On the Height of a Homotopy

Erin W. Chambers* David Letscher†

Abstract

Given 2 homotopic curves in a topological space, there are several ways to measure similarity between the curves, including Hausdorff distance and Fréchet distance. In this paper, we examine a different measure of similarity which considers the family of curves represented in the homotopy between the curves, and measures the longest such curve, known as the *height* of the homotopy. In other words, if we have two homotopic curves on a surface and view a homotopy as a way to morph one curve into the other, we wish to find the longest intermediate curve along the morphing.

We prove that given a pair of disjoint embedded homotopic curves, among minimal height homotopies on the surface, there exists an ambient isotopy; in other words, the homotopy with minimum height never makes a "backwards" move and results in disjoint simple intermediate curves.

1 Introduction

There are many ways of measuring similarity between curves. Hausdorff distance is one common measure, which is (intuitively) the maximum distance that an adversary can force by picking a point on one curve and allowing you to choose any point on the other curve. While Hausdorff distance does measure closeness in space, it does not take into account the flow of the curve in space; two curves may have small Hausdorff distance but still not be "similar".

A second metric for measuring similarity between curves in Euclidean space is the Fréchet distance, which is the minimum length of a leash required to connect a man and dog as they travel, from one endpoint to the other, without backtracking, along the two curves. Fréchet distance is used in different applications as a more accurate measure of similarity, and algorithms have been developed to compute Fréchet distance in several different settings[1, 8, 9]. Several variants, such as geodesic Fréchet distance [3] and homotopic Fréchet distance [2], have also been introduced to generalize the notion of Fréchet distance to more general settings.

*Department of Mathematics and Computer Science, Saint Louis University, echambe5@slu.edu

†Department of Mathematics and Computer Science, Saint Louis University, letscher@slu.edu

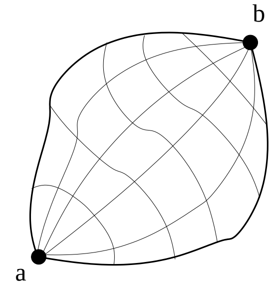

Figure 1: The height of the homotopy measures the maximum length of the curves "parallel" to α and β, while Fréchet distance measures the maximum length of the "transverse" curves.

In this paper, we examine a metric for measuring similarity between curves which is in many ways orthogonal to standard Fréchet distance. Any homotopy $H : I \times I \to S$ between two curves yields two families of curves: one set $H(s_0, t)$ (for fixed s_0) that run "between" the two curves being examined and the other $H(s, t_0)$ (for fixed t_0) that run "parallel" to the the curves being examined, see figure 1. Fréchet distance is the maximum length curve in the first family of curves, $H(s, \cdot)$, while the *height* of the homotopy is the maximum length curve in the second family, $H(\cdot, t)$.

Borrowing the concept of *thin position* from 3-manifold topology, we will show that among the minimal height homotopies between disjoint paths there is one that never "reverses direction" or "collides with itself". Thin position was developed by Gabai[6] and used by Thompson in the 3-sphere recognition algorithm[10]. The technique focuses on studying local properties of a sequence and then using local optimality conditions to prove global properties. We use this concept to prove the main theorem of the paper and also provide a characterization of move sequences that are of minimal "complexity".

2 Definitions

We will be working on a triangulated surface M, where curves lie along the edges of the triangulation and the edges of the triangulation are unweighted. In general, a *path* is a continuous map $p : [0, 1] \to M$. However, we restrict paths to follow the edges of the triangulation, where each edge is oriented consistently with traversing starting at $p(0)$ and ending at $p(1)$. The *length* of a path p, written $|p|$, is the number of edges (with multiplicity) in the path.

A path is a *geodesic* if it is impossible to perform a local reduction in its length. In other words (since the underlying graph is unweighted), a path is a geodesic if no edge in the path is immediately followed by its reversal and if no two cofacial edges appear consecutively along the path. Note that this is not the same as being a shortest path, as it is a purely local condition.

