



The coarse geometry of hexagon decomposition graphs

Funda Gültepe and Hugo Parlier

Abstract. We define and study graphs associated to hexagon decompositions of surfaces by curves and arcs. One of the variants is shown to be quasi-isometric to the pants graph, whereas the other variant is quasi-isometric to (a Cayley graph of) the mapping class group.

1 Introduction

Simplicial complexes related to curves and arcs have been used with considerable success in the study of surfaces and their deformation spaces. In particular, they have proved to be efficient tools for the study of the large-scale geometry of Teichmüller spaces and mapping class groups. For instance, curve graphs are rough models for the electrified Teichmüller metric, as is the marking graph for mapping class groups [13, 14], and the pants graph for the Weil-Petersson metric [3]. Flip graphs, which originally appeared in topics closer to computational or combinatorial geometry, also provide quasi-models for mapping class groups, but, like arc graphs, they require marked points to serve as basepoints for arcs. This is also true of other similar graphs, such as the polygonalization graph introduced recently [1].

Here, we introduce graphs that use both curves and arcs but, unlike standard arc graphs, do not require marked points. Instead, homotopy classes of arcs go between curves (with endpoints allowed to glide on the curves). In particular, these graphs work for surfaces with or without marked points, and so work for closed surfaces in particular, while mimicking some of the nice properties of flip graphs. Roughly speaking, the vertices of the graphs are collections of curves and arcs that decompose the surface into hexagons. Edge relations come from either flipping an arc or from adding or removing a curve. We have two versions, depending on whether we put weights on the curves.

On an orientable, finite-type and of negative Euler characteristic surface Σ , our first graph $\mathcal{H}(\Sigma)$ is defined purely topologically (see Section 2 for a precise definition), and hence, we call it the *topological hexagon decomposition graph*. We view it as a metric space by giving each edge unit length. Our main result is that this graph is quasi-isometric to the standard pants graph $\mathcal{P}(\Sigma)$.

Received by the editors February 7, 2024; revised September 9, 2024; accepted September 10, 2024.

The first author was partially supported by NSF grant DMS-2137611. Both authors were partially supported by ANR/FNR project SoS, INTER/ANR/16/11554412/SoS, ANR-17-CE40-0033.

AMS subject classification: 57K20, 57M15, 32G15.

Keywords: Moduli spaces, mapping class groups, pants graph.



Theorem 1.1 *For finite-type orientable surfaces, the topological hexagon decomposition graph $\mathcal{H}(\Sigma)$ is quasi-isometric to the pants graph $\mathcal{P}(\Sigma)$.*

By Brock's results on the pants graph [3], this means that $\mathcal{H}(\Sigma)$ is also a coarse model for the Weil-Petersson metric.

Our second type of graph is an augmented version of the above graph, where vertices are the same but with a collection of weights on each curve. Edges come from elementary moves as for $\mathcal{H}(\Sigma)$, but there is the additional consideration of how to determine the weight of an added curve. There would be multiple ways to do this, but we choose to associate weights by measuring lengths with a fixed hyperbolic surface X . We call the resulting graph the *geometric hexagon decomposition graph* and denote it $\mathcal{H}(X)$. In contrast to $\mathcal{H}(\Sigma)$, vertices in $\mathcal{H}(X)$ are of finite valency (Theorem 4.2), and there is a natural (coarse) action of the mapping class group $\text{Mod}(X)$ (which is also $\text{Mod}(\Sigma)$) on it. In contrast to the action of $\text{Mod}(\Sigma)$ on $\mathcal{H}(\Sigma)$, stabilizers are now finite, and hence we obtain the following:

Theorem 1.2 *For finite-type orientable surfaces, the geometric hexagon decomposition graph $\mathcal{H}(X)$ is quasi-isometric to the mapping class group $\text{Mod}(X)$.*

By work of Masur and Minsky [14], the geometric hexagon graph is also quasi-isometric to the *marking graph*. The proof of this theorem follows from a version of the Milnor-Schwartz lemma for quasi-actions. Most of the work goes into showing the graph has the properties that allow this to work. It is interesting to note that the topological graph, which requires only topological input, provides a coarse model for deformation spaces of hyperbolic surfaces, whereas the geometric graph provides a model for $\text{Mod}(\Sigma)$.

Organization. In a preliminary part (Section 2), we define the main building blocks we need – namely, hexagon decompositions with and without weights. In Section 3, we define and study the topological hexagon graph and show it is coarsely equivalent to the pants graph. In Section 4, we study the geometric version and show it is a coarse model for the mapping class group.

2 Preliminaries

Let Σ be a finite-type orientable surface with all connected components of negative Euler characteristic. For reasons of clarity, we think of boundary as consisting of curves (and not punctures or marked points). A *curve* will be an abbreviation for an isotopy class of an essential simple closed curve. If it is non-peripheral to boundary, it is an interior curve. On a surface with boundary, by *arc*, we will mean a simple arc between boundary elements up to homotopy with endpoints gliding on the boundary.

We will also need to consider multicurves which are collections of disjoint and distinct curves. Note that there are different uses of the word multicurve in the literature, but as all of our curves are simple, we avoid calling them systematically simple multicurves. Similarly, given a surface with boundary, a multiarc is a collection of distinct and disjoint arcs. Intersection numbers for curves and arcs are always defined as the minimal geometric intersection numbers between representatives in the respective free homotopy classes.

A multicurve or multiarc is said to be maximal if it is maximal with respect to inclusion. Note that a maximal multicurve is a pants decomposition, and a maximal multiarc decomposes the surface into hexagons. (Note that it is for exactly this reason that we prefer boundary curves to punctures or marked points; otherwise, we would have (ideal) polygons of different sizes in the complementary regions of our collection of arcs.)

We denote by $\kappa_c(\Sigma)$ the curve complexity of Σ – that is the number of curves in a pants decomposition. Similarly, we denote by $\kappa_a(\Sigma)$ the arc complexity – that is, the number of arcs in a hexagon decomposition. For example, if Σ is closed and of genus g , we have $\kappa_c(\Sigma) = 3g - 3$ and $\kappa_a(\Sigma \setminus \Gamma) = 4g - 4$ for any nonempty multicurve Γ .

Definition 2.1 A hexagon decomposition is a pair (Γ, \mathcal{A}) , where Γ is a nonempty multicurve and \mathcal{A} is a maximal multiarc on $\Sigma \setminus \Gamma$. In the case where Σ has boundary, we require that Γ contain all peripheral curves of Σ .

The complementary regions to a hexagon decomposition are of course hexagons. Note that the sides of the complimentary hexagons will necessarily alternate between sides consisting of arcs (from the multiarc) and those coming from subarcs of the multicurve.

