

SPHERES IN THE CURVE GRAPH AND LINEAR CONNECTIVITY OF THE GROMOV BOUNDARY

ALEX WRIGHT

ABSTRACT. We consider the curve graph in the cases where it is not a Farey graph, and show that its Gromov boundary is linearly connected. For a fixed center point c and radius r , we define the sphere of radius r to be the induced subgraph on the set of vertices of distance r from c . We show that these spheres are always connected in high enough complexity, and prove a slightly weaker result for low complexity surfaces.

CONTENTS

1.	Introduction	1
2.	Sufficient conditions for connectivity	5
3.	The essentially non-separating curve graph	9
4.	The Bounded Geodesic Image Theorem	13
5.	The key propositions	14
6.	The five times punctured sphere	17
7.	The medium and high complexity cases	21
	References	26

1. INTRODUCTION

1.1. Main results. Let $\Sigma = \Sigma_{g,n}$ be a connected surface with genus g and n punctures. We always assume g and n are such that Σ has a complete, finite volume hyperbolic metric and that $(g, n) \neq (0, 3)$. We define the complexity of Σ as $\xi(\Sigma) = 3g - 3 + n$. We say Σ is

- *exceptional* if $\xi(\Sigma) = 1$, i.e. $(g, n) \in \{(1, 1), (0, 4)\}$,
- *low complexity* if $\xi(\Sigma) = 2$, i.e. $(g, n) \in \{(1, 2), (0, 5)\}$,
- *medium complexity* if $\xi(\Sigma) = 3$, i.e. $(g, n) \in \{(2, 0), (1, 3), (0, 6)\}$,
- *high complexity* if $\xi(\Sigma) \geq 4$.

We let $\mathcal{C}\Sigma$ be the curve graph of Σ , and note the exceptional isomorphisms $\mathcal{C}\Sigma_{(1,1)} \cong \mathcal{C}\Sigma_{(0,4)}$, $\mathcal{C}\Sigma_{(1,2)} \cong \mathcal{C}\Sigma_{(0,5)}$, and $\mathcal{C}\Sigma_{(2,0)} \cong \mathcal{C}\Sigma_{(0,6)}$.

Fix an arbitrary vertex c in $\mathcal{C}\Sigma$, which we sometimes refer to as the center point. For $r \geq 0$ integral, we let $S_r = S_r(c)$ be the sphere of radius r in the graph $\mathcal{C}\Sigma$, which consists of all the vertices at distance r from c . We say that a subset of vertices of a graph is connected if the induced subgraph is connected.

Theorem 1.1. *For all r , we have the following:*

- *If Σ is high complexity, S_r is connected.*

- If Σ is medium complexity, $S_r(c) \cup S_{r+1}(c)$ is connected.
- If Σ is low complexity, $S_r(c) \cup S_{r+1}(c) \cup S_{r+2}(c)$ is connected.

We will also address the “sphere at infinity” as follows.

Theorem 1.2. *If Σ is non-exceptional, the Gromov boundary of $\mathcal{C}\Sigma$ is linearly connected.*

Here it is implicit that we are using a visual metric on the Gromov boundary. Recall that a metric space (X, d) is said to be linearly connected if there is a constant $L > 0$ such that for each pair $x, y \in X$ there is a compact connected set in X containing x and y of diameter at most $Ld(x, y)$. Linear connectivity is also called bounded turning or $LLC(1)$, and appears in questions relating to the existence of quasi-isometrically embedded hyperbolic planes in hyperbolic spaces [Tuk96, BK05, Mac08, MS20]. In our proof, the compact connected set produced will be a path from x to y , and as discussed in [Mac08, Section 1], this set can in general be taken to be an arc (embedded path) from x to y .

1.2. Previous results. In particular, we obtain a fundamentally new proof of the following result.

Corollary 1.3. *If Σ is non-exceptional, the Gromov boundary of $\mathcal{C}\Sigma$ is path connected and locally path connected.*

Corollary 1.3 was proven in general by Gabai [Gab09]. Connectivity was proven previously by Leininger and Schleimer for $\mathcal{C}\Sigma_{g,n}$ when $g \geq 4$ or when $g \geq 2$ and $n \geq 1$ [LS09], and path connectivity and local path connectivity were proven previously by Leininger, Mj, and Schleimer for $\mathcal{C}\Sigma_{g,1}$ when $g \geq 2$ [LMS11]. These results answer questions of Peter Storm recorded in [KL08a, Question 10] and [Min06, Section 2].

As part of their proof of quasi-isometric rigidity of curve complexes, Rafi and Schleimer used Corollary 1.3 to deduce that the union of a number of consecutive spheres depending on the hyperbolicity constant is connected [RS11, Proposition 4.1]. Schleimer also proved without using Corollary 1.3 that when $g \geq 2$, complements of balls in $\mathcal{C}\Sigma_{g,1}$ are connected [Sch11].

Following his work on path connectivity, Gabai later established higher connectivity results and, extending work of Hensel and Przytycki for $(g, n) = (0, 5)$, proved that boundaries of curve complexes when $g = 0$ are homeomorphic to Nöbeling spaces [HP11, Gab14]. In light of [Har86, AL13, BBM15], it is possible that the techniques in Section 6 could provide a starting point for an attempt to reprove and strengthen such higher connectivity results. We conjecture that Gabai’s higher connectivity results could be strengthened to linear higher connectivity, which may be of interest in light work such as [BK02].

The boundary of the curve complex is naturally homeomorphic to the space of ending laminations [Kla22]. Much previous work has taken place in the space of ending laminations, including the recent work of Chaika and Hensel proving connectivity results for spaces of uniquely ergodic and cobounded laminations [CH19].

1.3. Structure of the proof. In contrast to much of the work above, we work directly in the curve complex, using the Bounded Geodesic Image Theorem of Masur-Minsky [MM00], in the spirit of previous work such as the analysis of dead ends by Birman and Menasco [BM15] and the work of Schleimer cited above [Sch11].

The core approach of this paper might be described briefly as “push paths away from the center point c ”. To accomplish this, sometimes we use (a weaker statement than) Theorem 1.1 for smaller complexity surfaces, so we induct on the complexity of the surface. The base case for the induction on complexity is the five times punctured sphere, where our argument makes use of pentagons and reveals that some points of S_r might almost be thought of as “closer” to S_{r+1} than others.

In medium and high complexity our approach requires that we work with “essentially non-separating curves”, which are curves that are either non-separating or go around a pair of punctures. We prove a number of basic results on what we call the essentially non-separating curve graph, which may be useful for other purposes.

1.4. Motivation. Masur and Minsky famously proved that curve complexes are hyperbolic, and there are now many proofs of this fact [MM99, Aou13, Bow14, CRS14, HPW15]. Moreover, the curve complex exhibits very strong hyperbolicity features even beyond its Gromov hyperbolicity. For example, Dowdall, Duchin, and Masur proved that the “generic” pair of points on S_r are distance *exactly* $2r$ apart [DDM13], and wrote that “In this sense the curve graph is “even more hyperbolic than a tree.”” (This comment seems tailored to comparison to finite valence trees. One might also suggest that in this sense the curve graph is comparably hyperbolic to an infinite valence regular tree.)

Furthermore, most of the spaces closely analogous to curve complexes that arise in the modern study of hierarchical hyperbolicity are quasi-trees [BHS17, BHS19, Sis19]. Thus there seems to be great tension between tree-like behaviour and the connectivity theorems above. This tension is relevant for the elusive and much studied question of whether convex cocompact surface subgroups of mapping class groups exist (see for example [FM02, LS14, Mos06, Rei06]). Linear connectivity in particular is related to this question since convex co-compact surface subgroups gives rise to quasi-isometrically embedded hyperbolic planes in the curve complex [KL08a, Ham05], and proving linear connectivity is, roughly speaking, part way to establishing the existence of many such planes [Tuk96, Mac08].

The study of the topology of the Gromov boundary of the curve complex is far from complete; for example its topological dimension is not known in general [Gab14]. Nonetheless we feel it is worthwhile to draw attention to the rich additional structure on the boundary. Every nice enough hyperbolic space (including curve complexes) can be recovered up to quasi-isometry via a cone construction from its boundary with a visual metric [BS00, Theorem 8.2]. In contrast, the mere homeomorphism type of the boundary contains vastly less information about the space. See for example [Bou97] and [Mac12] for infinite families of hyperbolic groups whose boundaries are the Menger curve but are pairwise non-quasi-isometric to each other. In some situations, one might even go so far as to say that, far from being the end goal, the homeomorphism type of the boundary is the vessel that holds more refined information on the hyperbolic group. We feel that the situation for the curve complex is even more pronounced, since in addition to metric structure the boundary supports a rich collection of subsurface projection maps. Linear connectivity is, to our knowledge, one of the first results on the metric structure of the boundary, with another notable result being the work of Bestvina and Bromberg on its capacity dimension [BB19].

1.5. Open questions. Part of the motivation above concerns the following two questions, which we wish to state explicitly now.

Question 1.4. Do all non-exceptional curve complexes contain quasi-isometrically embedded hyperbolic planes?

The work of Leininger and Schleimer gives a positive answer in the case of $\Sigma_{g,1}$, $g \geq 2$ and surfaces that can be obtained via an appropriate branched cover of such surfaces [LS14].

Question 1.5. In non-exceptional cases, is the cobounded locus in the boundary of \mathcal{CS} quasi-arc connected?

Motivated by the existence of hierarchy paths and our desire to connect the study of the boundary to subsurface projections, we also propose the following.

Question 1.6. For each Σ that is non-exceptional, does there exist a $D > 0$ such that every pair of points in the Gromov boundary can be joined by a path in the Gromov boundary whose projection to the curve complex of each proper subsurface is an un-parametrized D -quasigeodesic?

We believe a positive answer to these questions would signal a significantly improved understanding of the geometry of the curve complex.

The curve graphs of the exceptional surfaces are both the Farey graph, which, being a quasi-tree, has hopelessly disconnected spheres. All of the individual spheres are disconnected in the low complexity case (Corollary 6.12), but our analysis leaves open the following.

Question 1.7. In the medium complexity case, are individual spheres always connected? In the low complexity case, are unions of two consecutive spheres always connected?

Finally, we note that in the study of $\text{Out}(F_n)$, one is interested in the boundary of the free factor complex FF_n , which is known to be path connected and locally path connected when $n \geq 18$ by [BCH21]. Since versions of our main tool, the Bounded Geodesic Image Theorem, are available in that context [SS12, BF14, Tay14], it would be interesting to see if our methods could be useful for questions such as the following.

Question 1.8. When is the Gromov boundary of the free factor complex linearly connected?

After this paper was initially released, it was discovered that an extension of our analysis can be used to give a positive answer to Question 1.7 [HKSS23], and our criterion for linear connectivity was improved and applied to fine curve graphs [LT24].

Acknowledgements. The author thanks David Gabai, Jonah Gaster, Sebastian Hensel, Jeremy Kahn, John Mackay, Howard Masur, Yair Minsky, Piotr Przytycki, Kasra Rafi, Saul Schleimer, Alessandro Sisto, and Richard Webb for helpful conversations. The author thanks Sayantan Khan for comments on a previous draft. The author was partially supported by NSF Grants DMS 1856155 and 2142712 and a Sloan Research Fellowship.

