, P_{2g-2}\}$ be a pants decomposition of S. To each pair of pants P_i, we associate a vertex v_i, and to each cuff c_{ij} joining the pants P_i to the pants P_j, we associate an edge v_i—v_j. Note that we may have edges of the form v_i—v_i because a pair of pants may have one of its cuffs glued to another of its cuffs. We may also have up to three cuffs of the form c_{ij} for given i, j. □

Proof of Mumford's Compactness Criterion. Let X_n be a sequence of hyperbolic surfaces with $\ell_{\mathrm{sys}(X_n)}(X_n) \geqslant \varepsilon$ for every n. We must find a convergent subsequence of X_n. Recall there exists a constant B_g, called the *Bers constant*, such that all hyperbolic surfaces of genus g have a pants decomposition all of whose cuffs have length no greater than B_g. Pick such a pants decomposition $\mathcal{P}_n$ on X_n.

By Proposition 18.2, there are only finitely many topological types of pants decomposition in genus g. Therefore there is some subsequence $X_{n'}$ such that every $\mathcal{P}_{n'}$ is of the same topological type. By our assumption, all cuffs of these pants have lengths contained in the closed interval $[\varepsilon, B_g]$. Furthermore, we may choose a (not necessarily injective) local system of Fenchel-Nielsen coordinates on $\mathcal{M}_g$ such that all the twist coordinates of these cuffs lie in $[-B_g, B_g]$. Therefore the subsequence $X_{n'}$ is contained in a region which is the continuous image of the compact set $[\varepsilon, B_g]^{3g-3} \times [-B_g, B_g]^{3g-3}$, and hence is compact. We conclude that $X_{n'}$ has a convergent subsequence $X_{n''}$, which is what we wanted to show. □

There is a metric $d_{\mathcal{T}_g}$ on $\mathcal{T}_g$ called the *Teichmüller metric*. In our proof that $\mathcal{C}(S)$ is an electrification of $\mathcal{T}_g$, we will only use the following two facts about $d_{\mathcal{T}_g}$.

Proposition 18.3. *The Teichmüller metric $d_{\mathcal{T}_g}$ is invariant under the action of* $\mathrm{MCG}(S)$ *on $\mathcal{T}_g$, and there is a constant C_g such that*

$$d_{\mathcal{T}_g}(X, Y) \leqslant 1 \quad \Longrightarrow \quad i(\mathrm{sys}(X), \mathrm{sys}(Y)) \leqslant C_g.$$

Remark 18.4. Note that the metric in which any two distinct points have distance 1 is invariant but doesn't at all satisfy the second property. This metric shows that we really do need to know at least a tiny bit about the metric beyond that it is MCG invariant.

Let $\varepsilon_0 > 0$ be such that any two curves of length less than or equal to ε_0 are disjoint on *any* hyperbolic surface. Fix $0 < \varepsilon \leqslant \varepsilon_0$. For a simple closed curve α on S, let $S_\alpha = \{X \in \mathcal{T}_g \mid \ell_\alpha(X) \leqslant \varepsilon\}$. Let $\mathcal{E}_g$ be the electrification of $\mathcal{T}_g$ along all the sets S_α, and let c_α denote the cone point of S_α in $\mathcal{E}_g$. We will denote the induced metric on $\mathcal{E}_g$ also by $d_{\mathcal{E}_g}$. Note that a hyperbolic surface can have more than one systole; indeed, it can have up to $3g-3$ systoles. Also note that, by our hypothesis on ε_0, we have $S_\alpha \cap S_\beta \neq \varnothing$ if and only if $i(\alpha, \beta) = 0$.

Since a surface may have multiple distinct systoles, let us fix a *choice* of function $\mathrm{sys} : \mathcal{T}_g \to \mathcal{C}(S)$, where $\mathrm{sys}(X)$ is a systole of X. Note that if $X \in S_\alpha$, we need not have $\mathrm{sys}(X) = \alpha$. Nevertheless, our hypothesis on ε_0 ensures that $i(\alpha, \mathrm{sys}(X)) = 0$, and hence $d_{\mathcal{C}(S)}(\alpha, \mathrm{sys}(X)) \leqslant 1$. Hence sys is "coarsely constant" on S_α. We extend sys to a function $\mathrm{sys} : \mathcal{E}_g \to \mathcal{C}(S)$ by setting $\mathrm{sys}(c_\alpha) = \alpha$.

Theorem 18.5. *The function* $\mathrm{sys} : \mathcal{E}_g \to \mathcal{C}(S)$ *is a quasi-isometry.*

Our proof of Theorem 18.5 will be an application of Proposition 18.7.

Definition 18.6. A map $f : A \to B$ is *coarsely Lipschitz* if there exist constants K, C such that

$$d_B(f(x), f(y)) \leqslant K d_A(x, y) + C \qquad \forall x, y \in A.$$

A map $\phi : A \to A$ is C-*coarsely equal to the identity* if $|d_A(\phi(x), \phi(y)) - d_A(x, y)| \leqslant C$ for every $x, y \in A$. If coarsely Lipschitz maps f and g are such that $f \circ g$ and $g \circ f$ are coarsely the identity, the we say they are *coarsely inverse* to each other.

Proposition 18.7. *If coarsely Lipschitz maps f and g are coarsely inverse to each other, then f and g are quasi-isometries.*

The proof of the proposition is left as an exercise.

Proof of Theorem 18.5. We define cone : $\mathcal{C}(S) \to \mathcal{E}_g$ by $\mathrm{cone}(\alpha) = c_\alpha$. By Proposition 18.7, we will be done when we show that cone and sys are coarsely Lipschitz maps that are coarsely inverse to each other. To show that a map $f : A \to B$ is coarsely Lipschitz, observe that it suffices to provide a global upper bound on $d_B(f(x), f(y))$ for any $x, y \in A$ with $d_A(x, y) \leqslant 1$.

We first show that cone is coarsely Lipschitz. By our observation above, it suffices to show that if α and β are simple closed curves on S with $d_{\mathcal{C}(S)}(\alpha, \beta) = 1$, i.e. $i(\alpha, \beta) = 0$, then $d_{\mathcal{E}_g}(c_\alpha, c_\beta) \leqslant 2$. Let $X \in \mathcal{T}_g$ such that $\ell_\alpha(X)$, $\ell_\beta(X) < \varepsilon$. Since $\varepsilon \leqslant \varepsilon_0$, our hypothesis on ε_0 implies that we can construct a pants decomposition of X such that α and β are cuffs of pants. This shows that $X \in S_\alpha \cap S_\beta$, and hence this intersection of sets is nonempty. It follows that $d_{\mathcal{E}_g}(c_\alpha, c_\beta) \leqslant 2$.

We now show that sys is coarsely Lipschitz. Suppose that $X, Y \in \mathcal{E}_g$, and $d_{\mathcal{E}_g}(X, Y) \leqslant 1$. If there is a geodesic from X to Y going through a point of the form $c_\alpha \in \mathcal{E}_g$, then X and Y both lie on edges to c_α and we get that $d_{\mathcal{C}(S)}(\mathrm{sys}(X), \mathrm{sys}(Y)) \leqslant 1$. Otherwise, let X' and Y' be the points in $\mathcal{T}_g \subset \mathcal{E}_g$ nearest to X and Y, respectively. Then $d_{\mathcal{C}(S)}(\mathrm{sys}(X), \mathrm{sys}(X')) \leqslant 1$, and similarly for Y and Y'. Furthermore, $d_{\mathcal{T}_g}(X', Y') \leqslant d_{\mathcal{E}_g}(X, Y) \leqslant 1$, as in Figure 82. (Note that a geodesic in $\mathcal{E}_g$ from X' to Y' cannot go through a cone point, since such a path would have length at least 2; such any geodesic in $\mathcal{E}_g$ stays in $\mathcal{T}_g$.) By Proposition 18.3, there is a constant C_g such that $i(\mathrm{sys}(X'), \mathrm{sys}(Y')) \leqslant C_g$, and hence also a constant C_g' such that

$$d_{\mathcal{C}(S)}(\mathrm{sys}(X'), \mathrm{sys}(Y')) \leqslant C_g'.$$

Therefore $d_{\mathcal{C}(S)}(\mathrm{sys}(X), \mathrm{sys}(Y)) \leqslant C_g' + 2$, and hence sys is coarsely Lipschitz by our observation above.

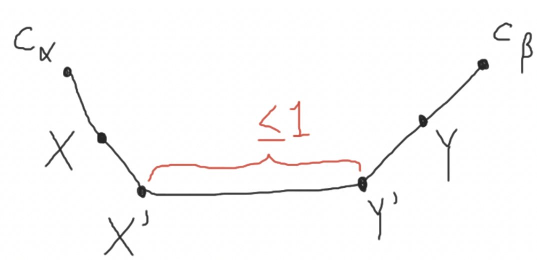

FIGURE 82. X and Y are close to points X' and Y' of bounded distance in $\mathcal{T}_g \subset \mathcal{E}_g$

In the next lecture, we will complete the proof by showing that cone and sys are coarsely inverse to each other.

19. The systole map (02/18, TY, GM)

We continue with the proof of Theorem 18.5, which states that sys : $\mathcal{E}_g \to \mathcal{C}(S)$ is a quasi-isometry. We already showed that sys and the candidate inverse map cone : $\mathcal{C}(S) \to \mathcal{E}_g$ are coarsely Lipschitz. It now suffices to show that sys and cone are coarsely inverse to each other.

Proof of Theorem 18.5 cont. First, note by definition that for $\alpha \in \mathcal{C}(S)$, we have that

$$\mathrm{sys}(\mathrm{cone}(\alpha)) = \mathrm{sys}(c_\alpha) = \alpha,$$

where c_α is the cone point of α.

Now, we need to consider cone $\circ$ sys. For a cone point $c_\alpha \in \mathcal{E}_g$, we have that

$$\mathrm{cone}(\mathrm{sys}(c_\alpha)) = \mathrm{cone}(\alpha) = c_\alpha,$$

again by definitions. The argument for a point of $\mathcal{E}_g$ that's on an edge is similar, up to a bounded error, as we can define the image under sys of such a point to be $\alpha \in \mathcal{C}(S)$.

We now need to show that for other points of the electrification, namely $X \in \mathcal{T}_g$, we have that $d_{\mathcal{E}_g}(\mathrm{cone}(\mathrm{sys}(X)), X)$ is bounded by a constant only depending on g. Suppose $\mathrm{sys}(X) = \alpha$, so we want to bound $d_{\mathcal{E}_g}(c_\alpha, X)$. It then suffices to bound $d_{T_g}(S_\alpha, X)$, since $d_{\mathcal{E}_g}(c_\alpha, X) \leqslant d_{T_g}(S_\alpha, X) + 1$.

We use ideas from Mumford's Compactness Criterion (Theorem 18.1) to prove this bound (alternatively, one could use knowledge of the Teichmüller metric).

If the length of α in X is less than ε, then note that $X \in S_\alpha$, and so we're done. Thus, assume otherwise.

We have a pants decomposition of X that includes α where all cuff lengths are contained in the closed interval $[\varepsilon, B_g]$, where B_g is a constant depending only on g. The existence of such a finite constant B_g can be shown similarly as the proof that the Bers constant is finite: any scc can be extended to a pants decomposition with cuff lengths are all bounded by a constant only depending on g and the length of the scc, and so since α has length bounded only in terms of g, we have a bound on every curve in the pants involving α, which is our B_g. WLOG, we can assume that all the twist coordinates of these cuffs lie in $[0, B_g]$, since doing a Dehn twist about the cuff changes the twist parameter by the length.

Let X' have the same Fenchel-Nielsen coordinates as X, but with the α length now set to ε, so both X, X' are in a fixed compact set of T_g, where all lengths are in $[\varepsilon, B_g]$, and all twists are in $[0, B_g]$. In addition, $X' \in S_\alpha$ by construction.

Then, the distance between X, X' is bounded above by the diameter of this compact set, and so $d_{T_g}(S_\alpha, X)$ is bounded as well. □

The main takeaway from this result is that we have a coarsely Lipschitz map sys : $T_g \to \mathcal{C}(S)$. Note that it is surjective, and it remembers a lot, as long as you believe that $\mathcal{C}(S)$ is big. However, the map also forgets a lot, as it fails to be a quasi-isometric embedding. In particular, it crushes each S_α to a point, so there are many things about T_g that you cannot see just by looking at the curve complex. Knowing $\mathrm{sys}(X) = \alpha$ doesn't tell you much about the shape of $X - \alpha$; all you know is that α is short.

The idea to remedy this is to use subsurfaces of S (e.g., thinking of $S - \alpha$ as a subsurface) and their curve complexes. In the future, we'll get a map

$$T_g \to \prod_{U,\text{ subsurface of } S} \mathcal{C}(U)$$

that is (speaking roughly in a way that will have to be corrected/clarified later) a quasi-isometry to its image; we'll now define $\mathcal{C}(U)$.

Let $U \subset S$ be a connected (for now), closed subsurface, whose boundary consists of non-essential curves. Equivalently, U is a component of the complement of a union of disjoint scc's. For now, we'll also exclude the following.

(1) Annulus, or a sphere with two boundary components. We'll later see that we actually need such subsurfaces, but they are annoying to deal with.
(2) Pants. For this and the annulus, the problem is that every scc is **peripheral**, meaning that it's homotopic to a boundary curve, and so if we defined the curve complex the usual way, it would be empty.
(3) Torus with one boundary component.
(4) Sphere with four boundary components. For this and the torus with one boundary component, there are non-peripheral simple closed curves, but any 2 such curves intersect. This will just require some slightly special definitions.

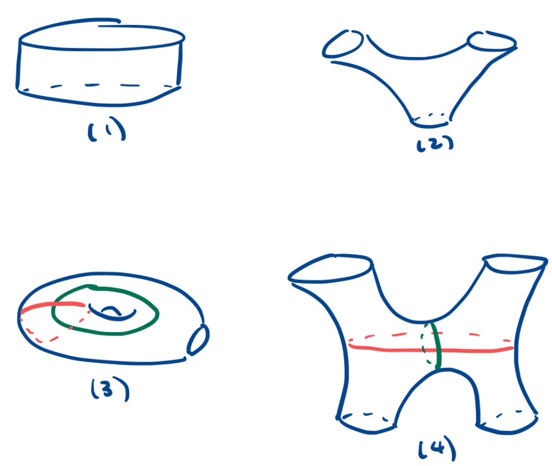

FIGURE 83. Subsurfaces that we are excluding for now.

Definition 19.1. The **curve complex of** U is the graph $\mathcal{C}(U)$ whose vertices are non-peripheral scc in U, with edges for disjoint curves.

Remark 19.2. We have that $\mathcal{C}(U) \hookrightarrow \mathcal{C}(S)$, as every curve of U gives a curve in S. However, the image has diameter 2, as a curve in the boundary of U has distance one from each curve coming from U.

Much of what we've proven about $\mathcal{C}(S)$ also holds for $\mathcal{C}(U)$, and can be proven similarly. For example, $\mathcal{C}(U)$ is connected, with $dist(\alpha, \beta) \leqslant O(\log i(\alpha, \beta))$ and it's δ-hyperbolic. It'll also turn out to be infinite diameter, though it wouldn't be if we included peripheral curves.

Optional Exercise 23. For $\epsilon > 0$ and α a scc, and let $S_\alpha(\epsilon)$ be the subset of Teichmüller space where $\ell_\alpha(X) \leqslant \epsilon$. Show that for any ϵ, even large ϵ, the diameter of $\mathrm{sys}(S_\alpha(\epsilon))$ is bounded only in terms of ϵ. (The bound gets worse as ϵ gets big. You might need the collar lemma for this. Once you do this, you should be able to check that the electrification of $\mathcal{T}_g$ along the $S_\alpha(\epsilon)$ is qi to $\mathcal{C}S$ for any ϵ, although the constants depend on ϵ.)

Optional Exercise 24. Almost everything we did regarding electrifying Teichmüller space worked for any metric on $\mathcal{T}_g$ that is mapping class group invariant and bounded on compact sets. But it doesn't work, for example, for a bounded metric on $\mathcal{T}_g$. The one place we used what the metric was is showing that if $d(X, Y) \leqslant 1$ then $d(\mathrm{sys}(X), \mathrm{sys}(Y))$ is bounded. Is that true for the WP metric? A naive guess is that it's true basically for the WP metric and anything bigger, but probably it wouldn't be true for a metric for which the metric completion of $\mathcal{M}_g$ is smaller than Deligne-Mumford.

Optional Exercise 25. Show that the image of a Teichmüller disc under the systole map is a quasi-tree. (This exercise is harder than most of the others. It requires in particular the non-trivial fact that a quadratic differential of area 1 and fixed genus always has a cylinder whose modulus is larger than some constant depending only on the genus, which is proven in [Vor05].)

20. Subsurface projections (02/21, SK, FAH)

Recall that we mentioned in an earlier class that (speaking roughly in a way that will have to be corrected/clarified later) the Teichmüller space quasi-isometrically embeds in $\prod_{U \text{ a subsurface}} \mathcal{C}(U)$. To understand the image of $\mathcal{T}_g$ in this product, we will also need to understand how the curve complexes of subsurfaces interact with each other.

Let S be the ambient surface, and U be a subsurface that is not one of the 4 exceptional types of subsurfaces discussed in the previous section. We define the *subsurface projection map*

$$\rho_U^S : \mathcal{C}(S) \to 2^{\mathcal{C}(U)}.$$

Here $2^{\mathcal{C}(U)}$ denotes the set of all subsets of $\mathcal{C}(U)$.

Remark 20.1. Despite the subsurface projection map taking values in subsets of $\mathcal{C}(U)$, one should really think of it as taking values in $\mathcal{C}(U)$. The reason for this is that for some choices of input, there is no reasonable value we can assign, in which case the output is the empty set, and for other input values, the output is only well-defined up to a set of bounded diameter. One can think of this as the output being a point at a large enough scale, which is very much in the spirit of coarse geometry.

Defining the subsurface projection maps. To define the map ρ_U^S for a curve α, we need to deal with three cases. Crucially, we assume there are no bigons between α and the boundary of U.

Case 1: α disjoint from U: If α is disjoint from U there's no reasonable definition of ρ_U^S; we define $\rho_U^S(\alpha)$ to be the empty set. (Warning: Later we might change our mind and decide we'd rather have it be all of $\mathcal{C}U$. Both conventions have their advantages.)

Case 2: α contained in U: In this case, $\rho_U^S(\alpha) = \{\alpha\}$.

Case 3: α intersects U, but it not fully contained in U: In this case, consider all the arcs obtained by intersecting α with U. Each of these arcs go from one boundary component of U to another (possibly the same) one. For each such arc α', we consider an ε-neighbourhood $N_\varepsilon(\alpha' \cup \partial U)$ of $\alpha' \cup \partial U$. This is again a submanifold with boundary: we consider the boundary $\partial\left(N_\varepsilon(\alpha' \cup \partial U)\right)$, which will be a collection of disjoint curves. We discard any curves that are non-essential or peripheral, and take the union of the resulting curves as we vary over all the arcs α'. This set of essential non-peripheral curves is defined to be $\rho_U^S(\alpha)$.

Examples of subsurface projections in Case 3. We can assume without loss of generality that each arc α' forms no bigons with any boundary component. If it does, it's easy to see that the resulting curves in U will be peripheral or non-essential, and thus be discarded. Figures 84, 85, and 84 illustrate examples of subsurface projection when the arc α' intersects the surface and starts and ends at boundary components. In each example, the red arc is the arc α', the green curve(s) are the boundary components of $N_\varepsilon(\alpha' \cup \partial U)$ which do not get discarded, and the blue curves are the ones that do.

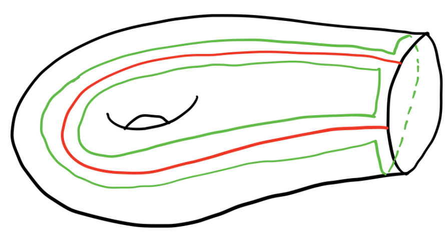

FIGURE 84. This arc results in two curves which are homotopic.

Non-emptiness of the subsurface projection in Case 3. Since the definition of the subsurface projection map in Case 3 involves discarding peripheral and non-essential curves, it is not clear that the resulting subset of $\mathcal{C}(U)$ is always non-empty. We prove that claim in this section.

Lemma 20.2. *Continue to assume α does not form any bigons with ∂U. Then for each arc α', at least one component of $\partial N_\varepsilon(\alpha' \cup \partial U)$ is essential and non-peripheral.*

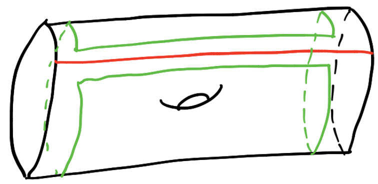

FIGURE 85. This arc results in only one curve.

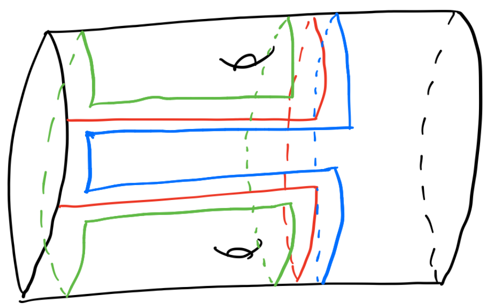

FIGURE 86. This arc results in two curves but one of them is peripheral.

Remark 20.3. The above lemma fails to be true for 2 of the 4 exceptional subsurfaces we are excluding, namely the $S_{0,2}$ (an annulus) and $S_{0,3}$ (a pair of pants). In the former case, the resulting curve is always non-essential, and in the latter case, the resulting curve is always peripheral. The lemma is however true for the other two exceptional subsurfaces; the proof works without any modification in those cases.

Proof of Lemma 20.2. We split the analysis up into two cases.

α' joins two distinct boundary components: Call the two distinct boundary components β and γ. Then $\partial(N_\varepsilon(\alpha' \cup \beta \cup \gamma))$ has one component, and that component, along with β and γ bounds subsurface of U that is homeomorphic to $S_{0,3}$ (a pair of pants) (see Figure 87). Suppose now that U had genus k and h boundary components. It follows from an Euler characteristic computation that the complementary subsurface to the pair of pants is a surface of genus k and $h-1$ boundary components. Since $(k, h-1)$ is not equal to $(0, 1)$, we know that the green curve is not non-essential, because that would correspond to bounding a disc. Similarly, since $(k, h-1)$ is not equal to $(0, 2)$, we know that the green curve is not peripheral, since that would correspond to the complementary subsurface being an annulus.

α' is connected to one boundary component: Let β be the boundary component α' is connected to. In this case $\partial(N_\varepsilon(\alpha' \cup \beta))$ will have two components γ and γ', and γ, γ' together bound a pair of pants, which we call P. Here, we

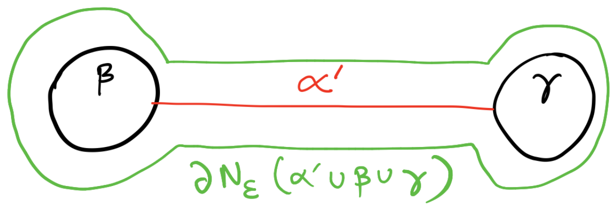

Figure 87. The green curve and the two boundary components bound a pair of pants.

again need to deal with two cases: whether deleting P from U results in two components or one component (see Figure 88). We first deal with the case

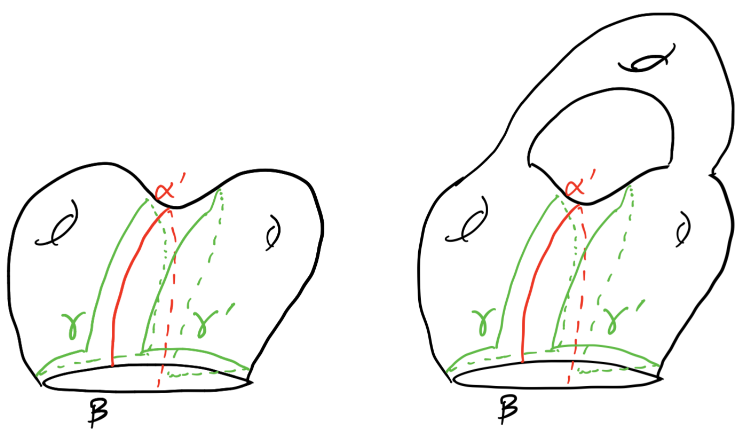

Figure 88. The two cases we can get if we delete P.

that there are two components left after deleting P. For both γ and γ' to be discarded, they must be peripheral or non-essential. That can only happen if the corresponding complementary subsurface is an annulus or a disc. If both the complementary subsurfaces were discs, U would be a disc. If both the complementary subsurfaces were annuli, U would be a pair of pants. And if one of the complementary subsurface was a disc, and one an annulus, U would be an annulus as well. Since U is none of the above, it must be the case that at least one of γ or γ' is not discarded.

Now suppose that there was only one component left after deleting P. By an Euler characteristic computation, that subsurface must have genus $k-1$ and have $h+1$ boundary components, where U had genus k and h boundary components. We see that since U is not one of the exceptional subsurfaces, the complementary subsurface to P is not an annulus proving that neither γ nor γ' get discarded.

We have dealt with both cases and therefore proved the lemma. $\square$

Optional Exercise 26. Let α be any simple closed curve, and let $\beta \in \mathcal{C}U$ be arbitrary. Show that there is a curve $\hat{\alpha}$ isotopic to α and an arc $\hat{\alpha}'$ of $\hat{\alpha} \cap U$ such that β is isotopic to a boundary component of $N_\epsilon(\hat{\alpha}' \cup \partial U)$.

21. Subsurface projections (02/23, SC, SK)

In the previous class, we had defined the *subsurface projection map*

$$\rho_U^S : \mathcal{C}(S) \to 2^{\mathcal{C}(U)}.$$

We now show that each subset of $\mathcal{C}(U)$ in the image of ρ_U^S has bounded diameter.

Lemma 21.1. *If $\gamma_1, \gamma_2 \in \rho_U^S(\alpha)$, then $i(\gamma_1, \gamma_2) \leqslant 4$.*

Since, by Lemma 10.10, $d_{\mathcal{C}(U)}(\gamma_1, \gamma_2) \leqslant 2\log_2(i(\gamma_1, \gamma_2)) + 2$, we get the following corollary.

Corollary 21.2. *$\rho_U^S(\alpha)$ has uniformly bounded diameter.*

Proof of lemma. By definition of the subsurface projection map, we may assume that γ_1 is a component of $N_\epsilon(\alpha_1 \cup \partial U)$ and γ_2 is a component of $N_{2\epsilon}(\alpha_2 \cup \partial U)$, where α_1 and α_2 are some arcs of $\alpha \cap U$. Then each γ_i has one or two arcs following α_i and one or two arcs following ∂U (Figure 89).

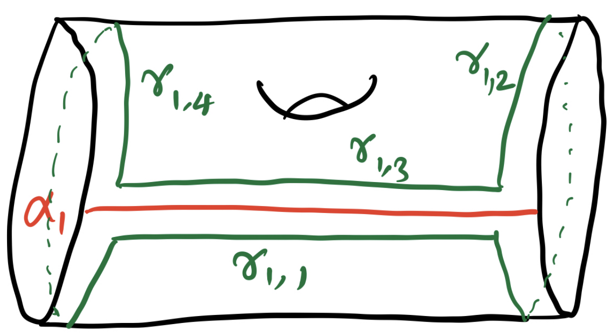

Figure 89. $\gamma_{1,1}$ and $\gamma_{1,3}$ are arcs of γ_1 following α_1 and $\gamma_{1,2}$ and $\gamma_{1,4}$ are arcs of γ_1 following ∂U

Since α is simple, α_1 and α_2 are disjoint. By letting ϵ be sufficiently small, we may assume that an arc of γ_1 following α_1 does not intersect an arc of γ_2 following α_2. Since arcs of γ_2 following ∂U are 2ϵ away from ∂U and arcs of γ_1 following ∂U are ϵ away from ∂U, we may also assume that an arc of γ_1 following ∂U does not intersect an arc of γ_2 following ∂U (Figure 90).

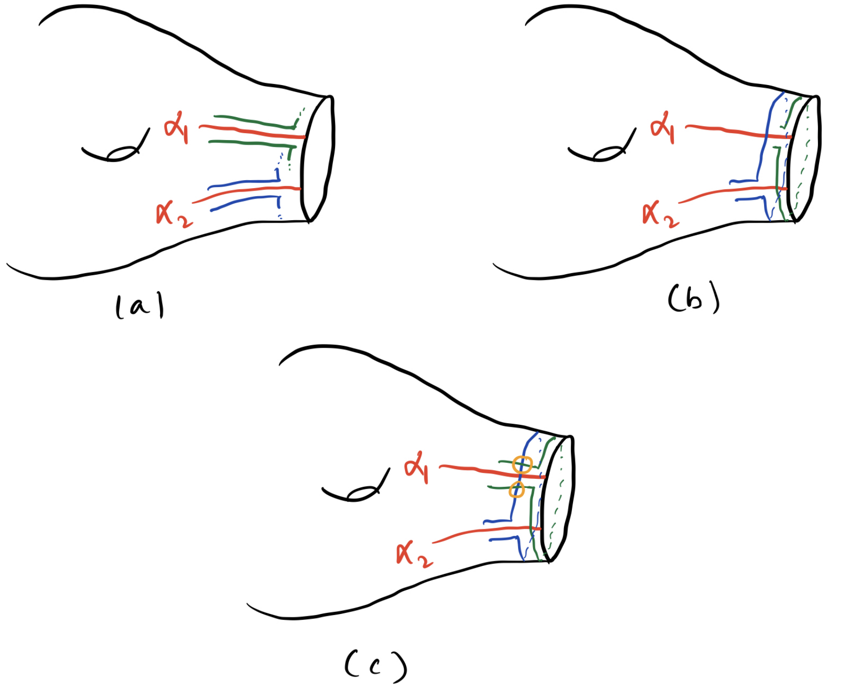

FIGURE 90. (a) Arcs of γ_1 following α_1 do not intersect arcs of γ_2 following α_2. (b) Arcs of γ_1 following ∂U do not intersect arcs of γ_2 following ∂U (c) Intersections between arcs of γ_1 following α_1 and an arc of γ_2 following ∂U

Thus, the only possible intersections can be between arcs of γ_1 following α_1 and arcs of γ_2 following ∂U (Figure 90). Since there are at most two arcs of γ_1 following α_1 and at most two arcs of γ_2 following ∂U, γ_1 and γ_2 can intersect at most 4 times. □

We observe that the proof of lemma 21.1 also goes through when $\gamma_1 \in \rho^S_U(\alpha)$ and $\gamma_2 \in \rho^S_U(\beta)$, for disjoint curves α and β, giving us the following corollary,

Corollary 21.3. *If $\gamma_1 \in \rho^S_U(\alpha)$ and $\gamma_2 \in \rho^S_U(\beta)$, for disjoint curves α and β, then $i(\gamma_1, \gamma_2) \leqslant 4$.*

We sometimes modify the subsurface projection map so that it takes values in $\mathcal{C}(U)$, by picking an arbitrary curve in $\rho^S_U(\alpha)$, for each α. Since $\rho^S_U(\alpha)$ has uniformly bounded diameter, this map is coarsely well defined. However, we need $\rho^S_U(\alpha)$ to be non-empty to make a choice for the image. So, keeping in mind Lemma 20.2, we restrict the domain to get a new map, still denoted by ρ^S_U,

$$\rho^S_U : \mathcal{C}(S) - (\mathcal{C}(S - U) \cup \partial U) \to \mathcal{C}(U).$$

For convenience, we henceforth denote $\mathcal{C}(S) - (\mathcal{C}(S - U) \cup \partial U)$ by $\mathcal{C}(S, U)$.

