The Geometry of Carrier Graphs in Hyperbolic 3-Manifolds

BY

WILLIAM MICHAEL SILER
B.S., Rhodes College, 2006
M.S., Indiana University, 2008

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Chicago, 2013

Chicago, Illinois

Defense Committee:

Peter Shalen, Chair and Advisor
Daniel Groves, Mathematics, Statistics, and Computer Science
Marc Culler, Mathematics, Statistics, and Computer Science
David Dumas, Mathematics, Statistics, and Computer Science
Benson Farb, University of Chicago

To my mom, who once had me give her an hour long, at-the-board lecture on topology, during which she took careful notes, because if topology was important to me, then it was important to her.

To my dad, who still asks me interesting math questions, despite (or because of) knowing how long, detailed, and enthusiastic the answers will be.

ACKNOWLEDGMENTS

First and foremost, I would like to thank my wife Sarah, who has been endlessly loving and encouraging. She has made my life better in innumerable ways, and because graduate school has been a major part of my life for the past seven years, this thesis would have been much harder to come by without her. I would also like to thank my mother-in-law Beverly, without whose expert nannying services, this last year of graduate school would have been impossible.

I am very fortunate to have had Peter Shalen as my thesis advisor. I have always thoroughly enjoyed our mathematical conversations and the clarity of thought that they bring me. I am never more excited about doing math than when I'm leaving his office with a collection of new ideas to explore.

I am grateful to my thesis committee, each of whom has given me helpful insights into my work and math in general. I would especially like to thank Daniel Groves, who has answered my questions on many occasions.

I would like to express my gratitude to Hao Liang, Jonah Gaster, and Jason DeBlois for discussing math with me countless times and for being good friends.

I would like to thank the University of Illinois at Chicago math department for allowing me to be a part of their mathematical community. Similarly, I am grateful to the Indiana University math department for the time I spent there, and especially grateful to Chuck Livingston at IU for showing me that topology really is as fascinating as I had suspected. Finally, I want to thank the Rhodes College math department, especially Chris Mouron, Chris Seaton, Eric Gottlieb,

ACKNOWLEDGMENTS (Continued)

Ivaylo Ilinkin, and Robert England. The department always felt like a real community, and they sparked and kindled my interest in math.

WMS

TABLE OF CONTENTS

LIST OF FIGURES

SUMMARY

A carrier graph is a map from a finite graph to a hyperbolic 3-manifold M, which is surjective on the level of fundamental groups. We can pull back the metric on M to get a notion of length for the graph. We study the geometric properties of the carrier graphs with minimal possible length. We show that minimal length carrier graphs exist for a large class of 3-manifolds. We also show that those manifolds have only finitely many minimal length carrier graphs, from which we deduce a new proof that such manifolds have finite isometry groups. Finally, we give a theorem relating lengths of loops in a minimal length carrier graph to the lengths of its edges. From this we are able, for example, to get an explicit upper bound on the injectivity radius of M based on the lengths of edges in a minimal length carrier graph.

CHAPTER 1

INTRODUCTION

We will be working in the context of hyperbolic 3-manifolds. A hyperbolic manifold M is a Riemannian manifold with constant sectional curvature -1. Equivalently, it is a quotient of $\mathbb{H}^n$ (the unique simply connected Riemannian manifold with constant sectional curvature -1) by a discrete, torsion-free group of isometries acting properly discontinuously. This group is isomorphic to $\pi_1(M)$. The study of 3-manifolds has long been closely tied to the study of their fundamental groups, and the work of Thurston and Perelman made it clear that geometry (especially hyperbolic geometry) has a very important role to play in this study, as well. The Mostow-Prasad ridigity theorem implies that if two finite volume hyperbolic manifolds with dimension at least 3 have isomorphic fundamental groups, then they are isometric. The topology of a manifold determines the fundamental group, and this rigidity theorem says that in the case of a finite volume, hyperbolic manifold of dimension at least 3, the fundamental group determines the topology and geometry of the manifold. Thus, topological and geometric properties of M and group-theoretic properties of $\pi_1(M)$ correspond to each other, at least abstractly. One of the major themes of low-dimensional topology in recent decades has been to make this correspondence more concrete.

Carrier graphs are a tool for using geometry to study generating sets of $\pi_1(M)$. A carrier graph is a finite graph X along with a map $f : X \to M$, with $f_* : \pi_1(X) \to \pi_1(M)$ surjective. This definition is due to White (21). Implicit in this definition is the requirement that $\pi_1(M)$

be finitely generated; we will assume this throughout. It is clear that M has a carrier graph; in particular, one could map a bouquet of circles into M so that each circle maps to a representative of an element of a finite generating set for $\pi_1(M)$. Define the rank of a group to be the minimal number of elements needed to generate the group. For any carrier graph, it is clear that $\mathrm{rank}(\pi_1(X)) \geq \mathrm{rank}(\pi_1(M))$, and for the carrier graph just given, $\mathrm{rank}(\pi_1(X)) = \mathrm{rank}(\pi_1(M))$. We will assume throughout the text that this rank equality holds for all carrier graphs we consider. The benefit of this assumption is that it implies that any simple closed curve in X maps to a homotopically nontrivial loop in M, for otherwise, we could break this loop in X and get a carrier graph with rank strictly smaller than that of $\pi_1(M)$, which is impossible.

Since M is a Riemannian manifold, we can measure lengths of paths in M. For an edge $e \subset X$, define the length of e, $\mathrm{len}_f(e)$, to be the length of the path $f|_e$ (we will assume that the edges are all mapped to rectifiable curves). Define the length of X, $\mathrm{len}_f(X)$, to be the sum of the lengths of all the edges in X. A *minimal length carrier graph* (MLCG) $f: X \to M$ is a carrier graph with the property that for all carrier graphs $g: Y \to M$, $\mathrm{len}_f(X) \leq \mathrm{len}_g(Y)$. It is not obvious that minimal length carrier graphs should exist, but White (21) proves that they do in closed, hyperbolic 3-manifolds. In Theorem 1, we will expand his proof to a much larger class of hyperbolic 3-manifolds. White showed that when they do exist, MLCGs have some nice geometric properties, which we will list in the next chapter. By exploiting some of these properties, he proved that when M is closed, an upper bound on the rank of $\pi_1(M)$ implies an upper bound on the injectivity radius of M, where the injectivity radius of M is half the length of its shortest non-trivial loop.

MLCGs have also been used to study the relationship between the rank of $\pi_1(M)$ and the Heegaard genus $\mathrm{hg}(M)$ of M, which is the smallest g such that M may be obtained by gluing two genus g handlebodies together along their boundary. A simple topological argument shows that $\mathrm{rank}(\pi_1(M)) \leq \mathrm{hg}(M)$, and in the 1960s, Waldhausen asked if equality holds for all 3-manifolds. Boileau and Zieschang (4) found non-hyperbolic examples of 3-manifolds with rank not equal to genus, and Tao Li (9) recently produced a hyperbolic 3-manifold with rank not equal to genus. These results (and how long it took them to be found) suggest that the relationship between rank and genus is subtle. Souto (20) and Biringer (2) used MLCGs to show that for a large class of fibered 3-manifolds M, the rank of $\pi_1(M)$ is equal to the Heegaard genus of M, where a fibered 3-manifold is one that is the total space of a fiber bundle over the circle. More specifically, Biringer's result, which generalized Souto's, was that if the genus of the fiber is fixed, then for every $\varepsilon > 0$, there are only finitely many fibered 3-manifolds with injectivity radius greater than ε for which the rank of π_1 is not equal to the Heegaard genus. Namazi and Souto (11) used MLCGs to show that if two handlebodies of genus $g \geq 2$ are glued together via a sufficiently high power of a certain, generic type of surface homeomorphism, then the resulting manifold will have rank and Heegaard genus equal to g.

Biringer and Souto (3) used MLCGs to show that there are only finitely many closed hyperbolic 3-manifolds with an upper bound on the rank of π_1 and a lower bound on the injectivity radius of M and a lower bound on the first eigenvalue of the Laplacian.

