PRESCRIBED ARC GRAPHS

MICHAEL C. KOPRESKI

Abstract. Given a compact surface Σ with boundary and a relation Γ on $\pi_0(\partial\Sigma)$, we define the *prescribed arc graph* $\mathcal{A}(\Sigma,\Gamma)$ to be the full subgraph of the arc graph $\mathcal{A}(\Sigma)$ containing only classes of arcs between boundary components in Γ. We prove that $\mathcal{A}(\Sigma,\Gamma)$ is connected and infinite-diameter (if Σ is not the sphere with three boundary components), and classify when it is δ-hyperbolic: in particular, $\mathcal{A}(\Sigma,\Gamma)$ is δ-hyperbolic if and only if Γ is not bipartite except in some sporadic cases, where δ may be chosen uniformly.

1. Introduction, definitions and main results

Given a compact surface Σ with boundary, we recall the *arc graph* $\mathcal{A}(\Sigma)$, whose vertices are isotopy classes of essential simple arcs and whose edges are determined by disjointness up to isotopy. We propose the *prescribed arc graph* as variant of the arc graph, defined to be the full subgraph $\mathcal{A}(\Sigma,\Gamma) \subset \mathcal{A}(\Sigma)$ spanned by the isotopy classes of arcs between boundary components in Γ, a symmetric relation on $\pi_0(\partial\Sigma)$:

Definition 1.1. Let Σ be a compact, orientable surface with boundary. Given Γ a graph with $V(\Gamma) = \pi_0(\partial\Sigma)$, let an essential simple arc α be Γ-*allowed* if α terminates on (not necessarily distinct) boundary components a_-, a_+ such that $(a_-, a_+) \in E(\Gamma)$. Define the Γ-*prescribed arc graph* $\mathcal{A}(\Sigma,\Gamma)$ as follows:

$$V(\mathcal{A}(\Sigma,\Gamma)) = \{a \text{ an unoriented isotopy class of } \Gamma\text{-allowed arcs}\}$$
$$E(\mathcal{A}(\Sigma,\Gamma)) = \{(a,a') : a, a' \text{ have disjoint representatives}\}$$

We will call Γ the *prescribing graph* for $\mathcal{A}(\Sigma,\Gamma)$. We note that if Γ is the complete graph with loops on $\pi_0(\partial\Sigma)$, then $\mathcal{A}(\Sigma,\Gamma)$ is the usual arc graph.

Henceforth let Σ_g^b denote the compact, orientable surface with genus g and b boundary components. We prove some initial results concerning the geometry of $\mathcal{A}(\Sigma,\Gamma)$, including the following:

Theorem 1.2. *Assume that $\chi(\Sigma) \le -1, E(\Gamma) \ne \varnothing$, and $\Sigma \ne \Sigma_0^3$. Then $\mathcal{A}(\Sigma,\Gamma)$ is connected and has infinite diameter.*

Theorem 1.3. *Assume that $\chi(\Sigma) \le -1, E(\Gamma) \ne \varnothing$, and $\Sigma \ne \Sigma_0^3$. If $\chi(\Sigma) \ge -2$ or if $\Sigma = \Sigma_0^{n+1}$ and Γ is a n-pointed star then $\mathcal{A}(\Sigma,\Gamma)$ is δ-hyperbolic. Otherwise, $\mathcal{A}(\Sigma,\Gamma)$ is (uniformly) δ-hyperbolic if and only if Γ is not bipartite.*

Email address: `kopreski@math.utah.edu`.

Section 2 proves Theorem 1.2 outside of some sporadic low-complexity cases, as well as establishing an upper bound on distances $d(a, b)$ in $\mathcal{A}(\Sigma, \Gamma)$ in terms of intersection number $i(a, b)$. Section 3 shows that $\mathcal{A}(\Sigma, \Gamma)$ is uniformly δ-hyperbolic if Γ is not bipartite, again excluding some sporadic cases.

To complete the proof of Theorem 1.3, we appeal to the existence of disjoint *witness subsurfaces*. Following standard convention (see [MS13, Def. 5.1]), we define witnesses to be subsurfaces that are cut by every Γ-allowed arc:

Definition 1.4. An essential connected proper subsurface $W \subset \Sigma$ is a (Γ-)*witness* for $\mathcal{A}(\Sigma, \Gamma)$ if every Γ-allowed arc intersects W.

Remark 1.5. We also assume every witness W is compact and *non-simple*: that is, $W \ncong \Sigma_0^3$.

Section 4 proves (again, ignoring some sporadic cases) that Γ is bipartite if and only if there exist a pair of disjoint Γ-witnesses; the latter implies a quasi-isometric embedding of $\mathbb{Z}^2$. Finally, Section 5 addresses the sporadic cases missing from the preceding sections.

In the spirit of [MS13], we may expect that $\mathcal{A}(\Sigma, \Gamma)$ is δ-hyperbolic if and only if there does not exist any pair of disjoint Γ-witnesses. This hypothesis in fact holds true, although we do not prove it directly. From the results in Sections 3 and 4 along with the sporadic cases in Section 5, we conclude:

Theorem 1.6. *Assume that $\chi(\Sigma) \leq -1, E(\Gamma) \neq \varnothing$, and $\Sigma \neq \Sigma_0^3$. Then $\mathcal{A}(\Sigma, \Gamma)$ is (uniformly) δ-hyperbolic if and only if Σ does not admit two distinct, disjoint Γ-witnesses.*

We remark that it is possible to show Theorem 1.6 without using Section 3, at the cost of the uniformity of δ: the forward direction is proven in Section 4, and the reverse direction may be shown by proving that (excluding a sporadic case) whenever Σ does not admit distinct, disjoint Γ-witnesses, $\mathcal{A}(\Sigma, \Gamma)$ is quasi-isometric to a twist-free multicurve graph, in the sense of [Vok22]. Hence $\mathcal{A}(\Sigma, \Gamma)$ is a hierarchically hyperbolic space with respect to projection to witness subsurfaces; since it has no orthogonal coordinates, it is δ-hyperbolic [BHS21], where δ depends on the complexity of Σ.

We also note that in the case where Γ is loop-free, *i.e.* every Γ-allowed arc terminates on two distinct boundary components, the proofs of the above (and in particular, those in Section 2) simplify considerably.

1.1. **Motivation.** Although perhaps overshadowed in recent decades by the curve complex $\mathcal{C}(\Sigma)$, the arc graph $\mathcal{A}(\Sigma)$ has been an object of intrinsic interest in the classical study of surfaces of finite topological type (*i.e.* compact, with finitely many marked points) and their mapping class groups. For example, the arc complex of a finite type surface Σ triangulates its Teichmüller space $\mathcal{T}(\Sigma)$ [Har86], with natural coordinates arising from horodisks in (cusped) hyperbolic metrics [BE88]; earlier work [Mos83] uses $\mathcal{A}(\Sigma)$ to study conjugacy classes of the mapping class group $\mathrm{Mod}(\Sigma)$. More recently, given $Q \subset \pi_0(\partial\Sigma)$, [MS13] investigates $\mathcal{A}(\Sigma, Q)$, the full subgraph of $\mathcal{A}(\Sigma)$ spanned by arcs terminating on Q, and [DFV18] studies $\mathcal{A}_2(\Sigma, Q)$, likewise defined but excluding loops. Both are prescribed arc graphs $\mathcal{A}(\Sigma, \Gamma)$, where Γ is the complete graph on Q with or without loops respectively.

Nonetheless, the current work arises instead from combinatorial objects defined for surfaces of infinite topological type (*i.e.* with infinite genus or infinitely many boundary components or marked points, and whose fundamental group is not finitely generated). For such surfaces, we have that $\operatorname{diam}(\mathcal{C}(\Sigma)) = 2$, and likewise $\operatorname{diam}(\mathcal{A}(\Sigma)) = 2$ whenever the number of marked points or boundary components is infinite. A number of authors have proposed suitable (*e.g.* infinite diameter, connected, and δ-hyperbolic) combinatorial models in the infinite type setting on which $\operatorname{Mod}(\Sigma)$ acts continuously, in analogue to $\mathcal{C}(\Sigma)$ and $\mathcal{A}(\Sigma)$ in the finite type setting. Danny Calegari's original 2009 blog post on mapping class groups of infinite type surfaces [Cal09] defines the ray graph on $\mathbb{R}^2 \setminus \text{Cantor set}$, which was shown to be infinite diameter, connected, and δ-hyperbolic by Juliette Bavard in [Bav16]. More recently, we note the omnipresent arc graph defined by Fanoni, Ghaswala, and McLeay [FGM21] and the grand arc graph defined by Bar-Natan and Verberne [BNV22]. We pay special attention to the grand arc graph $\mathcal{G}(\Sigma)$, in which one considers only arcs between distinct sets in a certain $\operatorname{Mod}(\Sigma)$-invariant partition of maximal ends, the sense of the partial order on $\operatorname{Ends}(\Sigma)$ defined in [MR20].

One observes that, in the definitions above, the combinatorial model is made "sparse enough" by restricting which arcs in Σ are considered. We propose the prescribed arc graph as a similar modification for the usual arc graph $\mathcal{A}(\Sigma)$ in the finite type setting, both inspired and motivated by the combinatorial models for infinite type surfaces discussed above. For example, let Σ be a surface of infinite type. Then for suitable compact exhaustion by finite type subsurfaces $\Sigma_i \subset \Sigma$ and prescribing graphs Γ_i on $\pi_0(\partial\Sigma_i)$, $\mathcal{A}(\Sigma_i, \Gamma_i)$ is meant to "approximate" $\mathcal{G}(\Sigma)$. To start, one may prove that for appropriate (Σ_i, Γ_i), $\mathcal{A}(\Sigma_i, \Gamma_i)$ coarsely embeds into $\mathcal{G}(\Sigma)$, hence *e.g.* $\operatorname{asdim} \mathcal{G}(\Sigma) \geq \operatorname{asdim} \mathcal{A}(\Sigma_i, \Gamma_i)$, through which we aim to show that the asymptotic dimension of $\mathcal{G}(\Sigma)$ is infinite. More generally, direct limits of suitable prescribed arc graphs $\mathcal{A}(\Sigma_i, \Gamma_i)$ on compact subsurfaces Σ_i may offer a diverse collection of combinatorial models for infinite type surfaces; since the δ in Theorem 1.6 is uniform, such objects are δ-hyperbolic if each $\mathcal{A}(\Sigma_i, \Gamma_i)$ is δ-hyperbolic. We will discuss these arguments in a future work.