Proposition 21.4. *$\rho^S_U : \mathcal{C}(S, U) \to \mathcal{C}(U)$ is Lipschitz, with the induced graph metric on $\mathcal{C}(S, U)$.*

Proof. It suffices to give a uniform bound on

$$d_U(\alpha, \beta) := d_{\mathcal{C}(U)}(\rho_U^S(\alpha), \rho_U^S(\beta))$$

for all $\alpha, \beta \in \mathcal{C}(S, U)$ with $d_{\mathcal{C}(S,U)}(\alpha, \beta) = 1$. For then, given arbitrary α and β, we could consider a path of curves from α to β at unit distances and use the triangle inequality.

If $d_{\mathcal{C}(S,U)}(\alpha, \beta) = 1$, then $d_{\mathcal{C}(S)}(\alpha, \beta) = 1$, which means α and β are disjoint. Corollary 21.3 then tells us that $i(\rho_U^S(\alpha), \rho_U^S(\beta)) \leqslant 4$, which in turn gives us that

$$d_{\mathcal{C}(U)}(\rho_U^S(\alpha), \rho_U^S(\beta)) \leqslant 2\log_2(4) + 2 = 6,$$

concluding the proof. □

Now consider the inclusion $\mathcal{C}(U) \hookrightarrow \mathcal{C}(S, U)$, which is 1-Lipschitz. Proposition 21.4 tells us that the map $\rho_U^S : \mathcal{C}(S, U) \to \mathcal{C}(U)$, which is the left inverse to this inclusion, is also Lipschitz. We thus get the following corollary,

Corollary 21.5. *The inclusion $\mathcal{C}(U) \hookrightarrow \mathcal{C}(S, U)$ is a quasi isometric embedding.*

Morally, this shows that the metric on $\mathcal{C}(U)$ can be recovered from the geometry of $\mathcal{C}(S)$. It also indicates that $\mathcal{C}(U)$ is "undistorted" in $\mathcal{C}(S, U)$.

Some might say that the fact that $\mathcal{C}(U)$ is hyperbolic, and that its geometry can be recovered from $\mathcal{C}(S)$ via $\mathcal{C}(S, U)$, is an indication of extra, hidden, hyperbolicity in $\mathcal{C}(S)$ that goes beyond the hyperbolicity of $\mathcal{C}(S)$ itself.

Moving towards stating the Behrstock inequality, we make some definitions.

Definition 21.6. Let U, V be subsurfaces of S.

- If $U \subseteq V$ (possibly after isotopy), we write $U \sqsubseteq V$ and say that U is *nested* in V.
- If $U \cap V = \varnothing$ (possibly after isotopy), we write $U \perp V$ and say that U is orthogonal to V.
- If neither of the above holds, we write $U \pitchfork V$ and say that U is transverse to V.

We note that if $i(\partial U, \partial V) \neq 0$, then $U \pitchfork V$, but the converse is not true (Figure 91).

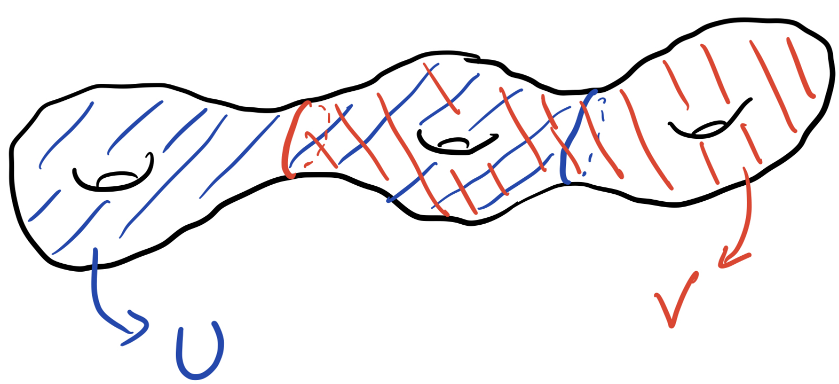

FIGURE 91. $U \pitchfork V$ but $i(\partial U, \partial V) = 0$

Definition 21.7. If $U \pitchfork V$ or $U \subsetneq V$, we define

$$\rho_V^U = \bigcup_{\gamma \in \partial U} \rho_V^S(\gamma).$$

We think of ρ_V^U as the projection of ∂U to $\mathcal{C}(V)$. Since any two curves of ∂U are disjoint in S, corollary 21.3 applies and we get that ρ_V^U has bounded diameter.

22. The Behrstock inequality (02/25, KS, AW)

Lemma 22.1 (Behrstock inequality)**.** *There exists $C > 0$ such that if $U \pitchfork V$ then*

$$\min\left(d_U(\alpha, \rho_U^V), d_V(\alpha, \rho_V^U)\right) \leqslant C.$$

Remark 22.2. The case when $\alpha \subset S - U$ or $S - V$ is problematic unless you make the convention that in that case $\rho_U^S(\alpha) = 2^{\mathcal{C}U}$.

Proof. Let us give the sketch of the proof, following [Man10, Lemma 2.5]; see also [Man13, Lemma 2.13] for the case of exceptional subsurfaces, which we exclude here.

We will show that if $d_U(\alpha, \rho_U^V)$ is big then $d_V\left(\alpha, \rho_V^U\right)$ is small. For concreteness, say $\rho_U^S(\alpha), \rho_U^V, \rho_V^U$ are points. Assume $d_U(\alpha, \rho_U^V)$ is big. Then $i(\rho_U^S(\alpha), \rho_U^V)$ is big, so the arc $\alpha \cap U$ intersects ∂V at least three times. Since we have at least three intersections, we must have a segment α'' of $\alpha \cap V$ that lies in U and is bounded by points of ∂V (see Figure 92). Since the segment α'' is disjoint from ∂U, we get that $i(\rho_V^S(\alpha), \rho_V^U) \leqslant 4$, so $d_{\mathcal{C}V}(\rho_V^S(\alpha), \rho_V^U)$ is bounded. ∎

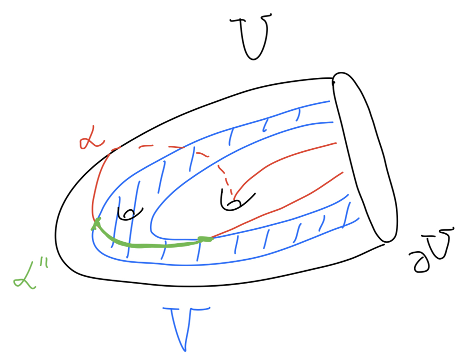

Figure 92. Green segment $\alpha'' \subset U$ bounded by points in ∂V

The statement of Behrstock inequality is illustrated in Figure 93.

Let us consider two related examples using closest point projections. Although these examples don't involve curve complexes, they give excellent motivation, and are useful in their own right.

Example 22.3. Let T be a tree, U, V be disjoint connected subsets of a tree. Let $\rho_V^U = \pi_V(U)$ be the closest point projection of U onto V. Then for any $\alpha \in T$

$$\min\left(d(\pi_U(\alpha), \rho_U^V), d(\pi_V(\alpha), \rho_V^U)\right) = 0.$$

To see that it is true, it is enough to pick y on the edge between U and V. WLOG α is on U side, then $\pi_V(\alpha) = \rho_V^U$ (see Figure 94).

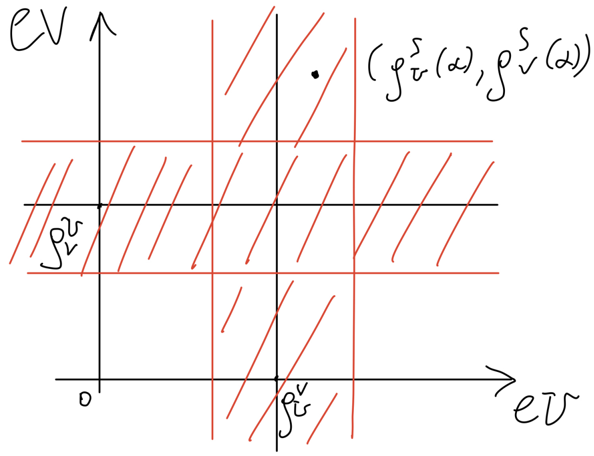

FIGURE 93. Behrstock inequality

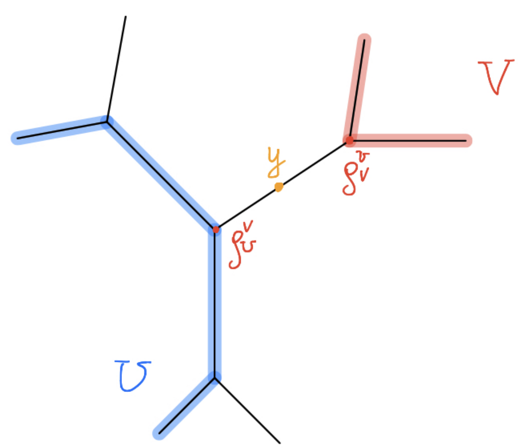

FIGURE 94. Behrstock inequality in a tree ($C = 0$)

Example 22.4. Let T be a δ-hyperbolic space, U, V quiasiconvex and far apart. By Lemma 9.3, $\rho^U_V = \pi_V(U)$ has a bounded diameter. For any x consider minimal length geodesic γ to V. We have two cases (see Figure 95).

Case 1: Suppose γ stays far from U, then $\pi_U(\gamma)$ is bounded, and $\pi_U(x)$ is close to $\pi_U(x') \in \pi_U(V)$. So it follows from Lemma 9.9 that $\pi_U(x)$ is close to $\pi_U(V)$.

Case 2: Suppose γ comes close to U, then $\pi_V(x) = \pi_V(x')$ is close to $\pi_V(x'')$ (since projection is coarse Lipschitz). So it follows that $\pi_V(x)$ is close to $\pi_V(U)$.

Optional Exercise 27. Recall that there are constants $C = C(\delta), D = D(\delta)$ such that if X is a δ-hyperbolic space, and U is a convex set, and γ is a geodesic segment that stays C away from U, then $\operatorname{diam}(\pi_U(\gamma)) < D$.

Suppose that U and V are both convex subsets such that the intersection of each with the C-neighbourhood of the other is bounded diameter. Show that $\pi_U(V)$ is bounded

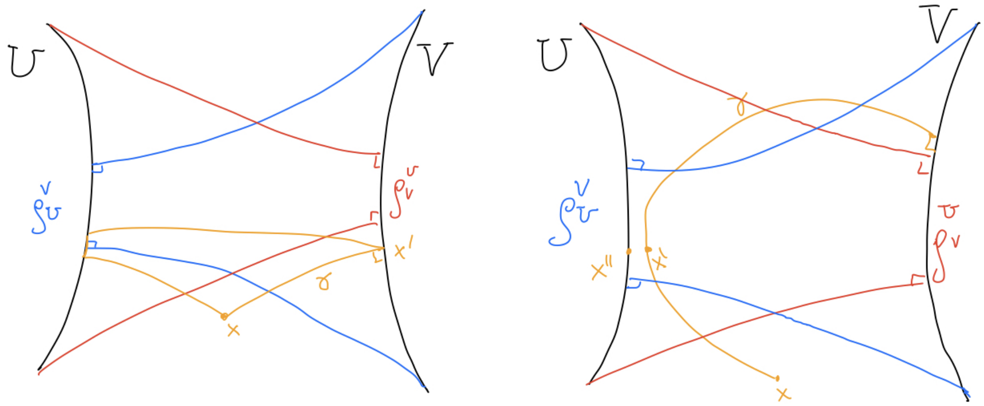

FIGURE 95. Behrstock inequality in a δ-hyperbolic space (left: case 1, right: case 2)

diameter and vice versa. Check that a version of Behrstock's inequality holds in this context, generalizing Example 22.4.

23. The Bounded Geodesic Image Theorem (03/07, SK, BZ)

We begin by recalling Lemma 9.9, which is a result about closest point projections in Gromov hyperbolic spaces.

Lemma 9.9. *Let X be δ-hyperbolic and $S \subseteq X$ a C-quasi-convex subspace. Then there exists a $B = B(\delta, C) > 0$ such that if γ is a geodesic segment with*

$$\gamma \cap N_{C+2\delta+1}(S) = \varnothing,$$

then $\operatorname{diam} \Pi_S(\gamma) \leqslant B$.

In this section, we prove an analogous result for subsurface projections, which suggests that subsurface projections behave like closest point projections.

Theorem 23.1 (Bounded Geodesic Image)**.** *There exists (possibly large) positive number E (depending only on the topology of S), such that if V is any subsurface of S, and γ any geodesic segment in $\mathcal{C}(S)$ such that the following holds*

$$\gamma \cap N_E(\rho_S^V) = \varnothing,$$

then $diam_{\mathcal{C}(V)}(\rho_S^V(\gamma)) \leqslant E$.

Remark 23.2. The conclusion of the theorem is also true if the hypothesis $\gamma \cap N_E(\rho_S^V) = \varnothing$ is replaced with the hypothesis that γ is contained in $\mathcal{C}(S) \backslash (\mathcal{C}(S \backslash V) \cup \partial V)$.

This is the weakest hypothesis that is possible: consider a geodesic segment of length 2 that starts and ends in $\mathcal{C}(V)$ but whose midpoint is in $\mathcal{C}(S \backslash V) \cup \partial V$. This geodesic can have arbitrarily large projection to $\mathcal{C}(V)$.

The stronger version of the theorem follows from the Theorem 23.1 by dividing γ into (at most) three segments: a "middle" segment of bounded length, and two segments that may have unbounded length but are at least E away from ρ_S^V. The middle segment had bounded image under ρ_S^V because ρ_S^V is Lipschitz, and the other two segments have bounded length by Theorem 23.1. (One can define the "middle" segment as the segment between the first and last intersections of γ with $N_E(\rho_S^V)$.)

To prove Theorem 23.1, we will prove an analogous statement about bicorns $\eta(x,y)$ interpolating between points x and y in $\mathcal{C}(S)$. We have already seen that bicorns stay within bounded Hausdorff distance of geodesics in $\mathcal{C}(S)$, so proving the result for the bicorn interpolation will suffice. (The optimal constant E for bicorn paths may be different than the optimal constant E for geodesics.) We first state a key lemma.

Lemma 23.3. *Let z be a bicorn between x and y, and let x' and y' denote the arcs of x and y forming z. If $\sum_{b\in\partial V} i(z,b) \geqslant 5$, the following inequality holds for some uniform constant k:*

$$\min\left(d_V(x,z), d_V(y,z)\right) \leqslant k.$$

Here $d_V(x,z)$ is shorthand for distance in $\mathcal{C}(V)$ of the subsurface projections of x and z, and $d_V(y,z)$ is the distance between the subsurface projections of y and z.

Before we prove this key lemma we show how it implies Theorem 23.1.

Remark 23.4. It may be helpful to note the following fact as a warm up: Suppose we are given a continuous map η from a closed interval $[x,y]$ to a metric space. Suppose that for all $z \in [x,y]$, we know that $\eta(z)$ is distance at most 1 from at least one of $\eta(x)$ or $\eta(y)$. Then the distance between $\eta(x)$ and $\eta(y)$ is at most 2. What we will do next will involve a coarse version of this.

Proof of Theorem 23.1. We assume that every point on $\eta(x,y)$ is more than distance E from ∂V in $\mathcal{C}(S)$. From the logarithmic inequality between curve complex distance and intersection numbers, we can conclude that for every $z \in \eta(x,y)$, $\sum_{b\in\partial V} i(z,b)$ is large. In particular, by picking a large enough E, we can ensure that the intersection number is at least 5.

Lemma 23.3 then tells us that $\rho_V^S(z)$ lies within distance k of either $\rho_V^S(x)$ or $\rho_V^S(y)$. Recall also that $\eta(x,y)$ is coarsely connected: In particular, Lemma 13.3 tells us that x and y can be joined by a path of bicorns with each having distance at most 2 to the next. Since the map ρ_V^S is coarsely (m,j)-Lipschitz for some m,j (Proposition 21.4), we can conclude that the projections of successive bicorns are distance at most $2m+j$ apart. Since all the bicorns are also distance at most k from either the projection of x or y, we can conclude that the projections of x and y are distance at most $2k+2m+j$ from each other, which proves the result. □

To complete the proof of the Theorem 23.1, we now need to prove Lemma 23.3.

Proof of Lemma 23.3. For simplicity, we first assume that z does not form bigons with ∂V. In this case, we have the following equality for the intersection number of z with ∂V.

$$\sum_{b\in\partial V} i(z,b) = \#\left(x' \cap \partial V\right) + \#\left(y' \cap \partial V\right)$$

Since the sum of the two terms is at least 5, one of them is at least 3. Without loss of generality, we assume it's the first term.

Consider the regions of x' in between successive intersections with ∂V, excluding the regions containing endpoints of x'. Since there are at least 3 intersections, there are at least 2 regions, and at least one of them is contained in V (see Figure 96). Call the

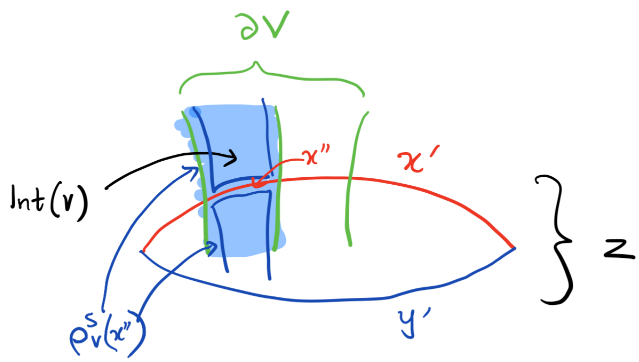

FIGURE 96. The projection of the arc x'' is a component of the projection of x and z.

corresponding arc x''. Observe now the curves in V obtained by projecting x''. One of the curves will be essential and non-peripheral, and be a component of both $\rho^S_V(x)$ and $\rho^S_V(z)$. Since we know that the diameters of $\rho^S_V(x)$ and $\rho^S_V(z)$ are uniformly bounded, and have non-empty intersection, the result follows. $\square$

Remark 23.5. The analysis in this section suggests a technical connection between BGI and Behrstock. A moral connection may also be possible, since BGI suggests that if $\gamma \cap N_E(\rho^V_S) = \varnothing$, then "$\gamma$ is transverse to V". (Imagine there was a subsurface I with $\mathcal{C}I = \gamma$, and I transverse to V. Think about what Behrstock would say.)

Optional Exercise 28. Show that if $Y \to X$ is a QI embedding of hyperbolic spaces, the image is quasi-convex in X. Conclude that there is a coarsely well defined closest point projection of $\mathcal{C}S - \alpha$ onto $\mathcal{C}(S - \alpha)$, for any (non-separating) curve α.

Optional Exercise 29. Show that the closest point projection of $\mathcal{C}S - \alpha$ onto $\mathcal{C}(S - \alpha)$ is coarsely equal to $\rho^S_{S-\alpha}$. (Exercise credit: Alessandro Sisto.)

24. HHS Axioms (03/09, JH, SC)

There is one technical point to clarify from the proof of Lemma 23.3, which is that we need to be careful to remove bigons before taking subsurface projections. The following lemma justifies the way we do this.

Lemma 24.1. *Suppose $x, y, \partial V$ do not form bigons. Let $z = x' \cup y'$ be a bicorn of x and y. Then V is isotopic to a subsurface V' that does not form bigons with z and such that every arc of $V' \cup z$ not containing a corner is an arc of $x \cap V$ or $y \cap V$.*

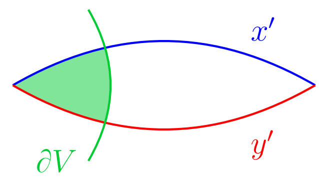

FIGURE 97. A problematic bigon.

Proof. The case we need to worry about is in Figure 97. Since $x, y, \partial V$ don't form bigons, any bigons be be at a corner of z. So, we can "push" the surface so that the boundary goes around the corner instead, as in Figure 98. Note that the shaded region does not need to be contained in V; it is just showing the bigon of ∂V and z. V could just as well be on the right hand side of the boundary. Also, if there were any genus in the shaded region, we would not have a bigon, so there is no problem. □

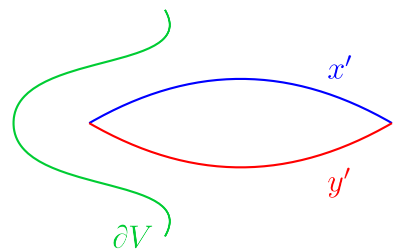

FIGURE 98. Pushing a surface to remove a bigon.

Remark 24.2. We could equally well push z off and leave V unchanged in this proof.

We now introduce the axioms for hierarchically hyperbolic spaces. Teichmüller space and the mapping class group are the most important examples and serve as motivation for the axioms. The basic idea with hierarchically hyperbolic spaces is that we have maps to hyperbolic spaces, which can be thought of as coordinates. However, there are some restrictions on these maps (mainly the Behrstock inequality), so the coordinates cannot be completely arbitrary. The result is that (in the best case scenario), we can completely reduce the study of a complicated space to the study of some hyperbolic spaces.

A good reference for this is [Sis19], and the best place to simply read the axioms all at once is [BHS19].

Definition 24.3. Suppose X is a q-quasi-geodesic space (i.e., every pair of points can be joined by a (q, q)-quasi-geodesic), with q fixed. We say X is a *hierarchically hyperbolic space*, or HHS, if there is an index set $\mathfrak{S}$, a $\delta \geqslant 0$, and a set $\{\mathcal{C}W : W \in \mathfrak{S}\}$ of δ-hyperbolic spaces subject to nine axioms.

Remark 24.4. Often, $\mathcal{C}W$ is called a curve complex even if X is not T_g. In all our examples, X will actually be a geodesic metric space rather than just a quasi-geodesic metric space; of course this doesn't matter because we are working coarsely. The set $\mathfrak{S}$ is like the set of subsurfaces, and $\mathcal{C}W$ is like the curve complex of $W \in \mathfrak{S}$.

(1) **Projections:**
There exist ξ, K and maps

$$\{\pi_W : X \to 2^{\mathcal{C}W} \mid W \in \mathfrak{S}\},$$

called *projections*, sending points in X to subsets of diameter $\leqslant \xi$, such that all these maps are (K, K)-coarse Lipschitz and have uniformly quasi-convex images.

Remark 24.5. This is the most important axiom, but we haven't discussed the π_W maps yet for our primary examples. Just to give a hint at how one might try to define them: for $X = T_g \ni x$ take $\alpha \in \operatorname{sys}(x)$, where $\operatorname{sys}(x)$ is the set of shortest curves. If α cuts W (i.e., α can't be isotoped out of W), we define $\pi_W(x) = \rho^S_W(\alpha)$; otherwise, we can try the next shortest curve, and so on. Eventually, we will get one that cuts W. We'll revisit this soon.

The purpose of the remaining axioms is to ensure that $\mathcal{C}W$ make useful coordinates and make it easy to work in coordinates.

(2) **Nesting:**
$\mathfrak{S}$ has a reflexive partial order $\sqsubseteq$ with a unique maximal element S (called the *whole surface*, or *top element*, or *main element*). We additionally require that if $V \sqsubsetneq W$, then we have two things:
- A subset ρ^V_W of $\mathcal{C}W$ of diameter at most ξ. For T_g, we will use ∂V as an element of the curve complex $\mathcal{C}W$.
- A projection map $\rho^W_V : \mathcal{C}W \to 2^{\mathcal{C}V}$.

We covered this for $W = S$, but in fact it works more generally. At this point, ρ^W_V is completely useless, but the other axioms give us information so that it is not useless.

(3) **Orthogonality:**
$\mathfrak{S}$ has a symmetric, anti-reflexive (i.e., nothing is orthogonal to itself) relation $\perp$, called *orthogonality*, such that
(3.1) $V \sqsubseteq W$ and $W \perp U$ implies $V \perp U$.
(3.2) The *container axiom* holds, which is slightly weaker than saying all V have a $V^\perp$ such that $W \perp V \iff W \sqsubseteq V^\perp$. There are examples where it was not clear how to prove the last equivalence, but it turns out that it suffices to assume something weaker. See exercise 31 for details.

For T_g, orthogonality will mean disjointness of subsurfaces, as we have seen before. In particular, it does not refer to subsurfaces that intersect orthogonally in the usual geometric sense.

Remark 24.6 (Chinmaya's observation). Note that anti-reflexivity and "$V \sqsubseteq W$ and $W \perp U$ implies $V \perp U$" implies (when $V = U$) that "$V \sqsubseteq W$" and "$V \perp W$" are mutually exclusive. That is, one subsurface cannot be both nested in and orthogonal to another.

(4) **Transversality and Consistency:**
If V and W are not orthogonal and neither is nested in the other, we say they are *transverse* and write $V \pitchfork W$.

(4.1) Behrstock Inequality:
There is a constant $\kappa_0 > 0$ such that if $V \pitchfork W$, then there are subsets $\rho_W^V \subseteq \mathcal{C}W$ and $\rho_V^W \subseteq \mathcal{C}V$ of diameter at most ξ satisfying

$$\min[d_W(\pi_W(x), \rho_W^V), d_V(\pi_V(x), \rho_V^W)] \leqslant \kappa_0$$

for all $x \in X$. Here, d_W denotes distance in $\mathcal{C}W$ and d_V distance in $\mathcal{C}V$.

(4.2) If $V \sqsubseteq W$, for all x we have

$$\min\big[d_W(\pi_W(x), \rho_W^V), \mathrm{diam}_{\mathcal{C}V}(\pi_V(x) \cup \rho_V^W(\pi_W(x)))\big] \leqslant \kappa_0.$$

This is a kind of functoriality requirement, which says that we can compose projections in the expected way with nested subsurfaces: projecting to a larger subsurface, then a smaller, nest subsurface is the same as just projecting straight to the smaller subsurface. This breaks down when subsurface projections don't make sense, which is when $\pi_W(x)$ and ρ_W^V are too close (closer than κ_0).

(4.3) Omitted here. This is a condition on some coarse points being close to others. See exercise 32 for details.

Together, 4.2 and 4.3 are called the consistency axioms, and they will eventually tell us which points are actually obtained as coordinates of points in X.

25. HHS Axioms (03/11, FAH, CK)

Recall that:

- X is a q-quasi-geodesic metric space.
- Axiom (1) provides "maps" $\pi_U \colon X \to \mathcal{C}U$ for every $U \in \mathfrak{S}$; we will sometimes call elements in S *domains.*
- Axioms (2) and (3) correspond to book-keeping for the relations $\sqsubseteq$ and $\perp$.
- Axiom (4) corresponds to the Behrstock inequality and functoriality, i.e., for $U \sqsubseteq V$, $\pi_U \approx \rho_V^U \circ \pi_V$ except when ρ_V^U is not meaningful.

Ok, let us get back to introducing new axioms:

(7) **Bounded geodesic image theorem:**
If $V \subsetneq W$ and γ is a geodesic (segment, ray, bi-infinite) in $\mathcal{C}W$, then

$$\gamma \cap N_E(\rho_W^V) = \varnothing \Rightarrow \mathrm{diam}_{\mathcal{C}V}(\rho_V^W(\gamma)) \leqslant E.$$

($E > 0$ is a uniform constant for X that will be introduced later in Axiom (6)).

(5) **Finite complexity:**
There is an upper bound for the length of properly nested chains

$$V_1 \subsetneq V_2 \subsetneq \cdots \subsetneq V_n.$$

The maximum such $n \in \mathbf{N}$ is called the *complexity* of the HHS.

Remark 25.1. We will later show that if $n = 1$ then the HHS is hyperbolic. The converse is not meaningful because the extra data, e.g., the "maps", of an HHS is not completely determined by the space itself. In this sense it would be more appropriate to refer to HHS spaces as metric spaces endowed with an HHS structure.

The remaining three axioms are best understood as very weak versions of two key theorems, which we now state. Historically, these theorems were originally part of the axioms but they were reduced to the weakest conditions that imply them to make it easier to check if a space is an HHS.

Definition 25.2. For $A \geqslant 0$ define the *threshold function* $[\cdot]_A \colon \mathbf{R}_{\geqslant 0} \to \mathbf{R}_{\geqslant 0}$ given by

$$[x]_A := \begin{cases} x & \text{if } x \geqslant A, \\ 0 & \text{else} \end{cases}.$$

Intuitively, we think of threshold functions as throwing away terms that are small, or coarsely zero.

Definition 25.3. Write $A \asymp_{C,D} B$ to mean

$$C^{-1}A - D \leqslant B \leqslant CA + D.$$

We read $A \asymp_{C,D} B$ as saying that A and B are coarsely equal with multiplicative error C and additive error D.

Theorem 25.4 (Distance Formula)**.** *For any HHS $(X, \mathfrak{S})$ there exists $s_0 > 0$ such that for every threshold $s \geqslant s_0$ there exist constants $K, C > 0$ such that for every $x, y \in X$,*

$$d_X(x, y) \asymp_{K,C} \sum_{W \in \mathfrak{S}} [d_W(x, y)]_s,$$

where we recall that $d_W(x, y) := d_{\mathcal{C}W}(\pi_W(X), \pi_W(Y))$.

Remark 25.5. Keep in mind that the game we are playing is to reduce the study of an HHS to the study of a bunch of hyperbolic spaces, hence a distance formula like this one seems quite useful. The presence of threshold functions in this formula should be thought of as throwing away noise.

Remark 25.6. Roughly speaking, the distance formula says that "product map"

$$X \to \prod_{U \in \mathfrak{S}} \mathcal{C}U$$

is almost a quasi-isometric embedding when $\prod_{U \in \mathfrak{S}} \mathcal{C}U$ is endowed with the ℓ^1-metric.

Remark 25.7. It is NOT obvious at all that the right hand side of the distance formula is either positive or finite.

The last remark motivates the introduction of two axioms, one to guarantee the right hand side of the distance formula is (coarsely) positive, the other to take steps towards the finiteness of the corresponding sum. So, more axioms:

(9) **Uniqueness:**
For every $s > 0$ there exists $\theta = \theta(s) > 0$ such that if $x, y \in X$ with $d(x, y) > \theta$ then there exists a domain $V \in \mathfrak{S}$ such that $d_V(x, y) > s$.

Remark 25.8. Axiom (9) implies that the right hand side of the distance formula is positive. It also implies that the image of $X \to \prod_{U\in\mathfrak{S}} \mathcal{C}U$ has coarsely unique preimages, hence the name of the axiom.

The next axiom is a statement in the direction that the right hand side of the distance formula is finite. Recall the constants ξ and K_0 from the previous lecture.

(6) **Large Links:**
There exists a constant $E \geqslant \max\{\xi, K_0\}, \lambda > 1$ such that given $W \in \mathfrak{S}$, $x, x' \in X$, and $N := \lambda\lfloor d_W(x, x')\rfloor + \lambda$, there exist domains $\tau_1, \ldots, \tau_N \sqsubsetneq W$ such that if $\tau \sqsubsetneq W$ and $d_\tau(x, x') \geqslant E$ then $\tau \sqsubseteq \tau_i$ for some $i \in \{1, \ldots, N\}$.[3]

Remark 25.9. Roughly speaking, Axiom (6) says that big terms in the right hand side of the distance formula are organized into a few domains (at most N of them). Said another way, it guarantees that there are at most distance many domains into which you are nested if you are a domain with big term in the distance formula. This makes steps towards ensuring there are not many big terms in the right hand side of the distance formula.[4]

To motivate the last axiom of an HHS we state the following theorem.

Theorem 25.10 (Realization theorem, imprecise version)**.** *The image of*

$$X \to \prod_{U\in\mathfrak{S}} \mathrm{Im}(\pi_U)$$

is as big as possible given the restrictions of Axiom (4).

Remark 25.11. Notice that the restrictions given by Axiom (4) are particularly nice because they only involve pairs of domains.

Now to the statement of the last axiom:

(6) **Partial realization:**
There exists a constant $\alpha > 0$ such that if $\{V_i\}_{i\in I}$ is a set of pairwise orthogonal domains (necessarily finite) and $p_i \in \mathrm{Im}(\pi_{V_i})$ for every $i \in I$, then there exists $x \in X$ with the following properties:
- $d_{V_j}(x, p_j) \leqslant \alpha$ for every $j \in I$.

