

On the complexity of optimal homotopies

Citation for published version (APA):

Chambers, E. W., de Mesmay, A., & Ophelders, T. A. E. (2018). On the complexity of optimal homotopies. In A. Czumaj (Ed.), *29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018* (pp. 1121-1134). Association for Computing Machinery, Inc. <https://doi.org/10.1137/1.9781611975031.73>

DOI:

[10.1137/1.9781611975031.73](https://doi.org/10.1137/1.9781611975031.73)

Document status and date:

Published: 07/01/2018

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

On the complexity of optimal homotopies

Erin Wolf Chambers *

Arnaud de Mesmay[†]

Tim Ophelders[‡]

Abstract

In this article, we provide new structural results and algorithms for the HOMOTOPY HEIGHT problem. In broad terms, this problem quantifies how much a curve on a surface needs to be stretched to sweep continuously between two positions. More precisely, given two homotopic curves γ_1 and γ_2 on a combinatorial (say, triangulated) surface, we investigate the problem of computing a homotopy between γ_1 and γ_2 where the length of the longest intermediate curve is minimized. Such optimal homotopies are relevant for a wide range of purposes, from very theoretical questions in quantitative homotopy theory to more practical applications such as similarity measures on meshes and graph searching problems.

We prove that HOMOTOPY HEIGHT is in the complexity class **NP**, and the corresponding exponential algorithm is the best one known for this problem. This result builds on a structural theorem on monotonicity of optimal homotopies, which is proved in a companion paper. Then we show that this problem encompasses the HOMOTOPIC FRÉCHET DISTANCE problem which we therefore also establish to be in **NP**, answering a question which has previously been considered in several different settings. We also provide an $O(\log n)$ -approximation algorithm for HOMOTOPY HEIGHT on surfaces by adapting an earlier algorithm of Har-Peled, Nayyeri, Salvatipour and Sidiropoulos in the planar setting.

1 Introduction

This paper considers computational questions pertaining to *homotopies*: in broad terms, a homotopy between two curves in a topological space is a continuous deformation between these two curves. This can be formalized either in a continuous setting, where it constitutes one of the fundamental constructs of algebraic topology, but also in a more discrete one, where the input is a simplicial, or more generally cellular description of a topological space; this latter setting will be the focus of this article. While considerably more restrictive than the more traditional mathematical settings, this setting

is nonetheless of key importance in applications areas such as graphics or medical imaging, where inputs are generally represented by triangular meshes built upon scanned point sets from an underlying 3D object.

Investigating homotopies from a computational perspective is a well-studied problem, dating back to the work of Dehn [13] on contractibility of curves, which has strong ties to geometric group theory. While deciding whether two curves in a 2-dimensional complex are homotopic is well-known to be undecidable in general (see for example Stillwell [29]), when the underlying space is a surface, efficient, linear-time algorithms have been designed to test homotopy [15, 17, 26]. In this article, we add a quantitative twist to this problem: the HOMOTOPY HEIGHT problem consists, starting with two disjoint homotopic curves on a combinatorial surface, of finding the homotopy of minimal height, that is, where the length of the longest intermediate curve in the homotopy is minimized. (We refer the reader to Section 2 for formal definitions.) The notion of homotopy height has obvious appeal from a practical perspective, as it quantifies how long a curve has to be to overcome a hurdle: for example, deciding whether a bracelet is long enough to slide off over a hand without breaking is essentially the question of homotopy height. From a computational side, deformations of minimal height minimize the necessary stretch and can be used to quantify how similar curves are, as in map or trajectory analysis.

1.1 Our results

We begin by considering two curves forming the boundary of a discrete annulus, and study the homotopy between these boundaries of minimal height. Our article leverages on recent results in Riemannian geometry [10, 11], and in particular on a companion article co-authored with Gregory Chambers and Regina Rotman [6] where we prove that in the Riemannian setting, such an optimal homotopy can be assumed to be very well behaved. Firstly, it can be assumed to be an isotopy, so that all the intermediate curves remain simple. Secondly, this isotopy can be assumed to only move in one direction and never sweeps any portion of the disk twice; we refer to this property as **monotonicity**, which we will define more precisely in Section 3.

These isotopy and monotonicity properties turn

*Dept. of Computer Science, Saint Louis University.

email: echambe5@slu.edu

[†]Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, 38000 Grenoble, France.

email arnaud.de.Mesmay@gipsa-lab.fr

[‡]Dept. of Mathematics and Computer Science, TU Eindhoven, the Netherlands.

email: t.a.e.ophelders@tue.nl

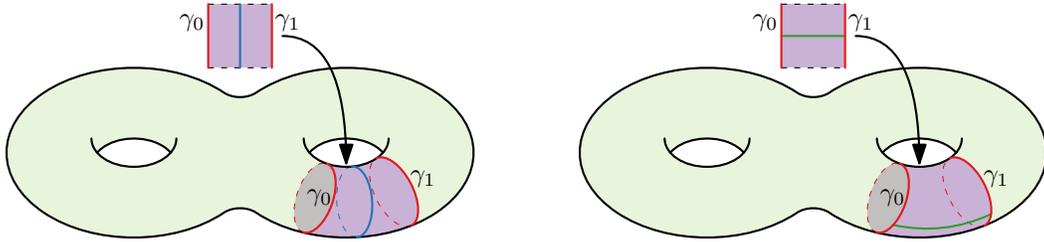


Figure 1: Left: the height of a homotopy between homotopic curves γ_0 and γ_1 measures the maximum amount an intermediate curve must stretch during the homotopy. Homotopies minimizing this amount of stretch measure the homotopy height. Right: the width of a homotopy measures the maximum length of a “slice” of the homotopy connecting the two boundary curves. Homotopies minimizing the length of this slice measure the homotopy width, also known as the homotopic Fréchet distance.

out to be a key ingredient for computational purposes, once we translate those results to the discretized setting. First, via some surgery arguments, it allows us to prove that HOMOTOPY HEIGHT is in **NP** (Theorem 5.1). The corresponding exponential time algorithm is to our knowledge the best exact algorithm for HOMOTOPY HEIGHT. We note that our setting is very general, and also implies **NP**-membership for a variant of HOMOTOPY HEIGHT in a more restricted setting that was considered in earlier papers [3, 9, 22], as well as for HOMOTOPIC FRÉCHET DISTANCE, where this was still open despite the recent articles investigating this distance [7, 22]. Then, further surgery arguments allow us to provide an $O(\log n)$ -approximation algorithm for HOMOTOPY HEIGHT (Corollary 6.2), by relying on an earlier $O(\log n)$ approximation-algorithm of Har-Peled, Nayeri, Salvatipour and Sidiropoulos [22] for homotopy height in a more restricted setting. Finally, we show that monotonicity directly implies an equivalence between the HOMOTOPY HEIGHT problem and a seemingly unrelated graph drawing problem which we call MINIMAL HEIGHT LINEAR LAYOUT. Therefore, we obtain that this problem is also in **NP** and we provide an $O(\log n)$ approximation for it.

1.2 Related work

Optimal homotopies (for several definitions of optimal) have been studied extensively in the mathematical community, mainly in Riemannian settings. This literature fits broadly in the setting of quantitative homotopy theory, initially introduced by Gromov [20], which aims at introducing a quantitative lens in the study of topological invariants on manifolds. Probably the most extensively considered notion of optimality is the study of homotopies minimizing the area swept; see for example [25] for an overview of some variants of this problem, or [30] for a discussion of how minimum area homotopies and homologies are connected in higher dimensions. The notion of controlling the width of a homotopy

has also been studied [5, 23], and more recent work on minimal height homotopies [10, 11] laid the foundation for the results in this paper.

On the computational side, the rise of Fréchet distance for measuring similarity between curves was a prime motivation for the notion of comparing two curves; see for example [1] for a survey. Generalizing the Fréchet distance to curves on surfaces led to the homotopic Fréchet distance, which is essentially the same as finding a minimum width homotopy given two input cycles on a surface. Polynomial time algorithms are known for the special case where the two input curves lie in the plane minus a set of obstacles [8]. Approximation algorithms exist for discrete settings where the two curves bound a disk [22].