Finally, we recall the results of Aramayona and Valdez [AV18], which classify the connectedness, diameter, and δ-hyperbolicity for a broad class of "sufficiently invariant" subgraphs of the arc and curve graph $\mathcal{AC}(\Sigma)$, for Σ a surface of infinite type. Here, arcs are assumed to terminate on a set of marked points Π, and a subgraph $\mathcal{G} \subset \mathcal{AC}(\Sigma)$ is called *sufficiently invariant* if there exists a subset $P \subset \Pi$ such that $\mathcal{G}$ is preserved setwise by the subgroup of the relative mapping class group $\operatorname{Mod}(\Sigma, \Pi)$ that preserves P setwise. An extension of our results to arc graphs on infinite type surfaces Σ (*e.g.* those arising from direct limits of prescribed arc graphs $\mathcal{A}(\Sigma_i, \Gamma_i)$ on compact surfaces Σ_i) would likewise extend the work of [AV18], in the sense that the resulting arc subgraphs would be invariant only up to subgroups of $\operatorname{Mod}(\Sigma, \Pi)$ that fix a particular symmetric relation Γ on Π.

1.2. **Trivial cases and marked points.** If Σ is the closed annulus or disk, $\mathcal{A}(\Sigma, \Gamma)$ is either empty or a singleton. We will ignore these trivial cases and assume Σ is neither. If $E(\Gamma) = \varnothing$ or $\Sigma = \Sigma_0^3$, then $\mathcal{A}(\Sigma, \Gamma)$ is likewise empty or finite; we typically exclude these cases as well.

We note that the omission of marked points or punctures from the definition of the prescribed arc graph is one of convenience, and likewise the choice that arcs terminate on boundary components instead of marked points. In particular, without loss of generality we will assume every surface considered is without marked points: if Σ is a marked surface and Σ' is obtained by deleting a neighborhood of each marked point, then $\mathcal{A}(\Sigma) \cong \mathcal{A}(\Sigma')$ and the prescribed arc graph may be taken to be the corresponding full subgraph of $\mathcal{A}(\Sigma')$. Moreover, Definition 1.1 could just as well have been taken with respect to arcs between marked points, with $V(\Gamma)$ corresponding to marked points instead of boundary components. We will occasionally appeal to this viewpoint when it is more appropriate.

We also define some convenient notation:

Notation. Let α be an oriented essential arc between boundary components of a surface Σ. We denote by α_-, α_+ the boundary component containing the initial, resp. terminal endpoint of α. If a is an isotopy class of such arcs, then $a_\pm := \alpha_\pm$ for a choice of representative arc α. Let $\partial\alpha := \{\alpha_\pm\}$ and $\partial a := \{a_\pm\}$.

Remark 1.7. For simplicity, we typically conflate arcs and their isotopy classes in proofs where well defined. For example, while $a \in V(\mathcal{A}(\Sigma, \Gamma))$ denotes an isotopy class of Γ-allowed arcs, we may apply topological operations such as intersection or concatenation, or ascribe properties like transversality and minimal position, by (implicitly) choosing a representative $\alpha \in a$. Conversely, given a Γ-allowed arc δ, we may view $\delta \in V(\mathcal{A}(\Sigma, \Gamma))$ by (implicitly) passing to the isotopy class $[\delta]$.

We conclude this section with an important final definition and some initial discussion useful for the remaining work.

1.3. **A canonical inclusion.** Consider any subgraph $\Gamma' \subset \Gamma$; without loss of generality, assume $V(\Gamma') = V(\Gamma) = \pi_0(\partial\Sigma)$, else append singletons for the remaining vertices. If a is an isotopy class of Γ'-allowed arcs, then likewise it is Γ-allowed. Let $\pi : \mathcal{A}(\Sigma, \Gamma') \to \mathcal{A}(\Sigma, \Gamma)$ be the simplicial map obtained by the inclusion of $V(\mathcal{A}(\Sigma, \Gamma'))$ into $V(\mathcal{A}(\Sigma, \Gamma))$; since π is simplicial, it is 1-Lipschitz. Moreover, if Γ is loop-free, then every Γ-allowed arc $a \in V(\mathcal{A}(\Sigma, \Gamma))$ is non-separating, and in particular, if $(u, v) \in E(\Gamma')$ then there exists a Γ'-allowed arc between u, v disjoint from a:

Lemma 1.8. *Suppose that Γ is loop-free and $\Gamma' \subset \Gamma$ contains an edge. Then $\pi : \mathcal{A}(\Sigma, \Gamma') \to \mathcal{A}(\Sigma, \Gamma)$ is 1-coarsely surjective.* □

Finally, we prove a technical lemma of importance in Sections 2 and 3:

Lemma 1.9. *Suppose that $\ell_0 \subset \Gamma$ is a loop edge and $w \subset \partial\Sigma$ is the boundary component in ℓ_0. Then the vertex set $X_2 := \{a \in V(\mathcal{A}(\Sigma, \Gamma) : d(a, \pi\, \mathcal{A}(\Sigma, \ell_0)) > 1\}$ is independent (contains no adjacent vertices) in $\mathcal{A}(\Sigma, \Gamma)$ and contains only arcs with an annular complementary component containing w.*

Proof. It suffices to show that if $a \in X_2$, then a has an annular complementary component containing w: we note that any two such non-isotopic arcs must intersect. Let $\Sigma_\pm$ denote the abstract closure of the complementary component(s) of a. Since a is an arc, the gluing map $\Sigma_\pm \to \Sigma$ is π_1-injective, hence any essential arc $\delta \subset \Sigma_\pm$ terminating on $w \cap \partial\Sigma_\pm$ descends to an essential ℓ_0-allowed arc in Σ disjoint from a, and $d(a, \pi\,\mathcal{A}(\Sigma, \ell_0)) = 1$: no such δ exists. Hence if a terminates on w then $\Sigma_\pm$ is a collection of disks, a contradiction since Σ is not a disk or annulus; if a does not terminate on w then one component is an annulus containing w. $\square$

Definition 1.10. We will call the vertex set X_2 *exceptional.*

2. Connectedness of $\mathcal{A}(\Sigma, \Gamma)$

We recall the definition of a *unicorn arc* introduced in [HPW15]:

Definition 2.1. Let a, b be distinct isotopy classes of essential, simple arcs. Then a *unicorn* of a, b is (the isotopy class of) a concatenation $\alpha_0 * \beta_0$ of subarcs $\alpha_0 \subset \alpha \in a$ and $\beta_0 \subset \beta \in b$.

For $c \in \partial a$ and $d \in \partial b$, let $\mathscr{U}(a^c, b^d) \subset \mathcal{A}(\Sigma)$ denote the full subgraph spanned by a, b and unicorns with an initial subarc along a from c and a terminal subarc along b to d. By [HPW15], $\mathscr{U}(a^c, b^d)$ is connected.

Definition 2.2. Given distinct Γ-allowed arcs a, b, a Γ-*unicorn* of a, b is a unicorn that is Γ-allowed. The Γ-*unicorn subgraph* $\mathscr{U}(a, b) \subset \mathcal{A}(\Sigma, \Gamma)$ is the full subgraph spanned by Γ-unicorns of a, b.

In addition, we impose that $a, b \in \mathscr{U}(a, b)$; we will say that a unicorn is *proper* if $x \neq a, b$. Unlike in the usual arc graph, $\mathscr{U}(a, b)$ is not always connected. For example, since any proper unicorn between a, b has one endpoint in ∂a and the other in ∂b, if ∂a shares no edges with ∂b in Γ then no such arc is Γ-allowed: $\mathscr{U}(a, b)$ consists of two singletons. In fact this is the only case for which $\mathscr{U}(a, b)$ is disconnected:

Proposition 2.3. *Let $a, b \in V(\mathcal{A}(\Sigma, \Gamma))$. If ∂a and ∂b share an edge in Γ, then $\mathscr{U}(a, b)$ is connected. Moreover, $d(a, b) \leq i(a, b) + 1$, where i denotes the geometric intersection number.*

Proof. Let $x \in \mathscr{U}(a, b)$, hence x is Γ-allowed and $(x_-, x_+) \in E(\Gamma)$. Then $x \in \mathscr{U}_x = \mathscr{U}(a^{x_-}, b^{x_+}) \subset \mathscr{U}(a, b)$, since every unicorn $y \in \mathscr{U}(a^{x_-}, b^{x_+})$ has $\partial y = \partial x \in E(\Gamma)$ hence is Γ-allowed, and all considered graphs are full in $\mathcal{A}(\Sigma)$. Since every $\mathscr{U}_x$ is connected and contains a, $\mathscr{U}(a, b)$ is connected. Finally, $|\mathscr{U}(a, b)| \leq i(a, b) + 2$ by construction, hence $\operatorname{diam}(\mathscr{U}(a, b)) \leq i(a, b) + 1$. $\square$

Remark. In particular, we note that if $a^+ = b^+$, then $\mathscr{U}(a^{a_-}, b^{b_+}) \subset \mathscr{U}(a, b)$: that is, any simple arc obtained from a, b by oriented surgery is a Γ-unicorn.

Despite that some unicorn subgraphs are disconnected, we may now show that $\mathcal{A}(\Sigma, \Gamma)$ is connected if Γ is loop-free. In particular, it suffices that $\mathcal{A}(\Sigma, e)$ is connected for some edge $e \in \Gamma$: $\pi\,\mathcal{A}(\Sigma, e) \subset \mathcal{A}(\Sigma, \Gamma)$ is thus a connected, 1-dense subgraph by Lemma 1.8. Observing that for any edge e and $a, b \in V(\mathcal{A}(\Sigma, e))$ $\partial a = \partial b$, $\mathcal{A}(\Sigma, e)$ is connected by Proposition 2.3.

If Γ contains a loop edge ℓ_0, a similar argument *almost* applies: by Proposition 2.3 $\mathcal{A}(\Sigma, \ell_0)$ is connected, and by Lemma 1.9 $V(\mathcal{A}(\Sigma, \Gamma)) \setminus X_2$ lies within the connected 1-neighborhood of $\pi\, \mathcal{A}(\Sigma, \ell_0)$. However, to show that the exceptional vertices X_2 are not isolated we must defer to Section 2.1, where we prove the distance bound in Proposition 2.3 in general. This result will be useful in addition for our discussion in Section 3.

2.1. **An upper bound for distance.** We will prove that except in certain low-complexity cases, $d(a, b) \leq i(a, b) + 1$, extending the result of Proposition 2.3 for generic arcs. We begin with arcs with one or two intersections.

Lemma 2.4. *Assume that* $\Sigma \neq \Sigma_0^3$ *and if* $\Sigma = \Sigma_1^2$ *then* Γ *is not two disjoint loop edges. If* a, b *are unoriented isotopy classes of* Γ*-allowed arcs and* $i(a, b) = 1$*, then* $d(a, b) = 2$ *in* $\mathcal{A}(\Sigma, \Gamma)$.

