Covering maps and hulls in the curve complex

by

Robert Tang

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Mathematics

June 2013

THE UNIVERSITY OF WARWICK

Contents

Acknowledgments

First, I thank my supervisor Brian Bowditch for his valuable guidance, patience and advice. I have learnt much from his intuitive approach and his knack for seeing "what's really going on". I must also thank Saul Schleimer for many helpful conversations, and I especially appreciate his ability to turn what appear to be difficult concepts at first into something very simple and enlightening. I thank him and Marc Lackenby for very thorough and careful comments on my thesis.

The past three and a half years would have been several orders of magnitude tougher if not for the friendship and support of my fellows students at the Warwick mathematics department. Special mention must be made to Alex, Linling, Mark, Richard, Tara, and Tom. I am very fortunate to have these people play a key part in my life, and I have felt very much at home after leaving everything familiar to me on the other side of the planet.

I am also grateful to the University of Warwick and the Warwick Mathematics Institute for providing financial support through a Warwick Postgraduate Research Scholarship.

Finally, I would not have grown to become the person I am today without the never-failing support and love from my parents. They have taught me to persevere through difficult times, to keep my feet on the ground, and to realise how lucky and privileged I am to have the opportunity to pursue my interests. Let this thesis stand as a testament to those lessons.

Declarations

I declare that the material in this thesis is, to the best of my knowledge, my own except where otherwise indicated or cited in the text, or else where the material is widely known. This material has not been submitted for any other degree, and only the material from Chapter 5 has been submitted for peer-reviewed publication.

Abstract

This thesis studies the coarse geometry of the curve complex using intersection number techniques. We show how weighted intersection numbers can be studied using appropriate singular Euclidean surfaces. We then introduce a coarse analogue of the convex hull of a finite set of vertices in the curve complex, called the *short curve hull*, and provide intersection number conditions to find nearest point projections to such hulls. We also obtain an upper bound for distances in the curve complex using a greedy algorithm due to Hempel.

Covering maps between surfaces also play a significant part in this thesis. We give a new proof of a theorem of Rafi and Schleimer which states that a covering map between surfaces induces a natural quasi-isometric embedding between their corresponding curve complexes. Our proof employs a distance estimate via a suitable hyperbolic 3–manifold which arises from work on the proof of the Ending Lamination Theorem. We then define an operation using a given covering map and intersection number conditions and show that it approximates a nearest point projection to the image of Rafi–Schleimer's map. We also prove that this operation approximates a circumcentre of the orbit of a vertex in the curve complex under the deck transformation group of a regular cover.

Introduction

The curve complex $\mathcal{C}(S)$ associated to a surface S was introduced by Harvey to understand the thin part of the Teichmüller space of S [Har81]. By definition, it is a simplicial complex whose vertices are isotopy classes of simple closed curves on S and whose simplices are spanned by collections of curves which can be realised disjointly simultaneously. It has proven to be a very useful tool in the geometric group theoretic study of surface mapping class groups – indeed, in all but a small finite number of cases, the mapping class group $\mathrm{Mod}(S)$ is naturally isomorphic to the automorphism group of the corresponding curve complex $\mathcal{C}(S)$ [Iva97], [Luo00]. It also plays a crucial role the proof of the monumental Ending Lamination Theorem [BCM12] where it captures the behaviour of short geodesics as one marches off "to infinity" in a hyperbolic 3–manifold.

Much is already known about the large scale geometry of the curve complex. The curve complex has infinite diameter [Kob88] and, in their celebrated paper, Masur and Minsky [MM99] prove that the curve complex is Gromov hyperbolic. Furthermore, Klarreich shows that the Gromov boundary of $\mathcal{C}(S)$ can be identified with the space of ending laminations for S [Kla].

We seek to understand various geometric notions occurring in the context of general Gromov hyperbolic spaces – including geodesics, quasi-convex sets, nearest point projections and quasi-isometric embeddings – using tools which naturally arise in low-dimensional topology and geometry. As a motivating example, let us consider geodesics in $\mathcal{C}(S)$. There are several methods of approximating geodesics, each with their own distinct advantages: Masur–Minsky employ projections of Teichmüller geodesics [MM99], other authors such as Hamenstädt use train-track splitting sequences [Ham06], while Bowditch uses intersection number conditions [Bow06b].

We now describe (a simplified version of) Bowditch's method in more detail. Let α_1 and α_2 be distinct curves in $\mathcal{C}(S)$. Given a pair $\mathbf{t} = (t_1, t_2)$ of non-negative reals satisfying $t_1 + t_2 = 1$, let $\gamma_{\mathbf{t}}$ be a curve which minimises

$$t_1 i(\alpha_1, \cdot) + t_2 i(\alpha_2, \cdot)$$

among all curves in $\mathcal{C}(S)$, where $i(\cdot,\cdot)$ denotes the *geometric intersection number*. Bowditch shows that by taking the union $\bigcup_{\mathbf{t}} \gamma_{\mathbf{t}}$ over all possible pairs $\mathbf{t}$, we obtain a subset of $\mathcal{C}(S)$ which is a bounded Hausdorff distance away from any geodesic segment joining α_1 and α_2 in $\mathcal{C}(S)$. We can easily generalise Bowditch's statement for an arbitrary set of curves $\alpha_1, \ldots, \alpha_n$ – indeed, in Chapter 4 we show that the analogous union behaves like a "coarse convex hull" for the α_j's considered as a set of vertices in $\mathcal{C}(S)$.

The appeal of this approach is that it can yield simple combinatorial methods of approximating naturally occurring objects in the curve complex. We continue with this theme to describe nearest point projection maps to the above "hulls" and also to subcomplexes which arise from surface covering maps using intersection number conditions.

Overview and main results

Chapter 1 covers background material regarding coarse geometric notions with a special emphasis on δ–hyperbolic spaces. All the results stated are either contained in the literature or can be deduced using elementary arguments.

In Chapter 2, we introduce some key players associated to a surface S, namely its *curve complex* $\mathcal{C}(S)$ and *mapping class group* $\mathrm{Mod}(S)$, and outline several key theorems. We also establish some basic notions regarding weighted intersection numbers and hyperbolic geometry.

In Chapter 3, we generalise Bowditch's construction of singular Euclidean surfaces which are used to estimate weighted intersection numbers [Bow06b]. The main result of this chapter is the proof that a quadratic isoperimetric inequality also holds for our generalised structures. We then apply a result of Bowditch to deduce that such surfaces possess essential annuli of definite width.

In Chapter 4, we introduce two notions of "hulls" for a given finite set of vertices in $\mathcal{C}(S)$: the *hyperbolic hull*, a purely coarse geometric object; and the *short curve hull*, defined using only intersection number conditions. With the aid of results in Chapter 3, we prove that these two objects agree up to bounded Hausdorff distance. We also give intersection number conditions for approximating a nearest point projection to such hulls.

Moving on the Chapter 5, we give a new proof of a theorem due to Rafi and Schleimer [RS09] which states that the natural *lifting map* between curve complexes induced by a finite index covering map of surfaces is a quasi-isometric embedding. Our proof employs a distance estimate via a suitable hyperbolic 3–manifold which arises from work towards the Ending Lamination Theorem [BCM12].

Chapter 6 continues the theme of covering maps and the curve complex. This time, we define a simple operation on curves and then prove that it approximates a nearest point projection to the image of Rafi–Schleimer's lifting map. Our proof employs the results obtained in Chapter 4. We also show that this operation approximates a *circumcentre* for the orbit of a curve in $\mathcal{C}(S)$ under the group of deck transformations for a regular cover.

Finally, in Chapter 7, we describe two methods for obtaining upper bounds for distance in the curve complex in terms of intersection numbers.

Chapter 1

Coarse geometry

In this chapter, we recall some basic definitions and notions concerning Gromov hyperbolic spaces. Many of the statements and results are either well known in the literature or relatively straightforward to deduce; we shall include them for completeness and to establish notation and terminology. We refer the reader to [BH99], [Gro87], [ABC$^+$91] and [Bow06a] for more background.

1.1 Notation

Let $(\mathcal{X}, d)$ be a metric space. Given any subset $A \subseteq \mathcal{X}$ and a point $x \in \mathcal{X}$, we define $d(x, A) := \inf\{d(x, a) \mid a \in A\}$. For $\mathsf{r} \geq 0$, let

$$\mathcal{N}_{\mathsf{r}}(A) = \{x \in \mathcal{X} \mid d(x, A) \leq \mathsf{r}\}$$

denote the r–*neighbourhood* of A in $\mathcal{X}$. For subsets $A, B \subseteq \mathcal{X}$ and $\mathsf{r} \geq 0$, write

$$A \subseteq_{\mathsf{r}} B \iff A \subseteq \mathcal{N}_{\mathsf{r}}(B)$$

and

$$A \approx_{\mathsf{r}} B \iff A \subseteq_{\mathsf{r}} B \text{ and } B \subseteq_{\mathsf{r}} A.$$

Define the *Hausdorff distance* between A and B to be

$$\mathrm{HausDist}(A, B) = \inf\{\mathsf{r} \geq 0 \mid A \approx_{\mathsf{r}} B\}.$$

To simplify notation, we will often write $a \in \mathcal{X}$ in place of a singleton set $\{a\} \subseteq \mathcal{X}$. We will always use the standard Euclidean metric on the reals unless otherwise specified. If a and b are real numbers then

$$a \approx_{\mathsf{r}} b \iff |a-b| \leq \mathsf{r}.$$

We will also adopt the following notation:

$$a \asymp_{\mathsf{r}} b \iff a \leq \mathsf{r}\, b + \mathsf{r} \quad \text{and} \quad b \leq \mathsf{r}\, a + \mathsf{r}.$$

1.2 Geodesics, quasiconvexity and quasi-isometries

Let $I \subseteq \mathbb{R}$ be an interval. A *geodesic* is a map $\gamma : I \to \mathcal{X}$ so that $d(\gamma(t), \gamma(s)) = |t-s|$ for all $t, s \in I$. A *geodesic segment* connecting points x and y in $\mathcal{X}$ is the image of a geodesic $\gamma : [0, d(x,y)] \to \mathcal{X}$ such that $\gamma(0) = x$ and $\gamma(d(x,y)) = y$. A metric space $\mathcal{X}$ is called a *geodesic space* if every pair of points can be connected by a geodesic segment.

Definition 1.1 (Quasiconvexity)**.** A subset $U \subseteq \mathcal{X}$ is Q–*quasiconvex* if any geodesic segment connecting any pair of points in U lies in $\mathcal{N}_{\mathsf{Q}}(U)$. We say a subset is *quasiconvex* if it is Q–quasiconvex for some $\mathsf{Q} \geq 0$.

Definition 1.2 (Quasi-isometric embedding, quasi-isometry)**.** A (one-to-many) map $f : \mathcal{X} \to \mathcal{Y}$ between metrics spaces is a Λ–*quasi-isometric embedding* if for all $x_1, x_2 \in \mathcal{X}$ and $y_1 \in f(x_1)$, $y_2 \in f(x_2)$ we have

$$d_{\mathcal{Y}}(y_1, y_2) \asymp_{\Lambda} d_{\mathcal{X}}(x_1, x_2).$$

In addition, if $\mathcal{N}_{\Lambda}(f(\mathcal{X})) = \mathcal{Y}$ then f is called a Λ–*quasi-isometry* and we say that $\mathcal{X}$ and $\mathcal{Y}$ are Λ–*quasi-isometric*. If $\mathcal{X}$ and $\mathcal{Y}$ are Λ–quasi-isometric for some $\Lambda \geq 1$ then we may simply say that they are *quasi-isometric*.

1.3 Gromov hyperbolic spaces

1.3.1 Thin triangles

We recall some basic results about Gromov hyperbolic spaces. Let $\mathcal{X}$ be a geodesic space.

Definition 1.3 (Geodesic triangle)**.** A geodesic triangle T in $\mathcal{X}$ consists of three points $x, y, z \in \mathcal{X}$ together with three geodesic segments $[x,y], [y,z], [z,x]$. The segments will be called the *sides* of the triangle T.

We will abbreviate $d(x, y)$ to xy.

Definition 1.4 (Gromov product)**.** Let x, y, z be points in $\mathcal{X}$. Define

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

This quantity is called the *Gromov product* of x and y with respect to z.

Given a geodesic triangle T on $x, y, z \in \mathcal{X}$, we construct a *comparison tripod* $\bar{T}$ as follows: Build a metric tree consisting of one central vertex whose valence is at most 3 and three vertices of valence one with three edges of lengths $\langle y, z\rangle_x$, $\langle z, x\rangle_y$ and $\langle x, y\rangle_z$. Label the central vertex o_T and the other endpoints of the edges $\bar{x}, \bar{y}$ and $\bar{z}$ respectively. We allow for the possibility of edges having length zero in this construction.

There exists a unique map

$$\theta_T \ : \ T \to \bar{T}$$

satisfying $\theta_T(x) = \bar{x}$, $\theta_T(y) = \bar{y}$ and $\theta_T(z) = \bar{z}$ which restricts to an isometric embedding on each edge of T. The elements of $\theta_T^{-1}(o_T)$ are called the *internal points* of T.

Recall that the *diameter* of a non-empty subset $U \subseteq \mathcal{X}$ is

$$\mathrm{diam}(U) := \sup\{d(x, y) \mid x, y \in U\}.$$

Definition 1.5 (Thin triangle, δ–hyperbolic space)**.** A geodesic triangle T is *δ–thin* if

$$\mathrm{diam}(\theta_T^{-1}(p)) \leq \delta$$

for all $p \in \bar{T}$. A geodesic space $\mathcal{X}$ is *δ–hyperbolic* if all of its geodesic triangles are δ–thin. We call $\mathcal{X}$ *(Gromov) hyperbolic* if it is δ–hyperbolic for some $\delta \geq 0$.

If T is a δ–thin geodesic triangle with vertices $x, y, z \in \mathcal{X}$ then its internal points decompose it into three pairs of δ–fellow travelling geodesic segments whose lengths are $\langle y, z\rangle_x$, $\langle z, x\rangle_y$ and $\langle x, y\rangle_z$.

The following result shows us that geodesic segments between two given points in a δ–hyperbolic space are essentially unique up to bounded error.

Lemma 1.6 (Stability of geodesics)**.** *Let x, y be points in a δ–hyperbolic space $\mathcal{X}$. Then any two geodesic segments γ_1, γ_2 joining x and y δ-fellow travel: if $u_1 \in \gamma_1$ and $u_2 \in \gamma_2$ are points such that $xu_1 = xu_2$ then $u_1u_2 \leq \delta$.*

Proof. Consider the geodesic triangle with vertices x, y and y whose non-degenerate sides are γ_1 and γ_2. The result follows from the definition of δ–thinness. $\square$

As an immediate corollary, we see that geodesics in a δ–hyperbolic space are δ–quasiconvex.

1.3.2 Equivalent notions of Gromov hyperbolicity

Assume $\mathcal{X}$ is a geodesic space.

Lemma 1.7 (Four point condition, [BH99] Proposition 1.22)**.** *If $\mathcal{X}$ is a δ–hyperbolic space then*

$$xy + zw \leq \max\{xz + yw,\, xw + yz\} + 2\delta$$

for all $x, y, z, w \in \mathcal{X}$.

Conversely, if the above inequality holds for all points x, y, z and w in a geodesic space $\mathcal{X}$, then $\mathcal{X}$ is δ'–hyperbolic for some $\delta' \geq 0$ depending only on δ. $\square$

Suppose $\mathsf{k} \geq 0$. A k–*centre* for a geodesic triangle $T \subseteq \mathcal{X}$ is a point in $\mathcal{X}$ which lies within a distance k of each side of T.

Lemma 1.8 ([Bow06a] Proposition 6.13)**.** *Any geodesic triangle in a δ–hyperbolic space possesses a δ–centre, namely, any of its internal points.*

Conversely, suppose $\mathcal{X}$ is a geodesic space and let $\mathsf{k} \geq 0$. If all geodesic triangles in $\mathcal{X}$ possess k–centres then $\mathcal{X}$ is δ–hyperbolic for some $\delta \geq 0$ depending only on k. $\square$

1.3.3 Nearest point projections to quasiconvex sets

Given a non-empty subset $U \subseteq \mathcal{X}$ and a point $x \in \mathcal{X}$, define

$$\mathrm{proj}_U(x) := \{p \in U \mid xp = d(x, U)\}$$

to be the set of *nearest point projections* of x to U in $\mathcal{X}$. If U is closed in $\mathcal{X}$ then $\mathrm{proj}_U(x)$ is always non-empty.

Nearest point projections to geodesic segments can be approximated by internal points:

Lemma 1.9. *Let $\mathcal{X}$ be a δ–hyperbolic space. Let T a geodesic triangle with vertices $x, y, z \in \mathcal{X}$. Let o_x be the internal point of T on $[y, z]$ and suppose $p \in [y, z]$ is a point such that $xp \leq d(x, [y, z]) + \epsilon$, for some $\epsilon \geq 0$. Then $o_x p \leq 2\delta + \epsilon$.*

Proof. Without loss of generality, suppose p lies on $[o_x, y]$. Let q be a point on $[x, y]$ such that $py = qy$. By δ–hyperbolicity, we have $pq \leq \delta$. Let o_z be the internal point on $[x, z]$ opposite y. Then

$$xo_z + o_z q = xq \leq xp + \delta \leq d(x, [y, z]) + \epsilon + \delta \leq xo_x + \epsilon + \delta \leq xo_z + 2\,\delta + \epsilon$$

and so $o_x p = o_z q \leq 2\delta + \epsilon$ □

Let us now assume that U is a closed, non-empty Q–quasiconvex subset of a δ–hyperbolic space $\mathcal{X}$.

Lemma 1.10. *Let p be a nearest point projection of $x \in \mathcal{X}$ to U. Let u be any point in U and let o_x, o_p and o_u be the respective internal points of a geodesic triangle with vertices x, p and u. Then $po_x \leq \delta + \mathsf{Q}$ and hence $po_u \leq \delta + \mathsf{Q}$.*

Proof. By quasiconvexity of U we have $d(o_x, U) \leq \mathsf{Q}$. Thus,

$$xo_u + o_u p = xp = d(x, U) \leq xo_x + d(o_x, U) \leq xo_u + o_u o_x + \mathsf{Q} \leq xo_u + \delta + \mathsf{Q}$$

and so $o_x p = o_u p \leq \delta + \mathsf{Q}$. □

Lemma 1.11. *For all $x \in \mathcal{X}$,*

$$\operatorname{diam}(\operatorname{proj}_U(x)) \leq 2\delta + 2\mathsf{Q}.$$

Proof. Let p and q be nearest point projections of x to U. Let o_x be the respective internal point opposite x of a geodesic triangle with vertices x, p and q. Applying 1.10, we deduce $pq \leq po_x + o_x q \leq 2\delta + 2\mathsf{Q}$. □

A consequence of Lemma 1.10 is that any geodesic from x to a point in U must pass within a distance of $\delta + \mathsf{Q}$ of every nearest point projection of x to U. It turns out that this property characterises nearest point projections to quasiconvex sets in hyperbolic spaces. For $\mathsf{r} \geq 0$, we define $\operatorname{entry}_U(x, \mathsf{r})$ to be the set of all points $q \in U$ such that for all $u \in U$, every geodesic connecting x to u passes within a distance of r of q. Such points will be called r–*entry points* of x to U.

Lemma 1.12. *Let $\mathsf{r} \geq 0$. Then for all $x \in \mathcal{X}$,*

$$\operatorname{entry}_U(x, \mathsf{r}) \subseteq_{2\mathsf{r}} \operatorname{proj}_U(x).$$

In particular, for $\mathsf{r} \geq 2\delta + \mathsf{Q}$ *we have*

$$\mathrm{entry}_U(x, \mathsf{r}) \approx_{2\mathsf{r}} \mathrm{proj}_U(x).$$

Proof. Suppose p is a nearest point projection and q is an r–entry point of x to U respectively. Then there is some point $y \in [x, p]$ so that $yq \leq \mathsf{r}$. Now

$$xy + yp = xp \leq xq \leq xy + yq \leq xy + \mathsf{r}$$

and so $pq \leq py + yq \leq 2\mathsf{r}$ which proves the first statement. The second statement follows from Lemma 1.10. $\square$

Furthermore, any geodesic segment from a point $x \in \mathcal{X}$ to $u \in U$ can be approximated by the concatenation of two segments: the first from x to any point $p \in \mathrm{proj}_U(x)$ and the second from p to u.