A path on a surface is *simple* if it is 1-1. Since (in a combinatorial sense) the same edge or vertex may be appear many times in a path, we will often examine paths that have been perturbed in an infinitesimally small neighborhood of the edges of the triangulation. We will say a path is *simple* if there exists such a perturbation to an embedded curve. Likewise, two paths will be considered *disjoint* if, after an infinitesimal perturbation, there have no points in common.

Two curves $\gamma_1, \gamma_2 : [0,1] \to M$ are *homotopic* if there is a continuous map $H : [0,1] \times [0,1] \to M$ such that $H(0,t) = \gamma_1(t)$ and $H(1,t) = \gamma_2(t)$. Essentially, this says that you can continuously deform one curves to another. However, since our paths lie completely in the edges of the triangulation, we will use an alternate mechanism to move from one path to another. We will study a *move sequence* from one path to another where each move is one of the following *elementary moves*.

- *Face lengthening* : A move from a single edge e_0 across a face to two edges e_1 and e_2.
- *Face shortening* : A move from two consecutive cofacial edges e_0 and e_0 across a face to a single edge e_2.
- *Spike* : Move across a single edge, so that an edge e followed by its reversal is included in the new path.
- *Reverse spike* : A reverse spike move, where an edge and its immediate reversal is removed from the path.

We will refer to the set of paths obtained by applying elementary moves one at a time as *intermediate paths*.

We can connect a move sequence with k moves to an obvious implied homotopy $H : [0,1] \to M$, where $H(s, 1/i)$ is equal to the i^{th} intermediate path in the sequence; the homotopy remains fixed between these paths except where the elementary move is being performed. If each of these intermediate paths is embedded then we will refer to the move sequence as *simple*. This is equivalent to saying that the implied homotopy can be perturbed to be an isotopy.

We would like to be able to say that, under appropriate conditions, moves sequences never backtrack and proceed monotonically from one path to another. To

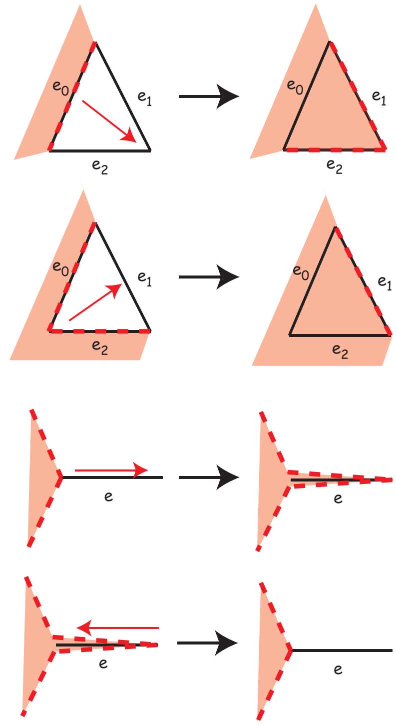

Figure 2: Elementary moves (top to bottom): face lengthening, face shortening, spike and reverse spike.

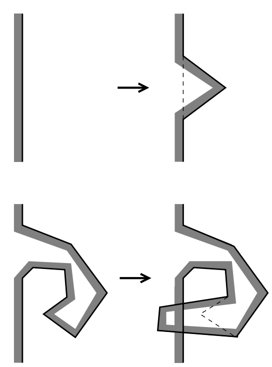

Figure 3: Two forward moves. Locally both move away from the region previous visited. Globally, the second results in a non-simple path.

be precise, consider a transverse orientation on a path that (locally) indicates where the path was previously. A move is considered *(locally) forward* if the move respects the transverse orientation. Figure 3 shows a forward move applied to a path where the transverse orientation is represented by shading; here, the shading is "behind" the curve, so the forward move goes away from the shaded side. Note that since this is purely local, a forward move may still cause the intermediate path to be non-simple. Also, note that move sequences consisting of only locally forward moves can have "spirals."

A move sequence is *embedded* if it is simple and only uses forward moves. This is equivalent to saying that after a perturbation, its associated homotopy is an isotopy that is an embedding everywhere except at the preimage of the two endpoints of the paths. Essentially, an embedded move sequence has intermediate paths that move smoothly across the disk, never crossing themselves or other intermediate paths.

