

CUBICAL GEOMETRY IN THE POLYGONALISATION COMPLEX

MARK C. BELL, VALENTINA DISARLO, AND ROBERT TANG

ABSTRACT. We introduce the polygonalisation complex of a surface, a cube complex whose vertices correspond to polygonalisations. This is a geometric model for the mapping class group and it is motivated by works of Harer, Mosher and Penner. Using properties of the flip graph, we show that the mid-cubes in the polygonalisation complex can be extended to a family of embedded and separating hyperplanes, parametrised by the arcs in the surface.

We study the crossing graph of these hyperplanes and prove that it is quasi-isometric to the arc complex. We use the crossing graph to prove that, generically, different surfaces have different polygonalisation complexes. The polygonalisation complex is not $CAT(0)$, but we can characterise the vertices where Gromov's link condition fails. This gives a tool for proving that, generically, the automorphism group of the polygonalisation complex is the (extended) mapping class group of the surface.




FIGURE 1. The polygonalisation complex of a hexagon $S_{0,0}^6 = \text{hexagon}$.

1. INTRODUCTION

In this paper we study the polygonalisation complex $\mathcal{P}(S)$ of a surface S with marked points. This is a cube complex encoding the combinatorics of the polygonalisations of S , that is, the multiarcs that decompose S into polygons. The polygonalisation complex contains (the barycentric subdivision of) the flip graph $\mathcal{F}(S)$ as a subcomplex. It can also be regarded as the barycentric subdivision of the contractible CW-complex naturally associated to the Ptolemy groupoid of S which appears in quantum Teichmüller theory [8] [19].

The mapping class group $\text{Mod}^\pm(S)$ acts geometrically on $\mathcal{P}(S)$. However, in general, $\text{Mod}^\pm(S)$ is not a CAT(0) group [13, Theorem 4.2] and so $\mathcal{P}(S)$ is not a CAT(0) cube complex. We show that, in spite of this, $\mathcal{P}(S)$ has many of the properties of CAT(0) cube complexes. In particular, it has a rich hyperplane structure that is closely related to the arc complex.

Theorem 3.8. *There is a natural one-to-one correspondence between the hyperplanes of $\mathcal{P}(S)$ and the arcs on S . Moreover, for each arc $\alpha \in \mathcal{A}(S)$, the corresponding hyperplane H_α is embedded, two-sided and separates $\mathcal{P}(S)$ into two connected components: $\mathcal{P}_\alpha(S)$ and $\overline{\mathcal{P}_\alpha(S)}$.*

One can encode the combinatorics of the hyperplanes in $\mathcal{P}(S)$ via its *crossing graph* $\text{Cr}(\mathcal{P}(S))$: this has a vertex for every hyperplane in $\mathcal{P}(S)$, and two hyperplanes are connected via an edge if and only if they cross.

Proposition 3.12. *Let H_α and H_β be the hyperplanes of $\mathcal{P}(S)$ corresponding to arcs $\alpha, \beta \in \mathcal{A}(S)$. Then H_α and H_β cross if and only if α and β are disjoint and do not bound a (folded) triangle.*

Together, Theorem 3.8 and Proposition 3.12 show that there is a natural embedding of $\text{Cr}(\mathcal{P}(S))$ into the arc graph $\mathcal{A}(S)$. This embedding is in fact a quasi-isometry (Corollary 3.16). In particular, when $\partial S = \emptyset$ the arc graph is Gromov hyperbolic [15, Theorem 20.2], hence the crossing graph is also. Proposition 3.12 also enables us to characterise the edges of the arc graph that do not appear in the crossing graph (Lemma 3.17). We show that $\mathcal{A}(S)$ can be recovered from the combinatorics of $\text{Cr}(\mathcal{P}(S))$. Applying rigidity results for arc graphs [12] [5] [14] we then obtain:

Theorem 3.18. *For all but finitely many pairs of surfaces, the complexes $\mathcal{P}(S)$ and $\mathcal{P}(S')$ are isomorphic if and only if S and S' are homeomorphic.*

We rephrase Gromov's link condition in terms of curves on S (Definition 4.1). We use this to extract the number of components of a given polygonalisation from the local combinatorics of $\mathcal{P}(S)$.

Corollary 4.5. *For each k , there is a combinatorial criterion that characterises the vertices of $\mathcal{P}(S)$ corresponding to polygonalisations with exactly k arcs.*

This gives a method for proving the rigidity of the polygonalisation complex via flip graph rigidity.

Theorem 4.7. *For all but finitely many surfaces, the natural homomorphism*

$$\text{Mod}^\pm(S) \rightarrow \text{Aut}(\mathcal{P}(S))$$

is an isomorphism.

We list the exceptions to these theorems in Appendix C. Additionally, we highlight these results do not follow from [1, Theorem 1.1] since $\mathcal{P}(S)$ does not satisfy the required rigidity axioms.

2. PRELIMINARIES

Let S be a connected, compact, orientable surface with a finite non-empty set of marked points. We assume that each boundary component (if any) contains at least one marked point. When $\partial S \neq \emptyset$, such surfaces are also known in the literature as ciliated [7, Section 2]. Let

$$E(S) := 6g + 3b + 3s + p - 6 \text{ and } F(S) := 4g + 2b + 2s + p - 4$$

where g is the genus of S , s is the number of marked points in its interior, b is the number of boundary components and p is the number of marked points on ∂S .

To avoid pathologies, from now on we will require that $F(S) \geq 3$. The *exceptional surfaces*, for which $F(S) < 3$, are listed and discussed in Appendix C.

2.1. **Objects.** We recall some standard objects that will appear throughout:

2.1.1. *Mapping class group.* The (*extended*) *mapping class group* $\text{Mod}^\pm(S)$ is the group of homeomorphisms of S relative to the marked points up to isotopy. We allow mapping classes to reverse orientation, permute the marked points and permute the boundary components of S .

2.1.2. *Triangulations.* An (*essential*) *arc* α on S is an embedded arc connecting marked points up to isotopy (relative to the set of marked points). Such arcs are not *null-homotopic* and are not *boundary-parallel*, that is, they do not cut off a monogon (Figure 2a) or cut off a bigon together with part of ∂S (Figure 2b) respectively. We write $\iota(\alpha, \beta)$ for the (*geometric*) *intersection number* of α and β . Arcs α and β have disjoint interiors, which we refer to simply as being *disjoint*, if and only $\iota(\alpha, \beta) = 0$. A *multiarc* is a set of distinct and pairwise disjoint arcs. For example, see Figure 3.



FIGURE 2. Inessential arcs in S .




FIGURE 3. A multiarc on S .

The set of multiarcs is a poset with respect to inclusion. An (*ideal*) *triangulation* of S is a maximal multiarc. An Euler characteristic argument shows that every triangulation has $E(S)$ arcs and $F(S)$ faces. Each complementary region of a triangulation is a triangle with vertices on the marked points of S . Triangles have embedded interior, but their boundary may be non-embedded. In particular, triangles can be *folded* as shown in Figure 4.





FIGURE 4. A folded triangle.

2.1.3. *The flip graph.* The *flip graph* $\mathcal{F}(S)$ has a vertex for each triangulation. Triangulations $\mathcal{T}, \mathcal{T}' \in \mathcal{F}(S)$ are connected via an edge (of length one) if and only if they differ by a *flip*. This move consists of replacing the diagonal of a quadrilateral inside the triangulation with the other diagonal, as shown in Figure 5.

Any arc of a triangulation is either flippable or appears as the arc β in Figure 4. In the latter case, the arc β is flippable after first flipping the arc α . Thus for any arc α there is a triangulation $\mathcal{T} \ni \alpha$ in which α is flippable.

FIGURE 5. Flipping the arc α of a triangulation.