Proof. Assume a, b are in minimal position. If $\partial a \cap \partial b \neq \varnothing$ then $\partial a, \partial b$ share an edge in Γ, and likewise by assumption if $\Sigma = \Sigma_1^2$, since then either Γ is a single loop edge or contains a non-loop edge. Hence applying Proposition 2.3, $d(a, b) \leq i(a, b) + 1 = 2$. We suppose $\partial a \cap \partial b = \varnothing$ and $\Sigma \neq \Sigma_1^2$ and consider three cases:

- *(i) Neither* a *nor* b *are loops (i.e.* $a_+ \neq a_-, b_+ \neq b_-$*).* Let N be a regular neighborhood of $a \cup b \cup b_\pm$, and let δ be a component of ∂N.[1] Then δ is essential and disjoint from a, b; since $\delta_\pm = a_\pm$, δ is Γ-allowed.
- *(ii) (Without loss of generality)* a *is a loop but* b *is not (*$a_+ = a_-, b_+ \neq b_-$*).* We proceed as above, and let δ_1, δ_2 denote the components ∂N. If both bound half-disks then $\Sigma = \Sigma_0^3$, a contradiction. Otherwise, at least one is essential and we conclude as above.
- *(iii) Both* a *and* b *are loops (*$a_+ = a_-, b_+ = b_-$*).* Let N be a regular neighborhood of $a \cup b \cup b_+$. Since a, b intersect once, $N \cup N(a_+)$ is a thrice-punctured torus; let $\delta = \partial N$. If δ is essential, then we conclude as above; else δ bounds a half-disk and $\Sigma = \Sigma_1^2$.

It follows that $d(a, b) = 2$. □

Lemma 2.5. *Assume that* $\Sigma \neq \Sigma_0^3$ *and if* $\Sigma = \Sigma_1^2$ *then* Γ *is not two disjoint loop edges. If* a, b *are unoriented isotopy classes of* Γ*-allowed arcs such that* $i(a, b) = 2$*, then* $d(a, b) \leq 3$ *in* $\mathcal{A}(\Sigma, \Gamma)$.

Proof. Assume a, b are in minimal position and that $\partial a \cap \partial b = \varnothing$, else conclude by Proposition 2.3 as above. Let $\{s, t\} = a \cap b$ and let α_0, β_0 denote the subarcs of a, b respectively between s and t; orient a such that s is nearer a_-, and likewise with b. Let $\alpha^\pm \subset a \setminus \mathring{\alpha}_0$ denote the subarc incident on $a_\pm$, and likewise for $\beta^\pm \subset b \setminus \mathring{\beta}_0$. We consider two cases:

[1] Remark: ∂N denotes the *topological* boundary, hence does not contain $b_\pm$ or any subarcs of $a_\pm$.

(i) Not both a, b are loops, or $\hat{\imath}(a,b) = \pm 2$, where $\hat{\imath}$ denotes the algebraic intersection number. Without loss of generality, assume that if only one of a, b is a loop, then it is a. Let δ be the concatenation of β^-, α_0, and β^+. We observe that δ and b are disjoint up to isotopy and $i(\delta, a) \leq 1$, with equality when $\hat{\imath}(a,b) = \pm 2$. Hence if δ is essential the claim is shown: since $\partial\delta = \partial b$, δ is Γ-allowed and $d(\delta, a) \leq 2$ by Lemma 2.4, hence $d(a,b) \leq d(a,\delta) + d(\delta, b) \leq 3$. Finally, to verify that δ is essential, we observe that either b (hence δ) is not a loop or both a, b are loops and $\hat{\imath}(a,b) = \pm 2$, hence $i(\delta, a) = 1$. In both cases, it follows that δ is non-separating.

(ii) Both a, b are loops and $\hat{\imath}(a,b) = 0$. Let N_a be a regular neighborhood of $a \cup b \cup b_\pm$ and N_b a regular neighborhood of $a \cup b \cup a_\pm$. Let $\mathscr{D}_a$ be the set of arc components of ∂N_a, and likewise $\mathscr{D}_b$ for ∂N_b. For any $\rho \in \mathscr{D} := \mathscr{D}_a \cup \mathscr{D}_b$, ρ is disjoint from a, b and either $\partial\rho \in \partial a$ or $\partial \rho \in \partial b$. Hence if ρ is essential, then ρ is Γ-allowed and $d(a,b) = 2$.

Assume every arc in $\mathscr{D}$ bounds a half-disk. Then (*e.g.* by an Euler characteristic argument) the simple closed curve $\alpha_0 \cup \beta_0$ separates $\Sigma = \Sigma_0^3 \cup_{\alpha_0 \cup \beta_0} \Sigma'$, where Σ' is the subsurface with boundary $\alpha_0 \cup \beta_0$ not containing $a_\pm$ and $b_\pm$. If Σ' is a disk, then a, b are not in minimal position, and if Σ' is an annulus, then $\Sigma = \Sigma_0^3$, both contradictions with our assumptions. Hence there exists a simple arc $\gamma \subset \Sigma'$ between $s, t \in \partial\Sigma'$ that does not bound a half-disk in Σ'. Let $\alpha' = \alpha^- * \gamma * \alpha^+$ and $\beta' = \beta^- * \gamma * \beta^+$. If α' bounds a half-disk in Σ, then likewise does γ in Σ', and similarly with β': both α' and β' are essential. Since $\partial\alpha' = \partial a$ and $\partial\beta' = \partial b$, α', β' are Γ-allowed, and we observe that a and α', α' and β', and β' and b are disjoint up to isotopy. $d(a,b) \leq 3$. $\square$

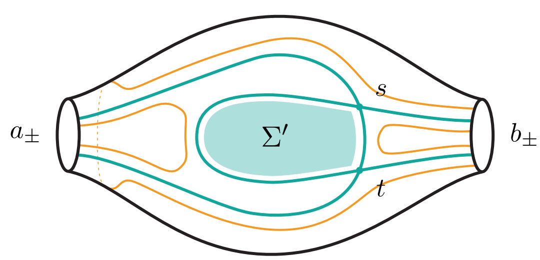

FIGURE 1. A maximal collection of disjoint arcs (orange, thin) in $\mathscr{D}$, each bounding a half disk.

For $i(a,b) > 2$, we construct a sequence of arcs with strictly decreasing intersection number and proceed by induction.

Proposition 2.6. *Assume that $\Sigma \neq \Sigma_0^3$ and if $\Sigma = \Sigma_1^2$ then Γ is not two disjoint loop edges. If a, b are unoriented isotopy classes of Γ-allowed arcs and $i(a,b) = n$, then $d(a,b) \leq n+1$ in $\mathcal{A}(\Sigma, \Gamma)$.*

Proof. As above, we may assume that $\partial a \cap \partial b = \varnothing$, else conclude by Proposition 2.3. Assume a, b are in minimal position. It suffices to find a Γ-allowed arc δ disjoint from b such that $i(\delta, a) < i(a, b)$, or likewise exchanging a and b; we then conclude by induction on intersection number, noting that the cases $i(a, b) \leq 2$ follow from Lemmas 2.4 and 2.5. Assume $i(a, b) \geq 3$ and let $s \in a \cap b$ be the first intersection along b; let t be the subsequent intersection along a. Let α_0, β_0 denote the subarcs of a, b respectively between s and t. Let $\beta^{\pm} \subset b \setminus \mathring{\beta}_0$ denote the subarc incident on $b_{\pm}$. We consider two cases:

(i) *s, t are intersections of the same sign.* Let $\delta = \beta^- * \alpha_0 * \beta^+$. Up to isotopy δ and b are disjoint and $i(\delta, a) \leq i(a, b) - 1$. Since $\delta_{\pm} = b_{\pm}$, if δ is essential then it is Γ-allowed. If b is not a loop, then δ is not a loop and hence essential. Suppose that b is a loop. Isotoping δ to be disjoint from b, we observe that the endpoints of δ separate the endpoints of b on $b_{\pm}$: δ is non-separating and hence essential.

(ii) *s, t are intersections of opposite sign.* Let δ as above; δ and b are disjoint up to isotopy and $i(\delta, a) \leq i(a, b) - 2$, hence if δ is essential, then we conclude. Suppose instead that δ bounds a half-disk D, hence δ and b are loops. Then since $\beta^- \cap a = \{s\}$, we have $\beta^+ \cap a = \{t\}$, else a, b share a bigon and are not in minimal position. Let N be a regular neighborhood of $a \cup D \cup b_{\pm}$, and let ξ be the arc component of ∂N not parallel to a. We observe that ξ is disjoint from a and that $i(\xi, b) \leq i(a, b) - 2$; $\partial \xi = \partial a$, hence if ξ is essential, then ξ is Γ-allowed. Suppose that ξ bounds a half-disk, and observe that b intersects ξ only on subarcs parallel to a. Hence b and ξ are disjoint, else b shares a bigon or half bigon with ξ incident only on subarcs in $\xi \cap a$, thus also with a. Since $a \setminus \mathring{\alpha}_0$ is parallel to ξ, a intersects b only at s, t, a contradiction since $i(a, b) \geq 3$. $\square$

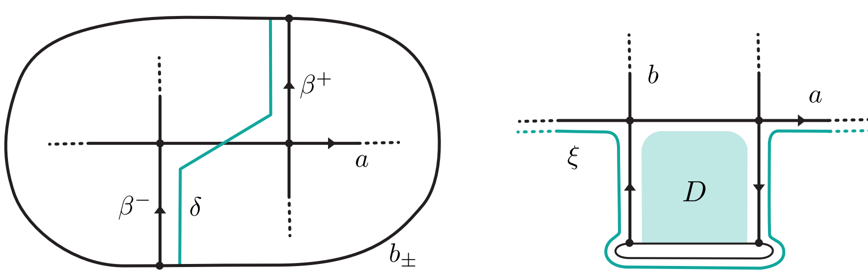

(a) (b)

FIGURE 2. (a) δ separates the endpoints of b in case *(i)*. (b) the arc ξ in case *(ii)* when δ bounds a half disk.

Remark. Note that if Γ is loop-free, then in the statement of the proposition and the two preceding lemmas we may omit our assumptions on (Σ, Γ).

