The Star Unfolding from a Geodesic Curve

by

Stephen Kiazyk

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2014

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

An *unfolding* of a polyhedron $\mathcal{P}$ is obtained by 'cutting' the surface of $\mathcal{P}$ in such a way that it can be flattened into the plane into a single polygon. For most practical and theoretic applications, it is desirable for an algorithm to produce an unfolding which is *simple*, that is, non-overlapping. Currently, two methods for unfolding which guarantee non-overlap for convex polyhedra are known, the *source* unfolding, and the *star* unfolding. Both methods involve computing shortest paths from a single source point on the polyhedron's surface.

In this thesis, we attempt to prove non-overlap of a variant called the *geodesic star unfolding*. This unfolding, much like the star unfolding, is computed by cutting shortest paths from each vertex to λ, a geodesic curve on the surface of a convex polyhedron $\mathcal{P}$, and also cutting λ itself. Non-overlap of this case was conjectured by Demaine and Lubiw [15]. We are unsuccessful in completely proving non-overlap, though we present a number of partial results, and discuss some areas for future study. We first develop a new proof for non-overlap of the star unfolding from a point. The original proof of non-overlap was given by Aronov and O'Rourke [7]. This new proof uses a partitioning of the unfolding around the ridge tree. Each edge of the ridge tree serves as a base edge to a pair of congruent triangles; in this way, the whole unfolding is decomposed into these pairs which are called *kites*. We prove non-overlap by showing that pairwise, no two kites in the unfolding overlap each other, by a method which bounds the surface angle of the source images to either side of any path through the ridge tree.

In addition to its simplicity compared to the previous proof, this new method easily generalizes to prove non-overlap for some cases of the star unfolding from geodesic curves. Specifically, we show non-overlap for two classes of geodesic curves, *geodesic loops*, and *fully-extended S-shaped geodesics*, by showing that the surface angle of the source images in those two cases are bounded. We also investigate a class of curves called *fully-extended C-shaped geodesics* for which the proof cannot hold directly. We show some specific cases where we are able to create a supplementary proof to show non-overlap, though non-overlap for the class as a whole remains unproven.

Acknowledgements

I would like to thank my advisor Anna Lubiw for all of her countless hours of instruction, guidance and support during my time as a student and researcher with her. Her assistance and insight has been invaluable in bringing this thesis to completion.

I would also like to thank my readers, Timothy Chan and Ian Munro, for taking the time to evaluate and provide feedback on on my work.

Thank you to all of the friends and acquaintances who helped shape my experience at the University of Waterloo. Special thanks (in alphabetic order, of course) goes out to: Cecylia Bokovich, Laura Cang, David Dufour, Hella Hoffmann, Marianna Rappoport, Aaron and Christina Moss, Rhiannon Rose, Dean Shaft, Valerie Sugarman, Jack Thomas, and Oliver Trujillo.

Finally, a special thanks to my parents, Chris and Barbara, and my siblings, Philip, Catherine, and Annette, for their endless love and encouragement.

Table of Contents

List of Figures

Chapter 1

Introduction and Background

An *unfolding* of a polyhedron $\mathcal{P}$ is obtained by cutting the surface of $\mathcal{P}$ in such a way that it can be flattened into the plane, forming a single polygon. A *folding* of a polygon is the reverse operation: to glue the edges of a polygon together in order to form a polyhedron.

Unfolding provides a way to decompose complex structures into a simpler form. From a practical standpoint, unfolding has uses in product manufacturing, where one might want to construct a 3-dimensional object from a sheet of metal or plastic, or in applying texture mapping in graphical applications, where 2-dimensional image coordinates must be assigned to points on a 3-dimensional model. Unfolding is also applied as a theoretical tool for the study of shortest paths constrained to the surface of a polyhedron. Unlike the general case of shortest paths in 3 dimensions with obstacles (which is NP-hard), the problem of shortest paths on a polyhedral surface is efficiently solvable, and thus many related problems can use unfolding as an intermediate step to reduce the problem to a simpler 2-dimensional variant; for example, see Chen and Han [11], and Agarwal et al. [4].

For applications of unfolding, we typically seek to create unfoldings which are simple (i.e. non-overlapping). While the concept is easy to understand and visualize, identifying and proving a method for unfolding which guarantees no overlap remains a difficult area of study. There are two major types of unfolding methods on polyhedra: *edge* unfolding, and *general* unfolding. An *edge* unfolding cuts only along the edges of a polyhedron. The earliest known examples of the edge unfolding of polyhedra are found in the works of German renaissance printer and painter Albrecht Dürer [19]. Implicit in his works is one of the most famous and long-standing questions in geometry: what classes of polyhedra admit a non-overlapping edge unfolding? In general, there exist non-convex polyhedra where every possible edge unfolding self-overlaps. For convex polyhedra, the question whether

there always exists at least one edge unfolding which does not self-overlap remains open. A *general* unfolding loosens the restrictions of the edge unfolding, allowing one to cut through the interior of faces as well. It remains unknown whether this can always be accomplished for non-convex polyhedra without overlap, however a solution can always be found for any convex polyhedron. More details, historical background, and references will be given in Section 1.1.

The methods for unfolding a convex polyhedron without overlap involve computing shortest paths on the surface of the polyhedron. A fundamental property is that when a shortest path crosses from one face to another, it becomes a straight line when you unfold those two faces to the plane. More generally, any path that forms a straight line when unfolded is called a *geodesic curve*. A geodesic curve is always a *locally shortest path*, i.e. it cannot be made shorter by local perturbations. Note, however, that every geodesic curve is not necessarily a shortest path. Throughout the thesis, we only consider *simple* geodesic curves, i.e. geodesic curves that do not self-intersect, except possibly at its start and endpoints. For more formal definitions related to geodesic curves, see Section 1.3.

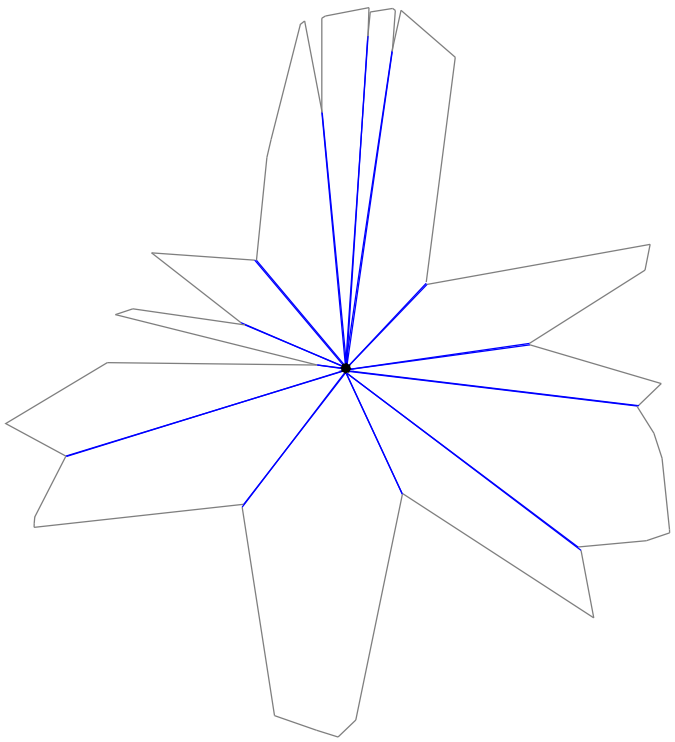

Figure 1.1: Source unfolding, all shortest paths are laid out around the source point.

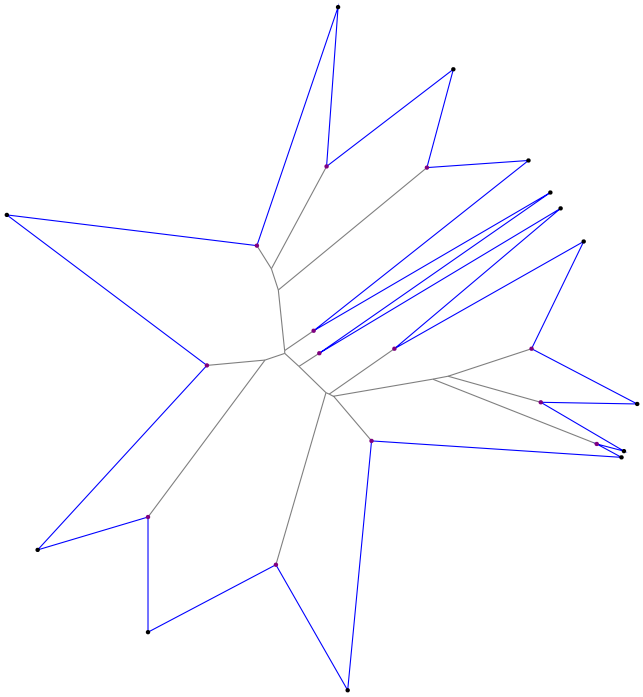

Figure 1.2: Star unfolding produced from the same source point as Figure 1.1. The shortest paths to the vertices (blue) are cut, and the perimeter of the source unfolding (grey) is glued to form the star unfolding.

There are two primary methods of unfolding which have been proven to avoid overlap on convex polyhedra: the *source* unfolding (Figure 1.1) and the *star* unfolding (Figure 1.2). The *source* unfolding simply selects a point x on the surface of a polyhedron $\mathcal{P}$, i.e. the source, and cuts along the *cut locus*, or *ridge tree*, which is the closure of all points with more than one shortest path to x (the term cut locus is a general term used in all Riemannian geometry, while the term ridge tree is specific to discrete polyhedral surfaces). The result lays out all points on the surface from which there is a unique shortest path to x. It is easy to show that the source unfolding does not overlap for convex polyhedra; simply observe that all shortest paths originating from x will form a straight line with the same origin point when unfolded into the plane (because they are geodesic curves), and thus, no two such paths can cross.

The *star* unfolding, by contrast, still uses a single source point $x \in \mathcal{P}$, but instead cuts a shortest path from x to every vertex in $\mathcal{P}$. The star unfolding can be thought of as 'dual'

to the source unfolding. In both cases, the surface is partitioned by the ridge tree and the shortest path cuts from each vertex to x. In the case of the source unfolding, the pieces of the partition are joined around x, while in the star unfolding, they are joined at the ridge tree (see Figures 1.1 and 1.2).

The primary goal of this thesis is to prove the following extended version of the star unfolding: given any simple geodesic curve $\lambda \in \mathcal{P}$, is the unfolding produced by cutting the curve λ and cutting shortest paths from λ to every vertex of $\mathcal{P}$ non-overlapping? We will call this the *Geodesic Star Unfolding*. While we were unsuccessful at proving this conjecture completely, we have found 3 partial results.

- We give a new proof of non-overlap for the point case of the Star Unfolding which is both shorter, and proves a slightly stronger claim than the original proof by Aronov and O'Rourke [7].
- We apply this proof to show non-overlap of a subset of geodesic cases, including the *quasigeodesic loop* case (defined below).
- We show other cases where we can restrict the specific regions in which there may exist overlap in a geodesic star unfolding.

1.1 Background

1.1.1 General Unfolding

Proof of non-overlap of the source unfolding on convex polyhedra is fairly straightforward [16, p. 355]. By comparison, the corresponding proof of the star unfolding is much more difficult. Aronov and O'Rourke first proved non-overlap of the star unfolding in 1992 [7], using a technique of Alexandrov's [5] to take a convex polyhedron $\mathcal{P}$, and produce a new polyhedron $\mathcal{P}'$ where the number of vertices is reduced by one. They argue inductively that the star unfoldings of $\mathcal{P}$ and $\mathcal{P}'$ are similar enough that non-overlap of $\mathcal{P}'$ implies non-overlap of $\mathcal{P}$. They reduce to the base cases of either a tetrahedron or a (degenerate) doubly-covered triangle, for which the star unfolding can be shown directly to not overlap.

Some papers since then have considered extensions of these basic methods (source and star unfolding) by changing the source point from a single point x to a geodesic curve λ on the surface of $\mathcal{P}$, and instead cutting either the ridge tree with respect to the curve (for

the source unfolding), or cutting the curve itself along with the set of shortest paths from each vertex back to the curve (for the star unfolding).

In 2009, Itoh, O'Rourke, and Vîlcu in 2009 [21] proved non-overlap for the special case of a star unfolding from a *quasigeodesic loop*, a special kind of closed curve. Normally this would disconnect the surface, but they showed there exists some segment of the quasigeodesic that can be re-glued such that the two pieces do not overlap, either with each other or themselves. In their discussion, they conjecture that quasigeodesic loops are not the widest class of curves for which a star unfolding method can lead to non-overlap.

Demaine and Lubiw in 2011 [15] presented an extension of the source unfolding, called the *Sun* unfolding, where source x can be either a geodesic curve on the surface of $\mathcal{P}$, or a closed convex curve composed of straight line segments and circular arcs (the former being a special case of the latter). In their discussion, they conjecture that the 'dual' version of the sun unfolding (in the same way that the star unfolding is dual to the source unfolding) does not self-overlap. They describe what this unfolding would look like, and state that a likely first step would be to determine if the star unfolding from any open geodesic is non-overlapping.

1.1.2 Edge Unfolding

It has been shown (for example, by Biedl et al. [9]) that in general there exist non-convex polyhedra where every possible edge unfolding results in overlap. This even holds for non-convex polyhedra whose underlying graphs are isomorphic to those of convex polyhedra (Bern et al. [8]), disproving an earlier conjecture by Shepard [26].

Conversely, the conjecture that every convex polyhedron has a non-overlapping edge unfolding remains open; no proof or counter-example has yet been found [16, p. 310]. Most of the current techniques used to find non-overlapping edge unfoldings are heuristics only, or are limited to a small set of specialized classes with limited or no applicability to other classes. Furthermore, little interplay between edge and general unfolding has been found; methods used for one tend not to translate effectively to the other.

Attempts have been made to identify limited classes of polyhedra for which non-overlapping edge unfolding can be proven. DiBiase [18] showed through a case-exhaustive proof that all polyhedra with 6 or fewer vertices admit a non-overlapping edge unfolding. Demaine and O'Rourke showed that 'dome' polyhedra, convex polyhedra with a single base face that shares an edge with every other face, can be edge-unfolded without overlap using the 'volcano' unfolding technique [16, p. 323].

One class that gained particular interest is prismatoids, and the slightly more restrictive prismoids. Demaine and O'Rourke [16, p. 319] showed non-overlapping edge unfoldings exist for all prismoids. Aloupis et al. [6] proved the existence of a non-overlapping edge unfolding for nested polyhedral bands. Polyhedral bands are formed by taking the intersection of the space between two parallel planes with some convex polyhedron $\mathcal{P}$ such that the intersected space contains no vertices of $\mathcal{P}$. These bands are not closed, in that there is no 'top' or 'bottom' face, only a circular strip of edge-adjacent faces. Nested polyhedral bands imply that the intersection of those two parallel planes with $\mathcal{P}$ are 'nested', that is the orthogonal projection of one plane is contained within the other. O'Rourke [23] later superseded this result to show non-overlap of the case of topless prismatoids, an even more general class, using a technique called the petal unfolding. Prismatoids are the convex hull of two parallel convex faces; topless implies that one of those two parallel faces is removed from the resulting polyhedron.

A variant on the edge unfolding problem that has been studied is the *zipper edge unfolding*. The goal of the zipper edge unfolding is to form a non-overlapping unfolding by taking a Hamiltonian path of cuts through the 1-skeleton of a polyhedron. Demaine et al. presented this idea in 2010 [13], showing that all Archimedean and Platonic solids have a zipper edge unfolding. Later, Demaine et al. [14] showed classes of domes and prismoids which have zipper edge unfoldings.

1.2 Motivation

A hope for further study into unfoldings is to develop a wider repertoire of techniques for proving non-overlap of different unfolding algorithms. In particular, it would be useful to have techniques which give useful bounds on the properties of an unfolding. For example, one might want to find an unfolding that minimized the total or maximum length of all cuts made to the surface, or minimized the 2-dimensional bounding box around the unfolded polygon.

A potential implication of non-overlap of the geodesic star unfolding discussed in this thesis would be a *general zipper unfolding*. Instead of finding a Hamiltonian path restricted only to the 1-skeleton of the polyhedron, we could instead find a quasigeodesic path which touches each vertex (the formal definition of quasigeodesic will be covered in Section 1.3.4). If the star unfolding from any quasigeodesic curve unfolds without overlap, and such a curve can be found, non-overlap of this method would follow directly.

1.3 Preliminaries

In this section we establish important definitions and observations related to polyhedra, shortest paths, and unfolding. This section can safely be ignored if one is already familiar with other results related to unfolding polyhedra. For all of the following, we assume that the polyhedra are *closed*, i.e. they have no boundary edges.

1.3.1 Surface Angle and Discrete Curvature

Definition 1.1. *Given a polyhedron $\mathcal{P}$, and a point $x \in \mathcal{P}$, the* surface angle *of x, $\angle(x)$ is the sum of the incident angles of all faces of $\mathcal{P}$ containing x. $\angle(x) = 2\pi$ for every non-vertex point of $\mathcal{P}$, and is strictly less than 2π at every vertex of $\mathcal{P}$. A related definition is the* curvature *of x, $\kappa(x) = 2\pi - \angle(x)$.*

Intuitively, curvature is the measure of how much surface angle is 'missing' to make the surface of $\mathcal{P}$ flat at point x.

The sum of the discrete curvature over every vertex $v \in \mathcal{P}$ on a 3-dimensional polyhedron is exactly 4π. The original proof of this fact goes back to Descartes' Theorem on total angular defect [17], but is also a special case of the Gauss-Bonnet Theorem for compact 2-dimensional Riemann manifolds [12, p.216].

1.3.2 Convexity of Polyhedra

Convexity of polyhedra is typically given as one of a set of equivalent conditions, for example:

Definition 1.2. *A polyhedron $\mathcal{P}$ is* convex *if and only if for each face $f \in \mathcal{P}$, $\mathcal{P}$ lies entirely to one side of the supporting plane of f.*

The most important property of convex polyhedra that we use in this thesis is as follows.

Observation 1.1. *Given a convex polyhedron $\mathcal{P}$, the discrete curvature at every point $p \in \mathcal{P}$ is greater than or equal to zero.*

This observation is a corollary to a theorem of Alexandrov [5, p. 210].

Note that the condition of Observation 1.1 also holds for some non-convex polyhedra as well. However, the manifold of every such non-convex polyhedron corresponds to the manifold of exactly one convex polyhedron, so unfolding results for the larger class are not more general than for the set of convex polyhedra.

1.3.3 Necessary Conditions for Unfolding

An unfolding must leave the surface flat, and in one connected piece, which leads to the following fact.

Lemma 1.1. *In order to produce an unfolding of a convex polyhedron $\mathcal{P}$, it is necessary and sufficient for the cuts to form a tree on the surface that reaches every vertex v_i of $\mathcal{P}$.*

Proof. The only 'non-flat' pieces of $\mathcal{P}$ are its vertices. Thus, we can flatten the surface by forming cuts that reach each vertex and then spreading the surface angle around the cut to one side or the other. A single cut to each vertex is sufficient, since the curvature at every vertex is positive (i.e. the surface angle around each vertex strictly less than 2π), per Observation 1.1. Multiple cuts would be necessary to flatten any point with negative curvature. A tree of cuts is necessary to retain connectivity of the flattened result as a single polygon (since any cycle would disconnect the surface). □

1.3.4 Geodesics and Quasigeodesics

A geodesic curve with start and endpoints a and b respectively on the surface of polyhedron $\mathcal{P}$ is a *locally shortest* path from a to b, that is, no local perturbation in λ will decrease its overall length. Every shortest path between a pair of points a and b on the surface of a convex polyhedron $\mathcal{P}$ is a locally shortest (and therefore, geodesic) curve. However the converse does not hold in general: a geodesic path between two points may not be a shortest path between them. We present a more formal definition of geodesic curves to help with our analysis later on.

Definition 1.3. *Given a polyhedron $\mathcal{P}$, a* geodesic curve $\lambda \in \mathcal{P}$ *is a curve along the surface of $\mathcal{P}$, such that at every interior point $p \in \lambda$, the surface angle to either side of p is* exactly π.

Figures 1.3 and 1.4 illustrate an important feature of geodesic curves: when the faces traversed by the curve are 'unrolled' into the plane, the image of the curve will form a straight line.