The *height* of a homotopy is the maximum length of any intermediate curve: $\max_{t\in[0,1]} |H(\cdot,t)|$; similarly, the *height* of a move sequence is the length of the longest curve in the sequence. We wish to determine the minimum height homotopy between two curves which form the boundary of a planar, unweighted triangulation; in other words, we want the morphing between these two curves that keeps the maximum length of an intermediate curve as small as possible. However, it is not immediately obvious that this homotopy is embedded or forward; our main result, stated formally and proven in the next section, is that some minimum height move sequence is embedded and proceeds uniformly from one path to another without spirals or other degeneracies.

To accomplish this, we need a more precise way to compare two move sequences. Given a sequence of moves, the *length spectrum* is the set of all lengths of the intermediate paths in the sequence. Two length spectrums can be compared by ordering each in decreasing order and comparing the two lists lexicographically. A move sequence is said to be in *thin position* if its length spectrum is lexicographically minimal among all possible move sequences between the same paths. A move sequence that is in thin position has minimal height. Furthermore, every subsequence of moves also has minimal height.

A move sequence is *locally thin* if you cannot decrease the lengths in its length spectrum by any of the following local improvements:

1. Remove a pair of sequential moves where the intermediate curves before and after the pair of moves are combinatorially identical.
2. Reverse the order of a path lengthening move followed by a path shortening move that are independent of each other.
3. Replace a pair of moves that accomplish the result as a single move. For example, a spike move followed by an adjacent face shortening move can be replaced by a single face lengthening move.

We will see that embedded locally thin move sequences share many properties with move sequences that are in thin position.

3 Unweighted Planar Triangulations

The setting for all our results in the next two sections is a planar, unweighted triangulation (so our underlying manifold is a disk), with two distinguished vertices a and b on the outer face of the graph. Our goal is to characterize the minimum height homotopy from one side of the outer face (a path from a to b along the outer face) to the other side of the outer face.

At each stage of a homotopy from one boundary curve to the other, we have a connected curve between a and b. Our goal is to argue that in a minimum height homotopy, these intermediate paths never move backwards - namely, once an elementary move occurs, it will never be in our interest to move back across that face or edge. We will show that any move sequence that contains a backwards move is not in thin position (which immediately implies that it cannot be a minimum height homotopy). Furthermove, the move sequence must be embedded or, equivalently, the homotopy induced by the move sequence can be infitensimally perterbed to be an embedded isotopy.

Theorem 1 *Given a move sequence from one side of the boundary of an unweighted planar triangulation to the other side, if the move sequence is in thin position, then it is embedded.*

Corollary 2 *There exists a minimum height moves sequences that is embedded.*

The proof of theorem 1 will follow from the following two propositions. The first shows that there are no backwards moves, and the second shows that any move sequence consisting of only forward moves is embedded.

Proposition 3 *Any move sequence in thin position will never contain a backwards move.*

The proof of this proposition relys on the observation that if there is a backwards move for a move sequence in thin position, then the move immediately prior to it must share an edge with the backwards move. In a case by case analysis, this pair of moves can be replaced by different moves that reduce the complexity of the move sequence.

Proposition 4 *A forward move sequence from an arc on the boundary of a disk to the complementary arc in the boundary is embedded.*

This can be proved using arguments involving covers of topological spaces.

4 Characterizing Move Sequences in Thin Position

In applications of thin position to 3-manifold topology, the local maximums and minimums of sequences that are in thin position have particularly nice properties. The same is true for move sequences.

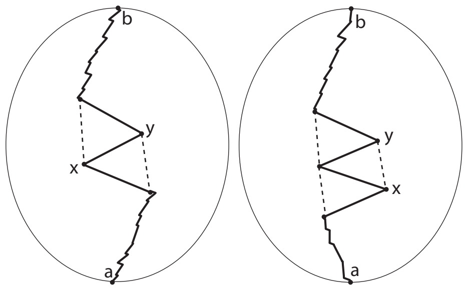

Figure 4: The configurations of paths at the local maximimums of a move sequence in (locally) thin position. The paths from a to x and y to b are geodesics.