Connectivity of $\mathcal{A}(\Sigma, \Gamma)$ follows as an immediate corollary:

Theorem 2.7. *Assume that either Γ is loop-free, or $\Sigma \neq \Sigma_0^3$ and if $\Sigma = \Sigma_1^2$ then Γ is not two disjoint loop edges. $\mathcal{A}(\Sigma, \Gamma)$ is connected.* $\square$

Remark. The connectedness of $\mathcal{A}(\Sigma,\Gamma)$ allows us to elaborate Lemma 1.9. Suppose $\ell_0 \subset \Gamma$ is a loop edge and $X_2 \subset V(\mathcal{A}(\Sigma,\Gamma))$ denotes the set of exceptional vertices not in the 1-neighborhood of $\pi\,\mathcal{A}(\Sigma,\ell_0)$, and additionally assume $\Sigma \neq \Sigma_0^3$ and if $\Sigma = \Sigma_1^2$ then Γ is not two disjoint loop edges. We observe that for any $v \in X_2$, $d(v,\pi\,\mathcal{A}(\Sigma,\ell_0)) = 2$. In particular, since X_2 has no adjacent vertices in $\mathcal{A}(\Sigma,\Gamma)$ by Lemma 1.9 and $\mathcal{A}(\Sigma,\Gamma)$ is connected, v must be adjacent to $a \notin X_2$, which lies in the 1-neighborhood of $\pi\,\mathcal{A}(\Sigma,\ell_0)$ by the definition of X_2. //

Proposition 2.6 also allows us to generalize Lemma 1.8:

Lemma 2.8. *Suppose that $\Gamma' \subset \Gamma$ contains an edge, and let $\pi : \mathcal{A}(\Sigma,\Gamma') \to \mathcal{A}(\Sigma,\Gamma)$ denote the simplicial map induced by the inclusion $\Gamma' \hookrightarrow \Gamma$. Then:*

(i) if Γ is loop-free, then π is 1-coarsely surjective.

Moreover, if $\Sigma \neq \Sigma_0^3$ and $\Sigma = \Sigma_1^2$ only if Γ is not two disjoint loop edges, then:

(ii) if Γ' has a non-loop edge, then π is 2-coarsely surjective; and otherwise

(iii) π is 3-coarsely surjective.

Proof. (i) is Lemma 1.8. By change of coordinates, for any $a \in V(\mathcal{A}(\Sigma,\Gamma)) \setminus V(\pi\,\mathcal{A}(\Sigma,\Gamma'))$, there exists a Γ'-allowed arc ρ that intersects a at most once if ρ is not a loop, and at most twice if ρ is a loop. We conclude by Proposition 2.6. $\square$

We note that Γ includes into the complete graph with loops on $\pi_0(\partial\Sigma)$, hence the induced map $\pi : \mathcal{A}(\Sigma,\Gamma) \to \mathcal{A}(\Sigma)$ is a coarse surjection by Lemma 2.8. We obtain:

Theorem 2.9. *Assume $\Sigma \neq \Sigma_0^3$. If $\mathcal{A}(\Sigma)$ has infinite diameter, then likewise does $\mathcal{A}(\Sigma,\Gamma)$.* $\square$

Remark 2.10. While Lemma 2.8 implies that $\mathcal{A}(\Sigma,\Gamma)$ lies canonically as a coarsely dense subset of $\mathcal{A}(\Sigma)$, it is usually significantly distorted: we show that π is typically not a quasi-isometric embedding in Section 4.1.

3. (Σ,Γ) for which $\mathcal{A}(\Sigma,\Gamma)$ is hyperbolic

In the following, we assume that $\Sigma \neq \Sigma_0^3$ and that if $\Sigma = \Sigma_1^2$, then Γ contains a non-loop edge. To show the hyperbolicity of $\mathcal{A}(\Sigma,\Gamma)$, we apply the following proposition of [Bow14]:

Theorem 3.1 (Guessing geodesics lemma)**.** *Suppose that Ω is a connected simplicial graph and there exists $\{\mathscr{U}_{a,b}\}_{a,b\in V(\Omega)}$ a family of connected subgraphs such that $a,b \in \mathscr{U}_{a,b}$ and $\Delta \geq 0$ such that*

(I) for $a,b \in V(\Omega)$, if $d(a,b) \leq 1$, then $\operatorname{diam}\mathscr{U}_{a,b} \leq \Delta$, and

(II) for $a,b,c \in V(\Omega)$, $\mathscr{U}_{a,c} \subset N_\Delta(\mathscr{U}_{a,b} \cup \mathscr{U}_{b,c})$.

Then Ω is δ-hyperbolic for some $\delta = \delta(\Delta) \geq 0$.

By Proposition 2.3, for $a,b \in V(\mathcal{A}(\Sigma,\Gamma))$, the Γ-unicorn subgraph $\mathscr{U}(a,b)$ is connected only when ∂a and ∂b share and edge in Γ. Our subgraphs will be chosen instead to be *augmented unicorn subgraphs* $\mathscr{U}^+(a,b)$ which lie within a uniformly bounded distance from some connected Γ-unicorn subgraph $\mathscr{U}(a',b')$.

We proceed by first proving δ-hyperbolicity in the case when Γ contains a loop edge, and then the loop-free case when Γ contains an odd cycle.

Proposition 3.2. *Assume that $\Sigma \neq \Sigma_0^3$ and if $\Sigma = \Sigma_1^2$ then Γ is not two disjoint loop edges. If Γ contains a loop edge, then $\mathcal{A}(\Sigma, \Gamma)$ is uniformly δ-hyperbolic.*

Proposition 3.3. *If Γ is loop-free and contains an odd cycle, then $\mathcal{A}(\Sigma, \Gamma)$ is uniformly δ-hyperbolic.*

Observing that Γ is bipartite if and only if it is loop-free and contains no odd cycles, we conclude the following:

Theorem 3.4. *Assume that either Γ is loop-free, or $\Sigma \neq \Sigma_0^3$ and if $\Sigma = \Sigma_1^2$ then Γ is not two disjoint loop edges. If Γ is not bipartite, then $\mathcal{A}(\Sigma, \Gamma)$ is uniformly δ-hyperbolic.* □

3.1. Γ **contains a loop edge.** We first consider the case when Γ is a single loop edge, in which case all Γ-allowed arcs share a boundary component and the usual unicorn subgraphs suffice.

Lemma 3.5. *If Γ is comprised of a single, looped edge, then $\{\mathscr{U}(a, b)\}_{a,b}$ form 1-slim triangles in $\mathcal{A}(\Sigma, \Gamma)$.*

Proof. Let $x \in \mathscr{U}(a, b)$, and suppose x is comprised of subarcs $\alpha' \subset a, \beta' \subset b$ with $x_- = \alpha'_- = a_-$ and $x_+ = \beta'_+ = b_+$. Let $c \in V(\mathcal{A}(\Sigma, \Gamma))$, and assume c is in minimal position with x.

It suffices to find $y \in \mathscr{U}(a, c) \cup \mathscr{U}(b, c)$ disjoint from x. Without loss of generality suppose that c last intersects x at $s \in \alpha'$. Let y be the concatenation of the subarc from x_- to s along a and the subarc from s to c_+ along c. Since c and x are in minimal position and $x_+, c_+ = b_\pm$, by Remark 2 $y \in \mathscr{U}(a, c)$ is Γ-allowed. y is disjoint with x up to isotopy. □

We note that if $d(a, b) = 1$ then a, b are disjoint, hence $\mathscr{U}(a, b) = \{a, b\}$ has diameter 1 and Theorem 3.1 is satisfied for $\Delta = 1$.

Corollary 3.6. *If Γ is comprised of a single, looped edge, then $\mathcal{A}(\Sigma, \Gamma)$ is δ-hyperbolic.* □

In fact, it suffices that Γ contain a loop edge. We assume that $\Sigma \neq \Sigma_0^3$ and if $\Sigma = \Sigma_1^2$, then Γ is not two disjoint loop edges. Let $\ell_0 \subset \Gamma$ be a loop edge in Γ and recall that by Lemma 1.9 the set of exceptional vertices X_2 are independent in $\mathcal{A}(\Sigma, \Gamma)$. Let $\mathcal{A}'$ be the metric graph obtained from $\mathcal{A}(\Sigma, \Gamma) \setminus X_2$, with the usual graph metric, by adjoining edges of length 2 between all neighbors of a vertex v for each $v \in X_2$. Since $\mathcal{A}'$ and $\mathcal{A}(\Sigma, \Gamma)$ are isometric outside of a collection of disjoint uniformly bounded subsets, they are uniformly quasi-isometric. We note that $\pi\,\mathcal{A}(\Sigma, \ell_0)$ embeds isometrically as a 1-dense subgraph in $\mathcal{A}'$.

We prove that $\mathcal{A}'$ is δ-hyperbolic. Let $\xi : \mathcal{A}' \to \pi\,\mathcal{A}(\Sigma, \ell_0)$ be a choice of nearest point projection. For $a, b \in V(\mathcal{A}(\Sigma, \Gamma))$, let

$$\mathscr{U}^+(a, b) := [a, \xi a] \cup \mathscr{U}_{\ell_0}(\xi a, \xi b) \cup [\xi b, b]$$

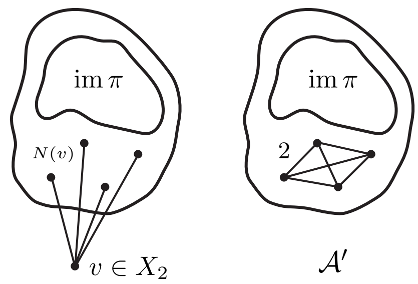

FIGURE 3. Obtaining $\mathcal{A}'$ from $\mathcal{A}(\Sigma, \Gamma)$.

where *e.g.* $[a, \xi a]$ denotes an edge and $\mathscr{U}_{\ell_0}$ denotes the π-image in $\mathcal{A}'$ of the respective ℓ_0-unicorn subgraph. Since π is a contraction, Lemma 3.5 implies that triangles in $\{\mathscr{U}_{\ell_0}(\xi a, \xi b)\}_{a,b}$ are likewise 1-slim, and thus for the augmented unicorns $\mathscr{U}^+(a, b)$ as well. To apply Theorem 3.1, we need only check that (I) is satisfied:

Lemma 3.7. *For any disjoint $a, b \in V(\mathcal{A}')$,* $\operatorname{diam} \mathscr{U}^+(a, b) \leq 10$.

Proof. Let w denote the boundary component in ℓ_0 and $a' = \xi a, b' = \xi b$. Since $d(a, a'), d(b, b') \leq 1$, a, a' and b, b' are disjoint up to isotopy. If both $a, b \in \pi\, \mathcal{A}(\Sigma, \ell_0)$, then $a' = a$ and $b' = b$ are disjoint: $\mathscr{U}_{\ell_0}(a', b') = \{a, b\}$. Assume not, and without loss of generality, let $b \notin \pi\, \mathcal{A}(\Sigma, \ell_0)$. Then $\partial b \neq \{w\}$; orient b such that $b_- \neq w$. It suffices to show that for any $x \in \mathscr{U}_{\ell_0}(a', b')$, $d(x, a) \leq 5$ in $\mathcal{A}'$.