More directly, minimum height homotopies have been studied from the computational perspective in various discretized settings [9, 22], although mainly to discuss the complexity of the problem. Indeed, as it was not known if the optimal height homotopy was even monotone, the complexity of the problem was completely open. Since the monotonicity result also holds in more geometric settings [6], a recent paper also examined one natural geometric setting, where the goal is to morph across a polygonal domain in Euclidean space with point obstacles; this work presents a lower bound that is linear in the number of obstacles, as well as a 2-approximation for the arbitrary weight obstacles and an exact polynomial time algorithm when all obstacles have unit weight [4]. The same problem also arises as a combinatorics question in lattice theory as a *b-northward migration*, where the authors leave monotonicity of such migrations as an open question [3].

1.3 Relations to graph searching and width parameters

This work also connects to sweep and search parameters in graph theory; see for example [18] for a survey of this topic. In each variant, the game consists of finding the

minimum number of searchers needed, where the goal is to find or isolate a hidden fugitive. For example, in the node searching variant, the fugitive hides on edges, all of which are originally contaminated, and the searchers clear an edge if two are on its incident vertices. In this variant, edges can be recontaminated if they are connected to a contaminated edge by a path without searchers, and the game ends when everything is decontaminated.

One key issue in these games is precisely that of monotonicity, or of determining whether in an optimal strategy, edges get recontaminated. In the node searching variant, monotonicity was established by Lapauh [24], and the argument was simplified by Bienstock and Seymour [2]. One important corollary to monotonicity for these games is that it immediately shows the problem lies in NP, since a strategy can be certified by the list of edges cleared.

Our homotopy problem is quite similar to these graph parameters; sweeping a disk while keeping the length small is intuitively quite similar to blocking in a fugitive. While our problem does display minor technical differences with the aforementioned variant – most notably, our setting is naturally edge-weighted and the cost is measured on the edges and not the vertices – the key difference is the one of *connectedness*, as node-searching games may allow for disconnected strategies. An important variant of node searching, called *connected node searching*, requires additionally that the decontaminated space remains connected, but makes no restriction on the uncontaminated space.

For graph searching problems, the main argument to establish monotonicity does not maintain connectivity [2], and it was proven that an optimal strategy for connected node searching may indeed be non-monotone [31]. By contrast, Theorem 3.4 establishes monotonicity of the optimal homotopy in our setting, and the arguments differ radically from the ones of Lapauh and Bienstock and Seymour. As such, we identify in this paper a new variant of graph searching which is somewhat tractable (i.e., in NP) and introduce a new proof technique to establish monotonicity results.

Finally, when monotonicity is established, graph searching parameters are very intimately related to width parameters of graphs. Minimum cut linear arrangement (also known as cut-width) is closely connected to the MINIMUM HEIGHT LINEAR LAYOUT problem, which we show to be equivalent to HOMOTOPY HEIGHT, but the key difference is that it may break the embedding of the graph. Thus, NP-hardness reductions for this problem [27] do not imply hardness for our problem. Connected variants of various width parameters give rise to *connected pathwidth* [14] and

connected treewidth [19], but in contrast to our homotopies, these parameters are only connected “on one side”, which makes them incomparable. We believe that the “doubly-connected” aspect of homotopy height makes it a worthwhile new graph parameter which could offer insights to other parameters in this area.

Outline of the paper. After introducing the preliminaries in Section 2, we lay the foundations of this work by explaining the structural theorems we rely on in Section 3. In Section 4 we establish surgery lemmas based on the idea of *retractions*. Then, in Section 5 we prove that HOMOTOPY HEIGHT is in NP. In Section 6 we draw connections with HOMOTOPIC FRÉCHET DISTANCE, and we leverage on these connections to provide an $O(\log n)$ -approximation algorithm for HOMOTOPY HEIGHT.

2 Preliminaries

Homotopy and Isotopy. Let Σ be a surface, endowed with a cellularly embedded graph G with n vertices such as in Figure 2, and let γ_0 and γ_1 be two simple cycles on G bounding an annulus.

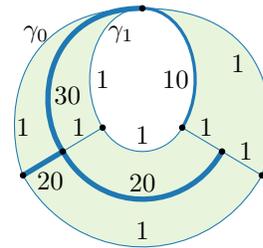


Figure 2: Example instance G , based on an example in [3].

A *discrete homotopy* h between γ_0 and γ_1 is a sequence of cycles $h(t_i)$ with $0 = t_0 \leq \dots \leq t_i \leq \dots \leq t_m = 1$, with $h(t_0) = \gamma_0$ and $h(t_1) = \gamma_1$ and any two consecutive paths $h(t_i)$ and $h(t_{i+1})$ are one *move* apart. The intermediate curves $h(t)$ are called *level curves* or *intermediate curves*. A move is either a face-flip, an edge-spike or an edge-unspike (flip, spike or unspike, for short). A *face-flip* for a face F replaces a single subpath p of $h(t_i) \cap \partial F$ with the path $\partial F \setminus p$ in $h(t_{i+1})$. An *edge-spike* for an edge $u \rightarrow v$ replaces a single occurrence of a vertex $u \in h(t_i)$ by the path $u \rightarrow v \rightarrow u$ consisting of two mirrored copies of that edge in $h(t_{i+1})$. Symmetrically, an *edge-unspike* replaces a path $u \rightarrow v \rightarrow u$ of $h(t_i)$ by the single vertex u in $h(t_{i+1})$. The *length* $\ell(h(i))$ of a path $h(i)$ is the sum of the weights of its edges (with multiplicity). The *height* of a homotopy h is the length of the longest path $h(t_i)$. An *optimal homotopy* is one that

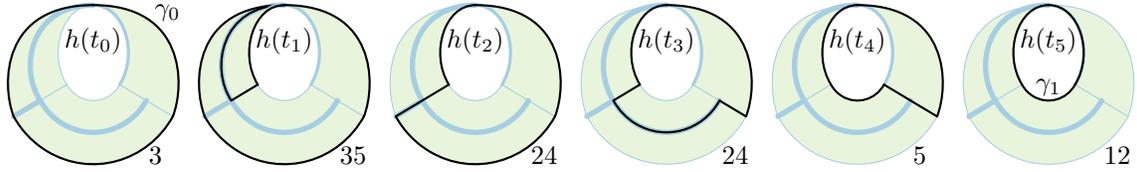


Figure 3: An optimal homotopy h of height 35 for the instance of Figure 2.

minimizes the height. The *homotopy height* between γ_0 and γ_1 is the height of an optimal homotopy between γ_0 and γ_1 . Figure 3 illustrates an optimal homotopy that uses only face-flips for the instance of Figure 2.

Cross Metric Surfaces. For most purposes, it is more convenient to think of this discrete model in a dual way, relying on the *cross-metric surfaces* [12] which are becoming increasingly used in the computational geometry and topology literature. In this dual setting, a cross-metric surface is a surface Σ endowed with a weighted (dual) graph G^* .

Assuming the primal surface is connected, we obtain this dual graph by gluing a disk to each boundary component, taking the dual graph, and puncturing the vertices corresponding to the added disks, without removing the adjacent edges. Such that these (dual) edges end at the boundary of the cross-metric surface instead of at a vertex, see Figure 4.

For a curve γ on Σ with a finite number of crossings with G^* , its length $\ell(\gamma)$ is the weighted sum of the crossings $\gamma \cap G^*$. Now, a homotopy between γ_0 and γ_1 is a homotopy in the usual sense, that is, a continuous map $h: S^1 \times [0, 1] \rightarrow \Sigma$ such that $h(\cdot, 0) = \gamma_0$ and $h(\cdot, 1) = \gamma_1$, except that we require that the values of t for which $h(\cdot, t)$ is not in general position with G^* are isolated, and each such curve has at most one such degeneracy¹ $h(x, t)$ with G^* . As before, the height of a homotopy is defined as the maximal length of an intermediate non-degenerate level curve $h(t)$. A

¹Any homotopy can be made so by a small perturbation without increasing the height, so we always consider this hypothesis fulfilled in the remainder of the article.

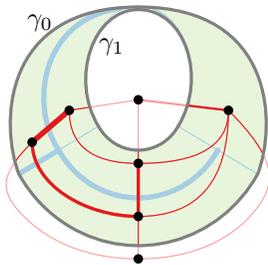


Figure 4: Dual representation of Figure 2.

homotopy is an *isotopy* if all the intermediate curves are simple.

Given a homotopy h^* in this setting, we obtain a discrete homotopy h on the primal graph G on Σ as follows. Pick a curve $h^*(t_i)$ in each maximal interval of non-degenerate curves in h^* (all curves in such interval have the same crossing pattern with G^* , and therefore the same length). Let $h(t_i)$ be the curve on G whose sequence of vertices and edges corresponds to the sequence of faces and edges of G^* visited by $h^*(t_i)$. This model is dual to the previous one, and Figure 5 illustrates how any move (flip, spike or unspike) connects two intermediate curves $h(t_i)$ and $h(t_{i+1})$. We say a discrete homotopy is an isotopy if it can be obtained from an isotopy in the dual setting.