Theorem 5 *If a move sequence is either in thin position or is embedded and locally thin then:*

1. *A path in the move sequence whose length is a local minimum is a geodesic.*
2. *A path in the move sequence whose length is a local maximum is geodesic everywhere except two or three points, and at these points the path has, up to symmetry, one of the configurations shown in Figure 4.*

5 Extensions and Open Questions

The case analysis used in the proof proposition 3 will also extend to weighted triangulations where the weights on each face satisfy the triangle inequality. Moreover, the same case analysis plus a few additional arguments can be used to prove similar results about move sequence between the boundary components of a triangulated annulus.

In sections 3 and 4, we have characterized the movement of any minimum height homotopy. The primary remaining open question, of course, is to find a polynomial time algorithm which, given two cycles on a combinatorial surface, computes a homotopy of minimum height (or at least the height of the minimum homotopy). Some initial work in this area has been done for the planar version of the problem, where the graph itself is a series parallel graph whose edges do not need to satisfy the triangle inequality [5]. One possible strategy for an algorithm in more general settings would rely on proving that the shortest path appears in a move sequence in thin position, and then recursively computing the minimum height homotopy in each half of the graph using our characterization of local minimum and maximum intermediate paths. However, our proofs do not give that the shortest path will appear in the minimum height homotopy, although we conjecture that it does.

It is not even clear that the problem is not NP-Complete, since it bears a close resemblance to finding the cut width of the dual graph. In fact, if we only allow face lengthening and face shortening moves, the problem is entirely equivalent to finding the cut width of the dual graph, which is NP-Hard even in planar graphs. (See [4] for a survey of cut width and other similar graph layout problems.)

Acknowledgements

The authors would like to thank Glencora Boradaile for initial comments on the proof in section 3.

References

[1] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *International Journal of Computational Geometry and Applications*, 5(1–2):75–91, 1995.

[2] E. W. Chambers, É. C. de Verdière, J. Erickson, S. Lazard, F. Lazarus, and S. Thite. Homotopic fréchet distance between curves, or walking your dog in the woods in polynomial time. *Computational Geometry: Theory and Applications*, 2009. To appear.

[3] A. Cook and C. Wenk. Geodesic Fréchet and Hausdorff distance inside a simple polygon. Tech. Rep. CS-TR-2007-004, University of Texas at San Antonio, 2007.

[4] J. Díaz, J. Petit, and M. Serna. A survey of graph layout problems. *ACM Comput. Surv.*, 34(3):313–356, 2002.

[5] D. Eppstein. Personal communication, 2008.

[6] D. Gabai. Foliations and the topology of 3-manifolds III. *Journal of Differential Geometry*, 26:479–536, 1987.

[7] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2001.

[8] A. Maheshwari and J. Yi. On computing Fréchet distance of two paths on a convex polyhedron. In *Proceedings of the 21st European Workshop on Computational Geometry*, pages 41–44, 2005.

[9] E. Sriraghavendra, K. Karthik, and C. Bhattacharyya. Fréchet distance based approach for searching online handwritten documents. In *ICDAR07*, pages 461–465, 2007.

[10] A. Thompson. Thin position and the recognition problem for S^3. *Mathematical Research Letters*, 1:613–630, 1994.

Appendix

In the proof of proposition 3, we will use more than just the length spectrum of the move sequence to show that there are no backwards moves. For a move sequence s, define the *complexity* as the vector $C(s)$ as the quadruple

$$\begin{aligned} C(s) = \quad & (\text{height spectrum, number of moves,} \\ & \text{number of backwards moves,} \\ & \text{time of first backwards move}) \end{aligned}$$

We will examine move sequences where $C(s)$ is minimal when ordered lexicographically.

Lemma 6 *In the move sequence s minimizing $C(s)$, the elementary move immediately preceding the first backwards move must involve some edge which is used in the first backwards move.*

Proof. If the move directly before the first backwards move did not involve an edge involved in the first backwards move, then we may form s' by swapping the order of the these two moves without changing anything else. This gives a complexity $C(s')$ with the first three coordinates identical but the time of the first backwards move earlier, which contradicts our choice of s. □

Lemma 7 *A move sequence minimizing $C(s)$ contains no backwards moves.*

Proof. Consider a move sequence s minimizing $C(s)$. If it contains a backwards move, we will show that $C(s)$ can be made smaller by modifying or removing the first backwards move.