[3]The original formulation of the Large Links axiom also stipulates that "$d_W(x, \rho_W^{\tau_i}) \leqslant N$ for every $i \in \{1, \ldots, N\}$". Jacob Russell pointed out to us that this can be derived from the BGI and consistency axioms so does not need to be included as an axiom.

[4]Technically speaking, it is possible to show finiteness of the distance formula without Large Links, as in Remark 33.4 below. A more sophisticated and correct motivation for Axiom (6) is the "Passing up lemma" [BHS19, Lemma 2.5].

- For each $j \in I$ and each $V \in \mathfrak{S}$ with $V_j \subsetneq V$, $d_V(x, \rho_V^{V_j}) \leqslant \alpha$.
- If $W \pitchfork V_j$ for some $j \in I$ then $d_W(x, \rho_W^{V_j}) \leqslant \alpha$.

Remark 25.12. Roughly speaking, Axiom (6) guarantees that $X \to \prod_{i \in I} \mathrm{Im}(\pi_{V_i})$ is coarsely onto. The last two conditions of Axiom (6) should be thought of as some extra information that one can arrange for the realized $x \in X$.

Optional Exercise 30. Show that any space of infinite diameter cannot admit a HHS structure where all $\mathcal{C}U$ are uniformly bounded. (This is easy, but it's good to pause and realize which axiom makes this impossible.)

Optional Exercise 31. This is a long exercise on axiom (3b), which was omitted in class. It is sometimes called the "container axiom". This is its statement:

(3b) Suppose U is nested in T, and that there is at least one domain that is both nested in T and orthogonal to U. Then there exists some W strictly nested in T such that everything that is both nested in T and orthogonal to U is nested in W. (This W is sometimes called a "container" for everything nested in T and orthogonal to U.)

I find that statement hard to parse. It is intended to be a weak version of the following:

(3b-strong) Suppose U is nested in T, and that there is at least one domain that is both nested in T and orthogonal to U. Then there exists some W nested in T such that sometime is both nested in T and orthogonal to U if and only if it is nested in W. (This W might be called "the perp of U in T".)

There is now also another version of this axiom, called "Bounded pairwise orthogonality":

(3b-weak) There is an upper bound for the size of a set of pairwise orthogonal domains. (Anything in the index set $\mathfrak{S}$ is called a domain; for us domains are subsurfaces.)

Check that (3b-strong) $\implies$ (3b) $\implies$ (3b-weak). (The first implication is basically trivial, and, although you can probably do the second implication yourself, you can also find it as [BHS19, Lemma 2.1].)

Remark 25.13. Remark: While it is not literally true that (3b-weak) imples (3b), the appendix of [ABD21] shows that this is true up to a quite harmless and insignificant change to the HHS structure: so substituting (3b-weak) for (3b) essentially gives an equivalent definition of HHS. It is an open problem if one similarly has that (3b) essentially implies (3b-strong). This is purely a problem about the set $\mathfrak{S}$ with the nesting and orthogonality relations, so if you like combinatorics maybe you can solve it!

Optional Exercise 32. This is a long exercise on axiom (4c), which was omitted in class. Let's start with a quick summary:

- If V is nested in U, we get a map ρ_V^U.
- If U is nested in V, or if they are transverse, we get a coarse point ρ_V^U.
- If U and V are orthogonal, there is no ρ_V^U at all.

Axiom (4c) concerns the case when they are coarse points. There are a few statements you'd want to be true, for example:

(i) If U nested in V nested in W, then ρ_W^U is coarsely equal to ρ_W^V.

(ii) If U and V are both nested in W and U and V are orthogonal, then ρ^U_W is coarsely equal to ρ^V_W.

Part A: Check (i) and (ii) for subsurfaces. (Here ρ^U_W is equal to the set of boundary components of U that aren't peripheral for W.)

There is also a variant of (i):

(iii) Suppose U is nested in V. Suppose U is nested in W, and that V is transverse to W. Then ρ^U_W is coarsely equal to ρ^V_W.

Part B: Draw an example of three subsurfaces U, V, W with these relations to each other.

Part C: Verify (iii) for subsurfaces. (When V is transverse to W, then ρ^V_W is the the projection of the boundary of V to W (as in our proof of Behrstock).)

I will finally state (4c): Both (i) and (iii) hold.

What about (ii)? It's omitted because it's automatic!

Part D: Prove (ii) from the HHS axioms. (This appears as [BHS17a, Lemma 1.8] and, as part of a slightly more general statement, as [DHS17, Lemma 1.5].)

Optional Exercise 33. Suppose that an HHS X has only a single domain S, and that the map $\pi_S : X \to \mathcal{C}S$ is surjective. Prove that X is hyperbolic.

Optional Exercise 34. Show that the infinite sum in the distance formula is actually finite, for sufficiently large "thresholds".

Optional Exercise 35. Consider $X = \mathbb{R}$ (the real line). Let $\{S\}$ be the set of domains (there's just one of them), and let $\mathcal{C}S$ be the upper half plane model of the hyperbolic plane. Let $\pi_S : X \to \mathcal{C}S$ be defined by $\pi_S(x) = x + i$ (so the real line maps to the boundary of a horoball).

How that this data satisfies all the axioms of an HHS, except that $\pi_S(X)$ is not quasi-convex. Show that the distance formula is false for this example.

Optional Exercise 36. Let $X = [0, \infty)$. Let $I_0 = [0, 1]$, and for all $k > 0$ let I_k be the interval of length k starting $1/2^k$ to the right of the endpoint of I_{k-1}. So all these intervals are disjoint and they have unbounded lengths.

Let $\mathcal{C}S = [0, 1)$, and let $\pi_S : X \to \mathcal{C}S$ be the map that contracts each I_k to a point. (Really a point, not just something bounded diameter.)

For each i in $\{0, 1, 2, \ldots\}$ let $\mathcal{C}i = I_i$, and let π_i be the closest point projection to I_i. All these i are nested in S and transverse to each other.

Define ρ^i_S be the point in $[0, 1)$ obtained from collapsing I_i. Define ρ^i_j be the closest point projection of I_i to I_j. (This is just one of the endpoints.)

Show that this data satisfies all the axioms of an HHS except for the Large Links axiom. (Exercise credit: Jacob Russell.)

Optional Exercise 37. Consider a tree X. The edges can have different lengths if you'd like. For each edge e, define $\pi_e : T \to e$ to be the closest point projection. Define $\mathfrak{S} = \{S\} \cup \{\text{edges of } X\}$. Define $\mathcal{C}S$ to be a point, and $\mathcal{C}e = e$ for all edges e. Define different edges to be transverse to each other; and the only nesting will be that all edges are nested in S.

Show that this satisfies all the axioms of an HHS except possibly large links and uniqueness, and that the distance formula holds with no error when the threshold is 0. (If all edges have length 1, actually large links will vacuously hold, but uniqueness will always fail. If there are only finitely many edges of length less than any given constant, then actually uniqueness will hold, but large links will fail unless the tree has an extremely special structure like having all but one vertex be a leaf.)

This example is not too hard, and is helpful since it gives excellent intuition for a lot of arguments with HHSes (even though the given structure doesn't satisfy all the axioms).

If you're an expert on CAT(0) cube complexes: I suspect you can do something similar for any CAT(0) cube complex, where now the domains would be S plus certain equivalence classes of edges, and now there is orthogonality.

Optional Exercise 38. For any valid index set, prove that any infinite set of domains contains an infinite subset consisting of pairwise transverse domains.

Remark 25.14. A slightly shorter but equivalent version of the axioms is provided by [BHS19, Proposition 1.11], which has the feature that the "downward" maps $\rho^V_U : \mathcal{C}U \to \mathcal{C}V$ with $U \subsetneq V$ do not need to be defined.

26. Toy HHSes (03/14, CK, TY)

We will now consider some toy examples of hierarchically hyperbolic spaces.

Example 26.1. Suppose X is a δ-hyperbolic quasi-geodesic metric space. Let $\mathfrak{S} := \{S\}$ and $\mathcal{C}S := X$. $\mathcal{C}S$ is clearly hyperbolic. Define the projection $\pi_S : X \to \mathcal{C}S$ by $\pi_S := id$. It is straightforward to see that this is an HHS structure on X. We will call this the trivial HHS structure on a hyperbolic space X. We often abbreviate this to saying that $(X, \mathfrak{S})$ is an HHS.

Example 26.2. Consider $X = X_1 \times X_2$ with X_1, X_2 both δ-hyperbolic. Let $\mathfrak{S} := \{S, A_1, A_2\}$ (these are formal variables) with the nesting relation given by $A_i \sqsubseteq S$. Also, $A_1 \perp A_2$, so there is no transversality.

Define $\mathcal{C}A_i := X_i$ for $i = 1, 2$ and $\mathcal{C}S := \{*\}$. Let $\pi_{A_i} : X \to \mathcal{C}A_i$ be the projection to the i^{th} coordinate for $i = 1, 2$. Let $\pi_S : X \to \mathcal{C}S$ be the constant map to the only point in $\mathcal{C}S$.

One can check that this is an HHS structure on X. The maps $\rho^S_{A_i}$ can be defined arbitrarily (say, $\rho^S_{A_i}(*) := \mathcal{C}A_i$ for $i = 1, 2$) and they will still satisfy all axioms. Partial realization is satisfied by choosing $x = (p_1, p_2)$, given $p_i \in A_i$ with $i = 1, 2$.

Example 26.3. This is the main example for this lecture and models a lot of the phenomena in HHS structures. Consider a δ-hyperbolic geodesic space X and pick C big depending on δ. How big C needs to be will be clarified as we go. To begin with, for large enough C, we know by the bounded geodesic image theorem (Theorem 23.1) that if γ is a geodesic and $A \subset X$ is convex, then there is some fixed $\xi = \xi(\delta)$ so that $\gamma \cap N_{\frac{C}{2}-1}(A) = \varnothing$ implies that $\operatorname{diam}(\pi_A(\gamma)) \leqslant \xi$.

Now let $\{A\}_{A \in \mathcal{A}}$ be a collection of convex (or quasi-convex) subsets of X, all at least at a distance C apart from each other. Let $\mathfrak{S} := \{S\} \sqcup A$, where S is an artificial index. Let $\mathcal{C}A := A$ for all $A \in \mathcal{A}$ and let $\mathcal{C}S := Cone_{\mathcal{A}}(X)$, the electrification of X along $\mathcal{A}$. Define the relation on $\mathfrak{S}$ by having $A \sqsubseteq S$ for all $A \in \mathcal{A}$, $A \pitchfork B$ for $A \neq B$ and $A, B \in \mathcal{A}$, and no orthogonality.

Define the projection $\pi_S : X \to \mathcal{C}S$ as the inclusion of X into $Cone_{\mathcal{A}}(X) = \mathcal{C}S$. Define $\pi_A : X \to 2^{\mathcal{C}A}$ to be the closest point projection. The image of a point has a uniformly bounded diameter by our result on closest point projections from January 26, Lemma 9.3.

If $A \neq B$, then define $\rho_A^B := \pi_A(B)$. Remember that $A \pitchfork B$.

Let $\rho_S^A \in \mathcal{C}S$ be the cone point of A. Define $\rho_A^S : \mathcal{C}S \to 2^{\mathcal{C}A}$ by:

- $\rho_A^S(x) := \pi_A(x)$, if $x \in X \subset \mathcal{C}S$.
- $\rho_A^S(x) := \rho_A^B$ if x is in the cone of $B \neq A$.
- $\rho_A^S(x) := A$ if x is in the cone of A. It doesn't actually matter what we do in this case. This is because $d_{Haus}(x, \rho_S^A) \leqslant 2$ by going through the cone point, and so the minimum in axiom $4(b)$ enters the trivial case for $K_0 > 2$. Additionally, the condition of the bounded geodesic image axiom is not affected by this choice, since for $E > 1$, any geodesic avoiding $N_E(\rho_S^A)$ also avoids x.

We now briefly sketch why this structure, abbreviated by $(X, \mathfrak{S})$, satisfies all the HHS axioms, defering three of them to next class.

1. **(Projections)** By Proposition 7.6, $Cone_{\mathcal{A}}(X)$ is hyperbolic. A is a (quasi-) convex subset of X, so it is also hyperbolic for any $A \in \mathcal{A}$. The map π_S is an inclusion and thus a contraction of quasi-geodesic metric spaces, so it is coarsely Lipschitz.

   The maps π_A are closest point projections, so by Lemma 9.3 again, the Hausdorff distance $d_{Haus}(\pi_A(x), \pi_A(y)) \leqslant D$ for some $D = D(\delta)$ whenever $d(x, y) \leqslant 1$ for $x, y \in X$. This implies that the map π_A is coarsely Lipshitz with respect to d_X and d_{Haus}.

2. **(Nesting)** The nesting order is a reflexive partial order and $\operatorname{diam}(\rho_S^A) \leqslant 2$ by going through the cone point, as required.

3. **(Orthogonality)** This holds trivially since there is no orthogonality.

4. **(Transversality and Consistency)** 4(a) follows from the discussion at the end of the lecture on Behrstock's inequality, along with the fact that the projections

ρ^B_A have uniformly bounded diameters by Lemma 9.9, using the (quasi-) convexity of A and B. 4(b) follows from the comments above in the definition of ρ^S_A.

5. **(Finite Complexity)** This system has complexity 2, by the design of the nesting relation.

8. **(Partial Realization)** The only set of pairwise orthogonal domains is a singleton, $\{V\}$.

   If $V = A$ for some $A \in \mathcal{A}$, then let $x = p \in X$ given any $p \in V = A \subset X$. Conveniently, $d(x,p) = 0$. This means that $d_S(x, \rho^A_S) \leqslant 1$ since ρ^A_S is a cone point. For any $B \pitchfork A$, $d_B(x, \rho^A_B)$ is uniformly bounded too, since $\pi_B(x) \in \rho^A_B$, which already has a uniformly bounded diameter.

   If $V = S$, then let $x = p$ if $p \in X \subset \mathcal{C}S$; otherwise let x be the cone point of the cone that p is in. Everything is satisfied vacuously here after noting that $d(x,p) \leqslant 1$.

27. Toy HHSes (03/18, SC, AW)

We continue with our sketch of how the structure defined in Example 26.3 satisfies the HHS axioms.

7. **(Bounded Geodesic Image)** The only indices with $V \subsetneq W$ are $W = S$ and $V = A$ for some $A \in \mathcal{A}$.

   Recall that, coming from our use of the Geodesic Guessing Lemma in the proof of Proposition 7.6, there is some $D = D(\delta)$ such that every geodesic γ in $\mathcal{C}S = Cone_{\mathcal{A}}(X)$ is distance at most D away from a geodesic $\hat{\gamma}$ in $X \subset \mathcal{C}S$.

   Assume that $d(\gamma, \rho^A_S)$ is big. This implies that $d(\hat{\gamma}, A)$ is also big as follows: $d_X(\hat{\gamma}, A)$ small would imply $d_{\mathcal{C}S}(\hat{\gamma}, \rho^A_S)$ small, which would imply $d_{\mathcal{C}S}(\gamma, \rho^A_S)$ small since γ and $\hat{\gamma}$ are close to each other. (Here we have added subscripts to emphasize the space the distance is being taken in.)

   So, by our results on closest point projections in hyperbolic spaces, namely Proposition 9.9, we have that $\operatorname{diam}(\pi_A(\hat{\gamma}))$ is bounded. Since ρ^S_A is *almost* defined to be π_A, one can show, with a bit of work, that $\rho^S_A(\gamma)$ has bounded diameter as well. For this, one should keep in mind that all $\pi_A(B) = \rho^B_A$ have uniformly bounded diameter, as was already used in the definition of ρ^S_A.

Before moving on to the large links and the uniqueness axioms, we note the following result, which is almost the same as the distance formula if we move the second term on the right to the left.

Lemma 27.1. *Given $\delta \geqslant 0$, $\exists_{C,R,D \geqslant 0}$ such that if X is δ-hyperbolic and $\{A\}_{A \in \mathcal{A}}$ are C-separated convex sets and $Y = Cone_{\mathcal{A}}X$, then $\forall x, y \in X$, if γ is a geodesic from x to y in X,*

$$d_Y(x,y) \asymp d_X(x,y) - [\textit{time } \gamma \textit{ spends in } \bigcup N_R(A)],$$

and if $\gamma \cap N_R(A) = \phi$, then $\operatorname{diam} \pi_A(\gamma) \leqslant D$.

Proof sketch. The fact that there exists a D for which the final statement holds for any R large enough is just a restatement of Lemma 9.9. This does not involve C.

A upper bound on $d_Y(x, y)$ of the desired form follows just from the definition of electrification, as long as C is significantly larger than R.

The key point is getting a lower bound on $d_Y(x, y)$. In general, it isn't even always clear if electrifications are finite diameter, but here the assumption that the different A are far apart is very powerful. The main idea in is to note that that geodesics in Y can only go through the cone points of those A that γ comes close to, assuming C is large compared to the Hausdorff distance bound from the geodesic guessing lemma (Figure 99).

In particular, for any R large enough, we can assume that C is large enough so that if γ doesn't intersect $N_R(A)$ then the geodesic from x to y doesn't go through the cone point of A; see Figure 99. We can assume C is much larger than R if desired, so even the $N_R(A)$ are far apart.

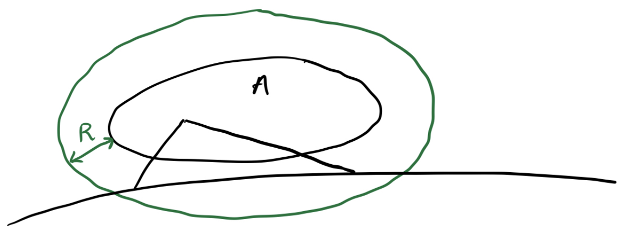

FIGURE 99. For a geodesic in Y to go through the cone point of A, the corresponding geodesic in X has to come close to A

A geodesic in the electrification consists of geodesics from one A to another, followed by hops through cone points. The last key point in the sketch is that if there is a geodesic segment of length ℓ from the boundary of $N_R(A)$ to the boundary $N_R(B)$ whose interior is disjoint from $N_R(A) \cup N_R(B)$ then the distance between A and B is at least ℓ minus a constant depending in R, assuming R is large enough. Indeed, the bounded geodesic image theorem says that the projection of the segment to A is close to the projection of B to A, and vice versa. Thus, as in Figure 100, the triangle inequality gives a lower bound for the distance from A to B.

We leave it as an exercise to combine these observations into a proof of the lemma. □

6. (Large Links) We need to show that for any two points x and x' in X, the number of A for which $d_A(x, x')$ can be large is bounded by $d_{\mathcal{CS}}(x, x')$. Lemma 27.1 tells us that $d_A(x, x')$ can be large only when the geodesic joining x and x' intersects the R-neighbourhood of A. Now, when $C \gg R$, the geodesic has to travel at least $C - 2R$ between the R-neighbourhoods of any two such A. This $C - 2R$ contributes to $d_{\mathcal{CS}}(x, x')$. Thus, $d_{\mathcal{CS}}(x, x')$ bounds the number of A for which $d_A(x, x')$ is large.

9. (Uniqueness) We use the same ideas as the large links axiom. For any two points x and y in X far apart, we need to show that there is at least one domain

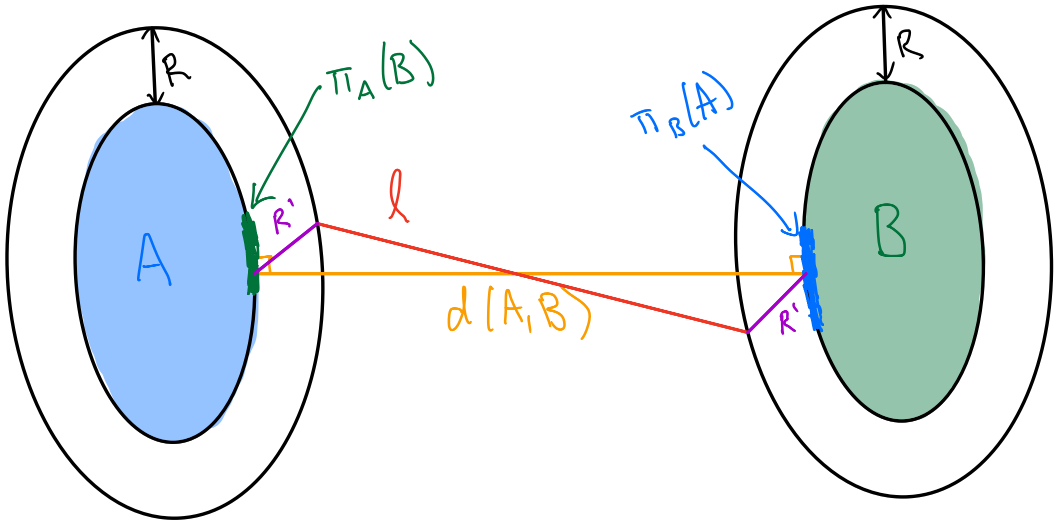

FIGURE 100. The triangle inequality gives $d(A,B) \geqslant \ell - 2R'$. An upper bound R' for the length of the purple geodesic segments can be obtained since the projections of the red geodesic segment are close to the projection of each set onto the other, and because those projections have bounded diameter and contain the endpoints of the orange minimal length path from A to B.

where the projections of x and y are far apart. If $d(x,y)$ is large but $d_{\mathcal{CS}}(x,y)$ is not, then the argument from the large links axiom tells us that the geodesic joining x and y can only go near $\lfloor d_{\mathcal{CS}}(x,y) \rfloor$ many A. Lemma 27.1 then tells us that the geodesic joining x and y must spend a large time near some A. One can then show that the projections of x and y to this A are far apart (Figure 101).

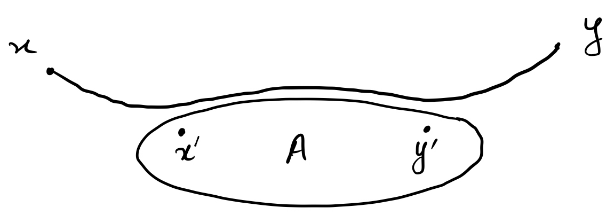

FIGURE 101. When the geodesic from x to y spends a large time near A, the projections x' and y' onto A of x and y respectively are far apart.

We now sketch a few other examples in lesser detail.

Example 27.2. We again look at a hyperbolic space, an infinite tree X, and consider infinite diameter convex subsets A, B whose intersection D is also convex, as shown in Figure 102. This is in contrast to the previous example, where the convex subsets were well separated. We let $\mathfrak{S}$ be the set of indices $\{S, A, B, D\}$ and define $\mathcal{C}S$ to be a point and $\mathcal{C}D$ to be $D \subset X$. Since the projections from A to B and vice versa are not independent, A and B cannot be orthogonal, and have to be transverse. Our definitions of $\mathcal{C}A$ and $\mathcal{C}B$ are then motivated by the properties that we want ρ^A_B and ρ^B_A to satisfy. A natural definition for ρ^A_B would be the closest point projection of A onto B, but this is not bounded in B, so we cone off $D \subset B$, and we look at the projection in the coned off space. Thus, we define $\mathcal{C}A := Cone_D A$ and $\mathcal{C}B := Cone_D B$. In general, whenever we want an image to be a coarse point, we cone off the relevant subset.

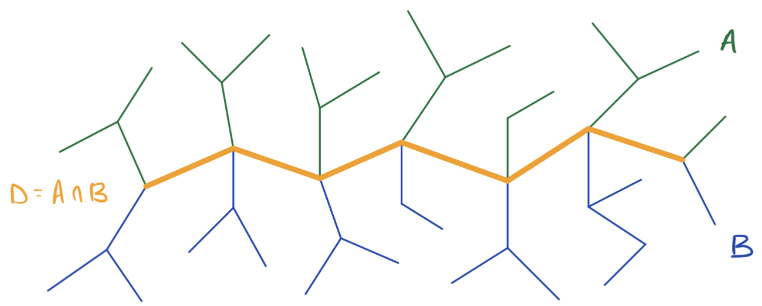

FIGURE 102. The tree X, and subsets A, B and D

For the most general result on HHS structures on spaces X that are already hyperbolic, see [Spr17].

Example 27.3. Let X_1 and X_2 be hyperbolic, and let X be the hyperbolic cone on $X_1 \times X_2$ minus the cone point. The hyperbolic cone on $X_1 \times X_2$ is hyperbolic, but one can show that it fails to be hyperbolic after removing the cone point. Let $\mathfrak{S}$ be $\{S, A_1, A_2\}$, and define $\mathcal{C}S$ to be the hyperbolic cone on $X_1 \times X_2$ and $\mathcal{C}A_i$ to be X_i. $\rho^{A_i}_S$ is the cone point of the hyperbolic cone, and $\rho^S_{A_i}$ is the projection onto the X_i coordinate of a point. The domains A_1 and A_2 are orthogonal, so we do not have sets $\rho^{A_1}_{A_2}$ and $\rho^{A_1}_{A_2}$.

We now look at some less typical examples.

Example 27.4. Let $\{(X_i, \mathfrak{S}_i)\}_{i\in I}$ be HHSes. Then so is $(\Pi_{i\in I} X_i, \bigcup_{i\in I} \mathfrak{S}_i \cup S_{new})$, with $\mathcal{C}S_{new}$ being a point. In this structure, each domain from $\mathfrak{S}_i$ is orthogonal to a domain from $\mathfrak{S}_j$, when $j \neq i$.

Example 27.5. Let X_1 and X_2 be hyperbolic. Then there are several HH structures on $Y = X_1 \bigvee X_2$, the wedge of X_1 and X_2 (Figure 103).

- (The trivial HHS) Since the wedge of two hyperbolic spaces is hyperbolic, we can have the trivial HH structure given by one domain S with $\mathcal{C}S = Y$.
- (X_1 on top) In this structure, we have $\mathfrak{S} = \{S, A\}$ with $\mathcal{C}S = X_1$, $\mathcal{C}A = X_2$ and $A \sqsubseteq S$. ρ^A_S is $\{p\}$ and ρ^S_A is the constant map to p.
- (Side by side) In this structure, we have $\mathfrak{S} = \{S, A, B\}$ with $\mathcal{C}S$ a point, $\mathcal{C}A = X_1$, $\mathcal{C}B = X_2$ and $A \pitchfork S$. The relevant projection maps are defined similarly.

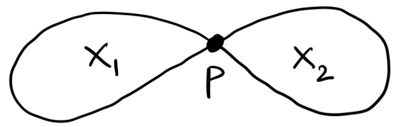

FIGURE 103. X_1 and X_2 wedged along a point p

Optional Exercise 39. Prove the distance formula for the HHS structure on a hyperbolic space coming from well-separated convex subsets.

Optional Exercise 40. Prove that the universal cover of a torus wedge a circle is an HHS. (This is a special case of cube complex or Right Angled Artin Group examples.)

Optional Exercise 41. Fix n intervals $[a_i, b_i]$ of finite length. Fix also a partial order $\ll$ on $\{1, 2, ..., n\}$.

Define X to be the subset of those points $(x_1, ..., x_n) \in \prod_i [a_i, b_i]$ for which $i \ll j$ implies either $x_i = a_i$ or $x_j = b_j$.

Show X has an "especially nice" HHS structure, in which each hyperbolic space is an interval or a point. Here "especially nice" should mean, ex, $\xi = 0, \kappa_0 = 0$, etc, so these constants in the axioms have their optimal, most restrictive values. (Feel free to ignore Large Links and Uniqueness. They don't have especially nice versions here, but in their usual form they are automatic if sort of dumb here because everything is bounded.)

Note that $A = (a_1, ..., a_n)$ and $B = (b_1, ..., b_n)$ are both in X. I think of $i \ll j$ as saying that, when going from A to B, the j coordinate must change before the i coordinate.

This example is important, since arbitrary HHSes have subsets that look a lot like the above. It seems related to the perspective on CAT(0) cube complexes in [AOS12, Section 2].

Optional Exercise 42. Let $\mathfrak{S}$ be a set with orthogonality and nesting relations as in the HHS axioms. For each $U \in \mathfrak{S}$, pick hyperbolic space $\mathcal{C}U$, with a base point p_U.

Define X to be the set of tuples $(x_U)_{U \in \mathfrak{S}}$ for which the set of U for which $x_U \neq p_U$ is a set of pairwise orthogonal domains.

An alternate description of this space is that one starts with the wedge of the $(\mathcal{C}U, p_U)$, and then glues in products as dictated by the orthogonality relation.

Show that $(X, \mathfrak{S})$ is an HHS where the orthogonality and nesting relations are as given, and the π_U are just the coordinate projections, and all ρ^U_V are either the point p_V or the map constantly equal to p_V.

Like the wedge example from class, this example is potentially misleading, since it has the strange feature that the space $\mathcal{C}U$ and the maps π_U don't uniquely determine the nesting relation, which wasn't even used in the construction of X. But it is handy to disprove overly optimistic conjectures, since it shows since any conceivably set with orthogonality and nesting relations actually appears for an HHS with all hyperbolic spaces unbounded.

28. The Farey Graph (03/21, TY, SK)

Our goal is to better understand the curve complex of the punctured torus $\mathcal{C}S_{1,1}$ by showing that it is isomorphic to the Farey graph. For the convenience of the reader, we will refer to *non-peripheral simple closed curves* as curves, since these are the only kinds of curves we care about in this section.

Lemma 28.1. *The vertices of $\mathcal{C}S_{1,1}$, i.e., the set of curves of the torus up to isotopy, are in bijection with $\mathbb{Q} \cup \{\infty\}$.*

Proof sketch. We construct a map from $\mathbb{Q} \cup \{\infty\}$ to the set of vertices. Think of $S_{1,1}$ as $(\mathbb{R}^2 - \mathbb{Z}^2)/\mathbb{Z}^2$, where the quotient comes from translation of the lattice. Given $q \in \mathbb{Z} \cup \{\infty\}$, take the line $x = qy$ in $\mathbb{R}^2$, so it has slope q^{-1}. This gives a curve on $\mathbb{R}^2/\mathbb{Z}^2$: it closes up because of the rational slope, and it's simple because it can't intersect itself (we get either parallel lines or complete overlaps). Also note that its translates that don't go through the punctures at $\mathbb{Z}^2$ are isotopic to each other, even on the punctured torus, via translations. Thus, we get a map $\mathbb{Q} \cup \{\infty\}$ to the set of vertices of $\mathcal{C}S_{1,1}$. One can take flat geodesic representatives to see that this map is surjective. □

Now, we want to understand when these curves intersect. We have an action of $\mathrm{SL}(2, \mathbb{Z})$ on $S_{1,1}$ coming from its action on $\mathbb{R}^2$, which preserves lattice points:

$$\begin{pmatrix} p & q \\ r & s \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} px + qy \\ rx + sy \end{pmatrix}.$$

This gives an action by fractional linear transformations on inverse slopes:

$$\frac{x}{y} \mapsto \frac{px + qy}{rx + sy} = \frac{p\frac{x}{y} + q}{r\frac{x}{y} + s}.$$

Remark 28.2. Note that x, y have no common factor if and only if $px + qy, rx + sy$ have no common factor; the forward direction is clear, and the backward direction can be seen via the inverse matrix, which also has integer coefficients. Thus, the action of $\mathrm{SL}(2, \mathbb{Z})$ on inverse slopes preserves fractions being in lowest terms. We consider a fraction being in lowest terms if no integer besides $0, 1$ divide both the numerator and denominator, so $\frac{1}{1}, \frac{0}{1}, \frac{1}{0}$ are in lowest terms.

We now want to determine the intersection number between the curves; we denote the curve by its inverse slope $\frac{a}{b}$. Note that two such curves do not form bigons, and so their intersection number is given by the number of points in their intersection. In addition, since the action of $\mathrm{SL}(2, \mathbb{Z})$ on $\mathbb{R}$ is by homeomorphisms, it preserves intersection number.