A weighted multicurve (Γ, w) is a multicurve $\Gamma = \bigcup_{k=1}^{|\Gamma|} \gamma_k$ together with a collection of weights $w = w(\Gamma) = (k_1, \dots, k_{|\Gamma|}) \in \mathbb{Z}^{|\Gamma|}$. Note that we only allow weights to be integers.

It will be useful to consider metric surfaces: for given Σ , we consider a fixed hyperbolic surface X homeomorphic to Σ , and if Σ has boundary, we ask that the boundary elements of X be realized by simple closed geodesics. (This is simply because we want to be able to measure the lengths of arcs in the easiest possible way.) Thus on X , curves are uniquely realized by simple closed geodesics, and arcs are uniquely realized by simple orthogeodesics (geodesics orthogonal to their terminal simple closed geodesics). A maximal multiarc is then realized by a maximal collection of disjoint orthogeodesics, and the complementary region is a collection of right-angled hexagons.

2.1 Pants graphs

Before introducing the hexagonal graphs, we recall that the pants graph $\mathcal{P}(\Sigma)$ is the graph where vertices are pants decompositions and two vertices are related by an edge if they differ by an elementary move. The two types of elementary moves are illustrated in Figures 1 and 2.

Brock's result [3] that the pants graph is quasi-isometric to Teichmüller space with the Weil-Petersson metric is one example among many of how to use a combinatorial model to study moduli type spaces and their isometry groups. We refer the reader to [2, 4, 5, 6, 8, 9, 10, 11, 12, 15, 16, 17] for others.

2.2 Flip graphs

As our graphs are also closely related to flip graphs, we include some useful facts about them that we will either use or emulate.

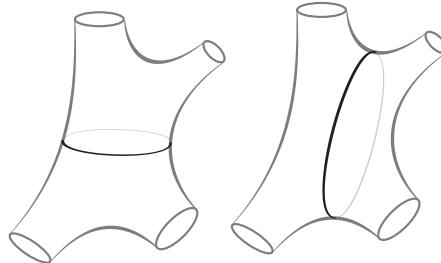


Figure 1: The first type of elementary move is on a four-holed sphere.

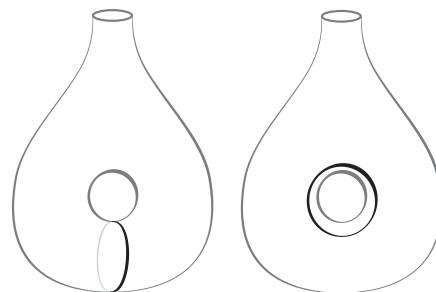


Figure 2: The second type of elementary move is on a one-holed torus.

We first briefly describe flip graphs. Consider a triangulated surface of finite type – that is, a decomposition of the surface along arcs into a finite number of triangles. (In the context of topological surfaces, we prefer to talk about triangulations as being collections of arcs with endpoints on marked points as opposed to edges and vertices in order not to confuse them with the edges and vertices of the flip graphs.) Now one can obtain a new triangulation by removing one of the arcs of the triangulation and by replacing it by the “other possible diagonal arc” (this is called a flip). This gives a new triangulation with the same set of marked points and all but one of the previous arcs. For a fixed surface and set of marked points, this gives rise to a graph where vertices are triangulations and two vertices are joined by an edge if their corresponding triangulations are related by a flip.

We briefly review some useful facts about flip graphs. First of all, they are always connected [8]. In certain specific cases, they are finite graphs. For instance, if the surface is a polygon, and the marked points are on the boundary, the flip graphs are the well-studied associahedron graphs; see, for instance, [22] for results about their diameters. There are other cases where the graphs are finite, such as when the surface is a sphere with three marked points. However, whenever the surface contains a cylinder bounded by two nontrivial arcs, the corresponding flip graph is always infinite, such as in the case of a sphere with four marked points.

The geometry of flip graphs is of particular interest, as, by the Milnor-Schwartz lemma, it is always quasi-isometric to the underlying mapping class group and can be

thus used as a tool to study this important homeomorphism group [17]. However, it always requires that there be at least one marked point, and so it cannot be directly used to study, for instance, mapping class groups of closed surfaces.

3 The topological graph and its properties

3.1 The definition

The topological version only depends on Σ and does not depend on a choice of metric; hence, we denote it $\mathcal{H}(\Sigma)$.

Vertices will be given by hexagon decompositions, and hexagon decompositions are related by an edge if they satisfy one of two relations. The first one is a type of flip relation between maximal multiarcs.

The flip relation. $H = (\Gamma, \mathcal{A})$ and $H' = (\Gamma', \mathcal{A}')$ are related by a flip if $\Gamma = \Gamma'$ and \mathcal{A} and \mathcal{A}' differ by a single arc. This is analogous to the flip operation between triangulations as described previously. See Figure 3 for an illustration.

The second one relates to transformations of the multiarc.

Curve addition.

Let $H = (\Gamma, \mathcal{A})$ be a hexagon decomposition with Γ non-maximal. A curve α is said to be compatible with H if $\alpha \cap \Gamma = \emptyset$ and $i(a, \alpha) \leq 1$ for all $a \in \mathcal{A}$.

For α compatible with H , we obtain a new hexagon decomposition $H' = (\Gamma', \mathcal{A}')$ as follows. We set $\Gamma' = \Gamma \cup \alpha$. Now we consider \mathcal{A}' the collection of arcs obtained from \mathcal{A} which are either disjoint from α or which are obtained by surgering (cutting) an arc of \mathcal{A} at its intersection point with α , which results in two arcs with endpoints on Γ' . Note that surgery could result in multiple copies of the same (homotopy class of) arc.

The next lemma shows that \mathcal{A}' is maximal, and hence, H' is a hexagon decomposition.

Lemma 3.1 *The multiarc \mathcal{A}' is maximal on $\Sigma \setminus \Gamma'$.*

Proof The complementary region of H on Σ is a collection of hexagons. The curve α , being disjoint from Γ , crosses these hexagons in the sides corresponding to arcs of \mathcal{A} . The complementary regions to α on each hexagon are thus a collection of quadrilaterals and hexagons. The quadrilaterals correspond to a pair of subarcs of \mathcal{A} .

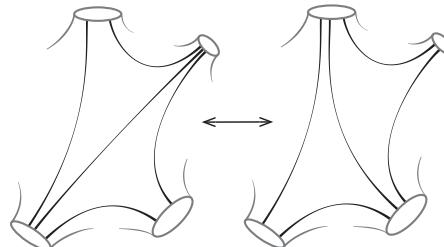


Figure 3: The illustration of a flip. Note that the curves of Γ the arcs terminate on need not be distinct.

that will be freely homotopic on $\Sigma \setminus \Gamma'$, where $\Gamma' = \Gamma \cup \alpha$, and so the resulting multiarc decomposes $\Sigma \setminus \Gamma'$ into hexagons. ■

We say that $H = (\Gamma, \mathcal{A})$ and $H' = (\Gamma', \mathcal{A}')$ are related by a curve addition if H' is obtained by adding a compatible curve α as described above.