Case I: The first backwards move is a face lengthening move, which replaces an edge e_0 with two edges e_1 and e_2 that share common face with e_0. By Lemma 6, the move immediately preceding the first backwards move must involve e_0.

If the elementary move directly before the first backwards was a Face I move, we may suppose without loss of generality that it moved from e_1 across the face to e_0 and e_2. The first backwards move must then go from e_0 to e_2, e_1, so the sequence of edges in the leash after the move will be e_2, e_2, e_1. To simplify this, we eliminate these two moves, and replace them with an Edge I move, which inserts a spike along edge e_2. This modification does not increases any lengths of the intermediate curves and reduces the number of moves overall, which contradicts our choice of s as minimizing $C(s)$. See Figure 5.

If the elementary move immediately before the first backwards move is a face shortening move, then we must go from e_1, e_2 to e_0 and then immediately have the backwards move from e_0 to e_1, e_2. We can simply cancel both of these moves out, which does not increase the any of the path lengths but which decreases the number of moves, again contradicting our choice of s. See Figure 6.

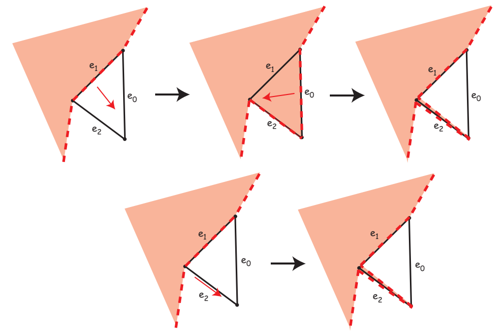

Figure 5: If the first backwards is a face lengthening move which occurs immediately after a face lengthening move (top), we can replace the two moves with an Edge I move (bottom) and get a better move sequence.

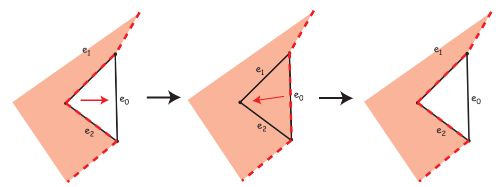

Figure 6: If the first backwards is a face lengthening move which occurs immediately after a face shortening move, we can simply cancel the two moves to get a move sequence which is shorter and has one fewer backwards move.

It is not possible for the move immediately preceding the first backwards move to be an spike or reverse spike move, since the move must involve the copy of e_0 which will move backwards, and any copy of e_0 inserted by an spike or reverse spike move will not border the face which the first backwards move goes across.

Case II: Suppose the first backwards move is a face shortening move, which replaces the edges e_0 and e_1 with the cofacial edge e_2. By Lemma 6, the previous move must involve either e_0 or e_1.

Again we have several possible choices for the move immediately preceding our first backwards move. If the previous move is a face lengthening move, we have two possibilities, both of which can be improved: see Figures 7. If the previous move is face shortening , then we can reduce the number of moves; see Figure 8. The previous move cannot be a spike move, since the there would be no face adjacent to one of the edges for a backwards face shortening move to go across. If the previous move is a reverse spike move, then we can simply swap the order of the two moves to make the first backwards move occur earlier; see Figure 9.

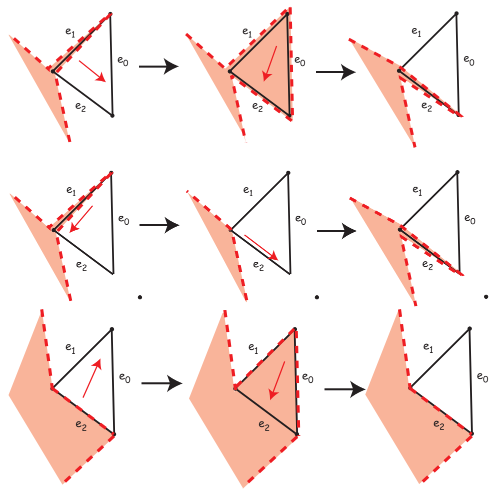

Figure 7: If the first backwards move is a face shortening move which occurs immediately after a face lengthening move, we have two possibilities (row 1 or row 3); in either case, we can switch the moves to make the first backwards move occur earlier (row 2) or to decrease the total number of moves.