All of the results just mentioned, except for White's original result, rely heavily on a single theorem (or sometimes a slight variation) about MLCGs, which we call the nested subgraph theorem. This is Theorem 4.2 of (2). Though White did not use this theorem in his original carrier graph paper, his result follows immediately from it. Since this one theorem about MLCGs has helped bring so much new insight into hyperbolic 3-manifolds (especially the old problem of rank versus genus), I believe it may be fruitful to consider MLCGs as an object of study in their own right. Hence, for my thesis research, I have studied some of the geometric properties of MLCGs.

After proving that MLCGs exist, an obvious question is whether they are unique. In Chapter 3, we will make this question more precise and show that the answer is no. However, we will also show in Theorem 3 that for a very large class of hyperbolic 3-manifolds M, although M may contain more than one MLCG, it can only have finitely many. This will follow from Theorem 2, which says that if two MLCGs are homotopic, they must be equal. From Theorem 3, we will derive a new proof that the isometry groups for these manifolds are finite.

Chapter 4 is devoted to proving Theorem 5, which connects the lengths of loops in a MLCG $f: X \to M$ to lengths of its edges. Roughly speaking, it says that if all of the loops in X are long, then the edges of X cannot be short. A simple loop in X is part of a minimal cardinality generating set for $\pi_1(X)$, and so it maps to part of a minimal cardinality generating set for $\pi_1(M)$. Thus, the contrapositive of Theorem 5 implies that if X contains a sufficiently short edge e, then $\pi_1(M)$ admits a minimal cardinality generating set containing a short element, and as the length of e approaches 0, so does the length of this generator.

The results and examples in Chapter 3 are available on the arXiv in (18), and the results and examples of Chapters 2 and 4 have been published in (19).

CHAPTER 2

BASIC PROPERTIES AND EXISTENCE

As mentioned in the introduction, White showed that MLCGs have some nice geometric properties. Specifically, he showed

Theorem (White (21))**.** *If M is a closed, hyperbolic 3-manifold, then M has a minimal length carrier graph. In addition, if $f: X \to M$ is a minimal length carrier graph for any hyperbolic 3-manifold M (closed or not), then:*

1. *the edges of X map to geodesics;*
2. *X is trivalent;*
3. *every edge of X has positive length; in particular, the images of vertices are still trivalent;*
4. *edges adjacent to the same vertex meet at an angle of $2\pi/3$.*

We can extend White's proof of the existence of minimal length carrier graphs to a larger class of 3-manifolds. For definitions of compression body and NP-end, see (5) and (10), respectively. An NP-end is essentially a topological end of $M \setminus \{\text{cusps of } M\}$.

Theorem 1. *Let M be an orientable, hyperbolic 3-manifold such that $\pi_1(M)$ is finitely generated and nonabelian. If M does not contain a minimal length carrier graph, then it has a compact core which is a compression body. Furthermore, M has a π_1-surjective, simply degenerate NP-end.*

Proof. Let l be the infimum of the lengths of all carrier graphs for M, and let $f_i : X_i \to M$ be a sequence of carrier graphs with geodesic edges and whose lengths approach l. In (21), White applies the Arzelà-Ascoli theorem to such a sequence and shows that the resulting limit is a minimal length carrier graph. For his argument to work, it is sufficient for there to be a compact set K containing $f_i(X_i)$ for all i. Suppose that M does not have a minimal length carrier graph and thus, that this condition does not hold for any subsequence of $\{f_i\}$.

Because it converges, the sequence $\{\mathrm{len}_{f_i}(X_i)\}$ has an upper bound L. There is a compact submanifold $C \subset M$ for which the inclusion map is a homotopy equivalence (see (17)), and in fact, by the topological version of the tameness theorem ((1), (6)), we can pick C so that $\overline{M \setminus C} = \partial C \times [0, \infty)$. The radius L neighborhood C_L of C is also compact, and so for some i_0, $f_{i_0}(X_{i_0}) \not\subset C_L$. Since $\mathrm{len}_{f_{i_0}}(X_{i_0}) \leq L$, we have that $f_{i_0}(X_{i_0}) \cap C = \emptyset$. Thus, $f_{i_0}(X_{i_0})$ is contained in $S \times [0, \infty)$ for some component S of ∂C. Since f_{i_0} is a carrier graph, it follows that the map $\pi_1(S) \to \pi_1(C)$ induced by inclusion is surjective, and in particular, the map $\pi_1(\partial C) \to \pi_1(C)$ is surjective. Since C is compact and has a π_1-surjective boundary component, C must be a compression body. For a proof of this well-known fact, see (5) Lemma 2.2.2.

Unless $C \cong (\text{surface}) \times I$, it has only one π_1-surjective boundary component S, and so $f_i(X_i)$ is eventually contained in the end $S \times [0, \infty)$. If $C \cong (\text{surface}) \times I$, then it is possible that $f_i(X_i)$ lies in each of the two ends infinitely often. In that case, we will assume that we have passed to a subsequence that lies entirely in one end $S \times [0, \infty)$.

Let M_0 be the manifold obtained by removing standard neighborhoods of the cusps of M. As can be seen by using the upper halfspace model of hyperbolic space, each cusp has a neighborhood isometric to the horoball $\{(x, y, z) \in \mathbb{H}^3 | z \geq Z\}$ modulo a group of translations, where $Z > 0$ depends on the cusp. There is a vertical projection to the horosphere $\{(*, *, Z) \in \mathbb{H}^3\}$ that descends to a retraction $\rho: M \to M_0$. The composition of ρ with any f_i is still a carrier graph for M. Since f_i cannot map entirely into a cusp (because then $\pi_1(M)$ would be abelian) and $\text{len}_{f_i}(X_i) < L$, there is an upper bound for the depth of $f_i(X_i)$, i.e. the z-coordinate, in any cusp. This, combined with the upper bound on the length of $f_i(X_i)$, implies an upper bound on the length of $\rho \circ f_i$. In addition, a point in a cusp with distance d from the boundary of the cusp gets moved a distance d by ρ. Since no point of $f_i(X_i)$ can have depth greater than L in a cusp, $\rho(f_i(X_i))$ is contained in a radius L neighborhood of $f_i(X_i)$. Thus, $\{\rho \circ f_i\}$ is a sequence of bounded length carrier graphs that eventually leaves every compact set, and $\rho \circ f_i$ misses fixed neighborhoods of every cusp. This implies that $\rho(f_i(X_i))$ is contained in an NP-end for large i.

We now have that M has a π_1-surjective, NP-end $\mathcal{E}$ which contains $\rho(f_i(X_i))$ for infinitely many i. According to the tameness theorem, there are two possibilities for the geometry of $\mathcal{E}$: it is geometrically finite or simply degenerate. Suppose $\mathcal{E}$ is geometrically finite. Then $\mathcal{E}$ has a flaring geometry. In particular, if $\mathcal{E} = S' \times [0, \infty)$, then it is easy to show that the injectivity radius at any point away from any cusps in $S' \times [t, \infty)$, goes to infinity as $t \to \infty$. Since $\rho(f_i(X_i))$ is not contained in any standard cusp neighborhood, contains nontrivial loops, and is exiting $\mathcal{E}$, the length of the carrier graph $\rho \circ f_i$ must being going to infinity. This contradicts these graphs having bounded length. Hence, $\mathcal{E}$ is simply degenerate. $\square$

