┃ GAFA Geometric And Functional Analysis

CONTRACTIBLE, HYPERBOLIC BUT NON-CAT(0) COMPLEXES

RICHARD C. H. WEBB

Abstract. We prove that almost all arc complexes do not admit a CAT(0) metric with finitely many shapes, in particular any finite-index subgroup of the mapping class group does not preserve such a metric on the arc complex. We also show the analogous statement for all but finitely many disc complexes of handlebodies and free splitting complexes of free groups. The obstruction is combinatorial. These complexes are all hyperbolic and contractible but despite this we show that they satisfy no combinatorial isoperimetric inequality: for any n there is a loop of length 4 that only bounds discs consisting of at least n triangles. On the other hand we show that the curve complexes satisfy a linear combinatorial isoperimetric inequality, which answers a question of Andrew Putman.

1 Introduction

In general the mapping class group $\mathrm{Mod}(S_{g,p})$ cannot act properly by semisimple isometries on a complete CAT(0) space [KL96, BH99, Bri10], in particular, it is not a CAT(0) group. However, the Teichmüller space with the Weil–Petersson metric is CAT(0) [Tro86, Wol86, Wol87], furthermore so is its completion [BH99, Corollary II.3.11], and the mapping class group acts on this by semisimple isometries [DW03]. When $g \geq 3$, Bridson [Bri10] used this action to show that any non-trivial homomorphism

$$\mathrm{Mod}(S_{g,p}) \to \mathrm{Mod}(\Sigma),$$

must send Dehn twists to roots of multitwists. This result was then used by Aramayona–Souto [AS12] to classify—under topological assumptions on $S_{g,p}$ and Σ—all non-trivial homomorphisms from $\mathrm{Mod}(S_{g,p})$ to $\mathrm{Mod}(\Sigma)$. Despite the fact that the mapping class group is not CAT(0) in general, the study of its algebra has been enhanced by its action on CAT(0) spaces.

Perhaps the most striking application of actions on CAT(0) spaces has been provided by CAT(0) cube complexes and their role in the proof of the virtual Haken conjecture, see [Ago13]. Not only did this uncover a difficult topological consequence but there were many exciting *algebraic* consequences for the fundamental groups of closed, hyperbolic 3-manifolds such as largeness, LERF, linearity over $\mathbb{Z}$,
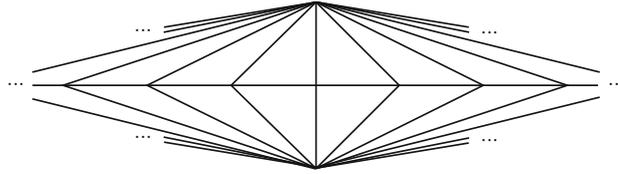
Figure 1: A natural triangulation of the suspension of $\mathbb{R}$ produces a non-locally compact, 2-dimensional simplicial complex that satisfies no combinatorial isoperimetric inequality: there are loops of combinatorial length four that require arbitrarily many triangles to deform them to a point. From an appropriate embedding in the euclidean plane it can be endowed with a complete CAT(0) metric.

bi-orderability, conjugacy separability, see for example [AFW15] and the references therein.

One theme of this paper concerns the problem of finding CAT(0) metrics on a given space. The spaces of interest in this paper are locally infinite complexes, such as the arc complex of a surface and the free splitting complex of a free group. We say that a complex K (equipped with a metric) has *finitely many shapes* if there are only finitely many isometry classes of simplex in K.

Bridson in his thesis [Bri91] initiated the study of non-positive curvature for metric complexes that are not necessarily locally finite. He showed that a metric simplicial complex with finitely many shapes (modelled on euclidean, spherical or hyperbolic simplices) is a complete geodesic metric space; he gave a list of equivalent global and/or local conditions for such a complex to be CAT(0); and he introduced a technique for deciding whether a complex can be endowed a CAT(0) metric equivariant with respect to some group action and applied this to Culler–Vogtmann's outer space.

We show the following

Theorem 1.1. *Whenever K is a (not necessarily locally compact) flag simplicial complex equipped with a CAT(0) metric with finitely many shapes then K satisfies a quadratic combinatorial isoperimetric inequality.*

Any CAT(0) space satisfies a quadratic *coarse* isoperimetric inequality. We remark that this observation does not suffice to prove Theorem 1.1. We point the reader to Figure 1 for an elementary example.

As far as the author is aware, Theorem 1.1 is not in the literature and is new. Perhaps this is because the focus of the isoperimetric inequality has mainly been on finitely presented groups, where the optimal combinatorial isoperimetric inequality and the optimal coarse isoperimetric inequality (the *Dehn function*) are equivalent for Cayley complexes. Cayley complexes are always locally compact with a cocompact group action—both of these properties are necessary for the equivalence between the coarse and the combinatorial.

It is known that in certain cases of low complexity the arc complex is CAT(0), and the disc complex/free splitting complex is CAT(0) when the genus/rank is 2.

In Section 3 we use Theorem 1.1 to show that the majority of arc complexes do not admit a CAT(0) metric with finitely many shapes. The same holds for all but finitely many disc complexes of handlebodies and all but finitely many free splitting complexes.

Theorem 1.2. *Let K be one of the following complexes.*

- *The arc complex $\mathcal{A}(S_{g,p})$ where $g \geq 2$ and $p \geq 2$ (or $p \geq 6 - 2g$ when $g = 0$ or 1).*
- *The disc complex $\mathcal{D}_n$ of a handlebody of genus $n \geq 5$.*
- *The free splitting complex $\mathcal{FS}_n$ of a free group of rank $n \geq 5$.*

Then there is a family of loops c_N of combinatorial length 4 in $K^{(1)}$ such that the following holds. Whenever P is a triangulation of a surface with one boundary component and $f \colon P \to K^{(2)}$ is a simplicial map where $f|_{\partial P}$ maps bijectively onto c_N then P must have at least N triangles. In particular K does not admit a CAT(0) metric with finitely many shapes.

In other words, each of the above complexes does not satisfy any combinatorial isoperimetric inequality at all. Note that all of the above complexes are contractible [Har85, Hat91, McC91, Hat95], hyperbolic [MS13, HM13], and such that any finite group action must fix some point [Ker83, Har86, HOP14], and so the usual and more well-known obstructions to being CAT(0) do not apply. Our obstruction to being CAT(0) with finitely many shapes is new. We leave the question of whether the complexes in the remaining cases are CAT(0) an open question.

On the other hand the above complexes are in stark contrast to the curve complexes and the arc-and-curve complexes, as we show

Theorem 1.3. *The curve complex $\mathcal{C}(S)$ and arc-and-curve complex $\mathcal{AC}(S)$ satisfy a linear combinatorial isoperimetric inequality.*

In [MM99, p. 104] Masur and Minsky state that it is an interesting question whether the arc-and-curve complexes $\mathcal{AC}(S)$ admit a CAT(κ) metric for some $\kappa \leq 0$. Note that any example is immediately also CAT(0). A bold guess based on Theorem 1.2 seems to be no. The proof of Theorem 1.1 also applies to higher-dimensional combinatorial isoperimetric inequalities, see Theorem 2.13, and at present it is unclear whether $\mathcal{AC}(S)$ should satisfy any at all. This avenue might lead to a negative answer to Masur and Minsky's question. We note that this question is important: if the answer is yes with an isometric action of a finite-index subgroup of Mod(S), then it might prove to be a useful tool for understanding Mod(S) and its quotients in a similar vein to [DHS].

We sketch the idea of the proof of Theorem 1.3 at the start of Section 4. The most important feature is the use of Masur and Minsky's tightening procedure i.e. the idea of the proof of the existence of tight geodesics [MM00, Lemma 4.5].

Presently the tightening procedure (and tight geodesics) is one of the missing pieces in the analogous Masur–Minsky theory for $Out(F_n)$. Such a theory could

prove useful for example in showing that the action on the free factor complex is acylindrical, which seems to be open, but known for the curve complex [Bow08]. However much intense progress has been made in discovering/creating analogues of hyperbolicity [BF14a, HM13, HH17] and subsurface projection [BF14b, Tay14]. Theorem 1.3 is a new phenomenon that is not yet observed in the $Out(F_n)$ theory. Because the tightening procedure leads to a proof of Theorem 1.3, there is an implied relationship between the two. The following question is therefore interesting: is there a cocompact complex for $Out(F_n)$, analogous to $\mathcal{C}(S)$, that satisfies a linear combinatorial isoperimetric inequality? Note that the free-and-cyclic splitting complex for F_n, for $n \geq 6$, does not satisfy any combinatorial isoperimetric inequality at all (see Remark 3.4, we thank the anonymous referee for suggesting this). The arc-and-curve complex $\mathcal{AC}(S)$ consists of geometric free and cyclic splittings of the fundamental group of S. It is interesting that Theorem 1.3 holds in the surface setting but does not in the free group setting.

We remark that it is currently not known whether the pants complex satisfies a quadratic combinatorial isoperimetric inequality. In light of Theorem 1.1 this is related to a question of Brock who asked whether there is a CAT(0) complex analogous to the pants complex (it is already a theorem of Brock that the pants complex is quasi-isometric to the Weil–Petersson metric [Bro03]). Recently Islambouli and Klug showed that any loop in the 1-skeleton of the pants complex of a closed surface determines a smooth 4-manifold [IK], and that any smooth 4-manifold arises by some such loop. They also define a signature that provides a lower bound on the number of triangles for any disc bounding that loop. It would seem that the structure of smooth 4-manifolds could give new insight on the pants complex, or possibly the other way round, or both.

