Communications in
Analysis and Geometry
Volume 30, Number 4, 843–868, 2022

Morse functions to graphs and topological complexity for hyperbolic 3-manifolds

Diane Hoffoss and Joseph Maher

Scharlemann and Thompson define the width of a 3-manifold M as a notion of complexity based on the topology of M. Their original definition had the property that the adjacency relation on handles gave a linear order on handles, but here we consider a more general definition due to Saito, Scharlemann and Schultens, in which the adjacency relation on handles may give an arbitrary graph. We show that for closed hyperbolic 3-manifolds, this is linearly related to a notion of metric complexity, based on the areas of level sets of Morse functions to graphs, which we call Gromov area.

1. Introduction

Mostow rigidity implies that for hyperbolic 3-manifolds, the hyperbolic metric is a topological invariant, so one might hope that the topological and metric complexities are related. We shall show that this is indeed the case for certain definitions of topological and metric complexity. We first describe the notions of complexity we shall use, and then give a brief outline of the arguments used to relate topological and metric complexity in the subsequent sections. In [HM16] we considered the linear version of these invariants, while in this paper we consider the more general case of invariants constructed from maps to graphs. It will be convenient to work with the collection of hyperbolic 3-manifolds which are covers of closed hyperbolic 3-manifolds, though not necessarily of finite volume. We remark that we do not consider finite volume manifolds with cusps, as in this case, the surfaces separating the 3-dimensional regions in the topological decomposition we construct from the metric might have essential intersection with the cusps. In other words, if the cusps are truncated to form a hyperbolic manifold with torus boundary components, then the dividing surfaces may be surfaces with essential boundary components on the boundary tori. However, the currently available versions of the topological decomposition results we use, due to

Scharlemann-Thompson [ST94] and Saito-Scharlemann-Schultens [SSS16], assume that the dividing surfaces are closed.

This paper is not entirely self-contained, and relies on the results of [HM16], however we review the main definitions and results from [HM16] for the convenience of the reader.

1.1. Metric complexity

In [HM16], we considered the following definition of metric complexity. Let M be a closed Riemannian 3-manifold, and let $f\colon M \to \mathbb{R}$ be a Morse function, i.e. f is a smooth function, all critical points are non-degenerate, and distinct critical points have distinct images in $\mathbb{R}$. We define the *area* of f to the maximum area of any level set $F_t = f^{-1}(t)$ over all points $t \in \mathbb{R}$. We define the *Morse area* of M to be the infimum of the area of all Morse functions $f\colon M \to \mathbb{R}$.

More generally, we may consider maps $f\colon M \to X$, where X is a trivalent graph. Recall that for a Morse function $f\colon M \to \mathbb{R}$ there are singularities of index $0, 1, 2$ and 3. The singularities of index 0 and 3 are known as birth or death singularities respectively, and the level set foliation near the singular point in M is locally homeomorphic to the level sets of the function $x^2 + y^2 + z^2$ close to the origin in $\mathbb{R}^3$. For singularities of index 1 and 2, the level sets near the singular point in M are locally homeomorphic to the level sets of the function $x^2 + y^2 - z^2$ close to the origin in $\mathbb{R}^3$.

In the case of index 1 or 2, there is a map from a small open ball containing the singular point to the leaf space of the level set foliation. As the singular leaf divides a small ball about the singular point into three connected components, the leaf space is a trivalent graph with a single vertex and three edges, and we call such a regular neighbourhood of the critical point a *trivalent singularity.* If X is a trivalent tree, we say a map $f\colon M \to X$ is *Morse* if it is a Morse function on the interior of each edge of X, and at each trivalent vertex v of X the pre-image under f is locally homeomorphic to a trivalent singularity. We say that the area of f is the maximum area of F_t, as t runs over all points $t \in X$. The *Gromov area* of M is the infimum of the area of $f\colon M \to X$ over all trivalent graphs X, and all Morse functions $f\colon M \to X$.

This definition of metric complexity is a variant of Uryson width, studied by Gromov in [Gro88], though we consider the area of the level sets instead of the diameter. Alternatively, one may consider it to be a variant of the definition of the waist of a manifold, but we prefer to call it area, as the dimension of our spaces is fixed, and the fibers have dimension two.

1.2. Topological complexity

We now describe the notions of topological complexity we shall consider. A *handlebody* is a compact 3-manifold with boundary, homeomorphic to a regular neighborhood of a graph in $\mathbb{R}^3$. Up to homeomorphism, a handlebody is determined by the genus g of its boundary surface. Every 3-manifold M has a *Heegaard splitting*, which is a decomposition of the manifold into two handlebodies. This immediately gives a notion of complexity for a 3-manifold, called the *Heegaard genus*, which is the smallest genus of any Heegaard splitting of the 3-manifold.

There is a refinement of this, due to Scharlemann and Thompson [ST94], which we now describe. A *compression body* C is a compact 3-manifold with boundary, constructed by gluing some number of 2-handles to one side of a compact (but not necessarily connected) surface cross interval and capping off any resulting 2-sphere components with 3-balls. The side of the surface cross interval with no attached 2-handles is called the *top boundary* of the compression body and denoted by $\partial_+ C$, and any other boundary components are called the *lower boundary* of the compression body, and their union is denoted by $\partial_- C$. A *linear generalized Heegaard splitting*,[1] which we shall abbreviate to *linear splitting*, is a decomposition of a 3-manifold M into a linearly ordered sequence of (not necessarily connected) compression bodies $C_1, \ldots C_{2n}$, such that the top boundary of an odd numbered compression body C_{2i+1} is equal to the top boundary of the subsequent compression body C_{2i+2}, and the lower boundary of C_{2i+1} is equal to the lower boundary of the previous compression body C_{2i}. Let H_i be the sequence of surfaces consisting of the top boundaries of the compression bodies C_{2i-1} and C_{2i}. The complexity $c(H_i)$ of the surface H_i is the sum of the genera of each connected component, and the complexity of the linear splitting is the collection of integers $\{c(H_i)\}$, arranged in decreasing order. We order these complexities with the lexicographic ordering. The *width* of the linear splitting is the maximum value (i.e. the first value) of $c(H_i)$ in the collection $\{c(H_i)\}$. The *linear width* of a 3-manifold M is the minimum width over all possible linear generalized Heegaard splittings. As a Heegaard splitting is a special case of a linear splitting, the Heegaard genus of M is an upper bound for the linear

[1] We warn the reader that these are often referred to as generalized Heegaard splittings in the literature; however we wish to distinguish them from a more general notion described subsequently, which is also occasionally referred to in the literature as a generalized Heegaard splitting.

width of M. A linear splitting which gives the minimum complexity of all possible linear splittings is called the *thin position* linear splitting.

There is a further refinement of this, described in Saito, Scharlemann and Schultens [SSS16]. A *graph generalized Heegaard splitting*, which we shall abbreviate to *graph splitting*, and is called a *fork complex* in [SSS16], is a decomposition of a compact 3-manifold M into compression bodies $\{C_i\}$, such that for each compression body C_i, there is a compression body C_j such that the top boundary of C_i is equal to the top boundary of C_j. Furthermore, for each component of the lower boundary of C_i, there is a compression body C_k, such that that component of the lower boundary of C_i is equal to a component of the lower boundary of C_k. We emphasize that different components of the lower boundary of C_i may be attached to lower boundary components of different compression bodies. Let $\{H_i\}$ be the collection of top boundary surfaces. The complexity of the graph splitting is the collection of integers $\{c(H_i)\}$, arranged in decreasing order. Again, we put the lexicographic ordering on these complexities. A graph splitting which realizes the minimum complexity is called a *thin position graph splitting.* The *width* of the graph splitting is the maximum integer (i.e the first integer) that appears in the complexity. The *graph width* of a 3-manifold M is the minimum width over all possible graph splittings of M. As a linear splitting is a special case of a graph splitting, the linear width of M is an upper bound for the graph width of M. The graph corresponding to the graph splitting is the graph whose vertices are compression bodies, with edges connecting pairs of compression bodies with common boundary components.