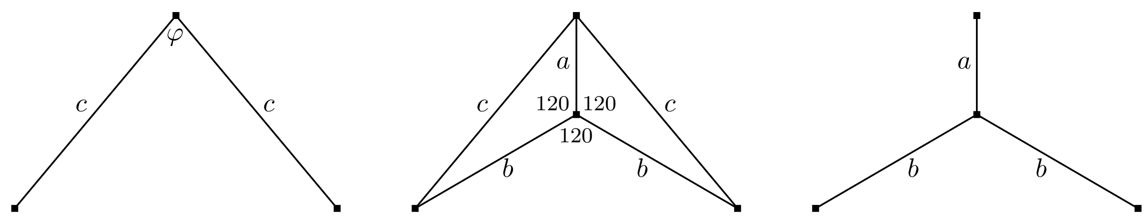

Figure 1. Shortening procedure

Suppose two geodesic segments in $\mathbb{H}^3$ with the same length meet at a shared endpoint with angle $\varphi < 2\pi/3$. We can replace the edges with a tripod (with the same endpoints) as shown in Figure 1. A little hyperbolic trigonometry shows (with the edge length labels from the figure) that $2c > 2b + a$, so the tripod is shorter. One can apply this procedure to a pair of geodesic edges that share an endpoint in a carrier graph. Generally, such a pair of edges will not be the same length, so we instead apply it to two "edge segments" (i.e. a subarc of an edge). Note that if an edge has both endpoints at the same vertex, we may apply the procedure with edge segments from the same edge, each with length up to half that of the full edge. This shortening procedure is the key tool that White uses to prove points 2-4 in the theorem above. We will need a stronger statement. Let $Sh(c, \varphi)$ be the reduction in length after performing the shortening procedure on edge segments of length c meeting at angle φ. If φ is understood, we will use the notation $Sh(c)$. With some hyperbolic trigonometry, one can see that a and b are functions of c and φ. So $Sh(c, \varphi) = 2c - 2b(c, \varphi) - a(c, \varphi)$.

Lemma 1. *For a fixed length c, if $\varphi < 2\pi/3$, $Sh(c, \varphi)$ is a strictly decreasing function of φ. For a fixed angle $\varphi < 2\pi/3$, $Sh(c, \varphi)$ is a strictly increasing function of c.*

Hence, smaller angles and longer edges produce a greater reduction in length in the shortening procedure.

Proof. Fixing c, we need to show that $2b+a$ is an increasing function of φ. We know c and its opposite angle, so the hyperbolic law of sines immediately shows that b is an increasing function of φ. Given b, we can compute a. Thus, treating b as a function of φ and a as a function of b, we have:

$$\frac{d}{d\varphi}(2b+a) = 2\frac{db}{d\varphi} + \frac{da}{db}\frac{db}{d\varphi}.$$

So it suffices to show that $\frac{da}{db} > -2$. Given that the angle opposite c is $2\pi/3$, the hyperbolic law of cosines says $\cosh c = \cosh a \cosh b + \frac{1}{2}\sinh a \sinh b$. Implicit differentiation with respect to b yields

$$\begin{aligned} 0 &= \sinh(a)a' \cosh(b) + \cosh(a)\sinh(b) + \frac{1}{2}\cosh(a)a'\sinh(b) + \frac{1}{2}\sinh(a)\cosh(b) \\ 0 &= a'(\sinh(a)\cosh(b) + \frac{1}{2}\cosh(a)\sinh(b)) + \cosh(a)\sinh(b) + \frac{1}{2}\sinh(a)\cosh(b) \\ a' &= -\frac{\cosh(a)\sinh(b) + \frac{1}{2}\sinh(a)\cosh(b)}{\sinh(a)\cosh(b) + \frac{1}{2}\cosh(a)\sinh(b)} \end{aligned}$$

We want $a' > -2$, which is equivalent to

$$\begin{aligned} \cosh(a)\sinh(b) + \frac{1}{2}\sinh(a)\cosh(b) &< 2\sinh(a)\cosh(b) + \cosh(a)\sinh(b) \\ \frac{1}{2}\sinh(a)\cosh(b) &< 2\sinh(a)\cosh(b) \end{aligned}$$

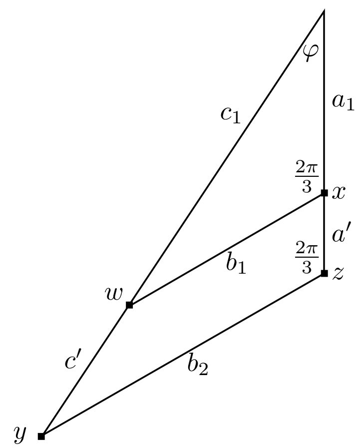

Figure 2. Shortening procedure triangles for $i = 1$ (smaller) and $i = 2$ (larger)

which is true if $a \neq 0$. However, in the shortening procedure, $a = 0$ if and only if $\varphi = 2\pi/3$; hence, the result follows.

Now fix $0 < \varphi < 2\pi/3$. Pick c_1 and c_2 with $0 < c_1 < c_2$. For $i = 1, 2$, let $a_i = a(c_i, \varphi)$ and $b_i = b(c_i, \varphi)$. Using some hyperbolic trigonometry, one can explicitly write down formulas for a and b as functions of c and φ. By (rather tediously) differentiating them, it is not hard to prove that a and b are increasing functions of c, so $a_1 < a_2$ and $b_1 < b_2$. Let $a' = a_2 - a_1$ and $c' = c_2 - c_1$. We wish to show that $2c_2 - 2b_2 - a_2 > 2c_1 - 2b_1 - a_1$. This is equivalent to $c' + b_1 > b_2 + \frac{1}{2}a'$. Figure 1 shows two symmetric triangles which each give the relationship between a, b and c. Figure 2 shows the corresponding triangles for a_i, b_i and c_i for $i = 1, 2$ on top of each other.

The quadrilateral $wxyz$ can be split into two triangles by inserting the diagonal $[x, y]$. By the triangle inequality, we have $c' + b_1 > \text{len}([x, y])$. The line segments $[x, y]$, $[y, z]$ and $[x, z]$

form a triangle in which the angle opposite $[x, y]$ is $2\pi/3$. As noted above, the shortening procedure works because of the observation that in such a triangle,

$$2\mathrm{len}([x, y]) - 2\mathrm{len}([y, z]) - \mathrm{len}([x, z]) > 0.$$

Since $\mathrm{len}([y, z]) = b_2$ and $\mathrm{len}([x, z]) = a'$, it follows that $c' + b_1 > b_2 + \frac{1}{2}a'$, which was our goal. $\square$

CHAPTER 3

UNIQUENESS

It is natural to ask whether minimal length carrier graphs are unique. We must take a little care regarding what it means to be unique. For example, if $f : X \to M$ is a minimal length carrier graph and $\eta : X \to X$ is a homeomorphism, then $f\eta$ is a minimal length carrier graph, as well, though intuitively, we would think of it as being the same graph as f. Thus, we define carrier graphs $f : X \to M$ and $g : Y \to M$ to be *strongly equivalent* if there exists a homeomorphism $\eta : X \to Y$ such that $f = g\eta$. Following Souto (20), we say f and g are *equivalent* if there exists a homotopy equivalence $\eta : X \to Y$ such that f and $g\eta$ are freely homotopic. The previously mentioned results about existence and basic properties of minimal length carrier graphs all still hold (with identical proofs) if you only consider carrier graphs within an equivalence class.

The question of uniqueness can now be stated in (at least) two more precise ways:

1. Must any two minimal length carrier graphs for M be strongly equivalent?
2. Must any two carrier graphs which both have minimal length within the same equivalence class be strongly equivalent?

The answer to both questions is no, according to the examples in Section 3.1. However, we will prove two weaker uniqueness results in Sections 3.2 and 3.3. In order to state them, we need one final (and very strong) notion of equivalence. Two carrier graphs $f, g : X \to M$ are *essentially equivalent* if $f = g\eta$ for some homeomorphism $\eta : X \to X$ that fixes vertices and leaves edges and their orientations invariant. In other words, f and g are the same except for reparameterizing the edges. Carrier graphs are *essentially distinct* if they are not essentially equivalent.