First, note that the curve $\frac{a}{b}$ goes around vertically $|b|$ times, and horizontally $|a|$ times. Then, since $\frac{1}{0}$ represents the horizontal curve, we see that

$$i\left(\frac{1}{0}, \frac{a}{b}\right) = b = \left|\det \begin{pmatrix} 1 & a \\ 0 & b \end{pmatrix}\right|.$$

Thus, by our comments from above, we have that if $\frac{a}{b}, \frac{c}{d}$ are in lowest terms, then

$$i\left(\frac{a}{b}, \frac{c}{d}\right) = \left|\det\begin{pmatrix} a & c \\ b & d \end{pmatrix}\right|.$$

Definition 28.3. The curve complex of the punctured torus $\mathcal{C}S_{1,1}$ is defined to have a vertex for curve, and an edge between two such curves if they have intersection number 1.

Definition 28.4. The **Farey graph** $\mathcal{F}$ has vertices $\mathbb{Q} \cup \{\infty\}$ and an edge between $\frac{a}{b}, \frac{c}{d}$ (written in lowest terms) if

$$\left|\det\begin{pmatrix} a & c \\ b & d \end{pmatrix}\right| = 1.$$

By our work above, we have shown that $\mathcal{C}S_{1,1} \cong \mathcal{F}$. We now introduce a model for $\mathcal{F}$ that will allow us to show that $\mathcal{C}S_{1,1}$ is hyperbolic and has infinite diameter. In particular, we'll show that $\mathcal{C}S_{1,1}$ is a **quasi-tree**, i.e., quasi-isometric to a tree.

One can draw $\mathcal{F}$ in the upper half plane $\mathbb{H}$ by drawing semicircles orthogonal to the real line for edges between points in $\mathbb{Q}$, and drawing vertical lines for edges between ∞ and rational numbers, as seen in Figure 104. For example, we have edges between $\frac{0}{1}, \frac{1}{0}$ (so a vertical edge coming from $\frac{0}{1} = 0$) and between $\frac{1}{2}, \frac{1}{1}$ (so a semicircle edge between the two numbers) because of the following calculations:

$$\left|\det\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\right| = 1, \qquad \left|\det\begin{pmatrix} 1 & 1 \\ 2 & 1 \end{pmatrix}\right| = 1.$$

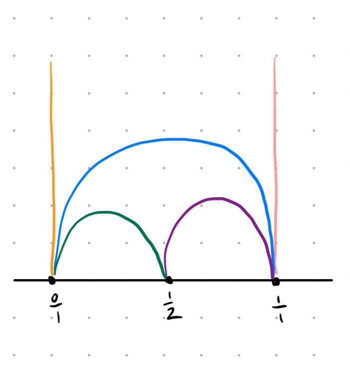

FIGURE 104. Farey graph $\mathcal{F}$ in $\mathbb{H}$

Note that this model for the Farey graph distorts the lengths of edges, as all edge lengths are actually 1.

Remark 28.5. If $\frac{a}{b} > \frac{c}{d}$, then $\frac{a}{b} - \frac{c}{d} = \frac{ad-bc}{bd}$, and so we have an edge joining $\frac{a}{b}, \frac{c}{d}$ if and only if

$$\begin{pmatrix} a & c \\ b & d \end{pmatrix} \in \mathrm{SL}(2, \mathbb{Z}).$$

Note also that the action of $\mathrm{SL}(2,\mathbb{Z})$ on itself via left multiplication is transitive, and so there's only one $\mathrm{SL}(2,\mathbb{Z})$-orbit of edges.

Lemma 28.6. *There are no crossing edges in our $\mathbb{H}$ picture of $\mathcal{F}$.*

Proof. By the transitivity of the $\mathrm{SL}(2,\mathbb{Z})$ action on the edges, which preserves (non)-crossing of edges, it suffices to show that no edge crosses the edge from 0 to ∞, i.e., the situation in Figure 105 is impossible.

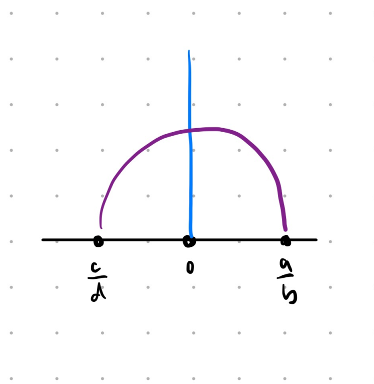

FIGURE 105. Impossible edge crossing of $\mathcal{F}$

If there were such an edge, it would be between $\frac{c}{d}$ and $\frac{a}{b}$, where $a, b, d > 0$ and $c < 0$. Then, $ad, -bc > 0$ and are both integers, and so $ad - bc \geqslant 2$. Thus, there in fact cannot be an edge between $\frac{a}{b}, \frac{c}{d}$. ☐

From our picture of $\mathcal{F}$, we get an ideal triangulation of $\mathbb{H}$, with each edge of the graph participating in exactly 2 triangles. We get this just from the fact that edges can't cross, $SL(2,\mathbb{Z})$ acts transitively on edges, and we've exhibited (in Figure 104) an edge above that participates in two triangles.

Lemma 28.7. *$\mathcal{F} \cong \mathcal{C}S_{1,1}$ is a quasi-tree. Therefore, $\mathcal{C}S_{1,1}$ is hyperbolic.*

Keep in mind that all edges have length 1 in $\mathcal{C}S_{1,1} \cong \mathcal{F}$. One typically shows that a space is a quasi-tree by applying the bottleneck criterion [Man05, Theorem 4.6]. One can compare that criterion to the fact that, in an actual tree, any path between points x and y must pass through every point on the geodesic from x to y.

Proof idea. By the bottleneck criterion, it suffices to show that for vertices v, w far apart, one can find points $p_{v,w}$ such that any path from v to w must pass through a bounded neighbourhood of $p_{v,w}$, the size of which is bounded independent of v and w. Suppose v, w are vertices of $\mathcal{F}$ that are not connected. Then by the planar structure of the graph in $\mathbb{H}$, any path from v, w must pass through the edges with endpoints between v, w, either through the edges' endpoints or along their interiors.

Thus, we have a bottleneck consisting of many diameter 1 sets (i.e., edges), and so it's a quasi-tree. ☐

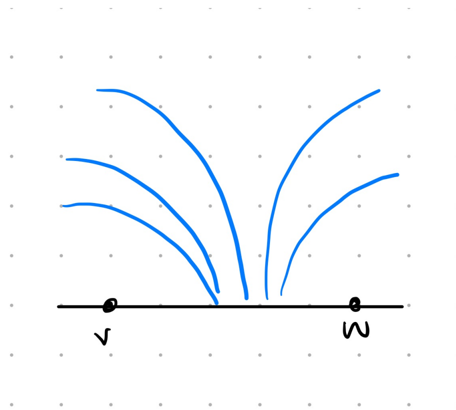

FIGURE 106. Bottleneck consisting of edges between v, w

Remark 28.8. The Teichmüller space $\mathcal{T}_{1,1}$ is a single Teichmüller disc. Exercise 25 asks you to show that, in general, any Teichmüller disc gives rise to a quasi-tree in the curve complex.

Lemma 28.9. *$\mathcal{C}S_{1,1}$ has infinite diameter.*

Proof idea. Given a semicircle edge, one can always find another semicircle edge contained in it, as seen in Figure 107, because $\mathcal{F}$ is an ideal triangulation. Any point on these two edges are at least distance 1 away from each other. We can continue this process to see that the diameter of $\mathcal{C}S_{1,1}$ has no bound, and therefore the graph has infinite diameter. □

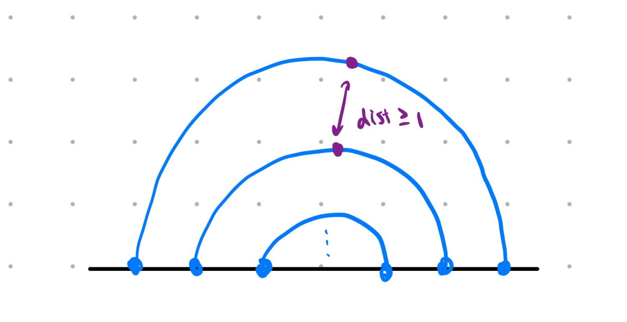

FIGURE 107. $\mathcal{F}$ has infinite diameter

29. Annular subsurfaces (03/23, BZ, PA)

Besides $S_{1,1}$, the other exceptional surfaces whose curve complexes we must define are $S_{0,2}$ (the annulus), $S_{0,3}$, and $S_{0,4}$. On $S_{0,2}$ and $S_{0,3}$, all simple closed curves are peripheral. We will give a special definition for $\mathcal{C}S_{0,2}$ that detects the twisting of curves on a surface, and we will simply define $\mathcal{C}S_{0,3}$ to be a point. We will first discuss $\mathcal{C}S_{0,4}$.

Consider the branched double-cover $\pi : \mathbb{R}^2/\mathbb{Z}^2 \to S_{0,4}$ that is the quotient by the involution $(x, y) \mapsto (-x, -y)$ on $\mathbb{T} := \mathbb{R}^2/\mathbb{Z}^2$. Observe that this involution fixes the four points $P := \{(0,0), (\frac{1}{2}, 0), (0, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2})\}$, and has fundamental domain $(0,1) \times (0, \frac{1}{2})$; see Figure 108.

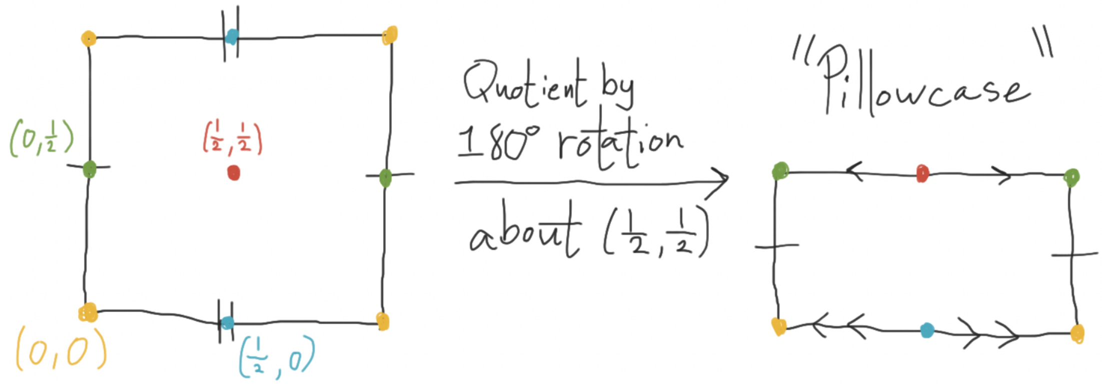

FIGURE 108. $S_{0,4}$ as the quotient of $\mathbb{R}^2/\mathbb{Z}^2$ by $(x, y) \mapsto (-x, -y)$

Definition 29.1. The curve complex $\mathcal{C}S_{0,4}$ of the sphere with four punctures has a vertex for each non-peripheral simple closed curve, and an edge from α to β if α and β have intersection number 2.

One can check that the map π induces an isomorphism $\mathcal{C}S_{1,1} \xrightarrow{\sim} \mathcal{C}S_{0,4}$. For instance, the condition that the curves in $\mathcal{C}S_{0,4}$ are non-peripheral and simple imply that they disconnect the sphere into two components each of which has two punctures. This implies that each curve in $\mathcal{C}S_{0,4}$ has two disjoint simple preimages on $\mathbb{T} - P$. If two curves in $\mathcal{C}S_{0,4}$ intersect twice, then any component of the preimage of one on $\mathbb{T} - P$ intersects any component of the preimage of the other exactly once.

Lastly, we give the definitions for two different versions of the curve complex of an annular subsurface $S_{0,2} \cong A \subset S$. Let $p : S_A \to S$ be the covering space associated with the subgroup $\pi_1(A) < \pi_1(S)$. Note that S_A is itself an annulus. Fix a hyperbolic metric on S, and consider the induced hyperbolic metric on S_A. We can compactify S_A to obtain a closed annulus $\overline{S}_A$, for example via the Gromov boundary construction. Equivalently, if $S_A = \mathbb{D}/\langle\gamma\rangle$ where $\mathbb{D}$ is the hyperbolic disk and $\gamma \in \mathrm{PSL}(2,\mathbb{R})$ has fixed points F on the boundary (so $|F| = 2$), then set $S_A := (\overline{\mathbb{D}} - F)/\langle\gamma\rangle$.

Definition 29.2. The curve complex $\mathcal{C}_{\mathrm{MCG}}A$ has a vertex for each arc (up to isotopy rel endpoints) from one boundary circle of $\overline{S}_A$ to the other. It has an edge for each pair of disjoint such arcs.

While there are uncountably many vertices of $\mathcal{C}_{\mathrm{MCG}}A$, each with uncountably many edges, the next lemma shows that $\mathcal{C}_{\mathrm{MCG}}A$ is not as pathological as this makes it sound.

Lemma 29.3. *$\mathcal{C}_{\mathrm{MCG}}A$ is quasi-isometric to $\mathbb{R}$.*

Proof sketch. Let x and y be points on each boundary circle, respectively, of $\overline{S}_A$. Define a map $\varphi : \mathbb{Z} \to \mathcal{C}_{\mathrm{MCG}}A$ as follows. Let $\varphi(0)$ be an arbitrary arc from x to y, and let

$\alpha \subset S_A$ be a simple closed curve such that $p(\alpha)$ is the core curve of $A \subset S$. We define $\varphi(n) = D_\alpha(\varphi(0))$, where D_α is the Dehn twist about α. For each $n \in \mathbb{Z}$, observe that $\varphi(n)$ is disjoint from $\varphi(n+1)$; see Figure 109. It follows that φ is a quasi-isometry. $\square$

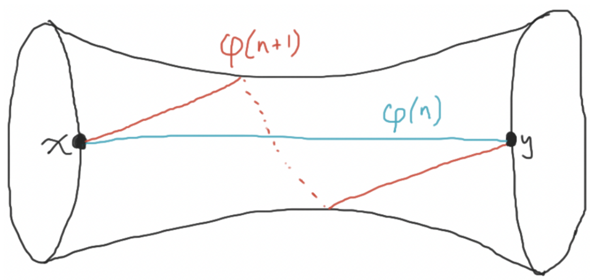

FIGURE 109. An arc is disjoint from its Dehn twist. One can obtain this picture with n arbitrary by taking a Dehn twist of the same picture with $n = 0$.

Definition 29.4. The curve complex $\mathcal{C}_{\text{Teich}}A$ is defined to be the set $\{x+iy \in \mathbb{C} \mid y \geqslant 1\}$, endowed with the hyperbolic metric.

This is a horoball in the hyperbolic plane. If we fix a coarse identification $\mathcal{C}_{\text{MCG}}A \cong \mathbb{R}$, we get a map

$$\begin{aligned} \mathcal{C}_{\text{MCG}}A &\to \mathcal{C}_{\text{Teich}}A \\ x &\mapsto x + i. \end{aligned}$$

This map is Lipschitz, but is not a quasi-isometric embedding.

30. Annular subsurfaces and Bers markings (03/25, KS, SC)

Definition 30.1. There is a map $\rho_A^S : \mathcal{C}S \to 2^{\mathcal{C}_{MCG}A}$, which can be composed with $\mathcal{C}_{MCG}A \to \mathcal{C}_{Teich}A$ to get a map $\rho_A^S : \mathcal{C}S \to \mathcal{C}_{Teich}A$.

Recall that we have fixed a hyperbolic structure on S. For any curve $\alpha \in \mathcal{C}S$, let us take its geodesic representative and consider all lifts to SA. Their closures give arcs in $\overline{SA}$. Define $\rho_A^S(\alpha)$ to be those arcs that go from one boundary to another. This has diameter not greater than one and is nonempty if α intersects core curve of A (see Figure 110). One can extend our previous proofs of Behrstock inequality and BGI to this setting. We now want to prove the next theorem.

Theorem 30.2. *T_g is HHS.*

Before we can prove this however, we will need to define the maps from Teichmüller space to the curve complexes, which will use the following definition.

Definition 30.3. Let $X \in T_g$. Let α_1 be the shortest curve. This is not unique, so we need to make a choice. Given $\alpha_1, \ldots, \alpha_k$ with $k \leqslant 3g-3$, let α_{k+1} be the shortest curve disjoint from $\alpha_1, \ldots, \alpha_k$ (so $\{\alpha_i\}$ is "greedy shortest pants"). For each i, let β_i be the

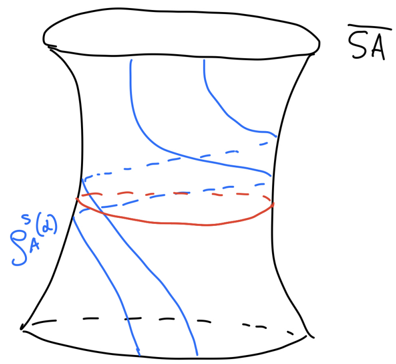

FIGURE 110. Core curve of A in red and $\rho_A^S(\alpha)$ in blue intersect in $\overline{SA}$.

shortest scc intersecting α_i minimally (once if α_i is not separating, twice if separating, see Figure 111). All these curves are simple. Then

$$\mu_X = \{\alpha_1, \dots, \alpha_{3g-3}, \beta_1, \dots, \beta_{3g-3}\}$$

is called Bers marking (see Figure 111).

Alternate definition: Sometimes, instead of assuming that β_j is a shortest curve intersecting α_j minimally, people arrange to have $i(\alpha_i, \beta_j) = 0$ if $i \neq j$, often as in the center and left part of Figure 111.

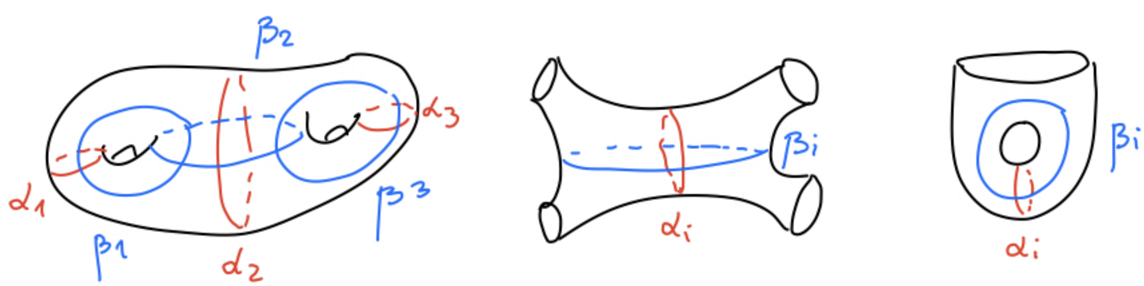

FIGURE 111. Examples of Bers marking

A crucial feature of μ_X is that μ_X "fills" X. In other words, for any scc γ there exists $\gamma' \in \mu_X$ with $i(\gamma, \gamma') \neq 0$.

Lemma 30.4. *There is a constant C such that for any $X \in T_g$ and any Bers marking μ_X, if $\gamma, \gamma' \in \mu_X$ then $i(\gamma, \gamma') \leqslant C$.*

Proof. We illustrate the proof in Figure 112. □

Definition 30.5. If U is not an annulus, define $\pi_U : T_g \to 2^{\mathcal{C}U}$ by

$$\pi_U(X) = \bigcup_{\substack{\gamma \in \mu_X \\ \gamma \text{ cuts } U}} \rho_U^S(\gamma)$$

for a Bers marking μ_X.

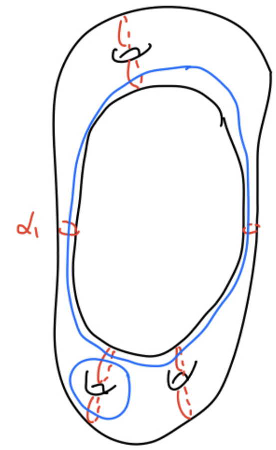

FIGURE 112. As β_i is shortest, it doesn't twist or visit the same place many times, for if it does, we could find a shorter curve satisfying the conditions required of β_i.

Remark 30.6. Lemma 30.4 and the logarithmic bound for distance in terms of intersection number in $\mathcal{C}U$ imply that $\pi_U(X)$ has uniformly bounded diameter.

31. Teichmüller space is an HHS (03/28, SK, TY)

Projection to annular subsurfaces. Recall that for annular subsurfaces U, we defined two different kinds of curve complexes. We denoted these curve complexes by $\mathcal{C}_{\text{MCG}}(U)$ and $\mathcal{C}_{Teich}(U)$ respectively. To unify the notation, now we'll use $\mathcal{C}(U)$ to denote $\mathcal{C}_{Teich}(U)$.

We can define a map $\pi_U^{MCG} : T_g \to 2^{\mathcal{C}_{MCG}(U)}$ by the same formula as for non-annular subsurfaces, namely

$$\pi_U^{MCG}(X) := \bigcup_{\substack{\gamma \in \mu_X \\ \gamma \text{ cuts } U}} \rho_U^S(\gamma)$$

for a Bers marking μ_X.

We record the length of the core curve γ of U using the following function.

$$y_U : T_g \to [1, \infty)$$
$$y_U(X) := \max\left(1, \frac{1}{\ell_\gamma(X)}\right)$$

Recall $\mathcal{C}(U)$ is the subset of the upper half plane with imaginary coordinate at least 1. The projection map $\pi_U : T_g \to \mathcal{C}(U)$ is given by first fixing an implicit quasi-isometric identification of $\mathcal{C}_{MCG}U$ and $\mathbb{R}$, and then using the following formula:

$$\pi_U(C) := \{x + iy_U(x) : x \in \pi_U^{MCG}(X)\}.$$

This is coarsely a point whose x coordinate measures the twisting in U. The y coordinate is only relevant in the case that the core curve of U is quite short, in which case it encodes how short this core curve is.

Facts about Teichmüller metric. To prove that T_g with the Teichmüller metric has an HHS structure, we need the following two results about the metric, which we state without proof.

Lemma 31.1 (Wolpert's lemma)**.** *Let x and y be points in T_g and let α be a simple closed curve. Then the ratio of lengths of α is x and y is bounded by $d_T(x, y)$ in the following manner:*

$$\exp(-d_T(x, y)) \leqslant \frac{\ell_\alpha(x)}{\ell_\alpha(y)} \leqslant \exp(d_T(x, y)).$$

Lemma 31.2 (Logarithmic lower bound)**.** *Let L_0 be a positive constant, x and y be points in the Teichmüller space, and α and β simple closed curves such that $\ell_\alpha(x) \leqslant L_0$ and $\ell_\beta(y) \leqslant L_0$. Then*

$$d_T(x, y) \succ \log\left(i(\alpha, \beta)\right).$$

Here $\succ$ denotes greater than, up to additive and multiplicative error that depends only on the genus and L_0.

A reference for the second lemma is [Raf07, Lemma 3.5], however it is also possible to prove the second lemma directly from the first lemma and the collar lemma.

Partial proof that Teichmüller space is an HHS. We now fill in some of the details of the proof of the fact that T_g with the Teichmüller metric is an HHS, by describing the index set $\mathfrak{S}$, the associated hyperbolic spaces $\mathcal{C}W$, and discussing the 9 axioms that define an HHS. References for this include [Raf07, Dur16], although are from before the modern definition of an HHS and we're not going to follow either. The paper [DDM14, Section 2] also has some nice discussion, in particular of the maps from Teichmüller space to the curve complexes. For the closely related case of the HHS structure on the mapping class group, a modern discussion with pointers to the literature can be found in [BHS19, Section 11].

The index set $\mathfrak{S}$ consists of all subsurfaces of the surface S (including the ones with multiple connected components), and we take the hyperbolic spaces $\mathcal{C}W$ to be the associated curve complexes (for disconnected subsurfaces, the associated curve complexes have bounded diameter, and thus are still hyperbolic).

(1) Projections: The projection maps are the projections of the Bers marking we described in this and the previous section. We also verified that the image of a point under these maps have uniformly bounded diameters. The image of T_g under any such map is uniformly quasi-convex by the virtue of being surjective.

A more delicate fact to verify is that these maps are all coarse Lipschitz. We sketch a proof of this fact in the thick part of Teichmüller space, and mention a theorem that helps with the thin part. It suffices to show that if $d_T(x, y) \leqslant 1$, and $\gamma \in \pi_W(x)$ and $\gamma' \in \pi_W(y)$, then $d_W(\gamma, \gamma')$ is bounded above by a constant independent of x, y, and W. Suppose now that x and y are in the thick part:

using Lemma 31.1, we can conclude that μ_x has similar lengths on x and y. Also, since y is in the thick part, μ_y is not too short on y, and since μ_x is not too short either, they must have low intersection number by Lemma 31.2. Thus, they are not too far off in the curve complex $\mathcal{C}W$.

In the thin part, getting a bound on the intersection number takes additional work. One approach is to use Minsky's product region theorem [Min96, Theorem 6.1].

(2) **Nesting:** A subsurface V is nested in W if it's a subsurface of W, after possibly isotoping. The coarse point ρ_W^V in $\mathcal{C}W$ is defined to be the boundary curves of V which are non-peripheral in W. The map ρ_V^W is subsurface projection.

(3) **Orthogonality:** Two subsurfaces are orthogonal if they are disjoint.

(4) **Tranvsersality and consistency:** The Behrstock inequality follows from the version of Behrstock inequality we proved for the curve complex. The proof of functoriality is omitted.

(5) **Finite complexity:** The complexity of T_g is linear in the Euler characteristic (or equivalently, the genus) of the surface.

(6) **Large Links:** This will be discussed next class.

(7) **Bounded geodesic image:** This follows from the Bounded Geodesic Image theorem for curve complexes which we proved in class.

(8) **Partial realization:** If V and W are orthogonal subsurfaces, and p and q are points in $\mathcal{C}V$ and $\mathcal{C}W$ respectively, it's easy to construct a point $x \in T_g$ that maps to p and q respectively. Namely, pinch both of those curves simultaneously; this can be done because the curves are in orthogonal subsurfaces. If p and q lie in transverse subsurfaces and satisfy the Behrstock inequality, one can construct a point $x \in T_g$ that maps to p and q, but we skip the proof of that.

(9) **Uniqueness:** This will be discussed next class.

32. More on Teichmüller space and more on Behrstock (03/30, JH, KS)

We are currently in the process of discussing the HHS structure on T_g. To finish the discussion, we need to address the large links and uniqueness axioms; in both cases we'll only give an idea of why they are true. For large links we'll restrict to the case $W = S$, where we need to prove the following:

(6) **Large Links:** For any $X, Y \in T_g$, there exist subsurfaces $T_1, T_2, \ldots, T_{\text{linear in } d_S(X,Y)}$, which are properly nested in S, such that

$$d_W(X, Y) \text{ large } \implies W \sqsubseteq T_i \text{ for some } i,$$

and the same holds for S replaced with any other element of $\mathfrak{S}$.

The main idea here is that if the distance between projections of two points in a lot of random (i.e., not properly nested) subsurfaces is large, then this forces the distance between the two points in S to be large. This is known as the passing up lemma, and is essentially the entire reason the large links axiom exists.

Remark 32.1. In a simplicial complex, the *link* of a vertex is everything distance 1 away from it. Analogously, the link of a non-separating curve α is the curve complex of the surface minus α. We can similarly view other curve complexes as links; for example

$\mathcal{C}W$ is the link of any collection of curves that fill the complement of W. So these $\mathcal{C}W$ are the links in "large links," and the axiom is about links where $d_W(X, Y)$ is large. In particular, it says that there are not too many large links.

Remark 32.2. Sometimes the large links axioms is stated with the linear bound replaced with the floor of the distance in S. The definitions are equivalent, but only up to changing curve complexes used in the HHS structure. This does matter and can be annoying, for example when a curve complex is a point (then the floor of the distance is 0, so there are no T_i's). We will proceed with the linear bound definition.

To see why (6) is true, recall the seemingly stronger but actually equivalent version of the bounded geodesic image theorem mentioned in Remark 23.2:

Theorem 32.3 (Strong Bounded Geodesic Image). *If $\alpha_1, \ldots, \alpha_n$ is a geodesic in $\mathcal{C}S$ and $d_W(\alpha_1, \alpha_n)$ is large, then some α_j does not cut W.*

The conclusion that α_j does not cut W is equivalent to saying $W \sqsubseteq S - \alpha_j$. So setting $T_j = S - \alpha_j$ proves (6).

Remark 32.4. Since large links and BGI are both about what can be large, they are almost always proved together.

Next, recall that the uniqueness axiom (9) is a result on coarse injectivity of the map $T_g \to \prod \mathcal{C}U$. We will give an idea of how this proof goes by proving a weaker statement.

Proposition 32.5. *Say X, Y are thick (i.e., they have no short curves) and $\mu_X = \mu_Y$, where μ_X and μ_Y are Bers markings. There is a universal upper bound for $d(X, Y)$.*

Proof. The curves in μ_X all have bounded intersection number with each other, so there are only finitely many mapping class group orbits of μ_X. Also, since X and Y are thick, there is an upper bound on the lengths of curves in μ_X, i.e., all curves in μ_X have length $\leqslant L_0$. So, it suffices to know

$$\{Y \in T_g : \ell_Y(\gamma) \leqslant L_0 \ \forall \gamma \in \mu_X\}$$

is compact, hence has bounded diameter. Even though there are infinitely many Bers markings that we could choose from, by mapping class group invariance, there are then only finitely many upper bounds that will arise. Compactness follows from the fact that μ_X is filling and the Thurston compactification of T_g. $\square$

Alternatively, we could think about Fenchel-Nielson coordinates (compare to the $9g - 9$ theorem).

From here, we just need to think about the case where $\mu_X \neq \mu_Y$, but they are close. In that case, construct an efficient path from one to the other by interpolating between markings.

Two alternative approaches for (9) are the following:

- Use Teichmüller geodesics and "balanced times," which is a technology that lets us show, for example, that if the Teichmüller geodesic is thick, then we're actually moving in the curve complex. So, if X and Y are far apart and the geodesic between them is thick, we're done immediately.

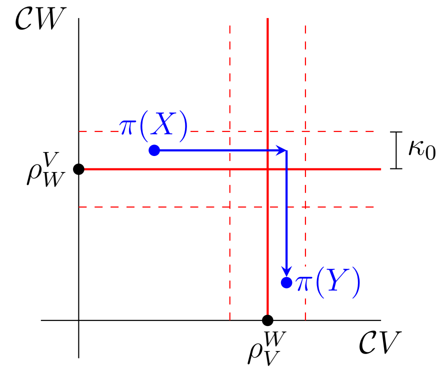

FIGURE 113. The Behrstock inequality, pictorally. Here $\pi = (\pi_V, \pi_W)$.

- One imagines there might be a way to reduce to the mapping class group case, where there is a full elementary proof of uniqueness [BHS19, Section 11].

Remark 32.6. The HHS structure on the mapping class group uses the Cayley graph as the underlying space, and if μ is a set of curves that fills U (perhaps specially chosen), $\pi_U(g)$ is the projection of $g\mu$ to $\mathcal{C}U$.

Switching gears, recall the Behrstock inequality: if $V \pitchfork W$, then

$$\min(d_W(\pi_W(X), \rho_W^V), d_V(\pi_V(X), \rho_V^W)) \leqslant \kappa_0.$$

We will pretend that ρ_W^V and ρ_V^W are points, for example by picking a point and adding to κ_0 to absorb additional error. The best way to remember this inequality is by the picture in Figure 113: it says that the projection of every point must lie within the dashed red cross, since the projection must be close to either ρ_W^V or ρ_V^W (or both).

This result is powerful because it says that in order to travel between points, it is necessary to travel in each direction separately, in a prescribed order. For example, in Figure 113, to get from $\pi(X)$ to $\pi(Y)$ while staying within the image of π (i.e., within the dashed red cross), we need to first travel along $\mathcal{C}V$, and then travel along $\mathcal{C}W$. If we're not worried about efficiency, we can travel around randomly before making the switch, but the path needs to start by traveling along $\mathcal{C}V$ and end with traveling along $\mathcal{C}W$.

A corollary of the Behrstock inequality is the following.