Let $s \in b \cap x$ be the first intersection along b; since b, b' are disjoint, $s \subset a'$. Let b_0 denote the subarc of b between b_- and s, and let a'_0 denote the subarc of x along a' between $a'_\pm$ and s. Let N be a regular neighborhood of $a'_0 \cup b_0 \cup b_-$. Then $\delta = \partial N$ intersects x at most once and must be essential, else Σ is an annulus: $d(x, \delta) \leq 2$. Since $\partial\delta = \{w\}$, δ is Γ-allowed and by Lemma 1.9 does not lie in X_2: $\delta \in V(\mathcal{A}')$. Finally, since a' is disjoint from a, δ intersects a only if a terminates on b_-, hence at most twice. Applying Proposition 2.6, $d(a, \delta) \leq 3$, hence $d(x, a) \leq 5$ as required. □

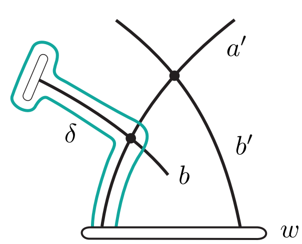

FIGURE 4. The arc δ in the proof of Lemma 3.7.

The proof of Proposition 3.2 follows from Theorem 3.1 and that $\mathcal{A}(\Sigma, \Gamma) \simeq_{\text{q.i.}} \mathcal{A}'$.

3.2. Γ **contains an odd cycle.** Assume that Γ is loop-free and contains an odd cycle C_n of length $n = 2k+1$. Fix $e_0 \in E(C_n)$. By Lemma 2.8, the canonical map $\pi : \mathcal{A}(\Sigma, e_0) \to \mathcal{A}(\Sigma, \Gamma)$ is 1-coarsely surjective. As above, let $\xi : \mathcal{A}(\Sigma,\Gamma) \to \pi\,\mathcal{A}(\Sigma, e_0)$ be a choice of nearest point projection, and for $a, b \in V(\mathcal{A}(\Sigma,\Gamma))$, let

$$\mathscr{U}^+(a,b) := [a, \xi a] \cup \mathscr{U}_{e_0}(\xi a, \xi b) \cup [\xi b, b].$$

We verify first that triangles in $\{\mathscr{U}_{e_0}(a,b)\}_{a,b\in V(\pi\,\mathcal{A}(\Sigma,e_0))}$ are slim in $\mathcal{A}(\Sigma, C_n)$; since π factors through the 1-Lipschitz induced map $\mathcal{A}(\Sigma, C_n) \to \mathcal{A}(\Sigma,\Gamma)$, the augmented unicorns $\mathscr{U}^+(a,b)$ likewise form slim triangles in $\mathcal{A}(\Sigma,\Gamma)$, satisfying (II) of Theorem 3.1.

Lemma 3.8. *Triangles in $\{\mathscr{U}_{e_0}(a,b)\}_{a,b\in V(\pi\,\mathcal{A}(\Sigma,e_0))}$ are 4-slim in $\mathcal{A}(\Sigma, C_n)$.*

Proof. Let $e_0 = (w_1, w_n)$, and let w_i denote the remaining boundary components in C_n, ordered by adjacency. Let $a, b, c \in V(\pi\,\mathcal{A}(\Sigma, e_0))$, and let $x \in \mathscr{U}_{e_0}(a,b)$. It suffices to find $y \in \mathscr{U}_{e_0}(a,c) \cup \mathscr{U}_{e_0}(c,b)$ such that $d(x,y) \le 4$. Orient a,b,c such that their initial (resp. terminal) points lie in w_1 (resp. w_n). If $x = a, b$, then we may conclude with $y = a, b$ respectively. Hence without loss of generality, let x be represented by the concatenation of subarcs $\alpha_0 \subset a$ and $\beta_0 \subset b$ such that α_0 is initial in a and β_0 is terminal in b, else exchange the roles of a, b.

Assume boundary components adjacent in C_n are separated by $c \cup x$, else there exists a C_n-allowed arc disjoint from x, c and $d(x,c) \le 2$. For $i \notin \{1, n\}$, let $D_i \supset w_i$ denote the closure of the complementary component of $c \cup x$ containing w_i, and let D_1, D_n denote the closure of the union of components intersecting w_1, resp. w_n. Note that the D_i need not be distinct, although $\mathring{D}_i, \mathring{D}_{i+1}$ are disjoint by assumption. We show the following claim, from which the result follows:

Claim 3.9. *If $D_i \cap \alpha_0 \ne \varnothing$ and $D_{i+1} \cap \alpha_0 \ne \varnothing$ then there exists $y \in \mathscr{U}_{e_0}(a,c)$ such that $d(y,x) \le 4$, and likewise for β_0, b.*

Proof of claim. We consider the case for α_0; the other case is analogous. Fix simple disjoint arcs ρ, δ from w_i and w_{i+1} respectively to α_0, both disjoint from c and disjoint from x except at one endpoint; if $i = 1$, then we allow ρ to be the point $\alpha_0 \cap a_-$. Let η' be the C_n-allowed concatenation of ρ, δ, and the subarc of α_0 between ρ and δ. Let $\gamma_0 \subset c$ denote the subarc between w_n and the first intersection with α_0, or c if no such intersection exists, and let y denote the e_0-unicorn formed by concatenating an initial arc of α_0 with γ_0, or c if $\gamma_0 = c$. We observe that y, η' and η', x both intersect at most once. Applying Proposition 2.6, $d(x,y) \le d(x,\eta') + d(\eta', y) \le 4$. □

In particular, we note that $\alpha_0 \cap D_1 \ne \varnothing$ and $\beta_0 \cap D_n \ne \varnothing$, and in general $\partial D_i \cap x \ne \varnothing$ else c is not simple, hence either $D_i \cap \alpha_0 \ne \varnothing$ or $D_i \cap \beta_0 \ne \varnothing$. Since n is odd, the hypothesis of the claim must hold for some i. □

To apply Theorem 3.1, it remains only to verify that (I) is satisfied:

Lemma 3.10. *For any disjoint $a, b \in V(\mathcal{A}(\Sigma,\Gamma))$, $\operatorname{diam} \mathscr{U}^+(a,b) \le 7$.*

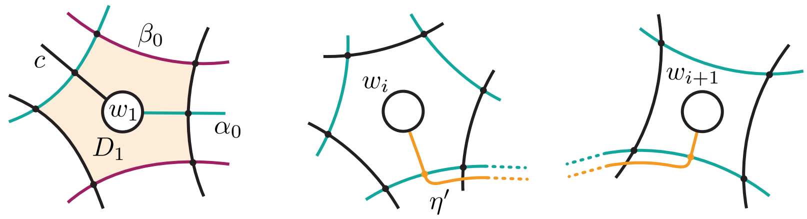

FIGURE 5. The region D_1; the arc η' in the proof of the claim.

Proof. Let w_i denote the boundary components in C_n, as in the proof of Lemma 3.8. It suffices to show that for any $x \in \mathscr{U}_{e_0}(\xi a, \xi b)$, $d(x, \{a, b\}) \leq 3$. Let $a' = \xi a, b' = \xi b$, oriented such that their initial (resp. terminal) points lie in w_1 (resp. w_n). We assume $x \neq a', b'$, else $d(x, \{a, b\}) \leq 1$. Hence, without loss of generality, let x be represented by the concatenation of subarcs $\alpha'_0 \subset a'$ and $\beta'_0 \subset b'$ such that α'_0 is initial in a' and β'_0 is terminal in b', else exchange the roles of a, b.

For $i \neq \{1, n\}$, let ρ'_i be the shortest path from w_i to $X = a \cup x \cup b$, and let p'_i denote the point of intersection $\rho_i \cap X$. If $p'_i \in a, b$, let ρ_i be the concatenation of ρ'_i with the shortest subarc between p'_i and x along a, b respectively; else let $\rho_i = \rho'_i$. Let p_i denote the endpoint of ρ_i along x, and note that if $p_i \in \alpha'_0$, then since a, a' are disjoint $p_i \in b \cup \alpha'_0$ and $\rho_i \cap X \subset a' \cup b$. a, b are likewise disjoint, hence ρ_i is disjoint from a, and analogously if $p_i \in \beta'_0$ then ρ_i is disjoint from b. Suppose that $p_2 \in \alpha'_0$, and let δ be the concatenation of ρ_2 with the subarc between p_2 and c_1 along α'_0. Since δ joins boundary components w_1, w_2 adjacent in C_n, δ is Γ-allowed; moreover, δ is disjoint up to isotopy from a and x, hence $d(x, a) = 2$. Thus assume $p_2 \notin \alpha'_0$ and, by a similar argument, $p_{n-1} \notin \beta'_0$.

We claim that if $p_i, p_{i+1} \in \alpha'_0$, then there exists a Γ-allowed arc δ such that δ is disjoint from a and intersects x at most once, and likewise for β'_0 and b. In particular, in the case for α'_0 let δ be the concatenation of ρ_i, ρ_{i+1}, and the subarc along α'_0 between p_i, p_{i+1}; δ satisfies the claim, and the remaining case for β'_0 is analogous. Finally, since n is odd and $p_2 \in \beta'_0$ and $p_{n-1} \in \alpha'_0$, the hypotheses of the claim are satisfied for some i. Applying Proposition 2.6, $d(x, \{a, b\}) \leq d(x, [\delta]) + d([\delta], \{a, b\}) \leq 3$. $\square$

Proposition 3.3 follows from Theorem 3.1.

4. (Σ, Γ) FOR WHICH $\mathcal{A}(\Sigma, \Gamma)$ IS NOT HYPERBOLIC

We aim to show the converse of Theorem 3.4, namely that, outside of some low-complexity cases, if Γ is bipartite then $\mathcal{A}(\Sigma, \Gamma)$ is not hyperbolic. In the case that Γ is bipartite we obtain two independent $\mathbb{Z}$-actions with positive translation length, each supported on a disjoint witness, and thereby a quasi-isometric embedding $\mathbb{Z}^2 \hookrightarrow \mathcal{A}(\Sigma, \Gamma)$. In considering such actions, we note that the usual mapping class group does not act upon $\mathcal{A}(\Sigma, \Gamma)$ in general: if Γ contains a non-loop edge and is not complete, or if Γ contains a loop edge but not every loop edge, then there exists

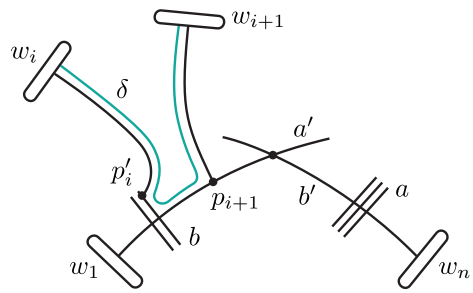

FIGURE 6. The arc δ in the proof of Lemma 3.10 when $p_i, p_{i+1} \in \alpha_0'$.

a mapping class whose induced permutation on $\pi_0(\partial\Sigma) = V(\Gamma)$ does not preserve adjacency in Γ. We make the following definition:

Definition 4.1. Let $\mathrm{Mod}(\Sigma, \Gamma) \leq \mathrm{Mod}(\Sigma)$ denote the subgroup of mapping classes $\varphi \in \mathrm{Mod}(\Sigma)$ whose induced map $\varphi_* : \pi_0(\partial\Sigma) \to \pi_0(\partial\Sigma)$ defines a graph automorphism on Γ.