2. SUFFICIENT CONDITIONS FOR CONNECTIVITY

2.1. Unions of spheres. Theorem 1.1 will be proven with the following sufficient condition for when unions of some number w of consecutive spheres are connected. As in the case of the curve graph, here we use the notation $S_r = S_r(c)$ for the sphere of radius r in a fixed arbitrary graph Γ , and we also use $B_\rho(z)$ to denote the set of vertices of distance at most ρ from z . A path means a sequence of vertices each adjacent to the next.

Lemma 2.1. *Let Γ be an arbitrary graph, and let $c \in \Gamma$ be arbitrary. Fix $w > 0$. Suppose that for any r the following hold.*

(1) *For every $z \in S_r(c)$ and $x, y \in S_{r+1}(c) \cap B_1(z)$ there exists a path*

$$x = x_0, x_1, \dots, x_\ell = y$$

with

$$x_i \in S_{r+1} \cup \dots \cup S_{r+w}$$

for $0 \leq i \leq \ell$.

(2) *For every adjacent pair $x, y \in S_r$ there exists a path*

$$x = x_0, x_1, \dots, x_\ell = y$$

with

$$x_i \in S_{r+1} \cup \dots \cup S_{r+w}$$

for $0 < i < \ell$.

Then for any r , the union $S_r \cup \dots \cup S_{r+w-1}$ is connected.

No bound is assumed for the length ℓ of the paths. Later, in Proposition 2.2, we will give a version of this lemma that implies linear connectivity of the Gromov boundary of a hyperbolic graph, which will require additionally that the paths above stay close to their endpoints.

For completeness, we provide a proof of the lemma.

Proof. We prove that $S_r \cup \dots \cup S_{r+w-1}$ is connected by induction on r . The base case of $r = 0$ is trivial, since every vertex can be connected to $S_0 = \{c\}$ by a geodesic.

For the inductive step, suppose that $S_r \cup \dots \cup S_{r+w-1}$ is known to be connected. Hence, any pair $a, b \in S_{r+1} \cup \dots \cup S_{r+w}$ can be joined by a path in $S_r \cup \dots \cup S_{r+w}$.

For any path P in $S_r \cup \dots \cup S_{r+w}$, let $V_r(P)$ denote the number of vertices on the path in S_r , and let $E_r(P)$ denote the number of edges of the path that go between two vertices of S_r . For any $a, b \in S_{r+1} \cup \dots \cup S_{r+w}$, let $f(a, b)$ denote the minimal value of $V_r(P) + E_r(P)$ over all paths P from a to b in $S_r \cup \dots \cup S_{r+w}$. So $f(a, b)$ measures the failure for a and b to lie in the same connected component of $S_{r+1} \cup \dots \cup S_{r+w}$, and our goal is to show that always $f(a, b) = 0$.

Suppose in order to find a contradiction that $S_{r+1} \cup \dots \cup S_{r+w}$ is not connected, and pick $a, b \in S_{r+1} \cup \dots \cup S_{r+w}$ such that $f(a, b)$ is positive but as small as possible. Consider a path P in $S_r \cup \dots \cup S_{r+w}$ with vertices

$$a = q_0, q_1, \dots, q_m = b$$

for which $V_r(P) + E_r(P) = f(a, b)$.

First suppose that this path contains a pair q_j, q_{j+1} of consecutive vertices in S_r . Applying assumption (2) with $x = q_j, y = q_{j+1}$, we obtain a path

$$q_j = x_0, x_1, \dots, x_\ell = q_{j+1}.$$

The path P' with vertices

$$q_0, q_1, \dots, q_j, x_1, \dots, x_{\ell-1}, q_{j+1}, \dots, q_m$$

has $V_r(P') = V_r(P)$ and $E_r(P') = E_r(P) - 1$, giving a contradiction.

Next suppose that the path does not contain a consecutive pair of vertices in S_r . So, if q_j denotes a point on the path in S_r , we know that $q_{j-1}, q_{j+1} \in S_{r+1}$. Applying assumption (1) with $x = q_{j-1}, y = q_{j+1}, z = q_j$, we obtain a path

$$q_{j-1} = x_0, x_1, \dots, x_\ell = q_{j+1}.$$

The path P' with vertices

$$q_0, q_1, \dots, q_{j-1}, x_1, \dots, x_{\ell-1}, q_{j+1}, \dots, q_m$$

has $V_r(P') = V_r(P) - 1$ and $E_r(P') = E_r(P)$, giving a contradiction. \square

2.2. The Gromov boundary. We now give a sufficient condition for the linear connectivity of the Gromov boundary of a graph.

Proposition 2.2. *Let Γ be a Gromov hyperbolic graph. Suppose that $c \in \Gamma$, and that the following hold.*

- (1) *Every point in the Gromov boundary of Γ can be represented by a geodesic ray starting at c .*
- (2) *Every vertex of Γ is adjacent to point of Γ that is 1 farther from c .*

Additionally suppose that there is some $D > 0$ such that the following hold for all $r \geq 0$.

- (3) *For every $z \in S_r(c)$ and $x, y \in S_{r+1}(c) \cap B_1(z)$ there exists a path*

$$x = x_0, x_1, \dots, x_\ell = y$$

with

$$x_i \in (\Gamma - B_r(c)) \cap B_D(x)$$

for $0 \leq i \leq \ell$.

- (4) *For every adjacent pair $x, y \in S_r$ there exists a path*

$$x = x_0, x_1, \dots, x_\ell = y$$

with

$$x_i \in (\Gamma - B_r(c)) \cap B_D(x)$$

for $0 < i < \ell$.

Then the Gromov boundary of Γ is linearly connected.

This can be compared to [BM91, Proposition 3.2] in the context of hyperbolic groups. A significant part of the proof is inspired by [MS20, Proof of Proposition 5.2].

Proof. Let \mathcal{R} be the set of half-infinite geodesic rays beginning at c . A point of \mathcal{R} can be viewed as a function $X : \{0, 1, 2, \dots\} \rightarrow \Gamma$ with $X(i) \in S_i(c)$ and $X(i)$ adjacent to $X(i+1)$.

The first assumption gives that the Gromov boundary is a quotient of \mathcal{R} by the equivalence relation of staying bounded distance apart. The second assumption gives that every point of Γ lies on an element of \mathcal{R} .

We use a visual metric d on the Gromov boundary. Visual metrics have the property that there exists $0 < a < 1$ such that for all E sufficiently large there exist $C > 0$ such that if r is the maximal number such that $d(X(r), Y(r)) \leq E$, then

$$(1/C)a^r \leq d([X], [Y]) \leq Ca^r.$$

Fix $E \geq 10D$, a and C so this estimate holds.

We will say that $X, Y \in R$ are r -adjacent if one of the following holds:

- (1) $X(r) = Y(r)$, or
- (2) $d(X(r), Y(r)) = 1$, or
- (3) $d(X(r), Y(r+1)) = 1$, or
- (4) $d(X(r+1), Y(r)) = 1$.

We will say they are at least r -adjacent if they are s -adjacent for some $s \geq r$. Similarly we will say they are more than r -adjacent if the same holds for some $s > r$. The key lemma is the following.

Lemma 2.3. *There exists a constant C_0 such that if X and Y are r -adjacent, then there is a sequence $X = X_0, X_1, \dots, X_n = Y$ such that X_i and X_{i+1} are at least $(r+1)$ -adjacent for all i , and*

$$\text{diam}(\{[X_0], [X_1], \dots, [X_n]\}) \leq C_0 a^r.$$

Proof. We proceed in cases as follows. In all cases the diameter bound follows immediately from the properties of the visual metric given above, so we merely indicate how to construct the desired sequence.

Case 1: $X(r) = Y(r)$. Assumption (3) from the proposition's statement with $z = X(r) = Y(r)$ gives the existence of a path x_0, \dots, x_ℓ from $x_0 = X(r+1)$ to $x_n = Y(r+1)$. For $0 < i < \ell$, let X_i be any point in \mathcal{R} that passes through x_i . Then

$$X = X_0, X_1, \dots, X_\ell = Y$$

is the desired path.

Case 2: $d(X(r), Y(r)) = 1$. In this case we first apply assumption (4) to get a path x_0, \dots, x_ℓ from $x_0 = X(r)$ to $x_\ell = Y(r)$, and again let X_i be any point in \mathcal{R} that passes through x_i . The geodesic ray X_1 can be chosen to pass through $X(r)$, and the geodesic ray $X_{\ell-1}$ can be chosen to pass through $Y(r)$.

By Case 1, we can find a sequence $X = W_0, W_1, \dots, W_k = X_1$ and a sequence $X_{\ell-1} = Z_0, \dots, Z_m = Y$. The concatenation

$$X = W_0, W_1, \dots, W_k = X_1, X_2, \dots, X_{\ell-1} = Z_0, \dots, Z_m = Y$$

is the desired path

Case 3: $d(X(r), Y(r+1)) = 1$. Let $Z \in \mathcal{R}$ be defined by $Z(i) = X(i)$ for $i \leq r$ and $Z(i) = Y(i)$ for $i > r$. So X and Z are i -adjacent and fall into Case 1, and Z and Y are $(i+1)$ -adjacent. By Case 1, we can find a sequence $X = W_0, W_1, \dots, W_k = Z$, and the concatenation

$$X = W_0, W_1, \dots, W_k = Z, Y$$

is the desired path.

Case 4: $d(X(r+1), Y(r)) = 1$. This is identical to the previous case. □

Let X_0 and Y_0 be points of \mathcal{R} . Let R be the maximal integer with

$$d(X_0(R), Y_0(R)) \leq E.$$

To prove the proposition, it suffices to build a path from X_0 to Y_0 of diameter at most La^R , for some constant L not depending on X_0 or Y_0 , and this is what we will do.

We will iteratively define functions $\gamma_r : I_r \rightarrow \mathcal{R}$, where

$$I_r \subset [0, 1], \quad r = 0, 1, 2, \dots$$

is an increasing nested family of finite sets and $\gamma_r|_{I_w} = \gamma_w$ for $w < r$.

First pick a path x_0, x_1, \dots, x_n from $x_0 = X_0(R)$ to $x_n = Y_0(R)$ of length $n \leq E$. Set

$$I_0 = \{i/n : i = 0, \dots, n\}.$$

Define $\gamma_0(0) = X_0$ and $\gamma_0(1) = Y_0$. For $0 < i < n$, define $\gamma_0(i/n)$ to be any ray in \mathcal{R} passing through x_i . Note that $\gamma_0(i/n)$ and $\gamma_0((i+1)/n)$ are at least $(R-E)$ -adjacent for all i , and that

$$\text{diam}(\gamma(I_0)) \leq C_1 a^R$$

for a constant C_1 not depending on X_0 and Y_0 .

Now, assume we have constructed I_r and γ_r in such a way that adjacent points of I_r map under γ_r to points that are at least $(R-E+r)$ -adjacent. We will define I_{r+1} and γ_{r+1} piece by piece on each interval of $[0, 1] - I_r$.

Suppose (a, b) is an interval of $[0, 1] - I_r$. So we have assumed that $\gamma_r(a)$ and $\gamma_r(b)$ are at least $(R-E+r)$ -adjacent. If these two points are more than $(R-E+r)$ -adjacent, we set

$$I_{r+1} \cap (a, b) = \{(a+b)/2\}$$

and $\gamma_{r+1}((a+b)/2) = \gamma_r(a)$. Otherwise they are $(R-E+r)$ -adjacent, and we use Lemma 2.3 to produce a path $\gamma_r(a) = X_0, X_1, \dots, X_n = \gamma_r(b)$. We set

$$I_{r+1} \cap (a, b) = \{a + i(b-a)/n : i = 1, \dots, n-1\},$$

and define $\gamma_{r+1}(a + i(b-a)/n) = X_i$. (The n here is not the same as the n above.)