2   Combinatorial isoperimetric inequalities

In this section we show the following, which is a generalisation of Theorem 1.1. The notions of *bounded shapes* and *thick shapes* are defined below.

Theorem 2.1. *Whenever K is a (not necessarily locally compact) flag simplicial complex equipped with a CAT(0) metric with bounded, thick shapes then K satisfies a quadratic combinatorial isoperimetric inequality.*

It is well known that CAT(0) spaces satisfy a quadratic *coarse* isoperimetric inequality (see [BH99, Chapter III.H.2.4]) but for complexes that are not locally compact this does not suffice to prove a quadratic *combinatorial* isoperimetric inequality, see Figure 1 for an example. In fact any two of the three hypotheses (CAT(0), bounded shapes, or thick shapes) of Theorem 2.1 are satisfied by some metric on the complex given in Figure 1 and so all three hypotheses are necessary.

The main theorems of this paper concern simplicial complexes and therefore for brevity we focus on this case. We expect that Theorem 2.1 is true for more general complexes K by using the methods given here but we have not checked the details.

2.1   Some background.   Now we explain the definitions in Theorem 2.1.

Let P and K be (simplicial) complexes. A *simplicial map* is a map $c\colon P \to K$ such that whenever $\Delta_P \subset P$ is a simplex then the image of the vertex set of Δ_P under c is a set of vertices of K that span a simplex $\Delta_K \subset K$. The map c may also be thought of as a continuous map between the spaces P and K: we may define c on any simplex $\Delta_P \subset P$ by extending linearly the images of the vertices of Δ_P that span $\Delta_K \subset K$.

The *i-skeleton* $K^{(i)}$ of K is the unique subcomplex of K consisting of all simplices of dimension at most i in K.

A *combinatorial loop* c in K is a sequence of vertices $(v_1, \ldots, v_j)$ of $K^{(1)}$ where v_j is adjacent (or equal) to v_1 and v_i is adjacent (or equal) to v_{i+1} whenever $1 \leq i \leq j-1$. The *combinatorial length* $l_C(c)$ of $c = (v_1, \ldots, v_j)$ is equal to j. A combinatorial loop c in K may also be thought of as a simplicial map $c\colon P \to K$ where P is a triangulation of S^1 with j 1-simplices.

We write D^2 for the closed unit disc with boundary S^1. Let c be a combinatorial loop in K. We say that *c can be capped off with at most n triangles* if there is a triangulation P of D^2 into at most n 2-simplices and there is a simplicial map $c'\colon P \to K^{(2)}$ such that $c'|_{S^1} = c$. In more informal words, the loop c in K can be deformed continuously to a point by pushing c past at most n triangles.

DEFINITION 2.2. Let K be a simplicial complex. A function $f\colon \mathbb{N} \to \mathbb{N}$ is called a *combinatorial isoperimetric bound* for K if every combinatorial loop c in K can be capped off with at most $f(l_C(c))$ triangles.

We say that K satisfies a *linear (or quadratic) combinatorial isoperimetric inequality* if there exists a combinatorial isoperimetric bound f for K such that $f(n) = O(n)$ (or $f(n) = O(n^2)$). We say that K *satisfies no combinatorial isoperimetric inequality* if no combinatorial isoperimetric bound of K exists.

REMARK 2.3. Another way of defining combinatorial isoperimetric inequalities is by using van Kampen diagrams, see [BH99, Chapter I.8A.4]. Van Kampen diagrams are combinatorial maps from planar 2-complexes (these are not necessarily discs, they may have separating edges) to K. Notwithstanding it is well known that the van Kampen perspective and Definition 2.2 are equivalent for complexes where the 2-cells have boundedly many sides, in particular, in this paper the two notions agree. As we will see in Section 3, for our purposes it is more convenient to use Definition 2.2.

REMARK 2.4. Sometimes the existence of a combinatorial isoperimetric bound restricted to combinatorial loops of short length implies the existence of a bonafide combinatorial isoperimetric bound. For instance suppose that there exists B such that whenever a combinatorial loop c satisfies $l_C(c) \leq 16\delta$ then c can be capped off using at most B triangles. Furthermore, suppose that $K^{(1)}$ is δ-hyperbolic (where each edge has length equal to 1). Then K satisfies a linear combinatorial isoperimetric inequality. This follows from the linear coarse isoperimetric inequality (see
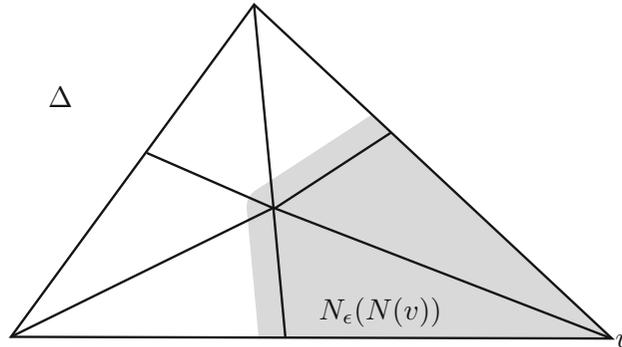
Figure 2: The first barycentric subdivision of a 2-simplex Δ. A closed neighbourhood $N_\epsilon(N(v))$ of $N(v)$ is shaded and is disjoint from the opposite 1-simplex; Δ is ϵ-thick.

[BH99, Chapter III.H 2.6 and 2.7]). This is well known for Rips complexes of hyperbolic groups but the focus in this paper is on *non-locally-compact* complexes. Note that for the complex in Figure 1 there is no such integer B despite $K^{(1)}$ being 2-hyperbolic.

Now we consider CAT(0) metrics on the topological spaces determined by simplicial complexes. We assume that the reader knows the definition of a CAT(0) metric space, see [BH99, Chapter II.1.1]. For cosmetic reasons, in places where no confusion will arise, we abuse notation by writing K for the simplicial complex, topological space, and the metric space (K, d_K).

DEFINITION 2.5. We say that (K, d_K) has *finitely many shapes* if there are only finitely many isometry types of metrics $d : \Delta \times \Delta \to \mathbb{R}$ on the simplices Δ of K that are induced by restriction of d_K.

We remark that we do not need to assume that the isometry types of simplices are euclidean.

We write $N_\epsilon(A)$ for the closed ϵ-neighbourhood of a subset $A \subset K$. Likewise we write $N_\epsilon(x) = N_\epsilon(\{x\})$ for an element $x \in K$.

DEFINITION 2.6. We say that K has *bounded* shapes if for every $\delta > 0$ there is a subdivision of each edge of K into at most $L = L(K, \delta)$ intervals, such that each interval has diameter at most δ.

We say that K has *thick* shapes if there exists some first barycentric subdivision $\tilde{K}$ of K such that the following holds. For each $v \in K^{(0)}$, set $N(v)$ to be the closure of the union of simplices $\tilde{\Delta}$ of $\tilde{K}$ such that $v \in \tilde{\Delta}$. Then for every simplex Δ of K, whenever $N_\epsilon(N(v)) \cap \Delta \neq \emptyset$ then $v \in \Delta^{(0)}$ (see Figure 2).

In the definition of thick shapes above, we only need the existence of some geometric realization of the barycentric subdivision with the desired properties.

In particular, having thick shapes means that disjoint simplices are uniformly far apart in K. Having bounded, thick shapes is a natural generalization of having finitely many shapes.

LEMMA 2.7. *If K has finitely many shapes then K has bounded, thick shapes.*

Sketch of proof. Clearly K has bounded shapes. For thick shapes, for each simplex of K we perform the first barycentric subdivision according to its isometry class, and then the finiteness provides some small enough $\epsilon > 0$ to exist. $\qquad\square$

REMARK 2.8. It is an interesting question whether the complexes in Theorem 3.1 admit a CAT(0) metric at all. If there is such a metric one can ask whether an interesting subgroup H of the automorphism group still acts by isometries. By Theorem 3.1 such a CAT(0) metric fails to have bounded, thick shapes, for example the edge lengths might be arbitrarily small (or large). In particular such a subgroup H cannot have finite index.

2.2 Proof of Theorem 2.1.

The rest of this section is devoted to proving Theorem 2.1. For a point x in the topological space K we write $\Delta(x)$ for the minimal simplex of K such that $x \in \Delta(x)$.

Proof of Theorem 2.1. There are different proofs but the proof given here determines an explicit way of capping off any combinatorial loop c in K. We give a careful proof because K is not necessarily locally compact.

Let $\epsilon > 0$ be such that K has ϵ-thick shapes. With respect to Definition 2.6, we choose $\delta = 0.5\epsilon$, so we may subdivide each edge of K into at most $L(K, \delta)$ intervals such that each interval has diameter at most δ.