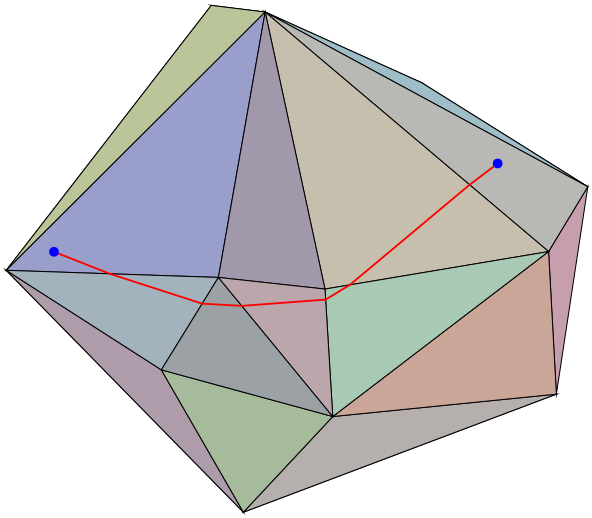

Figure 1.3: A geodesic curve on the surface of a polyhedron.

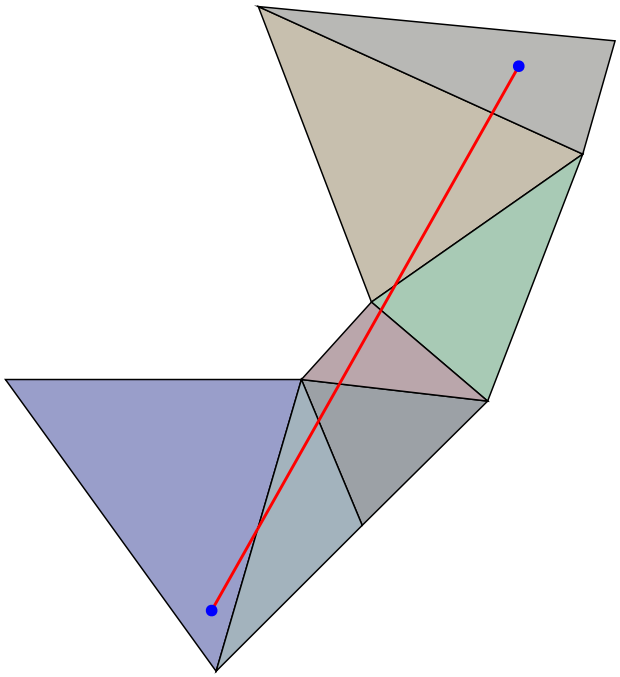

Figure 1.4: The geodesic curve of Figure 1.3, unrolled with its faces into the plane.

This definition requires that the curvature at every point along the geodesic λ is exactly 0 (except possibly at its two endpoints); a curve which passes through a vertex cannot be a geodesic. We give a slightly more relaxed definition to refer to curves that are 'geodesic-like', except they are allowed pass through vertices, called quasigeoedesic curves.

Definition 1.4. *Given a polyhedron $\mathcal{P}$, a* quasigeodesic curve $\lambda \in \mathcal{P}$ *is a curve along the surface of $\mathcal{P}$, such that at every interior point $p \in \lambda$, the surface angle to either side of p is $\leq \pi$.*

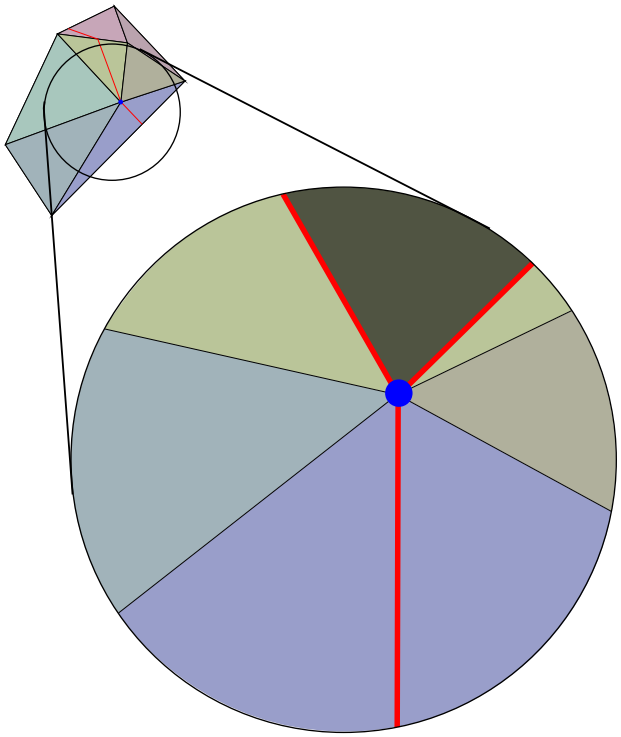

Figure 1.5: A quasigeodesic curve, with a zoomed in and flattened view of the surface near a crossed vertex. The darkened area above the vertex represents the gap formed by the curvature at vertex, and the two sides of the quasigeodeic after they pass through the vertex.

A quasigeodesic curve is geodesic at every point along the curve with the exception of where the curve passes through, or 'crosses' a vertex. Thus, like geodesic curves, a quasigeodesic curve will 'unroll' to a straight line when laid out with its incident faces into the plane, except where the curve meets a vertex (see Figure 1.5). Notice that a

quasigeodesic curve can only pass through vertices with positive curvature; a negative curvature vertex would necessarily have more than π surface angle to either side, and thus the curve would no longer be quasigeodesic. In this thesis, we will mostly restrict our arguments to only consider when the source curves are geodesic. In Section 3.3, however, we will discuss how our arguments generalize to quasigeodesic curves as well.

We will only consider geodesic curves that are *simple*, meaning there is no point of self-intersection between any two interior (i.e. non-endpoint) points of the curve. However, we will allow endpoints to lie incident either to each other or to the interior of the curve. The following definition covers the case when the two endpoints of the curve are co-incident.

Definition 1.5. *Given a polyhedron $\mathcal{P}$, a* (quasi-)geodesic loop *is a closed curve $\lambda \in \mathcal{P}$ that is (quasi-)geodesic at all points along λ, except at a single location called the* loop point*, o, which has surface angle $\geq \pi$ to at least one side. The (quasi-)geodesic loop will divide $\mathcal{P}$ into two sub-surfaces to either side of the curve. We will refer to these sub-surfaces by the angle at the loop point within that sub-surface. The side with the convex angle is the* interior *of the loop, and that with non-convex angle is the* exterior*. When λ has no distinguishable loop point (that is when the curve is (quasi-)geodesic at all points along the curve), then the curve is simply called a* closed (quasi-)geodesic*. Otherwise, if a geodesic curve is simple and also has no points of intersection with the curve endpoints, then the curve is an* open *(quais-)geodesic curve.*

The angle of the loop point on either sub-surface is the sum discrete curvature inside that sub-surface minus π. Notice that the curvature to either side must be at least π and at most 3π. These observations (which are also found in [21]) will become useful in Chapter 3. There are also (quais-)geodesic curves in which one or both of the endpoints are co-incident with the interior of the curve; these will be discussed in greater detail in Section 3.5.

Chapter 2

A new proof of the Point Star Unfolding

In this chapter we give a new proof of non-overlap for the star unfolding from a point. The original proof was given by Aronov and O'Rourke [7] and is somewhat long and complex compared to our result.

2.1 Preliminaries and Definitions

Note that most of the definitions and terminology used here are borrowed from [7].

Definition 2.1. *Let $\mathcal{P}$ be a convex polyhedron, and let x be a point on $\mathcal{P}$. The* star unfolding*, S_x, is a 2-dimensional polygon formed by cutting $\mathcal{P}$ along a shortest path from every vertex of $\mathcal{P}$ to x, and flattening the result into the plane. See Figure 2.1 for a perspective view of the shortest path cuts on a polyhedron, and Figure 2.2 for the resulting unfolding.*

The following theorem is the main result of [7]. We give a new proof of this theorem in this chapter.

Theorem 2.1. *(Aronov and O'Rourke [7, Theorem 9.1]) The star unfolding S_x does not overlap, i.e. it is a simple polygon.*

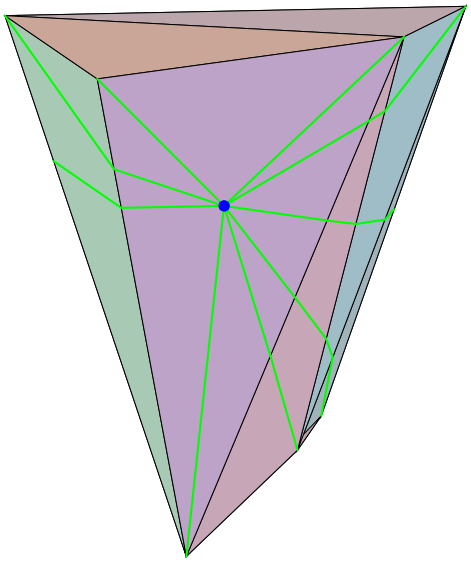

Figure 2.1: Shortest paths from all vertices to a source point.

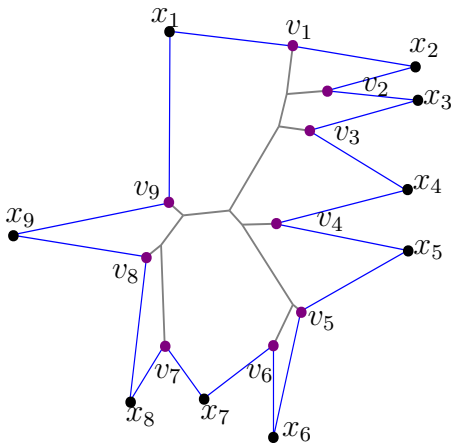

Figure 2.2: The star unfolding produced by the cuts from Figure 2.1. The ridge tree is shown in grey.

2.1.1 The Structure of the Star Unfolding

If $\mathcal{P}$ has n vertices, then the polygon S_x will (in general) have $2n$ vertices (the unfolding will have $2(n-1)$ vertices if x is located at a vertex of $\mathcal{P}$). The vertices of S_x will alternate around the boundary between points corresponding to the n vertices of $\mathcal{P}$, called *vertex images* and denoted v_i, and n 'copies' of x, called *source images* and denoted x_i (see Figure 2.2). The edges of S_x correspond to the shortest path cuts made from each vertex to x. Therefore, the two edges incident to any vertex image v_i are always the same length.

The exterior angle at each v_i is exactly $\kappa(v_i)$, the discrete curvature of that vertex on the surface of $\mathcal{P}$. This gap between consecutive source images is called the *cuvature gap* of v_i. The interior angle at each image x_i is the angle between the shortest paths to the two vertices adjacent to x_i in the unfolding; we refer to this angle as the *surface cut angle* of x_i, denoted $\angle(x_i)$. Clearly, the sum of the surface cut angles over all images of x is exactly the surface angle at x: when x is located at some vertex $v_j \in \mathcal{P}$, then $\sum_i x_i = 2\pi - \kappa(v_j)$, otherwise the sum will be exactly 2π.

2.1.2 The Ridge Tree

An important tool we will use in our analysis of the star unfolding is the *ridge tree* (also sometimes referred to as the *cut locus*), defined as follows.

Definition 2.2. *Given a polyhedron $\mathcal{P}$, and a source point $x \in \mathcal{P}$, the* ridge tree T_x *is the closure of the locus of all points on the surface of $\mathcal{P}$ that have more than one shortest path from x on the surface of $\mathcal{P}$. (An example ridge tree is shown by the grey tree inside the polygon of Figure 2.2).*

Aronov and O'Rourke proved that the ridge tree is a subset of the Voronoi diagram of the images of x as a corollary to their proof of non-overlap [7, Theorem 10.2]. While this fact is not required for our proof, it does provide useful intuition about its structure. What we do require about the structure of the ridge tree is that it is indeed a tree (this was addressed by Demaine and Lubiw [15]), and that it is composed strictly of straight (or rather, geodesic) line segments. This second point must be true since it is formed from the bisectors of point images in the plane.

We will use the ridge tree to construct a set of triangles from the unfolding polygon, shown in Figure 2.3. Each vertex of the ridge tree has 3 or more shortest paths to x on the surface of P. The triangles are formed by tracing the shortest paths from each ridge tree vertex to all of its nearest source images x_i. Notice that each edge of the ridge tree will give rise to exactly one triangle to each side. Because all shortest paths from a ridge tree vertex to the nearest source images x_i have the same length, the two triangles to either side of every ridge tree edge will have corresponding sides of equal length, and therefore are congruent. These pairs of triangles, which we will call *kites*, are central to our analysis.

Definition 2.3. *A* kite *of the star unfolding S_x with ridge tree T_x is the pair of congruent triangles to either side of some edge $e \in E(T_x)$. We will denote such a kite as kite(e). We will call the edge of the ridge tree the* base *of this kite (which is also its axis of symmetry), and the two images of x to each side its* apices.

Example pairs of triangles which form kites are highlighted in Figure 2.3. Observe that the kites form a partition of S_x. Also note that these kites may be either concave or convex. Our proof will show that no pair of distinct kites in the unfolding overlap.

Before moving to our proof, we require one more observation about the surface angle at the source images of x in relation to the ridge tree.

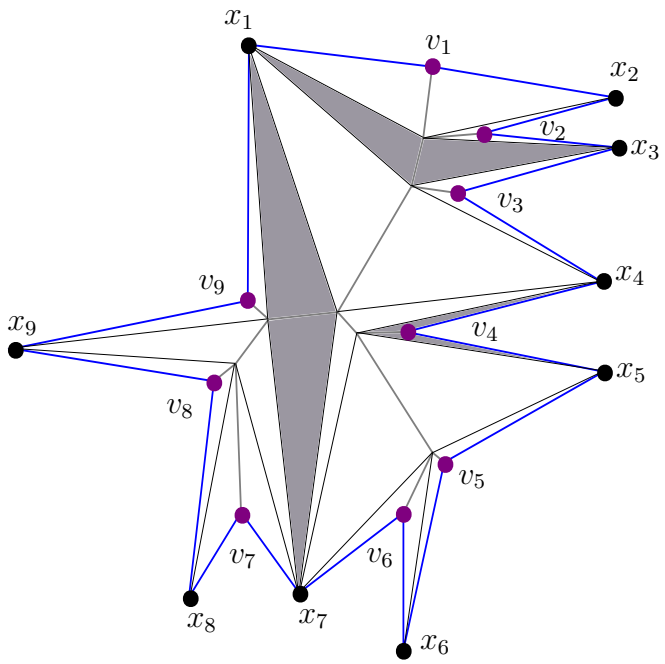

Figure 2.3: An example star unfolding broken down into triangles using the ridge tree. Some of the triangle pairs which form kites are highlighted.

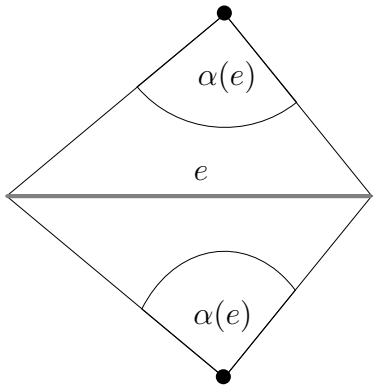

Figure 2.4: The source angle of a single edge e of the ridge tree.

Definition 2.4. *For any edge e of a ridge tree T_x, let the* source angle *of e, $\alpha(e)$, be the interior angle at either apex of kite(e) (See Figure 2.4).*

The source angle of a ridge tree edge e is a piece of the surface cut angle of x that lies inside the shortest paths from the endpoints e back to the image of x on one side or the other.

We will use the same notation for a path through the ridge tree. That is for a path σ, where $\sigma = e_0, \ldots, e_t$, the source angle of σ is $\alpha(\sigma) = \sum_{i=0}^{t} \alpha(e_i)$.

Observation 2.1. *For any path through the ridge tree σ, $\alpha(\sigma) \leq \pi$*

Proof. Notice that the sum of the angles at all of these triangles' apices comprises the entire surface angle at x. Since the surface angle at x is at most 2π, and the source angle is paired to both sides of each edge, the total source angle over all the edges through any path is at most $\frac{1}{2}$ of that. □

2.2 Proof of Non-Overlap

We will now present a new proof of Theorem 2.1. Our proof will show that any two kites, *kite*(u) and *kite*(v), do not overlap, and thus neither does the unfolding as a whole. To do so, we will establish a region of $\mathbb{R}^2$ which completely contains *kite*(u) while completely excluding *kite*(v). This region will be called a *W-wedge*, due to the shape of its boundary.

The *W-Wedge* is a region defined in terms of some edge y of the ridge tree T_x, a 'root' vertex r which is an endpoint of edge y, and an angle $\gamma \in [\alpha(y) + \frac{\pi}{2}, \frac{3\pi}{2}]$. The boundary of the W-wedge has two *inner legs*, call them l_1 and l_2, and two *outer legs*, L_1 and L_2. The inner legs are line segments connecting r to either apex of *kite*(y). The outer legs are rotated away from l_1 and l_2 respectively by an angle of γ on the side of *kite*(y), and extend either to their point of intersection, or as infinite rays if they do not intersect. The *inside* of the W-Wedge is the side of the boundary which contains *kite*(y), and the *outside* is the side which contains *kite*(z). Figure 2.5 illustrates the 3 possible scenarios for the division of $\mathbb{R}^2$ by the W-wedge, that is for one, the other, or both of the regions to be unbounded. Given a W-wedge W, we denote an object o as strictly inside W by $o \subseteq W$, and strictly outside of W by $o \subseteq \overline{W}$.

Lemma 2.1. *Let W be a W-wedge rooted at vertex r on edge y of a geodesic ridge tree, and let z be some other ridge tree edge incident to r. Suppose W has angle $\gamma \in [\alpha(y), 2\pi - \alpha(z)]$. Then kite(y) $\subseteq W$ and kite(z) $\subseteq \overline{W}$.*

Proof. The two inner legs of W lie exactly along the edges of *kite*(y) incident to r, and thus obviously do not intersect the kite. If the outer legs are rotated away from the inner

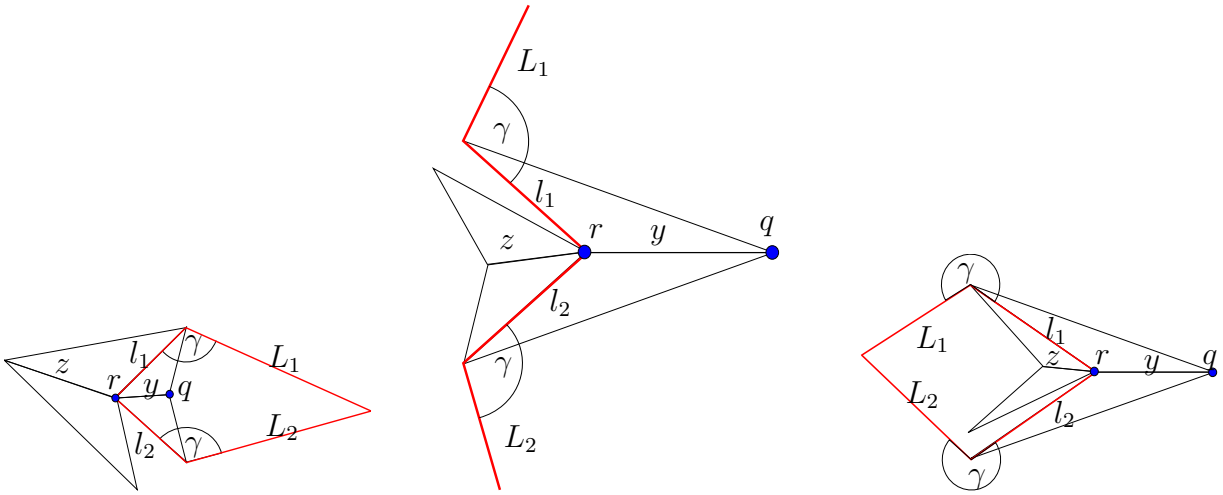

Figure 2.5: Illustration of W-wedges of vertex r on edge y, showing three possible cases. From left to right: a bounded region containing $kite(y)$, two unbounded regions, and a bounded region not containing $kite(y)$.

legs by $\frac{\pi}{2} + \alpha(y)$, then they are at least perpendicular to the edges of $kite(y)$ incident to the other endpoint of y (labelled q in Figure 2.5). Thus, the two rays either diverge (centre image of Figure 2.5), intersect such that they enclose $kite(y)$ (left image of Figure 2.5), or intersect such that the region containing $kite(y)$ is unbounded (right image of Figure 2.5). Either way, this implies $kite(y)$ lies to the inside of W. A symmetric argument for edge z shows that $kite(z)$ lies strictly outside W. □

The idea of our proof is to establish an initial W-wedge at one end of the path, and build subsequent W-wedges at each new vertex of the path. One side of each wedge is the inside region (as defined above), which will strictly contain all previous W-wedges, and therefore all kites visited so far. In each step, we will prove that the newly added kite initially lies entirely on the outside region of the current wedge (and therefore does not overlap any of the previous kites), then we expand the wedge to cover a strict superset of the previous wedge that includes the newly added kite. Figure 2.6 gives a general overview of the proof's construction.