Theorem 2. *Let M be a hyperbolic 3-manifold and let $f : X \to M$ and $g : X \to M$ be carrier graphs, which either each have minimal length within their equivalence classes or each have minimal length globally. If f and g are homotopic, then f and g are essentially equivalent.*

And although there may be more than one carrier graph of minimal length globally or within an equivalence class, we show

Theorem 3. *Let M be a hyperbolic 3-manifold and suppose that $\pi_1(M)$ is nonabelian and M does not have a π_1-surjective, simply degenerate NP-end. Then M has only finitely many essentially distinct minimal length carrier graphs, and each equivalence class of carrier graphs can have only finitely many essentially distinct minimal length representatives.*

3.1 Non-uniqueness examples

Proposition 1. *Let M be a hyperbolic 3-manifold with $\mathrm{rank}(\pi_1(M)) = 2$. Suppose that M has a minimal length carrier graph $f : X \to M$ and that M has a fixed-point free isometry h of finite order not divisible by 3. Then hf is a minimal length carrier graph not strongly equivalent to f.*

Proof. It is clear that hf is a minimal length carrier graph. Suppose it is strongly equivalent to f. Then there exists a homeomorphism $\eta : X \to X$ such that $hf = f\eta$. Note that η cannot fix any point $x \in X$, for then h would fix $f(x)$. As noted above, minimal length carrier graphs must be trivalent. There are only two trivalent graphs of rank 2: one that looks like a θ and one that looks like eye-glasses. The eye-glasses graph does not admit a fixed-point free homeomorphism; so X is the θ graph. Up to homotopy, η must be the homeomorphism that swaps vertices and cyclically permutes the edges.

Let m be the order of h. Then $h^m f = f\eta^m$, which is equivalent to $f = f\eta^m$. Since m is not divisible by 3, η^m cyclically permutes the edges of X. Hence, f must map each edge to the same image, which contradicts f being a carrier graph because $f_*(\pi_1(X))$ would be trivial. $\square$

We can get concrete examples from this proposition. For example, let M be the figure 8 knot complement. Then M is a two-fold cover of the Gieseking manifold, hence it has a fixed-point free isometry h of order 2, and the rank of $\pi_1(M)$ is easily found to be two (from, say, the Wirtinger presentation); so M has non-unique minimal length carrier graphs.

We can also get closed examples. Reid (16) shows how to produce, for any $p > 1$, a closed hyperbolic 3-manifold M with a regular, cyclic cover N of degree p such that $\mathrm{rank}(\pi_1(N)) = 2$. If p is not divisible by 3, then N with its order p deck transformation satisfies the hypotheses of Proposition 1 and thus has non-unique minimal length carrier graphs.

We can take these examples a bit further to get examples of carrier graphs which are minimal in the same equivalence class but are not strongly equivalent. Reid's manifolds are formed as follows. Let φ be a pseudo-Anosov homeomorphism of a punctured torus T and let M_φ be the mapping torus of φ. Let a, b be generators of $\pi_1(T)$. Reid forms a manifold, which we are calling N, by taking the obvious p-fold cyclic cover of M_φ and performing a certain Dehn filling on it. It is shown that N is a p-fold cyclic cover of a manifold obtained from Dehn filling M_φ and the preimage of the filling torus for M_φ is the filling torus of N (in particular, the deck transformations of N leave the filling torus invariant). By abuse of notation, we will use a and b to refer to the generators of the fiber subgroup of M_φ and its cover and to their images in the filled manifold N. Reid shows that a and b generate $\pi_1(N)$. Let H be the filling torus of N. Then $N \setminus H$ is fiber bundle over S^1 with fiber a compact surface of genus 1 and with 1 boundary component. Choose representatives α and β of a and b, respectively, that lie in a particular fiber Σ of $N \setminus H$. If h is an order p deck-transformation of N, then $h \circ \alpha$ and $h \circ \beta$ are loops in the fiber $h(\Sigma)$ which also generate $\pi_1(N)$. Notice that there is a submanifold homeomorphic to $\Sigma \times [0,1] \subset N$ containing α, β, $h \circ \alpha$ and $h \circ \beta$. The manifold $\Sigma \times [0,1]$ is a genus 2 handlebody and the pairs $\{\alpha, \beta\}$ and $\{h \circ \alpha, h \circ \beta\}$ each generate its fundamental group. It is a well-known fact that any two minimal cardinality generating sets for the fundamental group of a handlebody are Nielsen equivalent; hence, h preserves the Nielsen equivalence class of the generating pair $\{a, b\}$.

Let $f\colon S^1 \vee S^1 \to N$ be the carrier graph given by mapping one of the S^1s to α and the other to β, and let f' be a carrier graph of minimal length in the equivalence class of f. Then $h \circ f'$ has minimal length in the equivalence class of the graph coming from $h \circ \alpha$ and $h \circ \beta$. In (20), Souto shows how to associate an equivalence class of carrier graphs to a Nielsen equivalence class of generators for π_1 and vice versa. His discussion of this correspondence implies that since h preserves the Nielsen equivalence class of $\{a, b\}$, f' and $h \circ f'$ are equivalent. However, Proposition 1 implies that these carrier graphs are not strongly equivalent. Hence, minimal length carrier graphs are not unique even within an equivalence class.

3.2 Non-existence of homotopies of minimal length carrier graphs

For the proof of Theorem 2, we will need a lemma.

Lemma 2. *Let x, y and z be distinct points in $\mathbb{H}^n$. Let x' (resp. y') be the midpoint of the geodesic between x (resp. y) and z. Then $d(x', y') \leq \frac{1}{2}d(x, y)$, and equality is achieved exactly when the angle $\angle xzy$ is 0 or π.*

Proof. Let $a = d(x', z)$, $b = d(y', z)$, $c = d(x', y')$ and $\gamma = \angle xzy$. We wish to show that $2c \leq d(x, y)$. This is equivalent to $\cosh(2c) \leq \cosh(d(x, y))$. By the hyperbolic law of cosines,

$$\begin{aligned}\cosh(c) &= \cosh(a)\cosh(b) - \sinh(a)\sinh(b)\cos(\gamma)\\ \cosh d(x, y) &= \cosh(2a)\cosh(2b) - \sinh(2a)\sinh(2b)\cos(\gamma).\end{aligned}$$

Now we just follow our noses: $\cosh(2c) = 2\cosh^2(c) - 1$, so we need to show

$$2\cosh^2(c) - 1 \leq \cosh(2a)\cosh(2b) - \sinh(2a)\sinh(2b)\cos(\gamma).$$

The left side is equal to

$$2\cosh^2(a)\cosh^2(b) - 4\cosh(a)\cosh(b)\sinh(a)\sinh(b)\cos(\gamma) + 2\sinh^2(a)\sinh^2(b)\cos^2(\gamma) - 1.$$

Notice that

$$\sinh(2a)\sinh(2b)\cos(\gamma) = 4\cosh(a)\cosh(b)\sinh(a)\sinh(b)\cos(\gamma).$$

So our goal becomes

$$2\cosh^2(a)\cosh^2(b) + 2\sinh^2(a)\sinh^2(b)\cos^2(\gamma) - 1 \leq \cosh(2a)\cosh(2b).$$

Using the identity $\cosh(2x) = 2\cosh^2(x) - 1$ and some algebra, one can see that this is equivalent to

$$\sinh^2(a)\sinh^2(b)\cos^2(\gamma) + \cosh^2(a) + \cosh^2(b) \leq \cosh^2(a)\cosh^2(b) + 1.$$

It suffices to prove this inequality with the assumption that $\cos^2(\gamma) = 1$, or equivalently, $\gamma = 0, \pi$. In this case, it is an equality, which follows from the identity $\cosh^2(b) = \sinh^2(b) + 1$. $\square$

Proof of Theorem 2. Suppose f and g are homotopic and essentially distinct. Let $H : X \times [0,1] \to M$ be a homotopy from f to g. The space $X \times [0,1]$ can be triangulated as follows. Let $e \subset X$ be an edge. Suppose e has distinct endpoints. Then a homeomorphism from e to $[0,1]$ can be extended to a homeomorphism from $e \times [0,1]$ to $[0,1] \times [0,1]$ (sending e to $[0,1] \times \{0\}$). The latter space has a triangulation with two triangles obtained by splitting the square along one of its diagonals. This triangulation can be pulled back to a triangulation for $e \times [0,1]$. If e's endpoints are the same (i.e. e is a loop), then it can be triangulated in essentially the same way, but with $e \times \{0\}$ and $e \times \{1\}$ identified. These triangulations can be glued together in an obvious way to yield a triangulation of $X \times [0,1]$. The map H can be made simplicially hyperbolic with respect to this triangulation, i.e. it can be made to send edges to geodesic segments and 2-simplices to geodesic triangles. We will assume this has been done. Note that this does not change the ends of the homotopy (f and g), since they already have geodesic edges by virtue of having minimal length.