1.3. Results

In order to bound metric complexity in terms of topological complexity we will use the following result:

Theorem 1.1. *Let M be a Riemannian 3-manifold with a strongly irreducible Heegaard splitting. Then the Heegaard surface is isotopic to a minimal surface, or to the boundary of a regular neighborhood of a non-orientable minimal surface with a small tube attached vertically in the I-bundle structure.*

This was announced by Pitts and Rubinstein [PR86] (see also Rubinstein [Rub05]), and has now been proved by Ketover and Liokumovich [KL17], building on work of Colding and De Lellis [CDL03], De Lellis and Pellandrini [DLP10], and Ketover [Ket13].

In [HM16] we showed:

Theorem 1.2. *There is a constant $K > 0$, such that for any closed hyperbolic 3-manifold,*

$$K(\textit{linear width}(M)) \leqslant \textit{Morse area}(M) \leqslant 4\pi(\textit{linear width}(M)). \tag{1}$$

In this paper we show:

Theorem 1.3. *There is a constant $K > 0$, such that for any closed hyperbolic 3-manifold,*

$$K(\textit{graph width}(M)) \leqslant \textit{Gromov area}(M) \leqslant 4\pi(\textit{graph width}(M)). \tag{2}$$

We also expect there to be upper and lower bounds on topological complexity in terms of Uryson width, i.e. using diameter instead of area, but we do not expect them to be linear.

1.4. Related work in 3-manifolds

It may be of interest to compare our results with recent work of Brock, Minsky, Namazi and Souto [BMNS16] on manifolds with bounded combinatorics. Let $C_1, \dots C_n$ be a finite collection of homeomorphism types of compact 3-manifolds with marked boundary, which we shall refer to as *model pieces*, and fix a metric on each one. A 3-manifold M is said to have *bounded combinatorics* if it is a union of (possibly infinitely many) model pieces glued together by homeomorphisms along their boundaries, with certain restrictions on the gluing maps, which we do not describe in detail here. In particular, a manifold with bounded combinatorics is a manifold of bounded topological width. They show that such a manifold M is hyperbolic, with a lower bound on the injectivity radius, and the hyperbolic metric is K-bilipshitz homeomorphic to the induced metric on M arising from the metrics on the model pieces. A choice of foliation with compact leaves, containing the boundary leaves, on each model piece then shows that the metric complexity is linearly related to the topological complexity for this class of manifolds, where the linear constants depend on the collection of model pieces.

Note that in our context, a bound on the topological width of the manifold implies that the manifold is a union of compression bodies of bounded genus, and there are finitely many of these up to homeomorphism. Their result assumes restrictions on the gluing maps, but then shows the resulting manifold is hyperbolic, but the bilipshitz constant K depends on the width of M, i.e the genus of the compression bodies. We assume that the manifold M is closed and hyperbolic, and make no restriction on the gluing maps between the compression bodies, but we show that the linear constants relating topological and metric complexities are independent of the genus of the compression bodies.

1.5. Outline

In [HM16] we considered the linear case, in which the range of the Morse function $f\colon M \to \mathbb{R}$ is $\mathbb{R}$. Such a Morse function has the property that for each $t \in \mathbb{R}$, the pre-image $f^{-1}(t)$ is compact and separating. For the case in which the range of the Morse function $f\colon M \to X$ is a graph, one may consider the lifted Morse function $\widetilde{f}\colon \widetilde{M} \to \widetilde{X}$, where $\widetilde{X}$ is the universal cover of X, and $\widetilde{M}$ is the corresponding cover of M. This lifted Morse function has the property that for each $t \in \widetilde{X}$, each pre-image $\widetilde{f}^{-1}(x)$ is compact and separating, and so many of the arguments from [HM16] go through directly in this case. In particular, we construct polyhedral approximations to the level sets of $\widetilde{f}$, and show that they have bounded topological complexity, as we now describe.

A choice of Margulis constant μ determines a thick-thin decomposition for M, in which the thin part is a disjoint union of Margulis tubes. We also choose a Voronoi decomposition determined by a maximal ϵ-separated collection of points in M. This implies that every Voronoi cell has diameter at most ϵ, and, given μ, we may choose ϵ small enough such that every Voronoi cell that intersects the thick part contains an embedded ball of radius $\epsilon/2$. The thick-thin decomposition of M, and the Voronoi decomposition of M, lift to thick-thin decompositions and Voronoi decompositions of the cover $\widetilde{M}$. We give the details of this construction in Sections 2.1, 2.2 and 2.3.

A separating surface F in $\widetilde{M}$ determines a partition of the Voronoi cells, depending on which side of the surface the majority of the volume of the (metric) ball of radius $\epsilon/2$ inside the Voronoi cell lies. We will call the boundary between these two sets of Voronoi cells a *polyhedral surface* S, which is a union of faces of Voronoi cells, and we can think of this as a combinatorial approximation to the original surface F.

A key observation from [HM16] is that the number of faces of the polyhedral surface in the thick part is bounded by the area of F. This is because in the thick part of M, the metric ball of radius $\epsilon/2$ in each Voronoi cell is embedded, so moving the ball along a geodesic connecting the centers of the two Voronoi cells produces at some point a metric ball whose volume is divided exactly in two, giving a lower bound to the area of F near that point. There are bounds on the number of vertices and edges of any Voronoi cell in terms of ϵ, so a bound on the number of faces of S in the thick part gives a bound on the Euler characteristic of S. We are unable to control the number of faces in the thin part, so we cap off the part of S in the thick part with surfaces of bounded Euler characteristic contained in the thin part. This produces surfaces of bounded genus, which we call *capped surfaces*.

In this way, the lift of a Morse function $\widetilde{f}\colon \widetilde{M} \to \widetilde{X}$ gives rise to a collection of polyhedral surfaces in $\widetilde{M}$ of bounded genus. These surfaces are constant except at finitely many points of $\widetilde{X}$, which we call *cell splitters*, where a level set divides the ball contained in a Voronoi cell exactly in half. We give the details of the construction of the capped surfaces and the properties of the cell splitters in Sections 2.4 and 2.5.

The key step, in Section 2.6, is to show that we may construct these surfaces equivariantly in $\widetilde{M}$, so they project down to embedded surfaces in M, with the same bounds on their topological complexity.

Finally, in Section 2.7, by considering the local configuration near a cell splitter, we show that the regions between the capped surfaces may be constructed using a number of handles bounded in terms of the area of the level sets $\widetilde{f}^{-1}(t)$, and so this bounds topological complexity of the decomposition of M given by the capped surfaces in terms of metric complexity of M.

The bound in the other direction is a direct consequence of the bound from [HM16], though we review the argument in the Section 3 for the convenience of the reader.

1.6. Acknowledgements

The authors would like to thank the referee for many helpful comments and suggestions. The authors would also like to thank Dick Canary, David Futer, David Gabai, Joel Hass, Daniel Ketover, Sadayoshi Kojima, Yair Minksy, Yo'av Rieck and Dylan Thurston for helpful conversations, and the Tokyo Institute of Technology for its hospitality and support. The second author was supported by the Simons Foundation and PSC-CUNY. This material is based upon work supported by the National Science Foundation under Grant No. DMS-1440140 while the second author was in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester.