The strategy of the proof is as follows. Given a combinatorial loop c we abuse notation and write c for the same topological loop in K. We may parametrise the loop by $c \colon [0, l_C(c)] \to K$, not necessarily parametrised by arc length. We subdivide the loop c in K into intervals with endpoints $c(t_i)$ and $c(t_{i+1})$ at most 0.5ϵ apart (using the subdivision of the edges of K guaranteed by bounded shapes). There are at most $l_C(c) \cdot L(K, \delta)$ intervals. For each i, we construct combinatorial paths $(v_j^i)_j$ connecting $v_0^i = c(0)$ to $v_{n(i)}^i$ where $c(t_i) \in N(v_{n(i)}^i)$ (recall $N(v)$ from Definition 2.6). The sequence $(v_{n(i)}^i)_i$ defines a combinatorial loop in K, which is homotopic to the combinatorial loop c (we discuss this at the very end of the proof). It suffices to cap off $(v_{n(i)}^i)_i$ with a triangulation P'. We do this by constructing a triangulation $P(i)$ between $(v_j^i)_j$ and $(v_j^{i+1})_j$ using linearly many triangles in terms of $l_C(c)$ (and there are at most linearly many i in terms of $l_C(c)$). Gluing up the $P(i)$ gives a triangulation P' that caps off $(v_{n(i)}^i)_i$. Now we give the details.

Using the subdivision of the edges of K, we may fix $0 = t_0 < t_1 < \cdots < t_N = l_C(c)$ such that the intervals $c[t_i, t_{i+1}]$ correspond to paths of diameter at most 0.5ϵ in K. We have $N \leq l_C(c) \cdot L(K, \delta)$. We write $c_i \colon [0, d_K(c(0), c(t_i))] \to K$ for the

geodesic (using the CAT(0) metric on K) parametrised by arc length that connects $c(0)$ and $c(t_i)$.

Now we construct explicit combinatorial paths that start at $c(0)$ by using the geodesics c_i. Fix i and set $v_0^i = c(0)$ and $t_0^i = 0$. Now for $j \geq 1$ we define v_j^i and t_j^i inductively as follows. Note that $c_i(t_0^i) \in N(v_0^i)$ (recall $N(v)$ from Definition 2.6), which is the base case.

- We set $t_j^i > t_{j-1}^i$ to be minimal such that $v_{j-1}^i \notin \Delta(c_i(t_j^i))$. Then we pick any $v_j^i \in K^{(0)}$ such that $c_i(t_j^i) \in N(v_j^i)$. By ϵ-thick shapes we have $t_j^i - t_{j-1}^i \geq \epsilon$. We have v_j^i adjacent to v_{j-1}^i.
- On the other hand if no such t_j^i above exists then v_{j-1}^i is adjacent (or equal) to some vertex v such that $c(t_i) \in N(v)$. We set $v_j^i = v$ and $t_j^i = d_K(c(0), c(t_i))$. Set $n(i) = j$. We stop.

We note that this inductive process must terminate: Whenever $0 < j < n(i)$ then $t_j^i - t_{j-1}^i \geq \epsilon$ by ϵ-thick shapes. Hence for $j < n(i)$ we have $t_j^i \geq j\epsilon$. Thus

$$n(i) \leq \epsilon^{-1} d_K(c(0), c(t_i)) + 1 \leq 0.5 L(K, \delta) \cdot l_C(c) + 1.$$

Now for each i such that $0 \leq i \leq N-1$ we build triangles between the combinatorial paths $(v_j^i)_j$ and $(v_j^{i+1})_j$. We use the comparison triangle $T(i)$ in $\mathbb{E}^2$ with sides E_1, E_2 and E_3 with side lengths $t_{n(i)}^i$, $d_K(c(t_i), c(t_{i+1}))$ and $t_{n(i+1)}^{i+1}$ respectively. Note that $d_K(c(t_i), c(t_{i+1})) \leq 0.5\epsilon$. We sweep out $T(i)$ using intervals of length at most 0.5ϵ that are parallel to E_2. Write

$$R_i = \frac{t_{n(i+1)}^{i+1}}{t_{n(i)}^i}.$$

We may parametrise E_1 and E_3 by arc length such that $E_1(0) = E_3(0)$. Then the intervals parallel to E_2 are parametrised by their endpoints $E_1(s)$ and $E_3(R_i s)$ for $0 \leq s \leq t_{n(i)}^i$. In K the intervals' endpoints correspond to $c_i(s)$ and $c_{i+1}(R_i s)$.

We endow $\partial T(i)$ with a simplicial structure induced from the points $E_1(t_j^i)$ and $E_3(t_j^{i+1})$.

Now we define edges (written $e(m)$ below) between the points $E_1(t_j^i)$ and the points $E_3(t_k^{i+1})$ in order to triangulate $T(i)$. We justify doing this by finding edges (or proving equalities) between the appropriate vertices v_j^i and v_k^{i+1} in K. This is the hardest part of the proof and is the content of Lemmas 2.9 and 2.10. See Figure 3 for an example triangulation.   □

We lexicographically order the elements $(R_i t_j^i, i)$ and $(t_k^{i+1}, i+1)$ with the largest first. This ordering comes from the sweeping out of $T(i)$ by lines parallel to E_2: starting with E_2 itself and ending at the vertex opposite to E_2, we pass the vertices
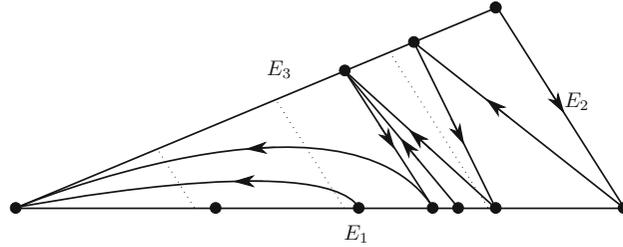
Figure 3: An example (topological) triangulation $P(i)$ of $T(i)$. Here, $n(i) = 6$, $n(i+1) = 3$, and there are 8 directed edges between E_1 and E_3. The edges pointing toward E_1 are constructed in Lemma 2.9. The edges pointing toward E_3 are constructed in Lemma 2.10. The first edge $e(1)$ always coincides with E_2. As we sweep out $T(i)$ starting from E_2, we encounter the edges $e(1), e(2), e(3) \ldots$ in that order.

of $T(i)$ namely $E_1(t_j^i)$ and $E_3(t_k^{i+1})$ and order them. We parametrise the above elements $(s(m), i(m))$ in order using a parameter m where

$$1 \leq m \leq n(i) + n(i+1) + 2.$$

We introduce another parameter $j(m)$ so that the element $(s(m), i(m))$ will correspond to $v_{j(m)}^{i(m)}$ and $t_{j(m)}^{i(m)}$.

Now for

$$1 \leq m \leq n(i) + n(i+1) - 1$$

we find exactly one outward edge from (or another vertex equal to) $v_{j(m)}^{i(m)}$ using the following two lemmas. This enables us to define edges $e(m)$ that cut $T(i)$ into pieces that will give the required triangulation $P(i)$ of $T(i)$, see Figure 3.

LEMMA 2.9. *If $i(m) = i + 1$ then v_j^i and $v_{j(m)}^{i+1}$ share an edge (or are equal) where j is largest such that $t_j^i \leq R_i^{-1} t_{j(m)}^{i+1}$.*

Proof. We argue that

$$v_j^i, v_{j(m)}^{i+1} \in \Delta(c_i(R_i^{-1} t_{j(m)}^{i+1})).$$

Using the comparison triangle $T(i)$ we have

$$d_K(c_i(R_i^{-1} t_{j(m)}^{i+1}), c_{i+1}(t_{j(m)}^{i+1})) \leq 0.5\epsilon.$$

Therefore we have $N_\epsilon N(v_{j(m)}^{i+1}) \cap \Delta(c_i(R_i^{-1} t_{j(m)}^{i+1})) \neq \emptyset$. By ϵ-thick shapes we have $v_{j(m)}^{i+1} \in \Delta(c_i(R_i^{-1} t_{j(m)}^{i+1}))$.

Finally, we have j largest such that $t_j^i \leq R_i^{-1} t_{j(m)}^{i+1}$. By definition of v_j^i and t_j^i we have $v_j^i \in \Delta(c_i(R_i^{-1} t_{j(m)}^{i+1}))$. □

By Lemma 2.9, as the corresponding vertices in K are either equal or adjacent, we construct an edge between $E_1(t_j^i)$ and $E_3(t_{j(m)}^{i+1})$ in $T(i)$, see Figure 3. This defines the required mth edge $e(m)$ inside $T(i)$ when $i(m) = i + 1$.

However if $i(m) = i$ then we use

LEMMA 2.10. *If* $i(m) = i$ *then* $v_{j(m)}^i$ *and* v_k^{i+1} *share an edge (or are equal) where* k *is largest such that* $R_i t_{j(m)}^i > t_k^{i+1}$.

Proof. We argue that for sufficiently small $\delta' > 0$ we have

$$v_{j(m)}^i, v_k^{i+1} \in \Delta(c_{i+1}(R_i t_{j(m)}^i - \delta')).$$

Indeed, for sufficiently small $\delta' > 0$ we have

$$d_K(c_i(t_{j(m)}^i), c_{i+1}(R_i t_{j(m)}^i - \delta')) \le \epsilon.$$

Therefore we have $N_\epsilon N(v_{j(m)}^i) \cap \Delta(c_{i+1}(R_i t_{j(m)}^i - \delta')) \ne \emptyset$ so by ϵ-thick shapes we have $v_{j(m)}^i \in \Delta(c_{i+1}(R_i t_{j(m)}^i - \delta'))$.