Lemma 1.13. *Given* $x \in \mathcal{X}$*, let* $p \in \mathrm{proj}_U(x)$*. Then for any* $u \in U$*,*

$$[x, u] \approx_{2\delta+\mathsf{Q}} [x, p] \cup [p, u]$$

and

$$xu \approx_{2\delta+2\mathsf{Q}} xp + pu.$$

Proof. By hyperbolicity, we have $[x, u] \approx_\delta [x, o_u] \cup [o_x, u]$ and $xu = xo_u + o_x u$. By Lemma 1.10,

$$\mathrm{diam}[o_u, p] = \mathrm{diam}[p, o_x] = po_x \leq \delta + \mathsf{Q}$$

and so $[o_u, p] \cup [p, o_x] \subseteq_{\delta+\mathsf{Q}} \{o_u, o_x\} \subseteq_\delta [x, u]$. $\square$

1.3.4 Circumcentres

Let $\mathcal{X}$ be a δ–hyperbolic space and suppose $U \subseteq \mathcal{X}$ is a non-empty finite subset.

Definition 1.14 (Radius, Circumcentre)**.** The *radius* of U is

$$\mathrm{rad}(U) := \min\left\{\mathsf{r} \geq 0 \mid \exists x \in \mathcal{X}, U \subseteq B_\mathsf{r}(x)\right\},$$

where $B_\mathsf{r}(x)$ is the closed ball of radius r centred at x. We call a point $x \in \mathcal{X}$ a *circumcentre* of U if $U \subseteq B_\mathsf{r}(x)$ for $\mathsf{r} = \mathrm{rad}(U)$ and write $\mathrm{circ}(U)$ for the set of circumcentres of U.

Lemma 1.15. *Let* x*,* y *and* z *be points in* $\mathcal{X}$*. Suppose* m *is a midpoint of some geodesic* $[x, y]$ *connecting* x *and* y*. Then* $\max\{xz, yz\} \approx_\delta \frac{1}{2}\, xy + mz$*.*

Proof. Without loss of generality, suppose $xz \geq yz$. Then m lies on $[x, o_z]$, where $o_z \in [x, y]$ is the internal point opposite z. By hyperbolicity, there is a point $q \in [x, z]$ such that $xq = xm$ and $qm \leq \delta$. Finally,

$$\max\{xz, yz\} = xz = xq + qz \approx_\delta xm + mz = \frac{1}{2}xy + mz$$

which completes the proof. □

Lemma 1.16. *Let c be a circumcentre of U and suppose $x \in \mathcal{X}$ is a point such that $U \subseteq B_{\mathsf{r}+\epsilon}(x)$, where $\mathsf{r} = \operatorname{rad}(U)$ and $\epsilon \geq 0$. Then $cx \leq 2\delta + 2\epsilon$ and hence* $\operatorname{diam}(\operatorname{circ}(U)) \leq 2\delta$*.*

Proof. Let m be a midpoint of c and x. Choose $u \in U$ so that um is maximal. Applying Lemma 1.15 and the definition of radius gives

$$\operatorname{rad}(U) \leq um \leq \max\{cu, xu\} - \frac{1}{2}cx + \delta \leq \operatorname{rad}(U) + \epsilon - \frac{1}{2}cx + \delta$$

and we are done. □

Lemma 1.17. *Suppose c is a circumcentre of U. Let x, $y \in U$ be points such that $xy \geq \operatorname{diam}(U) - 2\epsilon$, for some $\epsilon \geq 0$. Let m be the midpoint of a geodesic segment $[x, y]$. Then $c \approx_{2\delta+\epsilon} m$. Furthermore, we have*

$$\operatorname{diam}(U) \leq 2\operatorname{rad}(U) \leq \operatorname{diam}(U) + 2\delta.$$

Proof. Suppose x' and y' are points in U satisfying $x'y' = \operatorname{diam}(U)$. By Lemma 1.15, we deduce

$$um \leq \max\{x'u, y'u\} - \frac{1}{2}x'y' + \delta \leq \operatorname{diam}(U) - \frac{1}{2}\operatorname{diam}(U) + \delta$$

for all $u \in U$. Choosing u so that um is maximal yields

$$\operatorname{rad}(U) \leq um \leq \frac{1}{2}\operatorname{diam}(U) + \delta.$$

Next, observing

$$\operatorname{diam}(U) = x'y' \leq x'c + cy' \leq 2\operatorname{rad}(U)$$

completes the proof of the second claim. Finally,

$$cm \leq \max\{cx, cy\} - \frac{1}{2}xy + \delta \leq \operatorname{rad}(U) - \frac{1}{2}\operatorname{diam}(U) + \epsilon + \delta \leq 2\delta + \epsilon$$

where we have applied the second claim and Lemma 1.15 once more. □

1.3.5 Isometries and stable lengths

Suppose Γ is a group acting by isometries on a metric space $\mathcal{X}$. Define the *stable length* of an element $g \in \Gamma$ to be

$$\|g\| := \lim_{n\to\infty} \frac{1}{n}\, d(x, g^n x),$$

where x is any point in $\mathcal{X}$. It is straightforward to check that the $\|g\|$ does not depend on the choice of point $x \in \mathcal{X}$. Indeed, the stable length depends only on the conjugacy class of an element.

Following Gromov, we introduce the following terminology:

Definition 1.18. Let Γ act on a metric space $\mathcal{X}$ by isometries. Then an isometry $g \in \Gamma$ is:

- *elliptic* if any (hence all) of its orbits have finite diameter;
- *parabolic* if its orbits have infinite diameter but $\|g\| = 0$; or
- *loxodromic* if $\|g\| > 0$.

In the situation of groups acting isometrically on δ–hyperbolic spaces, the above list is in fact exhaustive.

Theorem 1.19 ([Gro87] Corollary 8.1B)**.** *Any isometry of a δ–hyperbolic space is either elliptic, parabolic or loxodromic.* □

Chapter 2

Surfaces and their relatives

In this chapter, we first provide some background information concerning curves on surfaces and intersection numbers. We then introduce the main player of this thesis, the *curve complex*, and state some key results concerning its large scale geometry. We also discuss the mapping class group and its action on the curve complex. Finally, we round off the chapter by recalling some basic hyperbolic geometry. All of the results stated in this chapter can be found in the literature.

2.1 Curves and surfaces

Throughout this thesis, all surfaces will be real two-dimensional orientable manifolds. We will use $S = (S, \Omega)$ to denote a closed, orientable, connected surface of genus $g \geq 0$ together with a set Ω of $m \geq 0$ marked points. Define the *complexity* of S to be $\xi(S) := 3g + m - 3$. We choose to work primarily with marked points instead of punctures or boundary components as this will be more convenient for many of our constructions. In some situations, though, it may become necessary to deal with the latter cases – we shall inform the reader when the need arises.

A *curve* on S is a continuous map $a : S^1 \to S - \Omega$, where $S^1 = \mathbb{R}/\mathbb{Z}$ is the circle. We will also write a for its image on S. A curve a is *simple* if it is an embedded copy of S^1.

Two curves a and b are *freely homotopic* if there exists a continuous map $\mathbf{F} : S^1 \times [0, 1] \to S - \Omega$ such that $\mathbf{F}(\theta, 0) = a(\theta)$ and $\mathbf{F}(\theta, 1) = b(\theta)$. We will also consider the curves a and $a'(\theta) = a(-\theta)$ to be freely homotopic; in other words, we will ignore orientations when speaking of free homotopy classes of curves.

We call a curve *trivial* or *peripheral* if it is freely homotopic to a curve bounding a disc or a disc with exactly one marked point respectively. A simple closed curve which is non-trivial and non-peripheral is called *essential*.

Let $\mathcal{C}^0(S)$ denote the set of free homotopy classes of essential simple closed curves on S. Unless explicitly stated otherwise, we will blur the distinction between curves and their free homotopy classes.

A *multicurve* on S is a finite collection of essential simple closed curves which can be realised disjointly simultaneously.

2.1.1 Intersection numbers

Naïvely speaking, the *intersection number* of two curves on S is the number of times they cross. We need to proceed with some caution to make this notion precise – especially in the case of non-simple curves.

Suppose a and b are curves on S. A pair $(\theta_1, \theta_2) \in S^1 \times S^1$ satisfying $a(\theta_1) = b(\theta_2)$ is called a *pre-intersection point* of a and b. Write $a \,\hat{\cap}\, b \subseteq S^1 \times S^1$ for the set of pre-intersection points of a and b.

Definition 2.1. Let α and β be free homotopy classes of closed curves on S. The *(geometric) intersection number* of α and β is

$$i(\alpha, \beta) := \min\{|a \,\hat{\cap}\, b| \mid a \in \alpha,\ b \in \beta\}.$$

We say that $a \in \alpha$ and $b \in \beta$ are in *minimal position* if $i(\alpha, \beta) = |a \,\hat{\cap}\, b|$.

If $i(\alpha, \beta) \neq 0$ then we say that α and β *intersect*; otherwise, we say that they are *disjoint*.

When a and b are both simple then $a \,\hat{\cap}\, b$ is in natural bijective correspondence with $a \cap b \subseteq S$ and so our definition agrees with the usual notion of intersection number. For non-simple curves, it is possible for distinct pre-intersection points of a and b to be sent to the same point in $a \cap b$, creating a *multiple point*. This can occur, for example, when b crosses a self-intersection point of a. Fortunately, one can always perturb a and b so that they are in *general position*, that is, they intersect transversely and with no multiple points. Thus, the intersection number of α and β is equal to the minimal number of intersection points over all their representatives which are in general position.

Suppose a and b are curves which intersect transversely. A *bigon* on S formed by a and b is a closed embedded disc whose boundary consists of exactly one subarc of a and one subarc of b which intersect at their endpoints. The following criterion allows us to detect whether two curves are in minimal position. We refer the reader to [FM12] (Proposition 1.7 and Corollary 1.9) for a proof.

Lemma 2.2 (Bigon Criterion)**.** *Suppose a and b are curves on S which intersect transversely. If a and b are in minimal position then they do not form any bigons. Furthermore, if a and b are simple then the converse also holds.* $\square$

We also remark that the definition of intersection number can be extended to multicurves.

Weighted intersection numbers

Let $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ be an n–tuple of distinct multicurves in $\mathcal{C}^0(S)$. A vector $\mathbf{t} = (t_1, \ldots, t_n) \neq \mathbf{0}$ of non-negative real numbers shall be referred to as a *weight vector*. Write $\mathbf{t} \cdot \boldsymbol{\alpha}$ for the formal sum $\sum_i t_i \alpha_i$. We will extend intersection number linearly over such sums:

$$i(\mathbf{t} \cdot \boldsymbol{\alpha}, \gamma) := \sum_i t_i i(\alpha_i, \gamma).$$

For notational convenience, define a function on weight vectors by setting

$$\|\mathbf{t}\|_{\boldsymbol{\alpha}} := \sqrt{i(\mathbf{t} \cdot \boldsymbol{\alpha})},$$

where

$$i(\mathbf{t} \cdot \boldsymbol{\alpha}) = \sum_{j<k} t_j t_k i(\alpha_j, \alpha_k)$$

is the *self-intersection number* of $\mathbf{t} \cdot \boldsymbol{\alpha}$. This serves as a rescaling factor for the singular Euclidean surface $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ appearing in Chapter 3.

2.1.2 Filling curves

A finite collection of curves $\{\alpha_1, \ldots, \alpha_n\}$ on S *fills* if every essential curve on S intersects at least one of the α_i. Equivalently, we can rephrase this as follows:

Lemma 2.3. *Assume S is a surface where $\xi(S) \geq 1$. Suppose a collection of curves $\{\alpha_1, \ldots, \alpha_n\}$ on S is realised so that they intersect minimally and pairwise transversely. Then the collection fills S if and only if every component of $S - (\alpha_1 \cup \ldots \cup \alpha_n)$ is an open disc with at most one marked point.* $\square$

The reverse direction is straightforward to verify. To prove the forward direction, one can use the complexity assumption to check that any complementary component which is not a disc with at most one marked point must contain an essential closed curve on S.

2.2 The curve complex

Harvey introduced the curve complex in [Har81] to study the geometry of Teichmüller space. This will be the primary object of study in this thesis. For what follows, we shall assume that S has complexity $\xi(S)$ at least 2; modifications to the definition are required for low-complexity cases but we shall not deal with them here. For an introduction to the curve complex, see [Sch].

Definition 2.4. The *curve complex* of S, denoted $\mathcal{C}(S)$, is a simplicial complex whose vertex set is $\mathcal{C}^0(S)$ and whose simplices are spanned by multicurves. In particular, two simple closed curves are connected by an edge in $\mathcal{C}(S)$ if and only if they have disjoint representatives on S.

The dimension of $\mathcal{C}(S)$ is equal to $\xi(S) - 1$. If one replaces marked points with boundary components then the top dimensional simplices of $\mathcal{C}(S)$ correspond to *pants decompositions* of S, that is, multicurves which cut S into a collection of pants.

We endow $\mathcal{C}(S)$ with the standard simplicial metric: each k–simplex is isometrically identified with a standard Euclidean k–simplex whose edge lengths are equal to 1. For our purposes, it suffices to study the 1–skeleton $\mathcal{C}^1(S)$ of the curve complex, often referred to as the *curve graph*. Indeed $\mathcal{C}^1(S)$ equipped with the induced path metric, denoted d_S, is naturally *quasi-isometric* to $\mathcal{C}(S)$. To simplify notation, we shall write $\mathcal{C}(S)$ in place of $\mathcal{C}^1(S)$ and $\alpha \in \mathcal{C}(S)$ to denote a curve (or multicurve).

Let α and β be curves in $\mathcal{C}(S)$. Their distance $d_S(\alpha, \beta)$ is equal to the length of a shortest edge-path in $\mathcal{C}(S)$ connecting α and β. Observe that α and β are disjoint if and only if $d_S(\alpha, \beta) \leq 1$. We also have $d_S(\alpha, \beta) = 2$ if and only if α and β intersect but do not fill S; and $d_S(\alpha, \beta) \geq 3$ if and only if they do fill. The following lemma allows us to bound distances in the curve graph using intersection numbers.

Lemma 2.5 ([Hem01], [Sch])**.** *Suppose α and β are curves in $\mathcal{C}(S)$. Then*

$$d_S(\alpha, \beta) \leq 2 \log_2 i(\alpha, \beta) + 2$$

whenever $i(\alpha, \beta) \neq 0$. □

As an immediate corollary, we see that $\mathcal{C}(S)$ is connected (this was originally observed by Harvey in [Har81]). In Chapter 7, we will give some improvements on the multiplicative constant for the logarithmic term; however, it is not possible to give a bound strictly better than a logarithmic one. It is also worth mentioning that one cannot give a lower bound on distance in $\mathcal{C}(S)$ purely in terms of intersection number – indeed, one can find pairs of non-filling curves which intersect an arbitrarily large number of times.

The curve graph is also locally infinite and has infinite diameter [Kob88]. Masur and Minsky proved the following celebrated theorem regarding the large scale geometry of the curve graph using methods from Teichmüller theory:

Theorem 2.6 ([MM99])**.** *Given any surface S with $\xi(S) \geq 2$, the curve graph $\mathcal{C}(S)$ is δ–hyperbolic for some $\delta > 0$.* □

In [Bow06b], Bowditch gives a combinatorial proof of hyperbolicity using intersection numbers. We will be extending many of the results established in his paper in Chapters 3 and 4.

Theorem 2.7 ([Bowb], [Aou], [CRS], [HPW])**.** *The constant $\delta > 0$ in Theorem 2.6 can be chosen independently of S.* □

Hensel, Przytycki and Webb in particular show that all geodesic triangles in $\mathcal{C}(S)$ possess 17–centres.

2.3 Mapping class groups

Assume $S = (S, \Omega)$ is a surface together with a finite (possibly empty) set of marked points Ω. Let $\mathrm{Homeo}^+(S)$ be the group of orientation-preserving self-homeomorphisms of S which preserve Ω setwise. Write $\mathrm{Homeo}_0(S) \triangleleft \mathrm{Homeo}^+(S)$ for the subgroup of homeomorphisms isotopic to the identity through isotopies fixing the marked points pointwise. The *mapping class group* of S is defined to be

$$\mathrm{Mod}(S) := \mathrm{Homeo}^+(S) \,/\, \mathrm{Homeo}_0(S),$$

the group of orientation-preserving self-homeomorphisms of S up to isotopy. A thorough introduction to mapping class groups can be found in [FM12].

Let us now assume $\xi(S) \geq 2$. Since self-homeomorphisms of S preserve disjointness of curves, it follows that $\mathrm{Mod}(S)$ acts on $\mathcal{C}(S)$ by isometries. In fact, bar a small number of exceptions, $\mathrm{Mod}(S)$ is naturally isomorphic to the full automorphism group of $\mathcal{C}(S)$ [Iva97], [Luo00]. Moreover, one can bound the number of $\mathrm{Mod}(S)$–orbits of vertices in $\mathcal{C}(S)$ in terms of $\xi(S)$ and thus the action is coarsely transitive.

2.4 Hyperbolic geometry

We now outline some basic notions regarding hyperbolic space. We will only be concerned with dimensions two and three, however to streamline the exposition we will proceed in full generality. Refer to [Bea95], [And05] and [Apa00] for more background.

The *upper half-space model* of n–dimensional hyperbolic space consists of the points

$$\mathbb{H}^n = \{(x_1, \ldots, x_n) \in \mathbb{R}^n \mid x_n > 0\}$$

equipped with the length element

$$ds := \frac{ds_{\mathbb{E}}}{|x_n|},$$

where $ds_{\mathbb{E}}$ is the standard Euclidean length element on $\mathbb{R}^n$. The *ideal boundary* of $\mathbb{H}^n$ is

$$\partial_\infty \mathbb{H}^n = (\mathbb{R}^{n-1} \times \{0\}) \cup \{\infty\},$$

which can be regarded as the one-point compactification of $\mathbb{R}^{n-1}$.

Geodesics in the upper half-space model for $\mathbb{H}^n$ are subarcs of circles or straight lines orthogonal to the hyperplane $\mathbb{R}^{n-1} \times \{0\} \subset \partial_\infty \mathbb{H}^n$. In particular, geodesic segments between pairs of distinct points are unique. We also state the following useful distance formula; see [Bea95] (Theorem 7.2.1) for a reference.

Theorem 2.8. *Let* $\mathbf{x} = (x_1, \ldots, x_n)$ *and* $\mathbf{y} = (y_1, \ldots, y_n)$ *be points in* $\mathbb{H}^n$. *Then*

$$\sinh\left(\frac{1}{2}\, d_{\mathbb{H}^n}(\mathbf{x}, \mathbf{y})\right) = \frac{\|\mathbf{x} - \mathbf{y}\|}{2\sqrt{x_n y_n}},$$

where $\|\cdot\|$ *is the standard Euclidean norm on* $\mathbb{R}^n$. □

Theorem 2.9. *Let* $g \in \mathrm{Isom}^+(\mathbb{H}^n)$ *be an orientation-preserving isometry of* $\mathbb{H}^n$. *Then exactly one of the following holds:*

- *g fixes at least one point in $\mathbb{H}^n$ and is therefore elliptic;*
- *g fixes exactly one point on $\partial_\infty \mathbb{H}^n$ and is therefore parabolic; or*
- *g fixes a pair of distinct points on $\partial_\infty \mathbb{H}^n$ and is therefore loxodromic.*

Furthermore, $\mathrm{Isom}^+(\mathbb{H}^n)$ *acts transitively on both* $\mathbb{H}^n$ *and* $\partial_\infty \mathbb{H}^n$. □

A group of isometries $\Gamma \leq \mathrm{Isom}^+(\mathbb{H}^n)$ is *parabolic* if all non-trivial elements are parabolic. One can show that there is a common fixed point for all non-trivial elements in a parabolic subgroup, and thus it makes sense to speak of the fixed point of such a group. It is often convenient to arrange, by conjugation, so that the fixed point is at ∞. A parabolic group Γ whose fixed point is at ∞ acts by Euclidean translations on the $\mathbb{R}^{n-1}$–factor of $\mathbb{H}^n = \mathbb{R}^{n-1} \times \mathbb{R}_+$ and trivially on the $\mathbb{R}_+$–factor.

A *horoball centred at* $\infty \in \partial_\infty \mathbb{H}^n$ is a set of the form

$$H = \{(x_1, \ldots, x_n) \in \mathbb{H}^n \mid x_n > \mathsf{h}\}$$

where h is some positive real. Its boundary $\partial H = \mathbb{R}^{n-1} \times \{\mathsf{h}\}$, called a *horosphere*, is a horizontal hyperplane at height h. We may speak of horoballs and horospheres centred at a point $z \in \partial_\infty \mathbb{H}^n$ by simply translating H by some isometry sending ∞ to z. Horoballs and horospheres centred at $z \in \partial_\infty \mathbb{H}^n$ are invariant under the action of any parabolic group with z as the fixed point.

A *hyperbolic n–manifold* is the quotient of $\mathbb{H}^n$ by a subgroup of $\mathrm{Isom}^+(\mathbb{H}^n)$ acting freely and discretely. In particular, any surface of negative Euler characteristic admits a hyperbolic metric on its interior.

Chapter 3

Singular Euclidean structures

In this chapter we give a generalisation of Bowditch's singular Euclidean surfaces which are used to to estimate weighted intersection numbers in [Bow06b]. We then show that such surfaces satisfy a quadratic isoperimetric inequality and apply a theorem of Bowditch to deduce that these surfaces contain essential annuli of definite width.