Set $I = \cup I_r$, and let γ be the extension of the γ_r to I . Tautologically we get

$$\text{diam}(\gamma(I)) \leq \text{diam}(\gamma(I_0)) + 2 \sum_{r=1}^{\infty} \max_{x \in I_r} \min_{y \in I_{r-1}} d(\gamma(x), \gamma(y)).$$

Given the bounds previously obtained, this becomes

$$\text{diam}(\gamma(I)) \leq C_1 a^R + 2 \sum_{r=1}^{\infty} C_0 a^{R-E+r},$$

which is bounded by a constant times a^R . Since taking closure doesn't change the diameter of the set, we get the same diameter bound for the closure $\overline{\gamma(I)}$.

A similar estimate shows that if (a, b) is a component of $[0, 1] - I_r$ then $\text{diam}(\gamma(I) \cap [a, b])$ has diameter at most a constant times a^r . This implies that γ is uniformly continuous as follows. Let $\epsilon > 0$. Find N such that if (a, b) is a component of $[0, 1] - I_N$ then $\text{diam}(\gamma(I) \cap [a, b]) \leq \epsilon/2$. Let δ be smaller than the largest interval of $[0, 1] - I_N$. If $x, y \in I$ have $|x - y| < \delta$, then we can find $z \in I_N$ within δ of both x and y , and we conclude that

$$d(\gamma(x), \gamma(y)) \leq d(\gamma(x), \gamma(z)) + d(\gamma(z), \gamma(y)) \leq \epsilon/2 + \epsilon/2 = \epsilon.$$

Since γ is uniformly continuous on I it extends to a continuous function on $\bar{I} = [0, 1]$. So $\overline{\gamma(I)}$ is connected, giving the result. \square

3. THE ESSENTIALLY NON-SEPARATING CURVE GRAPH

The purpose of this section is to define a class of curves (and pairs of disjoint curves) called “essentially non-separating”, which for our purposes will be just as good as non-separating curves, and to justify our later claims that we can arrange for various curves to be essentially non-separating. Our motivation for looking at essentially non-separating curves can be seen by looking ahead to Section 4.

3.1. Definitions. Our main interest is punctured surfaces, but since cutting curves leads to surfaces with boundary, the following definition will be helpful.

Definition 3.1. A cut surface is a connected surface of finite genus with a finite number of punctures and boundary components, together with a pairing of a subset of the boundary components¹, such that the surface admits a hyperbolic metric where the punctures are cusps and the boundary consists of geodesics, and the surface is not a pants.

A cut surface has type (h, m, p, u) if it has genus h , m punctures, $2p$ paired boundary components, and u unpaired boundary components.

Throughout this section, $\Upsilon = \Upsilon_{h,m,p,u}$ will denote a cut surface of type (h, m, p, u) . The condition that Υ is not pants excludes the cases when $h = 0$ and $m + 2p + u = 3$. The reader should have in mind that Υ is a component of the complement of a multi-curve on a surface Σ with no boundary components.

Definition 3.2. Let $\hat{\Upsilon}$ denote the surface of genus $h + p$ with m cusps and u boundary components obtained by gluing each of the p pairs. A multi-curve on Υ is called eventually non-separating if does not separate $\hat{\Upsilon}$.

Definition 3.3. A periphery of Υ is one of the punctures or boundary components.

Definition 3.4. A curve on Υ is called a pants curve if it bounds a genus 0 subsurface with no boundary components and two punctures.

Definition 3.5. A curve $\alpha \in \mathcal{C}\Upsilon$ will be called *essentially non-separating* if it is either eventually non-separating or a pants curve.

Definition 3.6. A multi-curve $\alpha \cup \beta$ with two non-isotopic components will be called *essentially non-separating* if α and β are essentially non-separating, and either

- (1) $\alpha \cup \beta$ is eventually non-separating, or
- (2) at least one of α or β is a pants curve, or
- (3) or $\alpha \cup \beta$ bounds a genus zero subsurface with no boundary components and 1 puncture.

Definition 3.7. The essentially non-separating curve graph $\mathcal{C}_0\Upsilon$ is the subgraph of $\mathcal{C}\Upsilon$ whose vertices are essentially non-separating, with an edge from α to β if $\alpha \cup \beta$ is essentially non-separating.

¹A pairing of a set is an equivalence relation where all equivalence classes have size two. The equivalence classes are called pairs. The boundary components in the subset with the equivalence relation are called paired, and the other boundary components are called unpaired.

3.2. Non-emptiness and 1-density. Our first result on the essentially non-separating curve graph concerns its non-emptiness.

Lemma 3.8. *Suppose $h > 0$ or $h = 0$ and either $m \geq 2$ or $p \geq 1$. Then $\mathcal{C}_0\Upsilon \neq \emptyset$.*

For the proof, one should keep in mind that Υ has been assumed not to be a pants.

Proof. If $h > 0$ then $\mathcal{C}_0\Upsilon$ has non-separating curves.

If $h = 0$ and $m \geq 2$, $\mathcal{C}_0\Upsilon$ has pants curves.

If $h = 0$ and $p \geq 1$, $\mathcal{C}_0\Upsilon$ contains separating but eventually non-separating curves. \square

Lemma 3.9. *Suppose that Υ is a subsurface of another cut surface Υ' , and assume that each pair of boundary components of Υ are either a pair of boundary components of Υ' or bound an annulus in Υ' . Then $\mathcal{C}_0\Upsilon$ is a subgraph of $\mathcal{C}_0\Upsilon'$.*

If $\alpha \in \mathcal{C}_0\Upsilon$ and $\beta \in \mathcal{C}_0\Upsilon'$ and β does not cut Υ and is not a boundary component of Υ , then there is an edge in $\mathcal{C}_0\Upsilon'$ from α to β .

Here we assume that each puncture of Υ is a puncture of Υ' , and as always we assume that the boundary components of Υ are essential in Υ' .

Proof. The first claim follows direction from the definitions, after noting that an eventually non-separating multi-curve in Υ is also eventually non-separating in Υ' .

The second claim is by definition if either α or β is a pants curve. So it suffices to note that if both α and β are essentially non-separating then so is $\alpha \cup \beta$. That follows from the fact that if Σ is a connected surface and U is a connected subsurface of Σ , and γ is a non-separating curve on U , then γ is non-separating on Σ , applied to $\Sigma = \hat{\Upsilon}' - \beta$ and $U = \hat{\Upsilon}$ and $\gamma = \alpha$. \square

Lemma 3.10. *Suppose either $h \geq 1$ or suppose $h = 0$ and one of the following hold:*

- (1) $m \geq 3$, or
- (2) $p \geq 2$ and $m + 2p + u \geq 5$, or
- (3) $p = 1$ and $m = 2$ and $u = 0$.

Then $\mathcal{C}_0\Upsilon$ is 1-dense in $\mathcal{C}\Upsilon$.

By definition 1-dense means that every vertex in $\mathcal{C}\Upsilon$ is either in $\mathcal{C}_0\Upsilon$ or adjacent to a vertex in $\mathcal{C}_0\Upsilon$.

Proof. Take $\alpha \in \mathcal{C}\Upsilon$ not in $\mathcal{C}_0\Upsilon$. So α is not eventually non-separating and is not a pants curve.

$\Upsilon - \alpha$ has two components. If $h \geq 1$, one of these has positive genus, and so there is a non-separating curve disjoint from α . So assume $h = 0$.

If there is a pair of punctures on one side of α , then α is disjoint from the loop around that pair. So assume $m \leq 2$, and that there is at most 1 puncture on each side of α .

If $p \geq 2$ and $m + 2p + u \geq 5$, there is a paired curve on one component that also has another periphery that is not the other half of the pair. A loop around the paired curve and this periphery gives an eventually non-separating curve.

The previous argument also works if $p = 1, m = 2, u = 0$, since there is a puncture on each side. \square

3.3. Connectivity. We now turn to the connectivity of $\mathcal{C}_0\Upsilon$, where we break up the work into a number of lemmas. The first few concern the case $h \geq 1$.

Lemma 3.11. *There exists C depending only on Υ such that the following holds. Suppose $\alpha, \beta \in \mathcal{C}\Upsilon$ are non-separating. Then there is a sequence*

$$\alpha = x_0, x_1, \dots, x_k = \beta$$

of non-separating curves, all of which are within distance C of a geodesic from α to β in $\mathcal{C}\Upsilon$, such that:

- (1) *if $h \geq 2$, then x_i is disjoint from x_{i+1} for all i , and moreover $x_i \cup x_{i+1}$ is non-separating, and*
- (2) *if $h = 1$, then the intersection number of x_i and x_{i+1} is equal to 1 for all i .*

Proof sketch. We start by noting that if x, y are disjoint non-separating curves on a surface of genus at least 2, then either $x \cup y$ is non-separating, or there is a curve z disjoint from both x and y such that both $x \cup z$ and $z \cup y$ are jointly non-separating. (To prove this, suppose that $x \cup y$ is separating. One component has positive genus, and we pick z to be non-separating on that component.) This observation implies that in the $h \geq 2$ case it suffices to find a path where all the x_i are non-separating, without worrying that the $x_i \cup x_{i+1}$ are non-separating, since one can then modify such a path by inserting curves between each pair x_i, x_{i+1} as required to get the result.

We also note that if x and y are disjoint non-separating curves on a torus, then there is a curve z that is non-separating and intersects each of x and y once. Thus we get that in the torus case, it similarly suffices to find a sequence where the intersection number between x_i and x_{i+1} is at most 1.

Let $\mathcal{P}\Upsilon$ be the pants complex of Υ . Let P_α be a pants decomposition containing the curve α , and let P_β be a pants decomposition containing the curve β . We will use the following fact: There is a path

$$P_\alpha = P_0, P_1, \dots, P_\ell = P_\beta$$

in the pants complex such that each curve of each P_i is bounded distance from the geodesic from α to β .

Let n_i be a non-separating curve in P_i . Pick $n_0 = \alpha$ and $n_\ell = \beta$. By definition of the pants complex, we have that either n_i and n_{i+1} are disjoint, or intersect once and fill a one-holed torus, or intersect twice and fill a four-holed sphere.

First suppose $h \geq 2$. For each i for which n_i and n_{i+1} fill a one-holed torus, let r_i be a non-separating curve in the complement of that torus, and replace n_i, n_{i+1} in the sequence with n_i, r_i, n_{i+1} .

For each i for which n_i and n_{i+1} fill a four-holed sphere $S_{0,4}$, note that since the n_i are non-separating, at least one of the 4 boundary components of $S_{0,4}$ must be non-separating. Let r_i be a non-separating boundary component of $S_{0,4}$ and again replace n_i, n_{i+1} in the sequence with n_i, r_i, n_{i+1} . This concludes the $h \geq 2$ case.

In the $h = 1$ case, it suffices to consider the pairs n_i, n_{i+1} that fill a $S_{0,4}$. In that case, again we see that at least one boundary component z must be non-separating, and we can replace n_i, n_{i+1} with n_i, z, n_{i+1} in the sequence. \square

Lemma 3.12. *If $h \geq 2$, then $\mathcal{C}_0\Upsilon$ is connected.*

Moreover, there exists C such that any two vertices of $\mathcal{C}_0\Upsilon$ can be joined by a path in $\mathcal{C}_0\Upsilon$ that stays within C of a $\mathcal{C}\Upsilon$ geodesic between the two points.