We now construct a new carrier graph $h{:}X \to M$. Let v be a vertex of X. Then H maps $\{v\} \times [0,1]$ to a geodesic arc in M. Let $h(v)$ be the midpoint of that geodesic. Having defined h on the vertices on X, we can define it on the edges. Let e be an edge of X with distinct endpoints v and w. Then $e \times [0,1]$ consists of two triangles, which share a common edge. Let m be the midpoint of the (geodesic) image of that edge under H. There are geodesic arcs e_1 and e_2 connecting $h(v)$ to m and m to $h(w)$ and lying within $H(e \times [0,1])$. Let h map e homeomorphically to the path formed by concatenating e_1 and e_2. See Figure 3. If e has only one endpoint, then $h(e)$ is formed similarly; the picture is the same as Figure 3, except that the left and right edges are identified. Thus, the map h is essentially the midpoint of the homotopy H. It is clear that h is homotopic to f and g, which implies that it is a carrier graph in the same equivalence class as f and g.

We will show that $\text{len}_h(X) < \frac{1}{2}(\text{len}_f(X) + \text{len}_g(X))$. Because $\text{len}_f(X) = \text{len}_g(X)$ (since they both have minimal length), h will be shorter than both, which will contradict minimality and

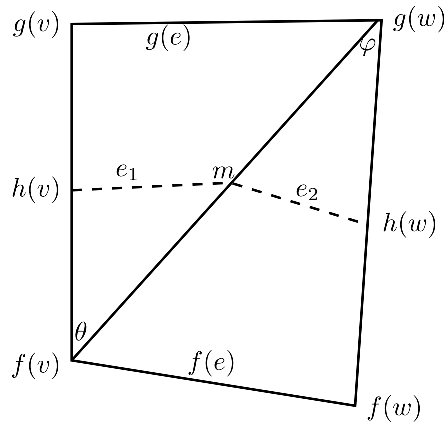

Figure 3. $H(e \times [0,1])$

complete the proof. Still referring to Figure 3, we can lift to $\mathbb{H}^3$ and apply Lemma 2 to the left and right triangles to get

$$\mathrm{len}_h(e) = \mathrm{len}(e_1) + \mathrm{len}(e_2) \leq \frac{1}{2}\mathrm{len}_g(e) + \frac{1}{2}\mathrm{len}_f(e) \tag{3.1}$$

with equality if and only if $\theta, \varphi \in \{0, \pi\}$. Summing over all edges, we get $\mathrm{len}_h(X) \leq \frac{1}{2}(\mathrm{len}_f(X) + \mathrm{len}_g(X))$. In order for this to be a strict inequality, we need for there to be at least one edge for which (Equation 3.1) is a strict inequality.

For any vertex $v \in X$, let $H_v = H(\{v\} \times [0,1])$. There must be some vertex v_0 such that H_v is not just a point. For otherwise, f and g would agree on the vertices of X and for each edge e of X, $f(e)$ and $g(e)$ would be geodesic segments homotopic relative to their endpoints. Hence, $f(e)$ would be the same as $g(e)$. If e has distinct endpoints, then it is clear that on e, f and g differ by an orientation preserving homeomorphism. If e is a loop, then perhaps f and g

map e to the same geodesic, but with opposite orientation. Since f and g are homotopic (via a homotopy that does not move the vertices), that would imply that the loop $f(e)$ is homotopic to its inverse. Since $\pi_1(M)$ is torsion-free and simple loops in X map to non-nullhomotopic loops in M, this cannot happen. Therefore, f and g must be essentially equivalent, which is a contradiction. Note that if H_v is not a single point, it is a geodesic path.

Pick an edge e with (not necessarily distinct) endpoints v and w, such that H_v is not a single point. If the inequality (Equation 3.1) for e is strict, then we are done. If it is an equality, then the angles θ and φ must each be either 0 or π. This implies that the angle between H_v and $f(e)$ is either 0 or π. The vertex v must have some other edge $e' \neq e$ adjacent to it. Since f is a minimal length carrier graph, the angle between $f(e')$ and $f(e)$ is $2\pi/3$. Thus, the angle between $f(e')$ and H_v cannot be 0 or π, and so for e', the inequality (Equation 3.1) must be strict. Hence, we get the desired contradiction

$$\mathrm{len}_h(X) < \frac{1}{2}(\mathrm{len}_f(X) + \mathrm{len}_g(X)) = \mathrm{len}_f(X).$$

□

3.3 Finiteness of minimal length carrier graphs

In the examples following Proposition 1, we found that minimal length carrier graphs were not unique because we can compose them with ambient isometries to get new carrier graphs. It is perhaps natural then to conjecture the following:

Conjecture 1. *If $f, g : X \to M$ are essentially distinct minimal length carrier graphs or carrier graphs of minimal length in the same equivalence class, then for some isometry $h \in \mathrm{Isom}(M)$, f and hg are essentially equivalent.*

It is well-known that a large class of hyperbolic 3-manifolds, including finite volume manifolds, have finite isometry groups. Thus, Theorem 3, that there are only finitely many minimal length carrier graphs, would follow immediately from this conjecture. On the other hand, the fact that the isometry group of M is finite (for a large class of hyperbolic 3-manifolds) follows from Theorem 3.

Theorem 4. *If M has nonabelian fundamental group and does not have a π_1-surjective, simply degenerate NP-end, then the isometry group of M is finite.*

Proof. Let $\mathcal{C}$ be the set of essential equivalence classes of minimal length carrier graphs in M. By Theorem 1, this set is nonempty, and by Theorem 3, this set is finite. It is clear that $\mathrm{Isom}(M)$ acts on $\mathcal{C}$, which gives a map from $\mathrm{Isom}(M)$ to the finite group of permutations of $\mathcal{C}$. Let K be the kernel of this map. It suffices to show that K is finite. Isometries of M that are in K fix a minimal length carrier graph $f : X \to M$ up to essential equivalence. In particular, for some vertex $v \in X$, they fix $f(v)$ and permute the images of the three edges attached to v. This gives a map from K to S_3, the permutation group on three elements. An element h of the kernel of this map would fix $f(v)$ and the three tangent vectors at $f(v)$ corresponding to the three edges of X attached to v. Since the angles between these edges are all $2\pi/3$, these tangent vectors span a plane in the tangent space. Lifting to $\mathbb{H}^3$, $\tilde{h}$ would fix some preimage of $f(v)$ and fix a hyperplane going through $f(v)$ pointwise (since it is an isometry and fixes the tangent plane). Hence, h must be the identity map. Thus, K injects into S_3, which means that K is finite. Therefore, $\mathrm{Isom}(M)$ is finite. □

Our proof of Theorem 3 uses Theorem 2 via the following proposition.