2. Gromov area bounds graph width

In this section we show that we can bound the topological complexity of the manifold in terms of its metric complexity, i.e. we show that graph width is bounded in terms of Gromov area.

Theorem 2.1. *There is a constant K, such that for any closed hyperbolic 3-manifold M,*

$$\text{graph width}(M) \leqslant K(\text{Gromov area}(M)).$$

Let $f \colon M \to X$ be a Morse function onto a graph X, such that the Gromov area of f is arbitrarily close to the Gromov area of M. Any metric graph is arbitrarily close to a trivalent metric graph, so we may assume the graph is trivalent. We now show that we may assume the level sets of f are connected.

Proposition 2.2. *Let M be a Riemannian manifold, and let $f \colon M \to X$ be a Morse function onto a trivalent graph X. Then there is a trivalent graph X', and a Morse function $f' \colon M \to X'$ with connected level sets, with Gromov area$(f') \leqslant$ Gromov area(f).*

Proof. The level sets of the function f give a singular foliation of M with compact leaves, which we shall call the *level set foliation*, and the leaves of this foliation are precisely the connected components of the pre-images of points in X. Consider the leaf space L of the level set foliation, i.e. the space obtained from M by identifying points in the same leaf. As all leaves are compact, the leaf space is Hausdorff. The leaf space is a trivalent graph, with vertices corresponding to vertex singularities, and the maximum area of the pre-images of the quotient map is less than or equal to the maximum area of the pre-images of f. Therefore, we may choose f' to be the leaf space quotient map $f' \colon M \to L$, which is a Morse function onto a trivalent graph, and has connected level sets, with the property that the area of the level sets of f' is bounded by the area of the level sets of f. □

In particular, this means that the vertices of X are precisely the critical values of the Morse function f in which a connected level set splits into two connected components, or the reverse of this, in which two connected components are joined together.

2.1. Morse functions to trees

We would like to work in the cover $\widetilde{M}$ of M corresponding to the universal cover $\widetilde{X}$ of the graph X, which will have the key advantage that all pre-image surfaces are separating in $\widetilde{M}$.

Let $p\colon \widetilde{M} \to M$ be the cover of M corresponding to the kernel of the induced map $f_*\colon \pi_1 M \to \pi_1 X$, and let $c : \widetilde{X} \to X$ be the universal cover of X, so $\widetilde{X}$ is a tree. Then the map $f \circ p\colon \widetilde{M} \to X$ lifts to a map $h = \widetilde{f \circ p} : \widetilde{M} \to \widetilde{X}$. Since each leaf F_t in M maps to a single point in X, the fundamental group of each leaf is contained in $\ker(f)$. Therefore, each leaf in M lifts to a leaf in $\widetilde{M}$, and as the cover is regular, the pre-image of a point $t \in \widetilde{X}$ is a disjoint union of homeomorphic copies of $F_{c(t)}$. In particular, the area bound for the leaves F_t in M is also an area bound for the leaves $H_t = h^{-1}(t)$ in $\widetilde{M}$.

$$\begin{array}{ccc} \widetilde{M} & \xrightarrow{\;h = \widetilde{f \circ p}\;} & \widetilde{X} \\ \Big\downarrow p & & \Big\downarrow c \\ M & \xrightarrow{\;\;f\;\;} & X \end{array}$$

As $\widetilde{X}$ is a tree, every point is separating, and so every pre-image surface $H_t = h^{-1}(t)$ is also separating.

2.2. Voronoi cells

We will approximate the level sets of f by surfaces consisting of faces of Voronoi cells. We now describe in detail the Voronoi cell decompositions we shall use, and their properties. The definitions in this section are taken verbatim from [HM16], but we include them in this section for the convenience of the reader.

A *polygon* in $\mathbb{H}^3$ is a bounded subset of a hyperbolic plane whose boundary consists of a finite number of geodesic segments. A *polyhedron* in $\mathbb{H}^3$ is a convex topological 3-ball in $\mathbb{H}^3$ whose boundary consists of a finite collection of polygons. A *polyhedral cell decomposition* of $\mathbb{H}^3$ is a cell decomposition in which every 3-cell is a polyhedron, each 2-cell is a polygon, and the edges are all geodesic segments. We say a cell decomposition of a hyperbolic manifold M is *polyhedral* if its pre-image in the universal cover $\mathbb{H}^3$ is polyhedral.

Let $X = \{x_i\}$ be a discrete collection of points in 3-dimensional hyperbolic space $\mathbb{H}^3$. The Voronoi cell V_i determined by $x_i \in X$ consists of all points of M which are closer to x_i than any other $x_j \in X$, i.e.

$$V_i = \{x \in \mathbb{H}^3 \mid d(x, x_i) \leqslant d(x, x_j) \text{ for all } x_j \in \widetilde{X}\}.$$

We shall call x_i the *center* of the Voronoi cell V_i, and we shall write $\mathcal{V} = \{V_i\}$ for the collection of Voronoi cells determined by X. Voronoi cells are convex sets in $\mathbb{H}^3$, and hence topological balls. By general position, we may assume that all edges of the Voronoi decomposition are contained in exactly three distinct faces, the collection of vertices is a discrete set, and there are no points which lie in more than four distinct Voronoi cells. We shall call such a Voronoi decomposition a *regular* Voronoi decomposition, and it is a polyhedral decomposition of $\mathbb{H}^3$. As each edge is 3-valent, and each vertex is 4-valent, this implies that the dual cell structure is a simplicial triangulation of $\mathbb{H}^3$, which we shall refer to as the *dual triangulation*. The dual triangulation may be realised in $\mathbb{H}^3$ by choosing the vertices to be the centers x_i of the Voronoi cells and the edges to be geodesic segments connecting the vertices, and we shall always assume that we have done this.

Given a collection of points $X = \{x_i\}$ in a hyperbolic 3-manifold M, let $\widetilde{X}$ be the pre-image of X in the universal cover of M, which is isometric to $\mathbb{H}^3$. As $\widetilde{X}$ is equivariant, the corresponding Voronoi cell decomposition $\mathcal{V}$ of $\mathbb{H}^3$ is also equivariant, and furthermore, gives rise to an equivariant Voronoi decomposition of any cover $\widetilde{M}$ of M. The distance condition implies that the interior of each Voronoi cell V is mapped down homeomorphically by the covering projection, though the covering projection may identify faces, edges or vertices of V_i under projection into M. By abuse of notation, we shall refer to the resulting polyhedral decomposition of M as the Voronoi decomposition $\mathcal{V}$ of M. By general position, we may assume that $\mathcal{V}$ is regular. The dual triangulation is also equivariant, and projects down to a triangulation of M, which we will also refer to as the dual triangulation, though this triangulation may no longer be simplicial.