Proof. We first note that every curve in $\mathcal{C}_0\Upsilon$ either is non-separating or is adjacent in $\mathcal{C}_0\Upsilon$ to a non-separating curve. This follows since a separating curve divides the surface into two components, at least one of which must have positive genus. So, it suffices to show that any two non-separating curves $x, y \in \mathcal{C}_0\Upsilon$ can be connected to each other. That follows from 3.11. \square

Lemma 3.13. *Suppose $h = 1$ and one of the following hold:*

- (1) $m \geq 2$,
- (2) $p \geq 1$ and $m + 2p + u \geq 3$.

Then $\mathcal{C}_0\Upsilon$ is connected.

Moreover, there exists C such that any two vertices of $\mathcal{C}_0\Upsilon$ can be joined by a path in $\mathcal{C}_0\Upsilon$ that stays within C of a $\mathcal{C}\Upsilon$ geodesic between the two points.

Proof. First note that every curve in $\mathcal{C}_0\Upsilon$ is equal to or adjacent in $\mathcal{C}_0\Upsilon$ to a non-separating curve. This follows since the complement of a separating curve must have a component that has genus 1, and that component has a non-separating curve. So, it suffices to show that any two non-separating curves $x, y \in \mathcal{C}_0\Upsilon$ can be connected to each other.

Let $x = x_0, x_1, \dots, x_\ell = y$ be the path produced by Lemma 3.11. Note that $\Upsilon - (x_i \cup x_{i+1})$ is connected for all i .

For each i , let y_i be a pants curve disjoint from $x_i \cup x_{i+1}$ (if $m \geq 2$) or a loop around one paired boundary component and another boundary component that isn't paired to the first (if $p \geq 1$ and $m + 2p + u \geq 3$). Then

$$x_0, y_0, x_1, y_1, x_2, y_1, \dots, x_d$$

is a path in $\mathcal{C}_0\Upsilon$. \square

We now turn to techniques that will work when $h = 0$. For any subset G of the peripheries of Υ , let $\mathcal{C}_G\Upsilon$ be the set of curves in $\mathcal{C}\Upsilon$ that bound a genus 0 surface containing two peripheries, both of which are in G . Thus $\mathcal{C}_G\Upsilon$ corresponds naturally to the simple arcs joining two different elements of G . Note that $\mathcal{C}_G\Upsilon$ is empty if $|G| < 2$.

Lemma 3.14. *There exists C depending only on Υ such that the following is true for any subset G of the peripheries. Suppose $\alpha, \beta \in \mathcal{C}_G\Upsilon$. Then there is a sequence*

$$\alpha = x_0, x_1, \dots, x_\ell = \beta$$

of curves in $\mathcal{C}_G\Upsilon$, such that for all i the arcs corresponding to x_i and x_{i+1} are disjoint except possibly at their endpoints, and such that each x_i lies within C of the geodesic from α to β .

The proof of Lemma 3.14 is contained in [HPW15]. The condition that the arcs are disjoint except possibly at their endpoints can be rephrased by saying that one of the following is true for all i :

- (1) x_i and x_{i+1} are disjoint,
- (2) exactly one periphery enclosed by x_i is also enclosed by x_{i+1} , and x_i and x_{i+1} have intersection number 2, or
- (3) x_i and x_{i+1} enclose the same peripheries and x_i and x_{i+1} have intersection number 4.

Lemma 3.15. *Suppose $h = 0$ and one of the following hold:*

- (1) $p + m \geq 5$, or
 (2) $p = 4$ and $u \geq 1$ and $m = 0$.

Then $\mathcal{C}_0\Upsilon$ is connected.

Moreover, there exists C such that any two vertices of $\mathcal{C}_0\Upsilon$ can be joined by a path in $\mathcal{C}_0\Upsilon$ that stays within C of a $\mathcal{C}\Upsilon$ geodesic between the two points.

While we have not claimed that any of the sufficient conditions in this section are necessary, we wish to emphasize that Lemma 3.15 in particular is likely not sharp.

Proof. If $p + m \geq 5$, let G be a set of peripheries containing one of each pair, together with all punctures. If $m = 0, p = 4, u \geq 1$, let G be a set containing one of each pair and exactly one unpaired boundary. In all cases $|G| \geq 5$, $\mathcal{C}_G\Upsilon \subset \mathcal{C}_0\Upsilon$, and there is an edge in $\mathcal{C}_0\Upsilon$ between any two disjoint curves in $\mathcal{C}_G\Upsilon$.

Sublemma 3.16. *Every curve in $\mathcal{C}_0\Upsilon$ is equal to or adjacent in $\mathcal{C}_0\Upsilon$ to an element of $\mathcal{C}_G\Upsilon$.*

Proof. Consider a curve $\alpha \in \mathcal{C}_0\Upsilon - \mathcal{C}_G\Upsilon$. Since pants curves are in $\mathcal{C}_G\Upsilon$, α must be eventually non-separating. Since $|G| \geq 5$, one side of α has at least three elements in G .

Suppose only a single pair W is separated by α . (This means α is “only barely” eventually non-separating.) One side of α has at least three elements of G , so we can find a loop β around two elements of G , neither of which is W . The curve α is adjacent to β in $\mathcal{C}_0\Upsilon$.

Finally suppose that more than one pair is separated by α . Again since one side of α has at least 3 elements of G , we can pick $\beta \in \mathcal{C}_G\Upsilon$ adjacent to α in $\mathcal{C}_0\Upsilon$. \square

Sublemma 3.17. *Suppose $\alpha, \beta \in \mathcal{C}_G\Upsilon$ are not disjoint but correspond to arcs which are disjoint except possibly at their endpoints. Then there is a curve $\gamma \in \mathcal{C}_G\Upsilon$ disjoint from both α and β .*

Proof. First suppose that α and β have intersection number 4. In this case they fill an annulus containing two elements of G . Since $|G| - 2 \geq 3$, one side of this annulus contains at least 2 elements of G , and we can pick γ to be a loop around two such elements of G .

Next suppose that α and β have intersection number 2. In this case they fill a disc containing three elements of G . The complement of this disc contains at least two elements of G , and we can pick γ to be a loop around two such elements of G . \square

The result now follows easily from Lemma 3.14 and the sublemmas. \square

4. THE BOUNDED GEODESIC IMAGE THEOREM

We start by recalling the main tool for our analysis, the Bounded Geodesic Image Theorem of Masur and Minsky [MM00].

Theorem 4.1. *For every connected surface U , there exists a constant M such that if V is a subsurface of U , and $a, b \in \mathcal{C}U$, and*

$$d_V(a, b) > M,$$

then every geodesic from a to b contains a curve which does not cut V .

Whenever we discuss a subsurface V of a surface U , we assume that every component of the boundary ∂V is essential and non-peripheral in U . The quantity $d_V(a, b)$ denotes the distance in \mathcal{CV} between the subsurface projections to V of a and b . One says that a curve cuts V if it cannot be isotoped outside of V .

It is known that the constant M can be taken to be independent of U [Web15]; we do not require that, but for notational simplicity we will assume we have fixed a constant M that works for all U .

In Section 3, we gave a number of definitions in the context of cut surfaces. For clarity, we now restate special cases of two of these definitions in the more typical context of surfaces with at most finitely many punctures.

Definition 4.2. A simple curve on a surface is called a pants curve if it bounds a genus 0 surface containing exactly two punctures. A curve is called essentially non-separating if it is non-separating or a pants curve.

One can rephrase the definition by saying that a pants curve is one whose complement contains at most one non-pants component. Following [Sch11, BM15, KL08b, Man13], we record the following corollary of Theorem 4.1.

Corollary 4.3. *Assume Σ is non-exceptional. Every $a \in S_r \cap \mathcal{CS}$ that is essentially non-separating is adjacent to a vertex of S_{r+1} . Moreover, that vertex can be taken to be essentially non-separating.*

We include a proof since it inspired our analysis and we will rely frequently on the techniques it uses.

Proof of Corollary 4.3. Note that the assumptions give that either $\Sigma - a$ is connected, or its complement consists of a three times punctured sphere and one other component. In the first case, set $V = \Sigma - a$, and in the second set V to be the component of $\Sigma - a$ that is not a three times punctured sphere. Note that V has infinite diameter curve complex.

Consider any $b \in \mathcal{CV}$ for which $d_V(b, c) > M$. Theorem 4.1 gives that any geodesic from b to c must contain a curve that fails to cut V . The only curve that fails to cut V is a itself, so we see that any geodesic from b to c must pass through a . Since a and b are adjacent, this implies that $b \in S_{r+1}$.

Lemma 3.10 allows us to choose b to be essentially non-separating. \square

Remark 4.4. Our use of the Bounded Geodesic Image Theorem could be avoided, by replacing assumptions like $d_V(b, c) > M$ with assumptions like $d_V(b, c) > Cd(b, c)$ and using that subsurface projections are coarse Lipschitz; compare to [Sch11, Lemma 2.2]. (Here C is the Lipschitz constant.)

5. THE KEY PROPOSITIONS

In this section we state two key propositions, which will be proven in the next two sections, and apply them to obtain Theorem 1.1 and Theorem 1.2.

5.1. The five times punctured sphere. The first of our key propositions is the following.

Proposition 5.1. *In $\mathcal{CS}_{0,5}$, for any c and any r the following hold.*

(1) For every $z \in S_r$ and $x, y \in S_{r+1} \cap B_1(z)$ there exists a path

$$x = x_0, x_1, \dots, x_\ell = y$$

with

$$x_i \in (S_{r+1} \cup S_{r+2} \cup S_{r+3}) \cap B_6(z)$$

for $0 \leq i \leq \ell$.

(2) For every adjacent pair $x, y \in S_r$ there exists a path

$$x = x_0, x_1, x_2, x_3, x_4 = y$$

with

$$x_i \in S_{r+1} \cup S_{r+2}$$

for $0 < i < 4$.

Proposition 5.1 implies the low complexity cases of Theorem 1.1 as follows.

Proof Theorem 1.1 in the cases $(g, n) \in \{(0, 5), (1, 2)\}$. Since $\mathcal{C}\Sigma_{0,5}$ and $\mathcal{C}\Sigma_{1,2}$ are isomorphic, it suffices to prove the result for $\mathcal{C}\Sigma_{0,5}$. The result follows using Lemma 2.1. \square

5.2. A subgraph avoiding dead ends. In the higher complexity case, for technical reasons including the existence of dead ends [BM15], we found it convenient to phrase many of our arguments in a rather specific subgraph of the curve graph. For $c \in \mathcal{C}\Sigma_{g,n}$, define the subgraph $\mathcal{C}_c\Sigma_{g,n}$ of $\mathcal{C}\Sigma_{g,n}$ as follows.

- A vertex $\alpha \in \mathcal{C}\Sigma_{g,n}$ will belong to $\mathcal{C}_c\Sigma_{g,n}$ if it is essentially non-separating or if $\alpha = c$.
- If α, β are vertices of $\mathcal{C}_c\Sigma_{g,n}$ that are joined by an edge in $\mathcal{C}\Sigma_{g,n}$, we join them by an edge in $\mathcal{C}_c\Sigma_{g,n}$ if either $d(c, \alpha) \neq d(c, \beta)$ or if $\Sigma_{g,n} - (\alpha \cup \beta)$ has at most one component that is not a three times punctured sphere.