Lemma 32.7. *If $d_W(x, y) > 2\kappa_0$ and $d_V(x, y) > 2\kappa_0$, then (up to swapping x and y throughout) all of the following hold:*

$$\begin{aligned} d_W(x, \rho_W^V) &\leqslant \kappa_0 & \qquad d_W(y, \rho_W^V) &> \kappa_0 \\ d_V(x, \rho_V^W) &> \kappa_0 & \qquad d_V(y, \rho_V^W) &\leqslant \kappa_0. \end{aligned}$$

After the discussion above, this result should be fairly intuitive: it says that if x and y are distance greater than $2\kappa_0$ in both $\mathcal{C}W$ and $\mathcal{C}V$, then x must be in the horizontal strip and y in the vertical strip, and neither is in the overlap square (since they are too far apart).

Proof. The triangle inequality and assumptions imply that at most one of

$$d_W(x, \rho_W^V) \leqslant \kappa_0 \qquad \text{and} \qquad d_W(y, \rho_W^V) \leqslant \kappa_0 \tag{2}$$

holds (by comparing both to ρ_W^V in $\mathcal{C}W$ and using the fact that $d_W(x, y) > 2\kappa_0$). Similarly, at most one of

$$d_V(x, \rho_V^W) \leqslant \kappa_0 \qquad \text{and} \qquad d_V(y, \rho_V^W) \leqslant \kappa_0 \tag{3}$$

holds. If neither expression in (2) held, then by the Behrstock inequality, both expressions in (3) must hold, so exactly one expression in (2) holds; similarly, exactly one expression in (3) holds. Exchanging x and y throughout if necessary, assume $d_W(x, \rho_W^V) \leqslant \kappa_0$, so $d_W(y, \rho_W^V) > \kappa_0$. By the Behrstock inequality, $d_V(y, \rho_V^W) \leqslant \kappa_0$ must hold, so $d_V(x, \rho_V^W) > \kappa_0$, and this is everything we wanted to prove. □

33. The Behrstock partial order (04/01, SC, JH)

In the previous lecture, we saw that if the projected distance of two points x and y in two transverse domains V and W is large, then, in any path from x to y, one of the coordinates in $\mathcal{C}V$ or $\mathcal{C}W$ has to change before the other can. Motivated by this, we define a partial order on domains where such a comparison makes sense.

Definition 33.1. For any $E > \kappa_0$, and $x, y \in X$, an HHS, define

$$\Omega_E(x, y) = \{V \in \mathfrak{S} | \quad d_V(x, y) \geqslant E\}.$$

If $V, W \in \Omega_E(x, y)$, say $V < W$ if $V \pitchfork W$ and $d_W(x, \rho_W^V) \leqslant \kappa_0$.

This gives us a partial order on $\Omega_E(x, y)$, where $V < W$ means we have to make progress in V before W. One can check that if $V < W$ and $W < U$, then $V < U$. Pictorially, this situation is described in Figure 114.

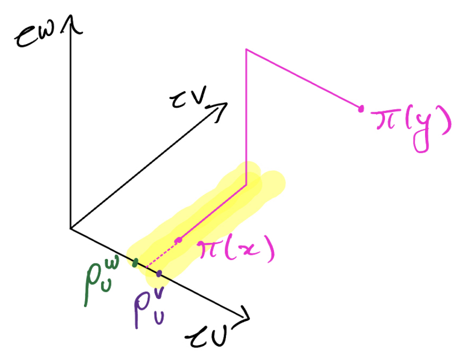

Figure 114. When $V < W$ and $W < V$ in $\Omega_E(x, y)$.

Figure 114 also gives us some extra information, such as $\rho_U^V \approx \rho_U^W$. More generally, if we have pairwise transverse domains $V_1, \ldots, V_k \in \Omega_E(x, y)$, we may reorder them so $V_1 < V_2 < \ldots < V_k$, and we have the schematic shown in Figure 115, which we shall term *Sisto's picture* since it appears in [Sis19].

Figure 115 is in general just representational, but it can be made literal for HHSes coming from coning off well separated convex subsets of a hyperbolic space. Before

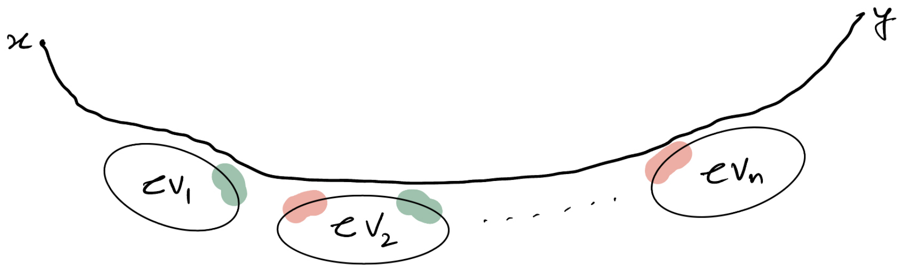

FIGURE 115. (Sisto's picture) The shaded regions have bounded diameter. The regions shaded in green contain $\rho^{V_i}_{V_j}$ for $i > j$ and the regions shaded in red contain $\rho^{V_i}_{V_j}$ for $i < j$.

proceeding, we remark that this is not total information, since we do not have any restriction if $V \perp W$, as depicted in Figure 116. However, if $V \subsetneq W$, then it can be shown that progress in $\mathcal{C}W$ usually happens when $\pi_W = \rho^V_W$.

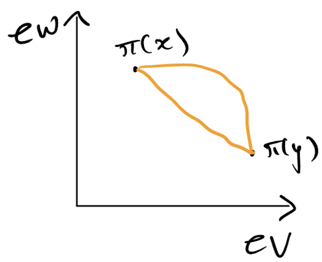

FIGURE 116. When $V \perp W$, there is no restriction on the $\mathcal{C}W$ and $\mathcal{C}V$ coordinates of paths from $\pi(x)$ to $\pi(y)$. In particular, both the paths shown are possible.

The next lemma, which we will not prove, shows us some more properties of domains in $\mathfrak{S}$. See [BHS19, Lemma 2.2] for a short proof.

Lemma 33.2. *For any HHS $(X, \mathfrak{S})$, there exists $M > 0$ depending only on the complexity of X such that if $V_1, V_2, \ldots, V_m$ are distinct and pairwise not transverse, then $m \leqslant M$.*

As a corollary, we get additional information about the partial order on $\Omega_E(x, y)$.

Corollary 33.3. *For large enough E, it is possible to partition $\Omega_E(x, y)$ into M subsets $\Omega_1, \Omega_2, \ldots, \Omega_M$, some of which may be empty, such that in each subset, all domains are pairwise transverse.*

The proof will use a very general combinatorial result on partial orders on finite sets called Dilworth's Theorem. For a proof, see wikipedia.

Proof. Because E is large, it is possible to use large links to show $\Omega_E(x, y)$ is finite. But, as the remark following this proof shows, this detail is actually not important.

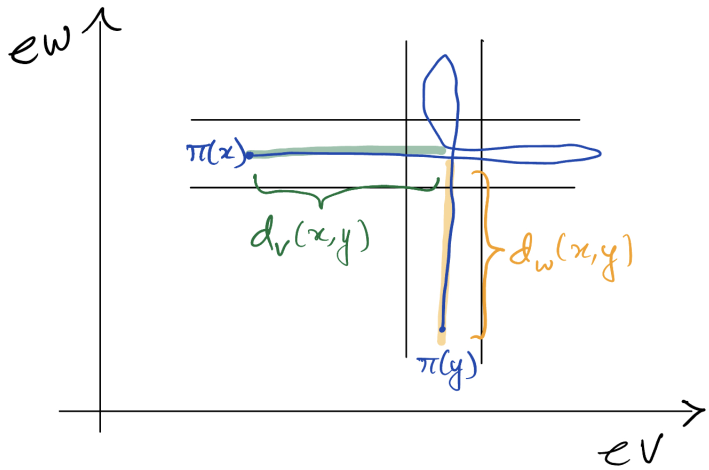

FIGURE 117. Any path from $\pi(x)$ to $\pi(y)$ must cover distances $d_V(x, y)$ in $\mathcal{C}V$ and $d_W(x, y)$ in $\mathcal{C}W$ as shown.

Consider the finite set $\Omega_E(x, y)$ with the partial order constructed above, and keep in mind that two elements of this subset are comparable in this partial order if and only if they are transverse. The previous lemma gives an upper bound on the size of an *anti-chain* of $\Omega_E(x, y)$, which is by definition a subset such that no two elements of the subset are comparable in the partial order. The conclusion is now the exact statement of one direction of Dilworth's Theorem, applied to the partial order on $\Omega_E(x, y)$. $\square$

Remark 33.4. An alternative to invoking large links in the previous corollary is just to prove the corollary for finite subsets of $\Omega_E(x, y)$. This is sufficient for the next lemma because the M is uniform, and at the end one can conclude that actually $\Omega_E(x, y)$ was finite all along.

As an application of this corollary, we get the lower bound in the distance formula.

Lemma 33.5. *For any x, y in an HHS X, and E large enough, $d(x, y) \gtrsim \sum_{V \in \mathfrak{S}}[d_V(x, y)]_E$.*

Proof sketch. Let Ω_i be the set of domains from Corollary 33.3 for which $\sum_{V \in \Omega_i}[d_V(x, y)]_E$ is maximum. Thus

$$\sum_{V \in \mathfrak{S}}[d_V(x, y)]_E \geqslant \sum_{V \in \Omega_i}[d_V(x, y)]_E \geqslant \frac{\sum_{V \in \mathfrak{S}}[d_V(x, y)]_E}{M}.$$

We conclude that $\sum_{V \in \mathfrak{S}}[d_V(x, y)]_E \asymp \sum_{V \in \Omega_i}[d_V(x, y)]_E$. So, it suffices to prove $d(x, y) \gtrsim \sum_{V \in \Omega_i}[d_V(x, y)]_E$.

To see this, we observe that any path in X (in particular, a quasi-geodesic) must progress (i.e. cover distances) in the curve complexes $\mathcal{C}V$ for $V \in \Omega_i$ in the order discussed (Figure 117). In addition, the path progresses in only one domain at a time and the progress in a domain does not happen much faster than the progress in X since the projection maps π_V are uniformly coarsely Lipschitz. $\square$

34. Hierarchy paths (04/04, TY, CK)

Recall Example 1.16 from the very first lecture, in which we saw that quasi-geodesics are not a very good notion in $\mathbb{R}^2$. Of course $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$ and $\mathbb{R}$ is hyperbolic, so $\mathbb{R}^2$ is an HHS. But one might reasonably feel like that's a somewhat degenerate example. The following lemma shows that this extends robustly to more complicated HHSes, in that quasi-geodesics can be arbitrarily bad.

Lemma 34.1. *Let $f : \mathbb{N} \to \mathcal{C}S$ be any path, so for each $i \in \mathbb{N}$, we have that $f(i), f(i+1)$ are disjoint curves of S. Then there is a quasi-geodesic ray $F : [0, \infty) \to T_g$ such that $\pi_s \circ F$ "traces out" this path f.*

By traces out, we mean that it may not be at unit speed, but the image is (very close to the) path.

Proof sketch. Pick a starting point $X_0 \in T_g$ with $f(0)$ having a really small length which is also the systolic length, say $\frac{1}{100}$. Then, let $X_1 \in T_g$ be a surface where $f(0), f(1)$ have length $\frac{1}{100}$ but with $n \gg 1$ twists at $f(0)$. So, if A_0 is an annular thickening of $f(0)$, then $d_{A_0}(X_0, X_1) \gg 1$.

In general, let X_i have both $f(i), f(i-1)$ with length $\frac{1}{100}$ and tons of twists around $f(i-1)$. To be a little more precise, if there are n_i twists, then since $d_{A_i}(X_i, X_{i-1})$ is comparable to (log of) the number of twists for an annular thickening A_i of $f(i)$, we can choose n_i so that $\log(n_i) \gg \sum_{j<i} \log(n_j)$. We also want to make sure that $d_{A_i}(X_i, X_{i-1})$ accounts for a definite fraction of $d_{T_g}(X_i, X_{i-1})$, so we have decent lower bounds on distances just by looking at progress in the A_i. □

Thus, we see that one can't really use quasi-geodesics on their own to work with HHSes. Instead, one uses the following notions.

Definition 34.2. If M is a metric space, a map $f : [0, \ell] \to M$ is a **D-unparameterized quasi-geodesic** if there exists a non-decreasing, surjective map $g : [0, L] \to [0, \ell]$ such that $f \circ g$ is a (D, D)-quasi-geodesic.

Intuitively, "the image of f is a quasi-geodesic and there is no backtracking."

Definition 34.3. A **D-hierarchy path** in an HHS $(X, \mathfrak{S})$ is a path $\gamma : [0, \ell] \to X$ that is a quasi-geodesic and such that $\pi_U \circ \gamma$ is a D-unparameterized quasi-geodesic for all $U \in \mathfrak{S}$. (Note that D doesn't depend on U.)

The third main theorem of HHSes (after the distance formula and realization theorem) is the following theorem on the existence of hierarchy paths.

Theorem 34.4. *Let $(X, \mathfrak{S})$ be a HHS. Then there exists D such that every pair of points in X can be joined by a D-hierarchy path.*

This takes a fair bit of work to prove, but the moral is that one should use hierarchy paths instead of quasi-geodesics for HHSes. (For the proof of all three main theorems, see [BHS19].)

Remark 34.5. Note that hierarchy paths are generally not unique. In one example - for the product HHS structure on $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$, either of the two "taxi-cab paths" between two points in $\mathbb{R}^2$ is a hierarchy path. In fact, many paths in the box demarcated by the two taxi-cab paths is a hierarchy path, since we only need both the x and the y components to be non-decreasing with time.

Instead of proving Theorem 34.4, we'll use it to prove the distance formula for an HHS (Theorem 25.4). We already saw that the lower bound comes from the Behrstock inequality, so we'll now show that Theorem 34.4 implies the upper bound in the distance formula. So we want to show that there exists $s_0 > 0$ such that for any $x, y \in X$, we have that

$$d_X(x,y) \leqslant \sum_U [d_U(x,y)]_s$$

for all $s \geqslant s_0$.

We'll use the following lemma.

Lemma 34.6. *For any s, D, there exists κ such that for any (D,D)-quasi-geodesic $\gamma : [0,M] \to Y$ in any metric space Y, if we have a sequence of times $s_1 \leqslant t_1 \leqslant s_2 \leqslant t_2 \leqslant \cdots \leqslant s_n \leqslant t_n$ satisfying $d(\gamma(s_i), \gamma(t_i)) \geqslant \kappa$ for all i, then $d(\gamma(0), \gamma(M)) \geqslant sn$.*

One can compare this to an analogous situation for geodesics: if we divide up a geodesic interval into n subintervals, with the minimum one have length s, then the length of the geodesic is at least sn. In particular, we expect κ in the lemma to be comparable to s.

Proof. By definition of quasi-geodesic,

$$d(\gamma(0), \gamma(M)) \geqslant M/D - D$$

and $t_i - s_i \geqslant (\kappa - D)/D$.

Since we have n intervals, we get $M \geqslant n(\kappa - D)/D$. Hence

$$d(\gamma(0), \gamma(M)) \geqslant n((\kappa - D)/D^2 - D/n).$$

If $s < (\kappa - D)/D^2 - D$, we get in particular the weak bound

$$d(\gamma(0), \gamma(M)) \geqslant ns.$$

Note that s can be made arbitrarily large by making κ large (D is fixed). □

Proof that Theorem 34.4 implies the distance formula. Let s be the desired threshold. Pick D as in Theorem 34.4, and let κ be as in Lemma 34.6. We then get a θ from the uniqueness axiom such that if $d_X(x,y) \geqslant \theta$, then there exists a domain U such that $d_U(x,y) \geqslant \kappa$.

Given a D-hierarchy path $\gamma : [0,M] \to X$, we have that $d(x,y) \asymp_D M$ because it's a quasi-geodesic. We'll want to divide the path into subintervals and apply uniqueness to each one.

If we have some $u, t \in [0,M]$ with $u - t = D\theta + D$, then since γ is a (D,D)-quasi-geodesic, we have that $d_X(\gamma(t), \gamma(u)) \geqslant \theta$. Then, by the uniqueness axiom, there exists U with $d_U(\gamma(t), \gamma(u)) \geqslant \kappa$. So, we have to divide $[0,M]$ into $\left\lfloor \frac{M}{D\theta + D} \right\rfloor$ intervals of length at least $D\theta + D$ each. Since D and θ are constants here, we have the following "quasi-equations:"

$$d(x, y) \asymp M \asymp \left\lfloor \frac{M}{D\theta + D} \right\rfloor,$$

where last quantity is the number of subintervals that we end up with. Note that θ and D are universal and independent of x and y.

Dividing $[0, M]$ into disjoint subintervals J_i of length $D\theta + D$ and using Lemma 34.6 to collect one contribution for each interval gives

$$\left\lfloor \frac{M}{D\theta + D} \right\rfloor s \leqslant \sum [d_U(x, y)]_s.$$

Combined with our above quasi-equations, this gives the desired result. □

Remark 34.7. Based on the previous proof, one can say that the lower bound in the distance formula is true because one can't have too much stuff happening at once. This intuition is also present in Rafi's work, where it takes a more geometric form: One thinks about moving along a Teichmüller geodesic, and observes that at each given time, only finitely many subsurfaces have a geometric shape which is compatible with progress in that subsurface at that time.

Optional Exercise 43. Prove the following special case of the existence of hierarchy paths, assuming the realization theorem: For any HHS X, show that there exists $D > 0$ such that if $x, y \in X$ with $\pi_U(x) = \pi_U(y)$ for all $U \neq S$, then there is a D-hierarchy path from x to y.

35. Infinite diameter and special quasi-geodesics (04/06, CK, SK)

We will sketch a proof of the fact that the curve complex of a surface has infinite diameter.

Consider the torus $T = \mathbb{R}^2/\mathbb{Z}^2$. By "the line of slope m on the surface containing a point p," we mean that one follows a line of slope m in $\mathbb{R}^2$ starting at p until it hits an edge e, and continues from the edge e' that e is glued to, with the same slope. This is the flat geodesic on the surface at p with tangent vector prescribed by m. It is a fact that every line of irrational slope is dense in T. See Figure 118.

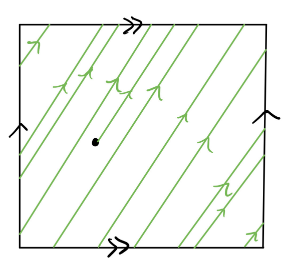

Figure 118. A line of irrational slope is dense in the torus.

Now consider the octagon surface S obtained by identifying opposite pairs of sides of a regular octagon. Remember that this corresponds to a CW-Complex with one 0-cell, four 1-cells and one 2-cell, giving an Euler characteristic of $2 - 2g = 1 - 4 + 1 = -4$, so that this represents a genus-2 surface.

It is a fact that for almost every slope, the line of that slope on the surface is either dense or hits a vertex (and most lines don't hit a vertex). The statement also holds for "flat" representatives of higher genus surfaces, where the flat representative is obtained by picking an arbitrary quadratic differential on the surface.

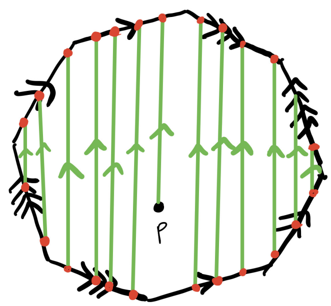

FIGURE 119. Rotate the polygon so that every vertical line is dense.

The existence of flat representatives with directions where vertical lines are dense can be used to construct an infinite diameter subset of the curve complex. Take a polygon defining S as above (for example, the octagon surface itself) and rotate it in $\mathbb{R}^2$ so that vertical lines are dense. This ensures there are no vertical simple closed curves. In particular, the vertical upwards line starting from $p \in S$ is dense. See Figure 119 above. Now perform the following iterative procedure, also illustrated by Figure 120 below:

- Pick a small horizontal interval I with left endpoint p.
- Let p_1 be the first point (as measured along l) where l meets I again. Define α_1 to be given by l from p to p_1 and then close up the curve along I.
- Let p_{i+1} be the first intersection of l with the sub-interval of I from p to p_i. Define α_{i+1} by following l from p to p_{i+1} and closing it up along I.

Since no vertical closed curves are simple, the curve l never intersects itself. We thus have a sequence of simple closed curves $\alpha_1, \alpha_2, \ldots$ such that:

- $\ell(\alpha_i) \to \infty$.
- Since l is dense, α_i are getting more and more dense in S. This can be made more precise, but will not be done in this lecture.
- For each s.c.c. γ, $i(\alpha_i, \gamma) > 0$ for i large enough.

Lemma 35.1. *$d_{\mathcal{C}S}(\alpha_n, \gamma) \to \infty$ for any fixed γ.*

This immediately implies the following proposition.

Proposition 35.2. $\operatorname{diam}(\mathcal{C}S) = \infty$.

Sketch of proof of lemma. Otherwise, after passing to a subsequence, we can assume $d_{\mathcal{C}S}(\alpha_n, \gamma) \leqslant k$ for some constant k. Then $d(\alpha_n, \gamma)$ takes only finitely many integer values. So, after replacing α_n with a sub-sequence, we can assume that for some fixed

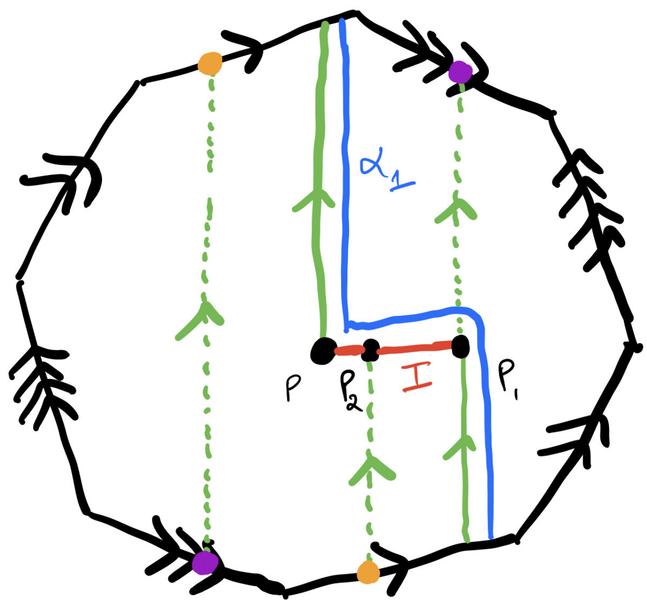

FIGURE 120. Determining p_1 and p_2.

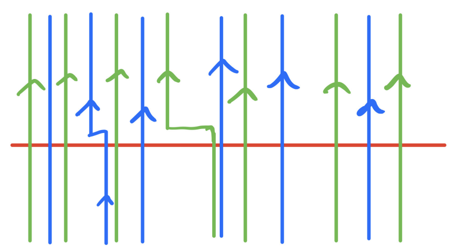

FIGURE 121. The curve $\alpha_n^{(1)}$ (in blue) must roughly follow α_n (in green) before closing up, and is thus long and dense if α_n is long and dense.

$d \leqslant k$, $d(\alpha_n, \gamma) = d \ \forall n \in \mathbb{N}$. For each n, let $\alpha_n, \alpha_n^{(1)}, \alpha_n^{(2)}, \ldots \alpha_n^{(d)} = \gamma$ be a geodesic in $\mathcal{CS}$ joining α_n and γ. So $d_{\mathcal{CS}}(\alpha_n^{(k)}, \gamma) = d - k$.

Consider the sequence $\{\alpha_n^{(1)}\}_{n=1}^{\infty}$. Notice that for big n, α_n is long and dense and for any n, $\alpha_n^{(1)}$ is disjoint from α_n since it is distance 1 from α_n. So, intuitively, $\alpha_n^{(1)}$ roughly follows the regions between two parallel segments of α_n and thus must loop back roughly as any times as α_n before closing up. See Figure 121 above. Hence, it can also be made arbitrarily long and dense for large enough n. We can repeat this for $\alpha_n^{(2)}$, this time using the fact that $\alpha_n^{(1)}$ is long and dense for large n and also that it is disjoint from $\alpha_n^{(2)}$. Repeating this process till $\alpha_n^{(d)}$, we get that $\alpha_n(d)$ can be made arbitrarily long and dense for large enough n. But this is clearly a contradiction, since $\gamma = \alpha_n^{(d)}$ is a fixed curve. This whole argument can be made precise, but will not be done in this lecture. It is somewhat forgiving since, for example, $\alpha_n^{(k)}$ does not have to be as long or as dense as $\alpha_n^{(k-1)}$; we mainly need that it gets longer and denser as $n \to \infty$. $\square$

Remark 35.3. We can prove Proposition 35.2 alternatively by showing that there is a coarse Lipschitz surjective map $\mathcal{CS} \to \mathbb{R}$ using "balanced times," outlined in the next few lines. We can isotope an s.c.c. α to a canonical length-minimizing piecewise linear path on the surface, and call its vertical and horizontal progress $x = \sum_i x_i$ and $y = \sum_i y_i$ respectively. See Figure 122 below. Then, the balanced time t_α is the time such that after travelling for time t along a Teichmuller geodesic, given by applying the matrix

$$\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$$

to the determining polygon, α is "as horizontal as it is vertical." That is, $e^{t_\alpha}x = e^{-t_\alpha}y$ or $t_\alpha = \frac{1}{2}\log(\frac{y}{x})$. One then shows the non-trivial fact that the map $\alpha \mapsto t_\alpha$ is coarse Lipschitz and surjective.

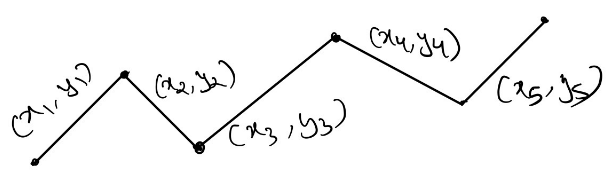

FIGURE 122. A piecewise linear path isotopic to α, used to determine the balanced time.

We end class with a fun lemma which, although we're not going to use it in this course, illustrates some useful techniques and is enough to prove that two different definitions of "convex co-compact" subgroups of the mapping class group are equivalent.

Lemma 35.4. *For a (D, D)-quasi-geodesic $\gamma : \mathbb{R} \to X$ into an HHS $(X, \mathfrak{S})$, TFAE:*
(1) $\pi_S \circ \gamma$ is a quasi-geodesic.
(2) $\exists s_0$ such that $\operatorname{diam}(\pi_U \circ \gamma) \leqslant s_0$ for all $U \neq S$.

This morally means that a path makes quasi-linear progress in the topmost curve complex iff it makes coarsely no progress in curve complexes of all nested subsurfaces.

Sketch of proof. We will use the distance formula for one direction and BGI for the other.

- (2) $\implies$ (1): We use the distance formula with a threshold greater than s_0. We claim that $|a - b| \asymp d_X(\gamma(a), \gamma(b)) \asymp d_S(\gamma(a), \gamma(b))$ for any $a, b \in \mathbb{R}$. Here, the first quasi-equality follows from the fact that γ is a quasi-geodesic. The second one follows from the distance formula and the fact that contributions from all other curve complexes are zero by (2). This establishes the fact that $\pi_S \circ \gamma$ is a quasi-geodesic.

- (1) $\implies$ (2): See Figure 123 below. Consider E given by BGI and κ_0 given by the transversality and consistency axiom. WLOG $E = \max(E, \kappa_0)$. Notice that $N_E(\rho^U_S) \subset \mathcal{C}S$ is uniformly bounded since ρ^U_S is uniformly bounded. Since $\pi_S \circ \gamma$ is a quasi-geodesic by (1), it spends a uniformly bounded amount of time in $N_E(\rho^U_S)$. Since π_S is uniformly coarsely Lipschitz and γ is also coarsely Lipschitz since it's a quasi-geodesic, the image of this uniformly bounded amount of time under $\pi_S \circ \gamma$ is also uniformly bounded.

  For any time t when γ is outside $N_E(\rho^U_S)$, BGI guarantees that the image of $(\rho^S_U \circ \gamma)(t)$ is in a bounded region. The functoriality condition (4.2) guarantees that $\pi_U \circ \gamma$ for the portion of γ outside of these times is bounded. ☐

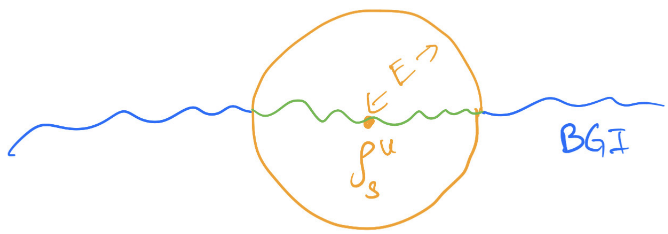

FIGURE 123. The casework for the projection of a quasi-geodesic from X to $\mathcal{C}U$ based on its distance from ρ^U_S in $\mathcal{C}S$, separated into the green and the blue case.

36. Historical comments and course highlights (04/08, KS, SC)

In this lecture we give a semi-historical overview. In pre 1980s coarse geometry existed in, for example, Mostow rigity and quasiconformal maps. Comparison geometry existed in, for example, CAT(0) spaces. In 1980s Gromov (and others, for example, Cannon) studied δ-hyperbolic spaces and especially δ-hyperbolic groups.

Definition 36.1. If G is a group with a finite generating set S, let $Cay(G, S)$ be the graph with

- a vertex for each element of G, and
- an edge from from g_1 to g_2 if $g_1 = sg_2$ for some $s \in S$.

If S_1 and S_2 are different finite generating sets, then $Cay(G, S_1)$ and $Cay(G, S_2)$ are quasi-isometric, so one says the Cayley graph of G (without specifying a generating set) is well defined up to quasi-isometry.

Definition 36.2. G is hyperbolic if $Cay(G, S)$ is.

Highlights of the theory of hyperbolic groups:

- Tons of groups are hyperbolic (hyperbolicity is generic in various models of random groups).
- Hyperbolic groups have solvable word problem, etc.

T_g, M_g, MCG are very important in geometry and topology, algebraic geometry, physics, dynamics, and 3-manifolds. To study M_g, one can compare it to $\mathbb{H}^3/\Gamma$. The idea is to cone off non-hyperbolic bits to get a hyperbolic space [Far98]. Motivated by this and the Ending Lamination Conjecture, Masur-Minsky showed $\mathcal{C}S$ is quasi-isometric to $Cone(T_g)$ and that $\mathcal{C}S$ is hyperbolic [MM99]. They proved the distance formula and BGI [MM99, MM00]. Many people (among them Minsky's former students Behrstock [Beh06] and Rafi [Raf07, Raf14]) added to the machine, which became central for questions like

- classifying hyperbolic 3 manifolds (see [Min03]) for an introduction);
- how hyperbolic T_g is?
- what do geodesics in T_g do?
- word problems and conjugacy problems in MCG;
- quasi-isometric rigidity;
- convex cocompact subgroups of MCG and connections to surfaces bundles, 4-manifolds, etc.

Progress on curve complexes continued: after several intermediate developments including the uniform hyperbolicity of curve graphs [Aou13, Bow14b, CRS14], in 2013 Hensel-Prytski-Webb posted a super-short proof of its hyperbolicity [HPW15], which lead the to bicorn approach that we followed in this course.

The same magical structure enjoyed by mapping class groups was found by Behrstock-Hagen-Sisto for many cube complexes, leading them to axiomatize the machine (2014), giving to the definition of an HHS [BHS17b, BHS19]. The main idea is that we want spaces that "would be" hyperbolic, except for product regions, whose factors again "would be" hyperbolic except for simpler product regions. We saw the following features for an HHS:

- Maps π_U to hyperbolic spaces $\mathcal{C}U$.
- If $U \sqsubseteq V$, $\mathcal{C}U$ holds information that was crushed (coned to a point called ρ^U_V)) in $\mathcal{C}V$), so it is reasonable to talk of a "hierarchy of hyperbolic spaces".
- $U \perp V$ corresponds to product behavior.
- For $U \pitchfork V$, $\mathcal{C}U$ and $\mathcal{C}V$ behave like convex subsets of hyperbolic space that are well separated or at least such that the projection of each to the other has been coned off.
- BGI.