The flip graph appears implicitly in works of Harer [10], Mosher [16] and Penner [17]. It is connected [10] and the mapping class group acts on it *geometrically*, that is, properly, cocompactly and by isometries. Thus by the Švarc–Milnor Lemma [4, Proposition I.8.19] this graph is quasi-isometric to $\text{Mod}^\pm(S)$. The geometry of the flip graph was recently studied by the second author and Parlier in [6]. By [14, Theorem 1.2] and [2, Theorem 1.1], we have:

Theorem 2.1. *The natural homomorphism*

$$\text{Mod}^\pm(S) \rightarrow \text{Aut}(\mathcal{F}(S))$$

is an isomorphism. □

2.1.4. *The arc graph.* The *arc graph* $\mathcal{A}(S)$ has a vertex for each arc on S . Arcs $\alpha, \beta \in \mathcal{A}(S)$ are connected via an edge (of length one) if and only if they are disjoint. The arc graph appeared first in work of Harer [10]. The arc graph of S can be naturally extended to a flag simplicial complex called the *arc complex* of S . The one-skeleton of the dual of the arc complex is the flip graph.

Theorem 2.2. *The graphs $\mathcal{A}(S)$ and $\mathcal{A}(S')$ are isomorphic if and only if S and S' are homeomorphic.*

Proof. Every isomorphism between $\mathcal{A}(S)$ and $\mathcal{A}(S')$ induces an isomorphism between $\mathcal{F}(S)$ and $\mathcal{F}(S')$. If neither S nor S' is the complement of a (possibly empty) multiarc on the four-times marked sphere or the twice-marked torus then the result follows from [2, Theorem 1.4]. Otherwise:

- If $\partial S = \emptyset = \partial S'$ then the result follows from [14, Theorem 1.1].
- If $\partial S = \emptyset \neq \partial S'$ then $E(S) = 6$ and $E(S') \leq 5$. Since maximal complete subgraphs of $\mathcal{A}(S)$ have exactly $E(S)$ vertices, it follows that $\mathcal{A}(S)$ and $\mathcal{A}(S')$ cannot be isomorphic.
- If $\partial S \neq \emptyset \neq \partial S'$ then the result follows from [5, Theorem 1.1].

In any case, the result holds. □

The mapping class group acts on $\mathcal{A}(S)$ cocompactly and by isometries but not properly. Generally, the automorphism group of $\mathcal{A}(S)$ is isomorphic to $\text{Mod}^\pm(S)$ by [12, Theorem 1.2] and [5, Theorem 1.2].

2.2. The polygonalisation complex. A *polygonalisation* P (or *ideal cell decomposition*) is a multiarc such that each complementary region of P is a *polygon*, a topological disk with at least three sides and no marked points in its interior. These are dual to fat graphs (also known as ribbon graphs) [18, Section 4.2].

Remark 2.3. If P is a polygonalisation and Q is a multiarc such that $P \subseteq Q$ then Q is a polygonalisation. Additionally, suppose that P, Q and R are polygonalisations:

- If $P, Q \subseteq R$ then $P \cup Q$ is a polygonalisation.
- If $P, Q \supseteq R$ then $P \cap Q$ is a polygonalisation.

Definition 2.4. The *polygonalisation complex* $\mathcal{P}(S)$ is a cube complex in which vertices correspond to polygonalisations. Two polygonalisations $P, Q \in \mathcal{P}(S)$ are connected by an edge (of length one) if and only if they differ by a single arc, that is, their symmetric difference $P \Delta Q$ is a single arc. Inductively, a k -cube is added whenever its $(k - 1)$ -skeleton appears. See Appendix A for examples.

Since any polygonalisation can be extended to a triangulation and $\mathcal{F}(S)$ is connected, $\mathcal{P}(S)$ is also connected. Again, $\text{Mod}^\pm(S)$ acts geometrically on $\mathcal{P}(S)$ and so they are quasi-isometric.

We now establish some definitions and notation that will be used throughout.

Definition 2.5. For an arc $\alpha \in \mathcal{A}(S)$, its *stratum* $\mathcal{P}_\alpha(S)$ is the subcomplex of $\mathcal{P}(S)$ induced by the subset $\{P \in \mathcal{P}(S) : \alpha \in P\}$. Similarly, define $\overline{\mathcal{P}}_\alpha(S)$ to be the subcomplex induced by $\{P \in \mathcal{P}(S) : \alpha \notin P\}$.

Definition 2.6. An arc α is *removable from* P if $\alpha \in P$ and $P - \{\alpha\}$ is also a polygonalisation. An arc α *addable to* P if $\alpha \notin P$ and $P \cup \{\alpha\}$ is a polygonalisation.

Observe that $\alpha \in P$ is removable if and only if its interior meets two distinct polygons of P . Let $\partial\mathcal{P}_\alpha(S)$ (resp. $\partial\overline{\mathcal{P}}_\alpha(S)$) be the subcomplex of $\mathcal{P}_\alpha(S)$ (resp. $\overline{\mathcal{P}}_\alpha(S)$) induced by the polygonalisations in which α is removable (resp. addable).

Notation 2.7. For $P, Q \in \mathcal{P}(S)$, write $P \succ Q$ if $P \supseteq Q$ and $|P - Q| = 1$.

For an edge $e = \{P, Q\}$ in $\mathcal{P}(S)$ let $\text{arc}(e) := P \Delta Q \in \mathcal{A}(S)$ denote the arc that appears in P but not Q (or vice versa).

Remark 2.8. For any edge $e = \{P, Q\}$ in $\mathcal{P}(S)$, observe that $\text{arc}(e) = \alpha$ if and only if $P \in \partial\mathcal{P}_\alpha(S)$ and $Q \in \partial\overline{\mathcal{P}}_\alpha(S)$ (or vice versa).

3. THE CUBE COMPLEX STRUCTURE OF $\mathcal{P}(S)$

We describe some of the key properties of the cube complex structure of $\mathcal{P}(S)$. We begin by recalling some of the standard terms for cube complexes. For a complete reference see [21].

An n -cube is the Euclidean cube $[-\frac{1}{2}, \frac{1}{2}]^n$. For every $0 \leq k \leq n - 1$ a k -*face* is a subspace obtained restricting $n - k$ coordinates to $\pm\frac{1}{2}$. A k -face is also a k -cube. A *cube complex* X is a cell complex obtained by gluing cubes along their faces by isometries.

The *link* of a vertex v of X is the complex $\text{link}(v)$ induced on a small sphere about v by X . A *flag complex* is a simplicial complex where $n + 1$ vertices span an n -simplex if and only if they are pairwise adjacent. A cube complex X is *non-positively curved* if it satisfies *Gromov's link condition*: The link of each vertex is a

flag complex [9, Section 4.2.C]. If X is non-positively curved and simply connected then X is CAT(0).

A *midcube* is a subspace of a cube obtained by restricting exactly one coordinate to 0. A *hyperplane* is the union of midcubes that meet parallel edges. The *carrier* $N(H)$ of a hyperplane H is the union of all the cubes such that intersect H in a midcube. Sageev showed that CAT(0) cube complexes come equipped with a family of “nice” hyperplanes.


Theorem 3.1 ([20, Theorem 1.1]). *If X is a CAT(0) cube complex then:*

- (1) *each midcube lies in an embedded hyperplane;*
- (2) *every hyperplane H is two-sided, that is, $N(H) \cong [-\frac{1}{2}, \frac{1}{2}] \times H$;*
- (3) *every hyperplane H separates X , that is, $X - H$ consists of two connected components H^+ and H^- called halfspaces;*
- (4) *every hyperplane H is a CAT(0) cube complex; and*
- (5) *every hyperplane H and its carrier $N(H)$ are convex in X .*

Despite the fact that generically $\mathcal{P}(S)$ is not CAT(0), we show that Properties 1, 2 and 3 still hold for $\mathcal{P}(S)$ (Theorem 3.8). However, generically, Properties 4 and 5 do not hold for $\mathcal{P}(S)$. For example, see Table 2 and $S_{0,0}^{2,1}$ in Table 1 respectively.