Thus, the vertices of $\mathcal{C}_c\Sigma_{g,n}$ are $\{c\}$ union the vertices of the essentially non-separating curve graph $\mathcal{C}_0\Sigma_{g,n}$, and disjoint curves are joined by an edge if they have different distances to c or if they are joined by an edge in $\mathcal{C}_0\Sigma_{g,n}$.

Lemma 5.2. *Every $a \in S_r$ can be joined to c by a geodesic $c = x_0, x_1, \dots, x_r = a$ where all $x_i, i < r$ are in $\mathcal{C}_c\Sigma$.*

Note that Lemma 5.2 implies that $S_r^c = \mathcal{C}_c\Sigma \cap S_r$ is the sphere of radius r with center c in the graph $\mathcal{C}_c\Sigma$.

Proof. Let $c = x_0, x_1, \dots, x_r = a$ be a geodesic from c to a with the minimal possible number of x_i not in $\mathcal{C}_c\Sigma$. Suppose in order to find a contradiction that some $x_j, 0 < j < r$ is not in $\mathcal{C}_c\Sigma$.

Let F be the subsurface filled by x_{j-1} and x_{j+1} , so F is connected. If some component of $\Sigma - F$ has at least two punctures, let γ be a loop in $\Sigma - F$ around two punctures. Otherwise, let γ be a non-separating curve in a component of $\Sigma - F$.

Replacing x_j with γ gives a contradiction to the fact that the geodesic was chosen to have the minimum number possible of vertices not in $\mathcal{C}_c\Sigma$. \square

To account for the difference between S_r and S_r^c , we have the following.

Lemma 5.3. *Assume Σ is not exceptional or low complexity. For all $a \in S_r$ not in $\mathcal{C}_c\Sigma$ there exists $a' \in S_r \cap \mathcal{C}_c\Sigma$ adjacent to a .*

Proof. By Lemma 5.2 there exists $b \in S_{r-1} \cap \mathcal{C}_c \Sigma$ adjacent to a . Suppose that a is not in $\mathcal{C}_c \Sigma$, so $\Sigma - a$ has two components, neither of which are a three times punctured sphere.

Let U be the component of $\Sigma - a$ not containing b . Let $a' \in \mathcal{C}_c \Sigma$ be a curve in $\mathcal{C}U$ for which $d_U(a', c) > M$, which can be found using Lemma 3.10.

By Theorem 4.1, every geodesic from a' to c must contain some curve e that fails to cut U . This e must be disjoint from (or equal to) a since it fails to cut U and a is on the boundary of U . If this e has distance less than $r - 1$ to c , that would imply that a had distance at most $r - 1$ to c , a contradiction. So e has distance at least $r - 1$ to c , and hence a' has distance at least r to c . Since a' is disjoint from $b \in S_{r-1}$ the distance from a' to c is at most r , so we get $a' \in S_r$. \square

5.3. The medium and high complexity cases. For convenience we continue to use the notation $S_r^c = S_r \cap \mathcal{C}_c \Sigma$. The second key proposition is the following.

Proposition 5.4. *In the curve complex $\mathcal{C}\Sigma$ with Σ medium or high complexity, for any c and any r we have the following.*

- (1) *For every $z \in S_r^c$ and $x, y \in S_{r+1}^c \cap B_1(z)$ there exists a path*

$$x = x_0, x_1, \dots, x_\ell = y$$

in the graph $\mathcal{C}_c \Sigma_{g,n}$ with

$$x_i \in (S_{r+1}^c \cup S_{r+2}^c) \cap B_2(z)$$

for $0 \leq i \leq \ell$. If Σ has high complexity, we can moreover obtain

$$x_i \in S_{r+1}^c \cap B_1(z)$$

for $0 \leq i \leq \ell$.

- (2) *For every adjacent pair $x, y \in S_r^c$ there exists a path*

$$x = x_0, x_1, x_2 = y$$

in $\mathcal{C}_c \Sigma_{g,n}$ with $x_1 \in S_{r+1}^c$.

We can now explain how Proposition 5.4 implies Theorem 1.1 in the medium and high complexity case.

Proof Theorem 1.1 in the medium and high complexity cases. Lemma 2.1 implies that the spheres S_r^c in the graph $\mathcal{C}_c \Sigma$ are connected in the high complexity case, and similarly for the union $S_r^c \cup S_{r+1}^c$ in the medium complexity case. Since Lemma 5.3 gives that every element of S_r is either in $\mathcal{C}_c \Sigma$ or adjacent to an element of S_r^c , this gives the result. \square

Remark 5.5. Roughly speaking, we will prove Proposition 5.4 by induction on the complexity of Σ . The proof for a given surface will make use of Theorem 1.1 for lower complexity surfaces. So it is important to note that really what we have just done is show that if Proposition 5.4 holds for a surface Σ , then Theorem 1.1 holds for that same Σ .

5.4. Linear connectivity. We conclude by explaining why the propositions imply linear connectivity.

Proof of Theorem 1.2. Fix $x \in \mathcal{C}\Sigma$. If $\Sigma = \Sigma_{g,n}$ has type $(g, n) \in \{(0, 5), (1, 2)\}$, we set $\Gamma = \mathcal{C}\Sigma$. Otherwise, we set $\Gamma = \mathcal{C}_c\Sigma$. Since Γ is 1-dense in $\mathcal{C}\Sigma$, the Gromov boundaries of Γ and $\mathcal{C}\Sigma$ are equal.

Recall that [Min10, Lemma 5.14] gives that any point of of the Gromov boundary can be represented by a geodesic ray starting at c in $\mathcal{C}\Sigma$. This can be easily modified to lie in Γ , for example using the observation in the proof of Lemma 5.2 first on the vertices of even index and then on the vertices of odd index. So assumption (1) in the statement of Proposition 2.2 is satisfied. Corollary 4.3 gives that assumption (2) is satisfied.

Propositions 5.1 and 5.4 show the remaining two assumptions are also satisfied, so Proposition 2.2 shows that Γ has linearly connected boundary. \square

6. THE FIVE TIMES PUNCTURED SPHERE

In this section, let $\Sigma = \Sigma_{0,5}$. The goal of this section is to prove Proposition 5.1.

6.1. Pentagons. We begin by recalling some basics on $\mathcal{C}\Sigma$ when $\Sigma = \Sigma_{0,5}$.

Lemma 6.1. *The curve complex $\mathcal{C}\Sigma$ does not contain cycles of length 3 or 4.*

Such cycles are called triangles and quadrilaterals respectively. We also note the following immediate consequence.

Corollary 6.2. *If x and y are vertices of $\mathcal{C}\Sigma$ with $d(x, y) \leq 2$, then there is a unique geodesic from x to y .*

Given the above, it is perhaps unsurprising that the geometry of $\mathcal{C}\Sigma$ is often studied using pentagons.

Definition 6.3. A pentagon in $\mathcal{C}\Sigma$ is a 5 tuple of curves $(a_1, a_2, a_3, a_4, a_5)$ such that the 5 punctures of Σ can be labeled by the elements of $\mathbb{Z}/5\mathbb{Z}$ in such a way that, for each i , the curve a_i

- (1) goes around punctures i and $i + 1$,
- (2) has intersection number 2 with a_{i-1} and with a_{i+1} , and
- (3) has intersection number 0 with a_{i-2} and with a_{i+2} .

Traversing these curves in the cyclic order $(a_1, a_3, a_5, a_2, a_4)$ gives a 5-cycle in the graph $\mathcal{C}\Sigma$. See Figure 1.

It is known that every 5-cycle in $\mathcal{C}\Sigma$ arises from a pentagon according to our definition [AL13, Theorem 3.1], but we do not require that result.

Warning 6.4. *We will call the tuple $(a_1, a_2, a_3, a_4, a_5)$ a pentagon, and the tuple $(a_1, a_3, a_5, a_2, a_4)$ a 5-cycle. Because of the different cyclic order, the terms pentagon and 5-cycle are not synonymous here.*

The following two lemmas produce pentagons that will be helpful for our analysis.

Lemma 6.5. *Suppose $a_1, a_3 \in S_{r-1}$ are disjoint and distinct. Then there are curves $a_2, a_4, a_5 \in S_r \cup S_{r+1}$ such that $(a_1, a_2, a_3, a_4, a_5)$ is a pentagon.*




FIGURE 1. The five curves in a pentagon (left), and the 5-cycle they define in $\mathcal{C}\Sigma$ (right).

Proof. Again, start with any $a_2, a_4, a_5 \in \mathcal{C}\Sigma$ such that $(a_1, a_2, a_3, a_4, a_5)$ is pentagon, and then replace it by its image under a large Dehn twist in a_1 and a large Dehn twist in a_3 to assume that

$$d_{a_1}(a_2, c) > M, \quad \text{and} \quad d_{a_3}(a_2, c) > M.$$

It suffices to show $a_2 \in S_{r+1}$, so consider a geodesic from a_2 to c . Theorem 4.1 gives that this geodesic must contain a curve that doesn't cut a_1 , and that it must contain a curve that doesn't cut a_3 .

First consider the possibility that the geodesic contains a single curve b that simultaneously fails to cut both a_1 and a_3 . The only such curves are a_1 and a_3 themselves, so without loss of generality assume $b = a_1$. We know $d(a_2, a_1) = 2$, and the only thing adjacent to both a_1 and a_2 is a_4 , so the geodesic must start a_2, a_4, a_1 . This gives the result.

Next assume the previous possibility does not occur. Say that, starting at a_2 , the geodesic first contains a curve b that doesn't cut a_1 and then contains a curve b' that doesn't cut a_3 , and neither of these curves are a_1 or a_3 . Note $d(a_2, b') \geq 2$. Since b' has distance 1 from $a_3 \in S_{r-1}$, we're done if $d(a_2, b') \geq 3$.

So assume $d(a_2, b') = 2$. Thus b comes immediately next to a_2 and hence is disjoint from both a_2 and a_1 . Thus $b = a_4$. Since b' is disjoint from $b = a_4$ and a_3 , we get that $b' = a_1$, a contradiction. \square

Lemma 6.6. *Suppose $a_3, a_4 \in S_r$ have intersection number 2 and are both disjoint from $a_1 \in S_{r-1}$. Then there are curves $a_2, a_5 \in S_r \cup S_{r+1}$ such that $(a_1, a_2, a_3, a_4, a_5)$ is a pentagon.*

Proof. Start with any a_2, a_5 such that $(a_1, a_2, a_3, a_4, a_5)$ is a pentagon, then replace it with its image under a large Dehn twist about a_1 to assume that

$$d_{a_1}(a_2, c) > M, \quad \text{and} \quad d_{a_1}(a_5, c) > M.$$

We wish to show $a_2 \in S_r \cup S_{r+1}$. Consider a geodesic from a_2 to c . Theorem 4.1 gives that this geodesic must contain a curve b that doesn't cut a_1 . First assume that $d(a_2, b) = 1$. In this case Corollary 6.2 gives that $b = a_4$, and we get $a_2 \in S_{r+1}$. Next assume that $d(a_2, b) \geq 2$. The fact that b doesn't cut $a_1 \in S_{r-1}$ guarantees that $d(c, b) \geq r - 2$. So we get $d(c, a_2) \geq r$ as desired. \square

6.2. First notes on spheres. The very first observation on spheres is that S_1 has no edges, and so is as disconnected as possible. In analyzing larger spheres, the following definitions will be useful, where we continue to assume that a center vertex c has been fixed.

Definition 6.7. A vertex $x \in \mathcal{C}\Sigma$ has unique backtracking if it has a unique neighbour y with $d(y, c) = d(x, c) - 1$.