We shall write $B(x, r)$ for the closed metric ball of radius r about x in M, i.e.

$$B(x, r) = \{y \in M \mid d(x, y) \leqslant r\}.$$

We shall write $\text{inj}_M(x)$ for the injectivity radius of M at x, i.e. the radius of the largest embedded ball in M centered at x. Then the injectivity radius of M, denoted $\text{inj}(M)$, is defined to be

$$\text{inj}(M) = \inf_{x \in M} \text{inj}_M(x).$$

We say a collection $\{x_i\}$ of points in M is *ϵ-separated* if the distance between any pair of points is at least ϵ, i.e. $d(x_i, x_j) \geqslant \epsilon$, for all $i \neq j$. Let $\{x_i\}$ be a maximal collection of ϵ-separated points in M, and let $\mathcal{V}$ be the corresponding Voronoi cell division of M. Since the collection $\{x_i\}$ is maximal, each Voronoi cell is contained in a metric ball of radius ϵ about its center. Furthermore, if the injectivity radius at the center x_i is at least 2ϵ, then as the points x_i are distance at least ϵ apart, each Voronoi cell contains a topological ball of radius $\epsilon/2$ about its center, i.e.

$$B(x_i, \epsilon/2) \subset V_i \subset B(x_i, \epsilon).$$

Definition 2.3. Let M be a closed hyperbolic 3-manifold. We say a Voronoi decomposition $\mathcal{V}$ is *ϵ-regular*, if it is regular, and it arises from a maximal collection of ϵ-separated points. For any cover $\widetilde{M}$ of M, we say that a Voronoi decomposition of $\widetilde{M}$ is *equivariant ϵ-regular*, if it is the lift of an ϵ-regular Voronoi decomposition of M.

A *normal surface* in a triangulated 3-manifold is a surface that intersects each tetrahedron in a union of normal triangles and quadrilaterals. One useful property of ϵ-regular Voronoi decompositions is that the boundary of any union of Voronoi cells is an embedded surface, in fact an embedded normal surface in the dual triangulation.

Proposition 2.4. [HM16, Proposition 2.2] *Let M be a closed hyperbolic manifold, and let $\mathcal{V}$ be an ϵ-regular Voronoi decomposition. Let P be a union of Voronoi cells in $\mathcal{V}$, and let S be the boundary of P. Then S is an embedded surface in M.*

Definition 2.5. We shall say a Voronoi cell V_i with center x_i is an *ϵ-deep* Voronoi cell if the injectivity radius at x_i is at least 4ϵ, i.e. $\mathrm{inj}_M(x_i) \geqslant 4\epsilon$. We shall also call centers, faces, edges and vertices of ϵ-deep Voronoi cells ϵ-deep.

In particular this implies that the metric ball $B(x_i, 3\epsilon)$ inside an ϵ-deep Voronoi cell is a topological ball. In the next section we will choose a fixed ϵ independent of the manifold M, and we will just say *deep* instead of ϵ-deep. We shall write $\mathcal{W}$ for the subset of $\mathcal{V}$ consisting of deep Voronoi cells.

Finally, we recall that there are bounds, which only depend on ϵ, on the number of vertices, edges and faces of a deep Voronoi cell.

Proposition 2.6. [HM16, Proposition 2.3] *Let M be a closed hyperbolic 3-manifold with an ϵ-regular Voronoi decomposition $\mathcal{V}$, and let $\mathcal{W}$ be the collection of deep Voronoi cells. Then there is a number J, which only depends on ϵ, such that each deep Voronoi cell $W_i \in \mathcal{W}$ has at most J faces, edges and vertices.*

2.3. Margulis tubes

We will use the Margulis Lemma and the *thick-thin* decomposition for finite volume hyperbolic 3-manifolds, and we now review these results.

Given a number $\mu > 0$, let $X_\mu = M_{[\mu,\infty)}$ be the *thick part* of M, i.e. the union of all points x of M with $\mathrm{inj}_M(x) \geqslant \mu$. We shall refer to the closure of the complement of the thick part as the *thin part* and denote it by $T_\mu = \overline{M \setminus X_\mu}$.

The Margulis Lemma states that there is a constant $\mu_0 > 0$, such that for any closed hyperbolic 3-manifold, the thin part is a disjoint union of solid tori, called *Margulis tubes*, and each of these solid tori is a regular metric neighborhood of an embedded closed geodesic of length less than μ_0. We shall call a number μ_0 for which this result holds a *Margulis constant* for $\mathbb{H}^3$. If μ_0 is a Margulis constant for $\mathbb{H}^3$, then so is μ for any $0 < \mu < \mu_0$, and furthermore, given μ and μ_0 there is a number $\delta > 0$ such that $N_\delta(T_\mu) \subseteq T_{\mu_0}$. For the remainder for this section we shall fix a pair of numbers (μ, ϵ) such that there are Margulis constants $0 < \mu_1 < \mu < \mu_2$, a number δ such that $N_\delta(\partial T_\mu) \subseteq T_{\mu_2} \setminus T_{\mu_1}$, and $\epsilon = \frac{1}{4}\min\{\mu_1, \delta\}$. We shall call (μ, ϵ) a choice of MV-constants for $\mathbb{H}^3$.

Let (μ, ϵ) be a choice of MV-constants, and consider an ϵ-regular Voronoi decomposition of M. The fact that $N_\delta(\partial T_\mu) \subseteq T_{\mu_2} \setminus T_{\mu_1}$ means that we may adjust the boundary of T_μ by an arbitrarily small isotopy so that it is transverse to the Voronoi cells, and we will assume that we have done this for the remainder of this section. Our choice of ϵ implies that the thick part X_μ is contained in the Voronoi cells in the deep part, i.e. $X_\mu \subset \bigcup_{W \in \mathcal{W}} W$, so in particular $\partial X_\mu = \partial T_\mu$ is contained in the deep part. Furthermore, each deep Voronoi cell hits at most one component of T_μ.

If $\widetilde{M}$ is a cover of M, we will write $\widetilde{X}_\mu$ for the pre-image of X_μ in $\widetilde{M}$, and similarly $\widetilde{T}_\mu$ for the pre-image of T_μ in $\widetilde{M}$. In the case in which $\widetilde{M}$ is a cover of M, the connected components of $\widetilde{T}_\mu$ are covers of the connected components of T_μ. If a connected component of $\widetilde{T}_\mu$ is compact, then it is a solid torus, otherwise it is the universal cover of a solid torus, and we shall refer to such a component as an *infinite Margulis tube*. Injectivity radius at a point can only increase under taking covers, so the lift of a deep Voronoi cell is also deep, and so $\widetilde{X}_\mu$ is contained in the deep part of $\widetilde{M}$. Although the thick part of $\widetilde{M}$ may be strictly larger than $\widetilde{X}_\mu$, by abuse of notation, we will sometimes refer to $\widetilde{X}_\mu$ as the thick part of $\widetilde{M}$ and $\widetilde{T}_\mu$ as the thin part of $\widetilde{M}$.

2.4. Cell splitters

The polyhedral surfaces we construct will be locally constant, except for a discrete collection of points in the trivalent graph Y, which roughly speaking correspond to points $t \in Y$ for which the level set $f^{-1}(t)$ divide a Voronoi cell in half. For technical reasons, we use points which divide a ball of fixed size in the Voronoi cell in half, as we now describe.

Let t be a point in a trivalent tree Y. Then the complement $Y \setminus t$ has at most three connected components: if t lies in the interior of an edge, then there are precisely two connected components, while if t is a vertex, there are precisely three connected components. It will be convenient to consider the closures of these components, which are the closed sets obtained by adding the point t to each connected component of $Y \setminus t$. We shall denote the closures of these connected components of $Y \setminus t$ by $Y_t^{c_i}$, and we shall refer to them as the *complements* of t.

Let $\widetilde{M}$ be a cover of a closed hyperbolic 3-manifold, and let $h\colon \widetilde{M} \to Y$ be a Morse function onto a trivalent tree Y. Given $t \in Y$, let $H_t^{c_i} = h^{-1}(Y_t^{c_i})$, and we shall refer to these as the *complements* of $H_t = h^{-1}(t)$ in $\widetilde{M}$. As before, there are either two or three complementary regions depending on whether t lies in the interior of an edge, or is a vertex in Y.