3 Isotopies and monotonicity of optimal homotopies

We begin by restating and explaining the two structural results that we will rely on. Introducing the relevant Riemannian background lies outside of the scope of this paper, so we will simply advise the uninitiated reader to picture a Riemannian surface as a surface embedded into \mathbb{R}^3 , where the metric on the surface is the one induced by the usual Euclidean metric of \mathbb{R}^3 . Thanks to the Nash-Kuiper embedding theorem (see [21]), this naive idea loses no generality. We refer to standard textbooks on the subject for more proper background on Riemannian geometry, for example do Carmo [16].

The first theorem shows that up to an arbitrarily small additive factor, the homotopy of minimal height between two simple closed curves can be assumed to be an isotopy.

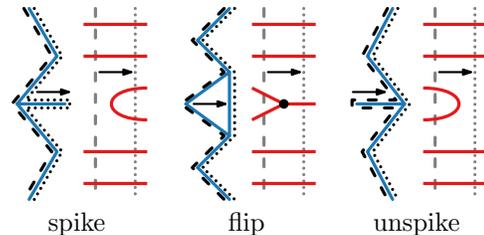


Figure 5: Three moves in the primal (left) and dual (right) representation.

THEOREM 3.1. ([10, THEOREM 1.1]) *Let Σ be two-dimensional Riemannian manifold with or without boundary, and let γ_0 and γ_1 be two non-contractible simple closed curves which are homotopic through curves bounded in length by L via a homotopy γ . Then for any $\varepsilon > 0$, there is an isotopy $\tilde{\gamma}$ from γ_0 to γ_1 through curves of length at most $L + \varepsilon$.*

Remark. The non-contractibility hypothesis is required because if M is not a sphere, contractible cycles with opposite orientations are homotopic but not isotopic. However, if we disregard the orientations, the result holds in full generality.

This theorem has the following discrete analogue:

THEOREM 3.2. *Let (Σ, G^*) be a cross-metric surface, and let γ_0 and γ_1 be two non-contractible simple closed curves on (Σ, G^*) which are homotopic through curves bounded in length by L via a homotopy γ . Then there is an isotopy $\tilde{\gamma}$ from γ_0 to γ_1 through curves of length at most L .*

The proof is exactly the same as the one of Theorem 3.1, except that it does not need the ε -slack: this was required to slightly perturb the curves so that they are simple but in the discrete setting it can be done with no overhead.

The second theorem shows that, *when the starting and finishing curves of a homotopy are the boundaries of the manifold*, there exists an optimal homotopy that is monotone, i.e., that never backtracks, once again up to an arbitrarily small additive factor. Formally, if γ is an isotopy (which we can assume the optimal homotopy to be, by Theorem 3.1) between γ_0 and γ_1 , for $0 \leq t \leq 1$, the curves γ_t and γ_1 bound an annulus A_t . Then the isotopy γ is **monotone** if for $t < t' < 1$, $\gamma_{t'}$ is contained in A_t .

THEOREM 3.3. ([6]) *Let M be a Riemannian annulus with boundaries γ_0 and γ_1 such that there exists a homotopy between γ_0 and γ_1 of height less than L . Then there exists a monotone homotopy between γ_0 and γ_1 of height less than L .*

Note that the ε -slack of Theorem 3.1 is also present here but is hidden in the open upper bound on the height. In this theorem, as was observed by Chambers and Rotman [11], crediting Liokumovitch, the hypothesis that the manifold is entirely comprised between both curves is necessary: see [11, Figure 5] for a counterexample.

In the discrete setting, the corresponding result is the following, where the definition of monotonicity is the same:

THEOREM 3.4. *Let (Σ, G^*) be a cross-metric annulus with boundaries γ_0 and γ_1 such that there exists a homotopy between γ_0 and γ_1 of height L . Then there exists a monotone isotopy between γ_0 and γ_1 of height L .*

The proof is exactly identical to the one in the Riemannian setting and it yields a slightly stronger result, since the cross-metric setting removes the need for perturbations and thus the need of an ε -slack.

Remark. Let us observe that the discrete theorems are in some way more general than the Riemannian ones: not only do they bypass the need for some ε -slack, but they also directly imply their Riemannian converses by the following reduction. Starting with a Riemannian surface, and a (non-monotone) isotopy between two disjoint curves, one can find a triangulation of the surface allowing, at an ε -cost, to approximate the isotopy using only elementary moves. Then, after making this isotopy monotone in the discrete setting, one can translate it back into a monotone isotopy in the Riemannian setting by interpolating between the face and edge moves.

4 Retractions and pausing at short cycles

In this section, we establish several technical lemmas which are necessary for our proofs in the next section. For simple closed curves β and γ bounding an annulus, denote that annulus by $A(\beta, \gamma)$. Let $\mathcal{S}(\beta, \gamma)$ be the set of closed curves in $A(\beta, \gamma)$ homotopic to boundaries β and γ , that do not intersect homotopic curves of shorter length. Then, for any point $p \in \alpha \in \mathcal{S}(\beta, \gamma)$, α is a shortest closed path through p in its homotopy class. Let $\mathcal{G}(\beta, \gamma)$ be the set of minimum length simple closed curves homotopic to the boundaries of $A(\beta, \gamma)$, then $\mathcal{G}(\beta, \gamma) \subseteq \mathcal{S}(\beta, \gamma)$.

We now introduce the concept of a retraction of a homotopy, which gives a way to shortcut a homotopy at a given curve, provided it is a curve of $\mathcal{S}(\beta, \gamma)$. This idea is implicit in Chambers and Rotman [11, Proof of Theorem 0.7], and we refer to their article for more details. For a monotone isotopy h between boundaries of an annulus A , and a homotopic annulus $A' \subset A$, define the **retraction** $h|^{A'}(t)$ of $h(t)$ to A' as the same curve with each arc of $h(t) \setminus A'$ replaced by the shortest homotopic path along the boundary of A' . Although paths along $\partial A'$ (dis)appear discontinuously as t varies, $h|^{A'}$ can be obtained in the form of a discrete homotopy by (un)spiking these paths as they (dis)appear. The resulting homotopy $h|^{A'}$ is a monotone isotopy.

LEMMA 4.1. *If $\alpha \in \mathcal{S}(\alpha, \gamma)$ and $A(\alpha, \gamma) \subseteq A(\beta, \gamma)$, and h is a monotone isotopy from β to γ of height L ,*

then $h|^{A(\beta,\alpha)}$ is a monotone isotopy from β to α with height at most L .

Proof. The retraction $h' = h|^{A(\beta,\alpha)}$ is a monotone isotopy from $h'(0) = \beta$ to $h'(1) = \alpha$. Let t' be the maximum t for which $h(t)$ intersects $A(\beta,\alpha)$. For $t \geq t'$, we have $h'(t) = \alpha$ and therefore $|h'(t)| = |\alpha| \leq |h(t)| \leq L$. For $t \leq t'$, each arc a of $h(t) \setminus A(\beta,\alpha)$ is replaced in $h'(t)$ by a homotopic path b along α with $|b| \leq |a|$, and thus $|h'(t)| \leq |h(t)| \leq L$. Hence $\text{height}(h') \leq L$.

LEMMA 4.2. *If $\alpha \in \mathcal{S}(\beta,\gamma)$, and h is a monotone isotopy from β to γ of height L , then there is a monotone isotopy from β to γ of height at most L having α as a level curve.*

Proof. We have $\alpha \in \mathcal{S}(\alpha,\beta)$ and $\alpha \in \mathcal{S}(\alpha,\gamma)$. So by Lemma 4.1, the monotone isotopies $h|^{A(\beta,\alpha)}$ from β to α and $h|^{A(\alpha,\gamma)}$ from α to γ have height at most L and can be composed to obtain a monotone isotopy from β to γ of height at most L with α as a level curve.