Definition 3.2 The *topological* hexagon decomposition graph $\mathcal{H}(\Sigma)$ is the graph whose vertex set is given by hexagon decompositions and H, H' are joined by an edge if they are either related by a flip or a curve addition.

Remark 3.3 The operation of curve addition leads naturally to the notion of curve removal, which is just the formal opposite operation, meaning that H can be obtained from $H' = (\Gamma', \mathcal{A}')$ by curve removal if H' is obtained by a curve addition to H . However, in contrast with curve addition, the removal of a curve $\alpha \in \Gamma'$ is not uniquely defined by a choice of curve α . Indeed, consider α compatible with H , and \tilde{H} obtained by Dehn twisting H along α . Observe that α is also compatible with \tilde{H} . If adding α to H results in H' , then adding α to \tilde{H} also results in H' . In particular, this means that $\mathcal{H}(\Sigma)$ contains vertices of infinite valency and finite valency. Infinite valency comes from curve removal, and certain curves are not removable. As such, finite valency corresponds to hexagon decompositions without any removable curves.

3.2 Connectivity

First, we show that these graphs are connected.

For this purpose, a useful quantity will be the following. For any multicurve Γ , the subsurface $\Sigma' \subset \Sigma$ obtained by cutting Σ along Γ has an associated flip graph, which is the subgraph of $\mathcal{H}(\Sigma)$ consisting of all vertices of the form (Γ, \mathcal{A}) , and edges coming from flip relations between them. By standard results on flip graphs, this subgraph is connected (see Section 2.2). Furthermore, we can take its quotient by homeomorphisms that fix curves in Γ individually to obtain a finite “modular” graph (see [8, 21]). The diameter of this modular subgraph is bounded by a constant that only depends on the topology of Σ' and in fact on its arc complexity. Note that this only depends on the homeomorphism type of Γ .

Notation 3.4 We consider the maximum of these diameters among all homeomorphism types of (nonempty) Γ . We denote the maximum diameter of all such modular subgraphs by $D = D(\Sigma)$.

It is a useful quantity because it means that from any given hexagon decomposition with any given multicurve Γ , in at most D moves, we can choose the topological type of arcs. Note that if two hexagon decompositions have the same pants decomposition as a multicurve, then the subgroup of the mapping class group which fixes the curves in the pants decomposition individually acts only by Dehn twists along the pants curves and so acts trivially on the two hexagon decompositions. In particular, they are related by at most D flip moves (without passing to a quotient).

Lemma 3.5 Any $(\Gamma, \mathcal{A}) \in \mathcal{H}(\Sigma)$ is connected to a $(P, \mathcal{A}_P) \in \mathcal{H}(\Sigma)$, where P is a pants decomposition and $\Gamma \subset P$, in at most K moves where K is a constant that only depends on topology.

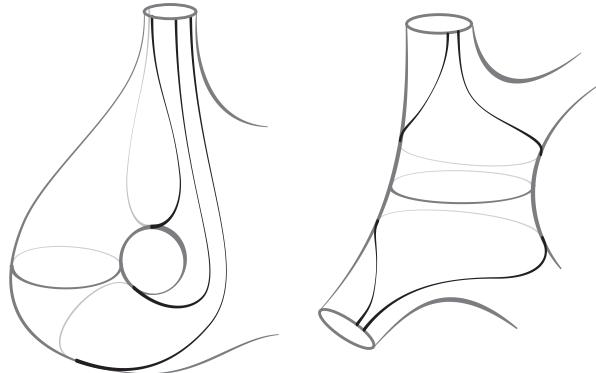


Figure 4: Arcs bounding embedded cylinders in the nonplanar and planar cases.

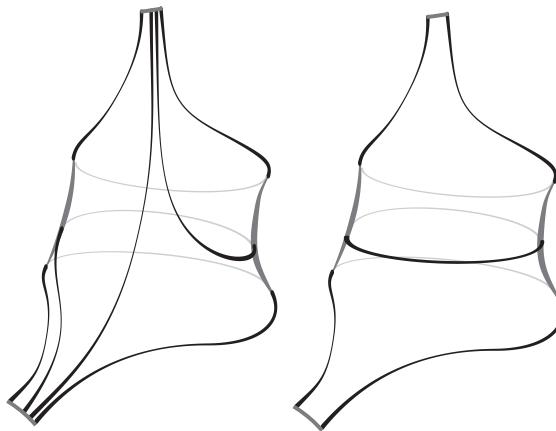


Figure 5: A hexagon decomposition of the cylinder and its core curve.

Proof The basic idea of the proof is to add curves one by one to Γ until we reach a full pants decomposition. To do so might require performing flip moves ahead of time.

If Γ is not a pants decomposition, there exists a connected component $\Sigma' \subset \Sigma \setminus \Gamma$ which is not a pair of pants. We focus our attention on this subsurface and show how to add a curve that lies inside Σ' to our multicurve Γ .

Observe that any such subsurface contains an embedded cylinder not peripheral to one of its boundary components. This cylinder can be bounded by two arcs and subarcs of boundary components.

Figure 4 portrays the two possible cases, depending on whether the surface is planar or not. We add arcs to the cylinder subsurface as in Figure 5.

By adding arcs, this can be completed into a hexagon decomposition. The core curve of the cylinder is now addable.

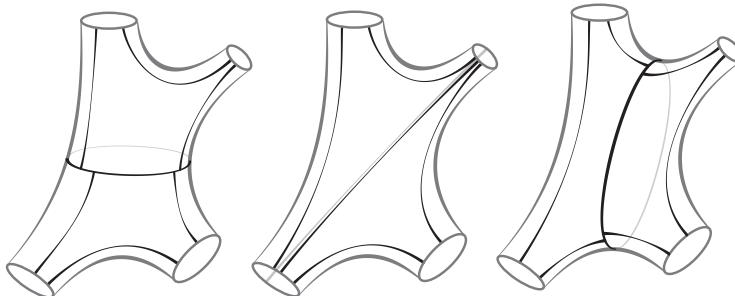


Figure 6: The first type of elementary move with hexagons.

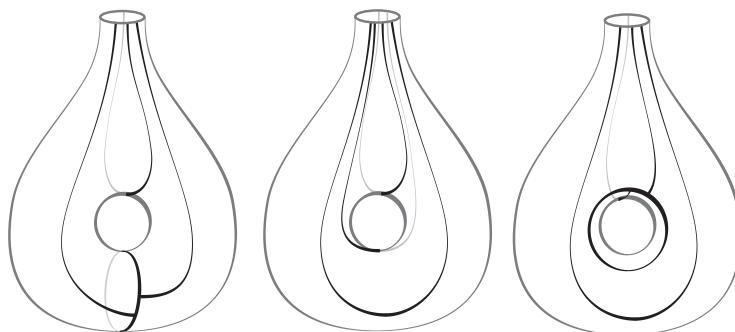


Figure 7: The second type of elementary move with hexagons.