3.1. The cubes of $\mathcal{P}(S)$. The following lemma will also be useful in Section 4.

Lemma 3.2 (Square lemma). *Suppose that $P_1, P_2, P_3, P_4 \subseteq \mathcal{P}(S)$ is an embedded 4-cycle. If $P_1 \succ P_2$ then $P_4 \succ P_3$. Moreover, $\text{arc}(\{P_1, P_2\}) = \text{arc}(\{P_3, P_4\})$.*



Proof. Suppose that instead $P_3 \succ P_4$. By considering the cardinalities of the P_i 's, we see that $P_1 \succ P_2 \prec P_3 \succ P_4 \prec P_1$. Since $P_2 \neq P_4$, by Remark 2.3 we have $P_1 \supseteq P_2 \cup P_4 \supsetneq P_2$. Since $|P_1| = |P_2| + 1$, we deduce $P_1 = P_2 \cup P_4$. Similarly, $P_3 = P_2 \cup P_4$ contradicting the fact that this cycle is embedded. Furthermore, if $\text{arc}(\{P_1, P_2\}) \neq \text{arc}(\{P_3, P_4\})$ then counting the number of components shows that $P_2 = P_4$. Again, this contradicts this cycle being embedded. \square

The n -cube graph C_n is the 1-skeleton of the standard n -cube. In particular, an embedded 4-cycle is a C_2 . By induction on n , the square lemma shows that any embedded cube graph C in $\mathcal{P}(S)$ contains a unique *source* C^+ and a unique *sink* C^- . That is, for any polygonalisation P in C we have that $C^- \subseteq P \subseteq C^+$.

For polygonalisations $P \subseteq Q$, let

$$[P, Q] := \{R \in \mathcal{P}(S) : P \subseteq R \subseteq Q\}.$$

In fact these are all the cubes that appear in $\mathcal{P}(S)$:

Lemma 3.3 (Characterisation of cubes). *If $P \subseteq Q$ then $[P, Q]$ forms the vertex set of an n -cube in $\mathcal{P}(S)$ where $n = |P| - |Q|$. Conversely, if C is an embedded n -cube in $\mathcal{P}(S)$ then the vertex set of C is $[C^-, C^+]$. \square*

When $\partial S = \emptyset$, this characterisation of cubes enables us to apply Penner's argument for the contractibility of the fatgraph complex [17, Theorem 2.5] to $\mathcal{P}(S)$. Hence, in this case the polygonalisation complex $\mathcal{P}(S)$ is contractible.

3.2. The hyperplanes of $\mathcal{P}(S)$. Let \sim denote the equivalence relation on the edges of $\mathcal{P}(S)$ generated by $e \sim e'$ if and only if e and e' are opposite edges of some square in $\mathcal{P}(S)$. We say that e and e' are *parallel* if $e \sim e'$.

Definition 3.4 ([20, Definition 1.4]). A *hyperplane* is the set of midcubes of $\mathcal{P}(S)$ that meet edges in $[e]$, for some equivalence class of edges. We write $\mathcal{H}(\mathcal{P}(S))$ for the set of all hyperplanes of $\mathcal{P}(S)$.

By the square lemma, if edges e, e' are parallel then $\text{arc}(e) = \text{arc}(e')$. Thus, the arc map descends to a well-defined map $\text{arc}: \mathcal{H}(\mathcal{P}(S)) \rightarrow \mathcal{A}(S)$. This map is surjective: for any $\alpha \in \mathcal{A}(S)$ there is a triangulation \mathcal{T} from which α is removable and so $\text{arc}(\{\mathcal{T}, \mathcal{T} - \{\alpha\}\}) = \alpha$. One of the aims of this section is to prove that this map is also injective. To achieve this we shall require some technical results.

Lemma 3.5 (Pentagon detour lemma). *Let \mathcal{T} and \mathcal{T}' be adjacent triangulations in $\mathcal{F}(S)$ and $Q := \mathcal{T} \cap \mathcal{T}'$. Suppose there is an arc $\alpha \in Q$ that is removable from \mathcal{T} and \mathcal{T}' but not from Q . Then there is a path $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3, \mathcal{T}_4 = \mathcal{T}'$ such that α is removable from each \mathcal{T}_i and from each $Q_i := \mathcal{T}_{i-1} \cap \mathcal{T}_i$.*

Proof. Let $\beta := \text{arc}(\{\mathcal{T}, Q\})$. Since α is not removable from Q but is removable from \mathcal{T} and \mathcal{T}' , it must appear as opposite sides of the unique square of Q . Since \mathcal{T} has at least three triangles, this square has another side γ that is removable from Q . The arcs β and γ are therefore a pair of chords of a pentagon in \mathcal{T} , as shown in Figure 6. The path $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3, \mathcal{T}_4 = \mathcal{T}'$ is the one obtained going around the natural five-cycle in $\mathcal{F}(S)$ defined by β and γ in the other direction. \square




FIGURE 6. A detour around a pentagon to stay in $\partial\mathcal{P}_\alpha(S)$.

Proposition 3.6. *For each arc $\alpha \in \mathcal{A}(S)$, the subcomplex $\partial\mathcal{P}_\alpha(S)$ is connected.*

Proof. Suppose $P, Q \in \partial\mathcal{P}_\alpha(S)$ are polygonalisations and let \mathcal{T} and \mathcal{T}' be triangulations containing P and Q respectively. Note that α is removable from \mathcal{T} and \mathcal{T}' . Furthermore, P and \mathcal{T} are in the same path-component of $\partial\mathcal{P}_\alpha(S)$ (and likewise for Q and \mathcal{T}'). Thus to show $\partial\mathcal{P}_\alpha(S)$ is connected it suffices to show that \mathcal{T} and \mathcal{T}' are in the same path-component.

There is a path $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n = \mathcal{T}'$ in $\mathcal{F}(S)$ such that $\alpha \in \mathcal{T}_i$ [6, Corollary 2.15]. Call \mathcal{T}_i *good* if α can be removed from it and *bad* otherwise. We describe how to modify this path to avoid any bad triangulations.

Suppose that \mathcal{T}_i is the first bad triangulation in this path. Note that $0 < i < n$ since both \mathcal{T}_0 and \mathcal{T}_n are good. Let β and γ be the arcs of \mathcal{T}_i that can be flipped to obtain \mathcal{T}_{i-1} and \mathcal{T}_{i+1} respectively. There are three possibilities to consider:

- If $\beta \cup \gamma$ is supported on two triangles then, since β cuts off a once-marked monogon, $\beta = \gamma$. Thus $\mathcal{T}_{i-1} = \mathcal{T}_{i+1}$ and so we simplify our path by removing \mathcal{T}_i and \mathcal{T}_{i+1} .
- If $\beta \cup \gamma$ is supported on three triangles then they fill a pentagon in \mathcal{T}_i . Thus we may replace \mathcal{T}_i with the good triangulations $\mathcal{T}', \mathcal{T}''$ as shown in Figure 7a.
- If $\beta \cup \gamma$ is supported on four triangles then the flips commute. Thus we may replace \mathcal{T}_i with the good triangulation \mathcal{T}' as shown in Figure 7b.