Lemma 2.2. *Given two distinct edges u and v of the ridge tree T_x of some star unfolding S_x, the two associated kites $kite(u)$ and $kite(v)$ do not overlap.*

Proof. Let $\sigma = p_0, \ldots, p_t$ be the vertex-path through the ridge tree from u to v. Let $e_i = p_{i-1}p_i$ be the edges of the path, such that $e_1 = u$, $e_t = v$.

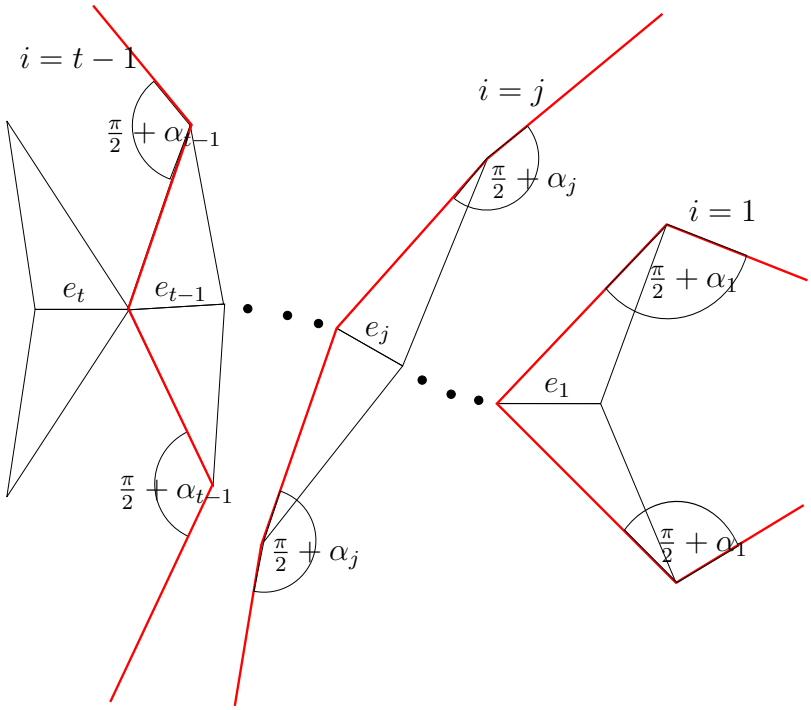

Figure 2.6: The proof proceeds from e_1 to e_t by establishing wedges which each successively contain all previous kites and wedges, while remaining disjoint from all future kites.

Let $k_i = kite(e_i)$, that is, the i^{th} kite along the path.
Let $\alpha_i = \sum_{j=1}^{i} \alpha(e_j)$ (note that $\alpha_i \leq \pi$ for all i by Observation 2.1).

Define W_i, for $i = 1, \ldots, t$, to be the W-wedge rooted at p_i on edge e_i, with inner legs along the two incident edges of k_i, and outer legs rotated out by $\gamma_i = \alpha_i + \frac{\pi}{2}$. (The extra angle of $\frac{\pi}{2}$ will be required during the inductive phase of our proof.)

Claim: $W_{i-1} \subseteq W_i$, and $k_i \subseteq (W_i \setminus W_{i-1})$, for $i = 2, \ldots, t$.

We will use Lemma 2.1 with $r = p_{i-1}$, $y = e_{i-1}$, and $z = e_i$ to show that $k_i \subseteq \overline{W}_{i-1}$. Note that the lemma applies since $\alpha(e_{i-1}) \leq \alpha_{i-1} \leq \pi - \alpha(e_i)$, and therefore $\alpha(e_{i-1}) \leq \frac{\pi}{2} + \alpha(e_{i-1}) \leq \gamma_{i-1} \leq 2\pi - \alpha(e_i)$. Then, we apply the two-step transformation indicated by Figure 2.7 to convert W_{i-1} to W_i. First, rotate the two inner legs of W_{i-1} away from the edges of k_{i-1} to k_i, while maintaining a fixed outer leg rotation angle γ_{i-1}. Observe that all the kite edges incident to p_{i-1} have the same length; thus there is no need to change the

length of the inner legs during this rotation. Notice that the intermediate wedge formed by this rotation covers a superset of W_{i-1}, because the angle γ_{i-1} is guaranteed to be in the range $[\frac{\pi}{2}, \frac{3\pi}{2}]$ (this is why we needed an extra $\frac{\pi}{2}$ of angle to each side to ensure adequate clearance during the rotation). Although Figure 2.7 shows consecutive kites sharing only a common vertex of the ridge tree, in all non-degenerate cases there will be at most three edges incident to the vertex, and so the kites will share an edge as well. Thus, typically only one of the inner legs of the wedge will rotate at a time.

The second piece of the transformation is to shift the root point of the wedge to p_i, while keeping the outer legs fixed. The outer leg rotation angle increases by $\alpha(e_i)$, and thus reaches a value of γ_i. The net effect of this step is to keep the outer leg positions fixed while moving the inner legs to the other side of $kite(e_i)$. This newly formed wedge is exactly W_i, and thus, by our construction, $k_i \subseteq W_i$.

With this claim in hand, we apply a simple inductive proof. First, observe that by Lemma 2.1, $k_1 \subseteq W_1$. Furthermore, by our claim, $W_1 \subseteq W_2 \subseteq \cdots \subseteq W_{t-1}$. Finally, once again by Lemma 2.1, $k_t \subseteq \overline{W}_{t-1}$. Thus k_1 and k_t lie on opposite sides of W_{t-1}, and therefore do not overlap.

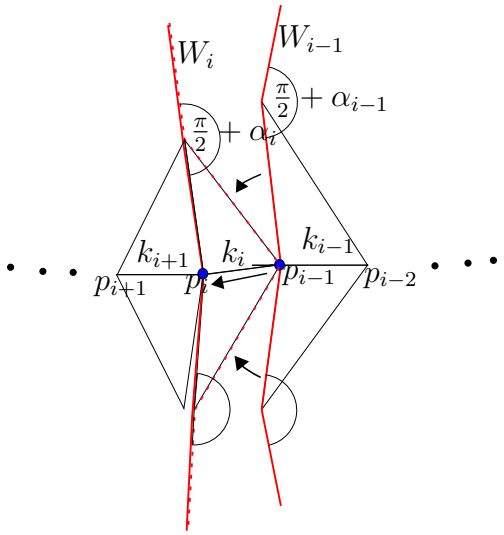

Figure 2.7: Diagram of the inductive step of the proof.

□

Repeating the argument of Lemma 2.2 for all pairs of kites establishes non-overlap and completes the proof of Theorem 2.1.

2.3 Discussion

2.3.1 Larger Sectors of Emptiness

In addition to non-overlap, a corollary to this new proof is a slightly stronger notion of what regions outside of the star unfolding are 'empty' compared to Aronov and O'Rourke's earlier proof [7]. In their proof, they show that the the closed sector of a disk centred at each vertex v_i, bounded by the radii of its two closest source images (call them x_i and x_{i+1}) is empty (see Figure 2.8).

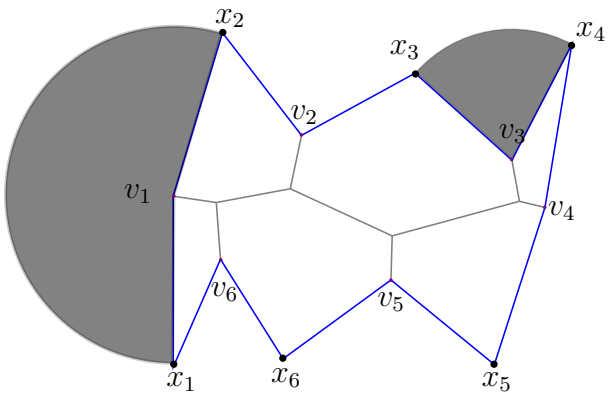

Figure 2.8: Sectors of emptiness established by [7].

Instead, we claim that a strictly larger region around each vertex, formed by rays leaving the two adjacent source images at angles of $\frac{\pi}{2}$, is empty, as a consequence to Lemma 2.2 (see Figure 2.9, note that the region of emptiness around v_1 is in fact unbounded to the left).

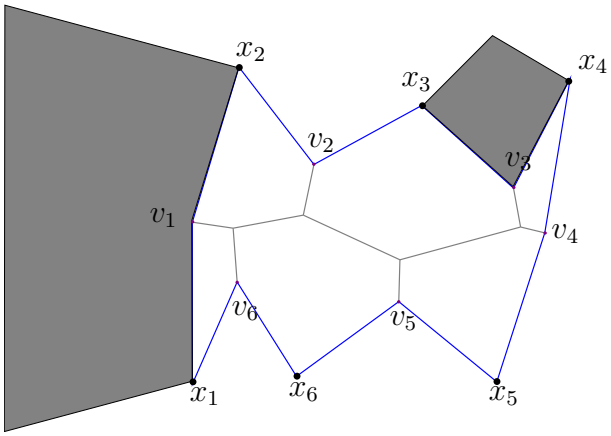

Figure 2.9: New sectors of emptiness implied by Lemma 2.2.

Lemma 2.3. *Given a star unfolding S_x, choose any vertex image v_i. The W-wedge, rooted at v_i with inner legs to its two adjacent images x_i and x_{i+1}, and outer legs rotated out by $\frac{\pi}{2}$ is empty; we will call it the* empty region *for v_i.*

Proof. We divide our proof into two cases: where v_i is a leaf of the ridge tree, and where v_i is an internal vertex of the ridge tree.

- v_i **is a leaf:** Suppose u is the edge of the ridge tree incident to v_i. Consider every possible path in the ridge tree that includes u as its first edge. The initial W-wedge established at u, W_1, contains the empty region on its inside. Every subsequence kite on the path will be to the outside of W_1 by the construction of Lemma 2.2, and therefore the region must remain empty.
- v_i **is an internal vertex:** Simply repeat the leaf case argument for each edge incident to v_i. Each initial W-wedge at v_i will contain a superset of the empty region to its inside. Take the intersection of the insides of all of those initial W-wedges. That intersection must also be empty, and must at least contain the empty region.

□

Chapter 3

The Geodesic Star Unfolding

In this chapter, we consider an extension of the star unfolding, called the *geodesic star unfolding*, where the source is a geodesic curve instead of a point. This extension was discussed by Demaine and Lubiw [15]. They conjectured that a dual to the Sun Unfolding (analogous to the duality of Source and Star unfoldings) would be non-overlapping, and claim that proving the result for any open geodesic would be a likely first step. In this Chapter we prove some special cases of that conjecture, using a modified version of Lemma 2.2. In Section 3.1, we discuss the general form of the geodesic star unfolding and its associated ridge tree structure. In Section 3.2, we present Lemma 3.4, our modification to Lemma 2.2 to the geodesic setting, which is subject to some restrictions on the polyhedron/cut pairs it can be applied to. Section 3.4 uses Lemma 3.4 to prove a missing result of Itoh, O'Rourke and Vîlcu [21] for *quasigeodesic loops*. Section 3.5 considers fully-extended geodesics, which are cuts that have been maximally extended in either direction, and proves non-overlap for a subset of those cases.

3.1 Definitions

Definition 3.1. *Let $\mathcal{P}$ be a convex polyhedron, and let $\lambda = ab \in \mathcal{P}$ be a simple (non-self-intersecting) geodesic on the surface of $\mathcal{P}$. The* geodesic star unfolding S_λ *is a 2-dimensional polygon formed by cutting λ, as well as a shortest path along the surface of $\mathcal{P}$ from every vertex v_i of $\mathcal{P}$ to λ, and flattening the result into the plane. See Figures 3.1 and 3.2 for an example perspective view and the resulting unfolding.*

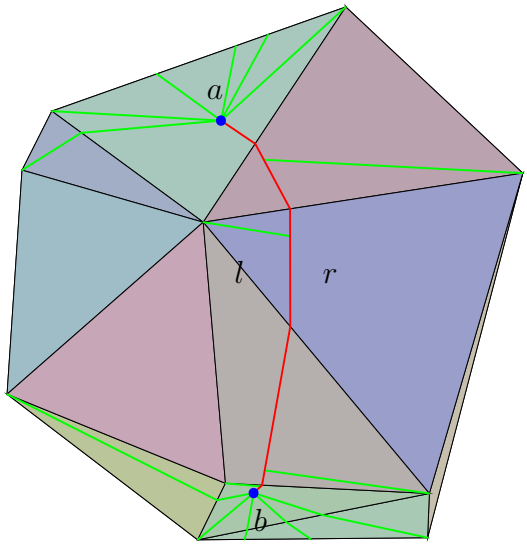

Figure 3.1: Shortest path cuts to each vertex from a geodesic, with a, b, l, and r labelled.

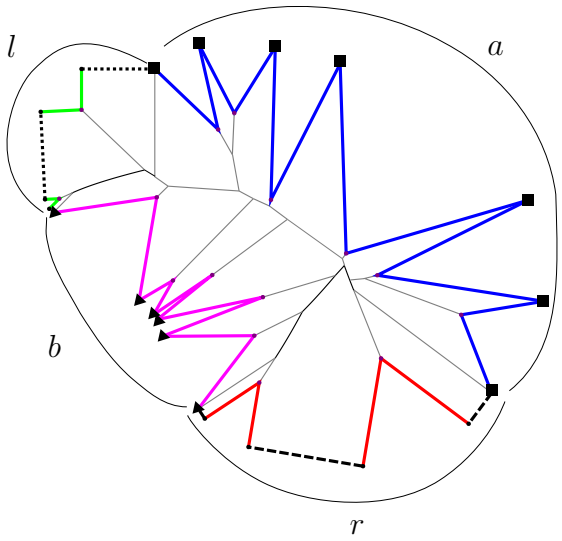

Figure 3.2: The geodesic star unfolding produced by the cuts from Figure 3.1, with the source images labelled.

3.1.1 Structure of the Geodesic Star Unfolding

The general structure of the geodesic star unfolding is more complex than the point case. Conceptually we will divide λ into four parts: its two endpoints a and b (for $a$bove and $b$elow), and its l and r ($l$eft and $r$ight) sides. This convention does not imply any specific orientation of the curve or the unfolding; it is strictly for notational purposes (any given example could be rotated or flipped to achieve any other possible labelling).

Shortest path cuts from the vertices can take on one of four cases (or rather, two pairs of two symmetric cases):

- The shortest path terminates at endpoint a or b, called *endpoint* cuts.
- The shortest path terminates perpendicular to an interior point of λ on either the l or r side, called *interior* cuts.

The fact that any shortest path which reports to an interior point along λ will intersect at a right angle is fairly intuitive, but we will state this lemma (adapted from Ieiri and Itoh [20]) for completeness.

Lemma 3.1. *Let λ be a geodesic curve on a convex polyhedron $\mathcal{P}$, $x \in \mathcal{P} \setminus \lambda$, and $x_0 \in \lambda$ be a point such that xx_0 is a shortest path from x to λ. The angle formed between xx_0 and λ is at least $\frac{\pi}{2}$ to each side.*

Proof. This is simply an application of Corollary 1 in [20]. □

A vertex is said to *report to* a, b, l or r respectively if the shortest path from the vertex terminates at that part of the geodesic cut.

A *source image* is either an endpoint (a or b), called a *point image*, or a sub-segment of λ corresponding to l or r, called a *segment image*, that lies between two consecutive shortest path vertex cuts. When we flatten the polyhedron after cutting from each vertex we obtain an alternating sequence of source images separated by vertex curvature gaps (each containing a vertex image) that forms a closed polygon (much the same as the point case). Note that all source images of the same type appear consecutively along the boundary of the unfolding (i.e. all a's, and then all l's, and so on). Also notice, in Figure 3.2 for example, that the source images at the transition points between the a and l, and the l and b sections of the unfolding can act as both endpoint images, and the start/end of a segment image. There are few restrictions on which types of cuts and/or source images will

exist for any given unfolding: the only requirement is that there must be at least two of the four types of cut present. Otherwise, there exist examples of all possible combinations.

Since the geodesic unfolding figures are a bit more complex, we will explicitly outline the conventions used in the figures (the colouring is helpful, but not necessary to understand the figures):

- Square dots represent images of endpoint a.
- Triangular dots represent images of endpoint b.
- Dotted lines represent interior sections of l.
- Dashed lines represent interior sections of r.
- Blue lines represent shortest path cuts from a.
- Pink lines represent shortest path cuts from b.
- Green lines represent shortest path cuts from l.
- Red lines represent shortest path cuts from r.

Before we move on, we present a useful fact regarding the contour of a geodesic star unfolding that we will draw upon in our analysis later on.

Lemma 3.2. *Consider a sequence of consecutive segment images along the boundary of an unfolding, where all segments correspond to the same side of the geodesic (i.e. either all l or all r). Join each consecutive pair of segment images as a chain in the following manner:*

- *If the curvature of the vertex between consecutive images is $\leq \pi$, simply join the consecutive endpoints of the two images by a straight line.*
- *Otherwise, if the curvature between them is $> \pi$, join them by the two edges of the unfolding between them (i.e. trace an edge to from each segment image to the vertex image in between them).*

This joined sequence forms a locally convex curve, i.e. a curve such that all bends to one side are $\leq \pi$ (see Figure 3.3).

Proof. This is simply due to the fact that all interior cuts are made at exactly right angles to the geodesic. In the first case, joining the two segments directly will give rise to two bends in the curve, each with interior angles in $[\frac{\pi}{2}, \pi]$. In the second case the joining will have 3 bends in the curve; the first and last bends will have angles $= \frac{\pi}{2}$, and the angle in between will be exactly 2π minus the curvature of the vertex, which will result in an angle $\leq \pi$. □

We will call the locally convex sequence formed by the construction of Lemma 3.2 the *convex chain* of a sequence of segment images.

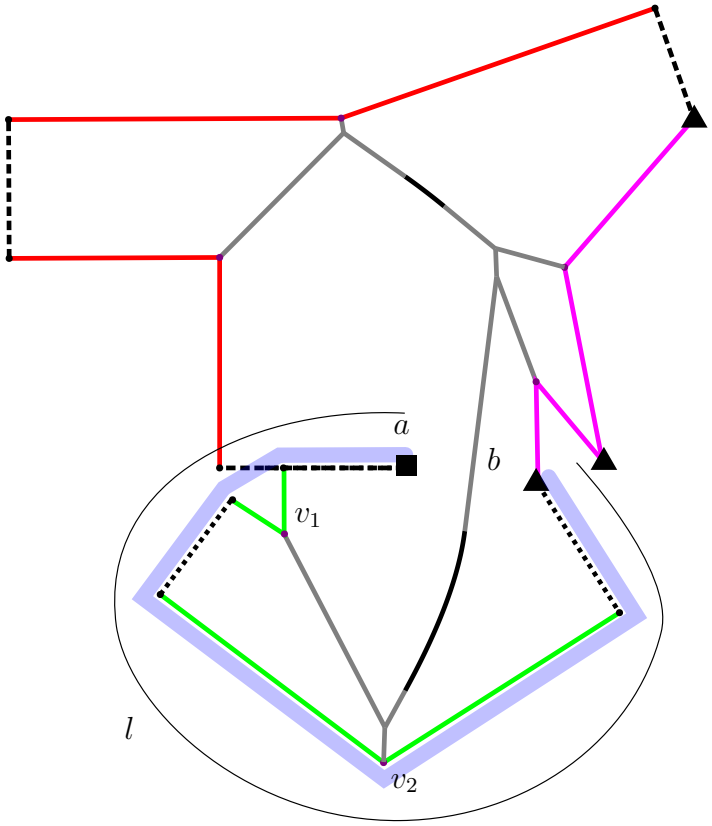

Figure 3.3: The convex chain for the l segment images is highlighted on the outside with a light shade. Vertex v_1 has low curvature ($\leq \pi$), so only produces two bends. Vertex v_2 has high curvature ($> \pi$), so it produces 3 bends.