Finally for $\delta' > 0$ sufficiently small we have by definition

$$v_k^{i+1} \in \Delta(c_{i+1}(R_i t_{j(m)}^i - \delta'))$$

because k is largest such that $R_i t_{j(m)}^i > t_k^{i+1}$. $\qquad\square$

By Lemma 2.10, as the corresponding vertices in K are either equal or adjacent, we construct an edge between $E_1(t_{j(m)}^i)$ and $E_3(t_k^{i+1})$ in $T(i)$. This defines the required mth edge $e(m)$ inside $T(i)$ for when $i(m) = i$, see Figure 3.

Now the edges $e(m)$ cut $T(i)$ into triangles that give a triangulation $P(i)$. This can be seen because the edges $e(m)$ and $e(m+1)$ share one vertex in common, and their endpoints correspond to values t_j^i and t_k^{i+1} that are monotonically decreasing in m. The existence of the edges in K provide a natural map $P(i) \to K$ that extends the natural map $\partial T(i) \to K$. This is a simplicial map because K is flag. We used exactly $n(i) + n(i+1) - 1$ triangles, which is at most $L(K, \delta) \cdot l_C(c) + 1$.

Now we can cap off the combinatorial loop $(v_{n(i)}^i)_{i=1}^N$ (it starts and ends at $c(0)$) with a triangulation P' that is constructed by gluing (for each i) $P(i)$ to $P(i+1)$ along the combinatorial path $(v_j^{i+1})_{j=0}^{n(i+1)}$. The number of triangles of P' is at most $L(K, \delta) \cdot l_C(c) + 1$ multiplied by N. This is at most some quadratic function in $l_C(c)$.

Finally, we need to glue triangles along the boundary of P' to construct a triangulation P that caps off the original combinatorial loop c. Write $c = (v_0, \ldots, v_{l-1})$ where $l = l_C(c)$. Each vertex v_i is equal to several consecutive vertices

$$v_{n(j)}^j, v_{n(j+1)}^{j+1}, \ldots, v_{n(k)}^k,$$

this is because the intervals $c[t_j, t_{j+1}], \ldots, c[t_{k-1}, t_k]$ are subsets of edges of K lying entirely within $N(v_i)$ (recall the construction of the $v_{n(j)}^j$ and Definition 2.6). So we attach triangles between these for each i. Lastly there is a triangle spanned by v_i, v_{i+1} and $v_{n(k)}^k$, see Figure 4. This adds only at most $l_C(c) + N$ triangles. $\qquad\square$
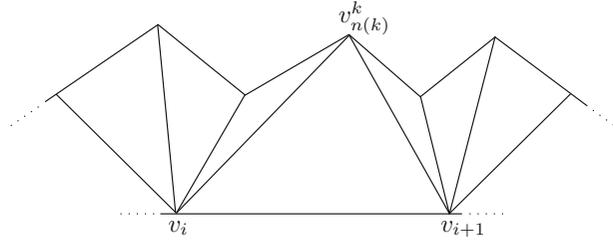
Figure 4: The last step. Adding triangles to P' to construct P whose boundary maps onto the original combinatorial loop c.

2.3 Higher-dimensional combinatorial isoperimetric inequalities. In this section we prove that a CAT(0) flag simplicial complex with finitely many shapes must satisfy quadratic higher-dimensional combinatorial isoperimetric inequalities. The proof is versatile enough to state it in greater generality so we state Theorem 2.13 below. The following definition is meant to be a d-dimensional analogue of the notion of bounded shapes, which was a constraint on the 1-simplices of K.

DEFINITION 2.11. Let $d \in \mathbb{N}$. We say that (K, d_K) has *d-bounded* shapes if given any $\delta > 0$ there is a subdivision of the d-skeleton $K^{(d)}$ into a collection of smaller simplices Δ such that:-

- each Δ has diameter at most δ, and,
- every simplex of $K^{(d)}$ is subdivided into a uniformly bounded number of simplices depending only on K and δ.

REMARK 2.12. If K has finitely many shapes then K has d-bounded, thick shapes for any $d \geq 1$.

Let X be a (simplicial) complex. We write $CX = (X \times [0,1])/(X \times \{0\})$ for the *cone* of X. We will abuse notation by writing $X = (X \times \{1\})$ for the corresponding subset of CX. When X is d-dimensional we write $|X|$ for the number of d-simplices in X.

Theorem 2.13. *Fix $d > 0$ and K a (not necessarily locally compact) flag simplicial complex equipped with a CAT(0) metric with d-bounded, thick shapes. Then there exists a quadratic function $f \colon \mathbb{N} \to \mathbb{N}$ such that whenever X is a connected d-dimensional complex and $g \colon X \to K$ is a simplicial map then there is a triangulation P of the cone CX of X and a simplicial map $g' \colon P \to K$ such that $g'|_X = g$, with $|P| \leq f(|X|^2)$.*

Proof. We follow the proof of Theorem 2.1, which is essentially a prerequisite. The first step is to subdivide $K^{(d)}$ into simplices of diameter at most 0.5ϵ, which is the only place in the proof where we require that K is d-bounded. Such a subdivision pulls back to X, and this triangulation of X we call P'', and there is a natural topological map $h \colon P'' \to K$. For each vertex u of P'' we pick $w \in K^{(0)}$ such

that $h(u) \in N(w)$—this defines a map $g'' \colon P'' \to K$ by setting $g''(u) = w$ and by ϵ-thickness we have that g'' is a simplicial map.

Let w_0 be an arbitrary vertex in the image (under g) of X in K. We wish to find a triangulation P''' of the cone CX where the induced triangulation on $X \times \{1\}$ is P'', and have a simplicial map $g''' \colon P''' \to K$ extending g''. As a start we write u_0 for the vertex endpoint of the cone CX (i.e. the point representing $X \times \{0\}/\sim$) and set $g'''(u_0) := w_0$.

We now want to construct many combinatorial paths between u_0 and $u \in P''$, and send them to combinatorial paths in K. We then fix a vertex u of P''. We take the geodesic in the CAT(0) metric between $h(u_0)$ and $h(u)$. We define these combinatorial paths just as before in Theorem 2.1. Now let $u_1, \dots, u_{d+1}$ span a simplex Δ in P'', we need to ensure that the combinatorial paths between u_0 and the $u_i \in \Delta$ extend to a triangulation of $C\Delta$ and a simplicial map $g'''_{C\Delta} \colon C\Delta \to K$. The construction and proof of this is a straightforward generalisation of Lemmas 2.9 and 2.10. This then provides a triangulation P''' of CX and a simplicial map $g''' \colon P''' \to K$. The number of $(d+1)$-simplices of P''' is at most a uniform multiplicative constant multiple of $|X|$, multiplied by the diameter of X, which is again at most $|X|$, hence $O(|X|^2)$.

However g''' is not necessarily equal to g restricted to $X \times \{1\}$. Just as we did at the end of the proof of Theorem 2.1 we add triangles to P''' to find the required P and g'. This adds on more $(d+1)$-simplices to CX, but at most some uniform constant multiple of $|X|$. $\square$

3   Contractible, hyperbolic complexes with no combinatorial isoperimetric inequality

In this section we prove the following theorem.

Theorem 3.1. *Let K be one of the following complexes.*

(1) *The arc complex $\mathcal{A}(S_{g,p})$ where $g \geq 2$ and $p \geq 2$, or, $g = 1$ and $p \geq 4$, or $g = 0$ and $p \geq 6$.*
(2) *The disc complex $\mathcal{D}_n$ of a handlebody of genus $n \geq 5$.*
(3) *The free splitting complex $\mathcal{FS}_n$ of a free group of rank $n \geq 5$.*

Then there is a family of loops c_N of combinatorial length 4 in $K^{(1)}$ such that the following holds. Whenever P is a triangulation of a surface with one boundary component and $f \colon P \to K^{(2)}$ is a simplicial map where $f|_{\partial P}$ maps bijectively onto c_N then P must have at least N triangles.

Combining this with Theorem 2.1 and Lemma 2.7 we have

COROLLARY 3.2. *Let K be as in Theorem 3.1. Then K does not admit a CAT(0) metric with finitely many shapes. Furthermore it does not admit one with bounded, thick shapes.* $\square$

3.1 Most arc complexes.

We write $S_{g,p}$ for the orientable surface with genus g and p punctures/marked points.

Proof of Theorem 3.1(1). The hypothesis on g and p ensures that there exists an essential, non-peripheral simple closed curve γ on $S = S_{g,p}$ such that $S - \gamma$ consists of two connected subsurfaces Y and Z, both of which have at least one puncture (which is also a puncture of S), and $\xi(Y), \xi(Z) \geq 1$ (here $\xi(S_{g,p}) = 3g + p - 3$, which coincides with the number of curves in a pants decomposition). Therefore $\mathcal{A}(Y)$ and $\mathcal{A}(Z)$ both have infinite diameter.

Pick an arbitrary integer $N \geq 2$. There exist arcs a_1^Y, a_2^Y, a_1^Z and a_2^Z of S such that

(1) $a_1^Y, a_2^Y \subset Y$ and $a_1^Z, a_2^Z \subset Z$, and
(2) $d_{\mathcal{A}(Y)}(a_1^Y, a_2^Y) \geq 3N$ and $d_{\mathcal{A}(Z)}(a_1^Z, a_2^Z) \geq 3N$.