Performing any of these modifications will reduce the number of bad triangulations. Hence, by induction, we may assume α is removable from each \mathcal{T}_i along this path. Taking the barycentric subdivision gives a path $\mathcal{T}_0, Q_1, \mathcal{T}_1, Q_2, \dots, \mathcal{T}_{n-1}, Q_n, \mathcal{T}_n$ in $\mathcal{P}_\alpha(S)$, where $Q_i := \mathcal{T}_{i-1} \cap \mathcal{T}_i$. Replacing subpaths with pentagon detours if necessary, we may assume that α is also removable from each Q_i . This yields a path in $\partial\mathcal{P}_\alpha(S)$ connecting \mathcal{T} to \mathcal{T}' as required. \square

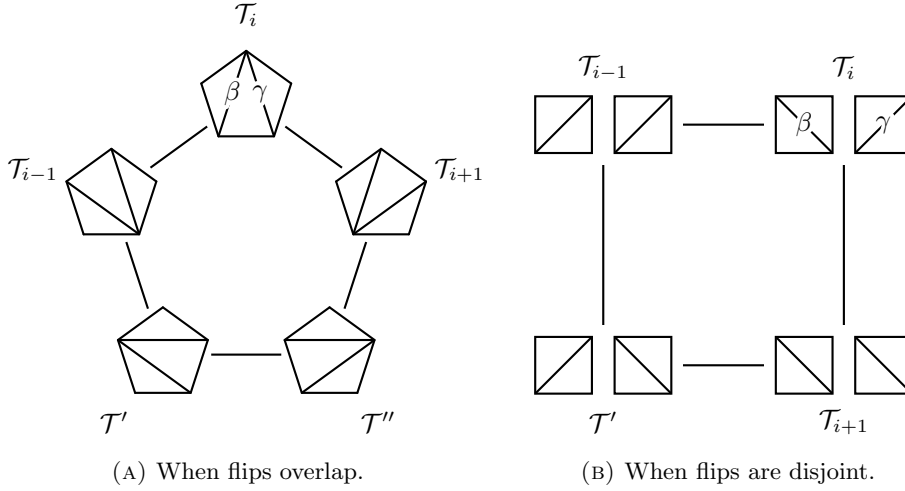


FIGURE 7. Diverting a path around a bad triangulation.

Observe that $P, Q \in \partial\overline{\mathcal{P}_\alpha(S)}$ are adjacent if and only if they form a square with $P \cup \{\alpha\}, Q \cup \{\alpha\} \in \partial\mathcal{P}_\alpha(S)$. Hence, by the above proposition, $\partial\overline{\mathcal{P}_\alpha(S)}$ is also connected.

Lemma 3.7. *Let $\alpha \in \mathcal{A}(S)$ be an arc. Then the subcomplexes $\mathcal{P}_\alpha(S)$ and $\overline{\mathcal{P}_\alpha(S)}$ in $\mathcal{P}(S)$ are both connected.*

Proof. Suppose that $P \in \mathcal{P}_\alpha(S)$ and $Q \in \overline{\mathcal{P}_\alpha(S)}$. For any edge-path in $\mathcal{P}(S)$ from P to Q , the arc α must be removed from a polygonalisation at some point along this path. Hence such a path in $\mathcal{P}(S)$ must cross H_α . Therefore there is a path within $\mathcal{P}_\alpha(S)$ from P to $\partial\mathcal{P}_\alpha(S)$ and a path within $\overline{\mathcal{P}_\alpha(S)}$ from Q to $\partial\overline{\mathcal{P}_\alpha(S)}$. Since $\partial\mathcal{P}_\alpha(S)$ and $\partial\overline{\mathcal{P}_\alpha(S)}$ are both connected, $\mathcal{P}_\alpha(S)$ and $\overline{\mathcal{P}_\alpha(S)}$ are too. \square

Theorem 3.8. *The map $\text{arc}: \mathcal{H}(\mathcal{P}(S)) \rightarrow \mathcal{A}(S)$ is a bijection. Moreover, for each arc $\alpha \in \mathcal{A}(S)$, the corresponding hyperplane H_α is embedded, two-sided and separates $\mathcal{P}(S)$ into two connected components: $\mathcal{P}_\alpha(S)$ and $\overline{\mathcal{P}_\alpha(S)}$.*

Proof. To show that the map $\text{arc}: \mathcal{H}(\mathcal{P}(S)) \rightarrow \mathcal{A}(S)$ is a bijection we must prove that edges e, e' in $\mathcal{P}(S)$ are parallel if and only if $\text{arc}(e) = \text{arc}(e')$. As described above, the forwards direction holds trivially by the square lemma.

For the backwards direction, suppose that e, e' are edges of $\mathcal{P}(S)$ such that $\text{arc}(e) = \text{arc}(e') = \alpha$. By Remark 2.8, we deduce that $e = \{P, Q\}$ and $e' = \{P', Q'\}$ for some $P, P' \in \partial\mathcal{P}_\alpha(S)$ and $Q, Q' \in \partial\overline{\mathcal{P}}_\alpha(S)$. By Proposition 3.6 there is a path p from P to P' in $\partial\mathcal{P}_\alpha(S)$. Now α can be removed from each polygonalisation in p to obtain a parallel path q from Q to Q' in $\partial\overline{\mathcal{P}}_\alpha(S)$. This gives a sequence of squares from which we can deduce that e and e' are parallel.

It is easy to verify that hyperplanes are embedded. If a hyperplane self-intersects then a self-intersection must occur in some square. Then all four edges in this square must correspond to adding/removing the same arc, which is impossible.

As in the proof of Lemma 3.7, any path from a vertex in $\mathcal{P}_\alpha(S)$ to a vertex in $\overline{\mathcal{P}}_\alpha(S)$ must cross H_α , hence H_α separates $\mathcal{P}_\alpha(S)$ and $\overline{\mathcal{P}}_\alpha(S)$. By Lemma 3.7, these are both connected and so H_α separates $\mathcal{P}(S)$ into $\mathcal{P}_\alpha(S)$ and $\overline{\mathcal{P}}_\alpha(S)$.

Finally, each hyperplane must be two-sided since it is separating. \square

Corollary 3.9. *Let $\alpha \in \mathcal{A}(S)$ be an arc. Then the hyperplane H_α separates polygonalisations $P, Q \in \mathcal{P}(S)$ if and only if $\alpha \in P\Delta Q$. In particular, there are at most $2E(S)$ hyperplanes separating any given pair of polygonalisations.* \square

Remark 3.10. The hyperplanes of $\mathcal{P}(S)$ cannot self-osculte, but they can inter-osculte. For example, see $S_{0,0}^{2,1}$ in Table 1.

3.3. The crossing graph of $\mathcal{P}(S)$.

Definition 3.11. Two hyperplanes H_α and H_β *cross*, denoted $H_\alpha \pitchfork H_\beta$, if

$$\mathcal{P}_\alpha(S) \cap \mathcal{P}_\beta(S), \mathcal{P}_\alpha(S) \cap \overline{\mathcal{P}}_\beta(S), \overline{\mathcal{P}}_\alpha(S) \cap \mathcal{P}_\beta(S) \text{ and } \overline{\mathcal{P}}_\alpha(S) \cap \overline{\mathcal{P}}_\beta(S)$$

are all non-empty.

There are several ways of characterising when hyperplanes cross.

Proposition 3.12. *The following are equivalent:*

- (1) *The hyperplanes H_α and H_β cross.*
- (2) *The arcs α and β are distinct and disjoint but do not form a folded triangle.*
- (3) *There is a triangulation \mathcal{T} containing $\alpha \neq \beta$ such that $\mathcal{T} - \{\alpha, \beta\}$ is a polygonalisation.*

Proof. We follow a cycle of implications:

1 \implies 2: There is a polygonalisation P containing α and β and so these arcs must be disjoint. If α and β form a folded triangle then every polygonalisation that contains α must contain β or vice versa. Therefore either $\mathcal{P}_\alpha(S) \cap \overline{\mathcal{P}}_\beta(S)$ or $\overline{\mathcal{P}}_\alpha(S) \cap \mathcal{P}_\beta(S)$ is empty and so H_α and H_β do not cross.