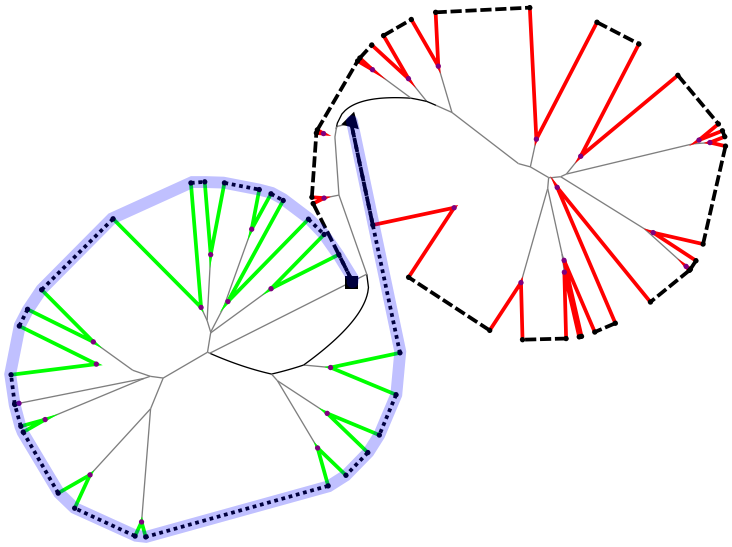

Figure 3.4: Note that although the chain is locally convex, globally it may not be.

This observation is not new; it was made by Itoh, O'Rourke and Vîlcu [21] (though in a slightly different form). Note that it does not guarantee that the chain itself lives on its own convex hull; rather, the sequence may 'spiral' in on itself (see Figure 3.4).

3.1.2 The Geodesic Ridge Tree

The ridge tree of a geodesic source λ is defined analogously to the point case: T_λ is the locus of all points on $\mathcal{P}$ with more than one shortest path to λ, plus all the vertices of $\mathcal{P}$. The cut locus from a geodesic source on a polyhedron is indeed a (finite) tree, as was established by Lemmas 4 and 5 of [15], thus it is correct to call this structure a ridge tree. An example ridge tree is drawn on the interior of Figure 3.2.

A key difference from the point case is that the edges of the geodesic ridge tree are no longer strictly straight-line segments. Recall that every edge of the ridge tree is the locus of points which are equidistant from two source images. Thus, when the two sources to either side of an edge are a point image and a segment image respectively, a parabolic ridge tree edge will result. The ridge tree edges between pairs of point or pairs of segment images will still be straight lines. Note that a segment image might include one of the endpoints of λ. For the purpose of our discussion, the endpoint is considered as its own, separate point image, rather than as part of the segment image.

Since ridge tree edges are no longer straight-line segments, the geodesic star unfolding can no longer be decomposed into kites as we did in the point case. We will instead define a decomposition into 'slices':

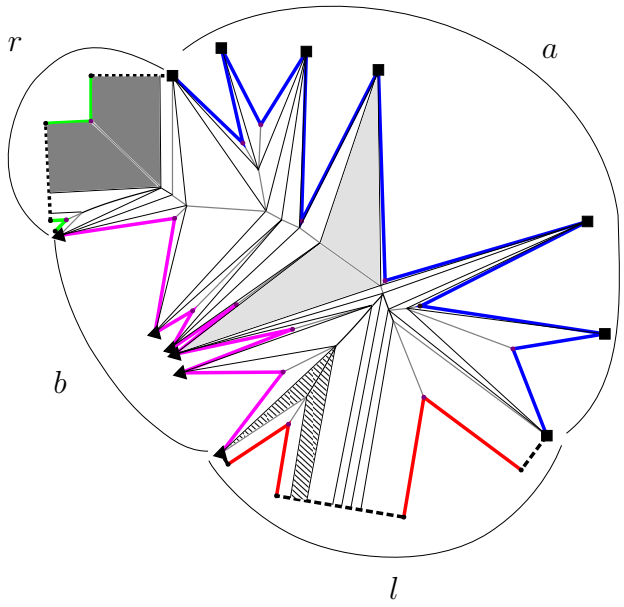

Figure 3.5: Decomposition of the unfolding from Figure 3.2 into slices. A *block-slice* is highlighted in dark-gray, *kite-slice* is highlighted in light-gray, and a *hybrid-slice* is highlighted with a lined pattern.

Definition 3.2. *Given an edge e of a geodesic ridge tree T_λ, its* slice*, denoted slice(e), is the union of all shortest paths from e to all source images of λ it reports to. We divide the slice to either side of e to obtain its* wings*. A wing may be either a* point-wing *or* segment-wing*, depending on whether the corresponding source image is a point image or a segment image. This results in slices that may take on one of three forms:*

- *slice(e) is a* kite-slice *if both wings are* point-wings *(light-gray region in Figure 3.5).*
- *slice(e) is a* block-slice *if both wings are* segment-wings *(dark-gray region in Figure 3.5).*
- *slice(e) is a* hybrid-slice *if one wing is a* point-wing *and the other is a* segment-wing *(black region in Figure 3.5).*

Along with this new definition, we also extend our definition of the *source angle* of an edge of the ridge tree.

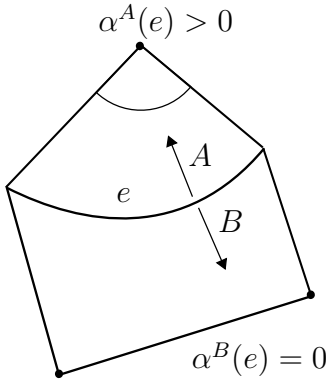

Figure 3.6: The source angle of the wings to either side of ridge tree edge e may differ.

Definition 3.3. *The source angle of a point-wing is the interior angle at its apex, while the source angle of a segment-wing is simply 0 (see Figure 3.6 for an example).*

We also extend our notation to cover paths through the ridge tree as well. Take path $\sigma = e_0, \ldots, e_t$, and arbitrarily label its two sides A and B respectively. Then the A source angle of σ is $\alpha^A(\sigma) = \sum_{i=0}^{t} \alpha^A(e_i)$ (and likewise for B).

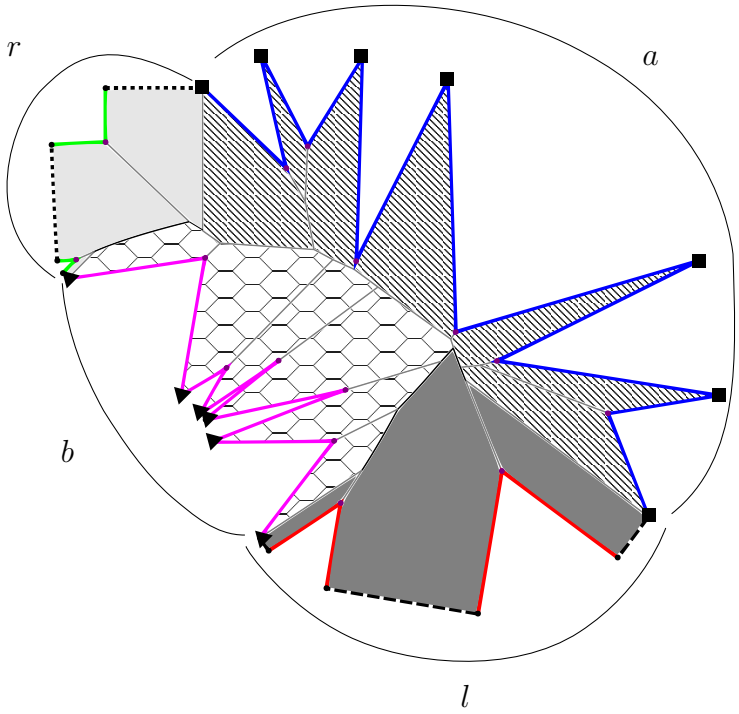

Figure 3.7: Example of an ab-adjacent ridge tree.

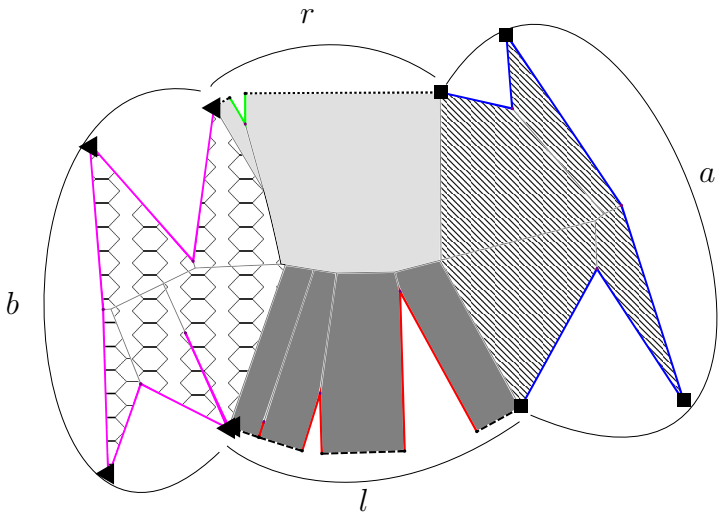

Figure 3.8: Example of an lr-adjacent ridge tree.

We can break down the unfolding into four pieces, by taking all regions of the unfolding containing shortest paths to a, b, l and r respectively (illustrated by Figures 3.7 and 3.8).

Observation 3.1. *Each region of the unfolding, corresponding to the pieces containing shortest paths to a, b, l or r respectively, must form a single, connected region (though they might be connected by a single point at certain places). We call these the* a region*,* b region*,* l region*, and* r region *respectively.*

This is simply because all source images of a particular type appear consecutively around the boundary of the unfolding, and the regions of shortest paths corresponding to individual source points themselves are single, connected regions. The boundary between the regions in their rotational order is a shortest path to the ridge tree from either a or b that is perpendicular to λ. To one side of each shortest path is an endpoint-reporting region (i.e. reporting to a or b), and to the other is an interior-reporting region (reporting to l or r). Thus the a and b regions each share a boundary with each of the l and r regions.

It is also possible for the a and b regions to share some edges of the ridge tree, or likewise the l and r regions to share edges of the ridge tree, though not both at the same time. Figure 3.7 is an example of the former case, and Figure 3.8 is an example of the latter. The other possibility is an extremely degenerate case where a, b, r and l share only a single vertex of the ridge tree, though we can treat this as one of the other two cases with an extremely small shared edge. This leads to the following definitions to characterize a given ridge tree:

Definition 3.4. *A ridge tree is called* ab-adjacent *if the a and b regions share an edge of the ridge tree,* lr-adjacent *if the l and r regions share an edge of the ridge tree, or* $ablr$-adjacent *if all four regions share a common vertex of the ridge tree.*

3.2 Applying the Point Case Proof to the Geodesic Setting

In this section we modify the proof of non-overlap of the star unfolding from a point source to some cases of the unfolding from a geodesic curve.

Note that in light of Definition 3.3, we must contend with a key fact: the source angle of the wings to either side of a ridge tree edge are no longer (necessarily) equal (see Figure 3.6). This in turn implies the source angle to one side of a path through the ridge tree T_λ is not necessarily at most π. Note that the sum of the source angles of all wings is still at most

2π, since the source angles of all point wings reporting to a and b respectively sum up to π each.

The proof proceeds in much the same way as Lemma 2.2, only we must keep track of the source angle to either side of the path through the ridge tree independently. We must give an extended definition of a *W-wedge* at vertex r of edge y from Section 2.2 to track the wedge angles between its inner and outer legs to either side of the edge separately. Thus, assuming we've already labelled the A and B sides of edge y, γ^A (resp. γ^B) is the angle of rotation of the inner leg from the outer leg on the A (resp. B) side of edge y.

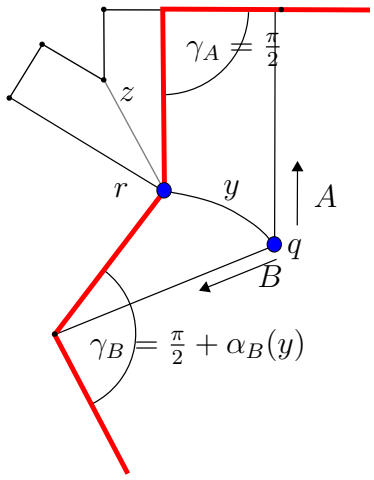

Figure 3.9: Example base case with a hybrid slice. Note the W-wedge is only rotated more than $\frac{\pi}{2}$ on the point-wing side of the edge.

Now, we give a lemma analogous to Lemma 2.1 concerning adjacent slices to either side of a W-wedge. The proof is nearly identical to that of Lemma 2.1, except that we must take into account the possibility of asymmetry in the slices.

Lemma 3.3. *Let W be a W-wedge rooted at vertex r on edge y of a geodesic ridge tree, and let z be some other ridge tree edge incident to r. Suppose W has angles $\gamma^A \in [\frac{\pi}{2} + \alpha^A(y), \frac{3\pi}{2} - \alpha^A(z)]$, and likewise $\gamma^B \in [\frac{\pi}{2} + \alpha^B(y), \frac{3\pi}{2} - \alpha^B(z)]$. Then slice(y) $\subseteq W$ and slice(z) $\subseteq \overline{W}$.*

Proof. Consider Figure 3.9. The two inner legs of W lie exactly along the edges of *slice*(y) incident to r, and the outer legs are rotated away from the inner legs by at least $\frac{\pi}{2} + \alpha^A(y)$ (resp. $\frac{\pi}{2} + \alpha^B(y)$). Thus, the outer legs are at least perpendicular to the edges of the wings

of *slice*(y) that are incident to the other endpoint of y (point q in Figure 3.9), regardless of whether they are segment wings (with source angle zero), or point wings (with non-zero source angle). Thus, the two rays either diverge, intersect and enclose *slice*(y) (left image of Figure 2.5), or intersect such that the *slice*(y) lies in an unbounded region. Therefore, *slice*(y) is to the inside of W. The argument is symmetric to show *slice*(z) is to the outside of W. □

We also give a lemma which extends the idea of Lemma 2.2, taking into account the limitations of the source angle on either side of the path.

Lemma 3.4. *Given two distinct edges u and v of the ridge tree T_λ of some geodesic star unfolding S_λ, the two associated slices slice(u) and slice(v) do not overlap, as long as over the path σ between u and v, the source angle of the path to each side is at most π.*

Proof. Let $\sigma = p_0, \ldots, p_t$ be the vertex-path through the ridge tree σ from u to v.
Let e_i be the edge of the ridge tree from p_{i-1} to p_i, such that $e_1 = u$, $e_t = v$.
Let $s_i = \mathit{slice}(e_i)$, be the i^{th} slice along the path.
Label the two sides of path σ A and B respectively, and define the A and B wings of each slice to be the wing lying to the appropriate side.
Let $\alpha_i^A = \sum_{j=1}^{i} \alpha^A(e_j)$
Let $\alpha_i^B = \sum_{j=1}^{i} \alpha^B(e_j)$
Define W_i, for $i = 1, \ldots, t$, to be the W-wedge rooted at p_i on edge e_i, with inner legs along the two incident edges of s_i, and outer legs rotated out by $\gamma_i^A = \alpha_i^A + \frac{\pi}{2}$, and $\gamma_i^B = \alpha_i^B + \frac{\pi}{2}$ respectively.

Once again, we claim $W_{i-1} \subseteq W_i$, and $s_i \subseteq (W_i \setminus W_{i-1})$, for $i = 2, \ldots, t$. Apply Lemma 3.3 to show $s_i \subseteq \overline{W}_{i-1}$. Then, apply the transformation shown in Figure 3.10. Notice that the value of γ^A did not change in the example in the figure, since it traversed a segment-wing which has source angle zero (indeed, the second phase of the transformation on that side is simply a translation). Either way, the newly established wedge still covers a strict superset of W_{i-1} in its interior region, including s_i. By induction, this will hold as long as both $\alpha_i^A \leq \pi$ and $\alpha_i^B \leq \pi$. □

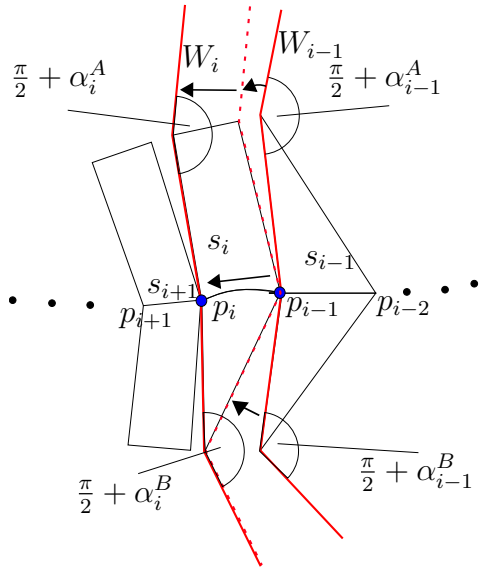

Figure 3.10: Sample induction step for a geodesic unfolding. Note that $\alpha_{i-1}^A = \alpha_i^A$ since the image to that side is a segment.

Definition 3.5. *Given ridge tree T_λ of some geodesic star unfolding S_λ, T_λ is* balanced *if for both sides of every path through the ridge tree the sum of the source angles is at most π.*

A geodesic star unfolding whose ridge tree is balanced is guaranteed by Lemma 3.4 to unfold without overlap. Unfortunately, it is not true that all geodesic ridge trees are balanced. Moving forward, we want to identify cases of the geodesic star unfolding which are balanced, or, in the case they are not, if other mitigating factors can be used to prove non-overlap.

3.3 Extension to Quasigeodesic Curves

In this section, we discuss extending the results of Section 3.2 to quasigeodesic curves. Although for simplicity we have restricted our discussion to the star unfolding from geodesic curves, all of our results apply to the star unfolding from a quasigeodesic curve as well. Recall that a quasigeodesic curve on the surface of a polyhedron $\mathcal{P}$ is a curve such that at each point along the curve the surface angle to each side is $\leq \pi$. Consider a quasigeodesic

curve λ on the surface of $\mathcal{P}$. We define the *quasigeodesic star unfolding* analogously to Definition 3.1; simply replace 'geodesic' with 'quasigeodesic'.

Suppose p is an interior point of the quasigeodesic curve λ where the surface angle to one side, say the r side, is α, where $\alpha < \pi$. Necessarily, p is a vertex of the polyhedron, otherwise the surface angle on the other side of the curve would be greater than π. We do not introduce an extra cut for this vertex in the unfolding, since it already lies on λ.

Using Lemma 3.1, we claim that no shortest path cut from any vertex v will report to point p on side r, since one of the two surface angles formed would be $< \frac{\pi}{2}$. Thus the quasigeodesic star unfolding from λ will have a vertex image with an angle α corresponding to the r side of p. We will call this vertex image in the unfolding p_r.

Observe that p_r is a leaf of the ridge tree and that the incident ridge tree edge e is a straight segment forming angles $\frac{\alpha}{2}$ with the segment images to either side of p_r. Thus e has two segment wings, and the only thing that distinguishes these wings from those that arise in the geodesic star unfolding is that each segment wing has a side of 0 length at both sides incident to p_r.

It remains to show that Lemmas 3.3 and 3.4 also apply to quasigeodesics.