3.1 Construction of $S(\mathbf{t} \cdot \boldsymbol{\alpha})$

Suppose $S = (S, \Omega)$ is a closed surface of genus g with a set of m marked points Ω so that $\xi(S) \geq 2$. Throughout this chapter, we shall fix an n–tuple $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ of distinct multicurves in $\mathcal{C}(S)$ and a weight vector $\mathbf{t} = (t_1, \ldots, t_n)$. For simplicity, assume that $\boldsymbol{\alpha}$ fills S and that all entries of $\mathbf{t}$ are positive. We will deal with the appropriate modifications for the non-filling case in Section 3.5.

Begin by realising the multicurves α_i on S so that they intersect generally and pairwise minimally. We can achieve this, for example, by placing a complete hyperbolic metric on the complement of the marked points in S, taking geodesic representatives of the α_i and perturbing slightly if required. The union of the α_i gives a connected 4–valent graph Υ on S. The closure of each component of $S - \Upsilon$ is a polygon with at most one marked point. The polygons together with Υ give S the structure of a 2–dimensional cell complex. By taking the dual 2–cell structure, we obtain a tiling of S by rectangles which are in bijection with the self-intersection points of $\boldsymbol{\alpha}$. We will insist that any marked point of S coincides with a vertex of this tiling.

Each rectangle R corresponding to an intersection of α_i with α_j is isometrically identified with a Euclidean rectangle of side lengths t_i and t_j so that α_i is transverse to the two sides of length t_i. This gives a singular Euclidean metric on S. We may arrange for each α_i to be locally geodesic in this metric by requiring $\alpha_i \cap R$ to be a straight line connecting the midpoints of opposite sides of R, for every rectangle R meeting α_i. Thus, each component of α_i is the core curve of an annulus of width t_i formed by taking the union of all rectangles R it meets.

The singular Euclidean surface defined above shall be denoted $S(\mathbf{t} \cdot \boldsymbol{\alpha})$. We remark that the metric depends on the realisation of $\boldsymbol{\alpha}$ on S up to isotopy, however, any such choice will work equally well for the purposes of proving the proposition below.

We will allow ourselves to homotope a curve $\gamma \in \mathcal{C}(S)$ to meet marked points in order to speak of geodesic representatives on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$. To be more precise, suppose c' is a simple closed curve on S representing γ. We say c is a *representative* of γ if there is a homotopy $\mathbf{F} : S^1 \times [0,1] \to S$ such that $\mathbf{F}(\theta, 0) = c'(\theta)$, $\mathbf{F}(\theta, 1) = c(\theta)$ and $\mathbf{F}(S^1 \times \{t\}) \subseteq S - \Omega$ for all $0 \leq t < 1$.

A geodesic representative c of γ on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ may not necessarily be an embedded copy of S^1. In these cases, there is a decomposition of the circle $S^1 = \cup I_k$ into a finite union of closed intervals with disjoint interiors so that $c : S^1 \to S(\mathbf{t} \cdot \boldsymbol{\alpha})$ embeds each I_k as a straight line segment whose endpoints are singular points or marked points.

We will use $l(\gamma)$ to denote the length of a geodesic representative of γ on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ with respect to the singular Euclidean metric.

Proposition 3.1. *The singular Euclidean surface $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ has the following properties.*

1. *$S(\mathbf{t} \cdot \boldsymbol{\alpha})$ has area $\|\mathbf{t}\|_{\boldsymbol{\alpha}}^2 = \sum_{j<k} t_j t_k i(\alpha_j, \alpha_k)$.*
2. *For all curves $\gamma \in \mathcal{C}(S)$, we have*

$$l(\gamma) \leq i(\mathbf{t} \cdot \boldsymbol{\alpha}, \gamma) \leq \sqrt{2} l(\gamma).$$

3. *There exists an essential annulus on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ whose width is at least $\mathsf{W}_0 \|\mathbf{t}\|_{\boldsymbol{\alpha}}$, where $\mathsf{W}_0 > 0$ is a constant depending only on $\xi(S)$.*

The first claim is immediate from the construction. Before proving the second and third claims, we will outline some consequences of this proposition which shall later be used to prove Lemma 4.5. It is worth mentioning that the third claim holds for a larger class of metrics satisfying a suitable isoperimetric inequality. We also remark that the metric on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ can be approximated by a non-singular

Riemannian metric. We will, however, choose to work with singular Euclidean metrics to simplify the exposition.

3.2 Short curves and wide annuli

We now state some facts arising from the interplay between weighted intersection numbers, lengths of curves and widths of annuli on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$. Most of the statements and proofs are covered in [Bow06b] but we will include them for completeness.

Let A be a closed Riemannian annulus. Define $\mathrm{width}(A)$ to be the length of a shortest arc joining its two boundary components and $\mathrm{length}(A)$ to be the length of a shortest core curve of A. The following is a consequence of the Besicovitch Lemma [Bes52] (see Lemma $4.5\frac{1}{2}$ in [Gro99] for a proof).

Lemma 3.2 (Annulus inequality)**.** *Let A be an annulus. Then*

$$\mathrm{width}(A) \times \mathrm{length}(A) \leq \mathrm{area}(A) \qquad \square$$

If γ is the core curve of an annulus A on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ then

$$\mathrm{width}(A) \times l(\gamma) \leq \mathrm{width}(A) \times \mathrm{length}(A) \leq \mathrm{area}(A) \leq \mathrm{area}(S(\mathbf{t} \cdot \boldsymbol{\alpha})) = \|\mathbf{t}\|_{\boldsymbol{\alpha}}^2$$

where we have applied the annulus inequality for the second comparison.

Lemma 3.3. *Let A be an annulus on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ with core curve γ. Then for all $\beta \in \mathcal{C}(S)$, we have*

$$\mathrm{width}(A) \times i(\gamma, \beta) \leq l(\beta).$$

Proof. Let $b : S^1 \to S(\mathbf{t} \cdot \boldsymbol{\alpha})$ be an map which realises β as a geodesic on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$. We may pull back the metric on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ via b to give S^1 the structure of a circle of length $l(\beta)$. The result follows by observing that $b^{-1}(A) \subseteq S^1$ contains at least $i(\gamma, \beta)$ disjoint arcs, each having length at least $\mathrm{width}(A)$. $\square$

By combining the above inequalities with Proposition 3.1, we can control weighted intersection numbers with the core curve of A:

Corollary 3.4. *Let A be an essential annulus on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ with core curve γ. Then*

$$i(\mathbf{t} \cdot \boldsymbol{\alpha}, \gamma) \leq \frac{\sqrt{2}\,\|\mathbf{t}\|_{\boldsymbol{\alpha}}^2}{\mathrm{width}(A)} \quad \text{and} \quad i(\gamma, \beta) \leq \frac{i(\mathbf{t} \cdot \boldsymbol{\alpha}, \beta)}{\mathrm{width}(A)}$$

for all $\beta \in \mathcal{C}(S)$. $\square$

3.3 A grid structure on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$

In this section, we describe a grid structure on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ and give a proof of the second claim of Proposition 3.1.

Definition 3.5 (Quarter-translation surface)**.** A *quarter-translation surface* is a topological surface S with a finite set of singularities ς together with an atlas of charts from $S - \varsigma$ to $\mathbb{R}^2$ whose transition maps are translations of $\mathbb{R}^2$ possibly composed with rotations through integral multiples of $\frac{\pi}{2}$. The singular points have cone angles which are integral multiples of $\frac{\pi}{2}$ and at least π.

Suppose S is a quarter-translation surface. We may pull back the standard Euclidean metric on $\mathbb{R}^2$ to give a singular Euclidean metric on S. Geodesics which do not meet any singular points or marked points with respect to this metric can only self-intersect orthogonally. We will also define an L^1 metric on S by pulling back the metric given infinitesimally by $|dx| + |dy|$ on $\mathbb{R}^2$. The following is immediate:

Lemma 3.6. *Let $l^2(\eta)$ and $l^1(\eta)$ denote, respectively, the Euclidean and L^1 lengths of a path η on S. Then $l^2(\eta) \leq l^1(\eta) \leq \sqrt{2} l^2(\eta)$.* □

Whenever we deal with quarter-translation surfaces, we will assume that we are working with the singular Euclidean metric unless otherwise specified.

The transition maps between the charts for S preserve a pair of orthogonal directions on $\mathbb{R}^2$ which we may take to be the standard horizontal and vertical directions. By pulling these back via the coordinate charts, we can equip S with a preferred (unordered) pair of orthogonal directions defined on the complement of the singular points. We shall refer to these as the *grid directions*. Geodesics which run parallel to a grid direction will be called *grid arcs*. Every non-singular point on S has an open rectangular neighbourhood, with sides parallel to the grid directions, on which the grid leaves restrict to give a pair of transverse foliations. Such a rectangle will be called an *open grid rectangle*.

It is straightforward to check that $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ is a quarter-translation surface. We will assume that the grid directions on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ run parallel to the sides of the rectangles used in its construction.

Lemma 3.7. *Given a curve $\gamma \in \mathcal{C}(S)$, let c be any of its geodesic representatives on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ with respect to the Euclidean metric. Then*

$$l^1(c) = i(\mathbf{t} \cdot \boldsymbol{\alpha}, \gamma).$$

Proof. If c is an embedded simple closed loop then we can isotope it to another geodesic representative which meets at least one singularity. Thus we can assume that S^1 decomposes as a finite union of intervals $\cup I_k$ with disjoint interiors such that $c : S^1 \to S(\mathbf{t} \cdot \boldsymbol{\alpha})$ embeds each I_k as a straight line segment with singularities or marked points at its endpoints.

We can homotope c to a closed path $c' : S^1 \to S(\mathbf{t} \cdot \boldsymbol{\alpha})$ so that each $c'(I_k)$ is an edge–path in the 1–skeleton of $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ with the same endpoints as $c(I_k)$. The homotopy can be performed in a way which preserves the l^1–length of the path and without creating new intersection points with any of the α_i. One can check that c intersects each α_i minimally and thus the same is also true of c'. Finally, we deduce

$$l^I(c'(S^1)) = i(\mathbf{t} \cdot \boldsymbol{\alpha}, \gamma)$$

by observing that every edge in the 1–skeleton of $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ transverse to α_i has length t_i. □

The second claim of Proposition 3.1 follows from the previous two lemmas.

3.4 An isoperimetric inequality

Let $S = (S, \Omega)$ be a closed singular Riemannian surface S with a finite set of marked points Ω. Let Δ be a closed disc and suppose $\iota : \Delta \to S$ is a piecewise smooth immersion which restricts to an embedding on its interior. Let D denote the image $\iota(\mathrm{int}(\Delta))$.

Definition 3.8. An open disc D arising in the above manner is called a *trivial region* on S if it contains at most one marked point.

Bowditch defines trivial regions as open discs on S containing at most one marked point without any conditions concerning piecewise smooth embeddings. Nevertheless, his proof of the following proposition still holds with our definition:

Proposition 3.9 ([Bow06b]). *Suppose $f : [0, \infty) \to [0, \infty)$ is a homeomorphism. Let ρ be a singular Riemannian metric on an orientable closed surface S with unit area. Let Ω be a finite set of marked points on S. We will assume $|\Omega| \geq 5$ whenever S is a 2–sphere. If* $\mathrm{area}(D) \leq f(\mathrm{length}(\partial D))$ *for any trivial region D then there is an essential annulus $A \subseteq S - \Omega$ such that* $\mathrm{width}(A) \geq \mathsf{W}_0$*, where* $\mathsf{W}_0 > 0$ *depends only on $\xi(S)$ and f.*

A little care is required to clarify what $\mathrm{length}(\partial D)$ means, especially when ∂D is not an embedded copy of a circle. In general, the boundary ∂D is an embedded

Eulerian graph on S whose edges are piecewise smooth arcs. We define length(∂D) to be the sum of the lengths of the arcs with respect to the metric on S.

This section will be devoted to proving the following lemma which, together with the above proposition, implies the third claim of Proposition 3.1.

Lemma 3.10. *Suppose D is a trivial region on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$. Then*

$$\operatorname{area}(D) \leq 4 \operatorname{length}(\partial D)^2.$$

Before launching into the details of the proof, we give a brief outline of our argument. First, we reduce the problem to that of studying embedded closed discs on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ whose boundary is a finite union of grid arcs. We then show that such a disc D can be given a tiling by grid rectangles. This tiling is dual to a collection of arcs on D, where each arc is parallel to a component of some $\alpha_i \cap D$. We shall call the union of all rectangles meeting a given arc a *band*. The key step is to observe that any two arcs in the collection intersect at most twice. Thus, the intersection of two distinct bands is the union of at most two rectangles arising from the tiling. Conversely, any rectangle from the tiling is contained in the intersection of two such bands. This then allows us to bound the area of the rectangles in terms of the length of ∂D.

3.4.1 Technical adjustments

Let us first make a couple of observations to simplify the problem.

Lemma 3.11. *Any trivial region D on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ can be perturbed to a trivial region D' whose boundary is a finite union of grid leaves. Moreover, D' can be chosen so that* $\operatorname{area}(D') \geq \operatorname{area}(D)$ *and* $\operatorname{length}(\partial D') \leq \sqrt{2} \operatorname{length}(\partial D)$. □

We will henceforth assume that the boundary of any trivial region on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ is a finite union of grid leaves.

Let $\iota : \Delta \to S(\mathbf{t} \cdot \boldsymbol{\alpha})$ be a piecewise smooth immersion whose restriction to $\operatorname{int}(\Delta)$ is an embedding with image D. Observe that $\iota : \partial\Delta = S^1 \to \partial D$ is an immersion of a circle which runs over each edge of ∂D at most twice. We will metrise Δ by pulling back the metric on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ via ι.

Lemma 3.12. *Suppose D and Δ are as given above. Then* $\operatorname{area}(\Delta) = \operatorname{area}(D)$ *and* $\operatorname{length}(\partial D) \leq \operatorname{length}(\partial \Delta) \leq 2 \operatorname{length}(\partial D)$. □

3.4.2 Tiling Δ by rectangles

The disc Δ inherits grid directions from $S(\mathbf{t}\cdot\boldsymbol{\alpha})$ via ι away from the preimage of the singular points. The boundary decomposes as a finite union $\partial\Delta = \cup I_k$ of closed grid arcs with disjoint interiors. We may assume that this decomposition is minimal, that is, it cannot be obtained from any other such decomposition by subdividing arcs. An endpoint of any grid arc I_k will be called a *corner point* of $\partial\Delta$. A corner point which does not coincide with a singularity or a marked point must be an orthogonal intersection point of two grid arcs.

It is worth noting that $\partial\Delta$ must contain at least two corner points and at least three if D contains no marked points. To see this, recall that the grid leaves on $S(\mathbf{t}\cdot\boldsymbol{\alpha})$ are parallel to some α_i. Any of the forbidden cases will imply that some α_i is trivial, peripheral, self-intersects or does not intersect some α_j minimally.

Let us refer to marked points, corner points and singularities collectively as *bad points*. Let $Z \subset \Delta$ be the union of $\partial\Delta$ with all grid arcs in Δ which have a bad point for at least one of their endpoints. Since there are finitely many bad points in Δ, it follows that Z is a finite embedded graph on Δ. A vertex $v \in \mathrm{int}\Delta \cap Z$ has valence k if and only if the cone angle at v is $\frac{k\pi}{2}$. If v is a vertex which lies on $\partial\Delta$ then it has valence $k+1$ if and only if the cone angle at v inside Δ is $\frac{k\pi}{2}$. It follows that every vertex v of Z has valence at least 2, and at least 3 if v is not a marked point.

Lemma 3.13. *There exists a tiling of Δ by finitely many grid rectangles with Z as its 1–skeleton.*

Proof. First note that there are finitely many connected components of $\Delta - Z$ since Z is a finite graph. Let R be such a component and let $\bar{R}$ be its completion with respect to its induced path metric. Observe that $\bar{R}$ is a closed planar region admitting a Euclidean metric with piecewise geodesic boundary, where the interior angle between adjacent edges of $\partial\bar{R}$ is $\frac{\pi}{2}$. By the Gauss–Bonnet formula, the sum of its interior angles must equal $2\pi\chi(R)$. Since the frontier of R in Δ meets at least one vertex of Z, the angle sum must be strictly positive. As R is planar, it follows that $\chi(R) = 1$ and therefore $\bar{R}$ is a Euclidean rectangle. Also note that Z is connected, for otherwise there would exist some component of $\Delta - Z$ with disconnected frontier.

The inclusion $R \hookrightarrow \Delta$ can be extended continuously to a map $\bar{R} \to \Delta$, sending each edge of $\partial\bar{R}$ isometrically to an edge of Z meeting the frontier of R. Thus R is a grid rectangle since the edges of Z, by construction, are parallel to the grid directions. Finally, the closures of distinct rectangles R and R' can only intersect in a union of vertices and edges of Z. $\square$

3.4.3 Controlling the area

Let $\mathcal{A}$ be the set of maximal grid arcs in Δ which intersect Z only at midpoints of edges of Z. This is a collection of arcs dual to the tiling of Δ by rectangles described in Lemma 3.13. (There cannot exist any closed curves in Δ dual to the tiling as this would imply that some α_i is trivial or peripheral.) Given an arc $a \in \mathcal{A}$, let $B = B(a)$ be the union of all rectangles in the tiling which meet a. We will call B a *band* and a a *core arc* of B. Define width(B) to be the length of any edge of Z crossed by a. Note that the set of bands in Δ is in bijection with $\mathcal{A}$.

Lemma 3.14. *The intersection of two distinct bands B and B' is the union of at most two rectangles whose side lengths are* width(B) *by* width(B'). *Conversely, each rectangle in the tiling lies in the intersection of a unique pair of distinct bands.*

Proof. Let a and a' be core arcs of B and B' respectively. If a and a' intersect at least 3 times then they must bound a bigon in Δ containing no marked points. Now, a and a' can both be properly isotoped in Δ to components of $\iota^{-1}(\alpha_i \cap D)$ and $\iota^{-1}(\alpha_i \cap D)$ for some α_i and α_j respectively. Moreover, the isotopies can be performed without passing through any singular points or marked points. This procedure cannot destroy any bigons since any right-angled bigon on Δ must contain at least one singularity. It follows that α_i and α_j also bound a bigon in D, contradicting minimality.

For the converse, simply take the bands corresponding to the unique pair of arcs which have an intersection point inside the given rectangle. □

We will refer to an edge of Z lying in $\partial\Delta$ simply as an *edge* of $\partial\Delta$.

Lemma 3.15. *Under the above hypotheses, we have*

$$\text{length}(\partial\Delta) = 2\sum_B \text{width}(B)$$

where the sum is taken over all bands B in Δ.

Proof. Let B be a band in Δ. Observe that $B \cap \partial\Delta$ consists of exactly two edges of $\partial\Delta$ whose length is equal to width(B). Conversely, each edge of $\partial\Delta$ lies in exactly one band. □

Lemma 3.16. *Let Δ be as above. Then*

$$\text{area}(\Delta) \leq \frac{1}{2}\,\text{length}(\partial\Delta)^2.$$

Proof. By Lemma 3.14, Δ is a union of rectangles, each of which lies in the intersection of a pair of distinct bands. Thus

$$\operatorname{area}(\Delta) = \operatorname{area}\left(\bigcup_{B\neq B'} B\cap B'\right) = \sum_{B\neq B'} \operatorname{area}(B\cap B').$$

Since the intersection of two distinct bands is the union of at most two rectangles whose side lengths are equal to the widths of the bands, we have

$$\operatorname{area}(B\cap B') \leq 2 \operatorname{width}(B) \times \operatorname{width}(B')$$

and hence

$$\operatorname{area}(\Delta) \leq 2 \sum_{B\neq B'} \operatorname{width}(B) \times \operatorname{width}(B') \leq 2 \left(\sum_{B} \operatorname{width}(B)\right)^2.$$

Finally, applying Lemma 3.15 completes the proof. □

Combining this result with Lemmas 3.11 and 3.12 completes the proof of Lemma 3.10. □

3.5 Non-filling curves

We now generalise the construction of $S(\mathbf{t}\cdot\boldsymbol{\alpha})$ to encompass non-filling curves. Assume $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ is an n–tuple of distinct multicurves and $\mathbf{t} = (t_1, \ldots, t_n) \neq \mathbf{0}$ is a weight vector satisfying $\|\mathbf{t}\|_{\boldsymbol{\alpha}} \neq 0$. Realise $\boldsymbol{\alpha}$ minimally on S to form a 4–valent graph Υ on S. Let $\Sigma \subseteq S$ be the (possibly disconnected) subsurface filled by Υ. This can be obtained by taking a closed regular neighbourhood of Υ on S and then attaching all complementary regions which are discs with at most one marked point. If $\boldsymbol{\alpha}$ fills S then $\Sigma = S$. In general, Σ will be a disjoint union of surfaces $\Sigma_1 \cup \ldots \cup \Sigma_s$. Observe that $s \leq \xi(S)$ since we can find a multicurve on S so that exactly one component is contained in each Σ_k (by choosing a suitable subset of all curves appearing in $\boldsymbol{\alpha}$, for example). Some of these components may be annuli – this occurs precisely when a multicurve α_i has a component disjoint from all other α_j. All other components will have genus at least one, or are spheres where the sum of the number of marked points and boundary components is at least four.