Definition 2.7. Let $\widetilde{M}$ be a cover of a closed hyperbolic 3-manifold, with an equivariant ϵ-regular Voronoi decomposition $\mathcal{V}$. Let $h\colon \widetilde{M} \to Y$ be a Morse function to a trivalent tree Y, and let V be a Voronoi cell with center x. Suppose that a point $t \in Y$ has the property that for each complementary region $H_t^{c_i}$, the volume of $H_t^{c_i} \cap B(x, \epsilon/2) \cap V$ is at most half the volume of the topological ball $B(x, \epsilon/2) \cap V$. Then we say that t is a *cell splitter* for the Voronoi cell V.

Proposition 2.8. *Let $\widetilde{M}$ be a cover of a closed hyperbolic 3-manifold, with an equivariant ϵ-regular Voronoi decomposition $\mathcal{V}$. Let $h\colon \widetilde{M} \to Y$ be a Morse function to a trivalent tree, and let V be a Voronoi cell with center x. Then there is a unique cell splitter $t \in Y$ for V.*

Proof. We first show existence. Let B be the topological ball $B(x, \epsilon/2) \cap V$, and let v be the volume of this ball. Consider $h(B) \subset Y$. If there is a vertex of Y which is a cell splitter, then we are done. Otherwise, suppose no vertex of $h(B)$ is a cell splitter. If t is a vertex in $h(B)$ which is not a cell splitter, then there is at least one complementary region $Y_t^{c_i}$ such that $H_t^{c_i} \cap B(x, \epsilon/2) \cap V$ has volume more than $\frac{1}{2}v$, and $Y_t^{c_i} \cap h(B)$ has at least one fewer vertex. So proceeding by induction, we may reduce to the case in which $h(B)$ contains an interval I with no vertices such that $h^{-1}(I) \cap B(x, \epsilon/2) \cap V$ has volume at least $\frac{1}{2}v$. In this case, let t_0 and t_1 be the endpoints of I, and consider $h^{-1}([t_0, s])$, for $s \in I$. When $s = t_0$, this has volume less than $\frac{1}{2}v$, and has volume greater than $\frac{1}{2}v$ when $s = t_1$. As the volume changes continuously with s, there is a point t' such that $H_{t'}$ divides B into two regions, each of which has volume exactly $\frac{1}{2}v$, so t' is a cell splitter for V.

We now show uniqueness. First suppose t is a cell splitter which is not a vertex. Then there are precisely two complementary regions $H_t^{c_1}$ and $H_t^{c_2}$, and each of $H_t^{c_1} \cap B(x, \epsilon/2) \cap V$ and $H_t^{c_2} \cap B(x, \epsilon/2) \cap V$ must have equal volume exactly $\frac{1}{2}v$. Any other point t' has a complementary region which contains at least one of these complements, and so has volume greater than $\frac{1}{2}v$, and so can not be a cell splitter.

Finally suppose t is a cell splitter which is a vertex. Then there are three complements $H_t^{c_1}, H_t^{c_2}$ and $H_t^{c_3}$, each of which has volume at most $\frac{1}{2}v$. As each region has volume at most $\frac{1}{2}v$, any two regions must have total volume at least $\frac{1}{2}v$. Any other point $t' \in Y$ must have a complementary region which contains at least two of the complements of H_t, and so has a complement with volume strictly greater than $\frac{1}{2}v$, and so can not be a cell splitter. □

Definition 2.9. We say that a Morse function $f\colon M \to Y$ to a tree Y is *generic* with respect to a Voronoi decomposition $\mathcal{V}$ if the cell splitters for distinct Voronoi cells V_i correspond to distinct points $t_i \in Y$. We say a point $t \in Y$ is *generic* if it is not a critical value for the Morse function, and is not a cell splitter.

We may assume that f is generic by an arbitrarily small perturbation of f, and we shall always assume that f is generic from now on. Finally, we remark that a trivalent vertex in Y is not necessarily a cell splitter.

2.5. Polyhedral and capped surfaces

Let $\widetilde{M}$ be a hyperbolic 3-manifold with an equivariant ϵ-regular Voronoi decomposition, and let Q be a 3-dimensional submanifold of $\widetilde{M}$, with boundary an embedded separating surface F. In this section we show how to approximate Q by a union of Voronoi cells, which in turn gives an approximation to F by an embedded surface S which is a union of faces of Voronoi cells.

We say a region R is *generic* if for every Voronoi cell V_i with center x_i, the region consisting of the intersection of $B(x_i, \epsilon/2)$ with the interior of V_i does not have exactly half its volume lying in R. We say a separating surface F in $\widetilde{M}$ is *generic* if it bounds a generic region.

Let P be the collection of Voronoi cells for which at least half of the volume of $B(x_i, \epsilon/2) \cap \text{interior}(V_i)$ lies in Q. We say the P is the *polyhedral region* determined by Q. The polyhedral region P may be empty, even if Q is non-empty. The boundary of P is a polyhedral surface S, which we shall call the *polyhedral surface* associated to $F = \partial Q$, and is a normal surface in the dual triangulation. We will use the following bound on the number of faces and boundary components of the intersection of the polyhedral surface S with the thick part of the manifold, in terms of the area of the corresponding surface F. If S is a surface, we will write $|\partial S|$ for the number of boundary components of S, and if S' is a subset of a polyhedral surface S, we will write $\|S'\|$ for the number of faces of S which intersect S'.

Proposition 2.10. [HM16, Proposition 2.10, 2.13] *Let (μ, ϵ) be MV-constants, and let $\widetilde{M}$ be a cover of a closed hyperbolic 3-manifold, with an equivariant ϵ-regular Voronoi decomposition $\mathcal{V}$ and thick part X_μ. Then there is a constant K, which only depends on the MV-constants, such that for any generic embedded separating surface F in $\widetilde{M}$, the corresponding polyhedral surface S satisfies:*

$$\|S \cap \widetilde{X}_\mu\| \leqslant K\, area(F),$$

and

$$\left|\partial(S \cap \widetilde{X}_\mu)\right| \leqslant K\, area(F).$$

In [HM16], this result is stated for the level set of a Morse function on a closed hyperbolic manifold, and one may then observe that every separating surface F is the level set of some Morse function, though in fact, the proof only uses the fact that F is separating. In [HM16, Proposition 2.10] the bound is stated in terms of $S \cap \mathcal{W}$, where $\mathcal{W}$ is the deep part defined above in Definition 2.5. As the pre-image of a deep Voronoi cell in M is deep in $\widetilde{M}$, so the pre-image $\widetilde{\mathcal{W}}$ of $\mathcal{W}$ is deep in $\widetilde{M}$. Then, as $S \cap \widetilde{X}_\mu \subset S \cap \widetilde{\mathcal{W}}$, the stated bound follows immediately.

As each polyhedral surface S is closed, each boundary component of the surface $S \cap \widetilde{X}_\mu$ is contained in $\widetilde{T}_\mu$, so $S \cap \widetilde{X}_\mu$ is a properly embedded surface in $\widetilde{X}_\mu$. We now wish to cap off the properly embedded surfaces $S \cap \widetilde{X}_\mu$ with properly embedded surfaces in $\widetilde{T}_\mu$ to form closed surfaces. We warn the reader that the following definition differs slightly from the definition in [HM16], as we extend the definition to include the case in which $\widetilde{T}_\mu$ has infinite components.