Note that the graph spanned by a_1^Y, a_1^Z, a_2^Y and a_2^Z in $\mathcal{A}(S)$ is a loop c_N whose combinatorial length is equal to 4.

Suppose that there is a triangulation P of a surface with one boundary component (for example a disc) and a simplicial map $f\colon P \to \mathcal{A}(S)^{(2)}$ such that $f|_{\partial P}$ coincides with c_N. Let us write $\tilde{a}_i^Y$ for $(f|_{\partial P})^{-1}a_i^Y$ and similarly $\tilde{a}_i^Z$ for $(f|_{\partial P})^{-1}a_i^Z$. Then $\tilde{a}_1^Y$, $\tilde{a}_2^Y$, $\tilde{a}_1^Z$ and $\tilde{a}_2^Z$ are the four vertices on the boundary of P.

We will colour the vertices of $\mathcal{A}(S)$ then pullback this colouring to give a colouring of the vertices of P. First we need some terminology. We say that an arc a *cuts* Y if every representative of a intersects Y. If a cuts Y then we can define $\kappa_Y(a)$ a simplex of $\mathcal{A}(Y)$ (see the definition of π'_Y in [MM00, p. 918]). If a does not cut Y then we say that a *misses* Y. If a cuts Y then we colour a red. If a is red then $\kappa_Y(a)$ is defined and it is a collection of essential arcs in Y. If a misses Y then we colour a blue. If a is blue then a must cut Z and therefore $\kappa_Z(a)$ is defined and it is a collection of essential arcs in Z. Each vertex is coloured either red or blue (but not both).

We claim that there is a red path between $\tilde{a}_1^Y$ and $\tilde{a}_2^Y$ in $P^{(1)}$ or a blue path between $\tilde{a}_1^Z$ and $\tilde{a}_2^Z$ in $P^{(1)}$. This follows directly from the proof of the 2-dimensional Hex Theorem given in [Gal79, p. 820]. We recall that beautiful proof here. We construct a graph G with four vertices (illustrated as squares in Figure 5) with additional vertices that correspond to the triangles of P. Then we add edges to the vertices according to the rule depicted on the left of Figure 5; one imagines one colour as land and the other colour as the sea, and these edges represent the cliff between the two. Each vertex in G has degree at most 2 but there are precisely four vertices of degree 1. Therefore there exists a path between two square vertices in G, see Figure 5. The required red or blue path can be constructed directly from this.

But by definition of a_1^Y, a_2^Y, a_1^Z and a_2^Z, any such monochromatic path has length at least $3N$. This is because a red path of length ℓ can be used via the projection κ_Y to construct a path of length at most ℓ in $\mathcal{A}(Y)$ between a_1^Y and a_2^Y, and similarly for a blue path in $\mathcal{A}(Z)$. Therefore there are at least $3N$ edges of P. If we count each
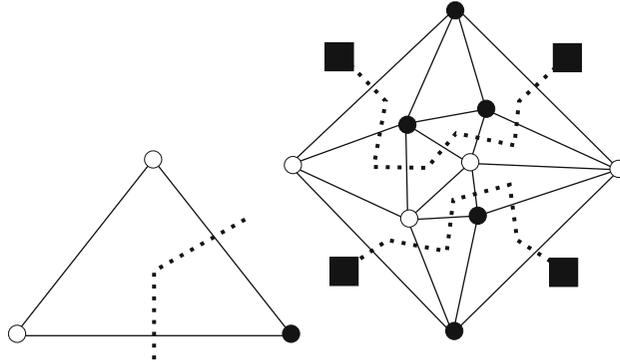
Figure 5: Left: How the subgraph of the dual graph of P is constructed locally. Right: An example construction.

triangle three times then we count each edge at least once, and so P must have at least N triangles. But $N \geq 2$ was arbitrary, and our loop had combinatorial length 4, so we are done. $\qquad\square$

3.2 All but finitely many disc complexes of handlebodies

Proof of Theorem 3.1(2). The surface $S = S_{0,p}$ has Euler characteristic $1 - (p-1)$, so the rank of its (free) fundamental group is $p-1$. The 3-manifold $H = S \times [0,1]$ is a handlebody of genus $p-1$ and so its boundary ∂H is a closed surface of genus $p-1$. Now let us set $n+1 = p \geq 6$.

Given $N \geq 2$, we take the arcs a_1^Y, a_2^Y, a_1^Z and a_2^Z for $S = S_{0,p}$ as constructed in the proof of Theorem 3.1(1). The loop c_N for the disc complex that we take is the one spanned by $a_1^Y \times [0,1]$, $a_2^Y \times [0,1]$, $a_1^Z \times [0,1]$ and $a_2^Z \times [0,1]$. Now we show that we require at least N triangles for any such triangulation P and simplicial map $f \colon P \to \mathcal{D}_n^{(2)}$ as in the statement of the theorem.

Let X be the subsurface of ∂H corresponding to $S \times \{0\}$. Then κ_X defines a map $\mathcal{D}_n^{(1)} \to \mathcal{A}(S)^{(1)}$ via $D \mapsto \kappa_X \partial D$, which is defined for all discs $D \in \mathcal{D}_n$ because the boundary ∂D of every essential disc D cuts $X = S \times \{0\}$. If D_1 and D_2 were disjoint then so are $\kappa_X \partial D_1$ and $\kappa_X \partial D_2$, and so by taking an arbitrary arc in each $\kappa_X \partial D$ we may define a 1-Lipschitz map from $\mathcal{D}_n^{(1)}$ to $\mathcal{A}(S)^{(1)}$, which we call g.

Note that g sends the loop c_N to the original loop spanned by a_1^Y, a_2^Y, a_1^Z and a_2^Z in the arc complex.

Therefore given $f \colon P \to \mathcal{D}_n^{(2)}$ the composition $gf \colon P \to \mathcal{A}(S)^{(2)}$ provides us with a simplicial map where $(gf)|_{\partial P}$ coincides with the original loop of length 4 in the arc complex. Therefore P has at least N triangles as in the proof of Theorem 3.1(1). $\qquad\square$

3.3 All but finitely many free splitting complexes of free groups.
We shall give two proofs. The first uses a result of Hamenstädt–Hensel, Theorem 3.1(1)

and Theorem 2.1. The second proof, suggested by the anonymous referee, is more similar to the proof of Theorem 3.1(1) but works directly for the free splitting complex.

First proof of Theorem 3.1(3). We shall use   □

PROPOSITION 3.3. *(Proposition 4.18 of* [HH15]*) Let S be a compact surface with genus g and $|\partial S| = b$. There is a canonical 1-Lipschitz embedding $\mathcal{A}(S) \to \mathcal{FS}_n$ where $n = 2g + b - 1$, and there exists a 1-Lipschitz left inverse $\mathcal{FS}_n \to \mathcal{A}(S)$.*

It is worth remarking that while the embedding is canonical their 1-Lipschitz left inverse is not because it depends on a choice of maximal arc system of S. For our purposes a 1-Lipschitz left inverse is crucial. It is important to note that canonical *coarsely* Lipschitz left inverses to this map have been given by Bowditch and Iezzi [BI18] and by Forlini [For] (who also addresses the arc-and-curve complex and the cyclic splitting complex).

We note that [HH15, Proposition 4.18] is formally stated only for the case $b = 1$. It is remarked [HH15, Remark 4.1] that minor modifications should enable the general case that we require here. The modifications required are as follows. In [HH15, Definition 4.7] the notion of *minimal position* is defined for an arc a relative to δ with respect to a sphere system Σ. This is then used in [HH15, Definition 4.14] to define A_0-*tight minimal position* for a map $\varphi \colon S \to W$ representing the homotopy class of S in the double W of $S \times [0,1]$. Here, A_0 is some maximal arc system of S. The minor modification is to define A_0-tight minimal position slightly: instead have A_0 a maximal arc system equipped with an ordering of the arcs $A_0 = (a_1, \ldots, a_m)$. Then we tweak the notion of minimal position [HH15, Definition 4.7(ii)] for a_i: we require that if $\tilde{\delta}$ is a lift of δ containing an endpoint q of $\tilde{a}_i$, then it intersects no sphere σ that $\tilde{a}_i$ intersects, unless $\tilde{\delta}$ contains an endpoint of some $\tilde{a}_j$ inbetween q and $\tilde{\delta} \cap \sigma$ (with $j < i$). With this definition A_0-tight minimal position exists and is unique [HH15, Lemmas 4.15, 4.16], and provides us with the 1-Lipschitz left inverse we require.

The 1-Lipschitz left inverse enables us to argue analogously to that of the case of the disc complex. We take the loops from Theorem 3.1(1) and embed them into $\mathcal{FS}_n$. Any disc bounding such a loop can be mapped via Proposition 3.3 back into $\mathcal{A}(S)$, so we are done.   □

Here is an alternative proof suggested by the anonymous referee.

Second proof of Theorem 3.1(3). Let $a_1, \ldots, a_n$ be a free basis for F_n, where $n \geq 5$. Let A be generated by $a_1, \ldots a_m$, B be generated by $a_m, \ldots a_n$ and C be generated by a_m. Here, we choose m such that the rank of A and B is at least three. Such m exists because $n \geq 5$. Then $F_n = A *_C B$, which is a cyclic splitting of F_n.