We now define a 2–dimensional complex $S(\mathbf{t}\cdot\boldsymbol{\alpha})$ as a quotient of S. Suppose Σ_k is an annular component of Σ whose core curve is a component of α_i. We identify Σ_k with $S^1 \times [0, t_i]$ and then collapse the first co-ordinate to give a closed interval

I_k of length t_i. Next, we collapse every complementary component of Σ in S to a marked point. These marked points will be called *essential*. We then apply the construction given in Section 3.1 to the image of each non-annular component of Σ in the quotient space. The resulting space is a finite collection of singular Euclidean surfaces and closed intervals identified along appropriate essential marked points. Note that this construction agrees with the one given in Section 3.1 for the case of filling curves.

Let c be a representative of a curve $\gamma \in \mathcal{C}(S)$ on S. Its image $\bar{c}$ on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ will be a closed curve or a union of paths connecting essential marked points. Define $l(\gamma)$ to be the minimal length of $\bar{c}$ over all representatives c of γ.

Proposition 3.17. *Suppose $\boldsymbol{\alpha}$ and $\mathbf{t}$ satisfy $\|\mathbf{t}\|_{\boldsymbol{\alpha}} > 0$. Then the first two claims of Proposition 3.1 hold for $S(\mathbf{t} \cdot \boldsymbol{\alpha})$.* □

The proof of the above proceeds in the same manner as for the case of filling curves. It remains to prove an analogue of the third claim.

We will refer to the image of a component Σ_k of Σ in $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ as a *component* of $S(\mathbf{t} \cdot \boldsymbol{\alpha})$. Let Y be a component of $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ with maximal area. Since Σ has at most $\xi(S)$ components, we have $\mathrm{area}(Y) \geq \frac{\|\mathbf{t}\|^2_{\boldsymbol{\alpha}}}{\xi(S)}$. Note that Y cannot be an interval since $\|\mathbf{t}\|_{\boldsymbol{\alpha}} > 0$. We may argue as in Section 3.4 to show the following.

Lemma 3.18. *Suppose Y has genus at least one, or is a sphere with at least five marked points. Then there is an essential annulus on Y whose width is at least $\frac{\mathsf{W}_0\|\mathbf{t}\|_{\boldsymbol{\alpha}}}{\sqrt{\xi(S)}}$, where W_0 is a constant depending only on $\xi(Y) \leq \xi(S)$.* □

Let us assume Y is a sphere with 4 marked points for the rest of this section. We will not prove the existence of annuli of definite width on Y. Instead, we show that it suffices to find wide annuli on a torus $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$ which double branch covers Y for the purposes of proving Lemma 4.5.

Define a hyperelliptic involution $h : \mathbb{T}^2 \to \mathbb{T}^2$ by setting $h(x, y) = (-x, -y)$ modulo $\mathbb{Z}^2$. The quotient map induced by the action of $\langle h \rangle$ is a double cover $P : \mathbb{T}^2 \to S^2$ branched over four points. The branch points correspond to the fixed orbits $(0, 0)$, $(0, \frac{1}{2})$, $(\frac{1}{2}, 0)$ and $(\frac{1}{2}, \frac{1}{2})$ of h. We will identify Y with the quotient space S^2 so that the marked points coincide exactly with the branch points. We then metrise $\mathbb{T}^2$ by lifting the singular Euclidean metric on Y via P. This metric can also be obtained by taking the preimages of the α_i contained in Y to $\mathbb{T}^2$ and then applying the construction as described in Section 3.1. It follows from the work in Section 3.4 that $\mathbb{T}^2$ enjoys the isoperimetric inequality stated in Lemma 3.10. Invoking Proposition 3.9 gives the following:

Lemma 3.19. *There exists an essential annulus on $\mathbb{T}^2$ of width at least $\frac{\mathsf{W}\|\mathbf{t}\|_{\boldsymbol{\alpha}}}{\sqrt{\xi(S)}}$ for some universal constant $\mathsf{W} > 0$.* $\square$

Remark 3.20. By following Bowditch's proof of Proposition 3.9 in [Bow06b] for the the case of the torus, one can show that setting $\mathsf{W} = \frac{1}{3\sqrt{2}}$ will satisfy the statement of the above lemma.

Observe that $h(\tilde{\gamma})$ is homotopic to $\tilde{\gamma}$ for any simple closed curve $\tilde{\gamma}$ on $\mathbb{T}^2$. Thus, the image of any simple closed on $\mathbb{T}^2$ under P is homotopic to a simple closed curve on Y.

Lemma 3.21. *Let A be an essential annulus on $\mathbb{T}^2$ with core curve $\tilde{\gamma}$. Let $\gamma \in \mathcal{C}(S)$ be the image of $\tilde{\gamma}$ on Y under P. Then*

$$i(\gamma, \beta) \leq \frac{2\, i(\mathbf{t} \cdot \boldsymbol{\alpha}, \beta)}{\text{width}(A)}$$

for all $\beta \in \mathcal{C}(S)$.

Proof. First, recall that $\beta \cap Y$ is either a simple closed curve or a union of paths connecting marked points of Y. The preimage $P^{-1}(\beta)$ is a finite union of (not necessarily disjoint) essential curves on $\mathbb{T}^2$. By perturbing γ to an embedded curve which misses the marked points of Y, we see that each point of $\gamma \cap \beta$ lifts to exactly two points on $\mathbb{T}^2$ under P, and so

$$i(\gamma, \beta) = \frac{i(P^{-1}(\gamma), P^{-1}(\beta))}{2} \leq i(\tilde{\gamma}, P^{-1}(\beta)).$$

Applying Lemma 3.3 to each curve in $P^{-1}(\beta)$ gives

$$\text{width}(A) \times i(\tilde{\gamma}, P^{-1}(\beta)) \leq l(P^{-1}(\beta)).$$

Next, we have

$$l(P^{-1}(\beta)) = 2\, l(\beta \cap Y) \leq 2\, l(\beta) \leq 2\, i(\mathbf{t} \cdot \boldsymbol{\alpha}, \beta)$$

where we have applied Proposition 3.17 for the final comparison. Finally, combining the above inequalities gives the desired result. $\square$

Chapter 4

Hulls in the curve complex

Let $S = (S, \Omega)$ be a connected compact surface S without boundary with a finite set of marked points Ω satisfying $\xi(S) \geq 2$. Throughout this chapter, we will fix an n-tuple $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ of distinct multicurves in $\mathcal{C}(S)$, where $n \geq 2$. We will assume that no pair α_i and α_j has a common component.

We shall establish a coarse equality between two subsets of $\mathcal{C}(S)$ determined by $\boldsymbol{\alpha}$ – its hyperbolic hull $\mathrm{Hull}(\boldsymbol{\alpha})$, defined purely in terms of the geometry of $\mathcal{C}(S)$; and $\mathrm{Short}(\boldsymbol{\alpha}, \mathsf{L})$ which is defined using only intersection numbers. We also give a combinatorial method of approximating nearest point projections to $\mathrm{Hull}(\boldsymbol{\alpha})$.

4.1 Hyperbolic hulls

Let $\mathcal{X}$ be a δ–hyperbolic space and suppose $U \subseteq \mathcal{X}$ is a set of points. The *hyperbolic hull* of U, denoted $\mathrm{Hull}(U)$, is the union of all geodesic segments in $\mathcal{X}$ connecting a pair of points in U.

Example 4.1. Let U be a finite subset of $\mathbb{H}^n$, where $n \geq 1$. Then $\mathrm{Hull}(U)$ is a uniformly bounded Hausdorff distance away from the convex hull of U in $\mathbb{H}^n$.

Lemma 4.2. *The hyperbolic hull of any non-empty set $U \subseteq \mathcal{X}$ is 2δ–quasiconvex.*

Proof. Let u and v be points in $\mathrm{Hull}(U)$. Let $x, y, z, w \in U$ be points, not necessarily distinct, so that $u \in [x, y]$ and $v \in [z, w]$. Let $[u, v]$ be any geodesic segment. By δ–hyperbolicity, we have

$$[u, v] \subseteq_\delta [u, y] \cup [y, v] \subseteq_\delta [u, y] \cup [y, z] \cup [z, v] \subseteq [x, y] \cup [y, z] \cup [z, w] \subseteq \mathrm{Hull}(U),$$

where $[u, y]$ and $[z, v]$ are assumed to be subarcs of $[x, y]$ and $[z, w]$ respectively. $\square$

Lemma 4.3. *Suppose $C \subseteq \mathcal{X}$ is a Q–quasiconvex set which contains U. Then* $\mathrm{Hull}(U) \subseteq_{\mathsf{Q}} C$.

Proof. This follows immediately from the definition of quasiconvexity. $\square$

In fact, the above properties characterises hyperbolic hulls up to finite Hausdorff distance.

Corollary 4.4. *Let $U \subseteq \mathcal{X}$ be non-empty. Suppose $C \subseteq \mathcal{X}$ is a Q–quasiconvex set such that*

1. *C contains U, and*
2. *for any Q'–quasiconvex set $C' \subseteq \mathcal{X}$ which also contains U, we have $C \subseteq_{\mathsf{r}} C'$ for some $\mathsf{r} = \mathsf{r}(\mathsf{Q}, \mathsf{Q}') \geq 0$.*

Then C and $\mathrm{Hull}(U)$ *agree up to finite Hausdorff distance.* $\square$

4.2 A hull via intersection numbers

4.2.1 Short curve sets

Let $\mathbf{t} = (t_1, \ldots, t_n) \neq \mathbf{0}$ be a weight vector. Given $\mathsf{L} \geq 0$, define

$$\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L}) := \{\gamma \in \mathcal{C}(S) \mid i(\mathbf{t} \cdot \boldsymbol{\alpha}, \gamma) \leq \mathsf{L} \left\|\mathbf{t}\right\|_{\boldsymbol{\alpha}}\}.$$

If $\left\|\mathbf{t}\right\|_{\boldsymbol{\alpha}} = 0$ then this set is contained in the 1–neighbourhood of $\boldsymbol{\alpha}$. If $\left\|\mathbf{t}\right\|_{\boldsymbol{\alpha}} > 0$ then $\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L})$ can be viewed as the set of curves of bounded length on $S(\mathbf{t} \cdot \boldsymbol{\alpha})$ after rescaling it to have unit area. Note that $\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L})$ remains invariant under multiplying $\mathbf{t}$ by a positive scalar.

The proof of the following lemma is essentially the same as in Bowditch's paper ([Bow06b] Lemma 4.1).

Lemma 4.5. *There exists a constant $\mathsf{L}_0 > 0$ depending only on $\xi(S)$ such that, for any $\mathsf{L} \geq \mathsf{L}_0$, the set* $\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L})$ *is non-empty. Moreover,*

$$\mathrm{diam}_{\mathcal{C}(S)}(\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L})) \leq 4 \log_2 \mathsf{L} + \mathsf{k}_0,$$

where k_0 is a constant depending only on $\xi(S)$.

Proof. If $\left\|\mathbf{t}\right\|_{\boldsymbol{\alpha}} = 0$ then $\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L})$ contains the α_i and is contained in the 1–neighbourhood of $\boldsymbol{\alpha}$ in $\mathcal{C}(S)$.

Next, suppose $\|\mathbf{t}\|_{\boldsymbol{\alpha}} > 0$ and $\boldsymbol{\alpha}$ fills S. By the third claim of Proposition 3.1, there exists an essential annulus A on $S(\mathbf{t}\cdot\boldsymbol{\alpha})$ whose width is at least $\mathsf{W}_0 \|\mathbf{t}\|_{\boldsymbol{\alpha}}$, where W_0 depends only on $\xi(S)$. Set $\mathsf{L}_0 = \frac{\sqrt{2}}{\mathsf{W}_0}$. Let γ be the core curve of A. By Corollary 3.4, we have

$$i(\mathbf{t}\cdot\boldsymbol{\alpha}, \gamma) \leq \frac{\sqrt{2}\,\|\mathbf{t}\|_{\boldsymbol{\alpha}}^2}{\mathrm{width}(A)} \leq \frac{\sqrt{2}}{\mathsf{W}_0}\|\mathbf{t}\|_{\boldsymbol{\alpha}}$$

and hence $\mathrm{short}\,(\mathbf{t}\cdot\boldsymbol{\alpha}, \mathsf{L}) \neq \emptyset$ for all $\mathsf{L} \geq \mathsf{L}_0$. Furthermore, if $\beta \in \mathrm{short}\,(\mathbf{t}\cdot\boldsymbol{\alpha}, \mathsf{L})$ is another curve then, by Corollary 3.4, we have

$$i(\gamma, \beta) \leq \frac{i(\mathbf{t}\cdot\boldsymbol{\alpha}, \beta)}{\mathrm{width}(A)} \leq \frac{\mathsf{L}\,\|\mathbf{t}\|_{\boldsymbol{\alpha}}}{\mathsf{W}_0\,\|\mathbf{t}\|_{\boldsymbol{\alpha}}} = \frac{\mathsf{L}}{\mathsf{W}_0}.$$

Applying Lemma 2.5 and the triangle inequality gives

$$\mathrm{diam}(\mathrm{short}(\mathbf{t}\cdot\boldsymbol{\alpha}, \mathsf{L})) \leq 2\left[2\log_2\left(\frac{\mathsf{L}}{\mathsf{W}_0}\right) + 2\right] = 4\log_2 \mathsf{L} + \mathsf{k}_0$$

where k_0 is a constant depending only on $\xi(S)$.

For the case where $\|\mathbf{t}\|_{\boldsymbol{\alpha}} > 0$ but $\boldsymbol{\alpha}$ does not fill S, it is immediate that $\mathrm{short}(\mathbf{t}\cdot\boldsymbol{\alpha}, \mathsf{L})$ is non-empty. To bound the diameter, invoke Lemmas 3.18, 3.19 and 3.21 then argue as above. □

Consequently, up to bounded error, we can view $\mathrm{short}(\mathbf{t}\cdot\boldsymbol{\alpha}, \mathsf{L})$ as a single curve in $\mathcal{C}(S)$ which has minimal intersection number with $\mathbf{t}\cdot\boldsymbol{\alpha}$.

4.2.2 The short curve hull

For $\mathsf{L} \geq 0$, define the L–*short curve hull* of $\boldsymbol{\alpha}$ to be

$$\mathrm{Short}(\boldsymbol{\alpha}, \mathsf{L}) := \bigcup_{\mathbf{t}} \mathrm{short}(\mathbf{t}\cdot\boldsymbol{\alpha}, \mathsf{L}),$$

where the union is taken over all weight vectors $\mathbf{t} \in \mathbb{R}^n_{\geq 0}$ (or, equivalently, by choosing one representative from each projective class).

We write $\mathrm{Hull}(\boldsymbol{\alpha}) \subseteq \mathcal{C}(S)$ for the hyperbolic hull of $\boldsymbol{\alpha}$ considered as a set of vertices in $\mathcal{C}(S)$.

Proposition 4.6. *Let $\boldsymbol{\alpha}$ be an n–tuple of multicurves in $\mathcal{C}(S)$. Then for any $\mathsf{L} \geq \mathsf{L}_0$,*

$$\mathrm{Short}(\boldsymbol{\alpha}, \mathsf{L}) \approx_{\mathsf{k}_1} \mathrm{Hull}(\boldsymbol{\alpha})$$

where k_1 *depends only on* $\xi(S)$*,* n *and* L*.*

This is essentially an extension of Bowditch's coarse description of geodesics using intersection numbers employed in his proof of hyperbolicity of the curve complex [Bow06b]. Let us begin with a reformulation of his result:

Lemma 4.7 ([Bow06b] Proposition 6.2)**.** *Let* $\boldsymbol{\alpha}' = (\alpha_1, \alpha_2)$ *be a pair of multicurves in* $\mathcal{C}(S)$*. Let* $[\alpha_1, \alpha_2]$ *denote any geodesic segment connecting* α_1 *and* α_2 *in* $\mathcal{C}(S)$*. Then for all* $\mathsf{L} \geq \mathsf{L}_0$*, we have*

$$\mathrm{Short}(\boldsymbol{\alpha}', \mathsf{L}) \approx_{\mathsf{k}_1'} [\alpha_1, \alpha_2].$$

where $\mathsf{k}_1' \geq 0$ *depends only* $\xi(S)$ *and* L*.* □

Proof of Proposition 4.6. By applying the previous lemma to all pairs of multicurves (α_i, α_j) in $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_2)$, we obtain the inclusion:

$$\mathrm{Hull}(\boldsymbol{\alpha}) \subseteq_{\mathsf{k}_1'} \mathrm{Short}(\boldsymbol{\alpha}, \mathsf{L}).$$

Fix a weight vector $\mathbf{t} = (t_1, \ldots, t_n)$ and assume, for notational simplicity, that the quantity $t_j t_k i(\alpha_j, \alpha_k)$ is maximised when $\{j, k\} = \{1, 2\}$. Let $\boldsymbol{\alpha}' = (\alpha_1, \alpha_2)$ and $\mathbf{t}' = (t_1, t_2)$. Since there are $\frac{n(n-1)}{2}$ distinct unordered pairs of indices $\{j, k\}$, it follows that

$$\|\mathbf{t}\|_{\boldsymbol{\alpha}}^2 = \sum_{j<k} t_j t_k i(\alpha_j, \alpha_k) \leq \frac{n(n-1)}{2} t_1 t_2 i(\alpha_1, \alpha_2) = \frac{n(n-1)}{2} \left\|\mathbf{t}'\right\|_{\boldsymbol{\alpha}'}^2.$$

Let γ be a curve in $\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L})$. Then

$$i(\mathbf{t}' \cdot \boldsymbol{\alpha}', \gamma) \leq i(\mathbf{t} \cdot \boldsymbol{\alpha}, \gamma) \leq \mathsf{L} \left\|\mathbf{t}\right\|_{\boldsymbol{\alpha}} \leq \mathsf{L}\sqrt{\frac{n(n-1)}{2} \left\|\mathbf{t}'\right\|_{\boldsymbol{\alpha}'}^2} \leq \frac{n\mathsf{L}}{\sqrt{2}} \left\|\mathbf{t}'\right\|_{\boldsymbol{\alpha}'}$$

which implies

$$\mathrm{short}(\mathbf{t} \cdot \boldsymbol{\alpha}, \mathsf{L}) \subseteq \mathrm{short}\left(\mathbf{t}' \cdot \boldsymbol{\alpha}', \frac{n\mathsf{L}}{\sqrt{2}}\right).$$

Invoking Lemma 4.7, we have

$$\mathrm{short}\left(\mathbf{t}' \cdot \boldsymbol{\alpha}', \frac{n\mathsf{L}}{\sqrt{2}}\right) \subseteq_{\mathsf{r}} [\alpha_1, \alpha_2] \subseteq \mathrm{Hull}(\boldsymbol{\alpha})$$

where $\mathsf{r} \geq 0$ is some constant depending on n, L and $\xi(S)$. This concludes the proof of Proposition 4.6. □

We can describe the above proof in terms of the geometry of $S(\mathbf{t}\cdot\boldsymbol{\alpha})$. Assume $S(\mathbf{t}\cdot\boldsymbol{\alpha})$ has unit area. One can obtain $S(\mathbf{t}'\cdot\boldsymbol{\alpha}')$ by homotoping the annuli consisting of rectangles traversed by α_i to the core curve α_i for each $i \neq 1, 2$. The maximality assumption on α_1 and α_2 ensures that the total area of the remaining rectangles is at least $\frac{2}{n(n-1)}$. We then scale $S(\mathbf{t}'\cdot\boldsymbol{\alpha}')$ by a factor of at most $\frac{n}{\sqrt{2}}$ to give it unit area. Finally, observe that the length of a curve γ on $S(\mathbf{t}\cdot\boldsymbol{\alpha})$ can only increase by a factor of at most $\frac{n}{\sqrt{2}}$ during this process.

4.3 Nearest point projections to hulls

In this section, we approximate nearest point projections to short curve hulls using only intersection number conditions.

Definition 4.8. Let $\beta \in \mathcal{C}(S)$ be a multicurve. A weight vector $\mathbf{t} = (t_1, \ldots, t_n)$ satisfying

$$t_j i(\alpha_j, \beta) = t_k i(\alpha_k, \beta)$$

for all j, k is called a *balance vector* for β with respect to $\boldsymbol{\alpha}$.