We also have 3 main theorems whose proofs you can now read in [BHS19]:

- Realization: The image of an HHS in $\prod \mathcal{C}U$ is a set of consistent tuples. We think of π_U as coordinates and can build points in X by specifying coordinates.
- Existence of hierarchy paths.
- Distance formula (most iconic feature of HHS).

There exist many more HHSes, for example, cube complexes, which we shall see more of in the next 4 lectures. Hierarchical hyperbolicity is an extremely special property that is only one of a zoo of different ways a space can have some hyperbolic behavior while failing to be literally hyperbolic; it requires a huge amount of very particular (and useful!) structure. So it is beautiful and surprising that a long and ever growing list of extremely important spaces are hierarchically hyperbolic. The result is a powerful unification of the study of these spaces.

There is now a well established mutually profitable relationship between different HHSes: Most prominently, the Masur-Minksy ideas got applied to give new information on cube complexes, and now cube complexes are turning out to be useful in the study of all HHSes, even Teichmüller space and the mapping class group.

37. Cube complex basics (04/11, AW, TY, guest lecture by Susse)

Definition 37.1. A **d-dimensional cube** is a metric space isometric to $[-\frac{1}{2}, \frac{1}{2}]^d$ (with the Euclidean metric). A **k-dimensional face** (also called a **k-face**) of a cube is obtained by setting $(n-k)$ coordinates to $\pm\frac{1}{2}$. A 0-cube is called a **vertex**, and a 1-cube is called an **edge**.

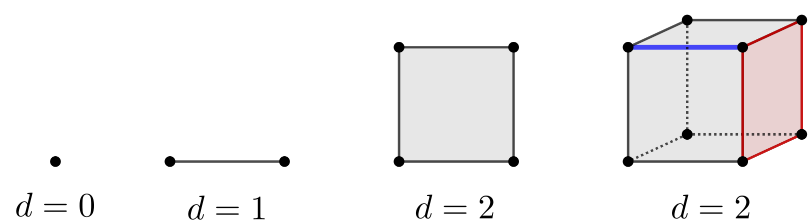

Figure 124. Some cubes. On the $d=3$ cube, a 1-face (edge) is highlighted in blue, and a 2-face in red.

Definition 37.2. A **cube complex** is obtained by gluing cubes isometrically along faces. It is **finite dimensional** if there is a bound on the dimensions of the cubes. It is **proper** if only finitely many cubes are glued along each face.

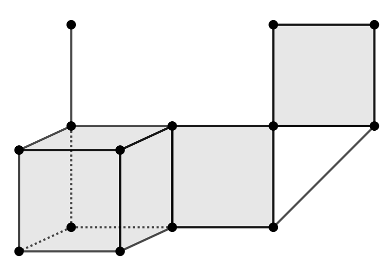

Figure 125. An example of a cube complex

Each cube complex has a natural path metric, where the length of a path γ is

$$\ell(\gamma) = \sum_C \ell(\gamma \cap C),$$

where the sum is over the maximal cubes C. The distance between points is

$$d(x,y) = \inf\{\ell(\gamma) : \ \gamma \text{ a path from } x \text{ to } y\}.$$

Theorem 37.3 (Bridson)**.** *In a finite dimensional cube complex, any two points can be connected by a geodesic.*

Definition 37.4. A geodesic metric space X is called **CAT(0)** if geodesic triangles in X are at most as fat as Euclidean triangles. More specifically, consider a triangle in X with vertices x, y, x and edges of length ℓ_1, ℓ_2, ℓ_3. Up to isometries, there is a unique triangle in Euclidean two space with edge lengths ℓ_1, ℓ_2, ℓ_3; sometimes this is called the comparison triangle. Call its vertices $\overline{x}, \overline{y}, \overline{z}$. For any point q on the edge from x to y,

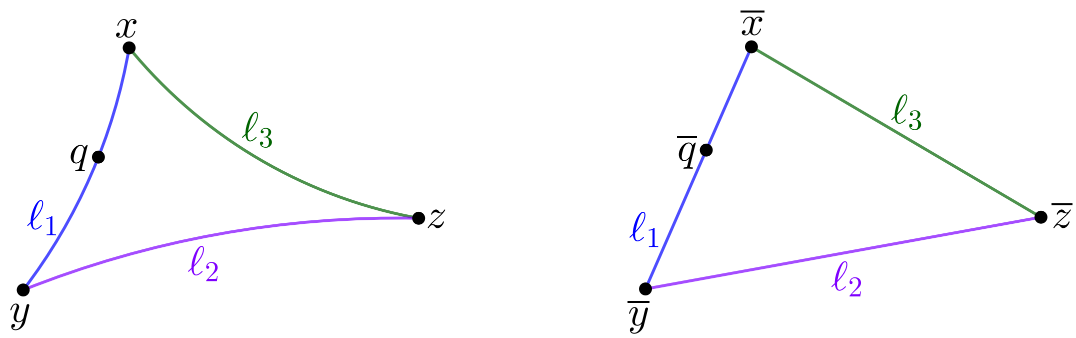

FIGURE 126. The definition of CAT(0). On the left is a triangle in X, and on the right is the comparison Euclidean triangle.

let $\overline{q}$ be the point on the edge from $\overline{x}$ to $\overline{y}$ defined by $d(x, q) = d(\overline{x}, \overline{q})$. The space X is CAT(0) if given any points in such a configuration, we have

$$d(z, q) \leqslant d(\overline{z}, \overline{q}).$$

This indicates that, as measured at q, the edge is "bowed in" towards z compared to the comparison triangle.

Definition 37.5. A space is called **non-positively curved** (NPC) if it is locally CAT(0).

An important question is: When is a cube complex X NPC or CAT(0)? Since the cubes themselves are Euclidean, any cube complex is locally CAT(0) in the interior of any maximal cube. One can worry though that "positive curvature" might be concentrated at a vertex. This motivates looking at the geometry near the vertices. Let $X^{(0)}$ denote the set of vertices for a cube complex X.

Definition 37.6. Given a vertex $v \in X^{(0)}$ of a cube complex X, its **link** is a combinatorial object roughly obtained as the ϵ-sphere based at v for any ϵ small. More formally, $\mathrm{Link}(v)$ is a delta-complex with a k-simplex for each corner of a $(k+1)$-dimensional cube at v, with these simplices glued in the same way that cubes incident to v are glued in X.

Theorem 37.7 (Gromov Link Condition)**.** *Let X be a cube complex. Then X is NPC if and only if the link of every vertex is a flag complex. Furthermore, X is CAT(0) if and only if it is NPC and simply connected.*

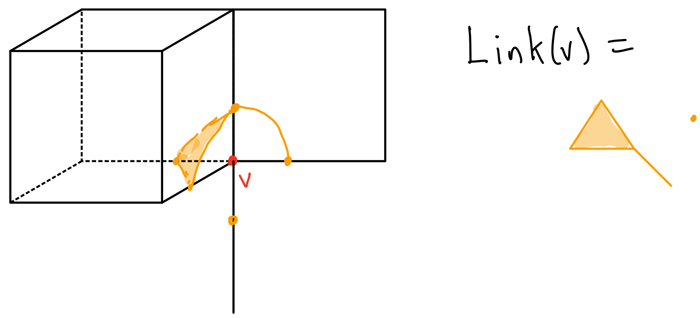

FIGURE 127. The definition of link. (The square and the cube really should be shaded in here, since they are part of the complex.)

A flag complex is a simplicial complex satisfying the following property: If it contains the 1-skeleton of a simplex, it contains the simplex. In other words, there are "no missing simplices"; see Figure 128 for examples and non-examples of flag complexes. (Recall that the difference between a simplicial complex and a delta-complex is that in a simplicial complex the vertices of each simplex must be distinct, and no two simplices can have the same vertex set.)

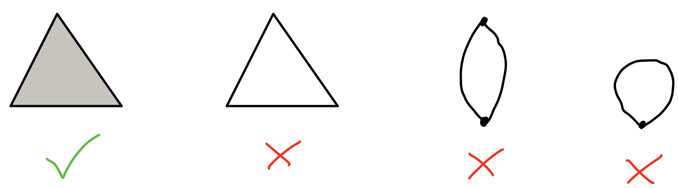

FIGURE 128. The left picture is a flag complex, but the right tree pictures are disallowed in a flag complex. Middle left contains the 1-skeleton of a triangle but does not contain a corresponding triangle, and the right two pictures are not simplicial complexes.

Example 37.8. Consider a cube X that isn't filled in, so it is a cube with 0 3-cubes, 6 2-cubes, 12 edges, and 8 vertices. The link of any of its vertices is a triangle that isn't filled it, and so isn't a flag complex. It is intuitive that this shouldn't be NPC, because it is topologically a sphere and spheres have positive curvature. One can also draw intuition from flat geometry, by noting that this cube complex is a singular flat surface with cone angles of $3 \cdot \frac{\pi}{2}$. In general, cone angles of less than 2π on a surface are often thought of as point masses of positive curvature.

Formally, one can show X isn't NPC directly by considering a small triangle near a vertex obtained by intersecting the cube with a plane. In this case, the comparison triangle can be thought of as the usual triangle lying on this plane. However, distance between points on this triangle are *larger* if one has to travel in X, because one can't cut across the plane and instead has to stay on the surface of the cube.

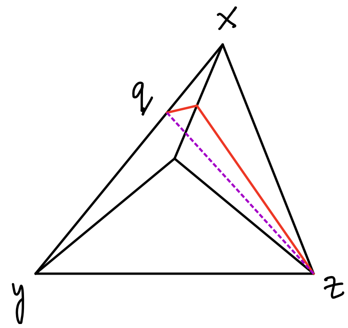

FIGURE 129. Cutting off the corner of a cube gives a visualization of a triangle in the space X of Example 37.8.

The summary is that X is NPC as long as it isn't missing any "corners of cubes". More broadly, the philosophy is that the combinatorics of a cube complex correspond to its geometry.

Let us mention a result in passing that illustrates the power of CAT(0) cube complexes. Suppose X is a proper, finite dimensional CAT(0) cube complex, and that a group G acts on X properly and co-compactly by isometries. Sageev and Wise showed that G satisfies the Tits alternative: every subgroup of G either contains the free group F_2, or is virtually Abelian. It is an open problem whether this is true if X is a CAT(0) space but not a cube complex. This result is perhaps one instance of theme, which is that CAT(0) cube complexes in some ways mimic symmetric spaces.

The next part of the lecture is on hyperplanes.

Definition 37.9. A **midplane** in $C = [-\frac{1}{2}, \frac{1}{2}]^d$ is one of the d sets $M_i = C \cap \{x_i = 0\}$.

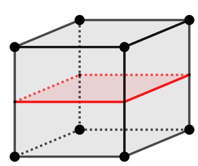

FIGURE 130. One of the three midplanes in a 3-cube.

Definition 37.10. A **hyperplane** in a cube complex is a connected subspace which intersects each cube in a midplane of that cube or in the empty set.

More concretely, one thinks of a hyperplane as being obtained by starting with a midplane, and then "following it around" by iteratively adding midplanes of adjacent cubes that one of the already added midplanes intersects.

Definition 37.11. Given a hyperplane H, its **carrier** $N(H)$ is the union of all the cubes that it intersects.

Fact 37.12. *If H is a hyperplane in a CAT(0) cube complex, $N(H) \cong H \times [-\frac{1}{2}, \frac{1}{2}]$.*

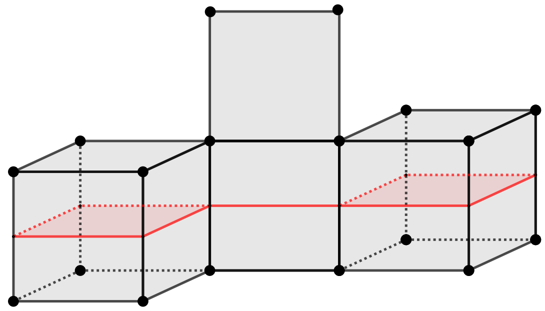

FIGURE 131. A hyperplane.

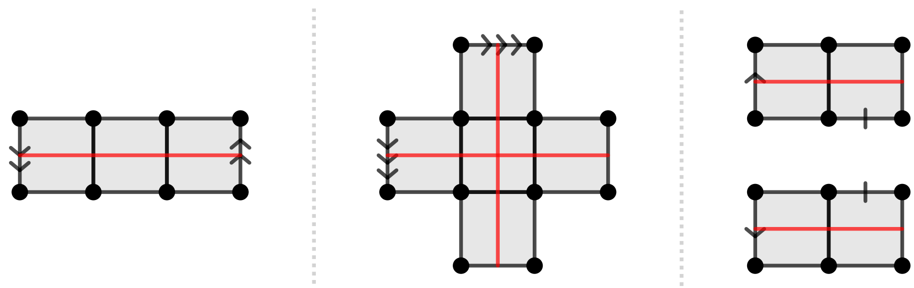

FIGURE 132. Three behaviours which can occur for hyperplanes in cube complexes, but cannot occur for hyperplanes in CAT(0) cube complexes: A one-sided hyperplane (left); a self-intersecting hyperplane (middle); and a self-osculating hyperplane (right). In all three examples, it is not the case that $N(H)$ is homeomorphic to $H \times [-\frac{1}{2}, \frac{1}{2}]$.

This fact can be derived by first showing that the behaviors illustrated in Figure 132 cannot occur.

Fact 37.13. *In a CAT(0) cube complex X,*

1. *For any hyperplane H, the space $X - H$ has exactly two components called **halfspaces**, denoted $\overleftarrow{H}, \overrightarrow{H}$. (There doesn't seem to be any canonical way in general to decide which halfspace gets the left arrow and which gets the right arrow.)*
2. *$N(H)$ is convex.*
3. *H is naturally a cube complex of dimension 1 less than X.*
4. *If $H_1 \cap H_2 = \varnothing$, then up to reordering we have $\overrightarrow{H}_1 \subset \overrightarrow{H}_2$ and $\overleftarrow{H}_2 \subset \overleftarrow{H}_1$.*

Let us focus now on the 1-skeleton $X^{(1)}$ with its natural graph metric. (Usually we use the Euclidean metric on a CAT(0) cube complex, but the graph metric on the 1-skeleton can be thought of as the restriction of the ℓ^1 metric on X.)

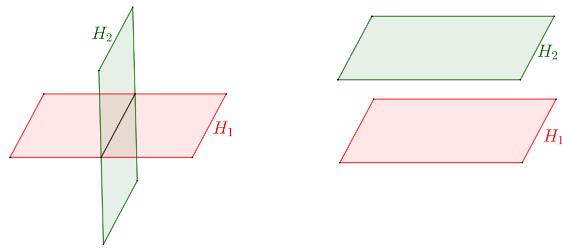

FIGURE 133. The two possible behaviors for how two hyperplanes relate to each other in a CAT(0) cube complex: Either they intersect (left), or they divide the space into an in between region and two halfspaces (right).

For any cube complex X, CAT(0) or not, a basic combinatorial observation is that if γ is a (simplicial) path in $X^{(1)}$, we have the formula

$$\ell(\gamma) = \sum_H \#(H \cap \gamma),$$

where the sum is over the hyperplanes H in X. The formula is true because the length $\ell(\gamma)$ is the number of edges in the path γ, and for each edge there is a unique hyperplane which intersects the edge in its midpoint.

We can say more if X is CAT(0): each geodesic crosses each hyperplane at most once.

Proposition 37.14. *Let X be a CAT(0) cube complex, $v, w \in X^{(0)}$, and let γ be a geodesic edge path from v to w. Then*

$$\#(H \cap \gamma) \leqslant 1.$$

You might compare this to the fact that in $\mathbb{R}^n$, each geodesic crosses each hyperplane (codimension 1 affine subspace) at most once. To cross more than once would be a waste of the geodesic's time!

Note that it is tempting to say that the proposition is an immediate corollary of the fact that $N(X)$ is convex and that geodesics in CAT(0) spaces are unique, but those results concern that Euclidean metric on X, and here we are using the restriction of the ℓ^1 metric.

Proof. Suppose to the contrary that H is a hyperplane with (at least) two points v', w' in $H \cap \gamma$. Suppose that we pick this data (H, v', w') with v' and w' as close together as possible along γ; we'll call this an "innermost" double intersection. It follows that the segment of γ from v' to w' does not intersect H at all, and intersects each hyperplane at most once.

Let γ' be the edge path that follows γ until just before v', then follows a minimal length path β from just before v' to just after w' in the 1-skeleton of $N(H)$, and then follows γ to w. It's possible to show that actually γ' is shorter than γ, giving a contradiction.

Here are a few more details, filled in by AW: We can prove this result by induction on the dimension of the cube complex. The 1-dimensional cube complexes are trees and the hyperplanes are midpoints of edges, so the base case is easy.

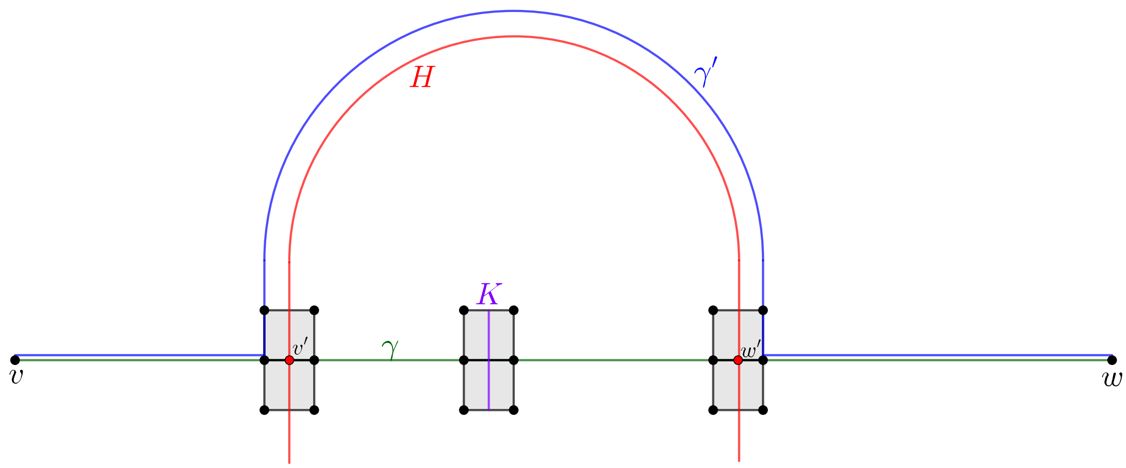

Figure 134

Let $\alpha \subset \gamma$ be the original path between the start and endpoints of β.

It's not hard to show that H is CAT(0) with its natural cube complex structure, and we know it has dimension one less than X, so by induction we can assume that β crosses each hyperplane at most once. Since hyperplanes are two-sided, it follows that α crosses each hyperplane that β does. Then, the length of α is at least as long as the length of β. But actually α is longer by at least 2, since it has to cross H twice, and β does not. This gives a contradiction. □

From the proposition, we get an upgraded version for the formula for the length of a path.

Corollary 37.15. *If X is a CAT(0) cube complex, for any $v, w \in X^{(0)}$, we have*

$$d_{X^{(1)}}(v, w) = \#\{H : x \text{ and } y \text{ are separated by } H\}.$$

For a reference on the results in this lecture, start with [Wis12], and, for more detail, see [Wis21]. For a reference that also includes the topics in the next lecture, see [Hag].

Optional Exercise 44. Show that geodesics are unique in CAT(0) spaces.

Optional Exercise 45. Show that geodesic CAT(0) spaces are contractible.

Optional Exercise 46. Prove the easier direction in Gromov's link condition: If X is a cube complex and the link of one of its vertices is not a flag complex, then X is not NPC.

38. Medians and gates (04/13, AW, KS, guest lecture by Susse)

Let X be a CAT(0) cube complex. We will work with the 1-skeleton $X^{(1)}$ with its graph metric, and all distances in this lecture will be with respect to this metric. The graph $X^{(1)}$ is typically not CAT(0), since it typically has many loops and CAT(0) spaces must be simply connected. But in addition to being more combinatorial (for example, distances between vertices are integers), the graph metric has some deep and useful geometric properties.

Definition 38.1. If X is a metric space, a point m is a median for a triple $x, y, z \in X$ of points if the following three equalities hold:

$$\begin{aligned} d(x,y) &= d(x,m) + d(m,y) \\ d(x,z) &= d(x,m) + d(m,z) \\ d(y,z) &= d(y,m) + d(m,z). \end{aligned}$$

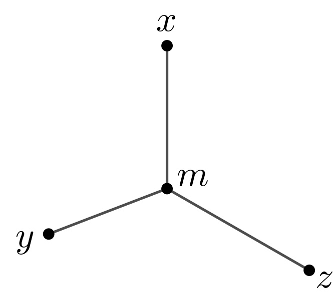

FIGURE 135. The definition of a median.

If X is a geodesic metric space, these equalities are equivalent to there being a geodesic between any pair of points in $\{x, y, z\}$ containing m. Note however that in most of the spaces we deal with, there may be more than one geodesic joining a given pair of points, as in Figure 136.

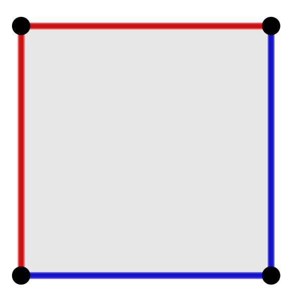

FIGURE 136. Both the red path and the blue path are geodesics from the bottom left to the top right vertices in the 1-skeleton of this square.

Definition 38.2. A space in which every three points has a median is called a median space. A graph in which every three vertices has a median is called a median graph. (Actually we should require *unique* medians, and Exercise 53 shows this is important, but we'll gloss over uniqueness in this lecture.)

Theorem 38.3 (Chepoi)**.** *If X is a CAT(0) cube complex, then $X^{(1)}$ is a median graph.*

Conversely, it is known that every median graph is the 1-skeleton of a CAT(0) cube complex. See [Hag, Theorem 1.17] for more details.

Proof. Fix a triple of vertices x, y, z; we want to show they have a median. Let $\mathcal{H}$ be the set of half spaces that contain at least two of the tree points x, y, z.

Recall that each hyperplane cuts X into two halfspaces. Since hyperplanes do not contain vertices, and x, y, z are vertices, exactly one of the two halfspaces arising from any given hyperplane will be in $\mathcal{H}$.

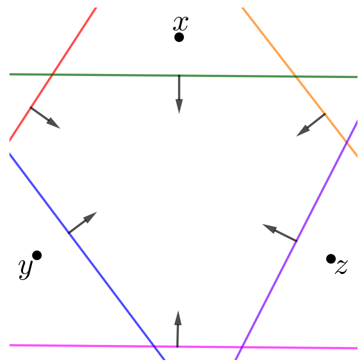

FIGURE 137. Half spaces in $\mathcal{H}$.

Define
$$M = X^{(0)} \cap \bigcap_{\vec{H} \in \mathcal{H}} \vec{H}$$
to be the vertices contained in all half spaces in $\mathcal{H}$. We will show M consists of a single point, and this point is a median for x, y, z.

Claim 1: M contains at least one point. Note that

(1) $\mathcal{H}$ does not contain an infinite descending chain
$$\vec{H}_1 \supsetneq \vec{H}_2 \supsetneq \cdots .$$

(2) Any two half spaces in $\mathcal{H}$ have non-empty intersection.

A result of Sageev gives that any collection of half spaces satisfying these two conditions contains at least one vertex in their intersection.

Claim 2: M contains at most one point. Suppose not, and let m and m' be distinct points in M.

Since m and m' are distinct vertices, there exists a hyperplane H separating them. Only one of the two half spaces associated to H can contain at least 2 of the 3 points $\{x, y, z\}$, so this immediately gives a contradiction using the definition of M.

Claim 3: The one point m in M is a median for x, y, z. By symmetry, it suffices to show
$$d(x, y) = d(x, m) + d(m, y).$$

Keeping in mind Corollary 37.15, we first make some observations on hyperplanes.

(1) There does not exist a hyperplane H that separates x from m and separates y from m. In other words, the situation shown in Figure 138 does not occur. This is because the associated half space containing x and y but not m would be in $\mathcal{H}$, proving that m is not in M and thus giving a contradiction.

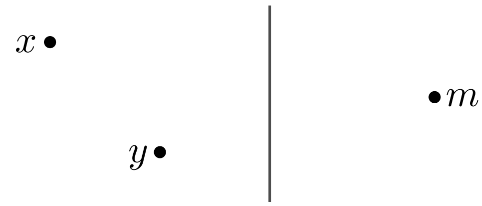

FIGURE 138

(2) It follows from (1) that every hyperplane that separates x and m must separate x from y, and similarly with x and y swapped.
(3) It also follows from (1) that every hyperplane that separates x and y must either separate x from m, or separate y from m, but not both.

Corollary 37.15 now implies Claim 3, because the above observations imply that the set of hyperplanes separating x and y is the disjoint union of the set of hyperplanes separating x and m and the set of hyperplanes separating m and y. □

Definition 38.4. A subcomplex K of $X^{(1)}$ is convex if every geodesic joining two points of K is contained in K.

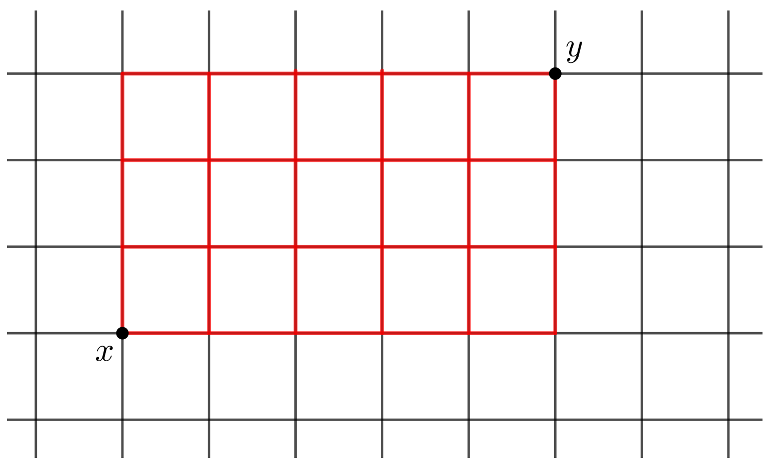

FIGURE 139. In this example using the standard cubulation of $\mathbb{R}^2$, the smallest convex subcomplex containing x and y is shown in red.

For example, if X is $\mathbb{R}^2$ with its standard cubulation, then $X^{(1)}$ is an infinite grid, and the only bounded convex sub-complexes rectangular regions. (The graph metric corresponds to the "taxi-cab" metric on $\mathbb{R}^2$.)

Lemma 38.5. *If K is convex, $x, y \in K$, $z \in X^{(0)}$, then the median $m = m(x, y, z)$ is in K.*

Note that only two of the three vertices x, y, z are required to lie in K. We emphasize again that convexity here is quite a strong assumption, since our definition requires that all geodesics joining two points of K stay in K, and there may be quite a number of such geodesics.

Proof. The definition of median gives that there is a geodesic from x to y that contains m. Because K is convex, this whole geodesic must lie in K, so in particular m lies in K. □

This has the following amazing consequence.

Corollary 38.6 (Helly property). *If $K_1, \ldots, K_m \subset X^{(1)}$ are convex, and $K_i \cap K_j \neq \varnothing$ for all i, j, then*

$$\bigcap_{i=1}^{m} K_i \neq \varnothing.$$

Proof. First assume $m = 3$. Pick

$$\begin{aligned} x &\in K_1 \cap K_2, \\ y &\in K_2 \cap K_3, \\ z &\in K_1 \cap K_3. \end{aligned}$$

Note that $m = m(x, y, z)$ is in K_1 because x and z are in K_1 and K_1 is convex. Similarly, m is in K_2 and K_3, proving that the triple intersection is non-empty.

The $m > 3$ case follows from the $m = 3$ case and induction, since one can replace $(K_1, \ldots, K_m)$ with $(K_1, \ldots, K_{m-2}, K_{m-1} \cap K_m)$ and use the $m = 3$ case to to show the assumptions continue to hold. □

We turn now to the topic of closest point projections to convex sets. If we used the Euclidean metric on X, the existence of such a projection would follow just from the fact that X is CAT(0). But instead we'll use the graph metric on the 1-skeleton. That means we have to do some extra work: since the 1-skeleton need not be CAT(0), we don't know from general theory that closest point projections exist. But, once we've done this work, we'll get a projection with some extra nice properties. We'll call such projections defined using the 1-skeleton gate maps, for reasons that will be clarified soon.

Proposition 38.7. *Let $K \subset X^{(1)}$ be a convex subcomplex. Let x be a vertex of X. Then there is a unique vertex y in K such that*

$$d_{X^{(1)}}(x, y) = \min\{d(x, t) : t \in K^{(0)}\}.$$

In other words, unique closest points exist for convex subgraphs of the 1-skeleton of a CAT(0) cube complex.

Proof. To start, note that since distances in a graph are integer, the minimum really is a minimum, not an infimum.

Suppose y and y' are both vertices in K that have minimal distance to x. Let $m = m(x, y, y')$ be the median. Since $y, y' \in K$ and K is convex, Lemma 38.5 gives that $m \in K$.

But, since $d(x, y) = d(x, m) + d(m, y)$, and since $m \in K$, the fact that y achieves the minimal distance implies that $m = y$. Similarly, $m = y'$, so $y = y'$. □

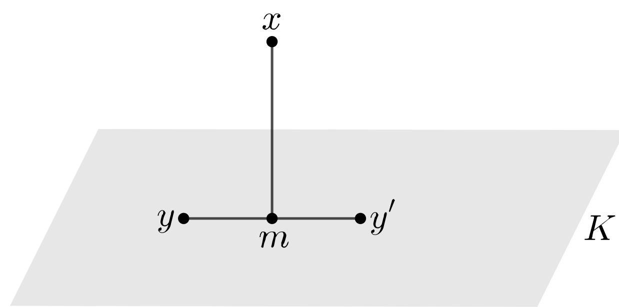

FIGURE 140. The proof of Proposition 38.7 rules out this situation.

Definition 38.8. In the situation above, we call y the projection of x onto K or the gate from x into K, and denote it $y = \mathfrak{g}_K(x)$. The map $\mathfrak{g}_K$ will be called either the projection map onto K, or the gate map.

The terminology "gate" is motivated by the following amazing observation.

Lemma 38.9. *In the setup above, for any $z \in K^{(0)}$, there is a geodesic from x to z that passes through $\mathfrak{g}_K(x)$.*

Thus, if you want to get from x to any point of K, you might as well first "enter" K at the gate $\mathfrak{g}_K(x)$.

Note however that it can sometimes be equally efficient to travel to given destination in K by entering K at a different location; it doesn't have to be the case that all geodesics from x to K go through the gate, just at least one to each vertex.

Proof. This is basically the same as the last proof. Set $y = \mathfrak{g}_K(x)$, and $m = m(x, y, z)$. As before, using that y achieves the minimal distance, we get that $m = y$. Then the median property gives that there is a geodesic from x to z pasing through $m = y = \mathfrak{g}_K(x)$. □

Proposition 38.10. *Let K be convex. Suppose that x and y are adjacent vertices in X. Then:*

(1) If $\mathfrak{g}_K(x) \neq \mathfrak{g}_K(y)$, then $d(x, \mathfrak{g}_K(x)) = d(y, \mathfrak{g}_K(y))$.
(2) $d(\mathfrak{g}_K(x), \mathfrak{g}_K(y)) \leqslant 1$.
(3) $d(\mathfrak{g}_K(x), \mathfrak{g}_K(y)) = 1$ if and only if the unique hyperplane separating x and y crosses K.
(4) If $\mathfrak{g}_K(x) \neq \mathfrak{g}_K(y)$, then a hyperplane separates x from $\mathfrak{g}_K(x)$ if and only if separates y from $\mathfrak{g}_K(y)$.

As is often the case, here it is implicit that we're using the graph metric on the 1-skeleton. The proposition might be summarized by saying that the gate map is 1-Lipschitz, and when two adjacent vertices have distinct projections we have many properties enjoyed by the "strip" picture in Figure 141. It would probably be ok to skip the proof and view the proposition as a black box.