LEMMA 4.3. *Let $\Pi = \{\pi_1, \dots, \pi_m\}$ be a set of paths from γ_0 to γ_1 without proper pairwise intersections, where each π_i is a shortest homotopic path in $A(\gamma_0, \gamma_1)$ between its endpoints. If h is a monotone isotopy from γ_0 to γ_1 of height L , then there exists a monotone isotopy of height at most L whose level curves all cross each π_i at most once (after infinitesimal perturbations).*

Proof. Denote by $c(a,b)$ the number of proper intersections of curves a and b , and by $c_{\Pi}(a) = \sum_{\pi \in \Pi} c(a, \pi)$ the total number of intersections of a with Π . Let $C_h = \max_t c_{\Pi}(h(t))$ be the maximum total number of intersections over all t , and let I_h be the set of maximal intervals (τ_0, τ_1) with $c_{\Pi}(h(t)) = C_h$ if $t \in (\tau_0, \tau_1) \in I_h$. If $C_h = m$, each level curve of h crosses each π_i exactly once and we are done, thus we assume in the following that $c_{\Pi}(h(0)) = c_{\Pi}(h(1)) = m < C_h$.

If $C_h > m$, we obtain a homotopy h' from h with $C_{h'} < C_h$ by, for each interval $(\tau_0, \tau_1) \in I_h$, replacing subhomotopy $h|_{(\tau_0, \tau_1)}$ of h by some $h^* = h'|_{(\tau_0, \tau_1)}$ with $C_{h^*} < C_h$.

Consider a single interval $(\tau_0, \tau_1) \in I_h$ and let $A = A(h(\tau_0), h(\tau_1))$. Then $\Pi \cap A$ consists of C_h subarcs of Π , each connecting the two boundaries of A . For $t \in (\tau_0, \tau_1)$, $h(t)$ intersects each such arc exactly once, and each $h(t)$ intersects these arcs in the same order. Among the components of $A \setminus \Pi$, there is a disk D_0 bounded by one arc of $h(\tau_0)$ and two arcs of $\pi_i \cap A$, and a disk D_1 bounded by one arc of $h(\tau_1)$ and one arc of π_j , such that these disks contain no other arcs of Π .

We can find $\alpha \in \mathcal{G}(h(\tau_0), h(\tau_1))$ that intersects any arc of $A \cap \Pi$ at most once (in the same order

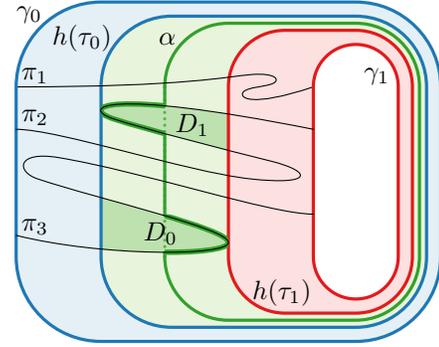


Figure 6: Choosing α such that $C_{h^*} < C_h$.

as $h(t)$), and does not intersect the interiors of D_0 and D_1 (because the two arcs of Π on their boundary form a shortest path). Then $c_{\Pi}(\alpha) < C_h$ and the retraction $h_0 = h|^{A(h(\tau_0), \alpha)}$ has $C_{h_0} < C_h$, since any arc $h_0(t)$ has fewer intersections than $h(t)$ has with Π (in particular with the boundary of D_1). Symmetrically, for $h_1 = h|^{A(\alpha, h(\tau_1))}$ we have $C_{h_1} < C_h$. Since the composition $h^* = h_0 h_1$ is a homotopy from $h(\tau_0)$ to $h(\tau_1)$ with $C_{h^*} < C_h$ and height at most L (by Lemma 4.2), we can use this as a replacement for $h|_{(\tau_0, \tau_1)}$ in h' . By induction, we obtain a homotopy of height at most L whose level curves all cross each π_i at most once.

5 Computing homotopy height in NP

In this section, we show that in the discrete setting, there is an optimal homotopy with a polynomial number of moves. First, we show that there is a homotopy that flips each face exactly once.

LEMMA 5.1. *For an annulus (Σ, G) bounded by γ_0 and γ_1 , there is a homotopy of minimum height between γ_0 and γ_1 that flips each face of G exactly once.*

Proof. By Theorem 3.4, some homotopy h of minimum height is a monotone isotopy. For two consecutive level curves $h(t)$ and $h(t')$ in a monotone isotopy, the move between $h(t)$ and $h(t')$ flips face F if and only if F lies in $A(h(t'), \gamma_1)$ or $A(h(t), \gamma_1)$ but not both. Because $A(h(0), \gamma_1)$ contains all faces, and $A(h(1), \gamma_1)$ contains none, each face is flipped at least once. By monotonicity, we have for $0 \leq t' < t \leq 1$, that $A(h(t'), \gamma_1) \supseteq A(h(t), \gamma_1)$. So, if face F does not lie in $A(h(t), \gamma_1)$, it will not be flipped again in $h|_{[t, 1]}$. Hence each face is flipped exactly once.

It remains to show that each edge is involved in a polynomial number of (un)spike moves; note that this does not directly follow from monotonicity, since a second spike of the same edge does not violate monotonicity (as can easily be seen in the dual setting).

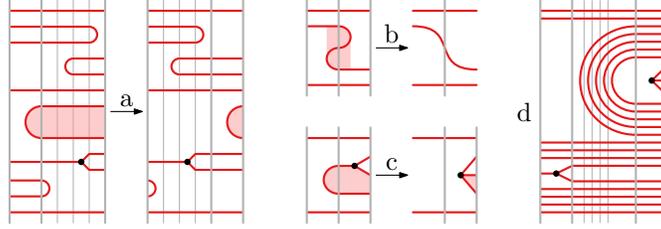


Figure 7: Delaying spikes (a). Canceling spikes with unspikes (b) or faces (c). Part of a reduced isotopy (d).

Postponing spikes. Before we bound the number of spike moves, we transform an optimal monotone isotopy h into one where each spike move is delayed as much as possible, and each unspike move happens as soon as possible. We explain this transformation in the dual setting.

Suppose a spike move occurs for edge e between $h(t_i)$ and $h(t_{i+1})$, then denote by s the (unique) arc of $A(h(t_i), h(t_{i+1})) \cap G^*$ both of whose endpoints lie on $h(t_{i+1})$. This arc is a subarc of the dual edge e^* . Consider the maximum $j > i$, for which the component s_j of $e^* \cap A(\gamma_0, h(t_j))$ containing s has both endpoints on $h(t_j)$, and for all $t_i < t \leq t_j$, curve $h(t)$ has exactly two crossings with s_j (so the only action performed on arc s_j was the spike between $h(t_i)$ and $h(t_{i+1})$). Then s_j and $h(t_j)$ enclose a disk D_j . If the interior of D_j contains no edges of G^* , we can delay the spike of e at least until just before t_j , as illustrated in Figure 7 (a), where D_j is shaded.

Depending on what happens in the move between $h(t_j)$ and $h(t_{j+1})$, we may transform the isotopy further. This move is either (1) an unspike attached to s_j , or (2) a face-flip connected to one endpoint or (3) both endpoints² of s_j , or (4) a face-flip or spike inside D_{j+1} . In cases (1) and (2), we cancel the spike against the unspike or flip, as illustrated in Figure 7 (b) and (c). We do not postpone the spike in cases (3) and (4). Symmetrically, unspike moves can be made to happen earlier. Observe that these operations cannot increase the height of a homotopy since each level curve in the resulting homotopy crosses a subset of the edges of some curve in the original homotopy.

Call a homotopy *reduced* if it is the result of applying the above rules to h until no spike can be canceled or postponed until after a flip or unspike, and no unspike can be canceled or be made to happen before any prior flip or spike. Starting from an optimal monotone isotopy, the reduced isotopy is also an optimal monotone isotopy. Lemma 5.2 captures a structural property of reduced homotopies.

²This happens only if the primal edge is adjacent to only one face of G .

LEMMA 5.2. *Between any two consecutive face-flips in a reduced isotopy lies a single (possibly empty) path of unspike moves followed by a (possibly empty) path of spiked moves.*

Proof. In a reduced homotopy, no unspike follows a spike move, and any spikes that remain ‘surround’ the next face-flip (if any), see Figure 7 (d). Symmetrically, all unspikes between two consecutive face-flips surround the previous face-flip (if any). From the primal perspective, these unspike moves form a path from the previously flipped face and spike moves form a path towards the next flipped face.

Any reduced homotopy starts with zero or more unspikes from γ_0 , after which a possibly empty path of spikes to the first face-flip occurs, then that face is flipped, and a possibly empty path of unspikes enabled by this flip occurs. Subsequently, a spiked path, face-flip, and unspiked path occur for the remaining faces. Finally, a sequence of spikes towards γ_1 may occur. We may assume that on γ_0 and γ_1 , any two consecutive edges are different, such that no immediate unspike moves are possible from γ_0 , and no immediate spike moves are possible to γ_1 . Otherwise we may by Lemma 4.1 perform those moves immediately without increasing the homotopy height.