Definition 6.8. A vertex $x \in \mathcal{C}\Sigma$ has no sidestepping if it does not have any neighbour y with $d(y, c) = d(x, c)$.

Definition 6.9. A vertex $x \in \mathcal{C}\Sigma$ is forward facing if it has unique backtracking and no sidestepping.

The absence of edges in S_1 immediately gives the following.

Lemma 6.10. *Every $x \in S_1$ is forward facing.*

As a warm up, we can also easily observe the following.

Lemma 6.11. *Every sphere S_r contains forward facing vertices.*

Proof. We assume $r > 1$. For any $y \in S_{r-1}$, pick x adjacent to y with $d_U(x, c)$ much larger than M , where U is the component of $\Sigma - y$ that isn't a three times punctured sphere. As noted in the proof of Corollary 4.3, we have that every geodesic from x to c goes through y , and in particular that $x \in S_r$. So x has unique backtracking.

To see that x has no sidestepping, suppose in order to obtain a contradiction that $z \in S_r$ is adjacent to y . Since $z \neq y$, we see that z cuts U . By the coarse Lipschitz property of subsurface projections, we see that $d_U(z, c)$ is approximately equal to $d_U(x, c)$, and in particular we get that $d_U(z, c) > M$. So any geodesic from z to c passes through y . Since $z \in S_r$ and $y \in S_{r-1}$, we see that actually z and y are adjacent. Thus x, y, z form a triangle, giving a contradiction. \square

Corollary 6.12. *For every r , the sphere S_r is disconnected.*

Proof. Lemma 6.11 gives that S_r has vertices that are not adjacent to any other vertex of S_r . So the graph S_r contains vertices with no edges. \square

We also make the following observations.

Lemma 6.13. *For any c , $S_1(c) \cup S_2(c)$ is connected.*

Proof. Consider two points in S_1 . Connect them with a path x_0, x_1, \dots, x_k in the Farey graph $\mathcal{C}(\Sigma - c)$ not passing through c .

For each $0 \leq i < n$, consider a 5 cycle $(x_i, y_i, z_i, x_{i+1}, c)$. The fact that S_1 does not have edges implies that $y_i, z_i \in S_2$. The path

$$x_0, y_0, z_0, x_1, y_1, z_1, \dots, x_{k-1}, y_{k-1}, z_{k-1}, x_k$$

is contained in $S_1 \cup S_2$ and proves the result. \square

Lemma 6.14. *Suppose $a \in S_r$ is forward facing. Then if $b, b' \in S_{r+1}$ are neighbours of a , then there is a path from b to b' in $(S_{r+1} \cup S_{r+2}) \cap B_2(a)$.*

Proof. The curve complex $\mathcal{C}(\Sigma - a)$ is the Farey graph, and removing any given vertex does not disconnect the Farey graph. So we can find a path $b = x_0, x_1, \dots, x_k = b'$ in $\mathcal{C}(\Sigma - a)$ not passing through the unique neighbour of a in S_{r-1} . All the vertices of this path are in S_{r+1} since a has no sidestepping.

Using Lemma 6.6, for each i we can find a 5-cycle $(x_i, y_i, z_i, x_{i+1}, a)$ with $y_i, z_i \in S_{r+1} \cup S_{r+2}$.

The path

$$x_0, y_0, z_0, x_1, y_1, z_1, \dots, x_{k-1}, y_{k-1}, z_{k-1}, x_k$$

lies in $(S_{r+1} \cup S_{r+2}) \cap B_2(a)$, proving the result. \square

Lemma 6.15. *For any c , $S_2(c) \cup S_3(c)$ is connected.*

Proof. Since every element of S_1 is forward facing, Lemma 6.14 gives that if $a \in S_1$, and $b, b' \in S_2$ are both neighbours of a , then there is a path from b to b' in $S_2 \cup S_3$.

Now, given two points of $S_2 \cup S_3$ that we wish to join with a path, we first join them by a path in $S_1 \cup S_2 \cup S_3$, which is possible by Lemma 6.13. Since S_1 does not have edges, this path cannot contain consecutive vertices in S_1 . The previous remark allows the path to be modified by replacing each entry of the path into S_1 by a path in $S_2 \cup S_3$. \square

6.3. The proof. Our key lemma is the following.

Lemma 6.16. *Suppose $a \in S_r$ and $b, b' \in S_{r+1} \cap B_1(a)$. Then there is a path $b, x_1, x_2, \dots, x_k, b'$ in \mathcal{CS} with*

- (1) $d(x_i, a) \in \{2, 3\}$,
- (2) $d(x_i, c) \geq r$,
- (3) $d(x_i, c) = r$ implies $d(x_i, a) = 2$, and the unique vertex adjacent to both x_i and a lies in S_{r-1} , and that x_i has unique backtracking, and
- (4) if a has unique backtracking then $d(x_i, c) = r$ also implies x_i has no sidestepping.

Proof. Consider a path x_1, \dots, x_k in $S_2(a) \cup S_3(a)$, starting next to b and ending next to b' , as must exist by Lemma 6.15. Without loss of generality, by applying a large power of the Dehn twist in a , we can assume $d_a(x_i, c)$ is much greater than M for all i . We will prove the lemma with this choice of path.

The first point is true by construction. For the second point, suppose that $d(x_i, c) < r$. By the Bounded Geodesic Image Theorem, a geodesic from x_i to c must pass through $B_1(a)$. But any point in the geodesic cannot be in $B_1(a)$ because $B_1(a) \subset S_{r-1} \cup S_r \cup S_{r+1}$. This proves the second point.

For the third point, we again use that any geodesic from x_i to c must contain a vertex z in $B_1(a)$; in fact z must be the vertex adjacent to x_i on the geodesic, and z must be in $B_1(a) \cap S_{r-1}$. Thus z lies on a geodesic from x_i to a of length 2. Since geodesics of length 2 are unique, there is only one possibility for z .

For the last point, suppose that y is a sidestep of x_i . There is a unique vertex z adjacent to x_i and a , and it is in S_{r-1} , so since $y \in S_r$ by assumption we get that $y \neq z$ and that y cuts a . By the coarse Lipschitz property of subsurface projections, we get that $d_a(y, c) > M$. The same analysis as in the previous paragraph shows that the geodesic from y to c must start with a curve in $S_{r-1} \cap B_1(a)$. Since a has unique backtracking, the curve z is the unique such curve, so we get that y is adjacent to z . But now y, x_i, z is a triangle, contradicting the fact triangles do not exist in \mathcal{CS} . \square

Lemma 6.17. *Suppose $a \in S_r$ and $b, b' \in S_{r+1}$ are neighbours of a . Then b and b' can be joined by a path in $(S_{r+1} \cup S_{r+2} \cup S_{r+3}) \cap B_6(a)$.*

Proof. We prove the lemma using a sequence of sublemmas.

Sublemma 6.18. *Suppose $a \in S_r$ is forward facing and $b, b' \in S_{r+1}$ are neighbours of a . Then b and b' can be joined by a path in $(S_{r+1} \cup S_{r+2} \cup S_{r+3}) \cap B_2(a)$.*

Proof. This is Lemma 6.14. (Here we don't need S_{r+3} .) \square

Sublemma 6.19. *Suppose $a \in S_r$ has unique backtracking and $b, b' \in S_{r+1}$ are neighbours of a . Then b and b' can be joined by a path in $(S_{r+1} \cup S_{r+2} \cup S_{r+3}) \cap B_4(a)$.*

Proof. Lemma 6.16 gives the existence of a path in $(S_r \cup S_{r+1} \cup S_{r+2} \cup S_{r+3}) \cap B_3(a)$ from b to b' , with the extra property that each vertex on the path in S_r is forward facing and is in $B_2(a)$. Since forward facing vertices in particular have no side stepping, the path does not have adjacent vertices in S_r . We can now replace the vertices in S_r with paths given by Sublemma 6.18 to obtain the result. \square

To complete the proof, note that Lemma 6.16 gives the existence of a path in $(S_r \cup S_{r+1} \cup S_{r+2} \cup S_{r+3}) \cap B_3(a)$ from b to b' , with the extra property that each vertex on the path in S_r has unique backtracking and is in $B_2(a)$.

Using Lemma 6.5 we can modify this path to get the additional assumption that no two adjacent vertices are in S_r , resulting in a path in $(S_r \cup S_{r+1} \cup S_{r+2} \cup S_{r+3}) \cap B_4(a)$. Sublemma 6.19 can then be used to modify the path to avoid the vertices in S_r . \square

Proof of Proposition 5.1. The first claim follows from Lemma 6.17, and the second claim from Lemma 6.5. \square

7. THE MEDIUM AND HIGH COMPLEXITY CASES

The goal of this section is to prove Proposition 5.4. Throughout this section, we assume Σ is medium or high complexity. We continue to use $\mathcal{C}_0\Sigma$ to denote the essentially non-separating curve graph, and $\mathcal{C}_c\Sigma$ to denote the result of adding the vertex c to $\mathcal{C}_0\Sigma$ as well as all edges between curves that are disjoint and have different distances to c (see Section 5.2 for details).

7.1. A note on the medium complexity case. To help motivate why the medium complexity case has a slightly different statement, we give the following observation, which will not be used.

Lemma 7.1. *If $\Sigma = \Sigma_{0,6}$ and c is a loop around a pair of punctures, then S_1^c is disconnected.*

Proof. Label the punctures by $\{1, 2, 3, 4, 5, 6\}$. Every element of $\mathcal{C}_c\Sigma$ is a loop around a pair of punctures, and we will call a loop around puncture i and puncture j an “ ij loop”.

Suppose c is a 12 loop. Note that every element of S_1^c is an ij loop with $i, j \in \{3, 4, 5, 6\}$. Note also that, in S_1^c , 34 loops can only be adjacent to 56 loops, 35 loops can only be adjacent to 46 loops, and 36 loops can only be adjacent to 45 loops. So S_1^c has at least three connected components. \square

7.2. The structure of the proof. For any $z \in S_r^c$, we define

$$\mathcal{O}(z) = \{a \in S_1(z) \cap \mathcal{C}_c \Sigma : d_U(a, c) > M\},$$

where U is the unique component of $\Sigma - z$ that isn't a pants.

Note that Theorem 4.1 implies that $\mathcal{O} \subset S_{r+1}^c$. We think of \mathcal{O} as the subset of $S_1(z)$ which is most obviously in S_{r+1}^c . In subsequent subsections we will prove the following two lemmas.

Lemma 7.2. *For $z \in S_r^c$, any $x \in S_1(z) \cap S_{r+1}^c$ can be connected to $\mathcal{O}(z)$ by a path in $S_1(z) \cap S_{r+1}^c$.*

Moreover, for any N , any point in $\mathcal{O}(z)$ can be connected by a path in $\mathcal{O}(z)$ to a point e with $d_U(e, c) > N$.

Lemma 7.3. *Suppose Theorem 1.1 holds for surfaces of smaller complexity than Σ . If Σ is high complexity, then $\mathcal{O}(z)$ is connected. If Σ is medium complexity, then any two points of $\mathcal{O}(z)$ can be joined by a path in $B_2(z) \cap (S_{r+1}^c \cup S_{r+2}^c)$.*

For both of these lemmas, it is implicit that we are working in (subgraphs of) the graph $\mathcal{C}_c \Sigma$. Before we give the proofs, we explain how the two lemmas imply Proposition 5.4.

Proof of Proposition 5.4. Suppose in order to find a contradiction that Proposition 5.4 is false, and let Σ be a minimal complexity surface for which it fails. Thus, Remark 5.5 gives that Theorem 1.1 can be assumed to hold for all surfaces of smaller complexity.