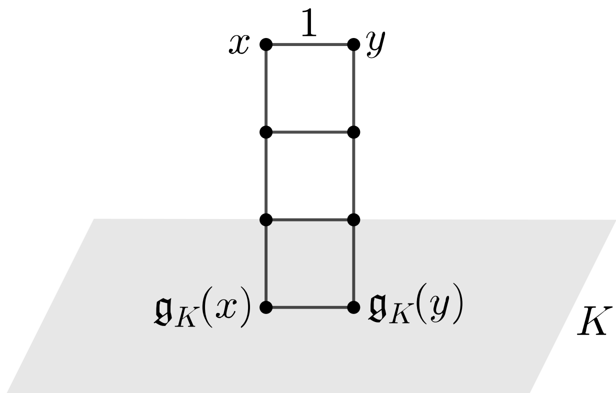

FIGURE 141. A strip.

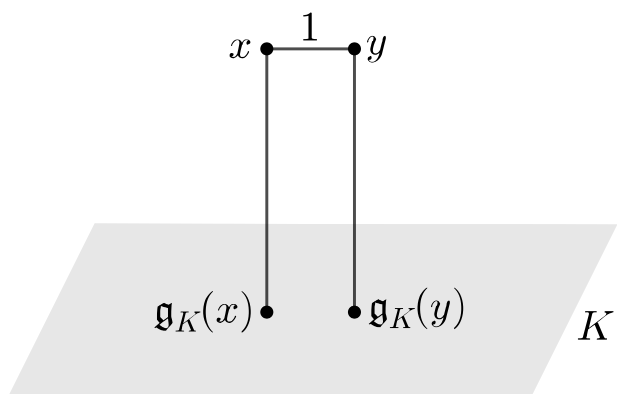

FIGURE 142. The proof of Proposition 38.10.

Proof. Since $\mathfrak{g}_K(x)$ is the closest vertex in K to x and $\mathfrak{g}_K(y)$ is a vertex in K of distance at most $d(x,y) + d(y, \mathfrak{g}_K(y)) = 1 + d(y, \mathfrak{g}_K(y))$ from x, we get

$$d(x, \mathfrak{g}_K(x)) \leqslant 1 + d(y, \mathfrak{g}_K(y)).$$

Furthermore, since $\mathfrak{g}_K(x)$ is the *unique* closest point to x in K, if $\mathfrak{g}_K(x) \neq \mathfrak{g}_K(y)$ we get that the above inequality must be strict, so actually

$$d(x, \mathfrak{g}_K(x)) \leqslant d(y, \mathfrak{g}_K(y)).$$

Since the same statements hold with x and y swapped, in general we get

$$|d(x, \mathfrak{g}_K(x)) - d(y, \mathfrak{g}_K(y))| \leqslant 1,$$

and if $\mathfrak{g}_K(x) \neq \mathfrak{g}_K(y)$ we get

$$d(x, \mathfrak{g}_K(x)) = d(y, \mathfrak{g}_K(y)).$$

This finishes the proof of (1).

By Lemma 38.9, there is a geodesic from x to $\mathfrak{g}_K(y)$ that passes through $\mathfrak{g}_K(x)$. Any geodesic from x to $\mathfrak{g}_K(y)$ has length at most $1+d(y, \mathfrak{g}_K(y))$. If $\mathfrak{g}_K(x) \neq \mathfrak{g}_K(y)$, it follows from (1) any geodesic from x to $\mathfrak{g}_K(y)$ has length at most $1 + d(x, \mathfrak{g}_K(x))$. Thus, the fact there there is such a geodesic that goes through $\mathfrak{g}_K(x)$ implies $d(\mathfrak{g}_K(x), \mathfrak{g}_K(y)) = 1$, finishing the proof of (2).

Now let H be the unique hyperplane separating x and y. First suppose it intersects K. Let z be a point of K on the same side of H as x (for example a vertex of an edge of K intersected by H). By Lemma 38.9, there is a geodesic from x to z passing through $\mathfrak{g}_K(x)$. Both x and z are on the same side of H, and Proposition 37.14 states that a geodesic can intersect a hyperplane at most once, so it follows that $\mathfrak{g}_K(x)$ must be on the same side of H as x and z. Symmetrically, $\mathfrak{g}_K(y)$ must be on the same side of H as y. Hence, H separates $\mathfrak{g}_K(x)$ and $\mathfrak{g}_K(y)$. Given (2), that shows $d(\mathfrak{g}_K(x), \mathfrak{g}_K(y)) = 1$.

Conversely, suppose $d(\mathfrak{g}_K(x), \mathfrak{g}_K(y)) = 1$, and let H be the unique hyperplane separating $\mathfrak{g}_K(x)$ and $\mathfrak{g}_K(y)$. It suffices to show that H separates x and y. For that, it suffices to show that x and $\mathfrak{g}_K(x)$ are on the same side of H, since combining this with the symmetric statement with x replaced with y proves that H separates x and y. Suppose to the contrary that x and $\mathfrak{g}_K(x)$ are on opposite side of H. There is a geodesic from x to $\mathfrak{g}_K(y)$ passing through $\mathfrak{g}_K(x)$. Our suppositions show that this geodesic must cross H at least twice, contradicting Proposition 37.14. This concludes the proof of (3).

Finally, suppose in order to find a contradiction that a hyperplane H separates x from $\mathfrak{g}_K(x)$, but that both y and $\mathfrak{g}_K(y)$ are on the same side of H. Case 1 is that y and $\mathfrak{g}_K(y)$ are on the same side of H as $\mathfrak{g}_K(x)$. In this case, H separates x and y but not $\mathfrak{g}_K(x)$ and $\mathfrak{g}_K(y)$, contradicting our analysis above. Case 2 is that y and $\mathfrak{g}_K(y)$ are on the same side of H as x, and again this contradicts our analysis above. □

Our last topic for this lecture is to give some hints on what the HHS structure will be on a CAT(0) cube complex X, and what difficulties might arise. To start, consider disjoint convex complexes K_1, K_2. Because of the previous proposition, we might expect a product region in between the projections of each to the other, as illustrated in Figure 143. (The idea is that the region should be parametrized by the set of hyperplanes crossing K_1 and K_2 times the set of hyperplanes separating K_1 and K_2; see the "Bridge Theorem" in [Hag, Theorem 1.22].) This may give some intuition.

Let's indicate what the set of domains should be. We can start with $\mathcal{H}$, the set of all hyperplanes. (Or rather, since it is better to have sub-cube-complexes here, take $\mathcal{H}$ to be the set of boundary components of carriers of hyperplanes.) This $\mathcal{H}$ is however simultaneously too large and too small to be the index set.

The sense in which it is too small is that, motivated by Figure 143, we should actually use the hyperclosure $\mathcal{F}$ of $\mathcal{H}$, which is defined to be the smallest collection of subcomplexes containing $\mathcal{H}$ that is closed under taking projections. So if $K_1, K_2 \in \mathcal{F}$, then $\mathfrak{g}_{K_1}(K_2) \in \mathcal{F}$.

The sense in which it is too large is that "parallel" objects should be identified. So actually the set of domains would be $\mathcal{F}$ modulo a "parallelism" equivalence relation.

Nesting will be inclusion. The main question will end up being whether the finite complexity axiom holds. This actually can fail for arbitrary CAT(0) cube complex, as illustrated in Figure 144. In other words, not all CAT(0) cube complexes have hyperclosure with finite complexity.

However, this phenomenon may be incompatible with cocompact group actions [HS20].

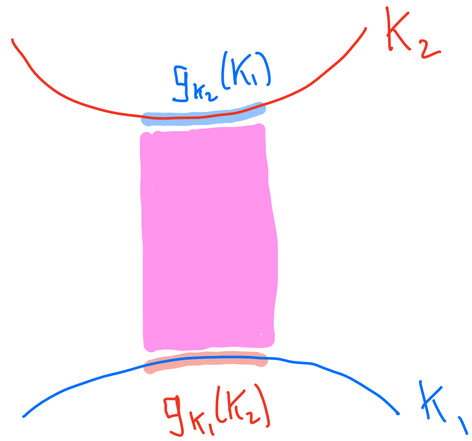

FIGURE 143. The highlighted in between region might be expected to be a product region.

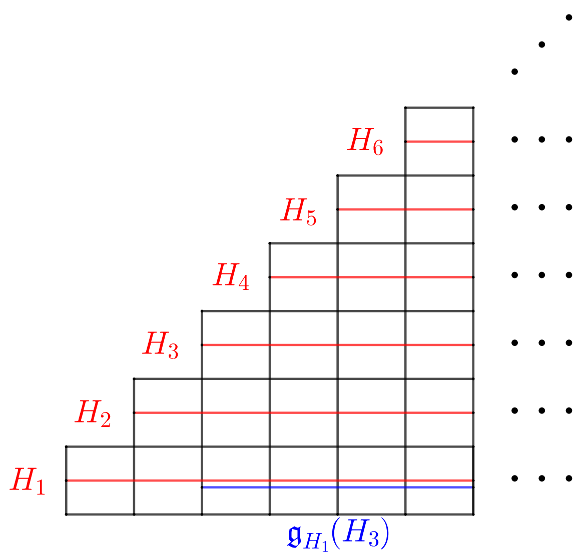

FIGURE 144. In this "staircase" example, $H_1 \supsetneq \mathfrak{g}_{H_1}(H_2) \supsetneq \mathfrak{g}_{H_1}(H_3) \supsetneq \cdots$ is an infinite nested chain.

Theorem 38.11 (Hagen-Susse)**.** *For all known examples of CAT(0) cube complexes that admit a proper cocompact group action, the hyperclosure has finite complexity.*

Optional Exercise 47. Show that, using the usual Euclidean metric on $\mathbb{R}^2$, a triple of points has a median if and only if they are collinear.

Optional Exercise 48. Show that $\mathbb{R}^2$ with the ℓ^1 (taxi-cab) metric is a median space.

Optional Exercise 49. Suppose X is a metric space with a unique geodesic joining any pair of points. Prove that X is a median space if and only if X is 0-hyperbolic.

Optional Exercise 50. Show that $\mathbb{R}^2$ with the Euclidean metric does not have the Helly property. Give a hands-on proof that a tree does satisfy the Helly property.

Optional Exercise 51. Show directly that the 1-skeleton of a d-cube is a median graph for any d.

Optional Exercise 52. Show that a median graph need not be a median space.

Optional Exercise 53. Consider the graph in Figure 145. Show that any three points has a median, but that there are triples of points with more than one median. Show that this graph is not the 1-skeleton of a CAT(0) cube complex.

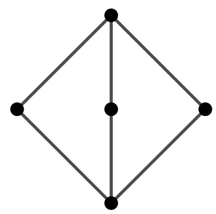

FIGURE 145

Optional Exercise 54. Suppose that K is a convex subcomplex. Show that $\mathfrak{g}(x)$ is the CAT(0) closest point in K to x.

39. THE CONTACT GRAPH (04/15, AW, GM, GUEST LECTURE BY HAGEN)

Fix a CAT(0) cube complex X. (For now, it could be infinite dimensional and/or locally infinite.)

Definition 39.1. Two hyperplanes h, v of X contact if $N(h) \cap N(v) \neq \varnothing$. Equivalently:

(1) h and v cross or osculate.

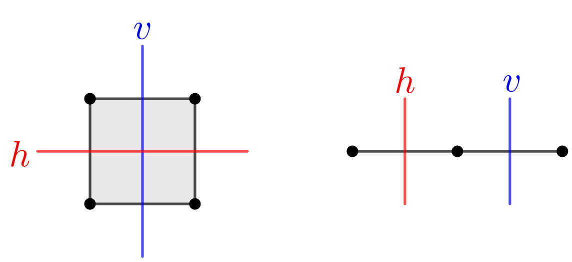

FIGURE 146. A crossing pair of hyperplanes (left) and an osculating pair (right).

(2) No hyperplane separates h and v.

Optional Exercise 55. Prove that (1) $\iff$ (2).

Definition 39.2. The contact graph CX is the graph with a vertex for each hyperplane and an edge between a pair of vertices if the corresponding hyperplanes contact.

Remark 39.3. This is almost the same thing as the graph with a vertex for each hyperplane carrier and an edge for each pair of carriers that intersect. It's not always exactly the same thing though because a subcomplex can be the carrier of two different hyperplanes (like when X is a single square, X is the carrier of both of its hyperplanes).

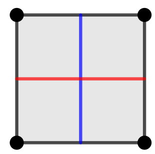

FIGURE 147. A single square is the carrier of both of its hyperplanes.

Theorem 39.4. *There exists a constant K ($K = 100$ works) such that CX is (K, K)-quasi-isometric to a tree.*

In particular, CX is hyperbolic. A reference for the theorem is [Hag14], but we'll also give a proof here.

Proof. We will verify the following condition:

($\star$) For all $h, v \in CX^{(0)}$ there exists a path α joining h to v such that if γ is another such path, then γ enters the ball $B_1^{CX}(\alpha(i))$ for all i.

Condition ($\star$) implies the theorem by Manning's Bottleneck Criterion [Man05]. We verify ($\star$) using two claims.

Claim A: If w is a hyperplane separating h, v then any path γ in CX from h to v enters $B_1^{CX}(w)$.

Proof of Claim A. Let γ be $h = h_0, h_1, \ldots, h_k = v$.

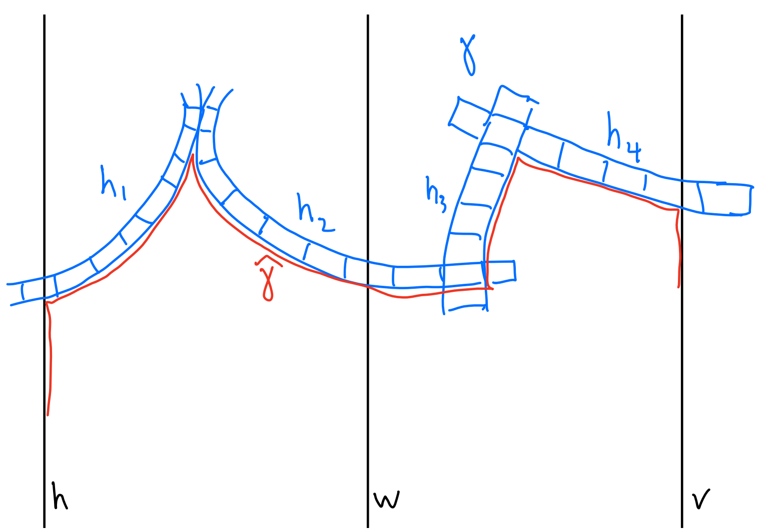

FIGURE 148

Choose $\hat{\gamma}_i$ to be paths in $N(h_i)$ that can be concatenated to form a path $\hat{\gamma} = \hat{\gamma}_0 \cdot \hat{\gamma}_1 \cdots \hat{\gamma}_k$ in X from h to v.

$\hat{\gamma}$ must intersect w, and more specifically it must contain end edge crossing w. So for some i, w crosses h_i or $w = h_i$. So $h_i \in B_1^{CX}(w)$. □

Claim B: There exists a path $\alpha : [0, n] \to CX$ such that $\alpha(0) = h$, $\alpha(n) = v$, and $\alpha(i)$ separates h and v for all $i \in \{1, \ldots, n-1\}$.

Proof of Claim B. If $h = v$ or h contacts v there is nothing to check.

Otherwise, Exercise 55 gives that there exists w such that w separates h from v. Choose such a hyperplane w with minimal distance to h; again using Exercise 55 on the pair h, w, we see that w contacts h.

Now induct on the number of hyperplanes separating h and v, or on the distance between h, and v. (Or, iterate, by next picking a hyperplane separating w and v that contacts w, etc., to build the path.) □

Remark 39.5. Claim B shows that CX is connected. (This can also be established using that it is the 1-skeleton of the nerve of the cover of X by hyperplane carriers.)

Claims A and B together imply ($\star$). □

Remark 39.6. We won't really use that CX is a quasi-tree; for us, showing CX a quasi-tree is just a convenient way to prove hyperbolicity. (There is also a 3 line proof of hyperbolicity using geodesic guessing.)

Example 39.7. If X is a tree, then hyperplanes correspond to edges and CX is the dual graph.

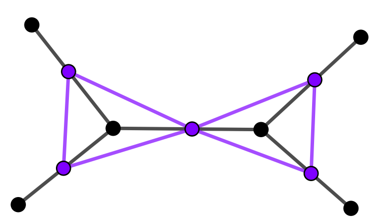

FIGURE 149. A tree in black and its dual graph in purple.

Optional Exercise 56. If $X = A \times B$, then the hyperplanes of X are partitioned into two subsets $\mathcal{H}_A, \mathcal{H}_B$, and $h \in \mathcal{H}_A, v \in \mathcal{H}_B$ implies $h \cap v \neq \emptyset$. It follows that CX is the graph theoretic join of CA and CB, and CX has diameter at most 2.

Remark 39.8. The contact graph of the staircase has diameter 3, even though the staircase isn't a product. (But, the staircase doesn't have much symmetry....)

We won't prove the following theorem, which establishes some "curve graph like" properties of CX. A reference is [Hag22].

Theorem 39.9. *Suppose X is a proper CAT(0) cube complex and G acts on X properly cocompactly with no invariant proper subcomplex. Then:*

(1) CX is infinite diamter unless X is a product.
(2) For all $g \in G$, one of the following holds:
  (a) g acts loxodromically on CX.
  (b) A power of g fixes some $h \in CX$.
  (c) X contains a subspace F that is CAT(0)-isometric to $[0, \infty) \times \mathbb{R}$ and is invariant under a power of g.

Note that (b) and (c) are not mutually exclusive, but if (a) holds then (b) and (c) do not. Case (a) corresponds to pseudo-Anosov mapping classes and case (b) corresponds to reducible mapping classes. Case (c) is a new phenemenon that doesn't occur for mapping class groups, but does occur for the staircase or even the standard cubulation of $\mathbb{R}^2$; the translation $g(x, y) = (x+1, y+1)$ on one of those spaces has the property that every orbit on CX is infinite, but since CX is finite diameter g cannot be loxodromic. (The staircase doesn't satisfy the assumptions of the theorem, since it doesn't have a cocompact action, but nonetheless this g satisfies the conclusion of Case (c).)

Theorem 39.9 was one of the first results suggesting parallels between cube complexes and mapping class groups. Similarly, the following result also suggested such parallels, since it provides a version of hierarchy paths. It is originally from [BHS17b, Proposition 3.1], but we will follow [Hag22, Section 4].

Theorem 39.10. *For any $x, y \in X^{(0)}$, there is a geodesic in $X^{(1)}$ whose image in CX tracks a geodesic.*

To parse this result, one should first note that all hyperplanes whose carriers contain a given point $x \in X^{(0)}$ are adjacent in CX, because they contact each other at x. Hence there is a coarse map $X^{(0)}$ to CX that maps points to subgraphs of diameter 1.

The proof will produce a geodesic in $X^{(1)}$ whose image in CX contains a geodesic in CX from the image of x to the image of y. As in the case of unparametrized quasi-geodesics, the geodesic in $X^{(1)}$ does not move at unit speed in CX, but at least its progress along the geodesic in CX is monotone and it does not backtrack.

Proof. Start by considering an arbitrary geodesic $h_0, h_1, \ldots, h_n$ in CX with $x \in N(h_0)$ and $y \in N(h_n)$. Choose a geodesic $\hat{\gamma}_i$ in each $N(h_i)$ such that the concatenation $\hat{\gamma} = \hat{\gamma}_0 \cdot \hat{\gamma}_1 \cdots \hat{\gamma}_n$ is a path from x to y.

There may be different options for the geodesic $h_0, h_1, \ldots, h_n$ in CX as well as the geodesics $\hat{\gamma}_i$ in $N(h_i)$. Consider the choice of all this data such that

$$(|\hat{\gamma}_0|, |\hat{\gamma}_1|, \ldots, |\hat{\gamma}_n|)$$

is LEX-minimal. This means that we prioritize $\hat{\gamma}_0$ being as short as possible, and among choices with $\hat{\gamma}_0$ as short as possible we prioritize choices with $\hat{\gamma}_1$ as short as possible, etc.

We now claim that $\hat{\gamma}$ is a geodesic in $X^{(1)}$. Since by assumption $h_0, h_1, \ldots, h_n$ is a geodesic, and each point of $\hat{\gamma}$ has image containing one of the h_i, this claim will prove the theorem. By Corollary 37.15, to show $\hat{\gamma}$ is a geodesic it is nescessary and sufficient to show it intersects each hyperplane at most once.

So suppose $\hat{\gamma}$ crosses a hyperplane w more than once, say at the edge e and then next at the edge f. Say $e \in N(h_i)$ and $f \in N(h_j)$. Note that we can assume $i \neq j$ because each carrier is convex and the $\hat{\gamma}_k$ are geodesics. So assume $i < j$. We will also assume that e is the first time γ crosses w.

Case 1: $j = i + 1$. In this case we will derive a contradiction. As in Figure 150, let A be the first vertex of $\hat{\gamma}_i$, B be the last vertex of $\hat{\gamma}_i$, and let C be the vertex of $f \subset \hat{\gamma}$ immediately after the second crossing with w.

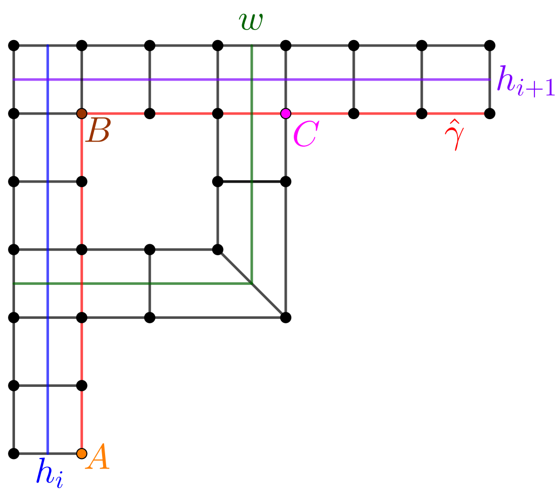

FIGURE 150. Case 1.

We now apply Lemma 38.5 to the median $m = m(A, B, C)$. Since A and B are in $N(h_i)$, we get that m must be in $N(h_i)$, and similarly m must be in $N(h_j)$. Moreover, A and C are in a half-space of w, and since half-spaces are convex, we get that m must be in this halfspace.

In particular, we have that $m \neq B$, since m and B are on different sides of w. By definition of median, there is a geodesic $\hat{\gamma}_i'$ from A to B that passes through m, so here we conclude that $d(A, m) < d(A, B)$. Since $m \in N(h_i)$ and $N(h_i)$ is convex, this geodesic $\hat{\gamma}_i'$ lies in $N(h_i)$. Since m is also in $N(h_j)$, this contradicts the LEX-minimality assumption.

Case 2: $j > i + 1$ and $w \neq h_{j+1}$. Since $h_0, h_1, \ldots, h_n$ is a geodesic in CX, it follows that that $j = i+2$ and $h_0, \ldots, h_i, w, h_{i+2}, \ldots, h_n$ is also a geodesic in CX. As Figure 151 indicates, it is possible to pick a new $\hat{\gamma}$ contradicting the LEX-minimality assumption.

Case 3: $j > i + 1$ and $w = h_{j+1}$.

Optional Exercise 57. Complete Case 3, using Figure 152 as a guide.

This concludes the proof. □

Optional Exercise 58. Define the crossing graph $C'X$ of X to be the subgraph of CX with the same vertex set, but with edges only when two hyperplanes cross. (So osculation gives edges in CX but not $C'X$.) Show that CX' is connected if and only if X does not have any cut points (vertices whose removal disconnects X). In the later case, show it is quasi-isometric to CX. Show that the link of a vertex h in $C'X$ is $C'h$.

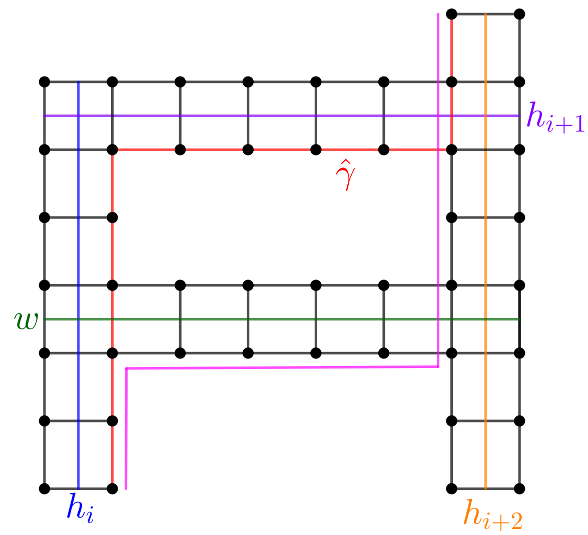

FIGURE 151. Case 2.

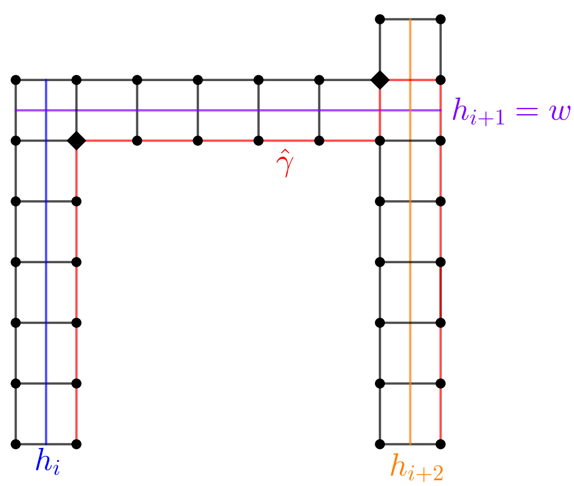

FIGURE 152. Case 3. The diamonds indicate the division of the path $\hat{\gamma}$ into the geodesics it is a concatenation of.

Optional Exercise 59. Suppose X has hyperplanes $h_0, h_1, \ldots, h_d$, such that for all i we have that $h_0, \ldots, h_{i-1}$ are on one side of h_i, and $h_{i+1}, \ldots, h_d$ are on the other. Suppose also that for each i, there does not exist a hyperplane crossing both h_i and h_{i+1}. Prove that the distance in $C'X$ between h_0 and h_d is at least d. Use this to show that $C'X$ has infinite diameter for the the cube complex illustrated in Figure 153. (This exercise can be easily adapted to CX if you prefer it to $C'X$.)

40. The HHS structure (04/18, AW, SK, guest lecture by Hagen)

Let X be a CAT(0) cube complex, and let $Y \subset X$ a convex subcomplex.

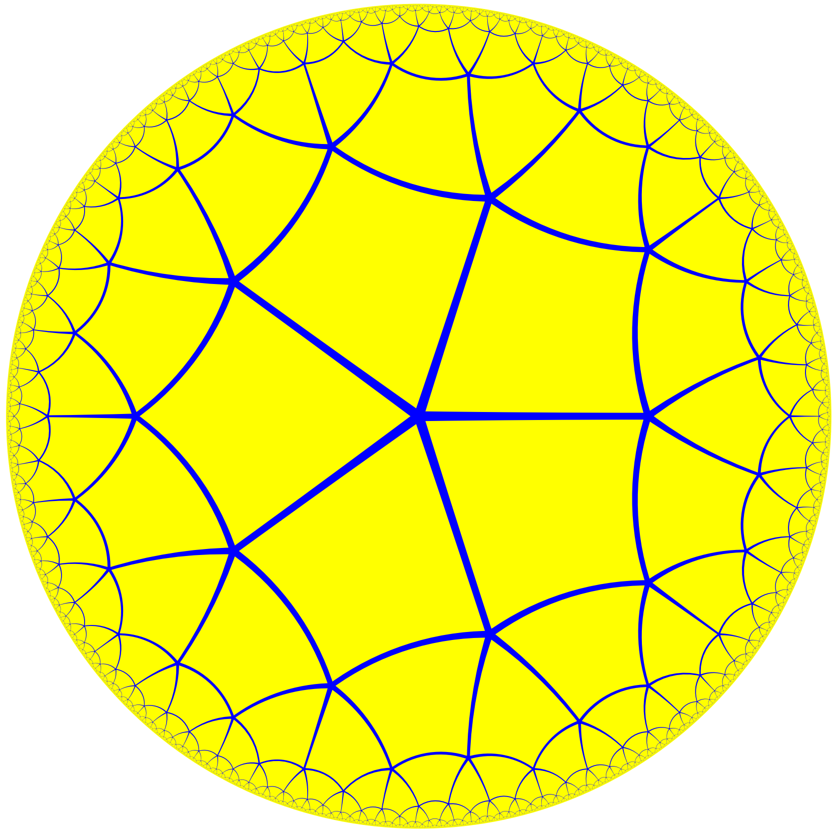

FIGURE 153. A cubulation of the hyperbolic plane (from Wikipedia).

Observation 40.1. *Y is a CAT(0) cube complex (since a missing cube would violate convexity). The hyperplanes of Y have the form $h \cap Y$, where h is a hyperplane of X intersecting Y.*

The assignment

$$h \cap Y \mapsto Y$$

induces a graph map

$$CY \to CX.$$

Lemma 40.2. *This map is injective, and its image is an induced graph.*

An induced graph is a subgraph containing all edges between its vertices. As a result of the lemma, we think of CY as a subgraph of CX.

Proof. It is injective simply because if $h \cap Y$ contacts $v \cap Y$, then h contacts v.

To see that it is an induced graph, we need to show that if h and v both intersect Y, and they contact somewhere in X, then in fact they contact in Y. This follows from the Helly property: since the convex sets $N(h), N(v)$, and Y pairwise intersect, Corollary 38.6 gives that $N(h) \cap N(v) \cap Y$ is nonempty. □

Observation 40.3. *There exists a uniform constant K such that CY is K-quasi-convex in CX.*

Proof. Consider $h, v \in CY$. Apply Theorem 39.10 to find a geodesic $h = h_0, h_1, \ldots, h_n = v$ in CX and a geodesic $\hat{\gamma} = \hat{\gamma}_0 \cdots \hat{\gamma}_n$ in $X^{(1)}$ formed from concatenating geodesics $\hat{\gamma}_i \in N(h_i)$, as in Figure 154.

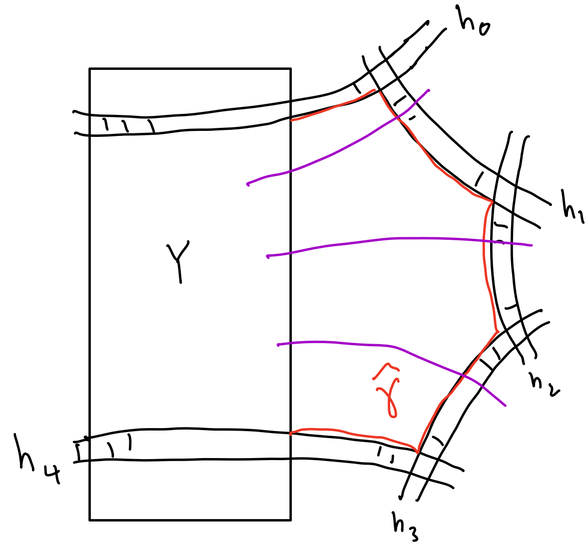

FIGURE 154. Each edge of $\hat{\gamma}$ (red) defines a unique hyperplane (purple) crossing it. Contrary to the image, the proof shows that in fact each edge of $\hat{\gamma}$ lies in CY, so these hyperplanes intersect CY. Each purple hyperplane thus represents a vertex in CX connected to both CY and $\hat{\gamma}$, showing that $\hat{\gamma}$ stays close to CY.

Since Y is convex, $\hat{\gamma} \subset Y$. So already the picture in Figure 154 is not accurate; actually $\hat{\gamma}$ lies in Y, possibly along the boundary of Y. Each edge of $\hat{\gamma}_i$ defines a hyperplane which thus intersects Y (illustrated in purple), so $h_0, h_1, \ldots, h_n$ lies in the 2-neighbourhood of CY.