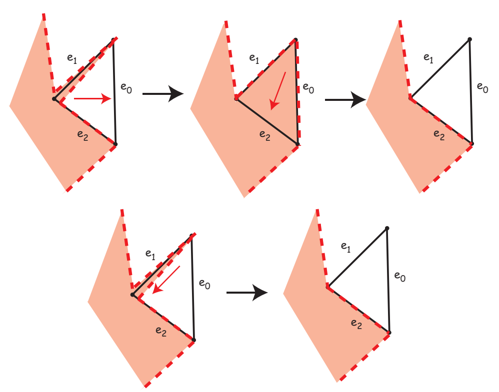

Figure 8: If the first backwards move is a face shortening move which occurs immediately after a face shortening move (top), we can replace the two moves with a single backwards reverse spike move, decreasing the number of moves and moving the first backwards move earlier (bottom).

Case III: Suppose the first backwards move is a spike move, which introduces two copies of some edge e_2. By Lemma 6, we know that the move immediately preceding the first backwards move must introduce the point which the backwards spike move comes from.

If the move immediately preceding the first backwards move is a face lengthening move, then we can simply replace it with a forwards move that results in the same sequence, reducing the number of backwards moves; see Figure 10. If the previous move was a face shortening move, we have three possibilities, any of which can be modified to make $C(s)$ smaller; see Figure 11. If the

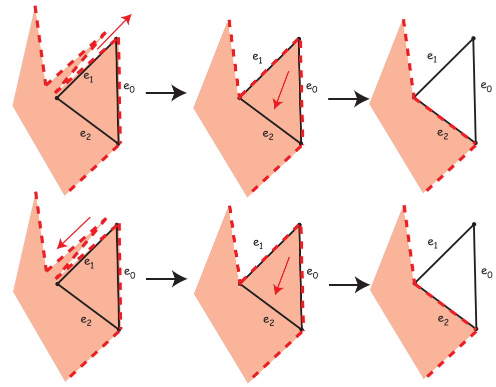

Figure 9: If the first backwards move is a face shortening move which occurs immediate after an reverse spike move (top), then we can replace with a backwards reverse spike followed by a reverse face shortening move, which moves the first backwards move ealier (bottom).

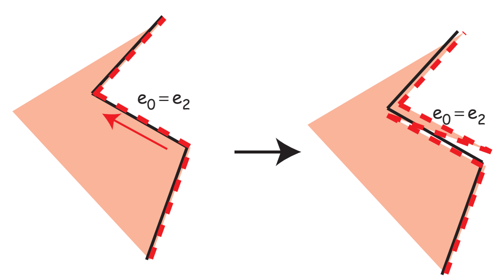

Figure 10: If the first backwards move is a spike move which immediately follows a face lengthening move (top), we can reduce the number of backwards moves (bottom).

move was a spike move, then we can replace the first backwards move by a forwards spike move instead, resulting in the same intermediate path. If the previous move was a reverse spike move, then we must be simply introducing and then removing a spike, so we can cancel the two moves out.

Case IV: Suppose the first backwards move is a reverse spike move, which takes two consecutive copies of an edge e on the intermediate path and removes them. By Lemma 6, we know that the move immediately before must have created one of the copies of e which will be removed.

Consider the edges immediately before and after the two copies of e on the intermediate path. If either of those edges is another copy of e, we can replace the backwards move with a forward reverse spike move which results in an equivalent intermediate path. So we may assume that the edge before and after the copies of e on the curve are different than e.