If β intersects all α_i then setting $t_i = i(\alpha_i, \beta)^{-1}$ yields the unique balance vector up to positive scale. If not, we can set $t_i = 1$ whenever $i(\alpha_i, \beta) = 0$ and $t_i = 0$ otherwise to produce a balance vector. We will use $\mathbf{t}_\beta$ to denote any balance vector for β. We also remark that the above definition is analogous to the notion of *balance time* for quadratic differentials as described by Masur and Minsky [MM99].

Proposition 4.9. *Assume $\mathsf{L} \geq \mathsf{L}_0$. Given a multicurve $\beta \in \mathcal{C}(S)$, let $\mathbf{t}_\beta$ be any balance vector with respect to $\boldsymbol{\alpha}$. Let γ be any nearest point projection of β to* $\mathrm{Hull}(\boldsymbol{\alpha})$. *Then*

$$\gamma \approx_{\mathsf{k}_2} \mathrm{short}(\mathbf{t}_\beta \cdot \boldsymbol{\alpha}, \mathsf{L}),$$

where $\mathsf{k}_2 \geq 0$ depends only on $\xi(S)$, n and L.

As was the case with Proposition 4.6, this is an extension of a result of Bowditch. His result was originally phrased in terms of centres for geodesic triangles, however, our statement agrees with it up to uniformly bounded error.

Lemma 4.10 ([Bow06b] Proposition 3.1 and Section 4)**.** *Let α_1, α_2 and β be multicurves in $\mathcal{C}(S)$. Let $\mathbf{t}'_\beta$ be a balance vector for β with respect to $\boldsymbol{\alpha}' = (\alpha_1, \alpha_2)$. Let γ be a nearest point projection of β to $[\alpha_1, \alpha_2]$. Then*

$$\gamma \approx_{\mathsf{k}'_2} \mathrm{short}(\mathbf{t}'_\beta \cdot \boldsymbol{\alpha}', \mathsf{L}),$$

where k_2' *depends only on* $\xi(S)$ *and* L. $\square$

If β is disjoint from some α_i then Proposition 4.9 follows immediately from Hempel's bound (Lemma 2.5). We will henceforth assume this is not the case.

Our first step is to reduce the problem of finding a nearest point projection to a hyperbolic hull to that of projecting to a suitable geodesic.

Lemma 4.11. *Let* U *be a subset of a* δ*–hyperbolic space* $\mathcal{X}$. *Fix a point* $w \in \mathcal{X}$. *Assume there exist* $x, y \in U$ *and* $\mathsf{R} \geq 0$ *such that*

$$d_{\mathcal{X}}([x,y],[z,w]) \leq \mathsf{R}.$$

for all $z \in U$. *Let* p *and* q *be nearest point projections of* w *to* $\mathrm{Hull}(U)$ *and* $[x,y]$ *respectively. Then*

$$p \approx_{\mathsf{R}'} q$$

where R' *depends only on* R *and* δ.

Proof. By Lemma 1.12, it suffices to show that for all $u \in \mathrm{Hull}(U)$, any geodesic $[w,u]$ must pass within a bounded distance of q. If u lies on a geodesic segment $[z,z']$ for some $z, z' \in U$ then $[w,u]$ must lie inside the 2δ–neighbourhood of $[w,z]$ or $[w,z']$. Hence, we only need to bound $d(q,[w,z])$ for all $z \in U$ in terms of δ and R.

Recall that geodesic segments are δ–quasiconvex. If z coincides with x or y then $d(q,[w,z]) \leq 3\delta$ by Lemma 1.10. Let us now suppose that x, y and z are distinct. Choose points $u \in [z,w]$ and $v \in [x,y]$ so that $uv = d_{\mathcal{X}}([x,y],[z,w]) \leq \mathsf{R}$. Then

$$q \subseteq_{3\delta} [w,v] \subseteq_{\mathsf{R}+\delta} [w,u] \subseteq [w,z]$$

where we have applied Lemma 1.13 for the first comparison. $\square$

In order to exploit the above result, we recall yet another lemma of Bowditch:

Lemma 4.12 ([Bow06b] Proposition 6.3)**.** *Let* $\mathsf{r} > 0$ *and suppose* $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \mathcal{C}(S)$ *are multicurves which satisfy*

$$i(\alpha_1,\alpha_4)\, i(\alpha_2,\alpha_3) \leq \mathsf{r}\; i(\alpha_1,\alpha_2)\, i(\alpha_3,\alpha_4).$$

Then

$$d_S([\alpha_1,\alpha_2],[\alpha_3,\alpha_4]) \leq \mathsf{R}$$

where $\mathsf{R} \geq 0$ *depends only on* r *and* $\xi(S)$. $\square$

Proof of Proposition 4.9. Let $\mathbf{t}_\beta$ be a balance vector for β with respect to $\boldsymbol{\alpha}$. To simplify notation, assume $t_j t_k i(\alpha_j, \alpha_k)$ is maximised when $\{j,k\} = \{1,2\}$. Let γ_{12} and γ_0 be nearest point projections of β to $[\alpha_1, \alpha_2]$ and $\mathrm{Hull}(\boldsymbol{\alpha})$ respectively. This implies

$$t_2 t_j i(\alpha_2, \alpha_j) \leq t_1 t_2 i(\alpha_1, \alpha_2)$$

for any $j = 1, \ldots, n$. Since β is assumed to intersect all the α_i, we have $t_i = i(\alpha_i, \beta)^{-1}$ (after rescaling) and so

$$i(\alpha_1, \beta)\ i(\alpha_2, \alpha_j) \leq i(\alpha_1, \alpha_2)\ i(\alpha_j, \beta).$$

Invoking Lemma 4.12 gives

$$d_S([\alpha_1, \alpha_2], [\alpha_j, \beta]) \leq \mathsf{R}$$

and so by Lemma 4.11 we deduce

$$d_S(\gamma_{12}, \gamma_0) \leq \mathsf{R}',$$

where R' depends only on $\xi(S)$.

Now suppose γ is a curve in $\mathrm{short}(\mathbf{t}_\beta \cdot \boldsymbol{\alpha}, \mathsf{L})$. Using the same reasoning as for the proof of Proposition 4.6, we see that

$$\gamma \in \mathrm{short}(\mathbf{t}_\beta \cdot \boldsymbol{\alpha}, \mathsf{L}) \subseteq \mathrm{short}\left(\mathbf{t}'_\beta \cdot \boldsymbol{\alpha}', \frac{n\mathsf{L}}{\sqrt{2}}\right)$$

where $\boldsymbol{\alpha}' = (\alpha_1, \alpha_2)$ and $\mathbf{t}'_\beta = (t_1, t_2)$. By Lemma 4.10, we deduce that

$$d_S(\gamma, \gamma_{12}) \leq \mathsf{k}'_2,$$

for some k'_2 depending only on n, L and $\xi(S)$. This together with the preceding inequality implies

$$d_S(\gamma, \gamma_0) \leq \mathsf{R}' + \mathsf{k}'_2$$

which concludes the proof of the proposition. $\square$

Chapter 5

Covering maps I: A quasi-isometric embedding

In [RS09], Rafi–Schleimer proved that a finite index covering map between surfaces induces a quasi-isometric embedding between their respective curve complexes. We will present a new proof of their theorem using methods from hyperbolic 3–manifold theory. The material in this chapter also appears in [Tan12].

5.1 The lifting map

To simplify matters when dealing with covering maps of surfaces, we will use boundary components in place of marked points. In this situation, we define the complexity a surface S to be $\xi(S) = 3\,\mathrm{genus}(S) + |\partial S| - 3$. We will also use $\mathrm{int}(S)$ to denote the interior of S. These modifications make no difference on the level of the curve complex.

Suppose $P : \Sigma \to S$ is surface covering map of degree $\deg P < \infty$. The preimage $P^{-1}(a)$ of a curve $a \in \mathcal{C}(S)$ is a disjoint union of simple closed curves on Σ. This induces a one-to-many *lifting map* $\Pi : \mathcal{C}(S) \to \mathcal{C}(\Sigma)$ where $\Pi(a)$ is defined to be the set of homotopy classes of the components of $P^{-1}(a)$. One can check that all lifts of a to Σ are essential and non-peripheral, thus $\Pi(a)$ spans a non-empty simplex in $\mathcal{C}(\Sigma)$.

Theorem 5.1 ([RS09])**.** *Let $P : \Sigma \to S$ be a finite degree covering map. Then the map $\Pi : \mathcal{C}(S) \to \mathcal{C}(\Sigma)$ defined above is a Λ–quasi-isometric embedding, where Λ depends only on $\xi(\Sigma)$ and* $\deg P$.

This was first proved by Rafi–Schleimer in [RS09] using Teichmüller theory and subsurface projections. We will give a new proof of this theorem using an estimate for distance in the curve complex via a suitable hyperbolic 3–manifold homeomorphic to $S \times \mathbb{R}$ with a modified metric. This allows us to naturally compare distances by taking a covering map between the respective 3–manifolds. The estimate (Theorem 5.7) arises from work towards the Ending Lamination Theorem [BCM12] and is made explicit in [Bowa]. More details will be given in Section 5.2.

Remark 5.2. A consequence of Theorem 5.1 and of Gromov hyperbolicity of $\mathcal{C}(\Sigma)$ is that the image $\Pi(\mathcal{C}(S))$ is quasiconvex.

5.1.1 One direction of the inequality

Lemma 5.3. *Let a and b be curves in $\mathcal{C}(S)$ and suppose $\alpha \in \Pi(a)$ and $\beta \in \Pi(b)$. Then*

$$d_\Sigma(\alpha, \beta) \leq d_S(a, b).$$

Proof. The preimages of disjoint curves on S under P must themselves be disjoint. Therefore, by choosing a preferred vertex in each $\Pi(a)$, we can send any edge-path connecting a and b in $\mathcal{C}(S)$ to an edge-path of equal length connecting α and β in $\mathcal{C}(\Sigma)$. (We can extend this to higher dimensional simplices if so desired.) □

5.2 Hyperbolic 3–Manifolds

5.2.1 Margulis tubes and cusps

Let X be a hyperbolic 3–manifold. The *injectivity radius* at a point $x \in X$ is equal to half the infimum of the lengths of all non-trivial loops in X passing through x. The infimum is realised by a locally geodesic arc in X with both endpoints at x. The *ϵ–thin part* (or just thin part) of X is the set of all points whose injectivity radius is less than ϵ. The *thick part* of X is the complement of the thin part.

The well-known Margulis lemma explicitly describes the geometry of the components of the thin part of a hyperbolic 3–manifold. We refer to [Mar07] and [Apa00] for more details. We first state the classification before briefly explaining each case.

Theorem 5.4 (Thin-thick decomposition)**.** *There is a universal constant $\epsilon_3 > 0$, called the Margulis constant, so that for all $\epsilon \leq \epsilon_3$, the components of the ϵ–thin part of any hyperbolic 3–manifold X are of the following three types:*

1. *Margulis tubes,*

2. *rank–one cusps, and*

3. rank–two cusps.

A *Margulis tube* is an r–neighbourhood of a simple closed geodesic in X, for some $\mathsf{r} \geq 0$, and is homeomorphic to a solid torus. A *Margulis cusp* C is the quotient of a horoball $H \subset \mathbb{H}^3$ by a discrete parabolic subgroup $\Gamma < \mathrm{Isom}^+(\mathbb{H}^3)$ isomorphic to $\mathbb{Z}$ in the rank–one case; or $\mathbb{Z} \oplus \mathbb{Z}$ in the rank–two case. Note that the fixed point of Γ on $\partial_\infty \mathbb{H}^3$ must coincide with the centre of H. Up to conjugation, we may assume ∂H is a horizontal plane in $\mathbb{H}^3$. The group Γ then acts by Euclidean translations on the $\mathbb{R}^2$–factor of $\mathbb{H}^3 = \mathbb{R}^2 \times \mathbb{R}_+$ and trivially on the $\mathbb{R}_+$–factor. Therefore, rank–one cusps and rank–two cusps are respectively homeomorphic to $A \times \mathbb{R}_+$ or $\mathbb{T}^2 \times \mathbb{R}_+$, where A is an infinite annulus and $\mathbb{T}^2$ is a torus. Moreover, the foliation of H by horizontal horospheres descends to a foliation of C by either Euclidean annuli or tori respectively. The boundary ∂C of a cusp is also a Euclidean annulus or torus.

Let us fix a value of $0 < \epsilon \leq \epsilon_3$ for the rest of this chapter. The tubes and cusps arising in the ϵ–thin part of X will be referred to as ϵ–tubes and ϵ–cusps respectively. We will be interested in understanding the geometry of the boundary of an ϵ–cusp in terms of ϵ. A *systole* on ∂C is a non-trivial loop of shortest length. We write $\mathrm{sys}(\partial C)$ to denote the length of any systole on ∂C.

Lemma 5.5. *Let C be an ϵ–cusp in X. Then*

$$\mathrm{sys}(\partial C) = 2 \sinh \epsilon.$$

Proof. Let σ be a systole on ∂C and suppose x is a point on σ. The cusp C lifts to a horoball H in the universal cover $\mathbb{H}^3$. We may assume that ∂H is a horizontal plane at height 1 in the upper half space model of $\mathbb{H}^3$. The loop σ lifts to a straight line segment $\tilde{\sigma}$ on ∂H whose endpoints, $\mathbf{x}$ and $\mathbf{y}$ say, are lifts of x. In particular, we have

$$\|\mathbf{x} - \mathbf{y}\| = \mathrm{length}(\tilde{\sigma}) = \mathrm{sys}(\partial C).$$

By Theorem 2.8, we deduce

$$\sinh\left(\frac{1}{2}\, d_{\mathbb{H}^3}\left(\mathbf{x}, \mathbf{y}\right)\right) = \frac{\|\mathbf{x} - \mathbf{y}\|}{2}$$

and so $\mathbf{x}$ and $\mathbf{y}$ minimise the hyperbolic distance among all pairs of distinct lifts of x in $\mathbb{H}^3$. Therefore the geodesic arc connecting $\mathbf{x}$ to $\mathbf{y}$ in $\mathbb{H}^3$ descends to a shortest non-trivial loop in X which passes through x. Finally, the result follows by noting that the injectivity radius at $x \in \partial C$ is equal to ϵ. $\square$

5.2.2 Electrified distance

Given a hyperbolic 3–manifold X, let $\Psi_\epsilon^{cusp}(X)$ denote its *non-cuspidal part*, that is, X with its ϵ–cusps removed. We will write $\Psi(X)$ for brevity. Define the *electrified length* of a path in $\Psi(X)$ to be its total length occurring outside the Margulis tubes of X. More formally, we take the one-dimensional Hausdorff measure of its intersection with the thick part of X. This induces a reduced pseudometric ρ_X on $\Psi(X)$ obtained by taking the infimum of the electrified lengths of all paths connecting two given points. One can show that the infimum is attained, for example, by taking a path which connects Margulis tubes by shortest geodesic segments.

The distance ρ_X shall be referred to as the *electrified distance* on $\Psi(X)$ with respect to its Margulis tubes or the reduced pseudometric obtained by electrifying the Margulis tubes.

5.2.3 A distance estimate

We are now ready to introduce a suitable hyperbolic 3–manifold X on which the electrified distance provides an estimate for the distance in the curve complex.

Proposition 5.6. *Fix a constant $\mu > 0$. Let S be a surface with $\xi(S) \geq 2$ and fix distinct curves a and b in $\mathcal{C}(S)$. Then there exists a hyperbolic 3–manifold $(X, d) \cong \mathrm{int}(S) \times \mathbb{R}$ with a preferred homotopy equivalence f to S such that the unique geodesic representatives of the two curves in X, denoted a^* and b^*, have d–length at most μ.*

The proposition is a consequence of Bers' Simultaneous Uniformisation Theorem together with a theorem of Sullivan relating the conformal structure at infinity of a hyperbolic 3-manifold with the metric on the boundary of its convex core. In fact, such a 3–manifold can be chosen so that there are no accidental cusps – that is, all cusps are in bijective correspondence with the boundary components of S. An outline proof is given in [Bow08] (Proposition 7.3); for further background see [EM06] and [Mar07]

A preferred homotopy equivalence between S and X is required to specify a particular identification of curves in $\mathcal{C}(S)$ with simple closed geodesics in X.

Theorem 5.7 ([Bowa])**.** *Suppose $(X, d) \cong \mathrm{int}(S) \times \mathbb{R}$ is a hyperbolic 3–manifold with a preferred homotopy equivalence to S. Let a and b be curves in $\mathcal{C}(S)$ whose corresponding geodesic representatives a^* and b^* in X have d–length at most $\mu \geq 0$. Then*

$$d_S(a, b) \asymp_{\Lambda'} \rho_X(a^*, b^*)$$

where the constant Λ' depends only on $\xi(S)$, μ and ϵ. □

This estimate follows from the construction of geometric models for hyperbolic 3–manifolds used in the proof of the Ending Lamination Theorem. We refer the reader to [Bowa] (Theorem 5.4) and [BCM12] for an in-depth discussion.

Remark 5.8. For Theorem 5.7 to make sense we need to ensure that a^* and b^* are indeed contained in $\Psi(X)$. This can be done using a pleated surfaces argument, such as in [Bow07], to show that closed geodesics of bounded length in X avoid cusps provided ϵ is sufficiently small.

5.3 The other direction of the inequality

5.3.1 Closed surfaces

We first prove Theorem 5.1 for closed surfaces. The required modifications for the general case shall be dealt with in Section 5.3.2.

Fix a length bound μ and a constant $\epsilon \leq \epsilon_3$. Let $P : \Sigma \to S$ be a covering map. Fix curves a, b in $\mathcal{C}(S)$ and choose $\alpha \in \Pi(a)$ and $\beta \in \Pi(b)$. From Lemma 5.3, we have $d_\Sigma(\alpha, \beta) \leq d_S(a, b)$ so it remains to prove the reverse inequality.

Let $(X, d) \cong \mathrm{int}(S) \times \mathbb{R}$ be a hyperbolic 3–manifold with a homotopy equivalence f to S as described in Section 5.2.3. We also assume that a^* and b^* have length at most $\frac{\mu}{\deg P}$ in X. There exists a covering map $Q : \Xi \to X$, where $\Xi \cong \mathrm{int}(\Sigma) \times \mathbb{R}$, and a homotopy equivalence $\tilde{f} : \Xi \to \Sigma$ such that $P \circ \tilde{f} = f \circ Q$. We lift the hyperbolic metric on X to Ξ via Q. Note that $Q(\alpha^*) = a^*$ and $Q(\beta^*) = b^*$, hence α^* and β^* have length bounded above by μ.

We may assume that X contains no accidental cusps, hence $\Psi(X) = X$ and $\Psi(\Xi) = \Xi$. Let ρ_X and ρ_Ξ be the respective pseudometrics on X and Ξ obtained by electrifying their ϵ–tubes.

Lemma 5.9. *The map Q is 1–Lipschitz with respect to ρ_Ξ and ρ_X.*

Proof. Let x be a point in the ϵ–thin part of Ξ. Then there exists a non-trivial loop passing through x of length at most 2ϵ. Since Q is a covering map, this loops descends to non-trivial loop of the same length in X passing through $Q(x)$. Therefore the ϵ–thin part of Ξ is sent inside the ϵ–thin part of X. It follows that the electrified lengths of paths cannot increase under Q. □

This result, together with Theorem 5.7, proves Theorem 5.1 for closed surfaces.

5.3.2 Surfaces with boundary

We now assume S has non-empty boundary. Recall that $\Psi_\nu^{cusp}(X)$ denotes X with its ν–cusps removed. For $\epsilon_3 \geq \nu > \nu' > 0$, observe that $\Psi_\nu^{cusp}(X) \subset \Psi_{\nu'}^{cusp}(X)$. Let $r : \Psi_{\nu'}^{cusp}(X) \to \Psi_\nu^{cusp}(X)$ be the nearest point projection map to $\Psi_\nu^{cusp}(X)$ with respect to the hyperbolic metric on $\Psi_{\nu'}^{cusp}(X)$. We will call r the *natural retraction* from $\Psi_{\nu'}^{cusp}(X)$ to $\Psi_\nu^{cusp}(X)$

Lemma 5.10. *Let X be a hyperbolic 3–manifold. Choose small constants $\nu > \nu' > 0$ and let ρ and ρ' be the pseudometrics on $\Psi_\nu^{cusp}(X)$ and $\Psi_{\nu'}^{cusp}(X)$ obtained by electrifying along their respective ϵ–tubes. Then the natural retraction*

$$r : (\Psi_{\nu'}^{cusp}(X), \rho') \to (\Psi_\nu^{cusp}(X), \rho)$$

is h–Lipschitz, where $\mathsf{h} = \frac{\sinh \nu}{\sinh \nu'}$.

Proof. Begin with a geodesic arc γ in $\Psi_{\nu'}^{cusp}(X)$. We will show that the length of γ can only increase by a bounded multiplicative factor under the retraction r.