Note that, as we are only interested in this configuration up to homeomorphism, such a hexagon decomposition can be reached in at most D moves, where D is the constant described above. Hence, a curve has been added in $D + 1$ moves. As there are at most $\kappa(\Sigma) - 1$ curves to add, we reach a full pants decomposition in at most $(D + 1)(\kappa(\Sigma) - 1)$ moves. ■

We now observe that moves in the pants graph can be emulated by moves in $\mathcal{H}(\Sigma)$. Consider two pants decompositions in $\mathcal{P}(\Sigma)$ that differ by an elementary move. To each, we consider a hexagon decomposition by adding arcs, and such the hexagon decompositions only differ in the subsurface where the elementary move takes place. This is either a four-holed sphere or a one-holed torus. We add arcs to the corresponding subsurface as in Figures 6 and 7.

Now we can perform the moves illustrated in the figure to emulate the elementary move and deduce connectedness.

Theorem 3.6 $\mathcal{H}(\Sigma)$ is connected.

Proof By Lemma 3.5, any vertex of $\mathcal{H}(\Sigma)$ is connected to a hexagon decomposition with a full pants decomposition. As elementary moves in the pants graph can be

emulated in $\mathcal{H}(\Sigma)$, connectivity of $\mathcal{H}(\Sigma)$ then follows from the connectivity of the pants graph. \blacksquare

We end this subsection with a basic result about curve addition. It is a consequence of the following observation which is a standard fact about normal coordinates for triangulations.

Observation 3.7 *Let $(\Gamma, \mathcal{A}) \in \mathcal{H}(\Sigma)$. Let γ be a curve disjoint from Γ . Then γ is uniquely determined by its intersection with arcs in \mathcal{A} .*

Note that it is also true for any simple multicurve.

Lemma 3.8 *For any $(\Gamma, \mathcal{A}) \in \mathcal{H}(\Sigma)$, there are at most K curves that can be added to (Γ, \mathcal{A}) , where K depends only on the topology of Σ .*

Proof Given (Γ, \mathcal{A}) , only the compatible curves can be added to Γ . By definition, we may add only the curves which intersect each arc in \mathcal{A} at most once; hence, a compatible curve can intersect each triangle in a triangulation at most twice. Hence, there are at most $2^{|\mathcal{A}|} = 2^{\kappa_a}$ ways to add a curve to Γ . \blacksquare

Note the strong contrast to curve removal.

3.3 A quasi-isometry with the pants graph

In this subsection, we prove the following.

Theorem 3.9 *The graphs $\mathcal{H}(\Sigma)$ and $\mathcal{P}(\Sigma)$ are quasi-isometric.*

Before passing to the proof, we describe two maps between $\mathcal{H}(\Sigma)$ and $\mathcal{P}(\Sigma)$.

The map $\phi : \mathcal{P}(\Sigma) \rightarrow \mathcal{H}(\Sigma)$

We define the map $\phi : \mathcal{P}(\Sigma) \rightarrow \mathcal{H}(\Sigma)$ as follows. Given $P \in \mathcal{P}(\Sigma)$, we add a collection of arcs to obtain an element of $\mathcal{H}(\Sigma)$. That is, $\phi(P) = (P, \mathcal{A})$. The choice of \mathcal{A} is arbitrary, but notice that any two choices are a bounded distance apart because the flip graph of a pair of pants is finite (see Section 2.2). In fact, they are at most distance D apart, where D is the diameter constant defined previously (see Notation 3.4 and the comment that follows). As an example of a choice of \mathcal{A} , one can add arcs to each pair of pants as in Figure 8.

The map $\psi : \mathcal{H}(\Sigma) \rightarrow \mathcal{P}(\Sigma)$ This second map is less obvious and takes a hexagon decomposition and associates to it a pants decomposition. To define it, we shall use

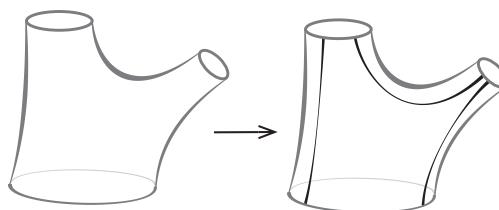


Figure 8: A local picture of the map ϕ .

the strategy of Lemma 3.5 – that is, to add curves until the multicurve is a full pants decomposition. It is somewhat loosely defined, but this is okay, as we will only need possible images to be at a bounded distance one from one another.

Recall the strategy: Given $H = (\Gamma, \mathcal{A}) \in \mathcal{H}(\Sigma)$, where Γ is not a pants decomposition, we add a curve to Γ . If it is not possible to add a curve, as shown in the proof of Lemma 3.5, by performing a minimal number of flips, we can add a curve, and then we repeat. Note that the number of necessary flips before adding a curve is at most D . The whole process ends in at most $(D + 1)(\kappa(\Sigma) - 1)$ steps. We define $\psi(H)$ to be the pants decomposition thus obtained. Observe that $\psi(\phi(P)) = P$.

We now proceed to show that these maps satisfy certain properties that will lead us to the conclusion of Theorem 1.1. We denote by $d_{\mathcal{H}}$ distance in $\mathcal{H}(\Sigma)$, and by $d_{\mathcal{P}}$ distance in $\mathcal{P}(\Sigma)$.

Lemma 3.10 *There exists a constant $C_1 = C_1(\Sigma)$ that depends only on the topology of Σ such that for all $P, P' \in \mathcal{P}(\Sigma)$, we have*

$$d_{\mathcal{H}}(\phi(P), \phi(P')) \leq C_1 d_{\mathcal{P}}(P, P').$$

Proof As explained previously, moves in $\mathcal{P}(\Sigma)$ can be emulated by moves in $\mathcal{H}(\Sigma)$ (see Figures 6 and 7), hence the result. ■

Note that if ϕ is defined exactly as in Figure 8, then C_1 above can be taken to be 2. Otherwise, it might require performing flips inside each pair of pants beforehand (but certainly at most D times the number of pairs of pants which are both functions of the topology).