2 \implies 3: Since α and β do not form a folded triangle, there is a triangulation \mathcal{T} containing α and β in which both are removable. Let $Q := \mathcal{T} - \{\alpha\}$. Suppose that β is not removable from Q . Then, as in the proof of Lemma 3.5, a polygon of Q is a square, two sides of which are β . Since \mathcal{T} has at least three triangles, this square has another side γ that is removable from Q . We flip γ in \mathcal{T} to obtain a new triangulation \mathcal{T}' in which both α and β are simultaneously removable.

3 \implies 1: The polygonalisations $\mathcal{T}, \mathcal{T} - \{\beta\}, \mathcal{T} - \{\alpha\}$ and $\mathcal{T} - \{\alpha, \beta\}$ show that the four sets required by Definition 3.11 are all non-empty. \square

Definition 3.13. The *crossing graph* $\text{Cr}(\mathcal{P}(S))$ is the graph with a vertex for each hyperplane in $\mathcal{P}(S)$. Two hyperplanes $H, H' \in \text{Cr}(\mathcal{P}(S))$ are connected via an edge (of length one) if and only if $H \pitchfork H'$.

By Proposition 3.12, we immediately deduce that the crossing graph $\text{Cr}(\mathcal{P}(S))$ embeds into the arc graph $\mathcal{A}(S)$. To get control over this embedding we will need to consider the paths in $\mathcal{A}(S)$ *without folds*, that is, the ones in which no consecutive pair of arcs form a folded triangle.

Lemma 3.14. *If $d(\alpha, \beta) = n \geq 3$ then there is a geodesic*

$$\alpha = \alpha_0, \alpha_1, \dots, \alpha_n = \beta$$

in $\mathcal{A}(S)$ that is without folds.

Proof. Let

$$\alpha = \alpha_0, \alpha_1, \dots, \alpha_n = \beta$$

be a geodesic in $\mathcal{A}(S)$ from α to β .

Suppose that $\{\alpha_i, \alpha_{i+1}\}$ form a folded triangle for some $0 < i < n - 1$. Then either $\iota(\alpha_{i-1}, \alpha_{i+1}) = 0$ or $\iota(\alpha_i, \alpha_{i+2}) = 0$, which contradicts this path being a geodesic. Hence we need only consider folded triangles formed by $\{\alpha_0, \alpha_1\}$ and $\{\alpha_{n-1}, \alpha_n\}$. We show that if the former occurs then there are arcs α'_1, α'_2 such that

$$\alpha_0, \alpha'_1, \alpha'_2, \alpha_3$$

is a geodesic without folds. We will take care to ensure that α'_1 and α'_2 are both arcs, that is, that they are not null-homotopic nor boundary-parallel. Similarly we can replace the end of the geodesic if $\{\alpha_{n-1}, \alpha_n\}$ forms a folded triangle.

Now note that if $\{\alpha_0, \alpha_1\}$ forms a folded triangle then α_0 must cut off a once-marked monogon. Furthermore, α_2 must have at least one endpoint on x , the inner marked point of this monogon, since it must intersect α_0 .

If α_2 has exactly one endpoint on x then let α'_1 and α''_1 be as shown in Figure 8a. If α'_1 or α''_1 is null-homotopic then $\alpha_1 = \alpha_2$, which contradicts this path being a geodesic. If α'_1 and α''_1 are both boundary-parallel then $F(S) = 2$, which is again a contradiction. Hence without loss of generality α'_1 is an arc. Since this arc is disjoint from α_0 and α_2 and does not form a folded triangle with either, $\alpha_0, \alpha'_1, \alpha_2, \alpha_3$ forms the required geodesic without folds.

On the other hand, if α_2 has two endpoints on x , we construct α'_1 by surgery as in Figure 8b. We note that α'_1 cannot be null-homotopic. If it were then α_2 cuts off a once-marked monogon and so α_3 must also be disjoint from α_1 as it is disjoint from α_2 . Again, this contradicts this path being a geodesic. Hence we need only consider the case in which α'_1 is boundary-parallel, since otherwise α'_1 is an arc and $\alpha_0, \alpha'_1, \alpha_2, \alpha_3$ is a geodesic without folds.

If α'_1 is boundary-parallel then α_2 must cut off an annulus with one marked point on each boundary component. Since it must meet α_1 , the arc α_3 must be contained in the annulus cut off by α_2 . However, as α_2 is essential, there is an arc α'_2 in the other connected component with one endpoint on x . By construction α'_2 does not cut off an annulus and it is disjoint from $\alpha_1, \alpha_2, \alpha_3$. Hence $\alpha_0, \alpha_1, \alpha'_2, \alpha_3$ is a geodesic and α'_2 has a single endpoint on x . Thus we can repeat the above argument to find a new arc α'_1 such that

$$\alpha_0, \alpha'_1, \alpha'_2, \alpha_3$$

is a geodesic without folds. □

By Proposition 3.12, we have $d(H_\alpha, H_\beta) \geq d(\alpha, \beta)$. On the other hand, Lemma 3.14 shows that if $d(\alpha, \beta) \geq 3$ then $d(H_\alpha, H_\beta) \leq d(\alpha, \beta)$. Thus, on the large scale, the natural map $j : \text{Cr}(\mathcal{P}(S)) \hookrightarrow \mathcal{A}(S)$ preserves distances. Furthermore, from Proposition 3.12, if $d(\alpha, \beta) = 1$ then $d(H_\alpha, H_\beta) \leq 2$. The final possibility is the following:

Lemma 3.15. *If $d(\alpha, \beta) = 2$ then $d(H_\alpha, H_\beta) \leq 4$.*




FIGURE 8. A detour to avoid a once-marked monogon.

Proof. Suppose that α, γ, β is a geodesic. There are now three cases to consider. First, if this geodesic is without folds then $H_\alpha, H_\gamma, H_\beta$ is a geodesic in $\text{Cr}(\mathcal{P}(S))$ and so $d(H_\alpha, H_\beta) = 2$. Second, if $\{\alpha, \gamma\}$ and $\{\gamma, \beta\}$ are both folded triangles then necessarily α and β both cut off once-marked monogons. Since $F(S) \geq 3$, there are disjoint arcs γ' and γ'' such that

$$\iota(\gamma', \alpha) = \iota(\gamma', \gamma) = 0 = \iota(\gamma'', \gamma) = \iota(\gamma'', \beta)$$

and the path $\alpha, \gamma', \gamma, \gamma'', \beta$ in $\mathcal{A}(S)$ is without folds. Hence by Proposition 3.12 it pulls back to a path $H_\alpha, H_{\gamma'}, H_\gamma, H_{\gamma''}, H_\beta$ in $\text{Cr}(P(S))$ and so $d(H_\alpha, H_\beta) \leq 4$. Third, if the geodesic contains a unique folded triangle then without loss of generality it is formed by $\{\alpha, \gamma\}$. In this case α must cut off a once-marked monogon (otherwise α and β would be disjoint). Hence again there is an arc γ' , disjoint from both α and γ , such that $\alpha, \gamma', \gamma, \beta$ is a path without folds in $\mathcal{A}(S)$. Again this pulls back to a path $H_\alpha, H_{\gamma'}, H_\gamma, H_\beta$ in $\text{Cr}(P(S))$ and so $d(H_\alpha, H_\beta) \leq 3$. \square

Corollary 3.16. *The map $j: \text{Cr}(\mathcal{P}(S)) \rightarrow \mathcal{A}(S)$ is a $(1, 2)$ -quasi-isometry. \square*

Hence $\text{Cr}(\mathcal{P}(S))$ has the same large scale geometry as $\mathcal{A}(S)$. For example, when $\partial S = \emptyset$, Masur–Schleimer showed that $\mathcal{A}(S)$ is hyperbolic [15, Theorem 20.2]. Hensel–Przytycki–Webb later showed that in this case in fact every geodesic triangle in $\mathcal{A}(S)$ has a 7-centre [11, Theorem 1.2] and so the same is true for $\text{Cr}(\mathcal{P}(S))$.