There is a loop S_1, S_2, S_3, S_4 of length 4 in the free splitting complex of F_n of the following form. We may choose free splittings S_1^A and S_3^A of A such that C is elliptic and the distance between S_1^A and S_3^A in the free splitting complex of A is at least

$3N$. We define S_2^B and S_4^B similarly, with C elliptic, and a lower bound on distance of $3N$. It is here that we require the rank of A and B to be at least three to ensure the lower bound on distance. Now S_1^A and S_3^A can be extended to free splittings S_1 and S_3 for $A *_C B$ in which B is elliptic. We can do this because C is elliptic for S_1^A and S_3^A. Similarly S_2^B and S_4^B can be extended to S_2 and S_4 in which A is elliptic. We can also combine the free splittings S_1^A and S_2^B to make a common refinement of S_1 and S_2, and so on, which is why S_1, S_2, S_3, S_4 form a loop in the free splitting complex of F_n.

Now suppose we cap off the loop $(S_i)_i$ with a disc in the free splitting complex of F_n. As we did in the case of the arc complex, we may colour a vertex S *red* if A does not act elliptically on S, otherwise we colour S *blue*. If S is blue then B cannot act elliptically on S, for if it did, then A and B both act elliptically with the same fixed point because they both contain C, but this implies that F_n has a global fixed point, which is a contradiction.

By the 2-dimensional Hex Theorem there exists a monochromatic path $(S_j')_j$ connecting S_1 to S_3 or S_2 to S_4. If the path is red, then for each $S = S_j'$ we can consider the A-minimal subtree S^A of S, i.e. the minimal subtree of S on which $A < F_n$ acts. Since A does not act elliptically on S^A, we can use the edges of the quotient S^A/A to obtain a 1-Lipschitz projection from $(S_j')_j$ to the free splitting complex of A, and so by definition of S_1 and S_3, this means that the path $(S_j')_j$ has length at least $3N$. This implies that at least N triangles are needed to cap off $(S_i)_i$, as required. The case where the monochromatic path $(S_j')_j$ is blue is similar, by projecting to the free splitting complex of B instead. $\square$

REMARK 3.4. The second proof above also applies to the free-and-cyclic splitting complex of F_n when $n \geq 6$. To see this let A be generated by $a_1, \ldots, a_m, a_{m+1}$, B be generated by $a_m, a_{m+1}, \ldots, a_n$ and C be generated by a_m and a_{m+1}. We require the rank of A and B to be at least four, which is why we require $n \geq 6$. As before we choose S_1^A and S_3^A to be splittings in which C is elliptic but far apart in the free-and-cyclic splitting complex for A, and similarly we choose S_2^B and S_4^B for B. Then as before we may extend these to S_1, S_2, S_3, S_4 a loop in the free splitting complex for F_n. If we cap this off with a disc in the free-and-cyclic splitting complex, then given a vertex S, we have that A and B cannot both be elliptic on S, and so we can define a projection to the free-and-cyclic splitting complex of A (or B) whenever A does not act elliptically on S (or when B does not act elliptically). The proof then follows using the 2-dimensional Hex Theorem as before. We thank the anonymous referee for suggesting this and the proof above.

4   Complexes satisfying a linear combinatorial isoperimetric inequality

In this section we prove

Theorem 4.1. *Let K be one of the following complexes.*

(1) *The curve complex $\mathcal{C}(S)$ with $S = S_{g,p}$ and $3g + p - 3 \geq 3$.*
(2) *The arc-and-curve complex $\mathcal{AC}(S)$.*

Then K satisfies a linear combinatorial isoperimetric inequality.

REMARK 4.2. For the arc-and-curve complex we only treat the case where $\xi(S) = 3g+p-3 \geq 2$ because the omitted cases are either straightforward or well known. It is a theorem of Harer [Har86] that the curve complex is homotopy equivalent to an infinite wedge of spheres of fixed dimension depending on g and p, therefore above we have omitted the curve complex of the five-times-punctured sphere $S_{0,5}$ and the two-times-punctured torus $S_{1,2}$ because they are not simply connected. However in these cases we are still able to state Theorem 4.10, which is key to proving Theorem 4.1.

Proof of Theorem 4.1(1). Let us discuss the strategy and outline of the proof and then defer the omitted details to Sections 4.1 and 4.2.

The curve graph $\mathcal{C}(S)^{(1)}$ is δ-hyperbolic [MM99] and therefore satisfies a linear coarse isoperimetric inequality [BH99, Proposition III.H.2.7]. In other words, given an arbitrary loop in $\mathcal{C}(S)^{(1)}$ we may decompose it into a net of linearly many loops of length at most 16δ i.e. linearly many short loops. To prove the theorem, it suffices to show that there exists an a priori bound on the number of triangles required to cap off a short loop. This is the content of the proof. It is not straightforward; the arc complex generally fails to have this property even for loops of length 4, see Theorem 3.1(1).

We show that after a uniformly bounded amount of homotoping a short loop (i.e. a homotopy across a bounded number of triangles) there is either a shortcut (see Definition 4.6) available that divides the loop into two/three smaller ones (then use induction on the length), or, the loop $(\gamma_i)_i$ is a *subpath* of a short, *tight loop* $(C_i)_i$ (this notation is nonstandard, by a *subpath* we mean that γ_i is a component of C_i for each i; tight loops are defined in Definition 4.7), see Lemma 4.9.

In Section 4.2 we show that there are only finitely many short, tight loops $c = (C_i)_i$ up to the action of $\mathrm{Mod}(S)$, see Theorem 4.10. Let $c' = (\gamma_i)_i$ be a subpath of c. Then there are only finitely many $\mathrm{Mod}(S)$-orbits of such $(\gamma_i)_i$ too. Since $\mathcal{C}(S)$ is simply connected, we have that each such $(\gamma_i)_i$ has some way of being capped off with a disc, but there are only finitely many orbits of such loop, and so an a priori bound on the number of triangles required exists and the proof is complete. $\square$

4.1 Homotoping and simplifying a loop via the tightening procedure

Masur and Minsky introduced the notion of a tight geodesic and showed that they exist between any pair of vertices in the curve complex (see [MM00, Lemma 4.5], they were originally called *tight sequences*). We refer to the idea in their proof as the tightening procedure. As we will see, the procedure is slightly more complicated for loops than it is for geodesics. Such a procedure for loops is new.

First we define what we mean by tightening a sequence, and then we give some remarks to make the definition clearer and explain the purpose.

DEFINITION 4.3. Suppose that $(C_i)_i$ is a sequence of curve systems (or multicurves) of S, such that C_i misses C_{i+1}, and such that each component γ_i of C_i cuts each component γ_{i+2} of C_{i+2}. We say that $(C_i)_i$ is *tight* at j if the curve system C_j consists precisely of the essential curves of the boundary of a closed regular neighbourhood of $C_{j-1} \cup C_{j+1}$. We will simply write $\partial(C_{j-1}, C_{j+1})$ for this curve system. If $(C_i)_i$ is not tight at j then we may *tighten* the sequence at j by replacing C_j by $\partial(C_{j-1}, C_{j+1})$.

REMARK 4.4.   (1) The reader may wonder why we are discussing curve systems and not simply curves. If one starts with a geodesic in the curve graph and then starts tightening the sequence then we might replace a curve by a curve system, and so eventually we may have to consider the process of tightening using curve systems also.

(2) Our assumption that C_i misses C_{i+1} simply means that the curve systems admit disjoint representatives on S. One should compare this to paths in the curve graph. On the other hand, if two isotopy classes do not miss then we say that they *cut*.

(3) We want each component γ_i of C_i to cut each component γ_{i+2} of C_{i+2}, otherwise there is an obvious shortcut (see Definition 4.6) from γ_i to γ_{i+2}. Such a shortcut can then be used to cut up our original loop into smaller ones.

(4) The curve system $\partial(C_{j-1}, C_{j+1})$ is defined by first taking representatives of C_{j-1} and C_{j+1} that intersect transversely and minimally. Then the set $C_{j-1} \cup C_{j+1}$ is well defined on S up to ambient isotopy (this is well known see for example [Web15, Lemma 2.2] for a proof). We can then discuss the closed regular neighbourhood $N = N(C_{j-1} \cup C_{j+1})$. Its boundary has at least one essential curve of S because if it did not then $S - N$ would be a collection of open discs and open peripheral annuli, but C_j is an essential curve system of S contained inside $S - N$, a contradiction. We write $\partial(C_{j-1}, C_{j+1})$ for the (non-empty) curve system obtained by taking all such essential curves of ∂N.

(5) We have that C_j misses $C_j' = \partial(C_{j-1}, C_{j+1})$, and so when we tighten the sequence $(C_i)_i$ at j we are in fact homotoping C_j across a simplex in $\mathcal{C}(S)$ to C_j'. Therefore if we take arbitrary components γ_i of C_i and γ_j' of C_j' then we are homotoping γ_j across two triangles to γ_j'. One of the triangles is spanned by γ_{j-1}, γ_j, and γ_j'; the other is γ_{j+1}, γ_j, and γ_j'. This fact is used in Lemma 4.9.

(6) The length of the sequence does not change. Furthermore, in an intuitive but not mathematical sense, $\partial(C_{j-1}, C_{j+1})$ is at most as complicated as C_j, and so tightening a sequence is a way of simplifying it.