If γ is contained in $\Psi_\nu^{cusp}(X)$ then we are done, so we shall assume γ meets some ν–cusp C of X. This cusp contains a ν'–cusp C' which does not meet γ. The cusps C and C' lift to nested horoballs H and H' in the universal cover $\mathbb{H}^3$. We can arrange for the horospheres ∂H and $\partial H'$ to be horizontal planes at heights 1 and $\mathsf{h} > 1$ respectively in the upper half-space model. Using Lemma 5.5 and basic hyperbolic geometry, we deduce

$$\mathsf{h} = \frac{\mathrm{sys}(\partial C)}{\mathrm{sys}(\partial C')} = \frac{\sinh \nu}{\sinh \nu'}.$$

Now define π_H and $\pi_{H'}$ to be the nearest point projections from $\mathbb{H}^3$ to ∂H and $\partial H'$ respectively. Taking a lift $\tilde{\gamma}$ of γ, we see that

$$\mathrm{length}(\tilde{\gamma} \cap H) \geq \mathrm{length}(\pi_{H'}(\tilde{\gamma} \cap H)) = \frac{1}{\mathsf{h}} \mathrm{length}(\pi_H(\tilde{\gamma} \cap H)).$$

Thus, by projecting each arc of γ inside a ν–cusp to the boundary of that cusp, we create a new path whose length is at most $\mathsf{h} \times \mathrm{length}(\gamma)$. $\square$

Let $Q : \Xi \to X$ be the covering map as described in Section 5.3.1. Observe that

$$\Psi_\epsilon^{cusp}(X) \subseteq Q(\Psi_\epsilon^{cusp}(\Xi)) \subseteq \Psi_{\epsilon'}^{cusp}(X)$$

where $\epsilon' = \frac{\epsilon}{\deg P}$. As before, let ρ_Ξ and ρ_X be the pseudometrics on $\Psi_\epsilon^{cusp}(\Xi)$ and

$\Psi_{\epsilon}^{cusp}(X)$ obtained by electrifying along their respective ϵ–tubes. Combining Lemma 5.10 with the proof of Lemma 5.9 gives us the following:

Lemma 5.11. *Let $r : \Psi_{\epsilon'}^{cusp}(X) \to \Psi_{\epsilon}^{cusp}(X)$ be the natural retraction. Then the composition $r \circ Q : \Psi_{\epsilon}^{cusp}(\Xi) \to \Psi_{\epsilon}^{cusp}(X)$ is h–Lipschitz with respect to ρ_{Ξ} and ρ_X, where the constant h depends only on* $\deg P$. □

Finally, $r \circ Q(\alpha^*) = a^*$ and $r \circ Q(\beta^*) = b^*$ and so the rest of the argument follows as in Section 5.3.1. This completes the proof of Theorem 5.1 for surfaces with boundary. □

Chapter 6

Covering maps II: Nearest point projections and circumcentres

In this chapter, we define an operation on curves using a given surface covering map and intersection number conditions. We then show that it approximates a nearest point projection map to the image of the associated lifting map defined in Chapter 5. We also show that our operation also approximates circumcentres of orbits in the case of regular covers.

6.1 A projection map via intersection numbers

Let $P : \Sigma \to S$ be a finite degree covering map of surfaces and let $\Pi : \mathcal{C}(S) \to \mathcal{C}(\Sigma)$ be the induced lifting map given by $\Pi(b) := P^{-1}(b)$. Define a map $\pi : \mathcal{C}(\Sigma) \to \Pi(\mathcal{C}(S))$ as follows: given a curve $\alpha \in \mathcal{C}(\Sigma)$, let $b \in \mathcal{C}(S)$ which has minimal intersection number with $P(\alpha)$ on S and set $\pi(\alpha) = \Pi(b)$. We will prove the following in Section 6.3

Theorem 6.1. *Let $P : \Sigma \to S$ be a finite degree covering map of surfaces and let Π and π be as above. Given a curve $\alpha \in \mathcal{C}(\Sigma)$, let γ be a nearest point projection of α to $\Pi(\mathcal{C}(S))$ in $\mathcal{C}(\Sigma)$. Then $\pi(\alpha) \approx_{\mathsf{k}_3} \gamma$, where k_3 is a constant depending only on $\deg P$ and $\xi(\Sigma)$.*

Consequently, the operation $\alpha \mapsto \pi(\alpha)$ is coarsely well-defined. One can check that the minimal value of $i(P(\alpha), \cdot)$ over all closed curves on S is attained by a simple closed curve. In Section 6.4, we will also prove the following.

Proposition 6.2. *Suppose further that P is regular and let G be its group of deck transformations. Let γ' be a circumcentre of the G–orbit of a curve α in $\mathcal{C}(\Sigma)$. Then*

$\pi(\alpha) \approx_{\mathsf{k}_4} \gamma'$, *where* k_4 *is some constant depending only on* $\deg P$ *and* $\xi(\Sigma)$.

6.2 Deck transformations

The *deck transformation group* $\mathrm{Deck}(P)$ of a covering map $P : \Sigma \to S$ is the group of all homeomorphisms $f \in \mathrm{Homeo}(\Sigma)$ satisfying $P \circ f = P$.

Lemma 6.3. *Let* $P : \Sigma \to S$ *be a finite degree covering map between surfaces of negative Euler characteristic. Then the natural quotient map from* $\mathrm{Deck}(P)$ *to* $\mathrm{Mod}(\Sigma) = \mathrm{Homeo}(\Sigma)/\mathrm{Homeo}_0(\Sigma)$ *is injective.*

Proof. We will only give a sketch proof. Endow $\mathrm{int}(S)$ with a hyperbolic metric and pull it back to $\mathrm{int}(\Sigma)$ via P. The group $\mathrm{Deck}(P)$ then acts on $\mathrm{int}(\Sigma)$ by isometries. The result follows since any isometry of a hyperbolic surface isotopic to the identity must in fact coincide with the identity. □

As a consequence, we may identify $\mathrm{Deck}(P)$ with its image in $\mathrm{Mod}(\Sigma)$. It is also worth mentioning that the above statement does not hold for covers of the torus or annulus.

6.3 Nearest point projections

6.3.1 Regular covers

We shall first deal with the case where $P : \Sigma \to S$ is regular. Let $G \leq \mathrm{Mod}(\Sigma)$ be the group of deck transformations of P. Given a curve $\alpha \in \mathcal{C}(\Sigma)$, observe that the set of lifts of $P(\alpha)$ to Σ via P is exactly $G\alpha$. Let $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ be an n–tuple of curves whose entries are the lifts of $P(\alpha)$ in any order. Note that $n \geq 1$ is some divisor of $\deg P$. Let $\mathbf{1}$ denote the vector of length n with all entries equal to 1.

Lemma 6.4. *Let* α *and* $\boldsymbol{\alpha}$ *be as above. Then* $\pi(\alpha) \in \mathrm{short}(\mathbf{1} \cdot \boldsymbol{\alpha}, \mathsf{L}_0|G|)$ *where* L_0 *is a constant depending only on* $\xi(\Sigma)$.

Proof. Let b be a closed curve on S. Each point of $b \cap P(\alpha)$ on S lifts to exactly $|G| = \deg P$ points of $P^{-1}(b) \cap G\alpha$ on Σ via P, hence

$$i(P^{-1}(b), \boldsymbol{\alpha}) = |G| \; i(b, P(\alpha)).$$

By Lemma 4.5, there exists a curve $\gamma \in \mathcal{C}(\Sigma)$ such that

$$i(\gamma, \alpha) \leq \mathsf{L}_0 \, \|\mathbf{1}\|_{\boldsymbol{\alpha}}$$

for some constant $\mathsf{L}_0 = \mathsf{L}_0(\xi(\Sigma))$. Now assume b has minimal intersection with $P(\alpha)$ out of all curves on S. It follows that

$$i(P^{-1}(b), \boldsymbol{\alpha}) = |G|\ i(b, P(\alpha)) \leq |G|\ i(P(\gamma), P(\alpha)) = i(G\gamma, \boldsymbol{\alpha}) \leq |G|\ i(\gamma, \boldsymbol{\alpha}).$$

Finally, by combining the preceding inequalities, we see that

$$i(\pi(\alpha), \boldsymbol{\alpha}) = i(P^{-1}(b), \boldsymbol{\alpha}) \leq |G|\ i(\gamma, \boldsymbol{\alpha}) \leq |G|\ \mathsf{L}_0 \left\|\mathbf{1}\right\|_{\boldsymbol{\alpha}}.$$

Thus $\pi(\alpha) \in \mathrm{short}(\mathbf{1} \cdot \boldsymbol{\alpha}, \mathsf{L})$ for $\mathsf{L} = \mathsf{L}_0|G|$. □

Lemma 6.5. *Let γ be any curve in $\Pi(\mathcal{C}(S))$ and let β be any of its nearest point projections to* $\mathrm{Hull}(\boldsymbol{\alpha})$. *Then $d_\Sigma(\pi(\alpha), \beta) \leq \mathsf{k}_4$, where k_4 depends only on* $\deg P$ *and* $\xi(\Sigma)$.

Proof. We may replace γ with the multicurve $G\gamma$ since their nearest point projections to $\mathrm{Hull}(\boldsymbol{\alpha})$ are a uniformly bounded distance apart. Since G acts transitively on $G\alpha$ and leaves $G\gamma$ invariant, it follows that $i(G\gamma, \alpha_i) = i(G\gamma, \alpha_j)$ for all i, j. Thus, $\mathbf{1}$ serves as a balance vector for $G\gamma$ with respect to $\boldsymbol{\alpha}$. By Proposition 4.9, we deduce that

$$\beta \approx_{\mathsf{k}_2} \mathrm{short}(\mathbf{1} \cdot \boldsymbol{\alpha}, \mathsf{L}),$$

where k_2 depends only on $\xi(\Sigma)$, n and $\mathsf{L} \geq \mathsf{L}_0$. Applying the previous lemma completes the proof. □

Proof of Theorem 6.1 for regular covers. Let α and $\boldsymbol{\alpha}$ be as above. Let γ be any curve in $\Pi(\mathcal{C}(S))$. Since $\mathrm{Hull}(\boldsymbol{\alpha})$ is quasiconvex, Lemmas 6.5 and 1.13 imply that any geodesic connecting α to γ in $\mathcal{C}(\Sigma)$ must pass within a distance r of $\pi(\alpha)$, where r depends only on $\deg P$ and $\xi(\Sigma)$. Therefore $\pi(\alpha)$ is an r–entry point of α to $\Pi(\mathcal{C}(S))$. Since $\Pi(\mathcal{C}(S))$ is also quasiconvex, Lemma 1.12 implies $\pi(\alpha)$ is a uniformly bounded distance away from any nearest point projection of α to $\Pi(\mathcal{C}(S))$. □

6.3.2 The general case

The main obstacle in proving the main theorem for a non-regular covering map $P : \Sigma \to S$ is the following: given a simple closed curve $\alpha \in \mathcal{C}(\Sigma)$ there may be some lifts of $P(\alpha)$ to Σ which are not simple. To address this issue, we pass to a suitable finite cover of Σ using a standard group theoretic argument.

Lemma 6.6. *Let $P : \Sigma \to S$ be a covering map of finite degree. Then there exists a cover $Q : \hat{\Sigma} \to \Sigma$ such that $F := P \circ Q$ is regular and* $\deg F \leq (\deg P)!$.

Proof. Let H be the finite index subgroup of $\Gamma = \pi_1(S)$ corresponding to the covering map P and let H_0 be the intersection of all Γ–conjugates of H. It is straightforward to check that H_0 is exactly the kernel of the action of Γ on the set of left cosets of H by left-multiplication. The desired result then follows. $\square$

The covering map F defined above is universal in the sense that any regular cover of S which factors through P must also factor through F.

Lemma 6.7. *Let $P : \Sigma \to S$ and $F : \hat{\Sigma} \to S$ be as above. If α is a simple closed curve on S then all lifts of $P(\alpha)$ to $\hat{\Sigma}$ via F are simple.*

Proof. Any lift of α to $\hat{\Sigma}$ via Q is also a simple lift of $P(\alpha)$ via F. Since F is regular, it follows that all other lifts of $P(\alpha)$ to $\hat{\Sigma}$ are simple. $\square$

Before continuing with the proof, we first show that that nearest point projections to quasiconvex sets are well-behaved under quasi-isometric embeddings. We remind the reader that we allow f to be a one-to-many function.

Lemma 6.8. *Let $f : \mathcal{X} \to \mathcal{X}'$ be a $\mathsf{\Lambda}$–quasi-isometric embedding of geodesic spaces, where $\mathcal{X}'$ is δ'–hyperbolic. Let C be a Q–quasiconvex subset of $\mathcal{X}$ and let $C' = f(C)$. Given a point $x \in \mathcal{X}$, let x' be a point in $f(x)$. Let p and q' be nearest point projections of x to C and x' to C' respectively. Let $q \in \mathcal{X}$ be a point so that $q' \in f(q)$. Then $p \approx_{\mathsf{K}} q$, where K depends only on δ', $\mathsf{\Lambda}$ and Q.*

Proof. First, note that $\mathcal{X}$ is δ–hyperbolic and C' is Q'–quasiconvex in $\mathcal{X}'$ for some constants $\delta = \delta(\mathsf{\Lambda}, \delta')$ and $\mathsf{Q}' = \mathsf{Q}'(\mathsf{Q}, \mathsf{\Lambda}, \delta)$. Let $c \in \mathcal{X}$ be a k–centre for x, p and q, where $\mathsf{k} = \delta$. Any point $c' \in f(c)$ is then a k'–centre for x', p' and q', where $\mathsf{k}' = \mathsf{k}'(\mathsf{k}, \mathsf{\Lambda})$ and $p' \in f(p)$. One can check that $xp \approx_{2\mathsf{k}} xc + cp$. By quasiconvexity of C, there is some point $y \in C$ satisfying $cy \leq \mathsf{k} + \mathsf{Q}$. Since p is a nearest point projection of x to C, we obtain

$$xc + cp - 2\mathsf{k} \leq xp \leq xy \leq xc + cy \leq xc + \mathsf{k} + \mathsf{Q}$$

which implies $cp \leq \mathsf{Q} + 3\mathsf{k}$. Similarly, we can deduce $c'q' \leq \mathsf{Q}' + 3\mathsf{k}'$. Since f is a $\mathsf{\Lambda}$–quasi-isometric embedding, it follows that $cq \leq \mathsf{\Lambda} \times c'q' + \mathsf{\Lambda}$ and hence

$$pq \leq pc + cq \leq \mathsf{K},$$

where $\mathsf{K} = \mathsf{Q} + 3\mathsf{k} + \mathsf{\Lambda}(\mathsf{Q}' + 3\mathsf{k}') + \mathsf{\Lambda}$. $\square$

Let $\Phi : \mathcal{C}(S) \to \mathcal{C}(\hat{\Sigma})$ and $\Psi : \mathcal{C}(\Sigma) \to \mathcal{C}(\hat{\Sigma})$ be the lifting maps induced by the covering maps F and Q respectively. Define $\phi : \mathcal{C}(\hat{\Sigma}) \to \Phi(\mathcal{C}(S))$ to be the projection map associated to F as described in Section 6.1 . We may assume that $\phi \circ \Psi = \Psi \circ \pi$.

Proof of Theorem 6.1. Given $\alpha \in \mathcal{C}(\Sigma)$, let $\hat{\alpha}$ be any of its lifts to $\hat{\Sigma}$ via Q. Note that $\phi(\hat{\alpha}) = \Psi(\pi(\alpha))$. Let $\hat{\gamma}$ be a nearest point projection of $\hat{\alpha}$ to $\Phi(\mathcal{C}(S))$ in $\mathcal{C}(\hat{\Sigma})$ and let $\gamma = Q(\hat{\gamma}) \in \Pi(\mathcal{C}(S))$. Since F is regular, we can apply Theorem 6.1 to deduce that

$$d_{\hat{\Sigma}}(\phi(\hat{\alpha}), \hat{\gamma}) \leq \hat{\mathsf{k}}_3,$$

where $\hat{\mathsf{k}}_3$ depends only on $\deg F$ and $\xi(\hat{\Sigma})$ which can in turn be bounded in terms of $\deg P$ and $\xi(\Sigma)$. By Theorem 5.1, Ψ is a Λ–quasi-isometric embedding, where $\Lambda = \Lambda(\deg F, \xi(\hat{\Sigma}))$, and so

$$d_{\Sigma}(\pi(\alpha), \gamma) \leq \Lambda\hat{\mathsf{k}}_3 + \Lambda.$$

By the previous lemma, γ is a uniformly bounded distance away from any nearest point projection of α to $\Pi(\mathcal{C}(S))$ in $\mathcal{C}(\Sigma)$ and we are done. □

6.4 Circumcentres and finite group actions

6.4.1 Circumcentres of orbits

We now show that $\pi(\alpha)$ also approximates a circumcentre of $G\alpha$ in $\mathcal{C}(\Sigma)$, where G is the deck transformation group of a regular cover $P : \Sigma \to S$. First, we give the following characterisation of circumcentres of orbits under finite group actions on δ–hyperbolic spaces:

Lemma 6.9. *Assume G is a finite group acting by isometries on a δ–hyperbolic space $\mathcal{X}$. Fix a point $x_0 \in \mathcal{X}$ and let c be a circumcentre for Gx_0. Given a point $z \in \mathcal{X}$, let p be any of its nearest point projection to* $\mathrm{Hull}(Gx_0)$*. Then*

$$pc \leq \mathrm{rad}(Gz) + 7\delta$$

and hence

$$zc \leq \mathrm{rad}(Gz) + d(z, \mathrm{Hull}(Gx_0)) + 7\delta.$$

Proof. We first claim that p lies within a distance δ of a geodesic segment $[u, v]$, where $u, v \in Gx_0$ are points such that $uv \geq \mathrm{diam}(Gx_0) - 2\delta$. Suppose p lies on a geodesic segment $[x, y]$ for some x and y in Gx_0. There exist some x' and y' in Gx_0 so that $xx' = yy' = \mathrm{diam}(Gx_0)$. If $x' = y'$ then the claim follows from hyperbolicity. Now assume $x' \neq y'$. By Lemma 1.7, we have

$$2\,\mathrm{diam}(Gx_0) = xx' + yy' \geq \max\{xy + x'y', xy' + x'y\} \geq 2\,\mathrm{diam}(Gx_0) - 2\delta.$$

If $xy + x'y' \geq 2\,\mathrm{diam}(Gx_0) - 2\delta$ then $xy \geq \mathrm{diam}(Gx_0) - 2\delta$ which implies the claim. If not then $xy' \geq \mathrm{diam}(Gx_0) - 2\delta$. The claim then follows by considering a geodesic triangle with x, y and y' as its vertices.

Now suppose $q \in [u, v]$ is a point so that $pq \leq \delta$. Then

$$d(z, [u, v]) \leq zq \leq zp + pq \leq d(z, \mathrm{Hull}(Gx_0)) + \delta \leq d(z, [u, v]) + \delta.$$

Invoking Lemma 1.9 we obtain $p \approx_\delta q \approx_{3\delta} o$, where $o \in [u, v]$ is the internal point opposite z. Setting $\mathsf{D} := \mathrm{rad}(Gz) \geq \frac{1}{2}\,\mathrm{diam}(Gz)$, observe that

$$d(z, x_0) = d(gz, gx_0) \approx_{2\mathsf{D}} d(z, gx_0)$$

for all $g \in G$. Therefore $zu \approx_{2\mathsf{D}} zv$ which implies $uo \approx_{2\mathsf{D}} ov$. It follows that $o \approx_{\mathsf{D}} m$, where m is the midpoint of $[u, v]$. Finally, applying Lemma 1.17 gives $p \approx_{4\delta} o \approx_{\mathsf{D}} m \approx_{3\delta} c$ and we are done. □

Proof of Proposition 6.2. Let γ' be a circumcentre for $G\alpha$. Combining Lemma 6.4 and Proposition 4.6, we deduce

$$d_\Sigma(\pi(\alpha), \mathrm{Hull}(G\alpha)) \leq \mathsf{k}_5',$$

where k_5' depends only on $\xi(\Sigma)$ and $\deg P$. Since $\pi(\alpha)$ is a G–invariant multicurve, the radius of its G–orbit is at most 1. Applying the previous lemma gives $d_\Sigma(\pi(\alpha), \gamma') \leq \mathsf{k}_5' + 7\delta + 1$ and we are done. □

6.4.2 Almost fixed point sets

Let G be a finite group acting by isometries on a δ–hyperbolic space $\mathcal{X}$. Given $\mathsf{R} \geq 0$, let

$$\mathrm{Fix}_{\mathcal{X}}(G, \mathsf{R}) := \{x \in \mathcal{X} \mid \mathrm{diam}(Gx) \leq \mathsf{R}\}$$

be the set of R–*almost fixed points* of G in $\mathcal{X}$.