Proposition 2. *Suppose M satisfies the hypotheses of Theorem 3 and let $L > 0$. There are only finitely many essentially distinct carrier graphs which are minimal length within their equivalence class and have length less than or equal to L.*

Proof. Suppose M has an infinite sequence of carrier graphs $f_i : X \to M$, each of minimal length within its equivalence class and each with length less than or equal to L. Being minimal length implies the graphs are trivalent. There are only finitely many trivalent graphs of a particular rank; so we may pass to a subsequence and assume that every X_i is homeomorphic to a particular graph X. We will continue to call this sequence f_i. From the proof of Theorem 1, it follows that there is some compact submanifold that contains $f_i(X)$ for all i. Additionally, the bound on the length of the f_i implies that the sequence is equicontinuous. We can now apply the Arzelà-Ascoli theorem to get that a subsequence of $\{f_i\}$ converges uniformly. Therefore, for some large i and j, f_i is sufficiently close to f_j that the two maps must be homotopic. This contradicts Theorem 2. □

Proof of Theorem 3. Let $\mathcal{C}$ be the set of minimal length carrier graphs for M (up to essential equivalence), and let L be the length of any element of $\mathcal{C}$. Elements of $\mathcal{C}$ clearly have minimal length within their equivalence classes. Thus, $\mathcal{C}$ is contained in the set of carrier graphs which are of minimal length in their equivalence classes and have length less than or equal to L. The latter set is finite, by virtue of Proposition 2.

Similarly, the set of carrier graphs of minimal length within a particular equivalence class is seen to be finite by letting L be the length of any minimal length representative and applying Proposition 2 in the same way. □

Remark. The nested subgraphs theorem of Biringer and Souto says that a MLCG $f : X \to M$ has a nested sequence of subgraphs

$$\emptyset = Y_0 \subset Y_1 \subset \cdots \subset Y_n = X$$

such that the length of any edge in $Y_{i+1} \setminus Y_i$ is bounded above in terms of the injectivity radius of M, the rank of $\pi_1(M)$, the length of Y_i, and the diameters of the convex cores of the covers of M corresponding to $f_*(\pi_1(Y_i^j))$ where $Y_i^1, \ldots, Y_i^k$ are the components of Y_i. If Y_i is a tree, the convex core of the cover is trivial. If Y_i has a single simple loop, then the convex core of the cover is a loop whose length is bounded above by the length of the loop. Using this fact and the nested subgraphs theorem, one can show that the length of a MLCG in a rank 2 manifold is bounded above in terms of only the injectivity radius of M. In fact, a bound can be computed explicitly from Biringer's proof of the theorem in (2). Knowing the injectivity radius of M and an upper bound on the length of any minimal length carrier graph allows one to compute how close two such graphs must be to force them to be homotopic. Since distinct minimal length carrier graphs cannot be homotopic (by Theorem 2), this gives a lower bound on how close such graphs can be to each other. Finally, a packing argument using the volume of M puts an explicitly computable bound on the number of minimal length carrier graphs that can occur in M. Because there are only finitely many manifolds with a lower bound on injectivity radius and an upper bound on volume and each manifold has finitely many minimal length carrier graphs, a bound on the number of such graphs in each manifold exists abstractly. It is interesting that it can actually be computed explicitly from a given injectivity radius and volume. The required computations are an excellent exercise in hyperbolic geometry.

CHAPTER 4

LENGTHS OF LOOPS AND EDGES

The main result of this chapter concerns bounding lengths of edges in a minimal length carrier graph. Essentially, our result states that if all the simple loops in a minimal length carrier graph are long, then it cannot contain a short edge. Contrapositively, if X has a (sufficiently) short edge, then it contains a short loop.

4.1 Technical lemmas

We will need two simple, technical lemmas.

Lemma 3. *Fix any $\varphi < 2\pi/3$. There exists $z > 0$ and $s_0 > 0$ such that if $s < s_0$ and $c/s > z$, then $Sh(c) > s$.*

Proof. Suppose the following claim is true: there exists $c_0 > 0$ and $y > 0$ such that for all $c \leq c_0$, $Sh(c)/c \geq y$. Let $z = 1/y$; then for any positive s and c with $c \leq c_0$ and $c/s \geq z$,

$$\frac{Sh(c)}{c} \geq \frac{1}{z}$$
$$Sh(c) \geq \frac{c}{z} \geq s$$

Let $s_0 = c_0/z$, and pick s with $0 < s < s_0$ and c with $c/s > z$. Let $c' = zs$. Then $c' < c_0$ and $c'/s = z$; so $Sh(c') \geq s$. Since $c > c'$, by Lemma 1, $Sh(c) > Sh(c') \geq s$.

To prove the claim, note that by continuity, it suffices to show that

$$\lim_{c\to 0^+} \frac{Sh(c)}{c} > 0.$$

Since $Sh(0) = 0$, this limit is, by definition, the derivative of Sh evaluated at 0. One can, without too much difficulty, write b and a explicitly as functions of c (and φ), then take their derivatives at 0 to get that the derivative of Sh evaluated at $c = 0$ is

$$2 - \frac{3}{2}B - \frac{1}{2}\sqrt{4 - 3B^2}$$

where $B = \frac{2}{\sqrt{3}} \sin\left(\frac{\varphi}{2}\right)$. This quantity is easily seen to be positive. Showing that $Sh'(c) > 0$ for $c > 0$ by this method is much harder, which is why Lemma 1 uses a geometric argument for that case. □

Lemma 4. *Let X be a finite, metric graph (e.g. a minimal length carrier graph). For any $m, l_0 > 0$, if X has an edge e of length less than l_0, then e is contained in a connected subgraph S with the following properties:*

1. *For every edge e' adjacent to S, but not contained in S,*

$$\frac{\mathrm{len}(e')}{\mathrm{len}(S)} > m;$$

2. If $|S|$ is the number of edges in S, then

$$\mathrm{len}(S) < l_0(m+1)^{|S|-1}.$$

Proof. We will build S inductively, one edge at a time, so that at each step, the second condition is satisfied. We will stop once the first condition is also satisfied. Let e_1 be an edge of X with $\mathrm{len}(e_1) < l_0$, and let $S_1 = e_1$. Note that S_1 satisfies the second property of the lemma. Having defined S_i (satisfying property 2), if the first condition of the lemma holds for S_i, then set $S = S_i$ and stop. Otherwise, let e_{i+1} be an edge adjacent to S_i, but not in it, with $\mathrm{len}(e_{i+1}) \leq m\,\mathrm{len}(S_i)$. Set $S_{i+1} = S_i \cup e_{i+1}$. Note that

$$\begin{aligned}
\mathrm{len}(S_{i+1}) &\leq \mathrm{len}(S_i) + m\,\mathrm{len}(S_i) \\
&= (m+1)\,\mathrm{len}(S_i) \\
&< (m+1)l_0(m+1)^{i-1} \\
&< l_0(m+1)^{(i+1)-1}.
\end{aligned}$$

So S_{i+1} satisfies the second property, as well. Since X has finitely many edges, this process must eventually stop, yielding S. □

4.2 A short edge implies a short loop

We are now ready to prove the following theorem.

Theorem 5. *Suppose M is a hyperbolic 3-manifold and $f : X \to M$ is a minimal length carrier graph. Let $k = \operatorname{rank} \pi_1(M)$ and assume $k > 1$. For every $r > 0$, there exists $l > 0$ such that if every circuit in X has length greater than r, then every edge in X has length at least l. The value of l depends only on r and k.*

By a *circuit* in a graph, we mean a simple, closed curve. Note that each circuit in X represents an element of a minimal cardinality generating set for $\pi_1(X)$, and thus for $\pi_1(M)$. We will refer to an element of a minimal cardinality generating set as a *basis element*. Thus, to meet the criterion that every circuit in X has length greater than r, it suffices, for example, for every basis element of M to have length greater than r. Contrapositively, the theorem states that if X has a sufficiently short edge, then X, and thus M, has a short basis element.

We would like to emphasize that an explicit formula for l is given in the proof.