Lemma 3.11 *There exists a constant $K = K(\Sigma)$, which only depends on the topology of Σ , such that the full image $\phi(\mathcal{P}(\Sigma))$ of the map ϕ is K -dense in $\mathcal{H}(\Sigma)$.*

Proof This follows from the proof of connectivity where we showed that any hexagon decomposition H can be transformed into a hexagon decomposition H' with a full pants decomposition P in at most $(D + 1)(\kappa(\Sigma) - 1)$ steps. Now the map ϕ takes P to a hexagon decomposition $\phi(P)$. We can now observe that H' and $\phi(P)$, which both share the same multicurve P , are related by at most D moves, just by flips in the flip graphs of the complementary pants. This proves the claim. ■

We now focus on the map ψ . The key step will be to show that it is “quasi” well-defined – that is, that any two possible choices of image are close to each other. The general idea is as follows: we show that any two choices of pants decompositions in the image have bounded intersection with a bound that depends only on the topology of the surface. From there, we will be able to conclude that any two potential image pants decompositions are a bounded distance away, again with a bound that only depends on the topology of the surface.

We begin by showing that an image pants decomposition has bounded intersection with the original multiarc.

Lemma 3.12 *Let $P = \psi(H)$. Then $i(P, H) \leq B$, where B is function of the topology of Σ .*

Proof As P contains the multicurve of H , we need to show it intersects the multiarc of H a bounded number of times.

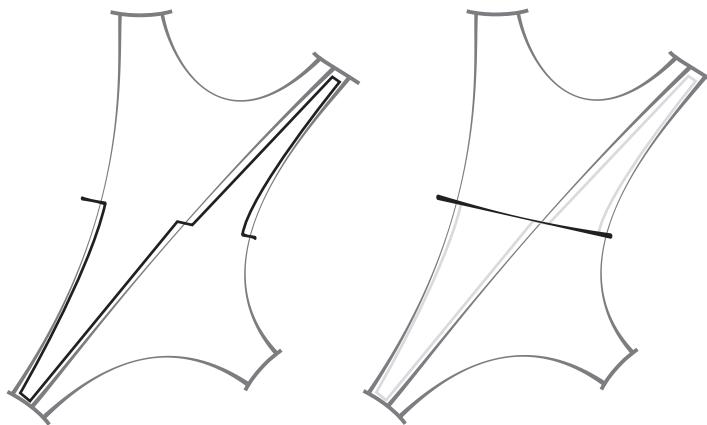


Figure 9: An arc of α_1 crossing two adjacent hexagons and its representation by a broken path on the left.

Given $H = (\Gamma_0, \mathcal{A}_0)$, we perform at most D flips on \mathcal{A}_0 to obtain a new collection of arcs, say \mathcal{A}'_0 . We then add a curve, say α_1 . Then because \mathcal{A}_0 and \mathcal{A}'_0 are related by at most D flips, the (total) intersection between \mathcal{A}_0 and \mathcal{A}'_0 is bounded above by a function of the topology of Σ , say I_1 . An explicit value for I_1 can be deduced from [8, Corollary 2.18].

We now make the following useful observation:

Observation: As α_1 intersects each arc of \mathcal{A}'_0 at most 1 time, it also intersects \mathcal{A}_0 at most I_1 times.

To show this, note that α_1 can be thought of as a concatenation of segments passing through the hexagons bounded by \mathcal{A}'_0 (at most one per hexagon). Now the homotopy class of each segment can be represented by a concatenation of at most 3 segments lying on the boundary of a hexagon (see Figure 9).

Then, by retracting back and forth segments if necessary, α_1 can be represented by a concatenation of arcs of \mathcal{A}'_0 and boundary segments of Γ , where each arc of \mathcal{A}'_0 appears at most one time. From this, we can deduce that the total intersection between α_1 and \mathcal{A}_0 is at most that of \mathcal{A}'_0 and \mathcal{A}_0 , and hence is at most I_1 . We now add α_1 to obtain a new hexagon decomposition, $(\Gamma_1, \mathcal{A}_1)$ with $\Gamma_1 = \Gamma_0 \cup \{\alpha_1\}$ and \mathcal{A}_1 the resulting hexagon decomposition when adding α_1 to \mathcal{A}'_0 . Observe that the total intersection between arcs in \mathcal{A}_1 and arcs in \mathcal{A}_0 is still bounded above by I_1 .

We now repeat the above process to add a second curve and so on. At each step, we might be required to perform flips (at most D), which may further increase the intersection between the arcs in \mathcal{A}_0 and the arcs \mathcal{A}'_k (the arcs at step k), but there is always a bound (denoted I_k) on this intersection which only depends on topology and the number of steps. Using the same trick as before, we add a curve α_k which intersects \mathcal{A}_0 at most I_k times. The process ends in at most $m \leq \kappa_c(\Sigma)$ steps, resulting in a pants decomposition where each curve intersects the original multiarc, and hence H , at most $B = I_m$ times. ■

The following lemma is our last missing ingredient.

Lemma 3.13 Consider $H = (\Gamma, \mathcal{A})$ on a surface Σ . For K a constant, and let P, Q be pants decompositions which both contain Γ , and such that $\max\{i(P, \mathcal{A}), i(Q, \mathcal{A})\} \leq K$. Then the pants distance between P and Q is bounded above by $R(K)$, a function of the topology of Σ and K .

Proof This basically follows from a finiteness argument. Any multicurve on Σ , disjoint from Γ , is uniquely determined by its intersection numbers with arcs in \mathcal{A} . As we have a bound on these numbers (K), there are finitely many such pairs of multicurves, and thus finitely many choices for P and Q . Among these, there is a pair at maximal distance. This quantity only depends on the topology of Σ , the topological type of Γ , and K . Taking again the maximum among all topological types of Γ (the number of these again only depends on the topology of Σ), we obtain an upper bound on the possible distance between P and Q , as claimed. ■

Now, using our upper bound of I_m on the intersection of a pants decomposition with $H = (\Gamma_0, \mathcal{A}_0)$, and applying the above lemma, we can conclude the following:

Proposition 3.14 Given H , any two possible choices of pants decompositions P and Q as images of H by ψ are at distance at most $R(I_m)$ in $\mathcal{P}(\Sigma)$.

Now consider H and H' at distance 1 in $\mathcal{H}(\Sigma)$. If they are related by a flip, by the argument given above, then their images by ψ are at a bounded distance apart, where this distance depends only the topology of Σ . If they are related by a curve addition, then their potential images intersect, and so any two choices will remain at a bounded distance apart. As such, we proved the following:

Lemma 3.15 There exists a constant $C_2 = C_2(\Sigma)$ that depends only on the topology of Σ such that for all $H, H' \in \mathcal{H}(\Sigma)$, we have

$$d_{\mathcal{P}}(\psi(H), \psi(H')) \leq C_2 d_{\mathcal{H}}(H, H').$$

Finally, we observe that ϕ and ψ are quasi-inverses. Indeed, we have $\psi(\phi(H)) = H$ for all $H \in \mathcal{P}(\Sigma)$, and $\phi(\psi(H))$ is at most a constant distance from H (where the constant only depends on topology). Putting these results together, we have now shown that the maps ϕ and ψ are quasi-isometries, hence showing Theorem 3.9.