Lemma 6.10. *The set* $\mathrm{Fix}_{\mathcal{X}}(G, 2\delta)$ *is non-empty. Moreover, for all* $\mathsf{R} \geq \delta$*,*

$$\mathrm{Fix}_{\mathcal{X}}(G, 2\,\mathsf{R}) \approx_{\mathsf{R}+\delta} \mathrm{Fix}_{\mathcal{X}}(G, 2\delta).$$

Proof. Let x be any point in $\mathcal{X}$ and let c be a circumcentre for Gx. Since Gx is G–invariant, all G–translates of c are also circumcentres for Gx. It follows from Lemma 1.16 that c is contained in $\mathrm{Fix}_{\mathcal{X}}(G, 2\delta)$.

Assume further that x is contained in $\mathrm{Fix}_{\mathcal{X}}(G, 2\,\mathsf{R})$, where $\mathsf{R} \geq \delta$. By Lemma 1.17, we have

$$xc \leq \mathrm{rad}(Gx) \leq \frac{1}{2}\,\mathrm{diam}(Gx) + \delta \leq \mathsf{R} + \delta$$

and hence $\mathrm{Fix}_{\mathcal{X}}(G, 2\,\mathsf{R}) \subseteq_{\mathsf{R}+\delta} \mathrm{Fix}_{\mathcal{X}}(G, 2\delta)$. The reverse inclusion is immediate. □

Thus, to understand the geometry of the set of 2R–fixed points of $\mathcal{X}$, for $\mathsf{R} \geq \delta$, it suffices to study the behaviour of $\mathrm{Fix}_{\mathcal{X}}(G, 2\delta)$. One can also check that $\mathrm{Fix}_{\mathcal{X}}(G, 2\mathsf{R})$ is quasi-convex for $\mathsf{R} \geq \delta$.

Lemma 6.11. *Let* $\mathsf{R} \geq \delta$*. Given a point* $x \in \mathcal{X}$*, let* c *be a circumcentre for its* G*–orbit. Let* p *be a nearest point projection of* x *to* $\mathrm{Fix}_{\mathcal{X}}(G, 2\mathsf{R})$*. Then* $c \approx_{\mathsf{k}_6} p$*, where* $\mathsf{k}_6 = 2\delta + 4\mathsf{R}$.

Proof. For all $g \in G$, we have

$$d(gx, p) \leq d(gx, gp) + d(gp, p) \leq d(x, p) + 2\mathsf{R} \leq d(x, c) + 2\mathsf{R} \leq \mathrm{rad}(Gx) + 2\mathsf{R}.$$

Applying Lemma 1.16 completes the proof. □

We will demonstrate below that when $\mathsf{R} < \delta$, Lemmas 6.10 and 6.11 need not hold; it is possible for $\mathrm{Fix}_{\mathcal{X}}(G, 2\mathsf{R})$ to be empty or lie very deeply inside $\mathrm{Fix}_{\mathcal{X}}(G, 2\delta)$.

Recall the standard cylindrical co-ordinate system on $\mathbb{R}^3$: a point specified by $(r, \theta, t) \in [0, \infty) \times \mathbb{R} \times \mathbb{R}$ in cylindrical co-ordinates represents the point $(r\cos\theta, r\sin\theta, t) \in \mathbb{R}^3$ under the standard Cartesian co-ordinate system.

Example 6.12 (Rocketship). A *rocketship* of length $l > 0$ with $n \geq 2$ fins, denoted $\mathfrak{R} = \mathfrak{R}(n, l)$, is the union of the following three subsets of $\mathbb{R}^3$ defined using cylindrical co-ordinates:

- the *nose* $\mathfrak{N} = \{(t, \theta, t) \mid 0 \leq t \leq 1,\ \theta \in \mathbb{R}\}$, a right circular cone of height 1 and base radius 1;
- the *shaft* $\mathfrak{S} = \{(1, \theta, t) \mid 1 \leq t \leq l + 1,\ \theta \in \mathbb{R}\}$, a right circular cylinder of height l and base radius 1; and

- the *fins* $\mathfrak{F}_n = \{(1, \frac{2k\pi}{n}, t) \mid t \geq l+1,\, k \in \mathbb{Z}\}$, a disjoint union of n closed rays.

We endow $\mathfrak{R}$ with the path metric inherited from $\mathbb{R}^3$ equipped with the standard Euclidean metric. One can show that $\mathfrak{R}$ is quasi-isometric to a tree and therefore δ–hyperbolic for some $\delta > 0$; this can be done by collapsing the radial component of the nose and shaft. Moreover, one can check that $\delta \geq \frac{\pi}{2}$ for l sufficiently large.

Observe that $G = \mathbb{Z}/n\mathbb{Z}$ acts isometrically on $\mathfrak{R}$ by rotations about the t–axis through integral multiples of $\frac{2\pi}{n}$. Let x be any point on $\mathfrak{F}_n$. Then a circumcentre c for Gx is a point of the form $(1, \frac{(4k+1)\pi}{2n}, l+1)$, for some $k \in \mathbb{Z}$. For $\mathsf{R} \geq 0$ sufficiently small, $\mathrm{Fix}_{\mathfrak{R}}(G, 2\mathsf{R})$ is contained entirely within the nose $\mathfrak{N}$. Therefore, c must be a distance at least l away from any nearest point projection of x to $\mathrm{Fix}_{\mathfrak{R}}(G, 2\mathsf{R})$. Furthermore, $\mathrm{Fix}_{\mathfrak{R}}(G, 2\delta)$ contains both $\mathfrak{N}$ and $\mathfrak{S}$, and so its Hausdorff distance from $\mathrm{Fix}_{\mathfrak{R}}(G, 2\mathsf{R})$ is at least l.

Let us return our attention to the curve graph. When G is the deck transformation group of a regular cover $P : \Sigma \to S$, the vertices in $\mathrm{Fix}_{\mathcal{C}(\Sigma)}(G, 1)$ coincide exactly with those of $\Pi(\mathcal{C}(S)) \subseteq \mathcal{C}(\Sigma)$. Combining Theorem 6.1 and Proposition 6.2, we deduce:

Corollary 6.13. *Any circumcentre for the G–orbit of a curve $\alpha \in \mathcal{C}(S)$ is a uniformly bounded distance away from any nearest point projection of α to $\Pi(\mathcal{C}(S))$.*

Therefore Lemma 6.11 still holds for $\mathrm{Fix}_{\mathcal{C}(\Sigma)}(G, 1)$, albeit with weaker control over the constant k_6. As the example above shows, this cannot be proved using purely synthetic methods assuming only δ–hyperbolicity of $\mathcal{C}(\Sigma)$. In conclusion: "There are no rocketships in the curve complex."

Chapter 7

Distance bounds

In this chapter, we describe two methods for obtaining distance bounds in the curve complex in terms of intersection number – the first due to Bowditch [Bowb] and the second arising from an algorithm of Hempel [Gor07]. We introduce the notion of *punctured graphs* to help us bound the length of Hempel's paths.

7.1 Bowditch's bound

In this section we review Bowditch's method given in [Bowb] for bounding the distance in the curve complex using intersection numbers. We will assume $S = S_{g,m}$ is a surface satisfying $\xi(S) \geq 2$ in this chapter.

For $g \geq 0$ and $m \geq 0$, define

$$\kappa(g,m) := \begin{cases} 2g-3 & : m = 0; \\ m-3 & : g = 0,\ m \text{ odd}; \\ 2g+m-4 & : \text{otherwise.} \end{cases}$$

The following is a well-known result in the literature, with proofs appearing in [Bowb] and [Aou]. We give minor improvements on the constants by taking a bit more care in the argument.

Lemma 7.1. *Suppose α and β are a pair of filling curves in $\mathcal{C}(S)$. Then $i(\alpha,\beta) \geq \kappa(g,m)+2$.*

Proof. Realise α and β in minimal position. Since α and β fill S, all complementary components of $\alpha \cup \beta$ on S are topological discs with at most one marked point. Thus, there is a 2–cell decomposition of S with $\alpha \cup \beta$ as its 1–skeleton. Let $\#F$, $\#E$ and $\#V$ respectively denote the number of 2–cells, edges and vertices in this decomposition. Observing that $\#F \geq \max\{1, m\}$, $\#E = 2i(\alpha, \beta)$ and $\#V = i(\alpha, \beta)$, we deduce

$$2 - 2g = \chi(S) = \#F - \#E + \#V = \max\{1, m\} - i(\alpha, \beta)$$

and hence

$$i(\alpha, \beta) = 2g + \max\{1, m\} - 2.$$

The result now follows for all cases except for when $g = 0$ and m is odd. On $S_{0,m}$, with $m \geq 4$, all curves are separating and hence the intersection number between any two curves is even. If m is odd then so is $m - 2$, so we can therefore conclude $i(\alpha, \beta) \geq m - 1$. $\square$

In particular, if $\alpha, \beta \in \mathcal{C}(S)$ are curves such that $i(\alpha, \beta) \leq \kappa(g, m) + 1$ then $d_S(\alpha, \beta) \leq 2$.

Lemma 7.2 ([Bow06b])**.** *Suppose $\alpha, \beta \in \mathcal{C}(S)$ are curves such that $i(\alpha, \beta) \leq \frac{ab}{2}$ for some $a, b \in \mathbb{N}$. Then there exists $\gamma \in \mathcal{C}(S)$ so that $i(\alpha, \gamma) \leq a$ and $i(\beta, \gamma) \leq b$.* $\square$

Bowditch sets $b = \kappa(g, m) + 1$ and repeatedly applies the previous two lemmas to deduce:

Lemma 7.3 ([Bowb])**.** *Suppose α and β are curves in $\mathcal{C}(S)$ so that*

$$i(\alpha, \beta) \leq 2\left(\frac{\kappa(g, m)}{2}\right)^k$$

for some $k \in \mathbb{N}$. Then $d_S(\alpha, \beta) \leq 2k$. $\square$

We may rearrange the above inequality to obtain:

Corollary 7.4. *Let α and β be curves in $\mathcal{C}(S)$. Then*

$$d_S(\alpha, \beta) \leq \frac{2}{\log \kappa(g, m)} \cdot \log i(\alpha, \beta) + 2$$

whenever $i(\alpha, \beta) \neq 0$. $\square$

7.2 Hempel's paths

Let us first recall a greedy algorithm due to Hempel [Gor07]:

Hempel's algorithm: Let α, β be curves in $\mathcal{C}(S)$. Set $\beta_0 = \beta$. Build a path $\beta_0, \beta_1, \ldots, \beta_L = \alpha$ as follows: Suppose we have found β_{n-1}. If α and β_{n-1} are disjoint then set $\beta_n = \alpha$ and STOP. Otherwise, choose β_n to be a curve which has minimal intersection number with α among all curves disjoint from β_{n-1}.

Lemma 7.5. *For $1 \leq n \leq L-2$, the curve β_n either bounds a disc with two marked points, or is non-separating.*

The above lemma will follow in the course of proving Proposition 7.15.

Define

$$\lambda(g,m) := \begin{cases} \frac{m+2}{4} & : g \geq 1,\ m = 3,4; \\ \frac{m-2}{2} & : g \geq 0,\ m \geq 5. \end{cases}$$

The rest of the chapter will be devoted to proving the following:

Lemma 7.6. *Suppose $\alpha, \beta \in \mathcal{C}(S)$ fill S and let $\gamma \in \mathcal{C}(S)$ be a curve which has minimal intersection number with α among all curves disjoint from β. Then*

$$i(\alpha,\gamma) \leq \frac{1}{\lambda(g,m)}\ i(\alpha,\beta)$$

whenever β is a non-separating curve or a pants curve.

Lemma 7.7. *Suppose $i(\alpha,\beta) \neq 0$. Then a path constructed by Hempel's algorithm from β to α has length equal to at most $\log_{\lambda(g,m)} i(\alpha,\beta) + 3$.*

Proof. Let $\lambda = \lambda(g,m)$. First note that $i(\alpha,\beta_1) \leq i(\alpha,\beta)$. By re-iterating Lemmas 7.5 and 7.6, we obtain

$$i(\alpha,\beta_n) \leq \frac{i(\alpha,\beta)}{\lambda^{n-1}}$$

and hence

$$n \leq \log_\lambda i(\alpha,\beta) - \log_\lambda i(\alpha,\beta_n) + 1.$$

If α and β_n fill S then, by Lemmas 7.1 and 7.6, we have

$$i(\alpha,\beta_n) \geq \kappa(g,m) + 2 \geq 2\lambda$$

which implies

$$n \leq \log_\lambda i(\alpha,\beta) - \log_\lambda 2\lambda + 1 < \log_\lambda i(\alpha,\beta).$$

Therefore, $d_S(\alpha,\beta_n) \leq 2$ whenever $n \geq \lceil \log_\lambda i(\alpha,\beta) \rceil$. $\square$

Corollary 7.8. *Let α and β be curves in $\mathcal{C}(S)$. Then*

$$d_S(\alpha,\beta) \leq \frac{1}{\log \lambda(g,m)} \cdot \log i(\alpha,\beta) + 3$$

whenever $i(\alpha,\beta) \neq 0$. $\square$

Upon fixing a value of g, we have $\lambda(g,m) > \sqrt{\kappa(g,m)}$ for m sufficiently large. Thus, the above bound performs better than Bowditch's.

7.3 Punctured graphs

A *punctured graph* $\mathcal{G} = (\mathcal{G}, V^\circ)$ is a finite metric graph $\mathcal{G} = (V, E)$ together with a distinguished set of vertices $V^\circ = V^\circ(\mathcal{G}) \subseteq V(\mathcal{G})$ called *punctured vertices.*

Let $S = S_{g,m}$ and suppose α and β are a pair of filling curves in $\mathcal{C}(S)$. We construct a punctured graph $\mathcal{G} = \mathcal{G}(\alpha, \beta)$ with the property that any curve $\gamma \in \mathcal{C}(S)$ disjoint from β gives rise to an edge-path in $\mathcal{G}$ whose length equals $i(\alpha, \gamma)$. This allows us to convert the problem of proving Lemma 7.6 to that of finding upper bounds for certain edge-paths in punctured graphs.

Let $\boldsymbol{\alpha} = (\alpha, \beta)$ and $\mathbf{1} = (1, 1)$ and build the singular Euclidean surface $S(\mathbf{1} \cdot \boldsymbol{\alpha})$ as described in Section 3.1. We will regard $S(\mathbf{1} \cdot \boldsymbol{\alpha})$ as a 2–dimensional complex tiled by squares. Let $\mathcal{G} = \mathcal{G}(\alpha, \beta)$ be the subgraph of the 1–skeleton of $S(\mathbf{1} \cdot \boldsymbol{\alpha})$ whose edges are dual to α. Note that $\mathcal{G}$ contains all vertices of $S(\mathbf{1} \cdot \boldsymbol{\alpha})$. The set of punctured vertices $V^\circ(\mathcal{G})$ is defined to be the marked points of S. (Recall that any marked points on S must coincide with some vertex of the square tiling.) Finally, we equip $\mathcal{G}$ with the induced path metric from $S(\mathbf{1} \cdot \boldsymbol{\alpha})$, thus each edge has unit length.

Lemma 7.9. *There exists a deformation retract of $S(\mathbf{1} \cdot \boldsymbol{\alpha}) - \beta$ onto $\mathcal{G}$.*

Proof. We define a homotopy $\mathbf{F} : (S(\mathbf{1} \cdot \boldsymbol{\alpha}) - \beta) \times [0,1] \to \mathcal{G}$ as follows. Each square R in $S(\mathbf{1} \cdot \boldsymbol{\alpha})$ can be isometrically identified with the unit square $[0,1] \times [0,1]$ so that $\beta \cap R$ is identified with $[0,1] \times \{\frac{1}{2}\}$. Define $\mathbf{F}_t(\cdot) = \mathbf{F}(\cdot, t)$ on $R - \beta$ by setting

$$\mathbf{F}_t(x,y) := \begin{cases} (x, (1-t)y) & : y < \frac{1}{2}; \\ (x, 1-(1-t)y) & : y > \frac{1}{2}. \end{cases}$$

One can check that $\mathbf{F}_t$ can be consistently extended to all of $S(\mathbf{1} \cdot \boldsymbol{\alpha}) - \beta$. $\square$

Now $\mathcal{G}$ has either one or two components, depending on whether β is non-separating or separating. Let S_β be the closure of $S(\mathbf{1} \cdot \boldsymbol{\alpha}) - \beta$. If β is separating, we will also write $S_\beta^\pm$ for the components of S_β and $\mathcal{G}^\pm$ for corresponding components of $\mathcal{G}$. If one component of S_β is a disc with two marked points, we will always refer to that component as S_β^- and call β a *pants curve.*

The *volume* $\mathrm{vol}(\mathcal{G})$ of $\mathcal{G}$ is the sum of all edge lengths in $\mathcal{G}$.

Lemma 7.10. *Let* $\mathcal{G} = \mathcal{G}(\alpha, \beta)$. *Then* $\mathrm{vol}(\mathcal{G}) = i(\alpha, \beta)$. *Moreover, if* β *is separating then* $\mathrm{vol}(\mathcal{G}^+) = \mathrm{vol}(\mathcal{G}^-) = \frac{1}{2}\, i(\alpha, \beta)$.

Proof. Observe that $\alpha - \beta$ consists of exactly $i(\alpha, \beta)$ disjoint arcs which connect two points on ∂S_β. By definition of $\mathcal{G}$, these arcs are in one-to-one correspondence with the edge set $\mathcal{G}$, simply by taking the corresponding dual edge, and hence the first claim follows. If β is separating, then the endpoints of each arc must lie on the same component of ∂S_β. Since each component of ∂S_β meets exactly $i(\alpha, \beta)$ of these endpoints, the second claim also holds. □

Define the *rank* (also called the *Betti number*) of a connected punctured graph $\mathcal{G}$ to be

$$\mathrm{rank}(\mathcal{G}) := |E(\mathcal{G})| - |V(\mathcal{G})| + 1.$$

Note that $\pi_1(\mathcal{G})$ is a free group on $\mathrm{rank}(\mathcal{G})$ generators.

Lemma 7.11. *If* β *is non-separating, then* $\mathrm{rank}(\mathcal{G}) = 2g - 1$ *and* $|V^\circ(\mathcal{G})| = m$. *If* β *is a pants curve then* $\mathrm{rank}(\mathcal{G}^+) = 2g$ *and* $|V^\circ(\mathcal{G}^+)| = m - 2$.

Proof. First recall that if Σ is a compact surface with boundary, then $\pi_1(\Sigma)$ is a free group of rank $2\,\mathrm{genus}(\Sigma) + |\partial\Sigma| - 1$. (We are ignoring the marked points at the level of the fundamental group.) If β is non-separating then S_β is a connected surface of genus $g - 1$ with two boundary components containing m marked points. Since $\mathcal{G}$ is homotopy equivalent to S_β, we have

$$\mathrm{rank}(\mathcal{G}) = \mathrm{rank}(\pi_1(\mathcal{G})) = \mathrm{rank}(\pi_1(S_\beta)) = 2g - 1.$$

Similarly, when β is a pants curve S_β^+ has genus g, one boundary component and $m - 2$ marked points, hence $\mathrm{rank}(\mathcal{G}^+) = 2g$. Finally, recall that the punctured vertices of $\mathcal{G}$ are exactly the marked points of S and we are done. □

Let $\gamma \in \mathcal{C}(S)$ be a curve disjoint from β. If $\gamma \neq \beta$ then any geodesic representative of γ on $S(\mathbf{1}\cdot\boldsymbol{\alpha})$ must in fact lie in $\mathcal{G}(\alpha, \beta)$. (Recall that we allow representatives of curves to meet marked points.) Thus, we may regard γ as a closed edge-path in $\mathcal{G}$ whose combinatorial length coincides with its Euclidean length on $S(\mathbf{1}\cdot\boldsymbol{\alpha})$.

Lemma 7.12. *Let* $\gamma \in \mathcal{C}(S)$ *be a curve disjoint from* β. *Then* $l(\gamma) = i(\gamma, \alpha)$, *where* $l(\gamma)$ *is the length of* γ *in* $\mathcal{G}$. *Moreover, the subgraph* $\mathcal{H}(\gamma)$ *of* $\mathcal{G}$ *spanned by the edges of* γ *must contain an embedded cycle or at least two punctured vertices.*

Proof. The first claim is immediate. For the second claim, suppose $\mathcal{H}(\gamma)$ contains no embedded cycles and has at most one punctured vertex. Then a regular neighbourhood $N(\mathcal{H}(\gamma))$ in $S(\mathbf{1} \cdot \boldsymbol{\alpha})$ is a disc containing γ with at most one marked point. This implies that γ is trivial or peripheral – a contradiction. $\square$

Lemma 7.13. *Any embedded loop σ in $\mathcal{G}$ arises as a representative of curve $\gamma \in \mathcal{C}(S)$ disjoint from β. Similarly, let τ be an embedded path connecting two punctured vertices in $\mathcal{G}$ which meets no other punctured vertex. Then there is a pants curve $\gamma' \in \mathcal{C}(S)$ disjoint from β which has a representative running over each edge of τ exactly twice.*

Proof. Given an essential loop $\sigma \subseteq \mathcal{G}$, take γ to be any boundary component of a regular neighbourhood $N(\sigma)$ in S_β. Since σ represents a non-trivial element of $\pi_1(\mathcal{G})$, it follows that γ is non-trivial in $\pi_1(S_\beta)$. Thus γ is either essential on S_β; or parallel to ∂S_β in which case $\gamma = \beta$ in $\mathcal{C}(S)$.