Bounding spike moves. We are now ready to bound the number of spike and unspike moves in an optimal homotopy. Call a homotopy h *good* if it is a minimum-height reduced monotone isotopy and it has a minimum number of moves. By Theorems 3.2 and 3.4, the height of h is the homotopy height between γ_0 and γ_1 .

Define an edge-spike of an edge e to be *between* existing copies of e , if the portion of the dual edge e^* crossed by the (dual) level curve, lies between two existing crossings of the level curve with e^* , such as in Figure 8. We show that such spikes never appear in h .

LEMMA 5.3. *If homotopy h is good, there are no spikes between existing copies of any edge e .*

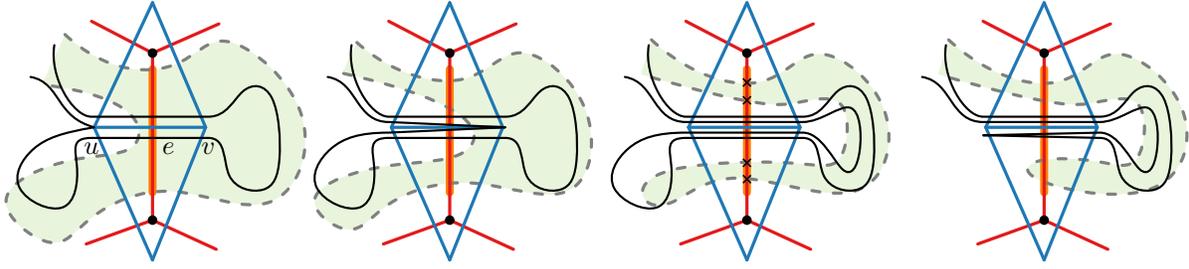


Figure 8: Development of a spike between existing copies of e . Part of the graph in red (dual) and blue (primal) and the level curve in gray dashed (dual) and black (primal, perturbed).

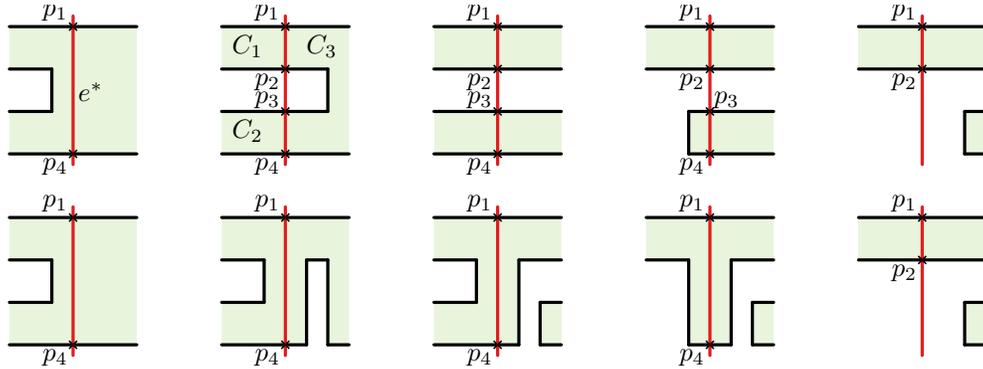


Figure 9: Top: the neighborhood of e^* throughout h . Bottom: the reconnected homotopy, reducing crossings with e^* . From left to right: the homotopy just before τ_0 , just after τ_0 , between τ_0 and τ_1 , just before τ_1 , and just after τ_1 .

Proof. Suppose the move from $h(t_i)$ to $h(t_{i+1})$ is the last move between existing copies of the same edge, and assume this move is a spike of edge $e = (u, v)$ from u to v . In the dual setting, consider the component π of $e^* \cap A(h(t_i), \gamma_1)$ that is crossed by the spike move (highlighted in Figure 8). Let $c(t)$ be the number of crossings of $h(t)$ with π , then for some τ_0 between t_i and t_{i+1} , $c(\tau_0) = 3$, and for some unique $\tau_1 > \tau_0$, $c(\tau_1) = 3$ again, and for $\tau_0 < t < \tau_1$, we have $c(t) = 4$ (because we assumed this was the last spike between existing copies of any edge).

For $\tau_0 < t < \tau_1$, label the four crossings of $h(t)$ with π by $p_1(t)$, $p_2(t)$, $p_3(t)$, and $p_4(t)$, in order along e^* , so the spike move at τ_0 creates p_2 and p_3 . Consider the three components $C_1(t)$, $C_2(t)$ and $C_3(t)$ of $A(h(t), \gamma_1) \setminus \pi$, such that C_1 touches p_1 and p_2 from the dual face of u , and C_2 touches p_3 and p_4 from the dual face of u , and C_3 touches e^* in two segments from the dual face of v . Because C_3 lies between C_1 and C_2 , h will first contract either component C_1 or C_2 , namely at $h(\tau_1)$. Assume without loss of generality that C_2 contracts first.

We modify $h|_{[\tau_0, \tau_1]}$ such that any level curve crosses π at most twice by reconnecting the neighbor-

hood of π , whose local structure evolves exactly as depicted in the top row of Figure 9. We essentially remove crossings p_2 and p_3 , and reconnect $\partial C_1(t) \cap h(t)$ with $\partial C_2(t) \cap h(t)$ using a (zero-length) segment along π in face u^* . On the other side, consider the arc of $\partial C_3(t) \cap h(t) \cap v^*$ with $p_4(t)$ as endpoint. We cut this arc in two subarcs a and b , where a has $p_4(t)$ as endpoint, and connect the other endpoint to the arc of $\partial C_3(t) \cap h(t)$ at the endpoint at $p_2(t)$ using a segment along π in v^* . Similarly, we connect the endpoint of that at $p_3(t)$ to the loose end of b . These reconnections are depicted in the bottom row of Figure 9. A more global view (corresponding to Figure 8) is illustrated in Figure 10.

Observe that the reconnected curves can be made to appear continuously in such a way that they form a monotone isotopy. Because level curves only changed in the neighborhood of π , where they were shortened by avoiding the crossings with π , we have an isotopy whose height is at most that of h , and in which at least one spike is removed. So, because h was optimal, we have constructed an optimal monotone isotopy with fewer moves. Therefore, the corresponding reduced isotopy also has fewer moves, contradicting that h was good.

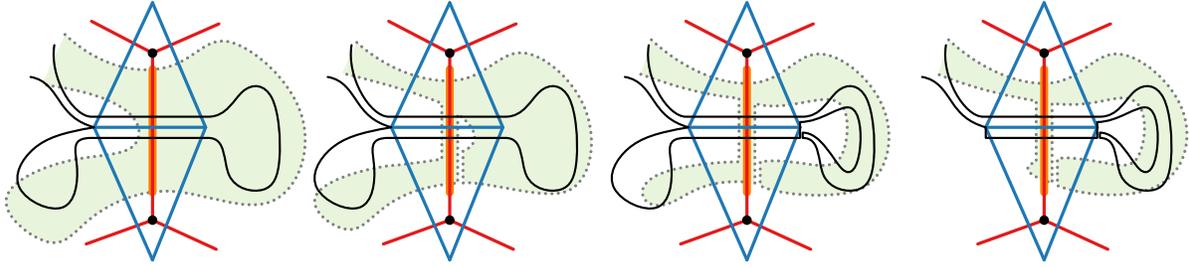


Figure 10: Figure 8 after a local surgery that avoids the spike between copies of e .