4 The geometric hexagon graph

By adding additional data to each hexagon decomposition, we construct a new graph which will give us a quasi-isometric model for the mapping class group. The graph will be associated to a (fixed) hyperbolic metric X which is homeomorphic to Σ , and denoted $\mathcal{H}(X)$.

4.1 Vertices and edges.

For this graph, we require that curves have an orientation. We fix an orientation on each curve (it can be arbitrary).

The vertices of $\mathcal{H}(X)$ are hexagon decompositions as before, but with an integer (called a weight) prescribed to each curve in the multicurve. Thus precisely, vertices are weighted hexagon decompositions $H = (\Gamma^w, \mathcal{A})$, where $\Gamma^w = (\Gamma, w)$ – that is, a

multicurve $\Gamma = (\gamma_1, \dots, \gamma_m)$ with a set of weights $w \in \mathbb{Z}^m$ – and \mathcal{A} a maximal multiarc on $X \setminus \Gamma$.

Note that if a curve α , with an orientation, is given weight k , this corresponds to α^{-1} (that is, α with the opposite orientation) with given weight $-k$. In other words, we can think of each unoriented curve as an equivalence class and (α, k) as a representative. Finally, weight 0 is not particular, other than it is the only weight for which $(\alpha, 0)$ and $(\alpha^{-1}, 0)$ correspond to the same class.

Edge relations are defined by extending the notion of edges as before.

- i) Flip relation: two weighted decompositions are related by an edge if they differ by a flip. Note that a flip retains the same weights.
- ii) Curve addition: as before, but now we need to define how to associate a weight to an added curve. This is explained below.
- iii) Let $w, w' \in \mathbb{Z}^m$ differ by only one coordinate that differs by 1 (expressed otherwise, $\|w - w'\|_1 = 1$). Then there is an edge between (Γ^w, \mathcal{A}) and $(\Gamma^{w'}, \mathcal{A})$. In other words, changing a single weight by 1 induces an edge. (This can be thought of as a type of twist relation.)

The weight of an added curve. This is where we use the hyperbolic metric X .

Let α be compatible to a given hexagon decomposition (Γ^w, \mathcal{A}) – that is, if $\alpha \cap \Gamma = \emptyset$ and $i(\alpha, a) \leq 1$ for every arc $a \in \mathcal{A}$.

Note that all arcs $a \in \mathcal{A}$ with $i(a, \alpha) = 1$ split into two arcs a_1, a_2 in $X \setminus \{\Gamma \cup \alpha\}$. These arcs are obtained by surgering a at its *geometric* intersection point with α – that is, the intersection point between the geodesic representatives. Let a_1 and a_2 be the resulting (homotopy classes of) arcs. Both a_1 and a_2 have exactly one of their endpoints on α .

We can equivalently think of the homotopy classes of arcs a_1 and a_2 as being realized by unique orthogeodesics (using the hyperbolic metric X). This gives rise to a converse construction of a , starting from the arcs a_1 and a_2 . To begin with, we give α an orientation. We orient a_1 so that it ends on α , and a_2 so that it begins on α . The (homotopy class of) arc a is obtained via the orthogeodesics a_1 and a_2 by first following a_1 , then following α (in the positive direction) for a certain length (which we denote by $t_\alpha(a)$) and then by following a_2 (see Figure 10).

The result is an oriented arc whose unoriented version is a . Note that the quantity $t_\alpha(a)$ does not depend on our choice of a_1 and a_2 , but only on the orientation of α .

Let α be a compatible curve that we add to (Γ^w, \mathcal{A}) . This results in $((\Gamma')^{w'}, \mathcal{A}')$, where

- (i) $\Gamma' = \Gamma \cup \{\alpha\}$
- (ii) \mathcal{A}' consists of all arcs of \mathcal{A} that do not intersect α , to which we add all arcs of \mathcal{A} split by α .
- (iii) The weights of the curves in Γ' are defined as follows:

For all curves that were previously in Γ , the weights remain unchanged.

The weight of α relative to an arc $a \in \mathcal{A}$ is defined as

$$w(\alpha, a) = \{ \lfloor t_\alpha(a) / \ell_X(\alpha) \rfloor \}.$$

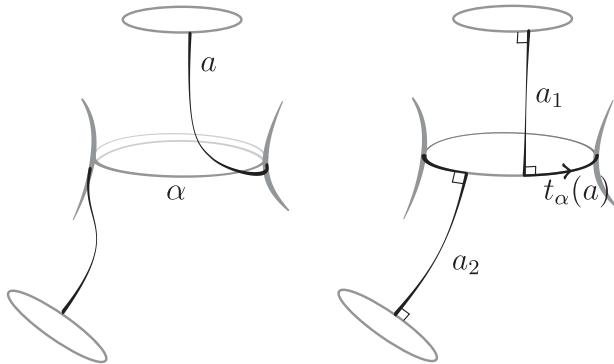


Figure 10: An arc a and its representation by a concatenation of arcs.

Here, $\ell_X(\alpha)$ is the length of the unique closed geodesic freely homotopic to α for the hyperbolic metric X , and $t_\alpha(a)$ is defined above. In the example of Figure 10, we have $w(\alpha, a) = 0$.

In order to make the weight independent of the choice of arc, we set

$$w(\alpha) = \max_{a \in \mathcal{A} \setminus \mathcal{A}'} \{ |t_\alpha(a)/\ell_X(\alpha)| \},$$

where a runs over all arcs of \mathcal{A} split by α .

Roughly speaking, the weight, associated to a choice of arc a , is defined as the (whole) twist parameter of α necessary to obtain the arc a by concatenating a_1 and a_2 with that number of copies of α . (Said otherwise, a is obtained by a concatenation of a_1 , then following α a certain number of times and then a_2 .) As there are multiple arcs to choose from, we take the maximum value to avoid any ambiguity. This is an arbitrary choice, and we could have chosen the minimum, for instance. In fact, in the next lemma, we show that for any two arcs, the values obtained will differ by at most one.

Lemma 4.1 *Weights of arcs differ by at most 1.*

Proof Let a, b be disjoint arcs in \mathcal{A} and α a simple closed curve that intersects both a and b once.

As a and b are disjoint, it follows that $t_\alpha(a)$ and $t_\alpha(b)$ cannot differ by more than $\ell_X(\alpha)$. Otherwise, one of the arcs would wrap (at least) one more time around α than the other, and the arcs would intersect. Hence,

$$|t_\alpha(a) - t_\alpha(b)| \leq \ell_X(\alpha),$$

and then by definition,

$$|w(\alpha, a) - w(\alpha, b)| = \frac{|t_\alpha(a) - t_\alpha(b)|}{\ell_X(\alpha)} \leq 1,$$

as claimed. ■

We finish this subsection with an upper bound on the degree of vertices.