For ease of notation, for a hyperplane $H \in \mathcal{H}(\mathcal{P}(S))$ let $\text{link}(H)$ denote the subgraph of $\text{Cr}(\mathcal{P}(S))$ induced by

$$\{H' \in \mathcal{H}(\mathcal{P}(S)) : H \pitchfork H'\}$$

We will now use the crossing graph to prove that, generically, different surfaces have different polygonalisation complexes.

Lemma 3.17. *Arcs α and β form a folded triangle, as shown in Figure 4, if and only if $\text{link}(H_\alpha) \subsetneq \text{link}(H_\beta)$.*


Proof. We use Proposition 3.12 repeatedly to determine whether one hyperplane lies in the link of another.

Suppose α and β form a folded triangle. If $H_\gamma \in \text{link}(H_\alpha)$ then α and γ are distinct, disjoint and do not form a folded triangle. Thus γ and β must also be disjoint and distinct. Additionally, since α only forms a folded triangle with β we have that γ and β do not form a folded triangle and so $H_\gamma \in \text{link}(H_\beta)$. On the other hand, let $\mathcal{T} \ni \beta$ be a triangulation without folded triangles. Then there is an arc $\gamma \in \mathcal{T}$ such that $\iota(\gamma, \beta) = 0$ but $\iota(\gamma, \alpha) \neq 0$. Thus $H_\gamma \in \text{link}(H_\beta)$ but $H_\gamma \notin \text{link}(H_\alpha)$ and so $\text{link}(H_\alpha) \subsetneq \text{link}(H_\beta)$.


Suppose α does not form a folded triangle with β . Let $\mathcal{T} \ni \alpha$ be a triangulation with one or zero folded triangles depending on whether α does or does not cut off a once-marked monogon respectively. There are now three cases to consider:

- (1) If $\iota(\beta, \mathcal{T}) = 0$ then $\beta \in \mathcal{T}$. Therefore, since β does not form a folded triangle with α , we have that $H_\beta \in \text{link}(H_\alpha)$ but $H_\beta \notin \text{link}(H_\beta)$.
- (2) If $\iota(\beta, \mathcal{T}) \neq 0$ and there is an arc $\gamma \in \mathcal{T}$ such that γ and α are distinct and do not form a folded triangle and $\iota(\gamma, \beta) \neq 0$ then $H_\gamma \in \text{link}(H_\alpha)$ but $H_\gamma \notin \text{link}(H_\beta)$.
- (3) If $\iota(\beta, \mathcal{T}) \neq 0$ and any arc of \mathcal{T} that intersects β is either α or forms a folded triangle with α then α and β are as shown in Figure 9. Hence there is an arc γ which appears as exactly one side of the unique square of $\mathcal{T} - \{\alpha\}$. Flipping γ in \mathcal{T} gives an arc $\gamma' \neq \alpha$. Since γ' does not cut off a once-marked monogon, is disjoint from α and $\iota(\gamma', \beta) \neq 0$ we have that $H_{\gamma'} \in \text{link}(H_\alpha)$ but $H_{\gamma'} \notin \text{link}(H_\beta)$.

In any case, the inclusion $\text{link}(H_\alpha) \subsetneq \text{link}(H_\beta)$ does not hold. \square



(A) When β meets a folded triangle.



(B) When β only meets α .

FIGURE 9. When β only meets α and arcs that form a folded triangle with α .

Theorem 3.18. *The cube complexes $\mathcal{P}(S)$ and $\mathcal{P}(S')$ are isomorphic if and only if S and S' are homeomorphic.*

Proof. If $\mathcal{P}(S)$ and $\mathcal{P}(S')$ are isomorphic as cube complexes then $\text{Cr}(\mathcal{P}(S))$ and $\text{Cr}(\mathcal{P}(S'))$ are isomorphic graphs. By Proposition 3.12 and Lemma 3.17, adding the edges

$$\{\{H, H'\} : H, H' \in \text{Cr}(\mathcal{P}(S)), \text{link}(H) \subsetneq \text{link}(H')\}$$

to $\text{Cr}(\mathcal{P}(S))$ produces a graph isomorphic to $\mathcal{A}(S)$. Similarly, by adding the edges

$$\{\{H, H'\} : H, H' \in \text{Cr}(\mathcal{P}(S')), \text{link}(H) \subsetneq \text{link}(H')\}$$

to $\text{Cr}(\mathcal{P}(S'))$ we obtain a graph isomorphic to $\mathcal{A}(S')$. This rule for adding edges is purely combinatorial, that is, it depends only on the graph structure. Since $\text{Cr}(\mathcal{P}(S))$ and $\text{Cr}(\mathcal{P}(S'))$ are isomorphic, we deduce $\mathcal{A}(S) \cong \mathcal{A}(S')$. By rigidity of the arc graph (Theorem 2.2), this in turn implies that S is homeomorphic to S' .

The reverse direction is straightforward. \square


In fact by examining the possible cases shown in Appendix C, we see that this theorem also holds when $F(S) = 2$.

4. FAILURE OF GROMOV'S LINK CONDITION

In general, $\mathcal{P}(S)$ contains many vertices that fail Gromov's link condition [9, Section 4.2.C], which we rephrase as follows. When $\partial S = \emptyset$ this is the only reason why $\mathcal{P}(S)$ fails to be CAT(0), since in this case $\mathcal{P}(S)$ is contractible.

Definition 4.1. A *positive curvature system* (based at $P \in \mathcal{P}(S)$) is a set of $k \geq 3$ edges e_1, \dots, e_k in $\mathcal{P}(S)$ incident to P such that:

- every subset of $\{e_1, \dots, e_k\}$ of size $k-1$ is contained in an embedded $(k-1)$ -cube, and
- the set $\{e_1, \dots, e_k\}$ is not contained in an embedded k -cube.

FIGURE 10. A positive curvature system based at P .

See Figure 10.

Positive curvature systems are useful as they allow us to determine the direction of inclusion of the edges involved.

Proposition 4.2 (Positive curvature criterion). *If e_1, \dots, e_k is a positive curvature system where $e_i = \{P, Q_i\}$ then $P \succ Q_i$.*

Proof. Let $\alpha_i = \text{arc}(e_i)$. For a contradiction, suppose that $P \prec Q_1$. Then α_1 is disjoint from all arcs of P . By assumption, $\{e_2, \dots, e_k\}$ lie in an embedded $(k-1)$ -cube C that contains P . Recall that C has a unique source C^+ and a unique sink C^- with respect to inclusion and that $C^+ - C^- = \{\alpha_2, \dots, \alpha_k\}$. Observe that for all $i = 2, \dots, k$, the edges e_1 and e_i are contained in a common $(k-1)$ -cube since $\{e_1, \dots, e_k\}$ is a positive curvature system. This means that α_1 is disjoint from α_i for all i . Note that $C^- \subseteq P \subseteq C^+$ and $C^+ - P \subseteq \{\alpha_2, \dots, \alpha_k\}$. We deduce that α_1 is disjoint from C^+ , and so $C^+ \cup \{\alpha_1\}$ is also a polygonalisation. The set $[C^-, C^+ \cup \{\alpha_1\}]$ contains $P \cup \{\alpha_1\}$, and is therefore the vertex set of a k -cube that contains the edges $\{e_1, \dots, e_k\}$. This contradicts the assumption that e_1, \dots, e_k is a positive curvature system. \square

Proposition 4.3. *Let e_1, \dots, e_k be edges incident to P with $k \geq 3$. Let $\alpha_i = \text{arc}(e_i)$. Then $\{e_1, \dots, e_k\}$ is a positive curvature system based at $P \in \mathcal{P}(S)$ if and only if each $\alpha_i \in P$ and there is a (possibly peripheral) simple closed curve c on S such that:*

- $\iota(c, \alpha_i) = 1$ for every i ;
- $\iota(c, \beta) = 0$ for every $\beta \in P - \{\alpha_1, \dots, \alpha_k\}$;
- c meets each polygon of P at most once.