Before we treat the harder and more general case let us engage with the loops of length 4 first to get a feel of where this proof is going. Note that any loop of length 3 already bounds a triangle in $\mathcal{C}(S)^{(2)}$, and any length less than this is straightforward.

LEMMA 4.5. *Let c be a loop of length 4 in $\mathcal{C}(S)^{(1)}$. Then c can be capped off with a disc of 2 or 4 triangles.*

Proof. We write $c = (c_i)_i$ where i ranges over the integers modulo 4. We may assume the vertices are distinct. If c_0 and c_2 miss then there is an edge between them and so c_0, c_1, c_2 form a triangle and so does c_2, c_3, $c_4 = c_0$, and we are done. Similarly we are done with two triangles if c_1 and c_3 are disjoint. Notice that we are dividing c into two loops of length 3 here via a shortcut.

Suppose instead that there are no such shortcuts. The plan instead is to consider tightening. Write $C = \partial(c_0, c_2)$. Then we observe that C misses c_1, see Remark 4.4(4) and (5), furthermore C also misses c_3.

We now give an awkward end to the proof for the purpose of indicating how it generalises to larger lengths. We tighten c at the index 1 to obtain a loop of curve systems (c_0, C, c_2, c_3). We observe that C misses c_3. This is a shortcut, so we should aim to divide the loop into two. Let γ be an arbitrary component of C. Then we may homotope c across two triangles in $\mathcal{C}(S)^{(2)}$ to obtain the loop $c' = (c_0, \gamma, c_2, c_3)$; the triangles mentioned are bounded by c_0, c_1 and γ, and, c_1, c_2 and γ. But γ misses c_3 and so we are in the case with a shortcut, and so c' can be capped off with 2 triangles, and in turn, c can be capped off with 4 triangles.   □

In the proof above we had to consider a loop of curve systems. From now on our short loop can be assumed to be a (periodic) sequence of curve systems $(C_i)_i$ where i ranges over the integers modulo n, and n is the length of the loop. Of course we are assuming that C_i misses C_{i+1}.

Let us make the following definition to tidy up the upcoming statement of Lemma 4.8.

DEFINITION 4.6. Let $c = (C_i)_i$ be a loop of curve systems of length $n \geq 5$. A *shortcut* for c is any one of the following

(1) a component γ_j of C_j and a component γ_{j+2} of C_{j+2} with $d_{\mathcal{C}(S)}(\gamma_j, \gamma_{j+2}) < 2$, or,

(2) a component γ_j of C_j and a component γ_{j+3} of C_{j+3} with $d_{\mathcal{C}(S)}(\gamma_j, \gamma_{j+3}) < 3$, and $n \geq 6$, or,

(3) a curve γ adjacent to γ_{j-1}, γ_{j+1}, and γ_{j+2}, some components of C_{j-1}, C_{j+1}, and C_{j+2} respectively.

We require $n \geq 6$ in the above second case for the following reason. If one has a loop of length 6 and finds a shortcut of the second type described above, then one can divide the loop into two smaller loops, and then we are done by induction—this doesn't quite work for $n = 5$. The third type is only introduced for handling the case $n = 5$ and enables us to divide the pentagon into two squares and one triangle. Of course the third type is a strong case of the second type when $n \geq 6$.

DEFINITION 4.7. Let $c = (C_i)_i$ be a loop of curve systems of length $n \geq 5$. We say that c is *tight* if

(1) there are no shortcuts for c, and,

(2) c is tight at all indices.

We will see in Section 4.2 that there are only finitely many tight loops of a given length up to the action of the mapping class group.

LEMMA 4.8 (*Tightening procedure for loops*). *Let* $c = (C_i)_i$ *be a loop of curve systems of length* $n \geq 5$. *Then there is a way of tightening* c *at most* n *times, during or at the end of which we either have that*

(1) *there is a shortcut for* c, *or,*
(2) c *is a tight loop of length* n.

Proof. We may assume that c has no shortcuts throughout.

The proof requires a discussion of the *subsurfaces filled by* C_{i-1} and C_{i+1}, written $F(C_{i-1}, C_{i+1})$, and their behaviour under the tightening procedure. This subsurface is defined to be the closure of: the union of a regular neighbourhood of $C_{i-1} \cup C_{i+1}$ union the complementary discs and once-punctured discs.

We tighten at index i, so we can assume $C_i = \partial(C_{i-1} \cup C_{i+1})$. Then we attempt to tighten at index $i+1$ by replacing C_{i+1} with $C'_{i+1} = \partial(C_i \cup C_{i+2})$. We wish to keep being tight at i. We may assume that there are no shortcuts and so every component of C'_{i+1} cuts every component of C_{i-1}. Therefore C'_{i+1} is contained in $F(C_{i-1}, C_{i+1})$ because it misses the boundary C_i. Hence $F(C_{i-1}, C'_{i+1}) \subset F(C_{i-1}, C_{i+1})$. Now if this is a strict inclusion then there exists γ that misses C_{i-1} and C'_{i+1} but is contained in $F(C_{i-1}, C_{i+1})$. Therefore γ cuts C_{i+1} because C_{i-1} and C_{i+1} fill $F(C_{i-1}, C_{i+1})$. If γ misses a component of C_{i+2} then there is a shortcut of the third type, namely via C_{i-1}, C'_{i+1}, and C_{i+2}, and we are done. So instead γ cuts every component of C_{i+2}. Then γ must be contained in $F(C_i, C_{i+2})$ because it misses the boundary C'_{i+1}. But then γ also misses C_{i+1}, a contradiction. Therefore if we assume there are no shortcuts throughout then after n tightenings we obtain a tight loop of length n. $\square$

We say that $(\gamma_i)_i$ is a *subpath* of $(C_i)_i$ if for each i we have that γ_i is some component of the curve system C_i. We now obtain

LEMMA 4.9. *Given any loop* $c = (\gamma_i)_i$ *of length* $n \geq 5$, *after a homotopy of* c *past at most* $2n$ *triangles, we either find a shortcut for* c *or it is a subpath of a tight loop of length* n.

Proof. We follow Lemma 4.8 starting with c. This lemma states that after at most n tightenings of the loop we either find a shortcut or the loop becomes a tight loop of length n. Suppose first that no shortcuts arise. Then for each curve system throughout the procedure, pick an arbitrary component. Then each loop of curve systems corresponds to a loop of curves, and whenever we tighten, then our loop of curves either stays the same or is homotoped past two triangles. Hence after homotoping c past at most $2n$ triangles we obtain the required subpath.

On the other hand if there is a shortcut during the tightening process, then we use those components of those curve systems that provide the shortcut. This then

provides a homotopy of c past at most $2n$ triangles ending up with a shortcut for the loop of curves. $\qquad\square$

4.2 Only finitely many kinds of short, tight loop.

In this section we adapt the argument in [Web15, Theorem 4.7] to show

Theorem 4.10. *For every surface S with $\xi(S) \geq 2$ and $n \geq 5$ there are only finitely many $\mathrm{Mod}(S)$-orbits of tight loop of length n.*

Proof. Let $\kappa_{C_i}(C)$ be the arc system (not counting parallel copies) of $S - C_i$ that is determined by the curve system C, defined when each component of C cuts C_i. The key observation that makes this proof work is, due to being tight at i (and having no shortcuts), that the arc system $\kappa_{C_{i-1}}(C_{i+1})$ determines C_i [Web15, Lemma 4.4].

We first claim that the geometric intersection number $i(C_{i-1}, C_{i+1})$ is bounded above in terms of S and n. So suppose that $n \geq 6$. Then C_{i-3} and C_i fill S. Up to the mapping class group, there are only finitely many possibilities for $\kappa_{C_i}(C_{i-3})$. Because this fills, we have only finitely many possibilities for $\kappa_{C_i}(C_{i-2})$. By considering $\kappa_{C_i}(C_{i+j})$ for $j = -3, -4, \ldots, 2$ (modulo n) we also have only finitely many possibilities for $\kappa_{C_i}(C_{i+2})$. But these determine C_{i-1} and C_{i+1}, and so our first claim follows.

Now suppose that $n = 5$. Then $C_{i-2} \cup C_{i+2}$ is a curve system and $\kappa_{C_i}(C_{i-2} \cup C_{i+2})$ fills $S - C_i$, because there is no shortcut. Then after a mapping class we can assume that the pair $\kappa_{C_i}(C_{i-2})$ and $\kappa_{C_i}(C_{i+2})$ is one of finitely many possibilities. But this then determines the pair C_{i-1} and C_{i+1}, so we are done as before. In fact, this settles the theorem in the case $n = 5$ by bounding the geometric intersection number of the collection of curve systems. So we assume $n \geq 6$ for the remainder of the proof.

Our second claim is that $i(C_{i-1}, C_{i+2})$ is bounded above in terms of S and n. Start off with $\kappa_{C_i}(C_{i-3})$, which intersects C_{i-1} a uniformly bounded number of times by the first claim above. Then because $\kappa_{C_i}(C_{i-3})$ fills $S - C_i$, we can use an inductive argument with $j = i - 3, i - 4, \ldots, i + 2$ (modulo n), to see that $i(\kappa_{C_i}(C_{i+2}), C_{i-1})$ is bounded above in terms of S and n. By our first claim C_{i+2} determines only a uniformly bounded number of parallel copies of $\kappa_{C_i}(C_{i+2})$, and so our second claim follows.