Lemma 3.5. *Lemmas 3.3 and 3.4 hold, even when λ is a quasigeodesic curve.*

Proof. Lemma 3.3 holds since the argument does not rely on the size of the curvature gap on either *slice*(y) nor *slice*(z); even with side lengths of 0, the partitioning by the W-wedge will still hold for the range of γ specified. Lemma 3.4 also does not pose a problem, simply because the quasigeodesic edges are always incident to a leaf of the ridge tree, and thus may only appear as the first or last elements in a sequence of ridge tree edges. Thus, the induction will never need to apply to the 0 length sides of the quasigeodesic slices, and the rest follows. □

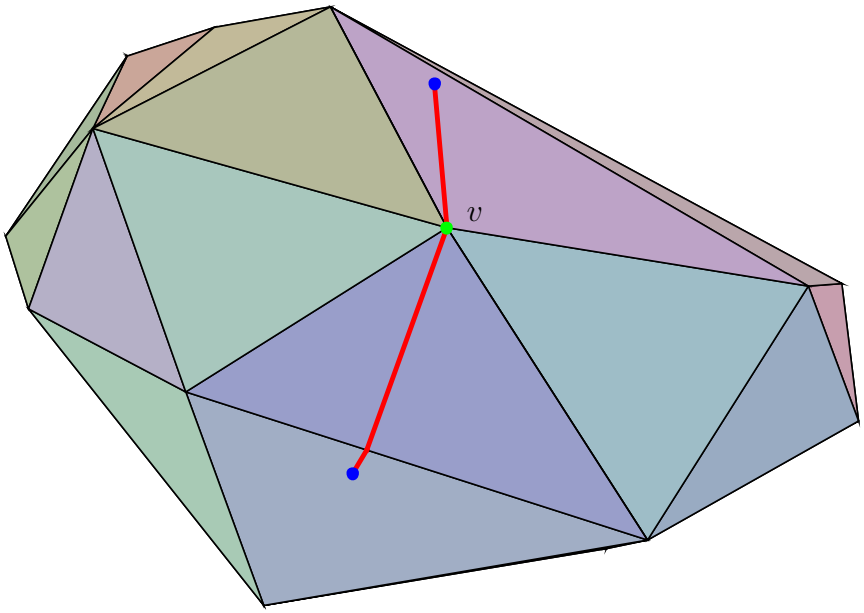

Figure 3.11: Perspective view of a quasigeodesic curve, passing through vertex v.

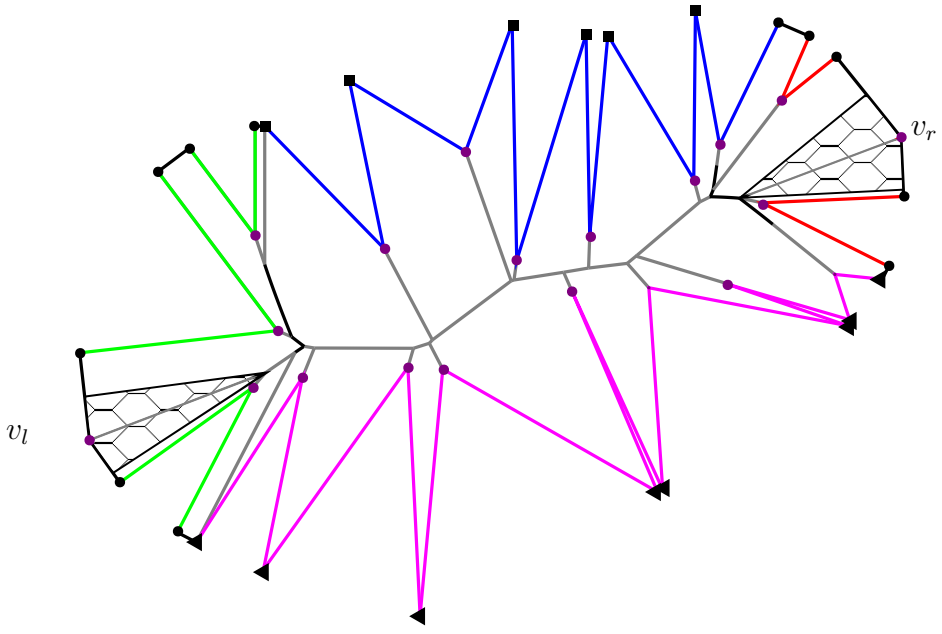

Figure 3.12: Quasigeodesic star unfolding of the curve from Figure 3.11. The two images of v and their associated slices are indicated.

3.4 Quasigeodesic Loops

In this section, we discuss non-overlap of the star unfolding from geodesic (and quasigeodesic) loops. Recall from Chapter 1 that when the two endpoints a and b of a geodesic or quasigeodesic curve λ coincide at point o, we call this a *(quasi-)geodesic loop* with *loop point* o. A (quasi-)geodesic loop cuts the surface of the polyhedron into two pieces. One piece must have a surface angle at o that is $\leq \pi$, and we call this the *inside* of the loop and in this section we will identify it with the r side of the curve. The other piece is called the *outside* and will be identified with the l side of the curve.

Itoh, O'Rourke, and Vîlcu [21] proved that for any quasigeodesic loop: (1) the inside unfolds without overlap; (2) the outside unfolds without overlap; and (3) the two unfolded pieces can be reattached (without overlap) along a common segment of the cut curve. Their proof of (2) relies on a Lemma [21, Lemma 7] about the star unfolding from a point, which they say will be proved in a companion paper.

Unfortunately, they discovered[1] that the Lemma is false. The Lemma claims that for for any star unfolding from a point x and for any polyhedron vertex v, the exterior angle at v in the unfolding yields an empty wedge. More precisely, if the exterior angle at v is (counterclockwise) x_i, v, x_{i+1} then the claim is that the counterclockwise wedge formed by the rays vx_i and vx_{i+1} does not contain any part of the unfolding. An example where this fails is shown in Figure 3.13.

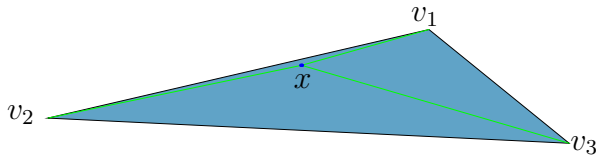

Figure 3.13: A counterexample occurs when we take a source point near or on the edge of a doubly-covered triangle.

[1] Private communication from J. O'Rourke

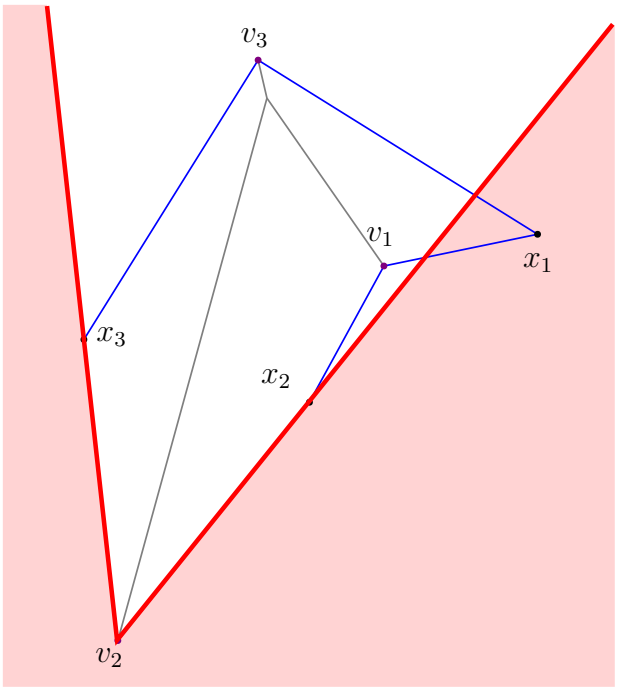

Figure 3.14: The unfolding of Figure 3.13. The exterior angle wedge formed by x_2, v_2, x_3 is not disjoint from the rest of the unfolding.

In this section we will consider the star unfolding from a geodesic curve where the two endpoints a and b are arbitrarily close together. In the limit when $a = b$ the unfolding consists of two pieces joined at the point $a = b$ with the angular bisectors at the point $a = b$ aligned in the unfolding. We call this the *conjoined star unfolding* from a geodesic loop. Our main result is that the conjoined star unfolding from a geodesic loop does not overlap. As shown by Lemma 3.5, our arguments for non-overlap of the geodesic star unfolding extend to the case of quasigeodesic source curves as well, so in that sense our result is no less general than [21]. Our result also implies that the outside piece unfolds without overlap, which repairs the missing step of Itoh, O'Rourke and Vîlcu's result.

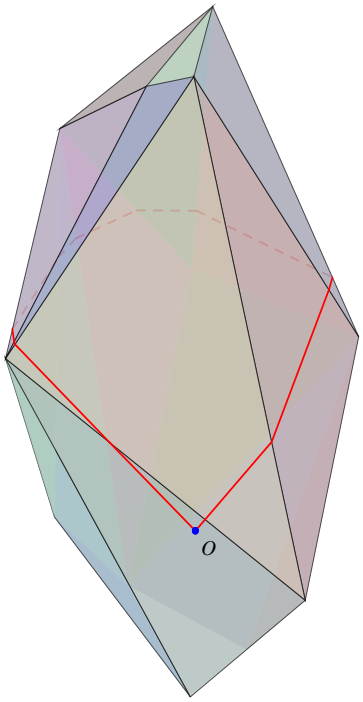

Figure 3.15: Geodesic loop on the surface of a polyhedron (partially transparent to view the full loop).

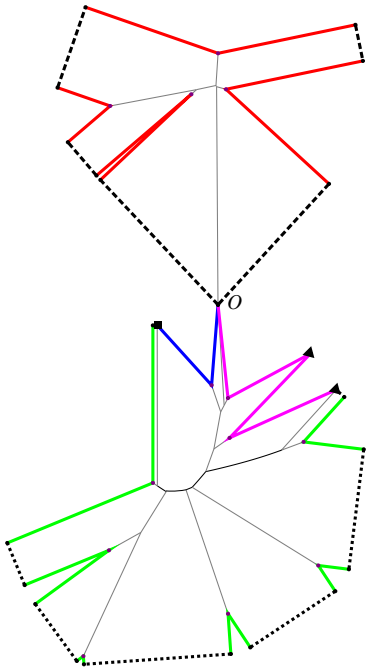

Figure 3.16: Unfolding of Figure 3.15.

An example of a geodesic loop and corresponding unfolding are given in Figures 3.15 and 3.16. Note that endpoint reporting vertices can only lie on the outside of the loop. No vertex inside of the loop can report to the loop point o itself, since the surface angle to the interior of the loop is $< \pi$.

Theorem 3.1. *The conjoined star unfolding from a (quasi-)geodesic loop does not overlap.*

Our proof of Theorem 3.1 will show that the ridge tree for a quasigeodesic loop is balanced, meaning we can apply Lemma 3.4 to prove non-overlap.

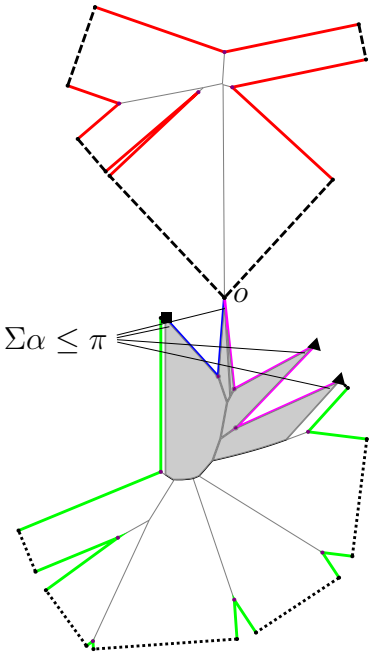

Figure 3.17: The total source angle not covered by u on a geodesic loop star unfolding is at most π.

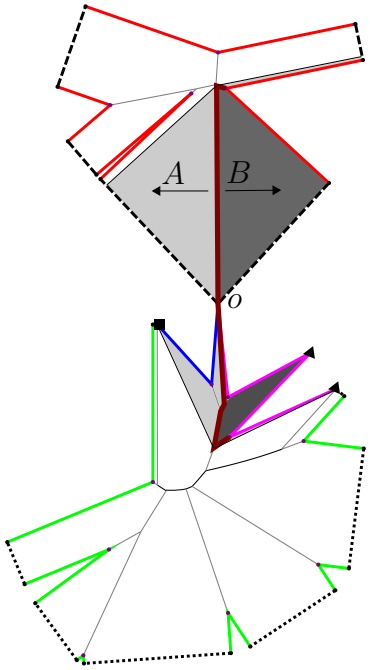

Figure 3.18: Path σ through the ridge tree that crosses u (shaded brown, A and B wings of the path shaded separately).

Lemma 3.6. *Every path through the ridge tree of the conjoined star unfolding from a (quasi-)geodesic loop has at most π source angle to either side.*

Proof. Without loss of generality, we will assume the convention established at the top of this section, namely that the r side of the (quasi-)geodesic is the one facing the inside (lower curvature) part of the loop, and l is on the outside (higher curvature).

The idea is to look at the segment of the ridge tree that lies between a and b as they approach each other. In the limit, there will be a single edge of the ridge tree, call it u, with at least $\frac{\pi}{2}$ source angle to either side passing through o. Note that o in this context is the eventual point of intersection if we extended the endpoints a and b in the direction of the (quasi-)geodesic curve.

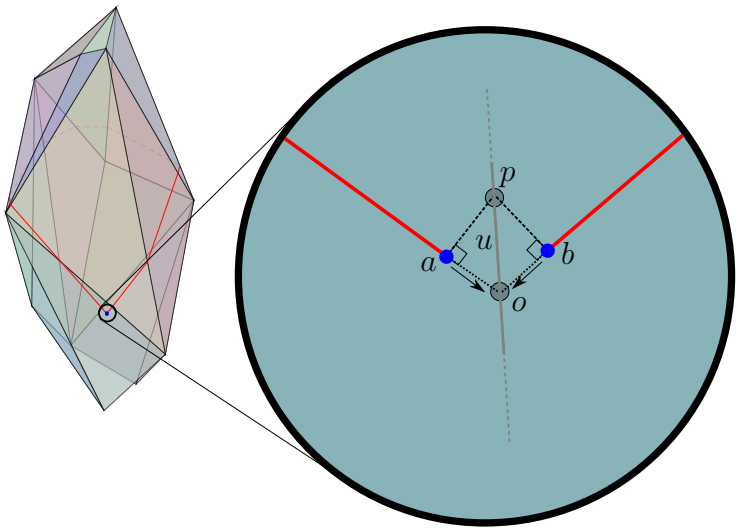

Figure 3.19: Zoomed-in view of the surface as a and b approach o. Assuming no vertices are inside the region $aobp$, there is at least $\frac{\pi}{2}$ source angle to either side of ridge tree edge u.

The ridge tree for a geodesic loop is necessarily ab-adjacent. Consider the segment of the ridge tree between a and b on the inside of the loop that touches loop point o, call it u. Observe what happens in the limit as a and b approach o. Consider the kite-shaped region of the surface delimited by a, b, o, and p, where p is the intersection between the rays perpendicular to the geodesic at a and b respectively (see Figure 3.19). Call this region ρ. Assuming we are zoomed in enough, there are no vertices or other ridge tree edges inside the region ρ, and therefore some sub-segment of u will reach from p to o (that is, exactly bisect a and b). Thus, we must conclude that $\alpha(u) \geq \frac{\pi}{2}$. Furthermore, this is the only edge of the ridge tree on the inside of the loop that has point-wings reporting to a or b. Consider any path σ through the ridge tree of the geodesic loop unfolding; it is restricted to one of the three following cases:

- The path σ does not include u, and remains entirely on the inside of the loop. In such a case, since we know u is the only edge with a non-zero source angle, it is guaranteed that $\alpha^A(\sigma) = \alpha^B(\sigma) = 0$.
- The path σ does not include u, and remains entirely on the outside of the loop. Since we proved that $\alpha(u) \geq \frac{\pi}{2}$, the total source angle of the wings to both sides of edge u $\geq \pi$. Thus, the remaining source angle of all wings along every other possible path

(i.e. not including u) must be $\leq \pi$ (see Figure 3.17), thus the path as a whole cannot have more than π source angle to either side.

- The path σ includes u. We want to show that the source angle to each side of σ is at most π. Any edge of σ that lies inside the geodesic loop only has segment wings to either side and these contribute 0 to the source angle of the path. Thus it suffices to look at the portion of σ starting at e and following edges that lie outside the geodesic loop. Call this subpath σ'. Ridge tree edge u has a point wing to either side, one reporting to a and one reporting to b. Label the two sides of σ' A and B. Suppose that the source angle of σ' on side A is greater than or equal to the source angle on side B. If side A only has point wings that report to a, then its source angle is at most π. So suppose that side A has a point wing that reports to b. As we walk around the path σ', the wings report in order to a, then l, then b (this simply a consequence of Observation 3.1, see Figure 3.20 for an example). So all the wings on the B side must report only to b, and the sum of their source angles is at most π. Since every wing on the B side is a point-wing reporting to b, every point wing on the A side must be paired with a point wing on the B side, thus each such pair has equal source angles. Thus the source angle of σ' on side A is bounded by the source angle of σ' on side B, which is bounded by π.

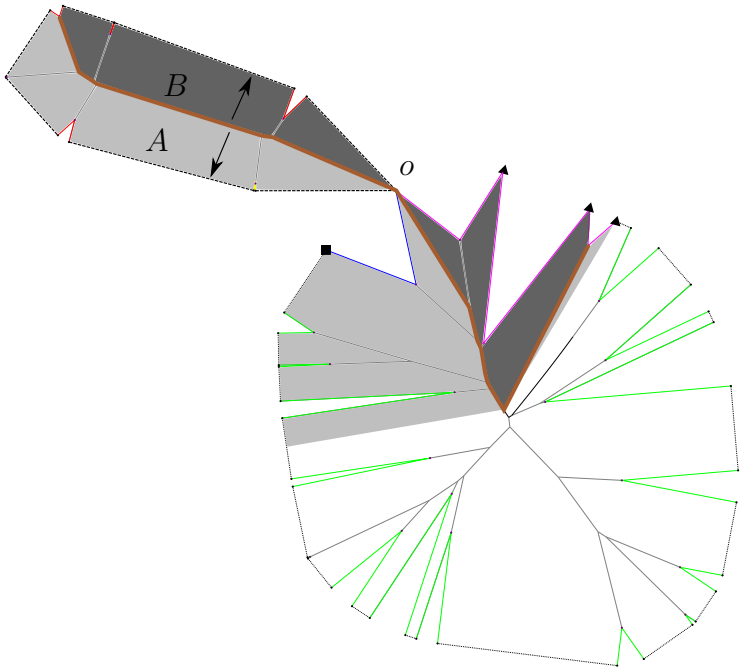

Figure 3.20: An illustration of the 3rd case, where the A side of the path has wings reporting to a, l, and b on the outside of the loop.

□

Proof of non-overlap then follows as a corollary to Lemma 3.6.

3.5 Fully Extended Geodesics

We proved in Lemma 3.4 that the geodesic star unfolding does not overlap when every path in the ridge tree has source angle at most π to each side (i.e. the ridge tree is balanced). In order to prove the complete conjecture that the star unfolding does not overlap, we must investigate cases where there are paths in the ridge tree which have source angle greater than π to one side. In this section, we explore one such case where we can easily identify if this property holds, when the geodesic cut λ has been *fully-extended*. This case is simpler in that no vertices report to the endpoints a and b of the geodesic curve. We will identify cases of fully-extended geodesic curves where the ridge tree is not balanced, and some sub-cases within those where we can still prove non-overlap.

Imagine taking a point $p \in \mathcal{P}$, and choosing two opposite directions (that is, rotated π away from each other), and then extending the endpoints of a geodesic curve λ from p as much as possible in both directions. Typically, it is not possible to continue extending this geodesic indefinitely — eventually there will be a point of self-intersection. The only exception to this rule is the case of an isosceles tetrahedron, where it is possible to extend a geodesic in either direction infinitely [21]. Note that the curve generated by this construction might not be unique. Rather, depending on which endpoint was 'extended' first, one path may block the other (see Figure 3.21).

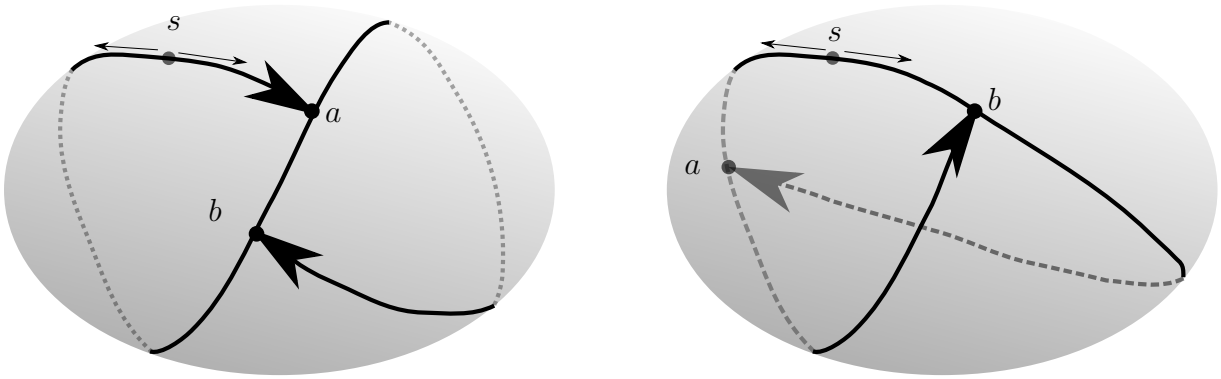

Figure 3.21: Schematic view of two fully-extended geodesics arising from the same origin point s.