Proof. By assumption $Q := P - \{\alpha_2, \dots, \alpha_k\}$ is a polygonalisation, but $P - \{\alpha_1, \dots, \alpha_k\}$ is not. Since α_1 is not removable from Q , its interior meets a single polygon \mathcal{R} of Q . Let c be a simple arc in \mathcal{R} connecting one side of α_1 to the other. By construction c is a simple closed curve in S disjoint from each arc in $P - \{\alpha_1, \dots, \alpha_k\}$. Furthermore c intersects each α_i at most once. We claim that c must intersect every α_i exactly once. For a contradiction, without loss of generality, suppose α_k is disjoint from c . Then by assumption $P - \{\alpha_1, \dots, \alpha_{k-1}\}$ is a polygonalisation. So c lies in some polygon and so is null-homotopic, a contradiction.

Conversely, suppose c is a curve satisfying the given conditions. Orient c and assume, without loss of generality, that the arcs it meets appear in the order $\alpha_1, \dots, \alpha_k$ (up to cyclic permutation). Let \mathcal{R}_i be the polygon of P which contains the segment of c appearing between α_{i-1} and α_i . The \mathcal{R}_i 's are well defined since c intersects each α_i exactly once, and are distinct since c meets each polygon of P at most once. Note that α_2 is removable from P since its interior meets distinct regions \mathcal{R}_1 and \mathcal{R}_2 . Proceeding inductively, we can remove α_i from $P - \{\alpha_2, \dots, \alpha_{i-1}\}$ for $2 < i \leq k$.

Each removal decreases the number of polygons meeting c by one. Therefore, the interior of α_1 meets only one polygon of $P - \{\alpha_2, \dots, \alpha_k\}$, and so $P - \{\alpha_1, \dots, \alpha_k\}$ is not a polygonalisation. Hence $\{e_1, \dots, e_k\}$ is a positive curvature system. \square

Propositions 4.2 and 4.3 show that $P \in \mathcal{P}(S)$ fails Gromov's link condition if and only if its dual fat graph contains an embedded cycle of length at least three. Hence, positive curvature systems are abundant throughout $\mathcal{P}(S)$. In fact, together with the square lemma and one additional piece of information, they can be used to determine the direction of inclusion of all edges.

Proposition 4.4. *There is a combinatorial criterion such that for each edge $e = \{P, Q\}$ in $\mathcal{P}(S)$ we have that $P \succ Q$ if and only if the criterion is satisfied by the $(E(S) + 2)$ -neighbourhood of e .*

Proof. Suppose that $P \succ Q$ and let $\alpha := \text{arc}(e)$. We may assume that e is not part of a positive curvature system as otherwise we recover that $P \succ Q$ automatically by Proposition 4.2.

Since α is removable from P it meets two distinct polygons \mathcal{R}_1 and \mathcal{R}_2 of P . We consider the different possibilities for how these polygons meet:

- Suppose that \mathcal{R}_1 and \mathcal{R}_2 share an additional arc $\beta \in P$. Let c be a simple closed curve contained in $\mathcal{R}_1 \cup \mathcal{R}_2$ that meets only α and β .
 - Suppose that \mathcal{R}_1 (respectively \mathcal{R}_2) has at least four sides. Let γ be a diagonal of \mathcal{R}_1 (resp. \mathcal{R}_2) that meets c . Then the edge of $\mathcal{P}(S)$ from $P \cup \{\gamma\}$ to $Q \cup \{\gamma\}$ is part of a positive curvature system. Therefore using Proposition 4.2 we have that $P \cup \{\gamma\} \succ Q \cup \{\gamma\}$ and so $P \succ Q$ by Lemma 3.2.
 - Otherwise, both \mathcal{R}_1 and \mathcal{R}_2 are triangles. Therefore there is a side γ ($\neq \alpha, \beta$) of \mathcal{R}_1 or \mathcal{R}_2 that is removable and doing so creates a polygon with at least four sides. By the preceding argument, $P - \{\gamma\} \succ Q - \{\gamma\}$ can be deduced from the positive curvature criterion and the square lemma. Hence, $P \succ Q$ can too.
- Otherwise, \mathcal{R}_1 and \mathcal{R}_2 meet only along α . In this case, any arc that can be added (respectively removed) from P can also be added (resp. removed) from Q . However there are also arcs that are disjoint from Q but intersect α . Hence $\deg(P) < \deg(Q)$.

Thus we can deduce that $P \succ Q$ either from: the positive curvature criterion applied to a parallel cube at most distance two away, or (if the positive curvature criterion is not definitive) the degrees of P and Q in $\mathcal{P}(S)$.

Furthermore, any cube in $\mathcal{P}(S)$ has diameter at most $E(S)$. Hence $P \succ Q$ can be determined from the combinatorics of the $(E(S) + 2)$ -neighbourhood of e . \square

Corollary 4.5. *For each k , there is a combinatorial criterion that characterises the vertices of $\mathcal{P}(S)$ corresponding to the polygonalisations with k arcs.*

Proof. This follows from the fact that a vertex P corresponds to a polygonalisation with k arcs if and only if

$$\exists P_1 \prec \dots \prec P_{E(S)-k} (P \prec P_1 \wedge \nexists Q (P_{E(S)-k} \prec Q)).$$

Since \prec can be expressed as a combinatorial criterion by Proposition 4.4, this statement is a combinatorial criterion also. \square

We will now use these criteria to prove the rigidity of $\mathcal{P}(S)$.

Lemma 4.6. *Every automorphism of $\mathcal{P}(S)$ induces an automorphism of $\mathcal{F}(S)$ by restriction of the vertices that correspond to triangulations.*

Proof. Suppose that $\phi \in \text{Aut}(\mathcal{P}(S))$ is an automorphism. By Corollary 4.5, the map ϕ restricts to an automorphism of the graph

$$\mathcal{F}' := \{P \in \mathcal{P}(S) : |P| \geq E(S) - 1\}$$

induced by the set of triangulations and triangulations with one arc missing (recall that $|P|$ equals the number of arcs in P). Moreover ϕ maps triangulations to triangulations. Since \mathcal{F}' is the barycentric subdivision of $\mathcal{F}(S)$, it follows that ϕ induces an automorphism $\phi' \in \text{Aut}(\mathcal{F}(S))$. \square

Theorem 4.7. *The natural homomorphism*

$$\rho : \text{Mod}^\pm(S) \rightarrow \text{Aut}(\mathcal{P}(S))$$

is an isomorphism.


Proof. If $f \in \text{Ker}(\rho)$ then it must induce the identity on $\mathcal{F}(S)$. By Theorem 2.1 the map f must be the trivial mapping class and so ρ is injective.