Proof. Let $k = \operatorname{rank} \pi_1(X)$. For $\varphi = \cos^{-1}(-1/3)$, let z and s_0 be as in Lemma 3. Choose $m > 0$ big enough so that $(1/2)(m-2)/(4k-5) > z$. Let $r > 0$, suppose that every circuit in X has length greater than r, and set

$$l = \min\left\{ \frac{s_0}{(4k-5)(m+1)^{2k-4}}, \frac{r}{(m+1)^{3k-4}} \right\}.$$

Suppose X has an edge e with $\operatorname{len}_f(e) < l$. Lemma 4 says that e is contained in a subgraph $S \subset X$ with $\operatorname{len}(S) < l(m+1)^{|S|-1}$, where $|S|$ is the number of edges in S, and for every edge e' that touches S but is not contained in it, $\operatorname{len}(e') > m\operatorname{len}(S)$. For brevity, let $L = \operatorname{len}(S)$. It is easy to see that X has exactly $3k-3$ edges; so $|S| \leq 3k-3$ and hence, $L < l(m+1)^{3k-4} \leq r$. Since

X has no circuits with length less than r, S must be a tree. Then $|S| \leq (\#\text{ vertices in } X) - 1 = 2k-3$, and $L < l(m+1)^{2k-4}$.

We are going to create a new carrier graph Y that will be homeomorphic to the graph obtained from X by collapsing S to a point. We first construct the abstract graph Y. Let R be the complement of the interior of S in X. Then $R \cap S$ consists of the vertices of S that have at least one adjacent edge not contained in S. We are going to modify R by "splitting" each valence 2 vertex (i.e. the vertices in $R \cap S$ that have exactly one edge from S attached to them). More precisely, let v be such a vertex and let e_1 and e_2 be the edges in R that are attached to it. Suppose that e_1 and e_2 are distinct edges and let w_i be the other endpoint of e_i. Remove v, e_1 and e_2 from R. For $i = 1, 2$, add in a new vertex v_i and an edge e_i' connecting w_i and v_i. If $e_1 = e_2$, then remove e_1 and v and replace them with two new vertices v_1 and v_2 and an edge connecting them. In either case, we will say that v_1 and v_2 were split from v. Call the modified graph R'. To obtain Y, add edges connecting every valence 1 vertex of R' to a single new vertex, the *cone point*. Set $C = Y \setminus R'$. Note that Y will have some vertices of valence 2 exactly at the points in $R' \cap C$. See Figure 4.

We now describe a map $g\colon Y \to M$. The graph R is a quotient of R' obtained by identifying pairs of vertices in R' split from the same vertex in R. On R', define g to be the composition of f with the quotient map $R' \to R$. Fix a nonvertex point $p \in S$, let p' be the cone point of C, and set $g(p') = f(p)$. If $[v', p'] \subset C$ is an edge in C, let v be the endpoint in R that v' was split from. There is a unique, injective path $[v, p]$ in S from v to p. Let g map the edge $[v', p']$ to the path $f|_{[v,p]}$. There will be some valence two vertices in Y coming from the endpoints of

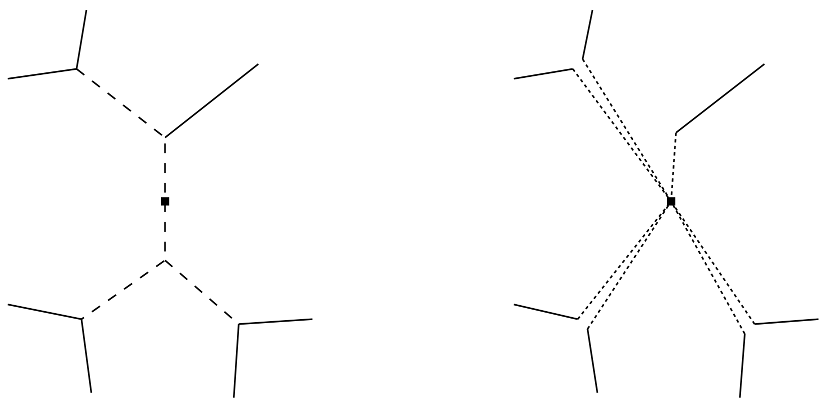

Figure 4. Collapsing a tree. Dashed lines are S, dotted lines are C, solid lines are part of R

R'. We will treat the two edges attached to such a vertex as a single edge, and though we may still refer to these endpoints, they will not be considered vertices. The map $g\colon Y \to M$ is still a carrier graph.

Notice that the length (with respect to g) of any edge in C is less than or equal to L. The number of such edges is at most twice the number of vertices in S, and the number of vertices in X is $2k-2$. Thus, $\text{len}_g(C) \leq (4k-4)L$. Since Y was formed by replacing S with C, $\text{len}_g(Y) \leq \text{len}_f(X) + (4k-5)L$. Some of the edges in Y map to non-geodesics, so we replace g by the map h, which is the same as g on the vertices of Y, but maps the edges to the geodesic arcs homotopic to the their images under g. The lengths of edges with respect to h will not be any longer than the lengths with respect to g, so we have $\text{len}_h(Y) \leq \text{len}_f(X) + (4k-5)L$.

Our goal will be to show that we can apply the shortening procedure to Y at the point $h(p')$ to reduce its length by more than $(4k-5)L$, thereby making it shorter than X. This will contradict X being a minimal length carrier graph and therefore, will show that X cannot contain the short loop e.

Corollary 7.2 of (7) says that if $Q_1, \ldots, Q_n \in \mathbb{H}^3$ are distinct from $P \in \mathbb{H}^3$, then

$$\sum_{1 \leq i < j \leq n} \cos \angle(Q_i, P, Q_j) \geq -n/2.$$

Since S is a tree and has at least one edge, p' has valence at least four. If we lift h on a small neighborhood of p' to $\mathbb{H}^3$ and apply this corollary to any four edges, we get that there are two edges attached to p' with angle at most $\cos^{-1}(-1/3)$ between them. Let η_1 and η_2 be two such edges and let $\varphi \leq \cos^{-1}(-1/3)$ be the angle between them. Note that we could have $\eta_1 = \eta_2$, if both endpoints from this edge are at p'.

We can get a lower bound for the lengths of the edges attached to p'. Let e_0 be an edge attached to p' and suppose e_0 has only one endpoint at p'. Then e_0 can be written as $e_0 = e_1 \cup e_2$, where $e_1 \subset R'$ and $e_2 \subset C$. From the construction of Y and g, we see that there is some edge $e_1' \subset X$ that touches S but is not contained in it, such that $g(e_1) = f(e_1')$. Hence, $\mathrm{len}_g(e_1) = \mathrm{len}_f(e_1') > mL$. Also, since e_2 is an edge in C, $\mathrm{len}_g(e_2) \leq L$. Under h, the image of e_0 comes from straightening $g(e_1 \cup e_2)$ into a geodesic. Applying the triangle inequality, we get that $\mathrm{len}_h(e_0) \geq (m-1)L$. Now suppose e_0 has both endpoints at p'. The argument here is similar: we write $e_0 = e_1 \cup e_2 \cup e_3$, where e_1 and e_3 are edges in C and have length at most L, and e_2 is in R' and has length at least mL. With two applications of the triangle inequality, we get $\mathrm{len}_h(e_0) \geq (m-2)L$.

We are going to apply the shortening procedure to the edges η_1 and η_2. If these are distinct edges, then we can use edge segments of length $\min\{\text{len}(\eta_1), \text{len}(\eta_2)\}$, which is at least $(m-1)L$, from each. If $\eta_1 = \eta_2$, then each edge segment can use up to half of the edge. Thus, we can use edge segments of length at least $(1/2)(m-2)L$. Let c be the length of the longest edge segments in η_1 and η_2 that we can do the shortening procedure on. We have $c \geq (1/2)(m-2)L$. The reduction in the length of Y coming from doing the shortening procedure on η_1 and η_2 is $Sh(c, \varphi)$. From Lemma 1, we have

$$Sh(c, \varphi) \geq Sh((1/2)(m-2)L, \cos^{-1}(-1/3)) > (4k-5)L.$$

The last inequality comes from Lemma 3 since, by our choices of m and L,

$$\frac{(1/2)(m-2)L}{(4k-5)L} > z$$

and

$$(4k-5)L < (4k-5)l(m+1)^{2k-4} \leq s_0.$$

After the shortening procedure we will have a carrier graph shorter than the minimal length carrier graph X, which is a contradiction. This implies that no edge of X can have length less than l. $\square$

4.3 Short loops, long generators

An immediate corollary of Theorem 5 is that a lower bound on injectivity radius gives a lower bound on the lengths of edges in a minimal length carrier graph, since there are no basis elements with length less than twice the injectivity radius. It is possible, *a priori*, that Theorem 5 does not give any more information than that corollary. In other words, it may be that if M has a small injectivity radius, then it must have a short basis element, too. We now give an example of a sequence of closed manifolds for which the injectivity radius goes to zero, but for which there is a lower bound on the length of a basis element and, thus, a lower bound on the length of any edge in their minimal length carrier graphs. This example was suggested independently by Ian Agol, Ian Biringer, and Juan Souto.