In general, fully-extended geodesics will contain two *geodesic loops*, one for either endpoint. We characterize fully-extended geodesics in terms of where the two endpoints intersect with the rest of the geodesic to form these loops. Call the point of intersection of a (resp. b) with λ the *loop point*, denoted $\lambda(a)$ (resp. $\lambda(b)$). In order to maintain the property that the unfolding remains a single connected polygon, we will consider a (resp. b) to lie infinitesimally close to $\lambda(a)$ (resp. $\lambda(b)$) without actually disconnecting the surface.

We begin by examining the local structure of the unfolding in a neighbourhood of one endpoint of the geodesic curve λ. Without loss of generality, we express the following claim in terms of endpoint a reaching side r of λ.

Claim 3.1. *In the star unfolding from a fully-extended geodesic, the neighbourhood around point a (resp. b) has the following properties (as shown in Figure 3.22)*

- *No vertex cut reports to the endpoint a.*
- *Near $\lambda(a)$ the curve λ unfolds to a straight line.*

- *Near a the two sides of λ will come together in the unfolding.*

Furthermore, if $\beta(a)$ is the surface angle at point $a = \lambda(a)$ on the r side of λ (i.e. on the inside of the geodesic loop), then the curvature of the piece of the surface inside the loop is $\beta(a) + \pi$.

Proof. The only thing we need to justify is that no vertex cut reports to endpoint a itself, the other two facts follow from it. Since a and $\lambda(a)$ effectively coincide, the surface angle to either side of that point is $\leq \pi$. By Lemma 3.1, the angle α made between a shortest path to λ and λ itself must be $\geq \frac{\pi}{2}$, and so cannot report to that exact point (see the lower magnification in Figure 3.22). The final point regarding the curvature inside the geodesic loop was established in [21].

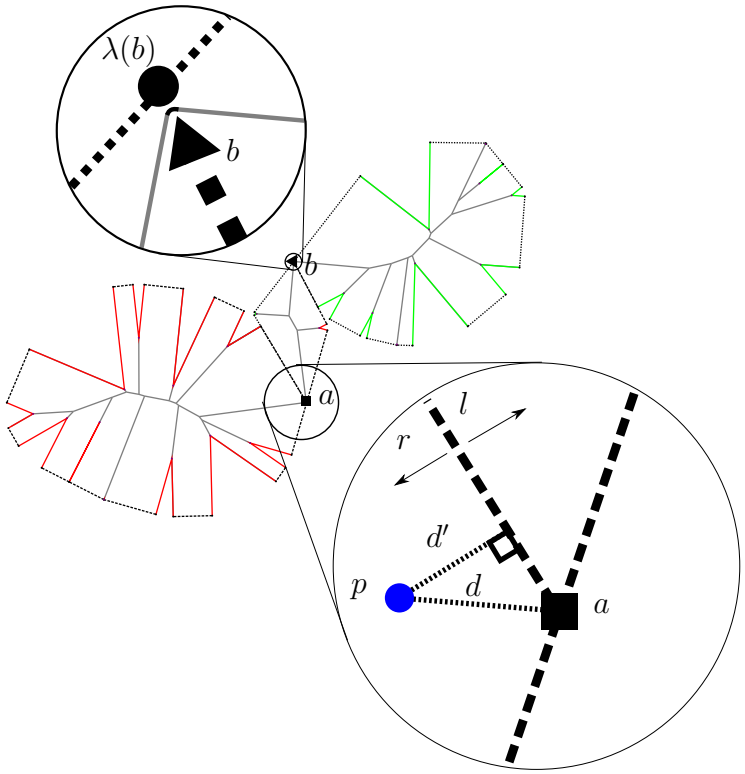

Figure 3.22:
Region around a: Pick any point p near endpoint a. No shortest path from p can report to the endpoint, since a shorter perpendicular path can always be found instead.
Region around b: A ridge tree edge passes between the infinitely small gap between b and $\lambda(b)$.

□

Note that we consider a (resp. b) to be infinitely close to, but not touching, $\lambda(a)$ (resp. $\lambda(b)$). Thus, there is a ridge tree edge which passes between a and $\lambda(a)$ (likewise for b, see the upper magnification in Figure 3.22). This edge has a point wing with source angle π on one side. These are the only two point-wings in a fully-extended geodesic ridge tree. The type of wing which lies on the other side of those ridge tree edges determines which class of fully-extended geodesics this instance falls into.

All geodesics, when fully extended, fall into one of the following 3 categories:

- *Closed Geodesic*: This is the somewhat degenerate case where the two endpoints will meet each other exactly 'head-on' (i.e. $\lambda(a) = b$, and vice versa), to form a perfect closed geodesic (with no distinguishable loop point).
- *S-Shaped Geodesic*: This is the case when a and b reach *opposite* sides of the geodesic curve, that is, one to l and one to r (see Figure 3.23 for a schematic, and Figures 3.24 and 3.25 for a full example).
- *C-Shaped Geodesic*: This case arises when a and b reach *the same* side of the geodesic (see Figure 3.23 for a schematic, and Figures 3.26 and 3.27).

For the remainder of this chapter, for the sake of brevity we will refer to fully-extended C- and S-shaped geodesics simply as C-shaped and S-shaped geodesics respectively.

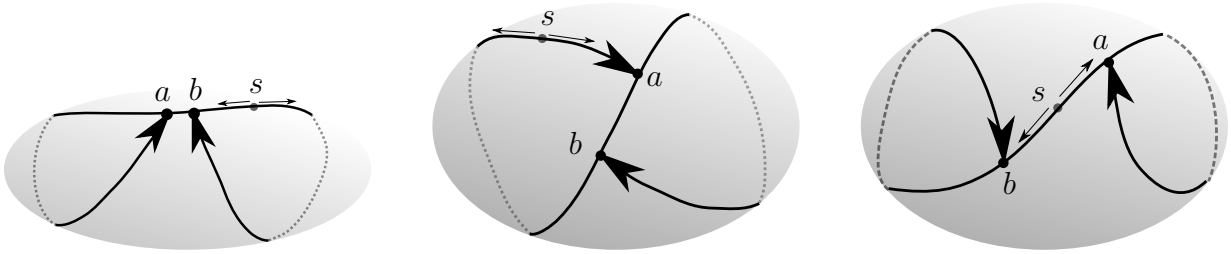

Figure 3.23: Schematic view of a C-shaped geodesic (left), and an S-shaped geodesic (centre). Also shown is an S-shaped geodesic where the two loops do not share any piece of the geodesic curve (right).

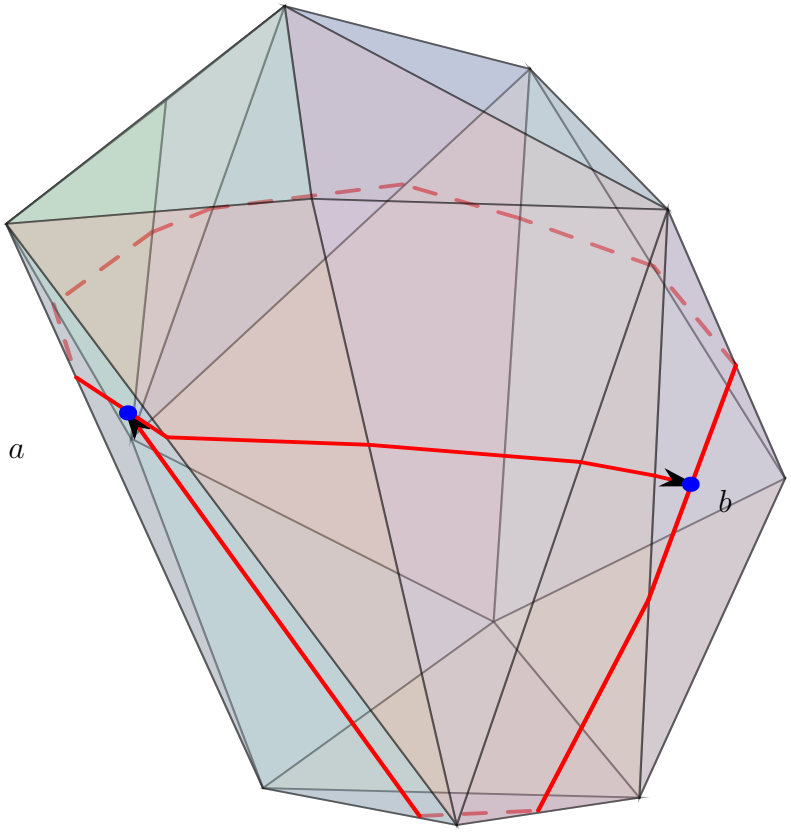

Figure 3.24: S-Shaped geodesic on a polyhedron (partially transparent to view the full curve). The arrows are directed to the endpoints, a and b, of the geodesic curve.

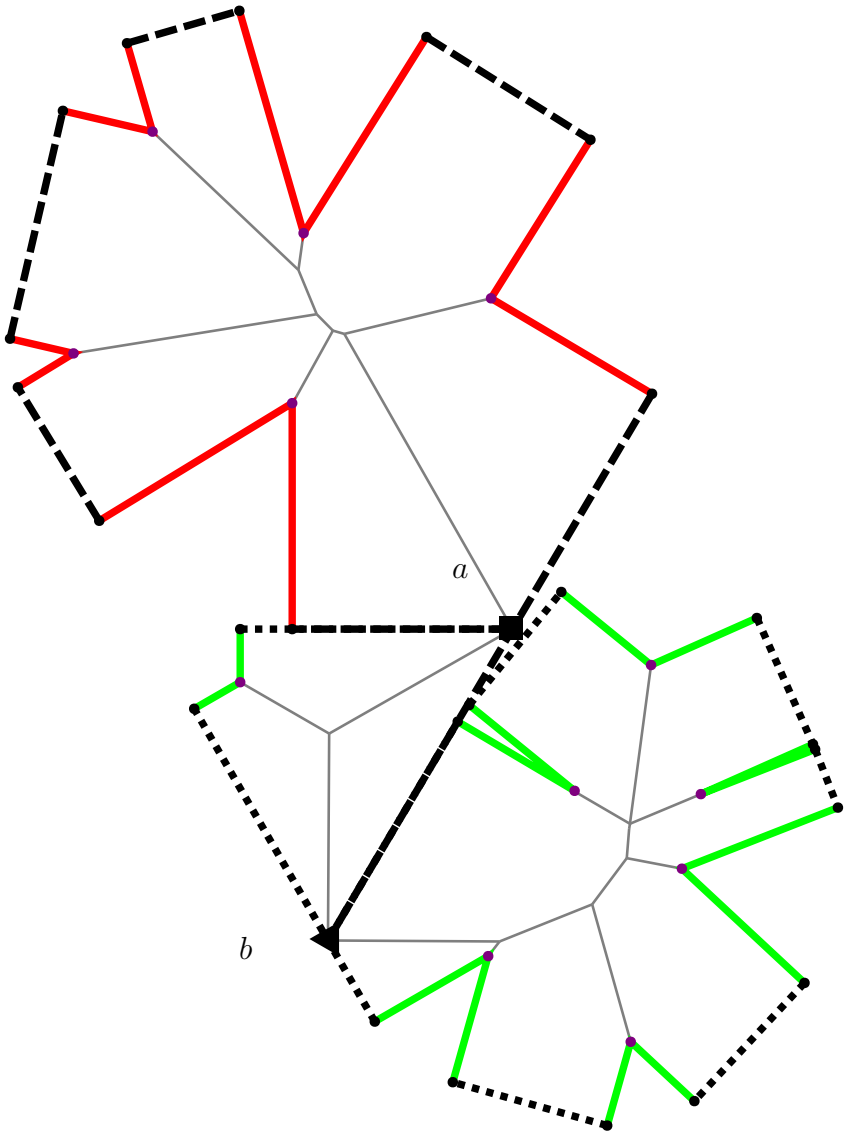

Figure 3.25: Unfolding of Figure 3.24. Note that a and b appear on opposite sides of every path possible path through the ridge tree.

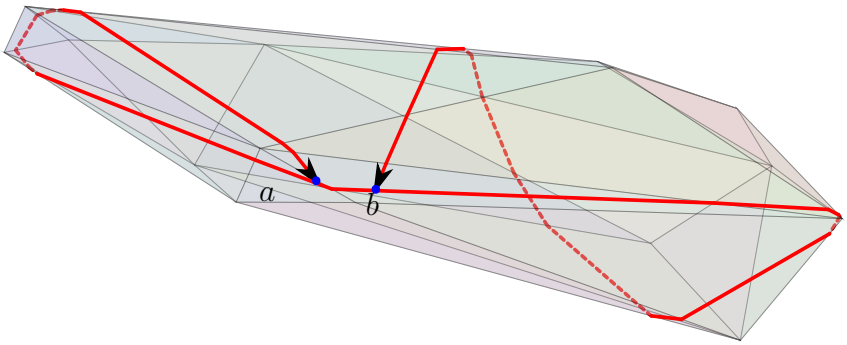

Figure 3.26: C-Shaped geodesic on a polyhedron (partially transparent to view the full curve).

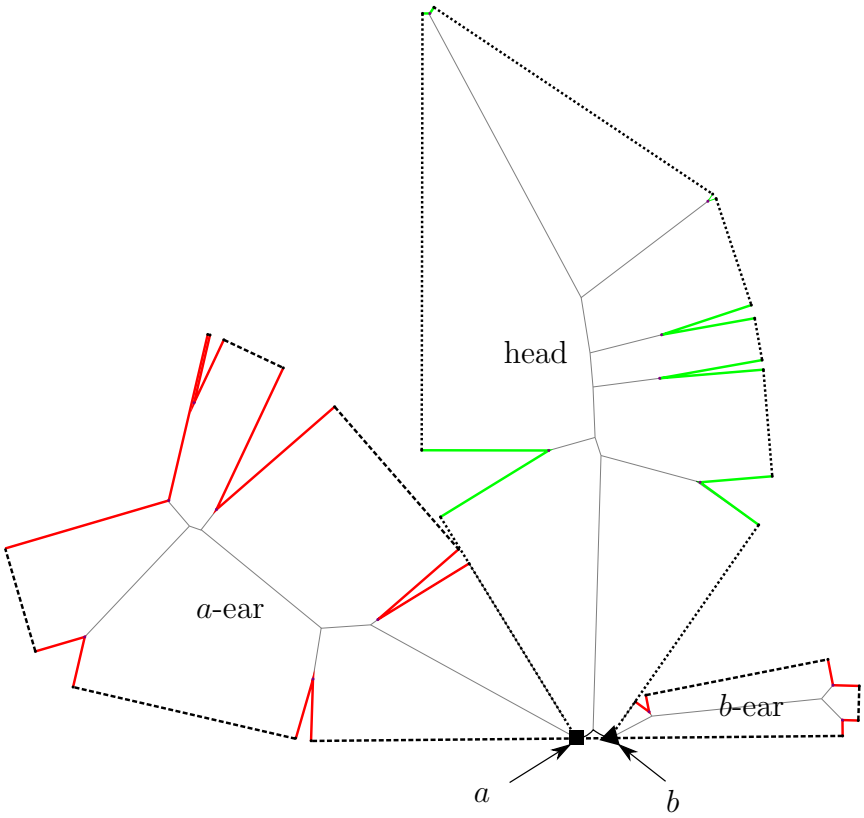

Figure 3.27: Unfolding of Figure 3.26. a and b are on the same side of any path passing from the interior of one ear to the other.

We will show that for a closed geodesic or an S-shaped geodesic, the ridge tree is always balanced, so the star unfolding is non-overlapping by Lemma 3.4. However, for the C-shaped geodesics, the ridge tree can be unbalanced—we explore this case further in Section 3.5.1.

First, we will look at closed geodesics.

Lemma 3.7. *The star unfolding from any fully-extended closed geodesic curve is non-overlapping.*

Proof. Since a closed geodesic is just a geodesic loop, this is just a special case of Lemma 3.6, □

Now, we give a simple proof for the non-overlap of fully-extended S-shaped geodesic curves. The key observation is that since the loop points appear on opposite sides of the geodesic curve, they will also appear on opposite sides of any path through the ridge tree, and thus the ridge tree is balanced.

Lemma 3.8. *The star unfolding from any S-shaped geodesic curve is non-overlapping.*

Proof. Consider a path through the ridge tree σ, and take one side of that path, say the A side. We want to show that $\alpha^A(\sigma) \leq \pi$. There are only two point-wings in the unfolding, one for source a and one for source b, and each of these point wings has a source angle of exactly π. The only way to have $\alpha^A(\sigma) > \pi$ is for both point wings to appear on the same side of σ. Suppose that this is the case. Since this is an S-shaped geodesic, a and b reach opposite sides of λ. We claim that this contradicts the property that the source images appear in cyclic order a, l, b, r around the boundary of the unfolding, since that same cyclic ordering (or a subsequence of it) would also have to occur around any path through the ridge tree (see Figure 3.28).

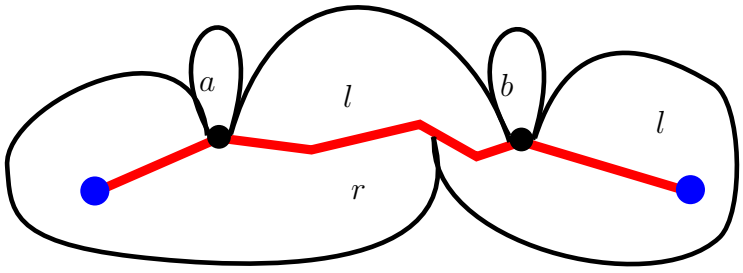

Figure 3.28: This depiction is not possible since the l region of the boundary would have to be disconnected to allow it.

□

3.5.1 Fully Extended C-shaped Geodesics

In this section, we examine the star unfoldings of fully extended C-shaped geodesic curves. As depicted in Figure 3.27, there may be a path in the ridge tree of the unfolding of a C-shaped geodesic that has more than π source angle to one side (in fact, all 2π surface angle could lie to one side of that path). However, we will show that the unfolding of a C-shaped geodesic has a very specific structure. Then, in the following two sections, we will use that structure to show the unfolding does not overlap in two special cases: (1) when each loop of the geodesic has only one vertex inside it (Section 3.5.2); and (2) when each loop of the geodesic has curvature $\leq 3\pi/2$ (Section 3.5.3).

First, some notation. We will call the two loops of the geodesic the *a-loop* and *b-loop* respectively. We define g_a (resp. g_b) to be the length of the geodesic loop from $\lambda(a)$ to a (resp. $\lambda(b)$ to b). Also, without loss of generality, we will assume that the r side of the curve is to the inside of both loops. We will show that the unfolding of a C-shaped geodesic is composed of 3 parts, which we will call the *head* and two *ears*. More precisely, the *a-ear* is the unfolding of the inside of the a-loop, the *b-ear* is the unfolding of the inside of the b-loop, and the *head* is the remainder of the unfolding. Figure 3.29 gives a general schematic of this form, see Figure 3.27 for a concrete example.

A useful fact about each of these pieces is that they are roughly convex; specifically, if we fill in the curvature gaps of all vertices with curvature $\leq \pi$ in the unfolding in the manner of Lemma 3.2, then each piece forms a convex polygon.

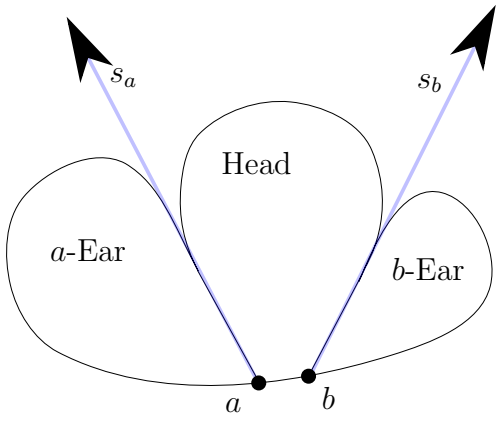

Figure 3.29: Schematic of the shape of the geodesic star unfolding of a C-shaped geodesic, with base rays indicated.