We have produced a geodesic joining any pair of points in CY that stays in the 2-neighbourhood of CY. Since any two geodesics with the same end points fellow travel in a hyperbolic space, this gives the result. □

Corollary 40.4. *If $\{Y_i\}$ is any collection of convex subcomplexes of X and $\widehat{CX}$ is obtained from CX by coning off each $CY_i \subset CX$, then $\widehat{CX}$ is a quasi-tree (with uniform constants).*

Optional Exercise 60. Prove the corollary using the Bottleneck Criterion.

It isn't essential to do this exercise, since all we need is that $\widehat{X}$ is hyperbolic and that also follows from Proposition 8.1.

Recall that the hyperclosure $\mathcal{F}$ of X is defined to be the smallest family of subcomplexes of X such that

(1) $X \in \mathcal{F}$,
(2) for each hyperplane h, the subcomplexes $h \times \{-\frac{1}{2}\}$ and $h \times \{\frac{1}{2}\}$ are in $\mathcal{F}$, and

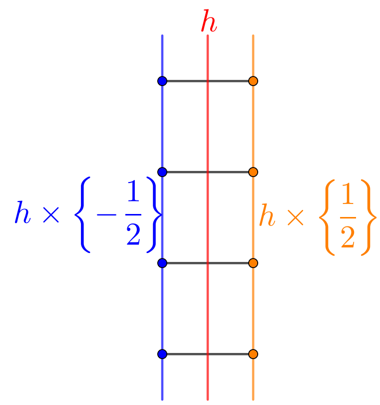

FIGURE 155. $N(h) = h \times [-\frac{1}{2}, \frac{1}{2}]$.

(3) if $F, F' \in \mathcal{F}$ then $\mathfrak{g}_F(F') \in \mathcal{F}$.

All the subcomplexes in $\mathcal{F}$ are convex.

Definition 40.5. Convex subcomplexes F, F' are parallel if the restriction

$$\mathfrak{g}_F : F' \to F$$

is a (cubical) isomorphism.

Optional Exercise 61. F and F' are parallel if and only if the set of hyperplanes intersecting F is equal to the set of hyperplanes intersecting F'.

The exercise implies that F and F' are parallel if and only if $CF = CF'$ as subgraphs of CX.

Definition 40.6. The index set $\mathfrak{S}$ is defined to be $\mathcal{F}$ modulo parallelism. If $[F], [F'] \in \mathfrak{S}$, say $[F] \sqsubseteq [F']$ if $F \subset F'$ up to parallism.

A fundamental *assumption* going forward is that $\sqsubseteq$-chains have bounded length. The staircase illustrated in Figure 144 shows this does not hold for all CAT(0) cube complexes, but Theorem 38.11 shows that it holds in all known examples of CAT(0) cube complexes which are universal covers of compact NPC cube complexes.

Definition 40.7. $[F] \perp [F']$ if, up to replacing F and F' with things in their parallelism class, the convex hull of $F \cup F'$ is isomorphic to $F \times F'$ as cube complexes.

Optional Exercise 62. $[F]$ and $[F']$ are orthogonal if and only if, for every hyperplane h crossing F and every hyperplane h' crossing F', we have that h crosses h'.

Our next topic is the factored contact graph.

Definition 40.8. For each $V \in \mathcal{F}$, let $\mathcal{F}_V = \{[W] \subsetneq [V]\}$. Define $\widehat{CV}$ to be the result of coning off all $CW \subset CV$ for each $[W] \in \mathcal{F}_V$.

Keep in mind that if V and V' are parallel, then $CV = CV'$ as subgraphs of CX. So really $\widehat{CV}$ only depends on $[V] \in \mathfrak{S}$.

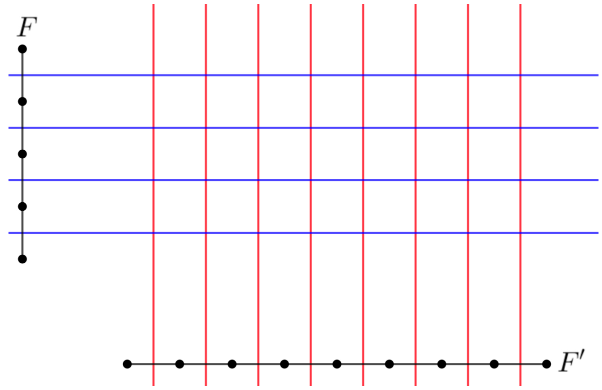

FIGURE 156

Definition 40.9. For any $[V] \in \mathfrak{S}$, define $\pi_{[V]} : X \to \widehat{CV}$ to be the composite

$$X \to V \to CV \hookrightarrow \widehat{CV}$$

of the gate map $\mathfrak{g}_V : X \to V$, the usual map $V \to CV$ defined by $x \mapsto \{h : x \in N(h)\}$ and the inclusion of CV into its electrification $\widehat{CV}$.

Optional Exercise 63. Check that the map $\pi_{[V]}$ does not depend on the choice of V in its parallelism class, possibly by showing that Figure 157 is accurate.

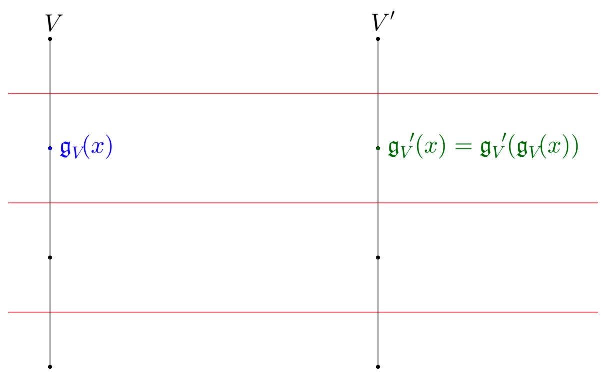

FIGURE 157

We're now in a good position to discuss the HHS axioms, although we won't have time for complete proofs. The context remains CAT(0) cube complexes where the hyperclosure $\mathcal{F}$ is assumed to have an upper bound on the length of a nested chain.

1. **(Projections)** The $\widehat{CV}$ are not only hyperbolic but are even quasi-trees, and the projection maps are the π_V above. The index set $\mathfrak{S}$ is the hyperclosure $\mathcal{F}$ modulo parallelism.

2. **(Nesting)** Nesting is as defined above.
   Suppose $V \subsetneq W$. We define ρ^V_W to be $\pi_W(V)$, which is a coarse point since $C(\mathfrak{g}_W(V))$ has been coned off in the passage from CW to $\widehat{CW}$.

The definition of the "downwards ρ maps" ρ_V^W is less important than you might think, since if one uses the equivalent variant of the HHS axioms given by [BHS19, Proposition 1.11], it is not required to define these maps. (But they are not hard to define: a subcomplex $T \subset W$ representing a vertex of $\widehat{CW}$ should map to the set of hyperplanes crossing both T and V, which via intersection can be interpreted as hyperplanes of V.)

3. **(Orthogonality)** Orthogonality is as defined above.

4. **(Transversality and Consistency)** We only discuss the Behrstock inequality here. Say $V \pitchfork W$, and define $\rho_W^V = \pi_W(V)$. This is a coarse point since $C(\mathfrak{g}_W(V))$ is coned off in the passage from CW to $\widehat{CW}$.

   We now sketch a proof of Behrstock. Suppose that $x \in X^{(0)}$ and $\pi_V(x)$ is far from ρ_V^W. It follows that $\mathfrak{g}_V(x)$ is far from $\mathfrak{g}_V(W)$. So there must be at least one hyperplane a separating $\mathfrak{g}_V(x)$ from $\mathfrak{g}_V(W)$. Note that previous arguments

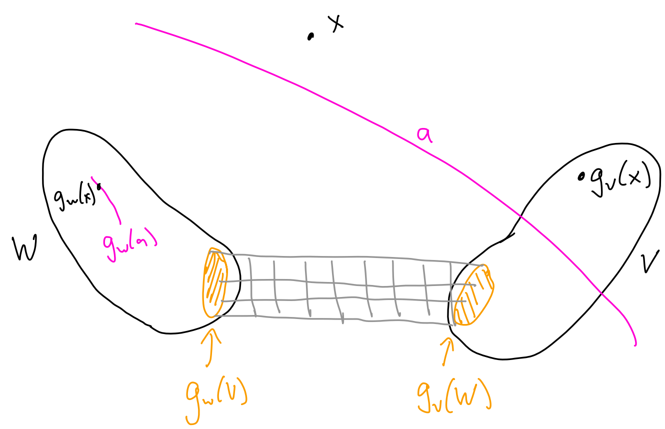

FIGURE 158

   we have done imply that x and $\mathfrak{g}_V(x)$ are on the same side of a. Since the hyperplanes intersecting both V and W are exactly those intersecting $\mathfrak{g}_V(W)$ and $\mathfrak{g}_W(V)$, we see that a does not intersect W. Since $\mathfrak{g}_V(x)$ and $\mathfrak{g}_V(W)$ are on opposite sides of a, it is also possible to show that x and W are on opposite sides of a. It follows that $\mathfrak{g}_W(x)$ is contained in $\mathfrak{g}_W(a)$, which is coned to a point in the passage from CW to $\widehat{CW}$. Since $\mathfrak{g}_W(V)$ intersects $\mathfrak{g}_W(a)$, we see that $\pi_W(x)$ is close to ρ_W^V in $\widehat{CW}$.

5. **(Finite Complexity)** This holds by assumption; we are excluding CAT(0) cube complexes like the staircase where it doesn't hold.

6. **(Large Links)** For both this axiom and also Uniqueness, it's helpful to have a version of Theorem 39.10 adapted to the coned off contact graphs rather than the original contact graphs.

Fix $x, y \in X$. A geodesic joining the projections of x and y to $\widehat{CX}$ can be thought of as a sequence $T_0, T_1, T_2, \ldots, T_n$, where each T_i is an element of the hyperclosure $\mathcal{F}$: indeed, the geodesic is a sequence of vertices, and each vertex is either a hyperplane or a cone-point over some element in the hyperclosure. Moreover, if T_i, T_{i+1} are hyperplanes, their carriers intersect. Otherwise T_i is a hyperplane and T_{i+1} corresponds to a cone-point (or vice versa), and T_{i+1} is a subcomplex in $\mathcal{F}$ crossed by the hyperplane T_i.

So we can make a path $\hat{\gamma}$, from x to y, by concatenating paths $\hat{\gamma}_i$ that lie in $N(T_i)$ (if T_i is a hyperplane) or in T_i itself (if it's a subcomplex corresponding to a cone-point).

Now very similarly to the proof of Theorem 39.10, we can choose the $\widehat{CX}$–geodesic and the $\hat{\gamma}_i$ so that $\hat{\gamma}$ is a geodesic in X. In short, you still get geodesics in X carried by geodesics in $\widehat{CX}$.

To address the Large Links axiom, we now claim that if $F \in \mathcal{F}$ is such that x, y project far apart on $\widehat{CF}$, then F is nested in one of the T_i.

To prove that, first note that $\mathfrak{g}_F(x)$ and $\mathfrak{g}_F(y)$ are far apart; this follows from the definition of π_F and the fact that all the composed maps in that definition are coarsely Lipschitz.

Note that if $j > i + 4$, it cannot be the case that a hyperplane a crosses both T_i and F and a hyperplane b crosses both T_j and F. This is because otherwise the path $T_0, \ldots, T_n$ could be replaced with the shorter path

$$T_0, \ldots, T_i, a, F, b, T_j, \ldots, T_n,$$

contradicting the fact that $T_0, \ldots, T_n$ was a geodesic in $\widehat{CX}$.

In particular, this means that at most 5 values of j can be such that a hyperplane crosses both T_j and F. (All that matters is that the number 5 is bounded.)

Because $\mathfrak{g}_F(x)$ and $\mathfrak{g}_F(y)$ are far apart, there have to be many hyperplanes crossing F and separating those two points. For any such hyperplane a, one can show that $\mathfrak{g}_F(x)$ is on the same side of a as x, and similarly for y. Thus, each of these hyperplanes has to cross the geodesic $\hat{\gamma}$ at some point.

If F is nested in some T_j, we're done. Otherwise, we can use the fact that the projections of each of the at most 5 relevant T_i to F get coned off in $\widehat{CF}$ to so that actually the $\widehat{CF}$ distance between the projections of x and y was must have been small.

7. (Bounded Geodesic Image Axiom) Due to time constraints, we'll only give a hint of where BGI comes from. Suppose $H \in \mathcal{F}$, and $\mathfrak{g}_H(x)$ and $\mathfrak{g}_H(y)$ are far apart. Then, using Proposition 38.10, we can find a hyperplane a that has x and $\mathfrak{g}_H(x)$ on one side, and y and $\mathfrak{g}_H(y)$ on the other. We can find a geodesic between the images of x and y in CX that contains a. Note that the image of a in $\widehat{CX}$ is contained in the coarse point ρ^H_X since a is contained in CH.

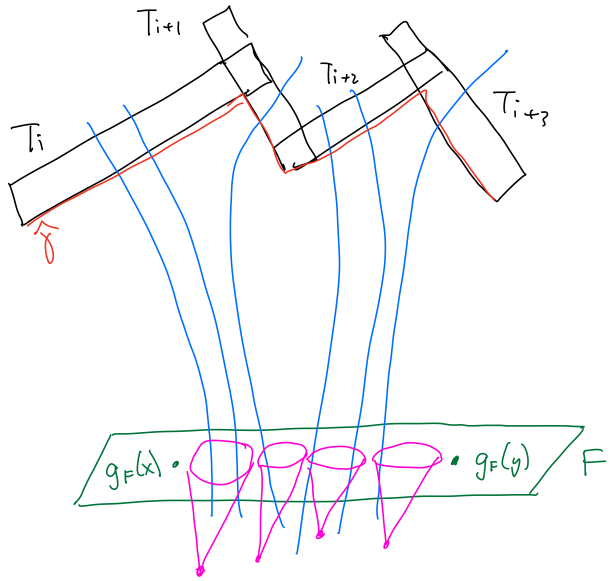

FIGURE 159. There may be many blue hyperplanes, but, in $\widehat{CF}$, one can hop over them all in at most 5 hops through cone points.

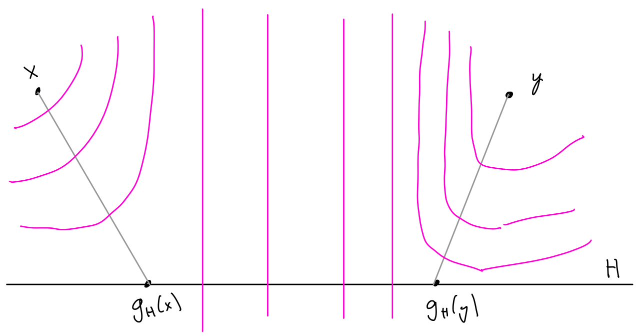

FIGURE 160

8. (Partial Realization) Suppose $[V_1], \ldots, [V_k] \in \mathfrak{S}$ are pairwise orthogonal, and choose $h_i \in \widehat{CV}_i$.

We will use the following fact: X contains an subcomplex isomorphism to $V_1 \times V_2 \times \cdots \times V_k$. For example, suppose $k = 3$. The definition of orthogonality gives subcomplexes isomorphic to $V_1 \times V_2, V_2 \times V_3, V_1 \times V_3$. To prove the fact, one would show one can translate them so that all three intersect at a common point. (One approach to do that is to apply the Helly property to the pairwise intersecting sets $V_1 \times V_1^{\perp}, V_2 \times V_2^{\perp}, V_3 \times V_3^{\perp}$.) Then one can imagine that to be CAT(0) requires filling in the product.

Now, for each i, choose $x_i \in h_i \cap V_i$, and let $x = (x_1, \ldots, x_k)$. Then $\pi_{V_i}(x) \in h_i$ for each i, proving the first part of partial realization.

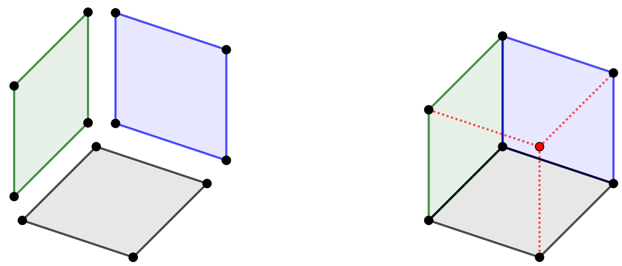

FIGURE 161

9. **(Uniqueness)** Using the notation from our discussion of Large Links, if the number of n of vertices $T_0, \ldots, T_n$ in the carefully chosen geodesic from $\pi_X(x)$ to $\pi_X(y)$ is large, we're done. Otherwise, a large part of $\hat{\gamma}$ must lie in some individual T_i, and we can induct or iterate to get the result.

We end by mentioning some options for further reading on the topic of this lecture. For more on the HHS structure on cube complexes, see [BHS17b]. For more on parallelism, the hyperclosure, and how it interacts with group actions, see [HS20]. For a recent concrete example of factor systems in action, see [Che20].

Appendix A. Hints for some exercises

Exercise 4: SK's idea: Consider many disjoint balls of radius 1 can you fit in a ball of radius R in both the spaces, as a function of R.

Exercise 5: For a geodesic n-gon in a δ-hyperbolic space, each point on an edge should be within roughly $\delta \log_2(n)$ of a different edge.

Exercise 6: For any space quasi-isometric to the given wedge, and for any $R > 0$, the space contains an infinite set of points all of distance at least R to each other, and all contained in a bounded set.

Exercise 11: First figure out what the half-infinite geodesic rays are.

Exercise 20: Suppose α and β are primitive geodesics. Say A is a lift of α to the unviversal cover $\mathbb{H}$. Consider all lifts of β whose endpoints are intertwined with those of A. Try to show that the intersection number is the number of $\mathrm{Stab}(A)$ orbits in this set.

Exercise 21: Start with the disconnected cover $S \times \{1, ..., D\}$, and try to modify it to make it connected. First modify the geodesic in the curve complex so it consists of separating curves, then cut the copies of the $(i-1)$-st and i-th curves on $S \times \{i\}$, and then glue all the pieces together in the right way.

Exercise 22: The intersection of all subgroups of $\pi_1(S)$ of index at most D is a finite index subgroup of $\pi_1(S)$.

Exercise 25: For each cylinder, we can look at the subset of Teichmüller disc where it has modulus at least some fixed large constant. This region is a horoball, and in this region the core curve of the cylinder is guaranteed to have small hyperbolic length. (These are basic facts in the study of quadratic differentials.) Each of these horoballs maps to a bounded diameter region of the curve complex. Considering two horoballs that overlap or come close to each other, as in Figure 162, gives a subset of the Teichmüller disc that disconnects the disc and yet maps to a bounded diameter region of the curve complex. This allows one to apply the bottleneck criterion.

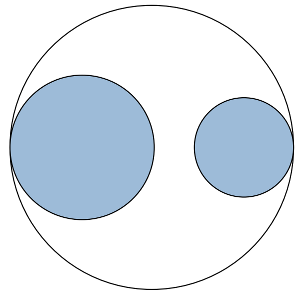

FIGURE 162. Two horoballs in this configuration, together with (say) a minimal length geodesic from one to the other, disconnects the hyperbolic plane.

Exercise 26: Compare to the point pushing map.

Exercise 28: The image of a QI embedding between geodesic metric spaces always contains a quasi-geodesic between any two points. Since geodesics stay close to quasi-geodesics in a hyperbolic space, this shows that the image is quasi-convex.

Exercise 29: Let β be a point of $\mathcal{C}S - \{\alpha\}$, and let γ be a geodesic from β to a point in the closest point projection of β to $\mathcal{S}(S-\alpha)$. The Bounded Geodesic Image Theorem says that the projection onto $\mathcal{C}(S-\alpha)$ can't change much along most of this geodesic. One can use that subsurface projections are Lipschitz to handle the last little bit of the geodesic.

Exercise 30: Consider the Uniqueness Axiom.

Exercise 33: Use the Uniqueness Axiom to directly construct a coarse inverse, and use this coarse inverse to show π_S is a quasi-isometry.

Exercise 34: BZ's idea: Use the Large Links Axiom plus some sort of induction or iteration. (If you also use that the π_U are Lipschitz, you should be able to show that the number of non-zero terms in the distance formula sum with sufficiently large threshold is at most a polynomial in the distance in the HHS, with the degree of the polynomial being something like the complexity of the HHS. Since each term is coarsely at most the distance between two points, this gives a not very good upper bound on the whole sum.)

Exercise 38: Use some sort of induction on complexity. Say you have already chosen $U_1, \ldots, U_k$ and they're all transverse to each other. There are infinitely many domains in your set, so one of the following three things happen: (a) there are infinitely many domains in your set transverse to all U_i, (b) there is some U_i so that there are infinitely many domains in your set nested in U_i, or (c) there are infinitely many domains in your set orthogonal to some U_i. If (a), just keep picking. If (b), throw out everything you've done so far and replace your set with the subset of domains nested in U_i. If (c), throw out everything you've done so far and replace your set with the subset of domains orthogonal to U_i. The point is that (b) and (c) reduce some notion of complexity, so those events can only happen finitely many times, and eventually you end up with (a) forevermore.

Exercise 43: You just need to show that, for any point on the geodesic from $\pi_S(x)$ to $\pi_S(y)$, you can define a point z with $\pi_S(z)$ equal to that point, and all other coordiantes the same as for x and y.

Exercise 44: Given two geodesics between a pair of points, arbitrarily divide one edge into two pieces to obtain a (degenerate) triangle.

Exercise 45: Use the previous exercise, and define a unit speed retraction along geodesics towards an arbitrarily chosen point.

Exercise 46: Compare to Example 37.8 and Figure 129.

Exercise 50: The sides of a triangle are convex sets.

Exercise 51: Given $(a_i), (b_i), (c_i) \in \{0,1\}^d$, try defining (m_i) by $m_i = 0$ if at least 2 of a_i, b_i, c_i are 0, and otherwise $m_i = 1$.

Exercise 52: Consider a single square and its 1-skeleton, which is a 4-cycle graph and hence is isometric to a circle.

Exercise 54: Let H be a hyperplane separating x and $\mathfrak{g}(x)$. Let $z \in K$. Then there is a geodesic from x to z through $\mathfrak{g}(x)$. Since a geodesic can cross each hyperplane at most once, we see that H separates x and z. Since this is true for all z, it follows that H separates x and K. Use this, and the fact that, even in the CAT(0) metric, it takes at least distance 1 to cross a hyperplane carrier.

Exercise 55: (1) $\implies$ (2) is clear. For the converse, suppose that v and h do not contact each other. Pick $x \in N(h)$, $y \in N(v)$ vertices that minimize $d(x,y)$. So $x = \mathfrak{g}_{N(h)}(y)$ and $y = \mathfrak{g}_{N(v)}(x)$. Pick w that separates x and y.

For any vertex z of $N(v)$ there is a geodesic from x to z through y. This geodesic must intersect w at least once, and since it is a geodesic it intersects w exactly once. Hence x and w are on opposite sides of w. Also using the symmetric statement, this shows w separates v and h.

Exercise 56: For the first part, consider a hyperplane in $A \times B$, and consider an edge e crossing the hyperplane. Edges of a product $A \times B$ are either a point in A cross and edge in B or the same with A and B swapped.

Exercise 59: Once you have done the first part, to show the given example of $C'X$ of infinite diameter, it might be helpful to start by finding h_0, h_1 as in the first part. (Maybe try drawing h_0 and h_1 on the figure.) Then pick an appropriate automorphism $g : X \to X$ of the cube complex sending h_0 to h_1, and set $h_i = g^i h_0$.

Exercise 61: First suppose F and F' are parallel. Let h be a hyperplane crossing F', and let $x, y \in (F')^{(0)}$ be vertices of an edge crossing h. Use Proposition 38.10 to show that h must separate $\mathfrak{g}_F(x)$ and $\mathfrak{g}_F(y)$ and hence show h intersects y. After that, let h be a hyperplane crossing F, and let $x, y \in (F')^{(0)}$ be such that $\mathfrak{g}_F(x)$ and $\mathfrak{g}_F(y)$ are vertices of an edge crossing h, and use a similar argument.

Next suppose F and F' have the same crossing hyperplanes. If $\mathfrak{g}_F$ is not surjective, consider a hyperplane not in its image. If it is not injective, consider a hyperplane separating two points that map to the same point.

Exercise 62: Because of Exercise 61, if the statement holds for one representative of the two parallism classes, it holds for all of them.

First suppose $[F]$ and $[F']$ are orthogonal and that there is a cubical isomorphism from the convex hull of $F \cup F'$ to $F \times F'$. Then it is easy to see the claim about hyperplanes intersecting, since an intersection can be found in $F \times F'$.

The harder direction is to assume the result about hyperplanes intersecting, and try to show that the convex hull is a product. For this, it may be nescessary to replace F and F' with parallel complexes, and MH's hint for this was to, ex, take the gate of F on the closest hyperplane to F that crosses F'.

Exercise 63: MH's hint: It may be helpful to first consider a gate map onto a product.

References

[ABD21] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham, *Largest acylindrical actions and stability in hierarchically hyperbolic groups*, Trans. Amer. Math. Soc. Ser. B **8** (2021), 66–104, With an appendix by Daniel Berlyne and Jacob Russell.

[AOS12] Federico Ardila, Megan Owen, and Seth Sullivant, *Geodesics in* CAT(0) *cubical complexes*, Adv. in Appl. Math. **48** (2012), no. 1, 142–163.

[Aou13] Tarik Aougab, *Uniform hyperbolicity of the graphs of curves*, Geom. Topol. **17** (2013), no. 5, 2855–2875.

[Beh06] Jason A. Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space*, Geom. Topol. **10** (2006), 1523–1578.

[BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999.

[BHS17a] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups*, Proc. Lond. Math. Soc. (3) **114** (2017), no. 5, 890–926.

[BHS17b] ———, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804.

[BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula*, Pacific J. Math. **299** (2019), no. 2, 257–338.

[Bow14a] Brian H. Bowditch, *Uniform hyperbolicity of the curve graphs*, Pacific J. Math. **269** (2014), no. 2, 269–280.

[Bow14b] ———, *Uniform hyperbolicity of the curve graphs*, Pacific J. Math. **269** (2014), no. 2, 269–280.

[Che20] Marissa Chesser, *Stable subgroups of the genus two handlebody group*, arXiv preprint arXiv:2009.05067 (2020).

[CRS14] Matt Clay, Kasra Rafi, and Saul Schleimer, *Uniform hyperbolicity of the curve graph via surgery sequences*, Algebr. Geom. Topol. **14** (2014), no. 6, 3325–3344.

[DDM14] Spencer Dowdall, Moon Duchin, and Howard Masur, *Statistical hyperbolicity in Teichmüller space*, Geom. Funct. Anal. **24** (2014), no. 3, 748–795.

[DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto, *Boundaries and automorphisms of hierarchically hyperbolic spaces*, Geom. Topol. **21** (2017), no. 6, 3659–3758.

[DT17] Spencer Dowdall and Samuel J. Taylor, *The co-surface graph and the geometry of hyperbolic free group extensions*, J. Topol. **10** (2017), no. 2, 447–482.

[Dur16] Matthew Gentry Durham, *The augmented marking complex of a surface*, J. Lond. Math. Soc. (2) **94** (2016), no. 3, 933–969.

[Far98] B. Farb, *Relatively hyperbolic groups*, Geom. Funct. Anal. **8** (1998), no. 5, 810–840.

[FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012.

[GM08] Daniel Groves and Jason Fox Manning, *Dehn filling in relatively hyperbolic groups*, Israel J. Math. **168** (2008), 317–429.

[Hag] Mark Hagen, *Cat(0) cube complex, median graphs, and cubulating groups*, `https://www.wescac.net/into_the_forest.pdf`.

[Hag14] Mark F. Hagen, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, J. Topol. **7** (2014), no. 2, 385–418.

[Hag22] Mark Hagen, *Large facing tuples and a strengthened sector lemma*, Tunis. J. Math. **4** (2022), no. 1, 55–86.

[Hem01] John Hempel, *3-manifolds as viewed from the curve complex*, Topology **40** (2001), no. 3, 631–657.

[HPW15] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, J. Eur. Math. Soc. (JEMS) **17** (2015), no. 4, 755–762.

[HS20] Mark F. Hagen and Tim Susse, *On hierarchical hyperbolicity of cubical groups*, Israel J. Math. **236** (2020), no. 1, 45–89.

[KR14] Ilya Kapovich and Kasra Rafi, *On hyperbolicity of free splitting and free factor complexes*, Groups Geom. Dyn. **8** (2014), no. 2, 391–414.

[Man05] Jason Fox Manning, *Geometry of pseudocharacters*, Geom. Topol. **9** (2005), 1147–1185.

[Man10] Johanna Mangahas, *Uniform uniform exponential growth of subgroups of the mapping class group*, Geom. Funct. Anal. **19** (2010), no. 5, 1468–1480.

[Man13] ———, *A recipe for short-word pseudo-Anosovs*, Amer. J. Math. **135** (2013), no. 4, 1087–1116.

[Min96] Yair N. Minsky, *Extremal length estimates and product regions in Teichmüller space*, Duke Mathematical Journal **83** (1996), no. 2, 249 – 286.

[Min03] Yair N. Minsky, *Combinatorial and geometrical aspects of hyperbolic 3-manifolds*, Kleinian groups and hyperbolic 3-manifolds (Warwick, 2001), London Math. Soc. Lecture Note Ser., vol. 299, Cambridge Univ. Press, Cambridge, 2003, pp. 3–40.

[MM99] Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149.

[MM00] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. **10** (2000), no. 4, 902–974.

[MS13] Howard Masur and Saul Schleimer, *The geometry of the disk complex*, J. Amer. Math. Soc. **26** (2013), no. 1, 1–62.

[MS20] John M. Mackay and Alessandro Sisto, *Quasi-hyperbolic planes in relatively hyperbolic groups*, Ann. Acad. Sci. Fenn. Math. **45** (2020), no. 1, 139–174.

[PS17] Piotr Przytycki and Alessandro Sisto, *A note on acylindrical hyperbolicity of mapping class groups*, Hyperbolic geometry and geometric group theory, Adv. Stud. Pure Math., vol. 73, Math. Soc. Japan, Tokyo, 2017, pp. 255–264.

[Raf07] Kasra Rafi, *A combinatorial model for the Teichmüller metric*, Geom. Funct. Anal. **17** (2007), no. 3, 936–959.

[Raf14] ———, *Hyperbolicity in Teichmüller space*, Geom. Topol. **18** (2014), no. 5, 3025–3053.

[Sch] Saul Schleimer, *Notes on the complex of curves*, `https://homepages.warwick.ac.uk/~masgar/Maths/notes2.pdf`.

[Sisa] Alessandro Sisto, *Guessing geodesics in hyperbolic spaces*, `https://alexsisto.wordpress.com/2012/07/16/bi-guessing-geodesics-in-hyperbolic-spaces/`.

[Sisb] ______, *Lecture notes on geometric group theory*, `http://www.macs.hw.ac.uk/~as2000/LectureNotesGGT.pdf`.

[Sis19] ______, *What is a hierarchically hyperbolic space?*, Beyond hyperbolicity, London Math. Soc. Lecture Note Ser., vol. 454, Cambridge Univ. Press, Cambridge, 2019, pp. 117–148.

[Spr17] Davide Spriano, *Hyperbolic HHS I: Factor systems and quasi-convex subgroups*, arXiv preprint arXiv:1711.10931 (2017).

[Vor05] Yaroslav Vorobets, *Periodic geodesics on generic translation surfaces*, Algebraic and topological dynamics, Contemp. Math., vol. 385, Amer. Math. Soc., Providence, RI, 2005, pp. 205–258.

[Wis12] Daniel T. Wise, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.

[Wis21] ______, *The structure of groups with a quasiconvex hierarchy*, Annals of Mathematics Studies, vol. 209, Princeton University Press, Princeton, NJ, [2021] ©2021.