Let H be a genus 2 handlebody, and let DH be the double of H. We would like to chose an essential, separating, simple, closed curve $\gamma \subset \partial H$ sufficiently complicated so that $M_\infty = DH \backslash \gamma$ is a finite volume, cusped hyperbolic 3-manifold. To see that such a curve exists, first note that by the Hyperbolization Theorem of Perelman ((12), (14), (13)), it sufficies to be able to identify M_∞ with the interior of a compact manifold N with torus boundary, where N contains no essential sphere, disk, annulus or torus. A slight modification of the proof of Corollary 3.7 in (8) shows that this will be the case if γ has distance at least 3 in the curve complex from any curve in ∂H bounding a disk in H. Furthermore, it follows from Theorem 2.7 of (8) that γ can be chosen to make this distance arbitrarily large. Assume that such a γ has been chosen and that M_∞ has the desired properties. Let $M_n = H \cup_{g^n} H$, where $g : \partial H \to \partial H$ is Dehn twist along γ. Notice that M_n can be obtained from DH by $1/n$ Dehn surgery on a neighborhood of

γ; call the core of the filling torus γ_n. It is clear that the map g^n acts trivially on $H_1(\partial H)$, so one can easily check that $H_1(M_n) = \mathbb{Z}^2$. Since there is an obvious genus 2 Heegaard splitting of M_n, we have $\operatorname{rank} \pi_1(M_n) = 2$. According to Thurston's hyperbolic Dehn filling theorem (see (15)), the sequence M_n converges geometrically to M_∞. Thus, the minimal injectivity radius of M_n approaches zero. Also, for large enough n, any sufficiently short (nontrivial) curve must be contained in the Margulis tube around γ_n, so that in $\pi_1(M)$, it represents a power of γ_n. Hence, we only need to show that γ_n cannot be a basis element. Note that γ_n and γ are freely homotopic in M_n, and γ is trivial in $H_1(M_n)$, since it is a separating curve in ∂H. Any generating pair for $\pi_1(M_n)$ must descend to a pair of generators for $H_1(M_n) = \mathbb{Z}^2$, so neither γ_n nor any sufficiently short curve can be a basis element.

CITED LITERATURE

1. Agol, I.: Tameness of hyperbolic 3-manifolds. arXiv:math/0405568, 2004.

2. Biringer, I.: Geometry and rank of fibered hyperbolic 3manifolds. Algebraic & Geometric Topology, 9:277–292, 2009.

3. Biringer, I. and Souto, J.: A finiteness theorem for hyperbolic 3-manifolds. J. Lond. Math. Soc. (2), 84(1):227–242, 2011.

4. Boileau, M. and Zieschang, H.: Heegaard genus of closed orientable Seifert 3-manifolds. Inventiones Mathematicae, 76:455–468, 1984.

5. Bowditch, B.: Notes on tameness. Enseign. Math. (2), 56(3-4):229–285, 2010.

6. Calegari, D. and Gabai, D.: Shrinkwrapping and the taming of hyperbolic 3-manifolds. Journal of the American Mathematical Society, 19(2):385–446, 2005.

7. Culler, M. and Shalen, P. B.: Margulis numbers for Haken manifolds. Israel Journal of Mathematics, 190:445–475, 2012.

8. Hempel, J.: 3-Manifolds as viewed from the curve complex. Topology, 40:631–657, 2001.

9. Li, T.: Rank and genus of 3-manifolds. arXiv:1106.6302, 2011.

10. Matsuzaki, K. and Taniguchi, M.: Hyperbolic manifolds and Kleinian groups. Oxford Mathematical Monographs. New York, The Clarendon Press Oxford University Press, 1998.

11. Namazi, H. and Souto, J.: Heegaard Splittings and Pseudo-Anosov Maps. Geometric and Functional Analysis, 19(4):1195–1228, 2009.

12. Perelman, G.: The entropy formula for the Ricci flow and its geometric applications. arXiv:math.DG/0211159, 2002.

13. Perelman, G.: Finite extinction time for the solutions to the Ricci flow on certain three-manifolds. arXiv:math.DG/0307245, 2003.

14. Perelman, G.: Ricci flow with surgery on three-manifolds. arXiv:math.DG/0303109, 2003.

15. Petronio, C. and Porti, J.: Negatively Oriented Ideal Triangulations and a Proof of Thurstons Hyperbolic Dehn Filling Theorem. Expo. Math., 18(1):1–35, 2000.

16. Reid, A. W.: Some remarks on 2-generator hyperbolic 3-manifolds. In Discrete Groups and Geometry, London Math. Soc. Lecture Note Ser., 209–219. Cambridge, Cambridge Univ. Press, 1992.

17. Scott, G. P.: Compact Submanifolds of 3-Manifolds. Journal of the London Mathematical Society (2), 7:246–250, 1973.

18. Siler, M.: A note on the uniqueness of minimal length carrier graphs. arXiv:1208.2241, 2012.

19. Siler, M.: Lengths of edges in carrier graphs. Geometriae Dedicata, 161:1–10, 2012.

20. Souto, J.: The rank of the fundamental group of certain hyperbolic 3manifolds fibering over the circle. Geometry & Topology Monographs, 14:505–518, 2008.

21. White, M. E.: Injectivity radius and fundamental groups of hyperbolic 3-manifolds. Comm. Anal. Geom., 10(2):377–395, 2002.

VITA

Education

Ph.D. (Expected) University of Illinois at Chicago, 2008-Present
M.A. Indiana University, 2006-2008
B.S. Rhodes College, 2002-2006

Research Interests

Topology, hyperbolic geometry, geometric and combinatorial group theory

Papers

A note on the uniqueness of minimal length carrier graphs. arXiv:1208.2241, August 2012.

Lengths of edges in carrier graphs. Geometriae Dedicata, 161:1-10, December 2012.

Presentations

"The Geometry of Carrier Graphs and Hyperbolic 3-Manifolds", Geometry, Topology and Physics Seminar, University of Pittsburgh, September 2012

"The Geometry of Carrier Graphs," Geometry, Topology, Dynamics Seminar, University of Illinois at Chicago, May 2011

"Explosion Points and Indecomposable Functions," Chico Topology Conference, California State University, Chico, May 2006

Awards

Hazel King Thompson Fellowship, Indiana University, 2006-2007

Conferences Attended

RTG Workshop on Recent Progress on Hyperbolic 3-Manifolds, 2012, University of Michigan
Geometric Group Theory and Logic: The Elementary Theory of Groups, 2011, University of Illinois at Chicago
Conference in Honor of William Jaco, 2010, Oklahoma State University
More Examples of Groups, 2009, Ohio State University
RTG Workshop on Geometric Group Theory, 2009, University of Michigan
Graduate Topology Conference, 2008, University of Illinois at Urbana-Champaign
Chico Topology Conference, 2008, California State University, Chico

Teaching Experience

Teaching assistant, University of Illinois at Chicago, 2008-Present

- Differential Equations, Calculus, Precalculus, Finite Math for Business Majors

Teaching assistant coordinator, University of Illinois at Chicago, 2010 and 2012

- Observed and evaluated new teaching assistants
- Provided feedback to teaching assistants and department.

Assistant instructor, Indiana University, 2007-2008

- Calculus, Calculus for Business Majors
- Graded graduate level class *Topology II*