Lemma 3.9. *The ears and head of a fully-extended C-shaped Geodesic unfolding are convex, assuming the curvature gaps with curvature $\leq \pi$ lying between consecutive segment images are filled in.*

Proof. The segment images around each ear are all of a single type (either r or l), and therefore the boundary of each ear is contained within a convex chain by Lemma 3.2. Because the sum curvature to the interior of the loops is less than 2π (see the proof of Claim 3.1), the chain must meet itself at a convex angle at the loop point $a = \lambda(a)$ (resp. $b = \lambda(b)$), thus the entire contour around the ear must be convex. For the head, notice that it is composed of one chain from r and one chain from l (these are the upper and lower chains respectively in Figure 3.27). These two chains will meet each other at the two loop points. The angle at which they meet on either side must be less than π, and therefore, form a convex angle (refer to the proof of Claim 3.1 and Figure 3.22 for an illustration). □

Define the *base ray* for endpoint a, denoted s_a of a C-shaped geodesic star unfolding to be the ray with an origin point at $a = \lambda(a)$ (resp. b), and directed along the segment of λ near a for which the two sides of λ come together in the unfolding (see Claim 3.1). Define s_b analogously for endpoint b (see Figure 3.29 for a schematic of the base rays).

Clearly, the head and a-ear lie strictly on opposite sides of s_a, and likewise the b-ear is on the opposite side of s_b. Thus, it is trivial to show that the head does not overlap with either ear. However, it is not so clear that the two ears do not overlap with each other.

Lemma 3.10. *If the two base rays s_a and s_b do not intersect, then the unfolding does not overlap.*

Proof. We already showed that the head does not overlap with the ears, so all that remains is to show the ears do not overlap with each other. The region in which the a-ear (resp. b-ear) may exist lies strictly to one side of s_a (resp. s_b). Consider the situation depicted in Figure 3.30. Since the base rays do not overlap, and we know that the contours of both ears belong to the same convex chain, then there is no way that they can possibly overlap. □

If the rays do intersect, as show in Figure 3.31, there is a 'zone' above the unfolding, above the point of intersection between the two rays (call it x), in which the two ears could potentially overlap if they extended far enough. Figures 3.32 and 3.33 show how to construct a concrete example with intersecting base rays.

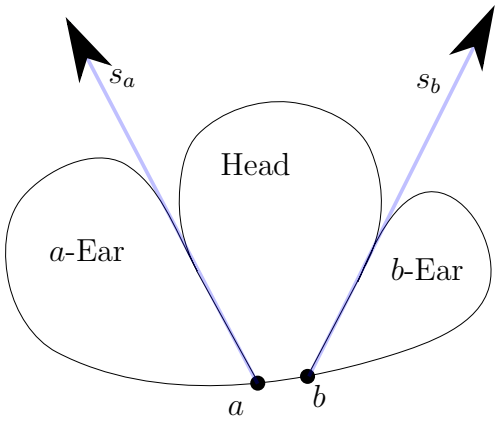

Figure 3.30: If the base rays do not overlap, then the unfolding cannot overlap.

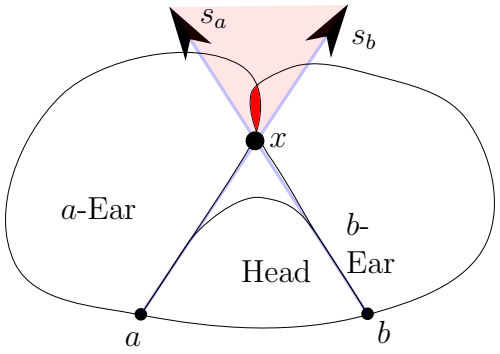

Figure 3.31: If the base rays do overlap, there exists a 'zone' of potential overlap above the unfolding (shaded red).

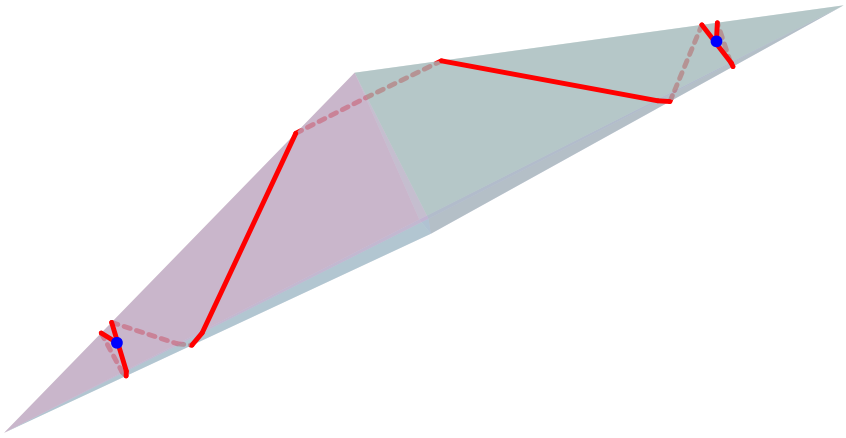

Figure 3.32: C-Shaped geodesic where the interiors of both loops have very high curvature.

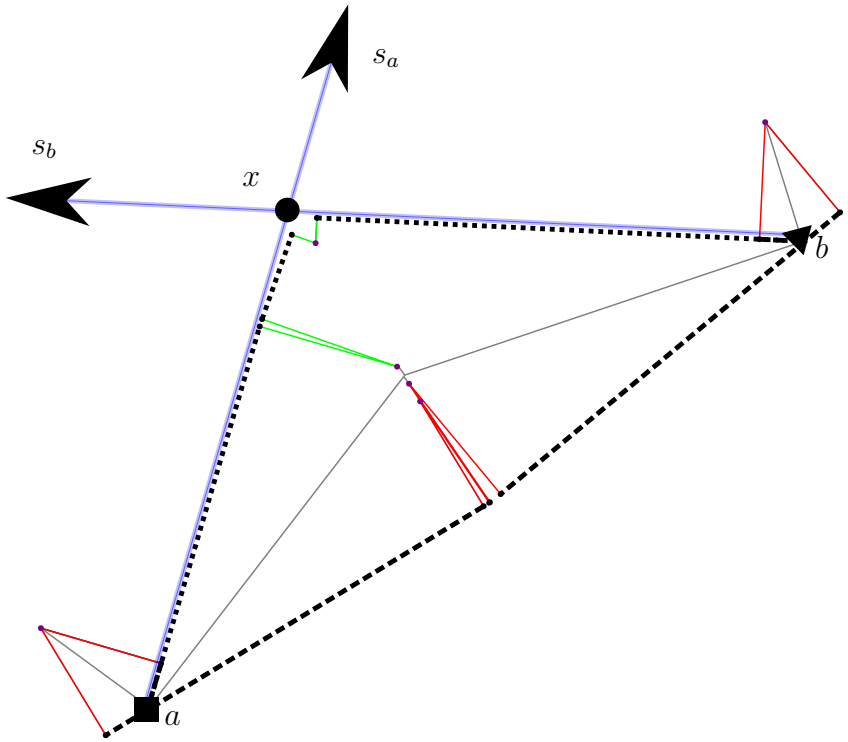

Figure 3.33: The base rays of this unfolding intersect.

We were unable, however, to construct an example where actual overlap between the

two ears occurs. As we can observe in our examples, the ears have a tendency to be 'small' relative to the middle section. We attempt to formalize this observation to create proofs of non-overlap in the following sections.

3.5.2 C-Shaped Geodesics With One Vertex In Each Loop

In this section we prove non-overlap for the simple case where only a single vertex is contained inside each loop.

Our technique, both in this section and the next, is to show that each ear is confined to a certain right-angle wedge. We will define notation for the a-ear (the b-ear is analogous). We will show that the a-ear is contained in a wedge formed by its base ray s_a together with a ray t_a at a right angle to s_a, intersecting s_a at point p_a (see Figure 3.34).

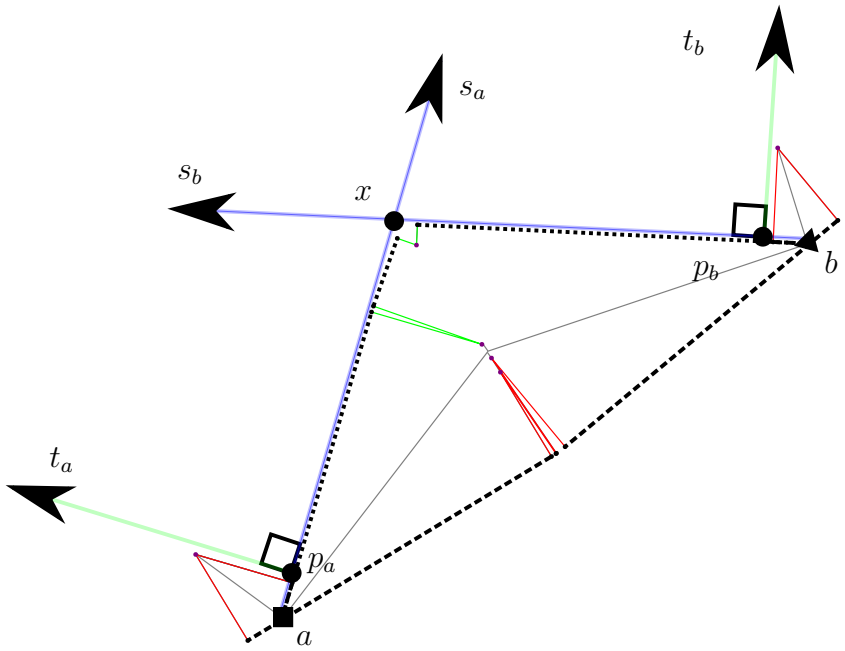

Figure 3.34: Even when the base rays intersect, we can sometimes still show the ears do not overlap by bounding them to a finite region.

The heart of the argument will be that if we can give a good upper bound on $|ap_a|+|bp_b|$ then we will be able to guarantee that the two ears do not overlap. As a first step in this

strategy, we will prove a lower bound on $|ax| + |bx|$ (recall, x is the point of intersection between s_a and s_b).

Lemma 3.11. $|ax| + |bx| \geq g_a + g_b$.

Proof. Clearly, $|\lambda| \geq g_a + g_b$. Let C be the convex chain of the l segments on the head between a and b, and denote its length $|C|$. Consider the triangle $\triangle(axb)$, clearly C lives inside this triangle. Futhermore, since the chain is convex, then by the triangle inequality, $|ax|+|bx| \geq C$. Finally, the length of the chain is clearly less than the sum of all the segment sources along it, which in turn is simply $|\lambda|$. Therefore, $|ax|+|bx| \geq |C| \geq |\lambda| \geq g_a+g_b$. □

Lemma 3.12. *If $|ap_a| + |bp_b| \leq g_a + g_b$ then the unfolding does not overlap.*

Proof. Suppose the two ears overlap at some point z. Take the projection of z onto s_a; call it z_a. Do the same onto s_b; call it z_b (see Figure 3.35). Clearly, if the two ears overlap at point z, then $|az_a| \leq |ap_a|$ and $|bz_b| \leq |bp_b|$. But, for every possible case, we will show that in fact $|az_a| + |bz_b| > |ax| + |bx|$, which will lead to a contradiction.

Consider the following cases:

(1) $|az_a| > |ax|$ **and** $|bz_b| > |bx|$. Then trivially $|az_a| + |bz_b| > |ax| + |bx|$.

(2) $|az_a| > |ax|$ **and** $|bz_b| \leq |bx|$. Take the point of intersection of s_b with zz_a, call it y (illustrated in Figure 3.35). Clearly $|z_ax| > |yx| > |z_bx|$. But then that implies $|z_ax|-|ax| > |bx|-|z_bx|$. Re-arranging the inequality, we get $|z_ax|+|z_bx| > |ax|+|bx|$

(3) $|az_a| \leq |ax|$ **and** $|bz_b| > |bx|$. Symmetric to the previous case.

(4) $|az_a| \leq |ax|$ **and** $|bz_b| \leq |bx|$. This is not possible since z lies in the convex wedge formed by the portions of s_a and s_b beyond x (and that is the only region where the two ears may intersect).

So, $|az_a| + |bz_b| > |ax| + |bx|$. But, by our assumption, $|ap_a| + |bp_b| \leq g_a + g_b$, and by Lemma 3.11, this is in turn $\leq |ax| + |bx|$, a contradiction. □

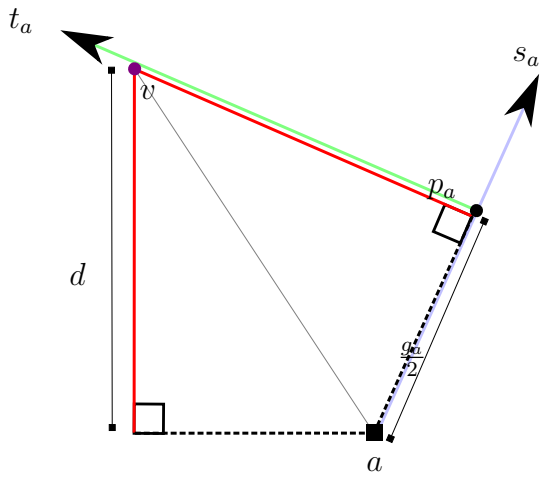

Figure 3.36: With a single vertex inside the a-loop, the wedge boundary ray t_a lies along the unfolded kite, and the distance from a to p_a is exactly $\frac{g_a}{2}$ from the loop point.

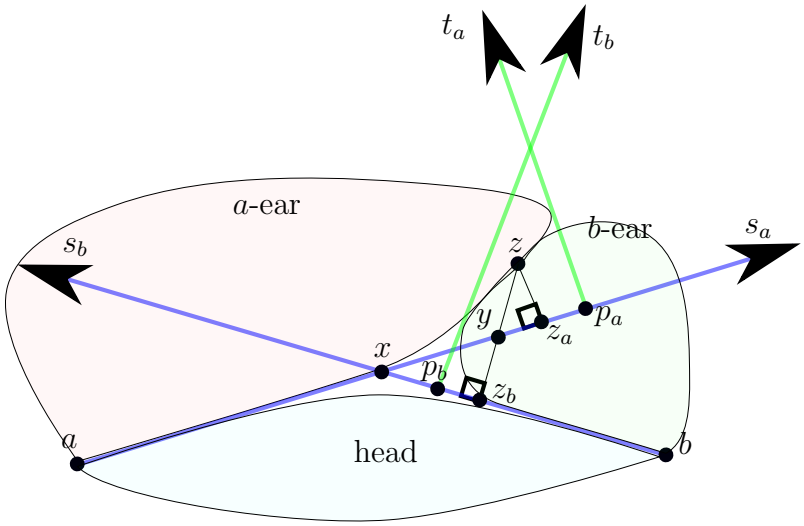

Figure 3.35: Proof of Lemma 3.12. Projecting the point of intersection z onto the two base rays. It is always true that $|az_a| \leq |ap_a|$ and $|bz_b| \leq |bp_b|$.

We now address the case where each loop has only one vertex.

Lemma 3.13. *If each loop of the C-shaped geodesic has just one vertex then the unfolding does not overlap.*

Proof. We will argue first for the a-ear. When there is a single vertex reporting to inside the loop, call it v, there is also exactly one ridge tree edge inside the loop, passing from a to v. Because there are only segment images, in the unfolding, the ridge tree edge is a straight line, and the unfolded region for the ear is a kite (see Figure 3.36). The ear has two segment images to either side, and since they are symmetric, they are of length $\frac{g_a}{2}$ each. Since the vertex is at a right angle from the segment images, p_a can be at exactly the endpoint of this segment image. Repeat this argument for the b side, to obtain that $|ap_a| = \frac{g_a}{2}$ and $|bp_b| = \frac{g_b}{2}$. Applying Lemma 3.12 completes the proof. □

3.5.3 C-shaped Geodesics with Low Curvature Loops

In this section we prove that the unfolding is non-overlapping if each loop of the C-shaped geodesic has curvature $\leq 3\pi/2$. We will use the techniques developed in the previous section to confine each ear to a bounded right-angle wedge. In particular, we will make use of Lemma 3.12 by bounding the distance $|ap_a| + |bp_b|$.

Our bound will be in terms of d, the maximum distance from the geodesic loop to any vertex (and furthermore, any point) on the inside of the loop. Without loss of generality, we will discuss each loop in isolation, calling its loop point o and denoting its length g_o. The following lemma makes use of an observation by O'Rourke and Vîlcu [24] regarding *conical development* of certain curves on convex polyhedra. Their proof shows that any convex curve 'lives on a cone', and continues to live on that same cone after iterations of Alexandrov's gluing theorem are applied to the vertices to one side of that curve. We show that the distance of the curve to the apex of this cone is the maximum distance to any other vertex inside the loop.

Lemma 3.14. *Let $\mathcal{P}$ be a convex polyhedron, and λ a geodesic loop with loop point o and length g_o on the surface of P, such that the curvature inside the loop is $\kappa \leq 2\pi$. Then the maximum distance from any vertex on the loop's interior to λ, call it d, is $\tan(\frac{\kappa}{2} - \frac{\pi}{2}) \cdot \frac{g_o}{2}$.*

Proof. Use Alexandrov's gluing theorem as described in [24] to merge every vertex to the interior side of the geodesic loop λ. Recall since the curvature inside the loop is $\leq 2\pi$, it is possible to merge all vertices inside the loop into one. By Lemma 1 of [24], there exists a uniquely determined cone on which λ and its interior side lives, and this cone remains the same after each merge. Consider unfolding this cone by cutting along its ridge tree, which in this case will be a geodesic path from its apex to the the loop point of λ. (Note that this is effectively the source unfolding from a geodesic curve of this 1-vertex subsurface, since we have cut it along the ridge tree.)