Since $\mathrm{Mod}(\Sigma, \Gamma)$ maps Γ-allowed arcs to Γ-allowed arcs and preserves disjointness, the following statement is immediate:

Proposition 4.2. $\mathrm{Mod}(\Sigma, \Gamma)$ *acts simplicially on* $\mathcal{A}(\Sigma, \Gamma)$. $\square$

Remark 4.3. The pure mapping class group $\mathrm{PMod}(\Sigma)$ is a subgroup of $\mathrm{Mod}(\Sigma, \Gamma)$.

4.1. **Subsurface projection.** For a witness subsurface $W \subset \Sigma$, we may restrict the usual definitions of subsurface projection for arc graphs to the prescribed arc graph $\mathcal{A}(\Sigma, \Gamma)$. In particular:

Definition 4.4. Let $W \subset \Sigma$ be a Γ-witness. Then let $\sigma_W = \rho_W \circ \pi : \mathcal{A}(\Sigma, \Gamma) \to \mathcal{A}(W)$ denote the *subsurface projection* to W, where $\pi : \mathcal{A}(\Sigma, \Gamma) \to \mathcal{A}(\Sigma)$ is the canonical inclusion and $\rho_W : \mathcal{A}(\Sigma) \to \mathcal{A}(W)$ is the usual subsurface projection.

Remark. We regard ρ_W as a choice of true function instead of a coarse function; we note for $a \in \mathcal{A}(\Sigma)$, the choice of image $\rho_W(a) \in \mathcal{A}(W)$ is canonical up to uniformly bounded diameter D. Thus ρ_W is coarsely 1-Lipschitz, and likewise since π is 1-Lipschitz, σ_W is coarsely 1-Lipschitz, *i.e.* $d_{\mathcal{A}(W)}(\sigma_W(a), \sigma_W(b)) \leq d_{\mathcal{A}(\Sigma,\Gamma)}(a, b)+2D$ for $a, b \in V(\mathcal{A}(\Sigma, \Gamma))$.

Since σ_W is coarsely 1-Lipschitz and coarsely equivariant with respect to homeomorphisms of pairs $(\Sigma, W) \to (\Sigma, W)$, we have the following:

Lemma 4.5. *Let* $W \subset \Sigma$ *be a* Γ*-witness for which there exists a loxodromic element* $\varphi \in \mathrm{PMod}(W)$ *acting on* $\mathcal{A}(W)$*. Then any extension* $\tilde{\varphi} \in \mathrm{Mod}(\Sigma, \Gamma)$ *acts with positive translation length on* $\mathcal{A}(\Sigma, \Gamma)$*.* $\square$

For a witness $W \subset \Sigma$, we are interested in pseudo-Anosov (without loss of generality, pure) mapping classes that act loxodromically on $\mathcal{A}(W)$ and furnish extensions to $\mathrm{Mod}(\Sigma, \Gamma)$ with positive translation length. It is for this reason we require $W \neq \Sigma_0^3$.

We remark that for $\Gamma' \subsetneq \Gamma$, the canonical inclusion $\pi : \mathcal{A}(\Sigma, \Gamma') \to \mathcal{A}(\Sigma, \Gamma)$ is often not a quasi-isometry. It suffices that the sets of Γ'-witnesses and Γ-witnesses are distinct:

Proposition 4.6. *Let $\Gamma' \subsetneq \Gamma$ and assume that $\mathcal{A}(\Sigma, \Gamma)$ is connected. If there exists a subsurface W that is a Γ'-witness but not a Γ-witness, then $\pi : \mathcal{A}(\Sigma, \Gamma') \to \mathcal{A}(\Sigma, \Gamma)$ is not a quasi-isometric embedding.*

Proof. Fix a Γ-allowed arc a disjoint from W. Let φ be a pseudo-Anosov mapping class in $\mathrm{PMod}(W)$, and let $\tilde{\varphi}$ be an extension of φ to Σ by identity on $\Sigma \setminus W$. Since φ_* acts loxodromically on $\mathcal{A}(W)$, by Lemma 4.5 $\tilde{\varphi}_*$ acts with positive translation length on $\mathcal{A}(\Sigma, \Gamma')$. Thus for any $b \in V(\mathcal{A}(\Sigma, \Gamma'))$, $\{\tilde{\varphi}_*^k(b)\}$ has infinite diameter in $\mathcal{A}(\Sigma, \Gamma')$, hence if π is a quasi-isometric embedding then likewise $\{\pi\tilde{\varphi}_*^k(b)\} = \{\tilde{\varphi}_*^k(\pi b)\}$ has infinite diameter in $\mathcal{A}(\Sigma, \Gamma)$, where the equality follows from the naturality of π. But a is fixed by $\tilde{\varphi}$, hence $\operatorname{diam}\{\tilde{\varphi}_*^k(\pi b)\} = 2d(a, \pi b) < \infty$, a contradiction. □

Example 4.7. Suppose that $\chi(\Sigma) \leq -3$ and Γ is loop-free. Then there is a Γ'-witness subsurface $W \subset \Sigma$ that is not a witness for $\mathcal{A}(\Sigma, \Gamma)$, hence π is not a quasi-isometry. Let $e \in E(\Gamma) \setminus E(\Gamma')$, fix a (non-separating) e-allowed arc a and a regular neighborhood $N \supset a \cup \partial a$, and let $W = \Sigma \setminus N$. Since $\chi(N) = -1$ and $\chi(W) = \chi(\Sigma) + \chi(N) \leq -2$, ∂W is an essential simple closed curve and W is essential; moreover, $W \ncong \Sigma_0^3$. $N \cap \partial\Sigma = \partial a = V(e)$ and $e \notin \Gamma'$, hence there are no Γ'-allowed arcs in $N = \Sigma \setminus W$. W is a Γ'-witness disjoint from $a \in V(\mathcal{A}(\Sigma, \Gamma))$. //

4.2. **Disjoint witness subsurfaces.** We show that outside of some low-complexity cases, bipartite Γ is equivalent to the existence of distinct, disjoint witnesses W_i.

Lemma 4.8. *Suppose that $\chi(\Sigma) \leq -3$ and if $\Sigma = \Sigma_0^{n+1}$ then Γ is not a n-pointed star. Then Γ is bipartite if and only if there exist two disjoint, distinct Γ-witnesses.*

Proof. We first prove the reverse direction. Suppose that there exist two disjoint Γ-witnesses, $W_1, W_2 \subset \Sigma$. Let $\mathscr{C}_{i,j} \subset \pi_0(\partial\Sigma)$ denote the boundary components of Σ contained in the j-th complementary component $C_{i,j}$ of W_i. Since W_i is essential, each $\mathscr{C}_{i,j}$ must be a loop-free, independent set in Γ, else there exists a Γ-allowed arc in $C_{i,j}$ disjoint from W_i. Let $\mathscr{W}_1 = \pi_0(\partial W_1) \cap \pi_0(\partial\Sigma)$ and observe that $V(\Gamma) = \mathscr{W}_1 \sqcup \bigsqcup_j \mathscr{C}_{1,j}$. Since W_1, W_2 are disjoint and connected, each lies within a unique complementary component of the other: without loss of generality, assume $W_1 \subset C_{2,1}$ and $W_2 \subset C_{1,1}$, and observe that $W_1 \cup \bigcup_{j \neq 1} C_{1,j}$ is connected and disjoint from $W_2 \subset C_{1,1}$, hence likewise lies in $C_{2,1} \supset W_1$. Then $\mathscr{W}_1, \bigcup_{j\neq 1} \mathscr{C}_{1,j} \subset \mathscr{C}_{1,1}$, hence $\mathscr{D} = \mathscr{W}_1 \cup \bigcup_{j \neq 1} \mathscr{C}_{1,j}$ is independent in Γ and $\mathscr{C}_{1,1} \sqcup \mathscr{D} = V(\Gamma)$ partitions $V(\Gamma)$ into two independent sets.

Conversely, suppose that $V(\Gamma)$ may be partitioned into two independent sets X_1, X_2, and without loss of generality assume both are non-empty. Moreover, we may assume $|X_1| = 1$ only if $|X_2| = 1$, and $|X_2| = 1$ only if Γ is an n-pointed star with center in X_2, else add any isolated vertices to X_2. We first show there is an essential simple closed curve separating X_1, X_2. By a change of coordinates

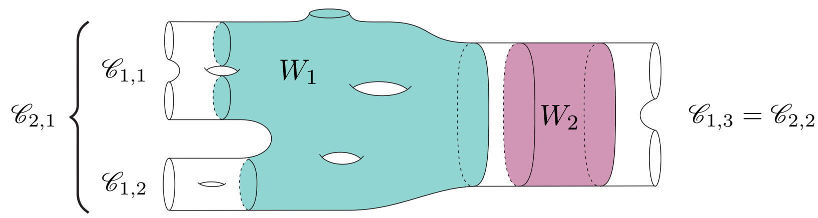

FIGURE 7. Disjoint witnesses and the collections of boundary components $\mathscr{C}_{i,j}$.

argument, there is a simple closed curve $\zeta \subset \Sigma$ separating X_1, X_2; let $Z_1 \supset X_1, Z_2 \supset X_2$ denote its complementary components. ζ is not essential if and only if it is peripheral if and only if (exactly) one of Z_1, Z_2 is an annulus. By our assumptions on X_1, X_2, in either case $|X_2| = 1$ and Γ is an n-pointed star, hence Σ has genus. Fix a subsurface $Z_2' \cong \Sigma_1^2$ containing X_2 and one genus. Let $\zeta' = \partial Z_2' \setminus \partial\Sigma$ and $Z_1' = \Sigma \setminus Z_2'$. Then ζ' is a simple closed curve separating X_1, X_2 and $\chi(Z_1') = \chi(\Sigma) - \chi(Z_2') = \chi(\Sigma) + 2 \leq -1$: neither of the complements of ζ' are annular, hence ζ' is essential. Up to replacing ζ with ζ', assume ζ is essential.