Definition 2.11. A separating surface F in $\widetilde{M}$ gives rise to a polyhedral surface S, which meets $\partial \widetilde{T}_\mu$ transversely, and intersects $\partial \widetilde{T}_\mu$ in a collection of simple closed curves which is separating in $\partial \widetilde{T}_\mu$. We replace S inside the thin part by surfaces $\{U_i\}$ which we now describe. For each torus component T_i in $\partial \widetilde{T}_\mu$ choose a subsurface U_i bounded by $S \cap \partial T_i$. For each infinite component T_i, choose a not necessarily connected surface U_i as follows: for each essential curve in the annulus ∂T_i choose a disc it bounds in T_i, and then let U_i be the union of these discs with the planar surface bounded by the remaining inessential curves. We call the resulting surface a *capped surface* $S^+ = (S \cap \widetilde{X}_\mu) \cup \bigcup_i U_i$.

We will use the following property of the capped surfaces.

Proposition 2.12. *Let (μ, ϵ) be MV-constants, and let $\widetilde{M}$ be a cover of a closed hyperbolic 3-manifold M, with thin part $\widetilde{T}_\mu$, and with an equivariant ϵ-regular Voronoi decomposition $\mathcal{V}$. Then there is a constant K, which only depends on ϵ, such that for any generic embedded separating surface F in $\widetilde{M}$, the corresponding capped surface S^+ satisfies:*

$$\mathit{genus}(S^+) \leqslant K\,\mathit{area}(F).$$

The proof of this result is essentially the same as the proof of [HM16, Proposition 2.14], and instead of repeating the entire argument, we explain the minor extension needed. The only difference is that [HM16, Proposition 2.14] is stated for closed hyperbolic manifolds, whereas Proposition 2.12 is stated for covers of such manifolds, so the pre-image of the thin part $\widetilde{T}_\mu$ may have infinite Margulis tubes. This makes no difference to the estimates of the number of faces and boundary components of the resulting polyhedral surface in terms of the area of the original surface. The extension of the definition of capped surface to the infinite case only involves capping off with planar surfaces, so the same genus bounds hold.

2.6. Disjoint equivariant surfaces

Each collection of points t_i in Y corresponds to a collection S_i^+ of capped surfaces. In this section we show that if the collection of points is equivariant, then we may arrange for the capped surfaces to be disjoint and equivariant.

Let $\widetilde{M}$ be a cover of a closed hyperbolic 3-manifold M with an equivariant ϵ-regular Voronoi decomposition $\mathcal{V}$. We say a subset $U \subset \widetilde{M}$ is *equivariant* if it is equal to the pre-image of its projection to M.

Let W be a discrete equivariant collection of points in $\widetilde{X}$, none of which are either cell splitters or critical points of the Morse function h. We say two points t_i, t_j in W are *adjacent* if the geodesic connecting them in the tree $\widetilde{X}$ does not contain any other point of W. We may choose W such that the geodesic in $\widetilde{X}$ connecting any pair of adjacent points in W contains either a single cell splitter, a single trivalent trivalent vertex of $\widetilde{X}$, or neither of these two types of points.

Consider the collection S of polyhedral surfaces S_t, as t runs over W. As the collection W is equivariant, S is also equivariant. Note that although each surface in S is individually embedded, each surface in S will share many common faces with other surfaces in S. We will now make this collection simultaneously equivariantly disjoint, so that we may push them down to M to obtain a collection of disjoint surfaces which will act as our splitting surfaces in a graph splitting of M.

Proposition 2.13. *Let M be a closed hyperbolic 3-manifold of injectivity radius at least 2ϵ, with an ϵ-regular Voronoi decomposition $\mathcal{V}$, and a generic Morse function $f : M \to X$ onto a trivalent graph X with connected level sets. Let $p\colon \widetilde{M} \to M$ be the cover of M corresponding to the kernel of the induced map $f_*\colon \pi_1 M \to \pi_1 X$, and let $c : \widetilde{X} \to X$ be the universal cover of X. Let W be a discrete equivariant collection of points in $\widetilde{X}$. Then the collection of polyhedral surfaces $\{S_w \mid w \in W\}$ in $\widetilde{M}$ is equivariantly isotopic to a disjoint collection of surfaces $\{\Sigma_w \mid w \in W\}$, and furthermore this equivariant isotopy may be chosen to be supported in a neighborhood of the 2-skeleton of the induced Voronoi decomposition of $\widetilde{M}$.*

Proof. We now give a recipe for constructing surfaces Σ_t, for $t \in W$. Each individual surface Σ_t will be isotopic to the original S_t, but the union of the surfaces Σ_t will be equivariantly disjointly embedded in $\widetilde{M}$.

We first show that there is a canonical ordering of the polyhedral surfaces Σ_t which share a common face. Let Φ be a face of a Voronoi cell in $\widetilde{M}$, and let $V(x_1)$ and $V(x_2)$ be the adjacent Voronoi cells. Let t_1 and t_2 be cell splitters for $V(x_1)$ and $V(x_2)$, so that $H_{t_i} = h^{-1}(t_i)$ is the surface which divides $B_{\epsilon/2}(x_i)$ precisely in half, for $i = 1, 2$.

We say a point in $\widetilde{X}$ is *regular* if it is not a cell splitter, and not a critical point for the Morse function h.

Claim 2.14. The collection of regular points in $\widetilde{X}$ corresponding to polyhedral surfaces Σ_t which contain the face Φ is precisely the regular points contained in the geodesic in $\widetilde{X}$ from t_1 to t_2.

Proof. The two embedded surfaces H_{t_1} and H_{t_2} divide $\widetilde{M}$ into three parts; call them A, B and C, with A the part only hitting H_{t_1}, and B the part hitting both H_{t_1} and H_{t_2}.

Let γ be the geodesic in $\widetilde{X}$ from t_1 to t_2. Each point t in γ corresponds to a surface H_t dividing $\widetilde{M}$ at most 3 parts, one of which contains A, and another containing C. Let P_t be the part containing A. Then, writing $|A|$ for the volume of a region A,

$$\left|B_{\epsilon/2}(x_1) \cap P_t\right| \geqslant \left|B_{\epsilon/2}(x_1) \cap A\right| \geqslant \frac{1}{2}\left|B_{\epsilon/2}(x_1)\right|$$

and

$$\left|B_{\epsilon/2}(x_2) \setminus P_t\right| \geqslant \left|B_{\epsilon/2}(x_2) \setminus C\right| \geqslant \frac{1}{2}\left|B_{\epsilon/2}(x_2)\right|.$$

Therefore the two Voronoi cells $V(x_1)$ and $V(x_2)$ lie in different partitions of the Voronoi cells determined by t, and so Φ lies in the polyhedral surface Σ_t.

Conversely, suppose t does not lie on the path γ, then t divides $\widetilde{X}$ into at most three parts, and γ is contained in exactly one of these parts. This means that H_{t_1} and H_{t_2} are contained in the same complementary component of H_t, and so Φ cannot be a face of Σ_t. □

It suffices to show that we can isotope the normal surfaces, preserving the fact that they are normal, so that they have disjoint intersection in the 2-skeleton of the dual triangulation.

Let e be an edge of the dual triangulation, with vertices x_1 and x_2, with corresponding cell splitters t_1 and t_2. A normal surface S_i intersects e if and only if the corresponding point w_i lies in the geodesic $[t_1, t_2]$ in $\widetilde{X}$ connecting t_1 and t_2. The points w_i in e therefore inherit an order from $[t_1, t_2]$, and we may isotope the normal surfaces by a normal isotopy so that they intersect the edge e in the same order. As the interiors of each edge have disjoint images under the covering translations, and the collection of edges is equivariant, we may do this normal isotopy equivariantly.