Lemmas 7.2 and 7.3 imply (1) holds.

To see that (2) holds, suppose $x, y \in S_r^c$ are adjacent. By definition of the graph $\mathcal{C}_c \Sigma$ and our complexity assumption on Σ , we know that there is a unique component V of $\Sigma - (x \cup y)$ that is not a pants.

By appealing to Lemma 3.8 and using that the mapping class group action on \mathcal{CV} has unbounded orbits, we can find a curve $x_1 \in \mathcal{CV}$ that is essentially non-separating and such that $d_V(x_1, c) > M$. Then Theorem 4.1 gives that every geodesic from z to c contains a curve not cutting V . The only curves not cutting V are x and y , so this shows that $x_1 \in S_{r+1}^c$. This completes the verification of (2), showing that in fact Proposition 5.4 actually does hold for Σ . \square

7.3. Getting into $\mathcal{O}(z)$. We now turn to the proof of Lemma 7.2, first giving another lemma on picking essentially non-separating curves.

Lemma 7.4. *Let Σ have medium or high complexity. Let $a, b \in \mathcal{C}\Sigma$ be disjoint and (individually) essentially non-separating. Then, there is a curve $d \in \mathcal{C}\Sigma$ such that d is disjoint from and not isotopic to a and b , and such that $a \cup d$ and $b \cup d$ are essentially non-separating. If V is a component of $\Sigma - (a \cup b)$ containing such a curve d , then the set of such d is coarsely dense in \mathcal{CV} .*

Proof. The last statement is immediate using the mapping class group action, so it suffices to find a single d . We proceed in cases, all of which are illustrated in Figure 2.

Case 1: a, b and $a \cup b$ are all non-separating. The existence of these curves implies $g \geq 2$, and we can pick d so that a, b and d bound a pants whose complement is connected.

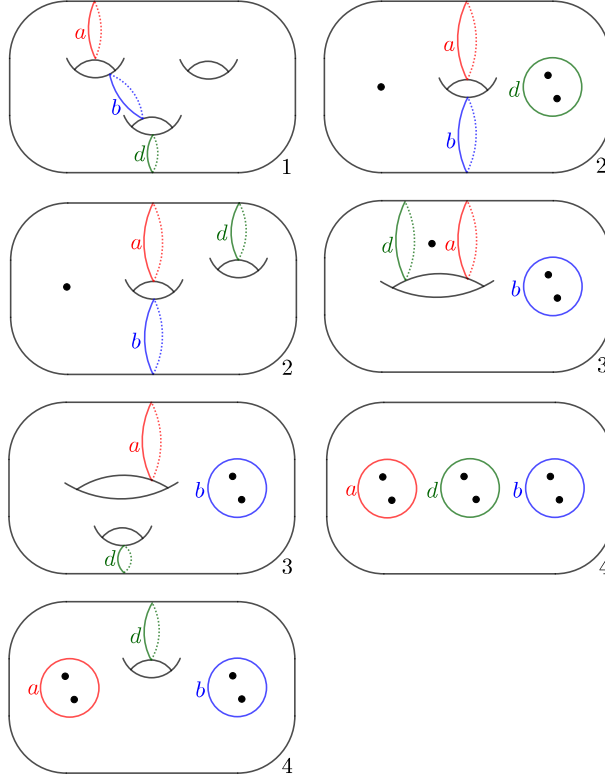


FIGURE 2. The proof of Lemma 7.4. The number below and to the right of each surface indicates which case it belongs to.

Case 2: a and b are both non-separating and $a \cup b$ is separating. If $m \geq 3$, then we can pick d to be a loop around two punctures on a component of $\Sigma - (a \cup b)$ with at least 2 punctures. If $m < 3$ then $g \geq 2$ and we can pick d to be a non-separating curve on a component of $\Sigma - (a \cup b)$ of positive genus.

Case 3: a is non-separating and b is a pants curve (or vice versa). If $m \geq 3$, we pick d so that a and d bound a once punctured annulus. If $m \leq 2$ then $g \geq 2$ and we can pick d so that d and $a \cup d$ are non-separating.

Case 4: a and b are both pants curves. If $m \geq 6$, we can pick d to also go around a pair of punctures. If $m \leq 5$ then $g \geq 1$ and we can pick d to be non-separating. \square

Proof of Lemma 7.2. Let $x_0 = x$, and let U continue to denote the unique component of $\Sigma - z$ that isn't a pants. We start with the following observation, whose notation has been chosen to match how it will be applied.

Sublemma 7.5. *Suppose $x_i \in S_{r+1}^c \cap B_1(z)$ and $z \in S_r^c$. Let V be a component of $\Sigma - (x_i \cup z)$, and suppose $x_{i+1} \in \mathcal{CV}$ satisfies $d_V(c, x_{i+1}) > M$. Then $x_{i+1} \in S_{r+1}$.*

Proof. First note that since x_i and z_i are essentially non-separating and \mathcal{CV} is non-empty, the boundary of V must contain both x_i and z . In particular, every curve

not cutting V is disjoint from (or equal to) x_i and z . Hence every curve not cutting V must lie in S_r or S_{r+1} .

Theorem 4.1 gives that the geodesic from x_{i+1} to c contains a curve not cutting V , giving the result. \square

Lemma 7.4 with $a = x_0$ and $b = z$ gives the existence of a curve x_1 disjoint from x_0 and z such that $x_0 \cup x_1$ and $x_1 \cup z$ are essentially non-separating. The coarse density statement in Lemma 7.4 and Sublemma 7.5 imply that we can pick x_1 to be in S_{r+1}^c as well.

We now note the following statement.

Sublemma 7.6. *For any $x_i \in S_{r+1} \cap \mathcal{CU}$ such that $z \cup x_i$ is essentially non-separating there exists $x_{i+1} \in S_{r+1} \cap \mathcal{CU}$ disjoint from x_i such that $x_i \cup x_{i+1}$ and $z \cup x_{i+1}$ are essentially non-separating and such that $d_U(x_{i+1}, x_0) = d_U(x_i, x_0) + 1$.*

Proof. Lemma 7.4 with $a = x_i$ and $b = z$ gives the existence of a curve x_{i+1} disjoint from x_i and z such that $x_i \cup x_{i+1}$ and $z \cup x_{i+1}$ are essentially non-separating. Furthermore, if V is the unique component of $\Sigma - (x_i \cup z)$ that isn't a pants, the coarse density statement in Lemma 7.4 implies we can assume that $d_V(x_{i+1}, x_0) > M$.

Theorem 4.1 implies that any geodesic from x_{i+1} to x_0 in \mathcal{CU} must pass through a curve in \mathcal{CU} not cutting V . The only such curve is x_i , so we conclude $d_U(x_{i+1}, x_0) = d_U(x_i, x_0) + 1$.

Sublemma 7.5 implies that $x_{i+1} \in S_{r+1}^c$. \square

Repeatedly using the sublemma we obtain a path x_0, x_1, x_2, \dots in $\mathcal{CU} \cap S_{r+1}^c$ which is a geodesic in \mathcal{CU} and hence eventually lies far from the projection of c to \mathcal{CU} .

To get the second statement, note that if x_0 is already in $\mathcal{O}(z)$ we can immediately use Sublemma 7.6 in the same way to get the desired result. \square

7.4. Connectivity of $\mathcal{O}(z)$. Our final task in this section is to prove Lemma 7.3. We start in high complexity.

Proof of Lemma 7.3 when Σ has high complexity. By Lemma 7.2, it suffices to consider two points $x, y \in \mathcal{O}(z)$ with $d_U(x, c), d_U(y, c) > M + C$, where C is a large constant, and show they can be connected by a path in $\mathcal{O}(z)$.

Using that Theorem 1.1 is true for U , we can start by finding a path

$$x = p_0, p_1, \dots, p_\ell = y$$

in \mathcal{CU} from x to y with $d_U(p_i, c) > M + C$ for all i .

Note that one of the following is true:

- (1) U has genus at least 2.
- (2) U has genus 1. If Σ has genus 1, then the high complexity assumption $\xi(\Sigma) \geq 4$ implies that Σ has at least 4 punctures, so U has at least 2 punctures. If Σ has genus 2, then Σ has at least 1 puncture, so U has a pair of boundary components and at least 1 puncture.
- (3) U has genus 0. If Σ has genus 0, it has at least 7 punctures, so U has at least 5 punctures. If Σ has genus 1, it has at least 4 punctures, so U has a pair of boundary components and at least 4 punctures.

Using Lemma 3.10, we can find an essentially non-separating $q_i \in \mathcal{CU}$ which is equal to or adjacent to p_i . (We pick $q_0 = x$ and $q_\ell = y$.)

One of Lemmas 3.12, 3.13, 3.15 applies to give that there is a path in the essentially non-separating curve graph from q_i to q_{i+1} that has bounded diameter in \mathcal{CU} . Assuming C is large enough, every element of this path is in $\mathcal{O}(z)$, so concatenating these paths gives the result. \square

The proof of the medium complexity case will be similar, but we first need the following.

Lemma 7.7. *Suppose $x_j, x_{j+1} \in \mathcal{O}(z)$ have distance much more than M to c . Suppose there exists a vertex w in the essentially non-separating curve graph $\mathcal{C}_0\Sigma$ not equal to z that is adjacent in $\mathcal{C}_0\Sigma$ to both x_j and x_{j+1} . Then $w \in S_{r+1}^c \cup S_{r+2}^c$, and w is adjacent in $\mathcal{C}_c\Sigma$ to x_j and x_{j+1} .*

Note that the assumption requires $w \cup x_j$ and $w \cup x_{j+1}$ to be essentially non-separating

Proof. The coarse Lipschitz property of subsurface projections implies that $d_U(c, w) > M$, so the geodesic from w to c must pass through z , proving the result. \square

Note that the medium complexity assumption that $(g, n) \in \{(2, 0), (1, 3), (0, 6)\}$ gives rise to the following possibilities for U in the notation of Definition 3.1:

$$(h, m, p, u) \in \{(1, 0, 1, 0), (1, 1, 0, 1), (0, 3, 1, 0), (0, 4, 0, 1)\}.$$

Recall that h denotes the genus of U , m the number of punctures, p the number of pairs of boundary components, and u the number of unpaired boundaries. In particular, U is always genus 0 or 1.

Definition 7.8. If U is genus 0, say $x_j, x_{j+1} \in \mathcal{C}_0U$ intersect nicely if they have intersection number 2. If U has genus 1, say $x_j, x_{j+1} \in \mathcal{C}_0U$ intersect nicely if they have intersection number 1. (This implies they are both non-separating.)

Corollary 7.9. *Suppose that Σ has medium complexity, and that $x_j, x_{j+1} \in \mathcal{C}_0U$ that are essentially non-separating and intersect nicely. Then there exists $w \in S_{r+1}^c \cup S_{r+2}^c$ adjacent to both x_j and x_{j+1} .*

Proof. The proof is illustrated in Figure 3.

First suppose $(g, n) = (2, 0)$. In this case let w be a non-separating curve disjoint from x_j, x_{j+1} that intersects z once.

Next suppose $(g, n) = (1, 3)$ and z is a pants curve. In this case let w be a pants curve disjoint from x_j, x_{j+1} that intersects z twice.

Next suppose $(g, n) = (1, 3)$ and z is non-separating. In this case, let w be a non-separating curve disjoint from x_j, x_{j+1} that intersects z once.