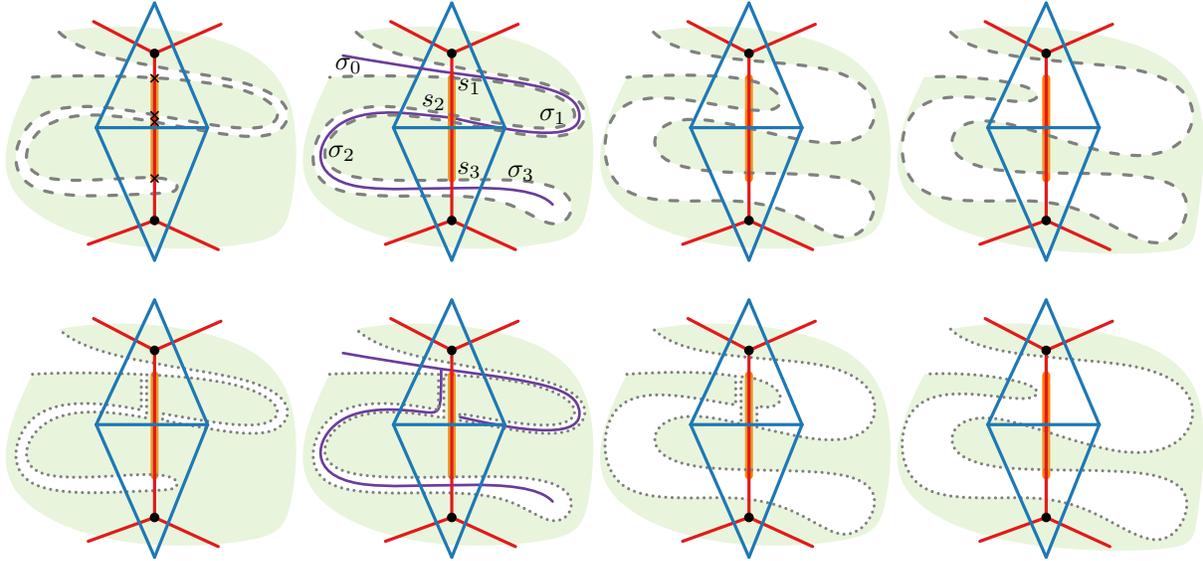


Figure 11: A local surgery to avoid five spikes of the same edge on a single spiked path.

Our final step towards bounding the number of edge spikes is to derive a contradiction if for some interval $[\tau_0, \tau_1]$ without face-flips, an edge e is spiked (or unspiked) 5 times in $h|_{[\tau_0, \tau_1]}$. The proof is similar to that of Lemma 5.3.

LEMMA 5.4. *For a good homotopy h , any subhomotopy $h|_{[\tau_0, \tau_1]}$ contains either a face-flip, or at most 4 spike (and at most 4 unspike) moves of the same edge.*

Proof. Suppose $h|_{[\tau_0, \tau_1]}$ contains no face-flip, then because h is reduced, the spike moves in $h|_{[\tau_0, \tau_1]}$ form a path σ of spike moves in G . Assume for a contradiction that some edge $e = (u, v)$ lies on σ at least 5 times. We say two spikes s_1 and s_2 are consecutive on e^* if no spike occurs on the arc of e^* between the first crossing of s_1 with e^* and the first crossing of s_2 with e^* .

Because by Lemma 5.3, h does not contain spikes between existing copies of edges, we can find three spikes s_1, s_2 and s_3 of e on σ where s_1 and s_2 as well as s_2 and s_3 are consecutive on e^* , and s_1 happens before s_2 and s_2 happens before s_3 . Let $\sigma_0, \sigma_1, \sigma_2$ and σ_3

be the subpaths of σ such that $\sigma = \sigma_0 s_1 \sigma_1 s_2 \sigma_2 s_3 \sigma_3$, also labeled in Figure 11.

To get rid of spike s_2 , we connect $\sigma_0 s_1 \sigma_1$ to $\sigma_2 s_3 \sigma_3$ in an alternative way. Figure 12 illustrates all possible ways s_1, s_2 and s_3 (in the dotted area) can be connected by σ , and how our method will reconnect σ without s_2 . Formally, to decide where this reconnection takes place, we consider the components of $A(h(\tau_1), \gamma_1) \setminus \pi$, where π is the arc of e^* between its intersections with s_1 and s_3 . There are three components, component C_1 touching π and σ_1 , component C_2 touching π and σ_2 , and component C_3 touching σ entirely, and touching π in two arcs. The component that h contracts first is either C_1 or C_2 (since C_3 lies between the other two).

First consider the case where C_1 is contracted first, then the path $\sigma_2 s_3 \sigma_3$ starts in the dual face of the endpoint of s_2 . Note that there is a (zero-length) path between the start or endpoint of s_1 and the endpoint of s_2 because s_1 and s_2 are adjacent along e^* . Use this zero-length path to connect $\sigma_2 s_3 \sigma_3$ to $\sigma_0 s_1 \sigma_1$ at the start or endpoint of s_1 and call the resulting tree λ .

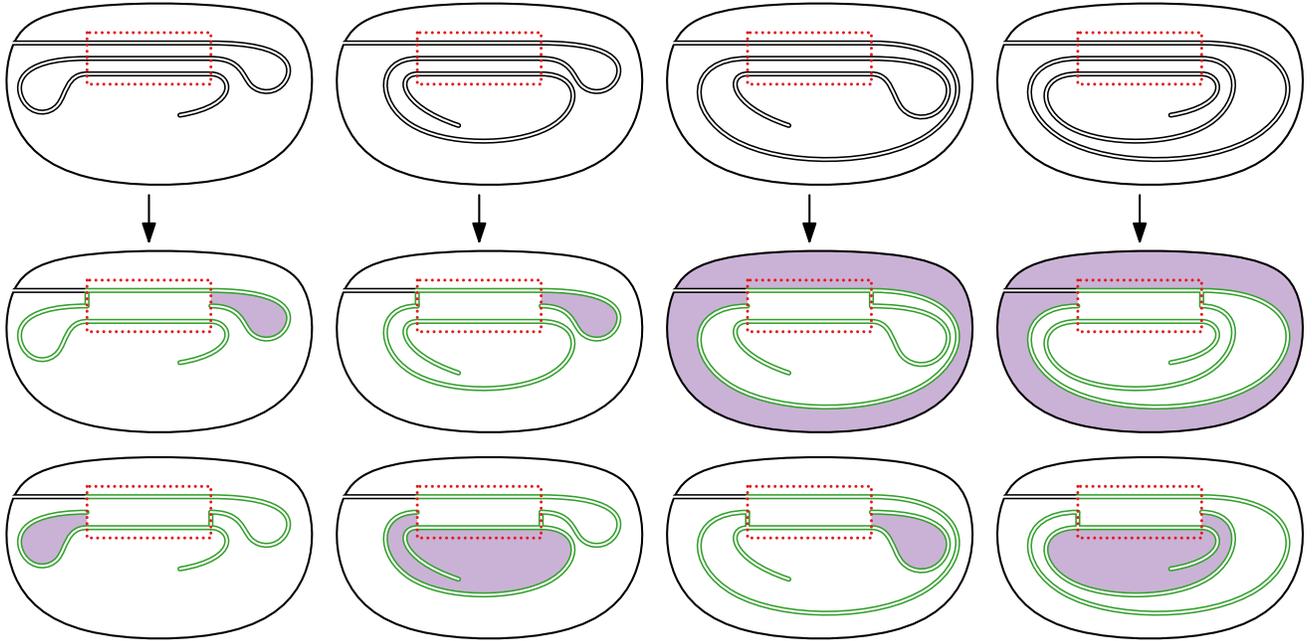


Figure 12: Cases for shortcutting spiked paths visiting the same edge often. The neighborhood of the repeated edge is dotted and the component contracted first is shaded.

We claim we obtain an optimal monotone isotopy h' from h by replacing the spiked path σ by the spiked tree λ , and removing the unspike move of e^* following the contraction of C_1 . Up until the creation of λ , the move sequence is the same as in h . Since λ contains a subset (all spikes except s_2) of the spikes of σ , the spiked tree can be created without surpassing the height of h . After the creation of σ in h and λ in h' , locally, the level curves of h and h' differ only in a small neighborhood of π , so that all moves of h except those crossing π can still be performed in h' . Because s_2 is the only spike along e^* that lies between s_1 and s_3 , the next move that crosses π is the unspike move, call it z , following the contraction of C_1 . The level curve of h' just before z is identical to the level curve of h just after z , so it is safe to omit move z in h' . All subsequent level curves of h and h' are identical, so we conclude that h' is an optimal monotone isotopy (with fewer moves). Therefore, the reduced monotone isotopy of h' has fewer moves, contradicting that h was good.

The proof for the case where C_2 contracts first, is symmetrical, except that the spiked tree λ is created differently. In this case, we define λ to be $\sigma_0 s_1 \sigma_1$, whose endpoint is connected to $\sigma_2 s_3 \sigma_3$ at the start or endpoint of s_3 . When spiking this tree, the direction of the spikes on σ_2 (and sometimes σ_3) is reversed, but this does not affect the proof.

Hence, in a good homotopy, no spiked path spikes

the same edge 5 times.