Since $\chi(Z_1) + \chi(Z_2) = \chi(\Sigma) \leq -3$, at least one of $\chi(Z_i) \leq -2$. Assume that $\chi(Z_1) \leq -2$ and let $W_1 = \overline{Z}_1$ and W_2 a closed annular neighborhood of ζ; since $\chi(W_1) \leq -2$, $W_1 \neq \Sigma_0^3$. Thus for either W_i, W_i is an essential non-pants subsurface and ζ is parallel to a component of ∂W_i, hence W_i separates X_1 and X_2: any Γ-allowed arc must intersect W_i. W_1, W_2 are homotopically disjoint and distinct Γ-witnesses. □

We conclude by adapting a well known fact (see [MS13, Lem. 5.9]) to our setting:

Proposition 4.9. *If $\mathcal{A}(\Sigma,\Gamma)$ admits two disjoint, distinct Γ-witnesses, then $\mathcal{A}(\Sigma,\Gamma)$ is not δ-hyperbolic and in particular there is a quasi-isometric embedding $\mathbb{Z}^2 \to \mathcal{A}(\Sigma,\Gamma)$.*

In particular, it suffices that the witnesses furnish commuting $\mathbb{Z}$ actions (by extensions of pseudo-Anosovs) with non-zero translation lengths. We omit the proof.

The main theorem of this section then follows immediately from Lemma 4.8 and Proposition 4.9:

Theorem 4.10. *Suppose that $\chi(\Sigma) \leq -3$ and if $\Sigma = \Sigma_0^{n+1}$ then Γ is not a n-pointed star. If Γ is bipartite, then $\mathcal{A}(\Sigma,\Gamma)$ is not δ-hyperbolic.* □

5. SPORADIC CASES

Assuming $\chi(\Sigma) \leq -1$, $E(\Gamma) \neq \varnothing$, and $\Sigma \neq \Sigma_0^3$, we enumerate the low complexity cases not addressed in the preceding sections:

(i) Infinite diameter. Since $\Sigma \neq \Sigma_0^3$, there exists a pseudo-Anosov class in $\mathrm{PMod}(\Sigma)$ acting loxodromically on $\mathcal{A}(\Sigma)$, hence $\mathcal{A}(\Sigma)$ is infinite diameter and by Theorem 2.9 likewise is $\mathcal{A}(\Sigma,\Gamma)$. //

(ii) Connectivity and δ-hyperbolicity. Theorem 2.7 implies the connectivity of $\mathcal{A}(\Sigma,\Gamma)$ except when $\Sigma=\Sigma_1^2$ and Γ is two disjoint loop edges. Similarly, if Γ is not bipartite then Theorem 3.4 implies δ-hyperbolicity for $\mathcal{A}(\Sigma,\Gamma)$ except when $\Sigma=\Sigma_1^2$ and Γ is two disjoint loop edges. We show connectedness and δ-hyperbolicity in this case in Lemma 5.2.

(iii) Non-hyperbolicity. If Γ is bipartite, then Theorem 4.10 implies that $\mathcal{A}(\Sigma,\Gamma)$ is non-δ-hyperbolic unless $\chi(\Sigma)\geq -2$ or $\Sigma=\Sigma_0^{n+1}$ and Γ is a n-pointed star. We show δ-hyperbolicity in the latter case in Section 5.2. For the former, it remains to consider only if $\Sigma=\Sigma_0^4$ or if $\Sigma=\Sigma_1^2$ and Γ is a non-loop edge; in the second case, we show δ-hyperbolicity in Lemma 5.3.

If $\Sigma=\Sigma_0^4$, then $\mathcal{A}(\Sigma,\Gamma)$ is a full subgraph of $\mathcal{A}(\Sigma_0^4)$, which is quasi-isometric to the Farey graph hence a quasi-tree; $\mathcal{A}(\Sigma,\Gamma)$ is likewise a quasi-tree and thus δ-hyperbolic. //

Collecting the results for these sporadic cases with the general results in Sections 2, 3, and 4, we have shown Theorem 1.2:

Theorem 1.2. *Assume that $\chi(\Sigma)\leq -1, E(\Gamma)\neq\varnothing$, and $\Sigma\neq\Sigma_0^3$. Then $\mathcal{A}(\Sigma,\Gamma)$ is connected and has infinite diameter.* □

We note that if $\Sigma=\Sigma_0^{n+1}$ and Γ is a n-pointed star with center c, then every witness subsurface must contain c: Σ does not admit disjoint Γ-witnesses. Similarly, if $\Sigma=\Sigma_0^4$ or Σ_1^2, then Σ does not admit two disjoint, distinct Γ-witnesses that are not Σ_0^3. Hence we conclude the following, proving Theorems 1.3 and 1.6:

Theorem 5.1. *Assume that $\chi(\Sigma)\leq -1, E(\Gamma)\neq\varnothing$, and $\Sigma\neq\Sigma_0^3$. Then*

(i) $\mathcal{A}(\Sigma,\Gamma)$ is (uniformly) δ-hyperbolic if and only if Σ does not admit two distinct, disjoint Γ-witnesses.

In particular, if $\chi(\Sigma)\geq -2$ or if $\Sigma=\Sigma_0^{n+1}$ and Γ is a n-pointed star then $\mathcal{A}(\Sigma,\Gamma)$ is δ-hyperbolic. Outside of these sporadic cases, (i) is equivalent to the following:

(ii) $\mathcal{A}(\Sigma,\Gamma)$ is (uniformly) δ-hyperbolic if and only if Γ is not bipartite. □

We conclude by addressing the cases when $\Sigma=\Sigma_1^2$ and Γ is either two loop edges or a non-loop edge, and when $\Sigma=\Sigma_0^{n+1}$ and Γ is a n-pointed star.

5.1. $\Sigma=\Sigma_1^2$ and Γ is two loop edges or a non-loop edge.

Lemma 5.2. *Suppose that $\Sigma=\Sigma_1^2$ and Γ is two disjoint loop edges. Then $\mathcal{A}(\Sigma,\Gamma)$ is connected and δ-hyperbolic.*

Proof. Let ℓ_1,ℓ_2 denote the two loop edges in Γ, and let w_i denote the boundary component in ℓ_i. Let $\pi_i:\mathcal{A}(\Sigma,\ell_i)\to\mathcal{A}(\Sigma,\Gamma)$ be the respective canonical map; note that $\operatorname{im}\pi_1,\operatorname{im}\pi_2$ partition $V(\mathcal{A}(\Sigma,\Gamma))$. We show $\mathcal{A}(\Sigma,\Gamma)$ is the graph product of $\mathcal{A}(\Sigma,\ell_1)$ and a single, non-loop edge, hence quasi-isometric to $\mathcal{A}(\Sigma,\ell_1)$: in particular, $\mathcal{A}(\Sigma,\Gamma)$ is obtained from $\mathcal{A}(\Sigma,\ell_1)\sqcup\mathcal{A}(\Sigma,\ell_1)$, with edges added between pairs of corresponding vertices. Since $\mathcal{A}(\Sigma,\ell_1)$ is connected and δ-hyperbolic by Proposition 2.3 and Corollary 3.6 respectively, the claim follows.

We note that each arc $a\in\operatorname{im}\pi_i$ is adjacent to exactly one arc $a'\in\operatorname{im}\pi_{3-i}$. In particular, if *e.g.* a is a ℓ_1-allowed arc, then cutting along a we obtain a three-holed sphere $\Sigma' = \Sigma_0^3$: there exists exactly one essential simple arc $a' \subset \Sigma'$ with endpoints on w_2, up to isotopy. Let $\tau : V(\mathcal{A}(\Sigma, \ell_1)) \to V(\mathcal{A}(\Sigma, \ell_2))$ be the bijection sending each ℓ_1-allowed arc a to its unique disjoint ℓ_2-allowed arc a'. Since each π_i is a graph inclusion, it suffices to show that τ is an isometry. In particular, it is enough to find a homeomorphism of pairs $(\Sigma, w_1) \to (\Sigma, w_2)$ that induces τ.

For convenience, observe that we can realize $\mathcal{A}(\Sigma, \Gamma)$ by replacing the boundary components w_i of Σ with marked points p_i on the torus. Fix a flat metric on $\Sigma \cong \mathbb{R}^2/\mathbb{Z}^2$ such that $p_1 = (0, 0)$ and $p_2 = (\theta, 0)$ for some irrational $\theta \in (0, 1)$, and define the homeomorphism $\xi : (a, b) \mapsto (a + \theta, b)$; we note that $\xi : p_1 \mapsto p_2$. For any ℓ_1-allowed arc a, we may assume a is a geodesic loop with non-zero slope and basepoint p_1, hence ξa is a geodesic loop with basepoint p_2 disjoint from a. Thus ξa is ℓ_2-allowed and by uniqueness $\xi a = a'$: ξ induces the map τ on $V(\mathcal{A}(\Sigma, \ell_1))$, as desired. □

Lemma 5.3. *Suppose that $\Sigma = \Sigma_1^2$ and Γ is a non-loop edge. Then $\mathcal{A}(\Sigma, \Gamma)$ is δ-hyperbolic.*

Proof. Let $\pi : \mathcal{A}(\Sigma, \Gamma) \to \mathcal{A}(\Sigma)$ denote the canonical inclusion to the usual arc graph. We show that π is a quasi-isometry: since $\mathcal{A}(\Sigma)$ is δ-hyperbolic, the lemma follows. By Lemma 2.8 π is 3-coarsely surjective; since it is simplicial, it is 1-Lipschitz. It remains to verify the lower quasi-isometry bound. We prove below:

Claim 5.4. *For any Γ-allowed arc $a \in \pi\,\mathcal{A}(\Sigma, \Gamma)$ and any arc $a' \in \mathcal{A}(\Sigma)$, there is a path P from a to a' with length $\ell(P) \leq 2d_{\mathcal{A}(\Sigma)}(a, a') - 1$ and $P \setminus \{a'\} \subset \pi\,\mathcal{A}(\Sigma, \Gamma)$.*

For $a' \in \pi\,\mathcal{A}(\Sigma, \Gamma)$, the claim then implies $d_{\mathcal{A}(\Sigma,\Gamma)}(a, a') \leq \ell(P) < 2d_{\mathcal{A}(\Sigma)}(a, a')$. □

Proof of Claim 5.4. We proceed by induction on $d = d_{\mathcal{A}(\Sigma)}(a, a')$. When $d \leq 1$, the claim is immediate. For $d = 2$, let b be an arc disjoint from a and a'. If b is Γ-allowed then $P = (a, b, a')$ satisfies the claim. Assume not, hence $b^+ = b^-$. Fix representatives along with a regular neighborhood $N \supset b$ such that N is disjoint from a, a'. Let c_0, c_1, c_2 be the boundary components of $\Sigma' = \Sigma \setminus N$, with $c_0 \subset \partial\Sigma$ and $\partial a = \{c_0, c_1\}$ in Σ'. Note that $\Sigma' \cong \Sigma_0^3$ and that by fixing representatives we may obtain from $\Sigma' \hookrightarrow \Sigma$ a simplicial embedding $\mathcal{A}(\Sigma') \hookrightarrow \mathcal{A}(\Sigma)$; denote by a_{ij} the image of the (unique) arc $\tilde{a}_{ij} \in \mathcal{A}(\Sigma')$ with $\partial\tilde{a}_{ij} = \{c_i, c_j\}$ for $i, j \in \{0, 1, 2\}, i \leq j$.