Theorem 4.2 *Any vertex in $\mathcal{H}(X)$ is of finite valency F , where F depends only on the topology of Σ .*

Proof Let (Γ^w, \mathcal{A}) be a vertex. We want to bound the number of edges that leave from this vertex. Edges leaving this vertex correspond to exactly one of the different types of elementary move:

- (i) Flipping arcs of \mathcal{A} . This contributes at most $|\mathcal{A}|$ edges.
- (ii) Each weight can change by ± 1 , thus contributing at most $2|\Gamma|$ edges.
- (iii) Curve addition: by Lemma 3.8, there are most $2^{|\mathcal{A}|}$ such curves.
- (iv) Curve removal: this is the part which requires some work, as curve removal is defined as the opposite of curve addition. We need to estimate how many different vertices can result in (Γ^w, \mathcal{A}) after a curve addition.

We consider $\alpha \in \Gamma$, and we look at how it might have appeared under a curve addition.

Note there are possibly an infinite number of hexagon decompositions of $X \setminus (\Gamma \setminus \alpha)$ that would topologically result in (Γ, \mathcal{A}) by adding α but only finitely many can result in (Γ^w, \mathcal{A}) with the correct weight along α . We will show this by bounding the number of possible arcs of a candidate hexagon decomposition.

Consider a hexagon decomposition such that adding α to it results in (Γ^w, \mathcal{A}) . Any arc in such a hexagon decomposition must have been obtained by concatenating an arc of \mathcal{A} , say a_1 , with an arc c belonging to α , and then another arc of \mathcal{A} , say a_2 . This comes from the fact that an added curve can intersect arcs of the hexagon decomposition at most once. There are at most $|\mathcal{A}|$ choices for both a_1 and a_2 , and thus at most $|\mathcal{A}|^2$ choices for the pair. Now because the weight of α is known, there are at most 2 choices for the arc c : thanks to Lemma 4.1, the arc must wrap the prescribed number of times around α , or possibly once less.

As such, there are at most $2|\mathcal{A}|^2$ choices for each arc. The full hexagon decomposition requires at most N new arcs where

$$N = \kappa_a(X \setminus (\Gamma \setminus \alpha)).$$

There are at most

$$\binom{2|\mathcal{A}|^2}{N}$$

choices. This was for a given choice of $\alpha \in \Gamma$. Multiplying by $|\Gamma|$, which is bounded above by the curve complexity of X , completes the upper bound on possible vertices that are related to (Γ^w, \mathcal{A}) by a curve removal.

The quantity F is the sum of all of these upper bounds. ■

4.2 $\mathcal{H}(X)$ is a model for the mapping class group

With these basic properties in hand, we will now be able to show the following.

Theorem 4.3 *Let X be a closed surface with hyperbolic metric. Then the mapping class group of X acts on $\mathcal{H}(X)$ by quasi-isometries and with finite quotient. As such, $\mathcal{H}(X)$ is a quasi-model for the mapping class group.*

To prove this theorem, we will use the following lemmas, the first of which is a version of the Milnor-Schwarz lemma for quasi-isometric actions. Although it is not as used as its isometric action counterpart, it is certainly well-known to experts. For a proof, see, for instance, [7, Theorem 4.40].

Lemma 4.4 [Milnor-Schwarz] *Let G be a group and X a geodesic metric space on which G acts with (\tilde{K}, \tilde{C}) uniform quasi-isometries. Assume that the action is properly discontinuous and cocompact. Then, G is finitely generated, and there are constants $K \geq 1, C \geq 0$ such that G is (K, C) -quasi-isometric to X .*

Let $\text{Mod}(X)$ ($= \text{Mod}(\Sigma)$) be the mapping class group of X (or of Σ). The action of the mapping class group on $\mathcal{H}(\Sigma)$ is straightforward: it sends multi-curves to multi-curves and multi-arcs to multi-arcs by homeomorphisms. Note that $(\Gamma, \mathcal{A}) \in \mathcal{H}(\Sigma)$ always has infinite stabilizers because of Dehn twists along curves in Γ . The weights on vertices will prevent this from happening in $\mathcal{H}(X)$.

Lemma 4.5 *$\text{Mod}(X)$ acts on $\mathcal{H}(X)$ is by quasi-isometries.*

Proof In addition to the simplicial action on $\mathcal{H}(\Sigma)$, we need to define how the mapping class group acts on weights. The action we describe is a quasi-action, and we describe it at a vertex (Γ^w, \mathcal{A}) for an individual weighted curve α , where $\alpha \in \Gamma^w$. Consider a mapping class $\varphi \in \text{Mod}(X)$. A weight needs to be attributed to $\varphi(\alpha)$ for a homeomorphism φ . If the homeomorphism is a Dehn twist along α , the weight changes by ± 1 depending on the direction of twist and orientation of α . More generally, we can compute the weight of the image as follows.

Let α be of weight $k \in \mathbb{Z}$. We consider a pair of disjoint simple arcs $a_1, a_2 \in \mathcal{A}$ with exactly one endpoint on either side of α . If no such pair exists, by performing a bounded number of flips (as in the proof of Lemma 3.5), one can ensure that this is the case. (Note that in the case of a one-holed torus, no such pair exists, and in that case, a similar argument works with a single arc, but we omit the details here. We are doing this to avoid handling the case where α is separating or nonseparating for all other surfaces.) Now we consider the arc a obtained by concatenating a_1 (so that it ends on α), then following α exactly k times around and then if necessary until the endpoint of a_2 , and then following a_2 . We now consider its image $\varphi(a)$ by the mapping class, which is an arc that intersects $\varphi(\alpha)$ exactly once. We will now compute the weight we will associate to $\varphi(\alpha)$ in order for φ to have a well-defined quasi-action. To do so, we cut $\varphi(a)$ in its intersection point to obtain a'_1 and a'_2 and measure the length of the arc t' needed to concatenate to get a (as in Figure 10). The weight of $\varphi(\alpha)$ is

$$\frac{t'}{\ell_X(\varphi(\alpha))}.$$

We now need to show that this weight does not depend (too much) on the choice of a_1 and a_2 .

The idea is similar to what happens in Lemma 4.1. There were different choices for a_1 and a_2 , but all resulting in arcs that are either disjoint, or possibly intersect at most once. Hence, as in the proof of Lemma 4.1, the weights associated to any such choice differ by at most ± 2 . The difference here is that the arcs might intersect one

time, whereas in the proof of the lemma, they were necessarily disjoint. This shows that the action is (quasi-)well-defined.

Now to show that the mapping class group acts by quasi-isometries, we need to show that two vertices at distance 1 are sent to vertices a (uniformly) bounded distance apart.