Let Φ be a triangle in the dual triangulation, with vertices x_1, x_2 and x_3, and corresponding cell splitters t_1, t_2 and t_3. As above, the collection of normal surfaces which intersect an edge $[x_i, x_j]$ of Φ corresponds to those w_i lying in the geodesic $[t_i, t_j]$ in $\widetilde{X}$. The union of the three geodesics $[t_i, t_j]$ forms a minimal spanning tree for the three cell splitters in $\widetilde{X}$. Let t_0 be the center of this tree, i.e. the unique point that lies in all three geodesics. Note that the tree may be degenerate, so t_0 may be equal to one of the other vertices.

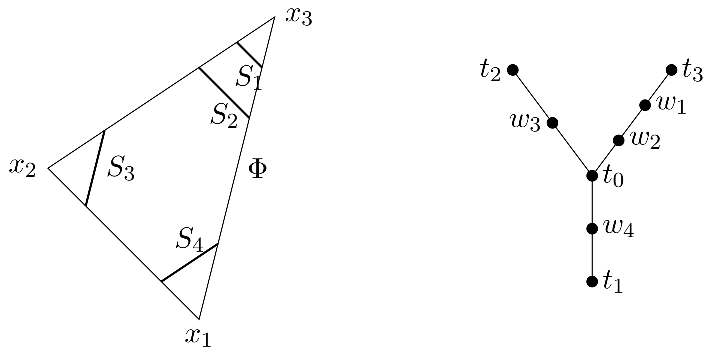

Figure 1: Example of normal surfaces intersecting a face of the dual triangulation.

Normal arcs parallel to the edge $[x_2, x_3]$ correspond to surfaces which hit both of the edges $[x_1, x_2]$ and $[x_1, x_3]$, so correspond points w_i which lie in both $[t_1, t_2]$ and $[t_1, t_3]$, and similarly for the other two cases. The intersection of these two geodesics in $\widetilde{X}$ is $[t_1, t_0]$, and so the corresponding surfaces appear in the same order on each of the edges in Φ, and so the arcs are disjoint. The same argument applies to each vertex of Φ. $\square$

As the resulting surfaces in $\widetilde{M}$ are disjoint and equivariant, they project down to disjoint surfaces in M.

We now show that the polyhedral surfaces, and their complements, project down homeomorphically into M. As the level set surfaces lift homeomorphically to $\widetilde{M}$, the area bound for the level sets of f is also an area bound for the level sets of h. Therefore, each polyhedral surface contains a bounded number of faces. The deck transformation group of the universal cover of a graph is equal to the fundamental group of the graph, which is a free group, so the orbit of any face consists of infinitely many disjoint translates. If two lie in the same connected component of a polyhedral surface, then that path corresponds to a covering translation, which has infinite order, so in fact the connected component contains infinitely many faces, which contradicts the fact that there is a bound on the number of faces in each component.

Each complementary region is compact, so the same argument applied to the complementary regions shows that they are all mapped down homeomorphically as well.

2.7. Bounded handles

We now bound the number of handles in a complementary region of the capped surfaces, which contains a single cell splitter. The following result will complete the proof of the left hand inequality in Theorem 1.3.

Proposition 2.15. *Let (μ, ϵ) be MV-constants, and let $\widetilde{M}$ be a cover of a closed hyperbolic 3-manifold, with an equivariant ϵ-regular Voronoi decomposition $\mathcal{V}$, and a generic Morse function $h\colon \widetilde{M} \to Y$, where Y is a tree. Let $\{u_i\}$ be a collection of points in Y, which separate the cell splitters in Y, and let $\{S_i^+\}$ be the corresponding collection of capped surfaces. If P is a complementary component of the capped surfaces in M, the region P has at most three boundary components, $S_{i_1}^+, S_{i_2}^+$ and $S_{i_3}^+$ say, where the final surface may be empty. Then P is homeomorphic to a manifold with a handle decomposition containing at most*

$$K \, \text{Gromov area}(M)$$

handles, where K depends only on the MV-constants.

We start with the observation that attaching a compression body P to a 3-manifold Q along a subsurface S of the upper boundary component of P, requires a number of handles which is bounded in terms of the Heegaard genus of P, and the number of boundary components of the attaching surface.

Proposition 2.16. [HM16, Proposition 2.16] *Let Q be a compact 3-manifold with boundary, and let $R = Q \cup P$, where P is a compression body of genus g, attached to Q by a homeomorphism along a (possibly disconnected) subsurface S contained in the upper boundary component of P of genus g. Then R is homeomorphic to a 3-manifold obtained from Q by the addition of at most $(4g + 2\,|\partial S|)$ 1-and 2-handles, where $|\partial S|$ is the number of boundary components of S.*

Proof (of Proposition 2.15). If P has two boundary components, then the argument is exactly the same as [HM16, Proposition 2.15], so we now consider the case in which P has three boundary components, which, without loss of generality we may relabel S_1^+, S_2^+ and S_3^+. Let t be the cell splitter corresponding to the region P, and let V be the corresponding Voronoi cell. As P has three boundary components, t must be a vertex of Y.

We first consider the case in which the Voronoi region V corresponding to the cell splitter t in $h(P)$ is disjoint from the thin part $\widetilde{T}_\mu$. Consider the three polyhedral surfaces S_1, S_2 and S_3, corresponding to the three capped surfaces, and let $\Sigma = \cup S_i \cup V$ be the union of the polyhedral surfaces, together with the Voronoi cell V. By Proposition 2.10, there is a constant K, which only depends on the MV-constants, such that the number of faces of Σ in the thick part is at most $3K_1$Gromov area(M), i.e.

$$\|\Sigma \cap \widetilde{X}_\mu\| \leqslant 3K_1 \text{Gromov area}(M),$$

where K_1 is the constant from Proposition 2.10. The number of boundary components of each surface $S_i \cap \widetilde{X}_\mu$ is also bounded by Proposition 2.10, and by Proposition 2.6, the Voronoi cell V has a bounded number J of vertices, edges and faces, where J depends only on the MV-constants. In particular, there is a constant A, depending only on the MV-constants, such that $P \cap X_\mu$ has a handle structure with at most A(Gromov area(M)) handles.

To bound the number of handles contained in P, we observe that P is a regular neighbourhood of the 3-complex obtained from capping off the boundary components of $\Sigma \cap \widetilde{X}_\mu$, using the parts of the capped surfaces in the thin part, i.e. the union of the components of $S_i^+ \cap \widetilde{T}_\mu$ over all three capped surfaces. Each component of $S_i^+ \cap \widetilde{T}_\mu$ has genus at most one, and the number of boundary components of $\Sigma \cap \widetilde{X}_\mu$ is bounded linearly in terms of Gromov area(M), therefore, there is a constant B, depending only on the MV-constants, such that the number of handles in P is at most B(Gromov area(M)), as required.

We now consider the case in which the region P has image $h(P)$ in Y which contains the cell splitter t, and the corresponding Voronoi cell V intersects $\widetilde{T}_\mu$. In this case, the connected components of $V \cap \widetilde{X}_\mu$ need not be topological balls, and there may be connected components of $P \cap \widetilde{T}_\mu$ whose boundary components are not parallel.

The connected components of $V \cap \widetilde{X}_\mu$ are handlebodies of bounded genus, as shown in the following result of Kobayashi and Rieck [KR11]. We state a simplified version of their result which suffices for our purposes, see [HM16] for further details.