THEOREM 5.1. *For γ_0 and γ_1 bounding an annulus with n faces and m edges, there is a homotopy of minimum height that has at most $O(mn)$ moves. Therefore, deciding whether their HOMOTOPY HEIGHT is at most L is in **NP**.*

Proof. Let n be the number of faces, and m the number of edges in G . As a direct consequence of Lemmas 5.1 and 5.4, there is a good homotopy that spikes each edge at most $4(n+1)$ times and unspikes each edge at most $4(n+1)$ times. So there is a homotopy of minimum height that has at most $8m(n+1) + n = O(mn)$ moves. Testing whether this homotopy indeed has height at most L can be done by computing the maximum length over its (polynomially many) level curves, each containing a polynomial number of edges, and comparing this maximum with L . Given a good homotopy, all of this can be done in polynomial time assuming addition and comparisons of numbers takes polynomial time.

We note that the HOMOTOPY HEIGHT problem can also be defined in slightly different settings, for example

- γ_0 and γ_1 are two paths with common endpoints s and t , such that $\gamma_0 \cup \gamma_1$ is the boundary of a combinatorial disk. Then γ_0 is homotopic to γ_1 with

fixed endpoints, and we are interested in computing the optimal height of this homotopy. This is the HOMOTOPY HEIGHT problem considered by E. Chambers and Letscher [9].

- There is a single curve γ forming the boundary of a combinatorial disk. This curve is contractible, and we are interested in computing the optimal height of such a contraction. This is one of the settings considered in [6].

In both these cases, the Theorems 3.2 and 3.4 have analogues establishing that some optimal homotopy is an isotopy and is monotone. The rest of our proof techniques then readily apply, and prove that the HOMOTOPY HEIGHT problem in these cases is also in **NP**. The next section investigates more distant variants.

6 Variants and approximation algorithms

6.1 Homotopic Fréchet distance

There is a strong connection between the problem of HOMOTOPY HEIGHT and the problem of HOMOTOPIC FRÉCHET DISTANCE, which we now recall. As in [22], our setting is the one of a disk D with four points p_0, q_0, q_1 and p_1 on the boundary, connected by four disjoint boundary arcs γ_0, γ_1, P and Q , with γ_0 from p_0 to q_0 ; γ_1 from p_1 to q_1 ; P from p_0 to p_1 ; and Q from q_0 to q_1 , see Figure 13, left. A homotopy between γ_0 and γ_1 is a series of elementary moves connecting curves of D with one endpoint on P and the other on Q , where the collection of curves starts at γ_0 and ends at γ_1 . The HOMOTOPIC FRÉCHET DISTANCE between P and Q is the height of a homotopy between γ_0 and γ_1 of minimal height. The common intuition for this distance is that it is the minimal length of a leash needed for a man on P to walk his dog along Q , where the leash may stretch but cannot be lifted out of the underlying space.

We note that this is slightly different than the original setting for homotopic Fréchet distance in the original work [8], where an exact algorithm is presented for the plane minus a set of polygonal obstacles. In the original work, the start and end leashes are not fixed, and in fact the bulk of the work is in determining an optimal relative homotopy class in order to find the best homotopy.

PROPOSITION 6.1. *The HOMOTOPIC FRÉCHET DISTANCE problem is in **NP**.*

Proof. We reduce HOMOTOPIC FRÉCHET DISTANCE to HOMOTOPY HEIGHT using the following construction. We add a vertex v and edges of weight K between this vertex and all the vertices of the paths P and Q , where K is a constant greater than the sum of the weights

of the edges of the disk, as well as all the intermediate triangles, see Figure 13, right. This results in a pinched annulus A , with two boundaries γ'_0 and γ'_1 obtained from the paths γ_0 and γ_1 , both completed into closed curves using the additional vertex v . We claim that an optimal homotopy between γ_0 and γ_1 translates into an optimal homotopy in A between γ'_0 and γ'_1 , and vice-versa. Indeed, by Lemma 4.3, there exists an optimal homotopy in A such that any intermediate curve crosses the shortest path between γ'_0 and γ'_1 exactly once, and in our case the shortest path is the zero length path starting and ending at the vertex v . Furthermore, if the weight K is big enough, the level curves of an optimal homotopy between γ'_0 and γ'_1 will always use exactly two of the edges of weight K , since two are needed but any more would be too expensive. Thus, an optimal homotopy between γ'_0 and γ'_1 translates directly into an optimal homotopy between γ_0 and γ_1 after cutting on v and removing the edges linked to v and vice-versa. The homotopy height is increased by exactly $2K$ in this translation.

Har-Peled, Nayyeri, Salvatipour and Sidiropoulos [22] provide an algorithm to compute in $O(n \log n)$ time a homotopy of height $O(d \log n)$, where d is a lower bound on the height of an optimal homotopy, and n is the complexity of Σ . In particular, one can set d to be the maximum of $\|\gamma_0\|, \|\gamma_1\|$, the diameter of Σ , and half of the total weight of the boundary of any face. This yields an $O(\log n)$ approximation for HOMOTOPIC FRÉCHET DISTANCE³. We show here that their algorithm can be adapted to yield an $O(\log n)$ approximation for HOMOTOPY HEIGHT.

PROPOSITION 6.2. *One can compute in $O(n \log n)$ time an $O(\log n)$ -approximation of HOMOTOPY HEIGHT.*

Proof. Starting with an annulus and two boundary curves γ_0 and γ_1 , we first compute a shortest path \mathcal{P} between the boundary curves γ_0 and γ_1 and cut along \mathcal{P} to obtain a disk D . This brings us to the setting of HOMOTOPIC FRÉCHET DISTANCE, and we can apply the aforementioned algorithm and obtain a homotopy h . In order to recover a homotopy between γ_0 and γ_1 , we glue back the disk along \mathcal{P} into an annulus, and the level curves of h are completed into closed curves by using subpaths of \mathcal{P} , this gives us a homotopy h' . It remains to show that this is an $O(\log n)$ approximation of the optimal homotopy. By Lemma 4.3, some optimal homotopy between γ_0 and γ_1 has level curves cutting

³This algorithm assumes triangular faces, but using our definition of d , one can extend the algorithm of [22] to also work with polygonal faces.

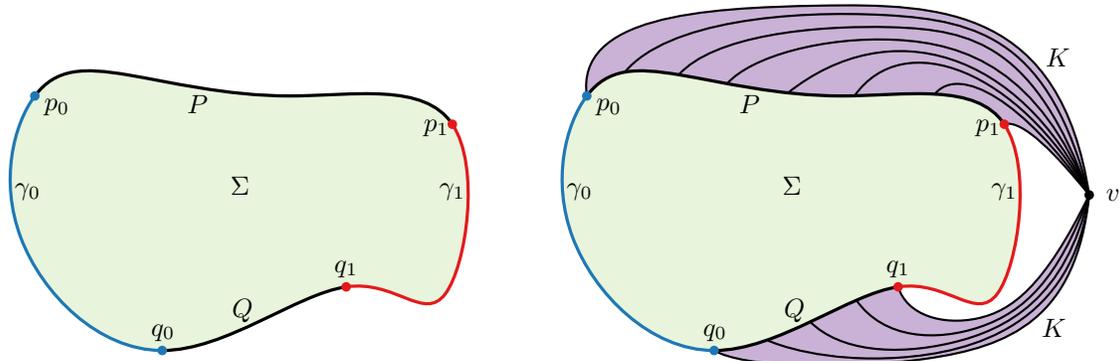


Figure 13: The setting of homotopic Fréchet distance.

\mathcal{P} exactly once. Thus, the height L of an optimal homotopy in the disk D is a lower bound for the height of an optimal homotopy in the annulus A . Furthermore, each level curve γ_t of h' consist of two subpaths, one being a level curve $h(t)$ of h and the other being a subpath \mathcal{P}'_t of \mathcal{P} . Since \mathcal{P} is a shortest path, \mathcal{P}'_t is also a shortest path between its endpoints, so it is shorter than $h(t)$ since they have the same endpoints. By construction, the length of $h(t)$ is $O(L \log n)$, and thus the length of γ_t is $O(2L \log n) = O(L \log n)$. This concludes the proof.

6.2 Minimum height linear layouts

We also show that a seemingly unrelated graph drawing problem is directly equivalent to the HOMOTOPY HEIGHT problem. A *linear layout* is an embedding of a planar graph where the edges have isolated tangencies with the vertical line, and all the vertices have distinct x coordinates. The MINIMUM HEIGHT LINEAR LAYOUT problem is the following one: Given a planar embedding of an edge-weighted graph G , find a homeomorphic linear layout of G in \mathbb{R}^2 such that the maximal weight of the vertical lines is minimized. Here, the weight of a vertical line is the sum of the weights of the edges that it crosses, and (similarly to the cross-metric setting), vertical lines crossing tangent to the edges or crossing vertices are not counted.