A regular neighbourhood $N(\tau)$ on S_β of a punctured path τ is a disc with two marked points. Taking $\gamma' = \partial N(\tau)$ gives the desired curve. $\square$

This result motivates the following definitions: A *systole* in $\mathcal{G}$ is an embedded loop of minimal length in $\mathcal{G}$. A *punctured path* in $\mathcal{G}$ is an embedded path in $\mathcal{G}$ connecting a pair of distinct punctured vertices with no punctured vertices in its interior. A *tense path* is a punctured path of minimal length.

Definition 7.14. For a punctured graph $\mathcal{G}$, define $\mathrm{girth}(\mathcal{G})$ to be the length of any systole in $\mathcal{G}$ and $\mathrm{pathgirth}(\mathcal{G})$ to be twice the length of any tense path in $\mathcal{G}$. Set $\overline{\mathrm{girth}}(\mathcal{G}) := \min\{\mathrm{girth}(\mathcal{G}),\ \mathrm{pathgirth}(\mathcal{G})\}$.

We summarise the results of this section below:

Proposition 7.15. *Let $S = S_{g,m}$. Given filling curves $\alpha, \beta \in \mathcal{C}(S)$, let $\mathcal{G} = \mathcal{G}(\alpha, \beta)$ be as above. Let $\gamma \in \mathcal{C}(S)$ be a curve which has minimal intersection number with α among all curves disjoint from β. Then $i(\gamma, \alpha) = \overline{\mathrm{girth}}(\mathcal{G})$ and γ is either a systole in $\mathcal{G}$, or runs over each edge of tense path exactly twice in $\mathcal{G}$.* $\square$

Consequently, the problem of bounding the intersection number $i(\alpha, \gamma)$ is equivalent to bounding the volume of systoles or tense paths in $\mathcal{G}$.

7.4 Girth, genus and doubles

Given a punctured graph $\mathcal{G}$, let $\mathrm{genus}(\mathcal{G})$ be the minimal integer $k \geq 0$ such that there exists an embedding of $\mathcal{G}$ into a surface of genus k. We can always take such a minimal genus surface Σ to be closed. If $\mathcal{G} \hookrightarrow \Sigma$ is a (smooth) embedding, then all complementary components of $\mathcal{G}$ in Σ are discs. The following result is well-known in the literature:

Lemma 7.16. *A minimal genus embedding $\mathcal{G} \hookrightarrow \Sigma$ has* $\operatorname{rank}(\mathcal{G}) - 2\operatorname{genus}(\mathcal{G}) + 1$ *complementary discs.*

Proof. This follows from an Euler characteristic calculation:

$$2 - 2\ \operatorname{genus}(\Sigma) = \chi(\Sigma) = |F| - |E| + |V| = |F| - \operatorname{rank}(\mathcal{G}) + 1,$$

where $V = V(\mathcal{G})$, $E = E(\mathcal{G})$ and F is the set of complementary regions. □

Let us write $\operatorname{face}(\mathcal{G}) := \operatorname{rank}(\mathcal{G}) - 2\operatorname{genus}(\mathcal{G}) + 1$. Since $\operatorname{face}(\mathcal{G}) \geq 1$, it follows that $\operatorname{genus}(\mathcal{G}) \leq \frac{1}{2}\operatorname{rank}(\mathcal{G})$.

Lemma 7.17. *For any connected punctured graph $\mathcal{G}$, we have*

$$\operatorname{girth}(\mathcal{G}) \leq \frac{2\operatorname{vol}(\mathcal{G})}{\operatorname{face}(\mathcal{G})}.$$

Proof. Suppose $\mathcal{G} \hookrightarrow \Sigma$ is a minimal genus embedding, with Σ closed. Let F be the set of complementary components of $\mathcal{G}$ in Σ. The set F can be viewed as the *faces* of a 2–cell decomposition of Σ with $\mathcal{G}$ as its 1–skeleton. By summing the lengths of ∂f for each face $f \in F$, we count each edge of $\mathcal{G}$ exactly twice, hence

$$\sum_f \operatorname{vol}(\partial f) = 2\operatorname{vol}(\mathcal{G}).$$

Observe that each ∂f is an essential loop in $\mathcal{G}$. Thus,

$$\operatorname{girth}(\mathcal{G}) \leq \min_f \operatorname{vol}(\partial f) \leq \frac{2\operatorname{vol}(\mathcal{G})}{|F|} = \frac{2\operatorname{vol}(\mathcal{G})}{\operatorname{face}(\mathcal{G})},$$

where we have applied Lemma 7.16 for the final equality. □

Definition 7.18. The *double* $\mathcal{DG}$ of $\mathcal{G}$ is the graph obtained by gluing two disjoint copies of $\mathcal{G}$ along its punctured vertices and declaring all vertices to be non-punctured. More formally,

$$\mathcal{DG} := \mathcal{G} \times \{0, 1\} \ /\ (v, 0) \sim (v, 1) \text{ for all } v \in V^\circ(\mathcal{G})$$

with $V^\circ(\mathcal{DG}) = \emptyset$.

Lemma 7.19. *Let $\mathcal{G}$ be a connected punctured graph with at least one punctured vertex. Then*

$$\mathrm{rank}(\mathcal{DG}) = 2\ \mathrm{rank}(\mathcal{G}) + |V^\circ(\mathcal{G})| - 1$$

and

$$\mathrm{girth}(\mathcal{DG}) = \overline{\mathrm{girth}}(\mathcal{G}).$$

Proof. The first claim follows from observing that $|V(\mathcal{DG})| = 2\ |V(\mathcal{G})| - |V^\circ(\mathcal{G})|$ and $|E(\mathcal{DG})| = 2\ |E(\mathcal{G})|$.

For the second claim, observe that the double of any systole σ in $\mathcal{G}$ is the union of two embedded loops in $\mathcal{DG}$, both of the same length as σ. The double of a tense path τ in $\mathcal{G}$ is an embedded loop in $\mathcal{DG}$ of twice the original length. Hence $\mathrm{girth}(\mathcal{DG}) \leq \overline{\mathrm{girth}}(\mathcal{G})$.

To prove the other direction, let σ be a systole in $\mathcal{DG}$. Let σ' be the pre-image of σ under the natural projection $\mathcal{G} \times \{0,1\} \to \mathcal{DG}$. If all edges of σ' are contained in one component of $\mathcal{G} \times \{0,1\}$ then σ' is isometric to an embedded loop in $\mathcal{G}$ and so $\mathrm{length}(\sigma) \geq \mathrm{girth}(\mathcal{G})$. If not, then σ' is the disjoint union of two punctured paths, one in each component of $\mathcal{G} \times \{0,1\}$, whose endpoints become identified upon projecting to $\mathcal{DG}$. These punctured paths must have equal length, for otherwise we can double the shorter one to obtain an essential loop in $\mathcal{DG}$ shorter than σ. Thus $\mathrm{length}(\sigma) \geq \mathrm{pathgirth}(\mathcal{G})$ and we are done. □

Corollary 7.20. *For a connected punctured graph $\mathcal{G}$ we have*

$$\overline{\mathrm{girth}}(\mathcal{G}) \leq \frac{4\,\mathrm{vol}(\mathcal{G})}{\mathrm{face}(\mathcal{DG})}$$

whenever $\mathcal{G}$ has at least one punctured vertex. □

Lemma 7.21. *Let $\mathcal{G}$ be a connected punctured graph with at least one punctured vertex. Then*

- $\mathrm{genus}(\mathcal{DG}) \leq 2\,\mathrm{genus}(\mathcal{G})$ *if* $\mathrm{face}(\mathcal{G}) = 1$*, and*
- $\mathrm{genus}(\mathcal{DG}) \leq 2\,\mathrm{genus}(\mathcal{G}) + \min\{\mathrm{face}(\mathcal{G}) - 1, |V^\circ(\mathcal{G})|\} - 1$ *otherwise.*

Proof. Let $\mathcal{G} \hookrightarrow \Sigma$ be a minimal genus embedding. First, suppose $\mathrm{face}(\mathcal{G}) = 1$. There exists a closed disc D on Σ so that $\mathcal{G} \cap \mathrm{int}D = \emptyset$ and $\mathcal{G} \cap \partial D$ is precisely $V^\circ(\mathcal{G})$ (we may take the complement an open regular neighbourhood of $\mathcal{G}$ on Σ and perturb near the punctured vertices, for example). By doubling $\Sigma - \mathrm{int}(D)$ along its boundary, we obtain an embedding of $\mathcal{DG}$ on a surface Σ' of twice the genus of Σ.

Now suppose $\mathrm{face}(\mathcal{G}) \geq 2$. Observe that any punctured vertex of $\mathcal{G}$ must meet the closure of at least one complementary region of $\mathcal{G}$ on Σ. We can choose an embedding so that all the punctured vertices of $\mathcal{G}$ are contained in $\bar{R}_1 \cup \ldots \cup \bar{R}_k$, the union of the closures of $1 \leq k \leq \min\{\mathrm{face}(\mathcal{G}) - 1, |V^\circ(\mathcal{G})|\}$ complementary regions. There exists a closed disc $D_i \subset \bar{R}_i$ so that $\mathcal{G} \cap \mathrm{int} D_i = \emptyset$ for each i and

$$\mathcal{G} \cap (\partial D_1 \cup \ldots \cup \partial D_k) = V^\circ(\mathcal{G}).$$

We can then double $\Sigma - (\mathrm{int} D_1 \cup \ldots \cup \mathrm{int} D_k)$ along its boundary to obtain a surface Σ' of

$$\mathrm{genus}(\Sigma') = 2\,\mathrm{genus}(\Sigma) + k - 1$$

admitting an embedding of $\mathcal{DG}$. $\square$

A graph $\mathcal{G}$ of rank at most 3 cannot contain a subgraph homemomorphic to K_5 (the complete graph on 5 vertices), or $K_{3,3}$ (the complete bipartite graph on 6 vertices, where 3 of the vertices are adjacent to the other 3). By Kuratowski's theorem, such graphs must be planar and so $\mathrm{genus}(\mathcal{G}) = 0$ whenever $\mathrm{rank}(\mathcal{G}) \leq 3$.

Corollary 7.22. *Let $\mathcal{G}$ be a connected punctured graph with at least one punctured vertex. Then* $\mathrm{genus}(\mathcal{DG}) = 0$ *if* $\mathrm{rank}(\mathcal{G}) \leq 1$*; and* $\mathrm{genus}(\mathcal{DG}) \leq 1$ *if* $\mathrm{rank}(\mathcal{G}) = 2$. $\square$

7.5 Bounding the ratio

Let $S = S_{g,m}$ with $\xi(S) \geq 2$. Suppose $\alpha, \beta \in \mathcal{C}(S)$ fill S and, furthermore, assume β is either non-separating or a pants curve. Let $\mathcal{G} = \mathcal{G}(\alpha, \beta)$ be the punctured graph defined in Section 7.3. Recall from Lemma 7.10 that $\mathrm{vol}(\mathcal{G}) = i(\alpha, \beta)$; and $\mathrm{vol}(\mathcal{G}^\pm) = \frac{1}{2}\, i(\alpha, \beta)$ if β is separating. By Lemma 7.11, we have $\mathrm{rank}(\mathcal{G}) = 2g - 1$ and $|V^\circ(\mathcal{G})| = m$ if β is non-separating; and $\mathrm{rank}(\mathcal{G}^+) = 2g$ and $|V^\circ(\mathcal{G}^+)| = m - 2$ if β is a pants curve.

Let $\gamma \in \mathcal{C}(S)$ be a curve with minimal intersection with α among all curves disjoint from β. Recall from Proposition 7.15 that $i(\alpha, \gamma) = \overline{\mathrm{girth}}(\mathcal{G})$.

Proof of Lemma 7.6.

If $g = 0$ and $m \geq 5$, then all curves on S are separating. From the above discussion, $\mathrm{rank}(\mathcal{G}^+) = 0$ and $|V^\circ(\mathcal{G}^+)| = m - 2$. By Lemma 7.19 and Corollary 7.22, we deduce $\mathrm{face}(\mathcal{DG}^+) = m - 2$ and $\mathrm{genus}(\mathcal{DG}^+) = 0$. Applying Corollary 7.20, we obtain

$$i(\alpha,\gamma) = \overline{\mathrm{girth}}(\mathcal{G}) = \mathrm{girth}(\mathcal{DG}) \leq \frac{4\,\mathrm{vol}(\mathcal{G}^+)}{\mathrm{face}(\mathcal{DG}^+)} \leq \frac{2\,i(\alpha,\beta)}{m-2} = \frac{i(\alpha,\beta)}{\lambda(0,m)}.$$

Now suppose $g = 1$ and $m \geq 3$. If β is non-separating then $rk(\mathcal{G}) = 1$ and $|V^\circ(\mathcal{G})| = m$. By a similar argument, we deduce $\mathrm{rank}(\mathcal{DG}) = 2 + m - 1$ and $\mathrm{genus}(\mathcal{DG}) = 0$. This yields $\mathrm{face}(\mathcal{DG}) = m + 2$ and thus $i(\alpha,\gamma) \leq \frac{4\,i(\alpha,\beta)}{m+2}$. Following the same reasoning for when β is a pants curve, we obtain $i(\alpha,\gamma) \leq \frac{2\,i(\alpha,\beta)}{m}$. Hence

$$i(\alpha,\gamma) \leq \max\left\{\frac{4\,i(\alpha,\beta)}{m+2},\ \frac{2\,i(\alpha,\beta)}{m}\right\} = \frac{4\,i(\alpha,\beta)}{m+2} = \frac{i(\alpha,\beta)}{\lambda(1,m)}.$$

For the general case, observe that if $(\mathcal{H}, V^\circ(\mathcal{H})) \subseteq (\mathcal{G}, V^\circ(\mathcal{G}))$ is a punctured subgraph, that is $\mathcal{H} \subseteq \mathcal{G}$ and $V^\circ(\mathcal{H}) \subseteq V^\circ(\mathcal{G})$, then $\overline{\mathrm{girth}}(\mathcal{G}) \leq \overline{\mathrm{girth}}(\mathcal{H})$. We may then reduce our problem to one of the above scenarios by removing suitable edges from $\mathcal{G}$. $\square$

In principle, we can apply the results of Section 7.4 to obtain improved bounds for when $\mathrm{genus}(S) \geq 2$. This boils down to the problem of finding upper bounds on the genus of graphs of a given rank. However, the problem of determining the genus of a finite graph is NP–hard; and determining whether a given graph has genus g is NP–complete [Tho89].

Bibliography

[ABC$^+$91] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, *Notes on word hyperbolic groups*, Group theory from a geometrical viewpoint (Trieste, 1990), World Sci. Publ., River Edge, NJ, 1991, Edited by Short, pp. 3–63. MR 1170363 (93g:57001)

[And05] James W. Anderson, *Hyperbolic geometry*, second ed., Springer Undergraduate Mathematics Series, Springer-Verlag London Ltd., London, 2005. MR 2161463 (2006b:51001)

[Aou] Tarik Aougab, *Uniform hyperbolicity of the graphs of curves*, preprint, (2012), arxiv:1212.3160.

[Apa00] Boris N. Apanasov, *Conformal geometry of discrete groups and manifolds*, de Gruyter Expositions in Mathematics, vol. 32, Walter de Gruyter & Co., Berlin, 2000. MR 1800993 (2002k:57087)

[BCM12] Jeffrey F. Brock, Richard D. Canary, and Yair N. Minsky, *The classification of Kleinian surface groups, II: The ending lamination conjecture*, Ann. of Math. (2) **176** (2012), no. 1, 1–149. MR 2925381

[Bea95] Alan F. Beardon, *The geometry of discrete groups*, Graduate Texts in Mathematics, vol. 91, Springer-Verlag, New York, 1995, Corrected reprint of the 1983 original. MR 1393195 (97d:22011)

[Bes52] A. S. Besicovitch, *On two problems of Loewner*, J. London Math. Soc. **27** (1952), 141–144. MR 0047126 (13,831d)

[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. MR 1744486 (2000k:53038)

[Bowa] Brian H. Bowditch, *The ending lamination theorem*, in preparation, Warwick (2011), `http://www.warwick.ac.uk/~masgak/preprints.html`.

[Bowb] ———, *Uniform hyperbolicity of the curve graphs*, preprint, Warwick (2012), `http://www.warwick.ac.uk/~masgak/preprints.html`.

[Bow06a] ———, *A course on geometric group theory*, MSJ Memoirs, vol. 16, Mathematical Society of Japan, Tokyo, 2006. MR 2243589 (2007e:20085)

[Bow06b] ———, *Intersection numbers and the hyperbolicity of the curve complex*, J. Reine Angew. Math. **598** (2006), 105–129. MR 2270568 (2009b:57034)

[Bow07] ———, *Length bounds on curves arising from tight geodesics*, Geom. Funct. Anal. **17** (2007), no. 4, 1001–1042. MR 2373010 (2009h:57027)

[Bow08] ———, *Tight geodesics in the curve complex*, Invent. Math. **171** (2008), no. 2, 281–300. MR 2367021 (2008m:57040)

[CRS] Matt T. Clay, Kasra Rafi, and Saul Schleimer, *Uniform hyperbolicity of the curve graph via surgery sequences*, preprint, (2013), arxiv:1302.5519.

[EM06] D. B. A. Epstein and A. Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces [mr0903852]*, Fundamentals of hyperbolic geometry: selected expositions, London Math. Soc. Lecture Note Ser., vol. 328, Cambridge Univ. Press, Cambridge, 2006, pp. 117–266. MR 2235711

[FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012. MR 2850125 (2012h:57032)

[Gor07] Cameron McA. Gordon, *Problems*, Workshop on Heegaard Splittings, Geom. Topol. Monogr., vol. 12, Geom. Topol. Publ., Coventry, 2007, pp. 401–411. MR 2408256 (2009i:57001)

[Gro87] M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263. MR 919829 (89e:20070)

[Gro99] Misha Gromov, *Metric structures for Riemannian and non-Riemannian spaces*, Progress in Mathematics, vol. 152, Birkhäuser Boston Inc., Boston, MA, 1999, Based on the 1981 French original [MR0682063

(85e:53051)], With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates. MR 1699320 (2000d:53065)

[Ham06] Ursula Hamenstädt, *Train tracks and the Gromov boundary of the complex of curves*, Spaces of Kleinian groups, London Math. Soc. Lecture Note Ser., vol. 329, Cambridge Univ. Press, Cambridge, 2006, pp. 187–207. MR 2258749 (2009a:30092)

[Har81] W. J. Harvey, *Boundary structure of the modular group*, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), Ann. of Math. Stud., vol. 97, Princeton Univ. Press, Princeton, N.J., 1981, pp. 245–251. MR 624817 (83d:32022)

[Hem01] John Hempel, *3-manifolds as viewed from the curve complex*, Topology **40** (2001), no. 3, 631–657. MR 1838999 (2002f:57044)

[HPW] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb, *Slim unicorns and uniform hyperbolicity for arc graphs and curve graphs*, preprint, (2013), arxiv:1301.5577.

[Iva97] Nikolai V. Ivanov, *Automorphism of complexes of curves and of teichmüller spaces*, International Mathematics Research Notices **1997** (1997), no. 14, 651–666.

[Kla] Erica Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space*, preprint, (1999), `http://ericaklarreich.com/research.html`.

[Kob88] Tsuyoshi Kobayashi, *Heights of simple loops and pseudo-Anosov homeomorphisms*, Braids (Santa Cruz, CA, 1986), Contemp. Math., vol. 78, Amer. Math. Soc., Providence, RI, 1988, pp. 327–338. MR 975087 (89m:57015)

[Luo00] Feng Luo, *Automorphisms of the complex of curves*, Topology **39** (2000), no. 2, 283–298. MR 1722024 (2000j:57045)

[Mar07] A. Marden, *Outer circles*, Cambridge University Press, Cambridge, 2007, An introduction to hyperbolic 3-manifolds. MR 2355387 (2008i:57001)

[MM99] Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149. MR 1714338 (2000i:57027)

[RS09] Kasra Rafi and Saul Schleimer, *Covers and the curve complex*, Geom. Topol. **13** (2009), no. 4, 2141–2162. MR 2507116

[Sch] Saul Schleimer, *Notes on the complex of curves*, Available online at `http://homepages.warwick.ac.uk/~masgar/Maths/notes.pdf`.

[Tan12] Robert Tang, *The curve complex and covers via hyperbolic 3-manifolds*, Geometriae Dedicata **161** (2012), 233–237.

[Tho89] Carsten Thomassen, *The graph genus problem is NP-complete*, J. Algorithms **10** (1989), no. 4, 568–576. MR 1022112 (91d:68054)