Proposition 2.17. [KR11] *Let μ be a Margulis constant for $\mathbb{H}^3$, M be a finite volume hyperbolic 3-manifold, let $0 < \epsilon < \mu$, and let $\mathcal{V}$ be a regular Voronoi decomposition of M arising from a maximal collection of ϵ-separated points. Then there is a number G, depending only on μ and ϵ, such that for any Voronoi cell V_i, there are at most G connected components of $V_i \cap X_\mu$, each of which is a handlebody of genus at most G, attached to T_μ by a surface with at most G boundary components.*

Recall that attaching a handlebody of genus G to a 3-manifold along a subsurface of the boundary with at most G boundary components requires at most $6G$ handles:

Proposition 2.18. [HM16, Proposition 2.16] *Let Q be a compact 3-manifold with boundary and let $R = Q \cup P$, where P is a compression body of genus g, attached to Q by a homeomorphism along a (possibly disconnected) subsurface S contained in the upper boundary component of P of genus g. Then R is homeomorphic to a 3-manifold obtained from Q by the addition of at most $(4genus + 2\,|\partial S|)$ 1- and 2-handles, where $|\partial S|$ is the number of boundary components of S.*

Therefore, adding a Voronoi cell which intersects ∂T_μ may be realized by at most $6G^2$ handles. As each Voronoi cell in M lifts to $\widetilde{M}$, and $\widetilde{T}_\mu$ is the pre-image of T_μ, this same bound holds for adding a Voronoi cell in $\widetilde{M}$ which intersects $\partial \widetilde{T}_\mu$.

If the Voronoi cell intersects $\widetilde{T}_\mu$, then there may be components of $P \cap \widetilde{T}_\mu$ whose boundary surfaces are not parallel. This case is considered in the proof of [HM16, Proposition 2.15], when the manifold has no infinite Margulis tubes, so it suffices to consider the case of a component of P contained in an infinite Margulis tube. However, the case of an infinite Margulis tube in which neither surface is an essential disc is the same as the ordinary Margulis tube case, and if both surfaces essential discs then they are parallel. Finally, if exactly one surface is an essential disc, then the other surface lies in the same homology class, via the component of P in the infinite Margulis tube, and so, after surgering inessential boundary components, is also an essential disc. However, the number of boundary components is at most K_1Gromov area(M), and so the total number of extra handles over all components of P in the infinite Margulis tubes is also bounded by K_1Gromov area(M).

We may choose the constant K to be the maximum of the constants arising from the two cases considered above, thus completing the proof of Proposition 2.16. □

3. Topological complexity bounds metric complexity

In this section we will show bounds for metric complexity in terms of topological complexity, i.e. the right hand inequality in Theorem 1.3, using Theorem 1.1.

We start by reminding the reader of the topological properties of thin position for generalized Heegaard splittings, as shown by Scharlemann and Thompson [ST94] for the linear case and Saito, Scharlemann and Schultens [SSS16] for the graph case.

Theorem 3.1. [ST94, SSS16] *Let H be a graph splitting that is in thin position. Then every even surface is incompressible in M and the odd surfaces form strongly irreducible Heegaard surfaces for the components of M cut along the even surfaces.*

We will use the following result due to Gabai and Colding [CG18, Appendix A], building on recent work of Colding and Minicozzi [CM16]. It is not stated explicitly in their paper, but see [HM16, Theorem 3.2] for further details.

Theorem 3.2. [CG18] *Let M be a hyperbolic manifold, with (possibly empty) least area boundary, with a minimal Heegaard splitting H of genus g. Then, assuming Theorem 1.1, the manifold M has a (possibly singular) foliation by compact leaves, containing the boundary surfaces as leaves, such that each leaf has area at most $4\pi g$.*

By Theorem 3.1, we may consider the compression bodies in the graph splitting in pairs, glued along strongly irreducible Heegaard splittings, and then Theorem 3.2 guarantees that each pair has a foliation with each leaf having area at most $4\pi g$. These foliations contain the boundary surfaces as leaves, and so the foliations on each pair extend to foliations of the entire manifold, as required.

References

[BMNS16] Jeffrey Brock, Yair Minsky, Hossein Namazi, and Juan Souto, *Bounded combinatorics and uniform models for hyperbolic 3-manifolds*, J. Topol. **9** (2016), no. 2, 451–501.

[CDL03] Tobias H. Colding and Camillo De Lellis, *The min-max construction of minimal surfaces*, Surveys in differential geometry, Vol. VIII (Boston, MA, 2002), Surv. Differ. Geom., VIII, Int. Press, Somerville, MA, 2003, pp. 75–107.

[CG18] Tobias Holck Colding and David Gabai, *Effective finiteness of irreducible Heegaard splittings of non-Haken 3-manifolds*, Duke Math. J. **167** (2018), no. 15, 2793–2832.

[CM16] Tobias Holck Colding and William P. Minicozzi II, *The singular set of mean curvature flow with generic singularities*, Invent. Math. **204** (2016), no. 2, 443–471.

[DLP10] Camillo De Lellis and Filippo Pellandini, *Genus bounds for minimal surfaces arising from min-max constructions*, J. Reine Angew. Math. **644** (2010), 47–99.

[Gro88] M. Gromov, *Width and related invariants of Riemannian manifolds*, Astérisque **163-164** (1988), 6, 93–109, 282 (1989) (English, with French summary). On the geometry of differentiable manifolds (Rome, 1986).

[HM16] Diane Hoffoss and Joseph Maher, *Morse area and Scharlemann-Thompson width for hyperbolic 3-manifolds*, Pacific J. Math. **281** (2016), no. 1, 83–102.

[KL17] Daniel Ketover and Yevgeny Liokumovich, *On the existence of unstable minimal Heegaard surfaces* (2017), available at `arXiv:1709.09744`.

[KR11] Tsuyoshi Kobayashi and Yo'av Rieck, *A linear bound on the tetrahedral number of manifolds of bounded volume (after Jørgensen and Thurston)*, Topology and geometry in dimension three, Contemp. Math., vol. 560, Amer. Math. Soc., Providence, RI, 2011, pp. 27–42.

[Ket13] Daniel Ketover, *Degeneration of Min-Max Sequences in 3-manifolds* (2013), available at `arXiv:1312.2666`.

[PR86] Jon T. Pitts and J. H. Rubinstein, *Existence of minimal surfaces of bounded topological type in three-manifolds*, (Canberra, 1985), Proc. Centre Math. Anal. Austral. Nat. Univ., vol. 10, Austral. Nat. Univ., Canberra, 1986, pp. 163–176.

[Rub05] J. Hyam Rubinstein, *Minimal surfaces in geometric 3-manifolds*, Global theory of minimal surfaces, Clay Math. Proc., vol. 2, Amer. Math. Soc., Providence, RI, 2005, pp. 725–746.

[SSS16] Martin Scharlemann, Jennifer Schultens, and Toshio Saito, *Lecture notes on generalized Heegaard splittings*, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2016. Three lectures on low-dimensional topology in Kyoto.

[ST94] Martin Scharlemann and Abigail Thompson, *Thin position for 3-manifolds*, Geometric topology (Haifa, 1992), Contemp. Math., vol. 164, Amer. Math. Soc., Providence, RI, 1994, pp. 231–238.

Department of Mathematics and Computer Science
University of San Diego, San Diego, CA 92110-2492, USA
E-mail address: dhoffoss@sandiego.edu

CUNY College of Staten Island and CUNY Graduate Center
Staten Island, NY 10314, USA
E-mail address: joseph.maher@csi.cuny.edu

Received September 11, 2017
Accepted September 26, 2019