Now consider the move immediately preceding the

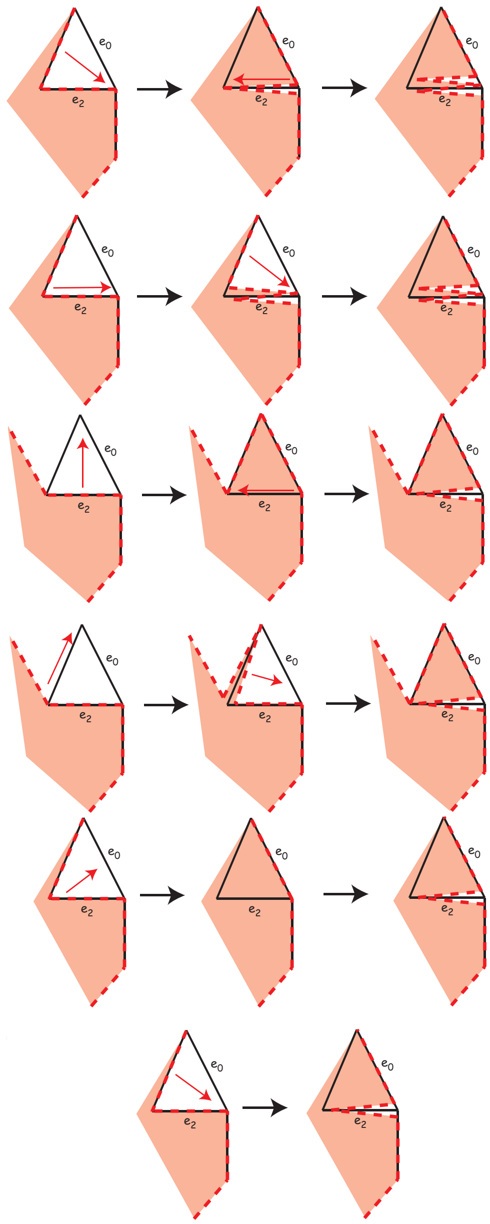

Figure 11: If the first backwards move is a spike move which is preceded by a face shortening move, we have 3 possible cases; see rows 1, 3, and 5 for the initial sequence and 2, 4, and 6 for the modified sequences.

first backwards move. Neither a face lengthening or face shortening move can result in this configuration, since either move would need to be a backwards move in order to leave such a spike. If the preceding move is a spike move, then we can simply cancel the two moves out. The only remaining case, where the preceding move is a reverse spike move, also cannot happen, since a move which removes a spike cannot leave duplicate copes of an edge involved afterwards. □

Proof of proposition 3. Suppose that a move sequence s is in thin position. Assume that s contains a backwards move. Then $C(s)$ is not minimal and one of the cases in the proof of the previous lemma can be used to reduce the complexity. These simplifications can be repeated until there are no backwards moves. Somewhere is this sequence of simplifications, a backwards move must be removed. All of the simplifications in the previous lemma that remove a backwards move reduce the length spectrum or reduces the number of moves. Note that the simplifications reduce the number of moves do so by removing values from the length spectrum, which reduces its length. So if a move sequence has backwards moves, then its length spectrum can be lowered. This contracts the assumption that s is in thin position. □

Note that the simplication moves used in proving that there are no backwards moves can be turned into an algorithm to simplify any move sequence. In particular, the cases above can be used to move the first backwards move earlier and eventually eliminate it; iterating this process will remove all backwards moves. At the same time, move sequences that are not locally thin can be simplified. This yields an algorithm to turn an arbitrary move sequence into one that is locally thin and embedded. In the original move sequence has n moves than this algorithm runs in $O(n^2)$ time.

To prove that the move sequence without backwards moves is embedded we need to use a couple of concepts from topology; see [7] or any other introductory text for a more detailed discussion of these issues. First, a map $f : X \to Y$ is a *local homeomorphism* if for every $x \in X$ there exists an open neighborhood U of x such that $f|_U$ is a homeomorphism. A *covering map* is a continuous map $p : C \to X$ such that for every $x \in X$ there exists an open neighborhood U of x such that $p^{-1}(U)$ is a disjoint union of open sets such that p restricted to any component maps homeomorphically onto U. In this situation C is called a *covering space* or *cover* of X.

Proof of proposition 4. Consider the homotopy constructed from a move sequence consisting of only forward moves. Using the construction discussed in the second section, this fails to be an embedded isotopy is several ways. To remedy this, we can modifiy the homotopy so that it is a local homeomorphism. When performing each move, we can perturb each curve towards their "right" in a arbitrarily small neighborhood of the paths. Figure 12 illustrates this pertubation. Note that this change is purely local and turns the homotopy into a local homeomorphism away from the start and end verticies. However, it is plausible that the curve could sprial back upon itself while still being a local homeomorphism.