Up to Dehn twists along c_1, c_2, arcs in $\mathcal{A}(\Sigma)$ disjoint from b are isotopic rel b to one of the a_{ij}. Let μ be a multitwist about $c_1 \cup c_2$ such that $\mu(a') = a_{ij}$ for some i, j and let ν be a twist about c_1 such that $\nu\mu(a) = a_{01}$. Let $b' = \mu^{-1}(a_{02})$, and note that a_{02} is disjoint from every $a_{ij} \neq a_{11}$. Since a_{02} and c_1 are disjoint, $\nu(a_{02}) = a_{02}$ and $\nu\mu(a) = a_{01}$ is disjoint from $\nu\mu(b') = \nu(a_{02}) = a_{02}$: a, b' are disjoint. Similarly, if $\mu(a') = a_{ij} \neq a_{11}$ then a', b' are disjoint. Otherwise, if $\mu(a') = a_{11}$ then let $b'' = \mu^{-1}(a_{01})$, which is disjoint from a' since $\mu(a') = a_{11}$ and $\mu(b'') = a_{01}$ are disjoint, and is disjoint from b' since $\mu(b'') = a_{01}$ and $\mu(b') = a_{02}$ are disjoint. Let $P = (a, b', a')$ or (a, b', b'', a') as appropriate, both of which satisfy the claim.

Assume that the claim holds for $d = k$, and suppose that $d = k + 1$. Fix $b \in V(\mathcal{A}(\Sigma))$ adjacent to a' and such that $d(a, b) = k$, hence there exists a path P' from a to b such that $\ell(P') \leq 2k - 1$ and $P' \setminus \{b\} \subset \pi\,\mathcal{A}(\Sigma, \Gamma)$. Let $b' \in \pi\,\mathcal{A}(\Sigma, \Gamma)$

be the pentultimate vertex in P'. By the $d \leq 2$ cases above, there exists P'' from b' to a' with $\ell(P'') \leq 3$ and $P'' \setminus \{a'\} \subset \pi\,\mathcal{A}(\Sigma,\Gamma)$. The concatenation P of $P' \setminus \{b\}$ and P'' satisfies $\ell(P) \leq 2k-2+3 = 2(k+1)-1$ and $P \setminus \{a'\} \subset \pi\,\mathcal{A}(\Sigma,\Gamma)$. □

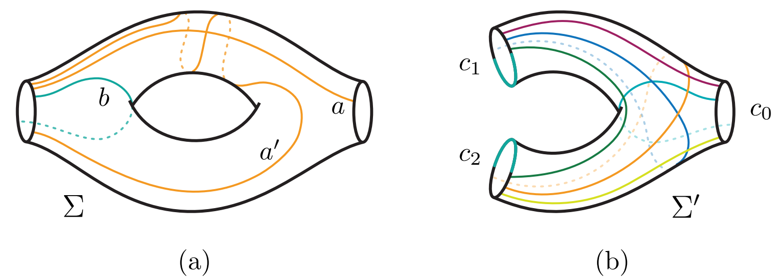

FIGURE 8. (a) Σ with arcs a, b, a'. (b) The subsurface Σ' obtained by cutting along b, with the six arcs a_{ij}.

5.2. $\Sigma = \Sigma_0^{n+1}$ **and** Γ **is a** n**-pointed star,** $n \geq 3$**.** The case when Σ is a $(n+1)$-holed sphere and Γ is a n-pointed star is analogous to the ray graph in [Bav16], but rather than a Cantor set of punctures we consider only finitely many. We reproduce Bavard's argument for δ-hyperbolicity below, suitably simplified for our purposes.

Lemma 5.5. *Let $\Sigma = \Sigma_0^{n+1}$ and Γ a n-pointed star with center c. Then $\mathcal{A}(\Sigma,\Gamma)$ is quasi-isometric to $\mathcal{A}(\Sigma,\ell_0)$, where ℓ_0 is a loop edge on c.*

Proof. It suffices to define a quasi-isometry $\xi : V(\mathcal{A}(\Sigma,\Gamma)) \to V(\mathcal{A}(\Sigma,\ell_0))$. Given an isotopy class of Γ-allowed arcs $a \in V(\mathcal{A}(\Sigma,\Gamma))$, we will always assume an orientation such that $a_- = c$. Let N be a regular neighborhood of $\alpha \cup a_+$ for $\alpha \in a$ a choice of representative; let $\hat{a}$ denote the isotopy class of ∂N. Since $n \geq 3$, $\hat{a}$ is essential; $\partial\hat{a} = \{c\}$, hence $\hat{a}$ is an isotopy class of ℓ_0-allowed arcs. Note that $\hat{a}$ is independent of the choice of α and N, and let $\xi : a \mapsto \hat{a}$.

If $\alpha \in a, \beta \in b$ are disjoint Γ-allowed arcs, then we may choose regular neighborhoods of $\alpha \cup a_+$ and $\beta \cup b_+$ respectively such that $i(\hat{a},\hat{b}) \leq 2$. Thus by Proposition 2.6 $d(\xi a, \xi b) \leq 3$, and in general ξ is 3-Lipschitz. For the lower quasi-isometry bound, consider $a, b \in V(\mathcal{A}(\Sigma,\Gamma))$ such that $d(\xi a, \xi b) = m$, and let $u_1 = \xi a, u_2, \ldots, u_m = \xi b$ be distinct ℓ_0-allowed arcs forming a geodesic path in $\mathcal{A}(\Sigma,\ell_0)$. We construct a path $v_1 = a, v_2, \ldots, v_m = b$ of Γ-allowed arcs in $\mathcal{A}(\Sigma,\Gamma)$, hence $d(a,b) \leq d(\xi a, \xi b)$, as desired.

We note that each u_i separates Σ into two complementary components homeomorphic to $\Sigma_{0,1}^k, \Sigma_{0,1}^{n-k}$ respectively, and that for $i \neq 1, m$, both u_{i-1}, u_{i+1} must lie in the same component: since (u_i) is geodesic, u_{i-1}, u_{i+1} cannot be adjacent in $\mathcal{A}(\Sigma,\ell_0)$, hence must intersect. Let $S_i \subset \Sigma$ denote the complementary component of u_i that contains u_{i-1} or u_{i+1}; let S_i' denote the the other component, which contains neither. We note that S_i must contain at least two boundary components of Σ (that are not c), else every ℓ_0-allowed arc in S_i belongs to u_i and $u_{i-1} = u_i$ or $u_{i+1} = u_i$, a contradiction since each u_i is distinct. Similarly, S_i' contains at least one boundary component $d_i \neq c$, since u_i is essential; we moreover claim that S'_i, S'_{i+1} are disjoint, else *e.g.* $u_i = \partial S'_i \subset S'_{i+1}$, a contradiction with our choice of S'_{i+1}. Finally, choose v_i to be an arc between c and d_i in S'_i. Since S'_i is disjoint from S'_{i+1}, v_i is disjoint from v_{i+1}, and note that we may choose $v_1 = a$ and $v_m = b$ since the complementary component of u_1 containing a contains only one boundary component of Σ, hence $a \not\subset S_1$, and likewise $b \not\subset S_m$. (v_i) is the desired path.

Coarse surjectivity follows from a similar argument: as above, a complementary component S of an (essential) ℓ_0-allowed arc u must contain at least one component $d \neq c$ of $\partial\Sigma$. Any choice of simple arc α from c to d in (the interior of) S is Γ-allowed and has a regular neighborhood $N \subset S$. Hence $a = [\alpha] \in V(\mathcal{A}(\Sigma, \Gamma))$ and $\xi a = \hat{a}$ is disjoint from u: ξ is 1-coarsely surjective. $\square$

By Corollary 3.6 $\mathcal{A}(\Sigma, \ell_0)$ is δ-hyperbolic, hence we may conclude that $\mathcal{A}(\Sigma, \Gamma)$ is likewise δ-hyperbolic. //

6. Acknowledgements

The author would like to thank Mladen Bestvina for numerous helpful discussions and the referees for their comments. This work was supported in part by NSF award no. 1840190: *RTG: Algebra, Geometry, and Topology at the University of Utah.*

References

[AV18] Javier Aramayona and Ferrán Valdez, *On the geometry of graphs associated to infinite-type surfaces*, Mathematische Zeitschrift **289** (2018), 309–322.

[Bav16] Juliette Bavard, *Hyperbolicité du graphe des rayons et quasi-morphismes sur un gros groupe modulaire*, Geometry & Topology **20** (2016), no. 1, 491–535.

[BE88] B.H. Bowditch and D.B.A. Epstein, *Natural triangulations associated to a surface*, Topology **27** (1988), no. 1, 91–117.

[BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Quasiflats in hierarchically hyperbolic spaces*, Duke Mathematical Journal **170** (2021), no. 5, 909 – 996.

[BNV22] Assaf Bar-Natan and Yvon Verberne, *The grand arc graph*, 2022.

[Bow14] Brian H. Bowditch, *Uniform hyperbolicity of the curve graphs*, Pacific Journal of Mathematics **269** (2014), no. 2, 269–280.

[Cal09] Danny Calegari, *Big mapping class groups and dynamics*, Blog post, 2009.

[DFV18] Matthew Gentry Durham, Federica Fanoni, and Nicholas G. Vlamis, *Graphs of curves on infinite-type surfaces with mapping class group actions*, Annales de l'Institut Fourier **68** (2018), no. 6, 2581–2612.

[FGM21] Federica Fanoni, Tyrone Ghaswala, and Alan McLeay, *Homeomorphic subsurfaces and the omnipresent arcs*, Annales Henri Lebesgue **4** (2021), 1565–1593.

[Har86] John Harer, *The virtual cohomology dimension of the mapping class group of an orientable surface*, Inventiones mathematicae **84** (1986), 157–176.

[HPW15] Sebastian Hensel, Piotr Przytycki, and Richard Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, Journal of the European Mathematical Society **17** (2015), no. 4, 755–762.

[Mos83] Lee Mosher, *Pseudo-anosovs on punctured surfaces*, Dissertation, Princeton, 1983.

[MR20] Kathryn Mann and Kasra Rafi, *Large scale geometry of big mapping class groups*, 2020.

[MS13] Howard Masur and Saul Schleimer, *The geometry of the disk complex*, Journal of the American Mathematical Society **26** (2013), no. 1, 1–62.

[Vok22] Kate M. Vokes, *Hierarchical hyperbolicity of graphs of multicurves*, Algebraic & Geometric Topology **22** (2022), no. 1, 113–151.