Finally, there are only finitely many $\mathrm{Mod}(S)$-orbits of the pair C_0 and C_3 by our second claim above. So fix C_0 and C_3 to be one such pair. But then we obtain two different paths of curve systems between C_0 to C_3. These are both so-called *tight filling multipaths* of length at most n, see [Web15, Section 3]. But there are only finitely many such paths connecting C_0 to C_3 in terms of S and n, see [Web15, Theorem 4.7] (or [BMM16, Appendix A] for an exposition). Briefly speaking the idea of that proof is to consider $\kappa_{C_0}(C_i)$. Whenever this fills then there are only finitely many possibilities for $\kappa_{C_0}(C_{i+1})$ (and $\kappa_{C_0}(C_{i-1})$). However $\kappa_{C_0}(C_2)$ (and $\kappa_{C_0}(C_{-2})$) do not fill, but in this case C_1 (and C_{-1}) are determined by those arc systems. We know precisely what C_3 is, so we can also deduce the finitely many possibilities for C_2 and C_4, and so on. $\qquad\square$

4.3 The arc-and-curve complexes. Now we show that Theorem 4.1(2) follows from Theorem 4.1(1), unless $\xi(S) = 2$, in which case it follows from Theorem 4.10.

Proof of Theorem 4.1(2). The arc-and-curve complex is quasi-isometric to $\mathcal{C}(S)$, and so it is also hyperbolic, therefore once again it suffices to show that for any loop of vertices $c = (v_i)_i$ of bounded length, there is a bounded number of triangles required for a disc to cap off c to deduce the theorem.

Let c have length n. Our first step is to homotope c into $\mathcal{C}(S)$, but across at most $3n$ triangles, and to a loop of length at most $2n$. This argument is well known. To see this define $c_i \in \mathcal{C}(S)$ for each vertex v_i of c. If v_i is a curve we set $c_i = v_i$. If v_i is an arc then we set c_i to be one of the peripheral curves of $S - v_i$ which is essential in S. Thus v_i and c_i are either equal or adjacent. Moreover for each i we can pick a curve c_i' which is peripheral in $S - v_i - v_{i+1}$ but which is essential in S. Such a curve c_i' is adjacent (or equal) to v_i, c_i, v_{i+1}, and c_{i+1}. Therefore we can push c past at most $3n$ triangles into $\mathcal{C}(S)$ to a loop of length at most $2n$.

Now we are done by Theorem 4.1(1), unless S is $S_{0,5}$ or $S_{1,2}$, in which case $\mathcal{C}(S)$ is not simply connected. So suppose that $\xi(S) = 2$ for the remainder of the proof. We have showed that we can assume that our loop $c = (v_i)_i$ has length at most $2n$ and consists only of curves v_i. If c has a shortcut (see Definition 4.6) then we can decompose c into two/three shorter loops in $\mathcal{C}(S)$ and be finished by induction on the length of c. If on the other hand c has no shortcuts then it is automatically a tight loop, because whenever $d_{\mathcal{C}(S)}(v_{i-1}, v_{i+1}) = 2$ then v_i is uniquely determined, and hence tight at i. By Theorem 4.10 there are only finitely many $\mathrm{Mod}(S)$-orbits of tight loops c of length at most $2n$. Now $\mathcal{AC}(S)$ is contractible [Hat91] and therefore simply connected, so each short, tight loop can be capped off by some number of triangles in $\mathcal{AC}(S)$, but there are only finitely many orbits of such loops, so there is an upper bound on the number of triangles required. $\square$

Acknowledgments

References

[AFW15] MATTHIAS ASCHENBRENNER, STEFAN FRIEDL, and HENRY WILTON. *3-manifold groups.* European Mathematical Society (EMS), Zürich (2015).

[Ago13] IAN AGOL. The virtual Haken conjecture. *Doc. Math.*, 18 (2013), 1045–1087 **With an appendix by Agol, Daniel Groves, and Jason Manning**

[AS12] JAVIER ARAMAYONA and JUAN SOUTO. Homomorphisms between mapping class groups. *Geom. Topol.*, (4)16 (2012), 2285–2341

[BF14a] MLADEN BESTVINA and MARK FEIGHN. Hyperbolicity of the complex of free factors. *Adv. Math.*, 256 (2014), 104–155

[BF14b] MLADEN BESTVINA and MARK FEIGHN. Subfactor projections. *J. Topol.*, (3)7 (2014), 771–804

[BH99] MARTIN R. BRIDSON and ANDRÉ HAEFLIGER. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften* . Springer, Berlin (1999).

[BI18] BRIAN H. BOWDITCH and FRANCESCA IEZZI. Projections of the sphere graph to the arc graph of a surface. *J. Topol. Anal.*, (2)10 (2018), 245–261

[BMM16] JOAN BIRMAN, DAN MARGALIT, and WILLIAM MENASCO. Efficient geodesics and an effective algorithm for distance in the complex of curves. *Math. Ann.*, (3-4)366 (2016), 1253–1279

[Bow08] BRIAN H. BOWDITCH. Tight geodesics in the curve complex. *Invent. Math.*, (2)171 (2008), 281–300

[Bri91] MARTIN R. BRIDSON. Geodesics and curvature in metric simplicial complexes. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 373–463. World Sci. Publ., River Edge, NJ (1991).

[Bri10] MARTIN R. BRIDSON. Semisimple actions of mapping class groups on CAT(0) spaces. In *Geometry of Riemann surfaces*, volume 368 of *London Math. Soc. Lecture Note Ser.*, pages 1–14. Cambridge Univ. Press, Cambridge (2010).

[Bro03] JEFFREY F. BROCK. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *J. Amer. Math. Soc.*, (3)16 (2003), 495–535

[DHS] F. DAHMANI, M. HAGEN, and A. SISTO. Dehn filling dehn twists. Preprint.

[DW03] GEORGIOS DASKALOPOULOS and RICHARD WENTWORTH. Classification of Weil-Petersson isometries. *Amer. J. Math.*, (4)125 (2003), 941–975

[For] MAXWELL FORLINI. Splittings of free groups from arcs and curves. Preprint.

[Gal79]   David Gale. The game of Hex and the Brouwer fixed-point theorem. *Amer. Math. Monthly*, (10)86 (1979), 818–827

[Har85]   John L. Harer. Stability of the homology of the mapping class groups of orientable surfaces. *Ann. of Math.*, (2)121 (1985), 215–249

[Har86]   John L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.*, (1)84 (1986), 157–176

[Hat91]   Allen Hatcher. On triangulations of surfaces. *Topology Appl.*, (2)40 (1991), 189–194

[Hat95]   Allen Hatcher. Homological stability for automorphism groups of free groups. *Comment. Math. Helv.*, (1)70 (1995), 39–62

[HH15]   Ursula Hamenstädt and Sebastian Hensel. Spheres and projections for Out(F_n). *J. Topol.*, (1)8 (2015), 65–92

[HH17]   Arnaud Hilion and Camille Horbez. The hyperbolicity of the sphere complex via surgery paths. *J. Reine Angew. Math.*, 730 (2017), 135–161

[HM13]   Michael Handel and Lee Mosher. The free splitting complex of a free group, I: hyperbolicity. *Geom. Topol.*, (3)17 (2013), 1581–1672

[HOP14]   Sebastian Hensel, Damian Osajda, and Piotr Przytycki. Realisation and dismantlability. *Geom. Topol.*, (4)18 (2014), 2079–2126

[IK]   G. Islambouli and M. Klug. Representing smooth 4-manifolds as loops in the pants complex. Preprint.

[Ker83]   Steven P. Kerckhoff. The Nielsen realization problem. *Ann. of Math.* , (2)117 (1983), 235–265

[KL96]   Michael Kapovich and Bernhard Leeb. Actions of discrete groups on nonpositively curved spaces. *Math. Ann.*, (2)306 (1996), 341–352

[McC91]   Darryl McCullough. Virtually geometrically finite mapping class groups of 3-manifolds. *J. Differential Geom.*, (1)33 (1991), 1–65

[MM99]   Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, (1)138 (1999), 103–149

[MM00]   H. A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, (4)10 (2000), 902–974

[MS13]   Howard Masur and Saul Schleimer. The geometry of the disk complex. *J. Amer. Math. Soc.*, (1)26 (2013), 1–62

[Tay14]   Samuel J. Taylor. A note on subfactor projections. *Algebr. Geom. Topol.*, (2)14 (2014), 805–821

[Tro86]   A. J. Tromba. On a natural algebraic affine connection on the space of almost complex structures and the curvature of Teichmüller space with respect to its Weil-Petersson metric. *Manuscripta Math.*, (4)56 (1986), 475–497

[Web15]   Richard C. H. Webb. Combinatorics of tight geodesics and stable lengths. *Trans. Amer. Math. Soc.*, (10)367 (2015), 7323–7342

[Wol86]   Scott A. Wolpert. Chern forms and the Riemann tensor for the moduli space of curves. *Invent. Math.*, (1)85 (1986), 119–145

[Wol87]   Scott A. Wolpert. Geodesic length functions and the Nielsen problem. *J. Differential Geom.*, (2)25 (1987), 275–296

 Department of Mathematics, The University of Manchester, Oxford Road, Manchester M13 9PL, UK.

richard.webb@manchester.ac.uk