The homotopy $H : [0,1] \times [0,1] \to S$ obtained by this proceedure is not a homeomorphism since $H(s, 0) =$

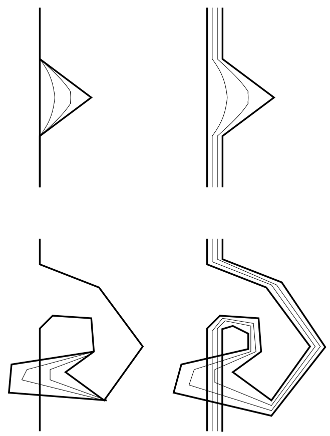

Figure 12: Perturbing the homotopy associated with a move sequence to a local homeomorphism. (Top) Moving across a face (Bottom) A local homeomorphism that is a not a homemophism.

$H(0,0)$ and $H(s,1) = H(0,1)$ for all s. However, it induces a continuous map $[0,1] \times [0,1]/\{(s,0) \sim (s',0),(s,1) \sim (s',1)\forall s\} \to S$. This induced map is a local homeomorphism from a disk onto itself and hence is a covering map. There are no non-trivial covers of a disk. Hence this map is a homemorphishm, which implies that H is an embedded isotopy and the move sequence must be embedded. □

Proof of theorem 5. If a locally minimum intermediate path is not a geodesic, then the move sequence can be simplified, reducing the lengths of the intermediate curves (and contradicting the fact that it is a local minimum).

At any locally maximum intermediate path, there must be adjacent edges that either form the tip of a spike or are edges of the same face. In fact, since we are at a local maximum in a move sequence, there have to be at least two of these configurations: one that simplifies the curve forwards and the other backwards. If any of these two configurations do not share an edge in common then the move sequence can be modified to interchange the order that these moves are done. This would reduce the complexity of the length spectrum, which cannot happen to a move sequence that is thin or locally thin. This also implies that there can be at most three points where the path fails to be a geodesic, as the existence of a forth point would allow two disjoint moves to opposite sides of the curve.

The previous arguement shows that there are either two or three points where the path fails to be a geodesic. These points are adjacent on the path, and the simplifications alternate sides of the path. There are are a total of 24 possible configurations. The cases differ on whether spike or face moves are involved and to which side they occur. Up to symmetry there are 9 possibilities that can be distinguished as you follow the path and see either spike moves or faces moves at these exceptional points. Note that reversing the sequence or swapping the sides that the moves occur on result in equivalent configurations and use the same arguments as below.

Case I (Spike-Spike): For this to occur a spike move followed by a reverse spike move must occur and they share an edge in common. The paths immediately before and after the local maximum would be combinatorially identitical. This pair of moves can be removed reducing the complexity of the length spectrum.

Case II (Spike-Face): Here, a spike move must occur followed by a face shortening move. These two moves share an edge in common. This pair of moves can be replaces by a single face lengthening move reducing the complexity of the length spectrum.

Case III (Face-Face): This is the first allowable case in the statement of the theorem; see Figure 4.

Case IV (Spike-Spike-Spike): This case is similar to case I above; again a pair of the spike moves can be removed, simplifying the move sequence.

Case V (Spike-Spike-Face): Either a spike move or a face move occurs immediately before the local maximum and a spike occurs immediately after. The same arguments as in case I and II can be used to remove a move in either case.

Case VI (Spike-Face-Spike): One of the spike moves occurs immediately before the local maximum, and the face shortening must occur after. These two moves can be removed and replaced by a single move as was done in case II.

Case VII (Spike-Face-Face): Either a spike or face lengthing move occurs immediately before the local maximum and a face shortening follows. Using the same arguments as in case I and II, both cases can be resolved.

Case VIII (Face-Spike-Face): A face lengthing occurs prior to the local maximum and is followed by a spike reversal. The same argument in case II can replace these two moves by a single one, simplifying the move sequence.

Case IX (Face-Face-Face): This can happen and is the second of the two allowable configurations in the statement of the theorem; see Figure 4.

□