THEOREM 6.1. *The MINIMUM HEIGHT LINEAR LAYOUT problem is equivalent to the HOMOTOPY HEIGHT problem.*

Proof. Indeed, a linear layout of a planar graph G naturally induces a discrete homotopy sweeping its dual graph G^* . More formally, we drill a small hole around the vertex dual to the outer face of G , and we view its complement as a disk D which is a cross-metric surface for the graph G . Since the hole was drilled in the middle of the face of G , its boundary has zero

length. We pick two arbitrary vertices s and t on it, which cuts the boundary into two paths L and R . Then we claim that a minimum height linear layout of G is equivalent to a homotopy of minimum height between L and R (where the endpoints are fixed)⁴. Indeed, whenever the sweep of \mathbb{R}^2 induced by the vertical lines crosses an edge or passes a vertex, by the dual interpretation of homotopies with cross-metric surfaces outlined in the preliminaries, it amounts to doing a face or an edge move, and thus the whole vertical sweep defines a homotopy between the two paths L and R . Furthermore, this homotopy is an isotopy, since the vertical lines are simple, and a monotone one since they only go in a single direction. Conversely, a discrete homotopy of optimal height between L and R can be “straightened” into a linear layout: by Theorem 3.4, one can assume such a homotopy h to be an isotopy and to be monotone, and therefore the succession of dual moves of h with respect to G are homeomorphic to a sweep of G by vertical lines, as pictured in Figure 14. An optimal homotopy amounts, via this homeomorphism, to finding a linear layout of minimal weight.

In particular, the MINIMUM HEIGHT LINEAR LAYOUT problem is in NP and admits an $O(\log n)$ -approximation algorithm.

⁴The somewhat artificial construction with L and R forces the homotopy to go through the outer face of G at all times.

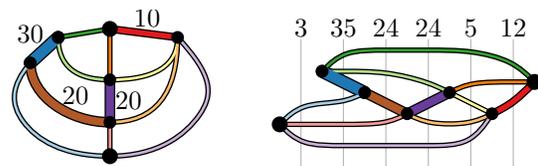


Figure 14: Dual representation of Figure 2 (left) and Figure 3 (right).

Acknowledgements. We are grateful to Tasos Sidiropoulos for his involvement in the early stages of this research, and to Gregory Chambers and Regina Rotman for helpful discussions. This research began while partially supported through the program “Simons Visiting Professorship” by the Mathematisches Forschungsinstitut Oberwolfach in 2015. Erin Chambers is supported in part by NSF grants IIS-1319944, CCF-1054779, and CCF-1614562. Arnaud de Mesmay is partially supported by the ANR project ANR-16-CE40-0009-01 (GATO). Tim Ophelders is supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 639.023.208.

References

- [1] H. ALT, *The computational geometry of comparing shapes*, in Efficient Algorithms, vol. 5760 of Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2009, pp. 235–248.
- [2] D. BIENSTOCK AND P. SEYMOUR, *Monotonicity in graph searching*, Journal of Algorithms, 12 (1991), pp. 239–245.
- [3] G. R. BRIGHTWELL AND P. WINKLER, *Submodular percolation*, SIAM J. Disc. Math, 23 (2009), pp. 1149–1178.
- [4] B. BURTON, E. W. CHAMBERS, M. VAN KREVELD, W. MEULEMANS, T. OPHELDERS, AND B. SPECKMANN, *Computing optimal homotopies over a spiked plane with polygonal boundary*, in Proceedings of the 25th Annual European Symposium on Algorithms, 2017, pp. 23:1 – 23:14.
- [5] A. CALDER AND J. SIEGEL, *On the width of homotopies*, Topology, 19 (1980), pp. 209–220.
- [6] E. W. CHAMBERS, G. R. CHAMBERS, A. DE MESMAY, T. OPHELDERS, AND R. ROTMAN, *Constructing monotone homotopies and sweepouts*. arXiv: 1704.06175, 2017.
- [7] E. W. CHAMBERS, É. COLIN 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, 43 (2010), pp. 295–311.
- [8] E. W. CHAMBERS, É. COLIN 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, 43 (2010), pp. 295 – 311. Special Issue on 24th Annual Symposium on Computational Geometry (SoCG’08).
- [9] E. W. CHAMBERS AND D. LETSCHER, *On the height of a homotopy*, in Proceedings of the 21st Canadian Conference on Computational Geometry, 2009, pp. 103–106.
- [10] G. R. CHAMBERS AND Y. LIKUMOVICH, *Converting homotopies to isotopies and dividing homotopies in half in an effective way*, Geometric and Functional Analysis, 24 (2014), pp. 1080–1100.
- [11] G. R. CHAMBERS AND R. ROTMAN, *Monotone homotopies and contracting discs on Riemannian surfaces*. arXiv:1311.2995, 2016.
- [12] É. COLIN DE VERDIÈRE AND J. ERICKSON, *Tightening nonsimple paths and cycles on surfaces*, SIAM Journal on Computing, 39 (2010), pp. 3784–3813.
- [13] M. DEHN, *Transformation der Kurven auf zweiseitigen Flächen*, Mathematische Annalen, 72 (1912), pp. 413–421.
- [14] D. DERENIOWSKI, *From pathwidth to connected pathwidth*, SIAM Journal on Discrete Mathematics, 26 (2012), pp. 1709–1732.
- [15] T. K. DEY AND S. GUHA, *Transforming curves on surfaces*, Journal of Computer and System Sciences, 58 (1999), pp. 297–325.
- [16] M. P. DO CARMO, *Riemannian geometry*, Birkhäuser, 1992.
- [17] J. ERICKSON AND K. WHITTLESEY, *Transforming curves on surfaces redux*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2013, pp. 1646–1655.
- [18] F. V. FOMIN AND D. M. THILIKOS, *An annotated bibliography on guaranteed graph searching*, Theoretical computer science, 399 (2008), pp. 236–245.
- [19] P. FRAIGNAUD AND N. NISSE, *Connected treewidth and connected graph searching*, in Latin American Symposium on Theoretical Informatics, Springer, 2006, pp. 479–490.
- [20] M. GROMOV, *Quantitative homotopy theory*, Prospects in Mathematics: Invited Talks on the Occasion of the 250th Anniversary of Princeton University (H. Rossi, ed.), (1999), pp. 45–49.
- [21] Q. HAN AND J.-X. HONG, *Isometric embedding of Riemannian manifolds in Euclidean spaces*, vol. 130 of Mathematical Surveys and Monographs, American Mathematical Society, Providence, RI, 2006.
- [22] S. HAR-PELED, A. NAYYERI, M. SALAVATIPOUR, AND A. SIDIROPOULOS, *How to walk your dog in the mountains with no magic leash*, Discrete & Computational Geometry, 55 (2016), pp. 39–73.
- [23] G. KOKAREV, *On geodesic homotopies of controlled width and conjugacies in isometry groups*, Groups, Geometry, and Dynamics, 7 (2013), pp. 911–929.
- [24] A. S. LAPAUGH, *Recontamination does not help to search a graph*, Journal of the ACM (JACM), 40 (1993), pp. 224–245.
- [25] H. LAWSON, *Lectures on minimal submanifolds*, vol. 1 of Mathematics lecture series, Publish or Perish, 1980.
- [26] F. LAZARUS AND J. RIVAUD, *On the homotopy test on surfaces*, in Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2012, pp. 440–449.
- [27] B. MONIEN AND I. SUDBOROUGH, *Min cut is NP-complete for edge weighted trees*, Theoretical Computer Science, 58 (1988), pp. 209 – 229.
- [28] K. REIDEMEISTER, *Elementare Begründung der Kno-*

- tentheorie*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, 5 (1927), pp. 24–32.
- [29] J. STILLWELL, *Classical topology and combinatorial group theory*, Springer-Verlag, New York, 1980.
- [30] B. WHITE, *Mappings that minimize area in their homotopy classes*, Journal of Differential Geometry, 20 (1984), pp. 433–446.
- [31] B. YANG, D. DYER, AND B. ALSPACH, *Sweeping graphs with large clique number*, in International symposium on algorithms and computation, Springer, 2004, pp. 908–920.