Finally suppose $(g, n) = (0, 6)$. In this case, let w be a pants curve disjoint from x_j, x_{j+1} that intersects z twice.

With this w as an input, Lemma 7.7 gives the result. \square

Proof of Lemma 7.3 when Σ has medium complexity. The proof starts as in the high complexity case, giving a path

$$x = p_0, p_1, \dots, p_\ell = y$$

in \mathcal{CU} from x to y with $d_U(p_i, c) > M + C$ for all i , and we still have the existence of an essentially non-separating curve $q_i \in \mathcal{CU}$ equal to or adjacent to p_i for all i .

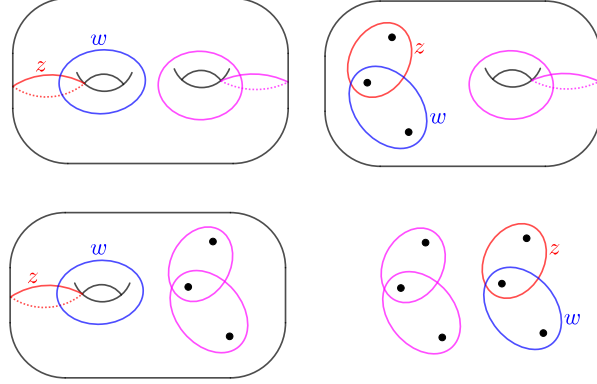


FIGURE 3. The proof of Corollary 7.9. The unlabelled curves are x_j and x_{j+1} .

(We pick $q_0 = x$ and $q_\ell = y$.) It suffices to show that for each fixed i , there is a path from q_i to q_{i+1} in $B_2(z) \cap (S_{r+1}^c \cup S_{r+2}^c)$.

If $h = 1$, then Lemma 3.11 gives the existence of a sequence

$$q_i = x_0, x_1, \dots, x_k = q_{i+1}$$

in \mathcal{C}_0U with all x_j , $0 < j < k$ non-separating, and $i(x_j, x_{j+1}) = 1$ for $0 < j < k - 2$. We can also assume that either x_0 is a pants curve and x_1 is disjoint from x_0 , or that x_0 is a non-separating curve and $i(x_0, x_1) = 1$, and similarly for x_k . Since Lemma 3.11 gives that the path has bounded diameter in $\mathcal{C}U$, we can assume that $d_U(x_i, c)$ is much greater than M . Now, Corollary 7.9 gives that when $i(x_j, x_{j+1}) = 1$ there is a $w \in S_{r+1}^c \cup S_{r+2}^c$ adjacent to both x_j and x_{j+1} . By interlacing w of this form, we get a path as desired.

If $h = 0$, first note that if $a, b \in \mathcal{C}_0U$ have intersection number 4 and go around the same pair of peripheries, then there is some $v \in \mathcal{C}_0U$ such that $i(a, v) = 2 = i(b, v)$. Also note that if a and b are disjoint, then there is some $v \in \mathcal{C}_0U$ such that $i(a, v) = 2 = i(b, v)$. This note and Lemma 3.14 gives the existence of a sequence

$$q_i = x_0, x_1, \dots, x_k = q_{i+1}$$

in \mathcal{C}_0U such that for each j the intersection number between x_j and x_{j+1} is 2. Again using Corollary 7.9 allows us to conclude. \square

REFERENCES

- [AL13] Javier Aramayona and Christopher J. Leininger, *Finite rigid sets in curve complexes*, J. Topol. Anal. **5** (2013), no. 2, 183–203.
- [Aou13] Tarik Aougab, *Uniform hyperbolicity of the graphs of curves*, Geom. Topol. **17** (2013), no. 5, 2855–2875.
- [BB19] Mladen Bestvina and Ken Bromberg, *On the asymptotic dimension of the curve complex*, Geom. Topol. **23** (2019), no. 5, 2227–2276.
- [BBM15] Joan Birman, Nathan Broaddus, and William Menasco, *Finite rigid sets and homologically nontrivial spheres in the curve complex of a surface*, J. Topol. Anal. **7** (2015), no. 1, 47–71. MR 3284389
- [BCH21] Mladen Bestvina, Jon Chaika, and Sebastian Hensel, *Connectivity of the Gromov boundary of the free factor complex*, arXiv:2105.01537 (2021).

- [BF14] Mladen Bestvina and Mark Feighn, *Subfactor projections*, J. Topol. **7** (2014), no. 3, 771–804.
- [BHS17] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804.
- [BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula*, Pacific J. Math. **299** (2019), no. 2, 257–338.
- [BK02] Mario Bonk and Bruce Kleiner, *Quasisymmetric parametrizations of two-dimensional metric spheres*, Invent. Math. **150** (2002), no. 1, 127–183.
- [BK05] ———, *Quasi-hyperbolic planes in hyperbolic groups*, Proc. Amer. Math. Soc. **133** (2005), no. 9, 2491–2494.
- [BM91] Mladen Bestvina and Geoffrey Mess, *The boundary of negatively curved groups*, J. Amer. Math. Soc. **4** (1991), no. 3, 469–481.
- [BM15] Joan S. Birman and William W. Menasco, *The curve complex has dead ends*, Geom. Dedicata **177** (2015), 71–74.
- [Bou97] M. Bourdon, *Immeubles hyperboliques, dimension conforme et rigidité de Mostow*, Geom. Funct. Anal. **7** (1997), no. 2, 245–268.
- [Bow14] Brian H. Bowditch, *Uniform hyperbolicity of the curve graphs*, Pacific J. Math. **269** (2014), no. 2, 269–280.
- [BS00] M. Bonk and O. Schramm, *Embeddings of Gromov hyperbolic spaces*, Geom. Funct. Anal. **10** (2000), no. 2, 266–306.
- [CH19] Jon Chaika and Sebastian Hensel, *Path-connectivity of the set of uniquely ergodic and cobounded foliations*, arXiv:1909.03668 (2019).
- [CRS14] Matt Clay, Kasra Rafi, and Saul Schleimer, *Uniform hyperbolicity of the curve graph via surgery sequences*, Algebr. Geom. Topol. **14** (2014), no. 6, 3325–3344.
- [DDM13] Spencer Dowdall, Moon Duchin, and Howard Masur, *Spheres in the curve complex*, In the tradition of Ahlfors-Bers. VI, Contemp. Math., vol. 590, Amer. Math. Soc., Providence, RI, 2013, pp. 1–8.
- [FM02] Benson Farb and Lee Mosher, *Convex cocompact subgroups of mapping class groups*, Geom. Topol. **6** (2002), 91–152.
- [Gab09] David Gabai, *Almost filling laminations and the connectivity of ending lamination space*, Geom. Topol. **13** (2009), no. 2, 1017–1041.
- [Gab14] ———, *On the topology of ending lamination space*, Geom. Topol. **18** (2014), no. 5, 2683–2745.
- [Ham05] Ursula Hamenstädt, *Word hyperbolic extensions of surface groups*, arXiv:0505244 (2005).
- [Har86] John L. Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. **84** (1986), no. 1, 157–176.
- [HKSS23] Helena Heinonen, Roshan Klein-Seetharaman, and Minghan Sun, *Optimal connectivity results for spheres in the curve graph of low and medium complexity surfaces*, arXiv:2309.10690 (2023).
- [HP11] Sebastian Hensel and Piotr Przytycki, *The ending lamination space of the five-punctured sphere is the Nöbeling curve*, J. Lond. Math. Soc. (2) **84** (2011), no. 1, 103–119.
- [HPW15] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, J. Eur. Math. Soc. (JEMS) **17** (2015), no. 4, 755–762.
- [KL08a] Richard P. Kent, IV and Christopher J. Leininger, *Shadows of mapping class groups: capturing convex cocompactness*, Geom. Funct. Anal. **18** (2008), no. 4, 1270–1325.
- [KL08b] ———, *Uniform convergence in the mapping class group*, Ergodic Theory Dynam. Systems **28** (2008), no. 4, 1177–1195.
- [Kla22] Erica Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space*, Groups Geom. Dyn. **16** (2022), no. 2, 705–723.
- [LMS11] Christopher J. Leininger, Mahan Mj, and Saul Schleimer, *The universal Cannon-Thurston map and the boundary of the curve complex*, Comment. Math. Helv. **86** (2011), no. 4, 769–816.

- [LS09] Christopher J. Leininger and Saul Schleimer, *Connectivity of the space of ending laminations*, Duke Math. J. **150** (2009), no. 3, 533–575.
- [LS14] ———, *Hyperbolic spaces in Teichmüller spaces*, J. Eur. Math. Soc. (JEMS) **16** (2014), no. 12, 2669–2692.
- [LT24] Yusen Long and Dong Tan, *Connectedness of the gromov boundary of fine curve graphs*, arXiv:2401.15383 (2024).
- [Mac08] John M. Mackay, *Existence of quasi-arcs*, Proc. Amer. Math. Soc. **136** (2008), no. 11, 3975–3981.
- [Mac12] ———, *Conformal dimension and random groups*, Geom. Funct. Anal. **22** (2012), no. 1, 213–239.
- [Man13] Johanna Mangahas, *A recipe for short-word pseudo-Anosovs*, Amer. J. Math. **135** (2013), no. 4, 1087–1116.
- [Min06] Yair N. Minsky, *Curve complexes, surfaces and 3-manifolds*, International Congress of Mathematicians. Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 1001–1033.
- [Min10] Yair Minsky, *The classification of Kleinian surface groups. I. Models and bounds*, Ann. of Math. (2) **171** (2010), no. 1, 1–107.
- [MM99] Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149.
- [MM00] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. **10** (2000), no. 4, 902–974.
- [Mos06] Lee Mosher, *Problems in the geometry of surface group extensions*, Problems on mapping class groups and related topics, Proc. Sympos. Pure Math., vol. 74, Amer. Math. Soc., Providence, RI, 2006, pp. 245–256.
- [MS20] John M. Mackay and Alessandro Sisto, *Quasi-hyperbolic planes in relatively hyperbolic groups*, Ann. Acad. Sci. Fenn. Math. **45** (2020), no. 1, 139–174.
- [Rei06] Alan W. Reid, *Surface subgroups of mapping class groups*, Problems on mapping class groups and related topics, Proc. Sympos. Pure Math., vol. 74, Amer. Math. Soc., Providence, RI, 2006, pp. 257–268.
- [RS11] Kasra Rafi and Saul Schleimer, *Curve complexes are rigid*, Duke Math. J. **158** (2011), no. 2, 225–246.
- [Sch11] Saul Schleimer, *The end of the curve complex*, Groups Geom. Dyn. **5** (2011), no. 1, 169–176.
- [Sis19] Alessandro Sisto, *What is a hierarchically hyperbolic space?*, Beyond hyperbolicity, London Math. Soc. Lecture Note Ser., vol. 454, Cambridge Univ. Press, Cambridge, 2019, pp. 117–148.
- [SS12] Lucas Sabalka and Dmytro Savchuk, *Submanifold projection*, arXiv:1211.3111 (2012).
- [Tay14] Samuel J. Taylor, *A note on subfactor projections*, Algebr. Geom. Topol. **14** (2014), no. 2, 805–821.
- [Tuk96] Pekka Tukia, *Spaces and arcs of bounded turning*, Michigan Math. J. **43** (1996), no. 3, 559–584.
- [Web15] Richard C. H. Webb, *Uniform bounds for bounded geodesic image theorems*, J. Reine Angew. Math. **709** (2015), 219–228.

Email address: alexmw@umich.edu