- (i) If two vertices x, y differ by a flip, then $\varphi(x)$ and $\varphi(y)$ also differ (topologically) by a flip. Furthermore, the weights associated to each curve can be computed using the same arcs (up to a single flip) and so, by the above argument, are comparable (they differ by at most 2).
- (ii) If two vertices x, y differ by a curve addition, then the weights of a curve in $\varphi(x)$ and $\varphi(y)$ can be computed using arcs that are either the same or that intersect at most once. Hence, as before, the weights are always comparable.

In all situations, the weights are uniformly comparable, and since the number of curves only depends on the curve complexity of the surface, this shows that if two vertices x, y are distance d apart, then $\varphi(x)$ and $\varphi(y)$ are distance at most Cd apart where C only depends on the topology of the surface. ■

The quotient of $\mathcal{H}(X)$ by $\text{Mod}(X)$ is exactly the quotient of $\mathcal{H}(\Sigma)$ by $\text{Mod}(\Sigma)$. This is because any two hexagon decompositions that only differ by weights are related by Dehn twists along their curves and so correspond to the same point in $\mathcal{H}(\Sigma)$. The quotient thus has a vertex set corresponding to the topological types of hexagon decompositions which is a finite set and has at most one edge between any two vertices. As such, the action is cocompact.

By Theorem 4.2, $\mathcal{H}(X)$ has finite degree, and so the action of $\text{Mod}(X)$ has finite stabilizers. Hence, the action of the mapping class group by quasi-isometries is proper and cocompact, and by Lemma 4.4, Theorem 4.3 follows.

5 Perspectives

As the graphs we introduce are shown to be quasi-isometric to other well-studied topological graphs, it is worth pondering why one would want to study them at all. We briefly point out some of their features and perspectives.

A less coarse model for the Weil-Petersson metric

Our proof that $\mathcal{H}(\Sigma)$ is quasi-isometric to the Weil-Petersson metric follows from Brock's theorem [3], which shows the corresponding result for the pants graph. There should, however, be a more explicit way of showing this. An essential ingredient in Brock's proof is the existence of a Bers constant – that is, a bound on the length of a shortest pants decomposition that depends only on the topology. In our setting, an equivalent theorem would be about the “length” of hexagon decompositions. One has to be careful when dealing with surfaces with short curves, but there is a way of quantifying length by focusing on the length of arcs in the thick part of the surface. This is done in [18] where the bounds are logarithmic in terms of the topology. This is in strong contrast to bounds on the lengths of pants decompositions which are linear in terms of the genus and cannot be better than square root (see [19, 20] and references therein). This suggests that the hexagon decomposition graph $\mathcal{H}(\Sigma)$ is

more evenly distributed in Weil-Petersson space and could be used to study some of its finer properties.

Flip graphs for closed surfaces

A second motivation comes from wanting to have a model which mimics some of the nice properties of flip graphs (which do not exist for closed surfaces). Flip graphs have proved to be an interesting tool for studying mapping class groups. One can hope that the geometric hexagon graphs $\mathcal{H}(X)$ will enable similar interpretations for closed surfaces.

Acknowledgements We thank Mark Bell, Mark Hagen, Chris Leininger, Pallavi Panda, and Binbin Xu for fruitful discussions and comments. We are extremely grateful to the referee for making multiple wise suggestions which have considerably improved our manuscript. Thanks!

References

- [1] M. C. Bell, V. Disarlo and R. Tang, *Cubical geometry in the polygonalisation complex*. Math. Proc. Cambridge Philos. Soc. 167(2019), no. 1, 1–22.
- [2] J. A. Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space*. Geom. Topol. 10(2006), 1523–1578.
- [3] J. F. Brock, *The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores*. J. Amer. Math. Soc. 16(2003), no. 3, 495–535.
- [4] J. Brock and B. Farb, *Curvature and rank of Teichmüller space*. Amer. J. Math. 128(2006), no. 1, 1–22.
- [5] J. Brock and H. Masur, *Coarse and synthetic Weil-Petersson geometry: quasi-flats, geodesics and relative hyperbolicity*. Geom. Topol. 12(2008), no. 4, 2453–2495.
- [6] J. Brock and D. Margalit, *Weil-Petersson isometries via the pants complex*. Proc. Amer. Math. Soc. 135(2007), no. 3, 795–803.
- [7] M. Cordes, H. Tobias and V. Tonić, *Foundations of geometric approximate group theory*. Preprint, 2020.
- [8] V. Disarlo and H. Parlier, *The geometry of flip graphs and mapping class groups*. Trans. Amer. Math. Soc. 372(2019), no. 6, 3809–3844.
- [9] F. Gültepe and C. J. Leininger, *An arc graph distance formula for the flip graph*. Proc. Amer. Math. Soc. 145(2017), no. 7, 3179–3184.
- [10] A. Hatcher, *On triangulations of surfaces*. Topol. Appl. 40(1991), no. 2, 189–194.
- [11] A. Hatcher and W. Thurston, *A presentation for the mapping class group of a closed orientable surface*. Topology 19(1980), no. 3, 221–237.
- [12] D. Margalit, *Automorphisms of the pants complex*. Duke Math. J. 121(2004), no. 3, 457–479.
- [13] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*. Invent. Math. 138(1999), no. 1, 103–149.
- [14] 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.
- [15] H. Masur and S. Schleimer, *The geometry of the disk complex*. J. Amer. Math. Soc. 26(2013), no. 1, 1–62.
- [16] H. Masur and S. Schleimer, *The pants complex has only one end*. Spaces of Kleinian groups, London Math. Soc. Lecture Note Ser., Vol. 329, Cambridge Univ. Press, Cambridge, 2006, 209–218.
- [17] L. Mosher, *Mapping class groups are automatic*. Ann. of Math. (2) 142(1995), no. 2, 303–384.
- [18] H. Parlier, *Interrogating surface length spectra and quantifying isospectrality*. Math. Ann. 370(2018), no. 3–4, 1759–1787.
- [19] H. Parlier, *A short note on short pants*. Canad. Math. Bull. 57(2014), no. 4, 870–876.
- [20] H. Parlier, *A shorter note on shorter pants*. Bull. Lond. Math. Soc. 56(2024), no. 4, 1483–1487.
- [21] H. Parlier and L. Pournin, *Flip-graph moduli spaces of filling surfaces*. J. Eur. Math. Soc. (JEMS) 19(2017), no. 9, 2697–2737.
- [22] L. Pournin, *The diameter of associahedra*. Adv. Math. 259(2014), 13–42.

Department of Mathematics and Statistics, University of Toledo, Toledo, OH, 43606

e-mail: funda.gultepe@utoledo.edu

URL: <http://www.math.utoledo.edu/~fgultepe/>

Department of Mathematics, University of Luxembourg, Esch-sur-Alzette, Luxembourg

e-mail: hugo.parlier@uni.lu

URL: <http://math.uni.lu/parlier/>