We now prove that ρ is surjective. By Lemma 4.6 any $\phi \in \text{Aut}(\mathcal{P}(S))$ induces some $\phi' \in \text{Aut}(\mathcal{F}(S))$. By Theorem 2.1, ϕ' is induced by some $f \in \text{Mod}^\pm(S)$. Let f_* be the automorphism of $\mathcal{P}(S)$ induced by f . We will now prove that $f_*(P) = \phi(P)$ for every $P \in \mathcal{P}(S)$. Denote by $\mathcal{U}(P)$ the set of all triangulations that contain P . By Lemma 4.6 the map ϕ preserves the set of all triangulations. By Proposition 4.4 the map ϕ preserves the direction of the inclusions. Hence, we have $\mathcal{U}(\phi(P)) = \phi(\mathcal{U}(P))$. Similarly, $\mathcal{U}(f_*(P)) = f_*(\mathcal{U}(P))$. Since f_* and ϕ coincide on every triangulation, we have $\mathcal{U}(f_*(P)) = \mathcal{U}(\phi(P))$. Note that $\mathcal{U}(Q) \supseteq \mathcal{U}(Q')$ if and only if $Q \subseteq Q'$. It follows that $f_*(P) = \phi(P)$, hence $\rho(f) = \phi$. \square

Acknowledgements. The authors are grateful for the hospitality of the University of Illinois, Indiana University, University of Oklahoma, ICERM and MSRI. The authors acknowledge support from U.S. National Science Foundation grants DMS 1107452, 1107263, 1107367 ‘‘RNMS: GEometric structures And Representation varieties’’ (the GEAR Network). The second author acknowledges support from IU Provost’s Travel Award for Women in Science. This material is based upon work supported by the National Science Foundation under Grant No. DMS 1440140 while the second and third authors were in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester.

APPENDIX A. EXAMPLES OF POLYGONALISATION COMPLEXES

Here we use the notation $S_{g,s}^{p_1, \dots, p_b}$ to denote the surface of genus g with s marked points in its interior and b boundary components with p_1, \dots, p_b marked points respectively. In the following examples we colour a polygonalisation P red, blue or green depending on whether $E(S) - |P|$ is zero, one or two respectively. Hence the subgraph induced by the red and blue vertices is homeomorphic to the flip graph of the surface. In each example we highlight one top-dimensional cube in each $\text{Mod}^\pm(S)$ -orbit. In each of these cases $F(S) = 3$ and the only omitted case with $F(S) = 3$ is $S_{1,0}^1$. Additionally, the polygonalisation complex of a hexagon, where $F(S_{0,0}^6) = 4$, is shown in Figure 1. Many of these complexes were generated using **fatter** [3].






TABLE 1. Examples of polygonalisation complexes.

APPENDIX B. EXAMPLES OF HYPERPLANES

Some examples of hyperplanes in polygonalisation complexes. There are examples of CAT(0)/non CAT(0) hyperplanes and examples in which the hyperplanes are not convex subsets.






TABLE 2. Examples of hyperplanes in polygonalisation complexes.

APPENDIX C. EXCEPTIONAL SURFACES

In Table 3 we list the exceptional surfaces, that is, the surfaces with $F(S) < 3$. Again we use the notation of Appendix A to describe these surfaces. This table shows the standing of Theorem 3.18 and Theorem 4.7 in these cases. Examination of $\mathcal{P}(S)$ in these cases also shows that variants of some of our results, such as Theorem 3.8¹, also hold when $F(S) < 3$.

S	$\mathcal{P}(S)$	$\text{Aut}(\mathcal{P}(S))$	$\text{Mod}^\pm(S) \rightarrow \text{Aut}(\mathcal{P}(S))$
$S_{1,1}$		Uncountable [14]	Homomorphism
$S_{0,3}$		$\text{Sym}(3)$	Epimorphism
$S_{0,1}^2$		\mathbb{Z}_2	Epimorphism
$S_{0,0}^4$		\mathbb{Z}_2	Epimorphism
$S_{0,0}^{1,1}$		$\mathbb{Z}_2 * \mathbb{Z}_2$	Homomorphism
$S_{0,1}^1$		$\mathbb{1}$	Epimorphism
$S_{0,0}^3$		$\mathbb{1}$	Epimorphism
$S_{0,2}$	\emptyset	$\mathbb{1}$	Epimorphism
$S_{0,0}^2$	\emptyset	$\mathbb{1}$	Epimorphism
$S_{0,0}^1$	\emptyset	$\mathbb{1}$	Epimorphism
$S_{0,1}$	\emptyset	$\mathbb{1}$	Epimorphism

TABLE 3. The exceptional surfaces.

¹For $S_{1,1}$ the hyperplanes are separating but the two complementary components are not $\mathcal{P}_\alpha(S)$ and $\overline{\mathcal{P}_\alpha(S)}$.

REFERENCES

1. Javier Aramayona, *Combinatorial superrigidity for graphs associated to surfaces*. [2]
2. Javier Aramayona, Thomas Koberda, and Hugo Parlier, *Injective maps between flip graphs*, *Ann. Inst. Fourier (Grenoble)* **65** (2015), no. 5, 2037–2055. MR 3449205 [4]
3. Mark Bell, *fatter (computer software)*, `pypi.python.org/pypi/fatter`, 2016, Version 0.1.3. [15]
4. Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR 1744486 (2000k:53038) [4]
5. V. Disarlo, *Combinatorial rigidity of arc complexes*, ArXiv e-prints (2015). [2, 4, 5]
6. V. Disarlo and H. Parlier, *The geometry of flip graphs and mapping class groups*, ArXiv e-prints (2014). [4, 7]
7. Vladimir V. Fock and Alexander B. Goncharov, *Dual Teichmüller and lamination spaces*, **11** (2007), 647–684. MR 2349682 [3]
8. Louis Funar, *Ptolemy groupoids actions on Teichmüller spaces*, Modern trends in geometry and topology, Cluj Univ. Press, Cluj-Napoca, 2006, pp. 185–201. MR 2250211 [2]
9. M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263. MR 919829 [6, 12]
10. John L. Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface*, *Invent. Math.* **84** (1986), no. 1, 157–176. MR 830043 [4]
11. Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, *J. Eur. Math. Soc. (JEMS)* **17** (2015), no. 4, 755–762. MR 3336835 [11]
12. Elmas Irmak and John D. McCarthy, *Injective simplicial maps of the arc complex*, *Turkish J. Math.* **34** (2010), no. 3, 339–354. MR 2681579 [2, 5]
13. Michael Kapovich and Bernhard Leeb, *Actions of discrete groups on nonpositively curved spaces*, *Math. Ann.* **306** (1996), no. 2, 341–352. MR 1411351 [2]
14. Mustafa Korkmaz and Athanase Papadopoulos, *On the ideal triangulation graph of a punctured surface*, *Ann. Inst. Fourier (Grenoble)* **62** (2012), no. 4, 1367–1382. MR 3025746 [2, 4, 17]
15. Howard Masur and Saul Schleimer, *The geometry of the disk complex*, *J. Amer. Math. Soc.* **26** (2013), no. 1, 1–62. MR 2983005 [2, 11]
16. Lee Mosher, *Mapping class groups are automatic*, *Ann. of Math. (2)* **142** (1995), no. 2, 303–384. MR 1343324 [4]
17. Robert C. Penner, *Decorated Teichmüller theory*, QGM Master Class Series, European Mathematical Society (EMS), Zürich, 2012, With a foreword by Yuri I. Manin. MR 3052157 [4, 6]
18. N. Yu. Reshetikhin and V. G. Turaev, *Ribbon graphs and their invariants derived from quantum groups*, *Comm. Math. Phys.* **127** (1990), no. 1, 1–26. MR 1036112 [5]
19. J. Roger, *Ptolemy groupoids, shear coordinates and the augmented Teichmüller space*, ArXiv e-prints (2012). [2]
20. Michah Sageev, *CAT(0) cube complexes and groups*, Geometric group theory, IAS/Park City Math. Ser., vol. 21, Amer. Math. Soc., Providence, RI, 2014, pp. 7–54. MR 3329724 [6, 7]
21. Daniel T. Wise, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012. MR 2986461 [5]

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS, URBANA, IL
E-mail address: `mcbell@illinois.edu`

MATHEMATICAL SCIENCES RESEARCH INSTITUTE, BERKELEY, CA
E-mail address: `v.disarlo@gmail.com`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF OKLAHOMA, NORMAN, OK
E-mail address: `rtang@math.ou.edu`