The resulting unfolded shape will form an isosceles triangle whose base is length g_o and whose congruent sides are the length of the cut we made from the apex to o. Now, consider taking a similar unfolding of the surface interior to λ at any intermediate state of the vertex merging process (see Figure 3.37 for an example). Since each merge strictly *adds* surface to the polyhedron, it must be the case that the each unfolding lives strictly inside the final isosceles triangle. The farthest possible point from λ anywhere on the surface is exactly the apex of the isosceles triangle. If the curvature inside the loop is κ, then the surface angle at the apex is $2\pi - \kappa$. Finally, we use trigonometry to measure the perpendicular distance from λ to the apex, giving our result that $d \leq \frac{g_o}{2} \cdot \tan(\frac{\kappa}{2} - \frac{\pi}{2})$ (see Figure 3.38). $\square$

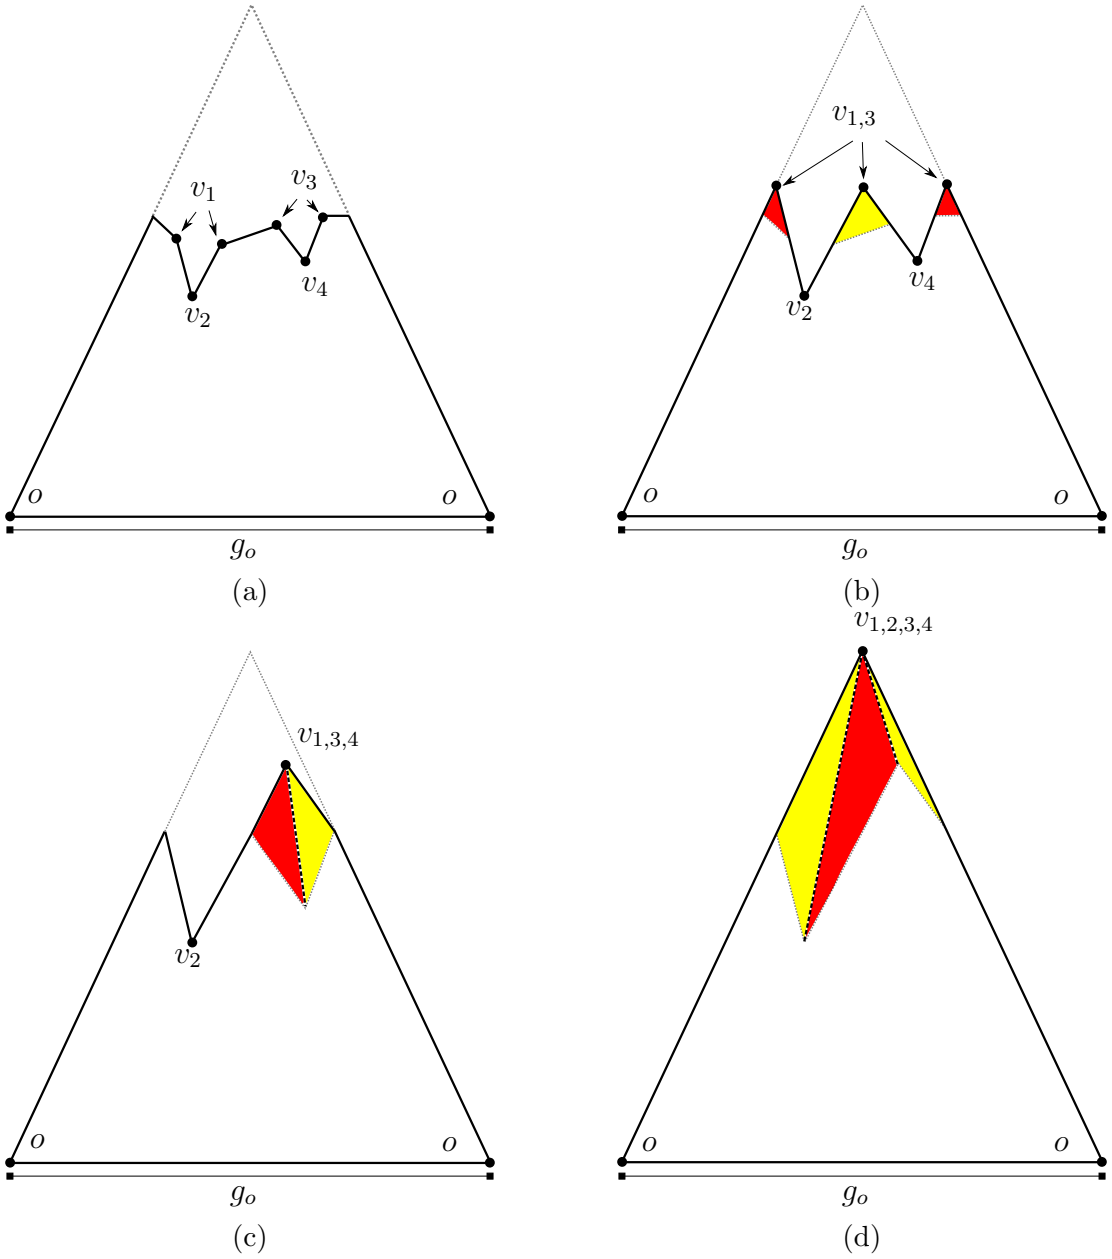

Figure 3.37: Example of merging 4 vertices on the interior of a geodesic loop by surface extension. (a) Depicts the initial configuration. (b) Merging v_1 and v_3 to form a new vertex $v_{1,3}$ with their combined curvature. The triangles that were 'sutured in' are highlighted (notice one triangle 'rolls over' to the other side). (c) and (d) Additional merges, to reach a completed cone/isosceles triangle.

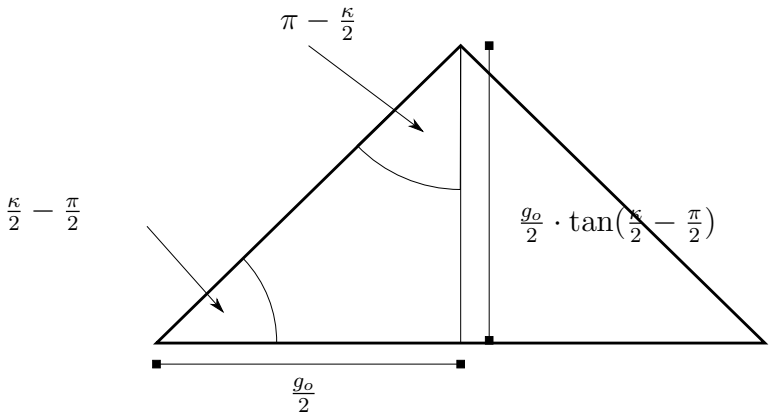

Figure 3.38: Calculating the maximum distance d from a lone vertex to the inside of the geodesic loop, based on the curvature inside the loop, κ, and the length of the loop, g_o.

Although it is most clear to express this distance d in terms of the 'farthest possible vertex', the real use for this fact is expressed by the following corollary:

Corollary 3.1. *The maximum distance from the geodesic to any point on the interior of a geodesic loop is $\leq d$, the maximum distance to any vertex inside the geodesic loop.*

Proof. Simply observe that at no point during the construction in Lemma 3.14 could the unfolded surface expand beyond the boundary of the isosceles triangle formed by the 1 vertex unfolding. Therefore, if we measure the perpendicular distance from the base of the unfolded image (which is the geodesic cut) to *any* point on the surface, the distance is at most d. □

Note that in particular, this implies that the maximum distance to any point on the *ridge tree* from the geodesic cut is $\leq d$ as well.

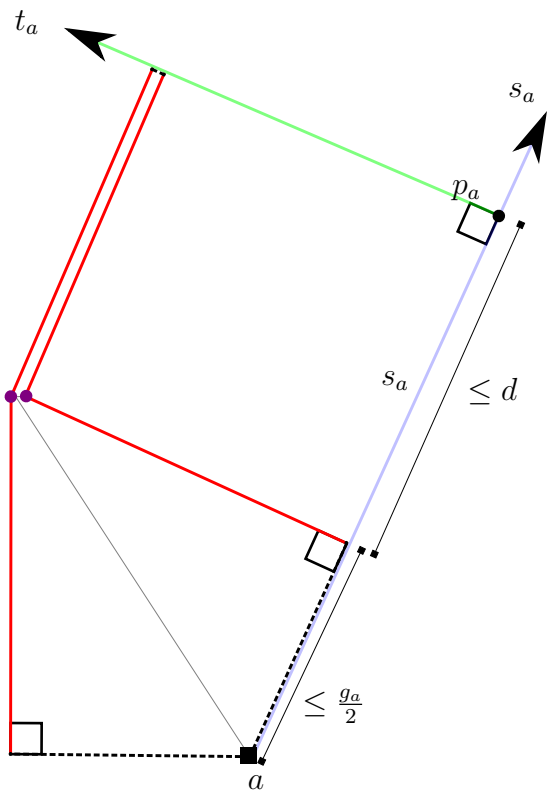

Figure 3.39: The total to distance p_a can increase by up to d with just one additional vertex.

We will compute an upper bound for $|ap_a|$ (resp. $|bp_b|$), and apply Lemma 3.11 to prove non-overlap. It is easy to see in the one-vertex case that the projection distance of each loop is at most $\frac{1}{2}$ of their respective loop lengths (see Figure 3.36), however, adding a second vertex can extend that distance significantly. In particular, that distance can increase by at most the maximum possible distance to any vertex inside the loop, as computed by Lemma 3.14 (see Figure 3.39). However, this is the strict limit to how far the projection can extend, as we formalize in the following lemma.

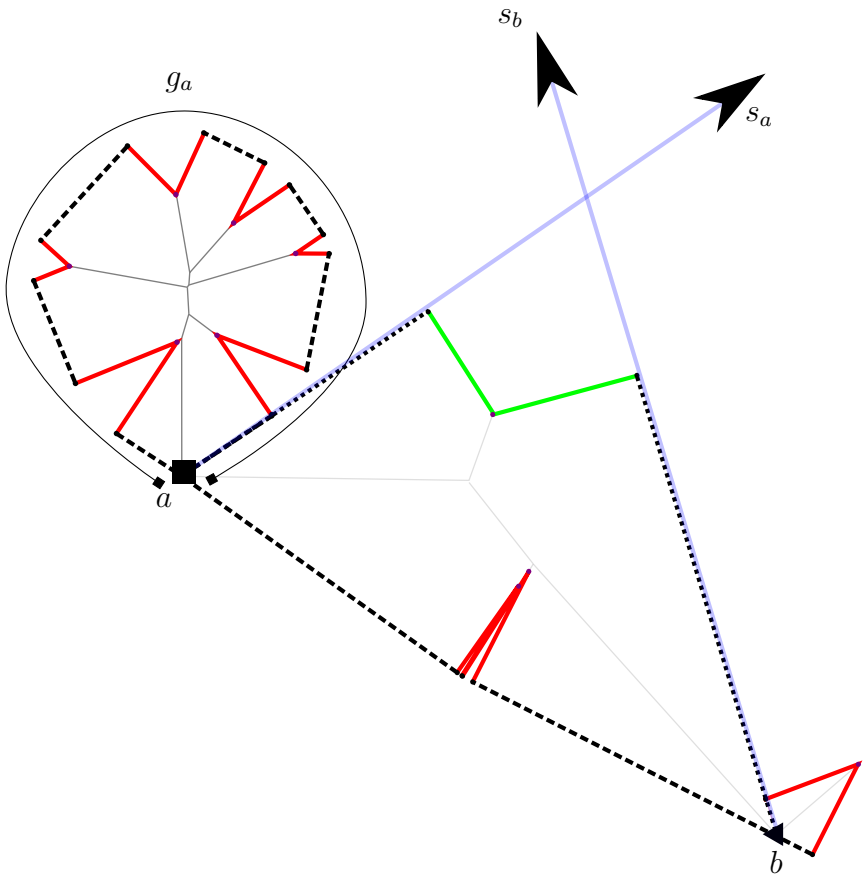

Figure 3.40: We will consider this single loop of a C-shaped geodesic unfolding in isolation.

Lemma 3.15. *Given a geodesic loop, with loop length g_a and a maximum possible vertex distance (as computed by Lemma 3.14) of d, then the projective distance $|p_a a|$ along the base ray s_a is at most $\frac{g_a}{2} + d$.*

Proof. Consider the situation depicted in Figure 3.41. We want to compute an upper bound for $|p_a a|$, in terms of g_a and d. First, observe that as we follow the contour of the unfolding in either direction, it stops moving any further along ray s_a when the sum of the curvature of the vertices reaches at least $\frac{\pi}{2}$; this occurs at the vertex labelled v' in Figure 3.42. Trace the unique path through the ridge tree from a to the ridge tree parent of v'. Each edge of this path is flanked by two congruent segment-wings, one to either side. As a result, the sum of the lengths of all segments to either side of any path is at most $\frac{g_a}{2}$. Thus the sum length of their projections onto s_a is $\leq \frac{g_a}{2}$ as well. The remaining distance is accounted

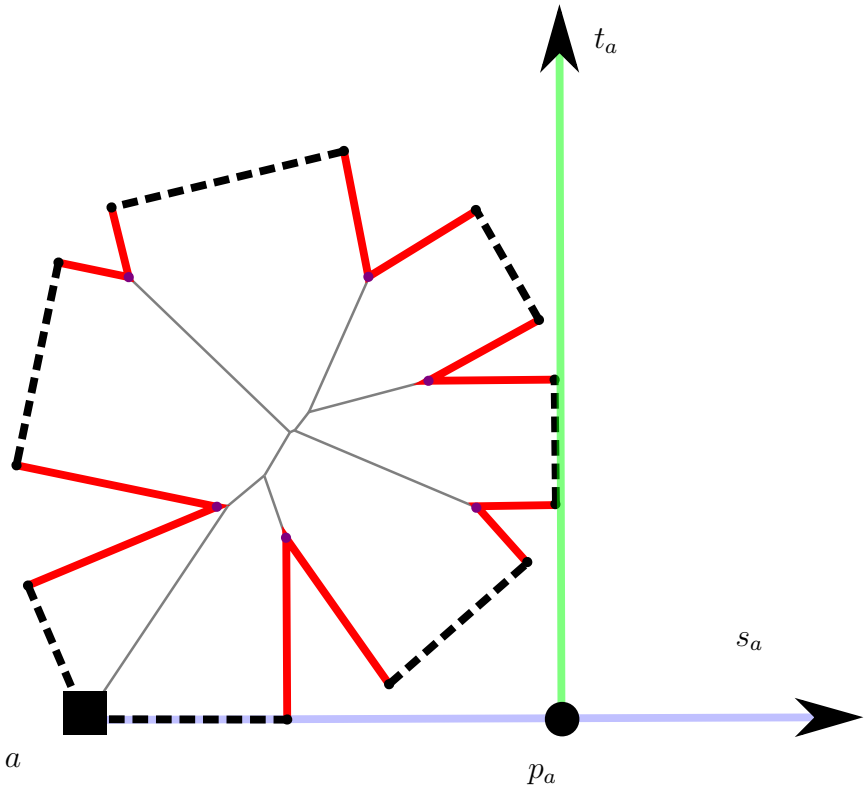

Figure 3.41: The loop of Figure 3.40 in isolation. Note that we wish to bound $|p_a a|$.

for by the curvature gaps between those projected images, as well as the final segment (see Figure 3.43). Consider cutting shortest paths from the segment images to each vertex along the path from a to v', and isolating those pieces. By Lemma 3.14, the distance of each such cut is $\leq d$. If we imagine those regions concatenated together (Figure 3.44), then the whole set can fit inside a circle of radius d. Assuming we truncate this to only consider the first quarter circle (which is the only pieces which matter when considering the projection onto s_a), then the total length of the projection is $\leq d$. Combining both shows that $|ap_a| \leq \frac{g_a}{2} + d$. □

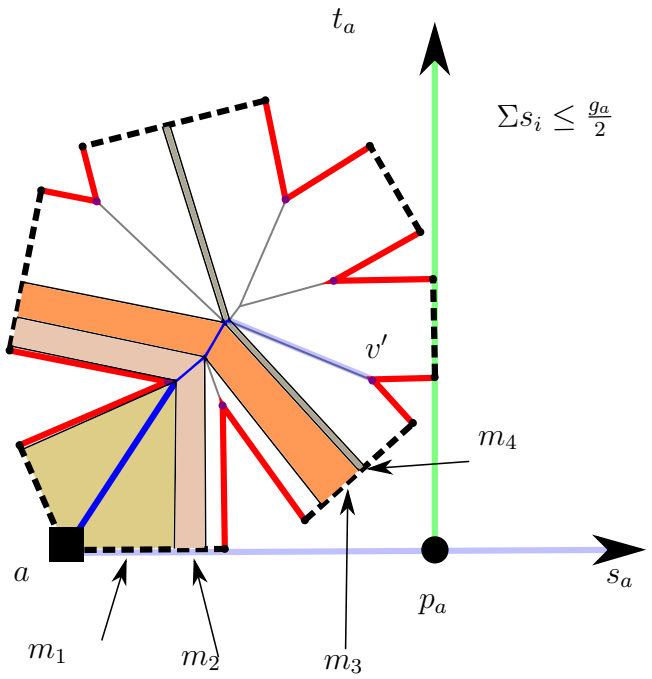

Figure 3.42: Each segment along the path is 'paired' with one on the opposite side; therefore their total length to one side is $\leq \frac{l}{2}$.

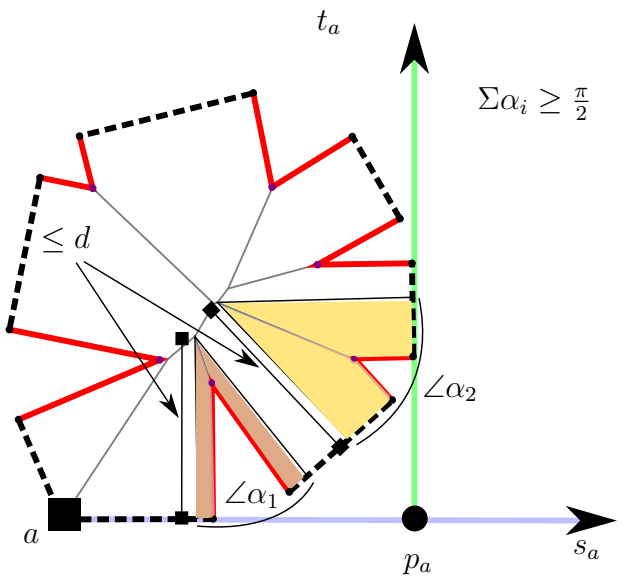

Figure 3.43: The remaining pieces form circular sectors, where the sum of the curvature gap angles are $\geq \frac{\pi}{2}$.

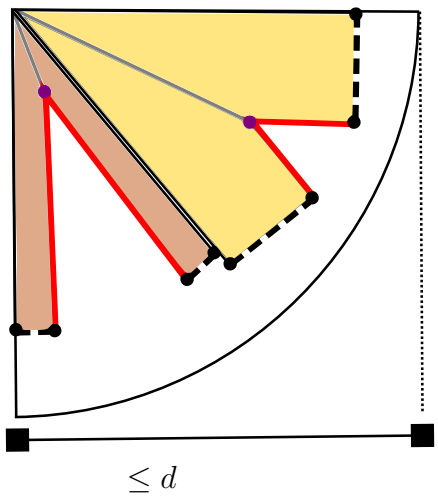

Figure 3.44: If we concatenate those slices together, and truncate anything rotated beyond $\frac{\pi}{2}$ of the base ray, they live inside a quarter-cicle with radius at most d.

When the curvature inside each loop is 'relatively low', Lemma 3.14 implies the maxi-

mum vertex distance is small as well. We formalize that idea in the following lemma.

Lemma 3.16. *If each loop of the C-shaped geodesic has curvature $\leq \frac{3\pi}{2}$ then the unfolding does not overlap.*

Proof. Suppose the curvature κ_a to the inside of the a-loop loop is $\leq \frac{3\pi}{2}$. Then $\frac{\kappa}{2} - \frac{\pi}{2} \leq \frac{\pi}{4}$, and therefore the maximum vertex distance $d_a \leq \tan(\frac{\kappa_a}{2} - \frac{\pi}{2}) \cdot \frac{g_a}{2} \leq \frac{g_a}{2}$. Therefore, $|ap_a| \leq \frac{g_a}{2} + d_a \leq g_a$. Apply the same argument on the b-loop to show that $|bp_b| \leq g_b$. Thus, $|ap_a| + |bp_b| \leq g_a + g_b$, and by Lemma 3.12, the unfolding does not overlap. □

Chapter 4

Conclusions

We have made steps towards a proof of non-overlap for the geodesic star unfolding; however a solution to the complete problem remains elusive.

4.1 Recapitulation

In Chapter 2, we introduced a new method for proving non-overlap of the star unfolding (Lemma 2.2). Our method usesd a decomposition the unfolding into *kites* formed by the two congruent triangles to either side of each ridge tree edge. We argued that for every path through the ridge tree, the kites along that path do not overlap, by establishing *W-wedges* which contain all previous kites along the path. Proceeding by induction on the length of the path, we showed that each subsequent kite further along the path is disjoint from the previous wedges, and therefore, cannot overlap with any other kite in the unfolding.

In Chapter 3, we showed that Lemma 3.4 can be applied with some limited success to the geodesic star unfolding, except that the partitioning of the unfolding is slightly more complex since the source images are both points and line segments. As long as the *surface angle* to either side of every path through the ridge tree is less than π, we proved the geodesic star unfolding does not overlap. In Sections 3.4 and 3.5, we showed two classes of geodesic curves for which this property, and therefore non-overlap of their unfoldings, is guaranteed, namely geodesic loops, and fully-extended S-shaped geodesics. In Section 3.5.1, we discussed a class of geodesic curves where this property does not hold at all, fully-extended C-shaped geodesics. Although no overlapping unfolding was found, we could not devise a way to prove non-overlap for this class of unfoldings. We gave a

concrete bound to the sizes of specific pieces of the unfolding contained inside geodesic loops of the geodesic cut, by bounding the maximum distance to any vertex. Using this bound, we were able to solve non-overlap for two specific sub-cases of this class of curves.

4.2 Open Questions and Future Directions

The most obvious question still outstanding from this work is whether the star unfolding from any open (or closed) geodesic is non-overlapping. We covered a few specific classes of curves for which non-overlap is guaranteed, but it seems unlikely further progress will be made without further insight. The following points were posed or investigated briefly during our research. Hopefully they might inspire improved results and new directions for research into this problem.

4.2.1 Extending the Geodesic from a Point

One direction for future study could be to look at taking an initial star unfolding from a point on the surface (that is, a geodesic with length zero), and investigating how the unfolding changes as we extend a geodesic from that point. The idea is that one endpoint would be *fixed*, remaining at a single location on the surface, while the other would be *moving*, that is extend in the direction of the geodesic. One could continue this transformation until the moving endpoint intersects with another part of the segment (i.e. becomes fully-extended in that direction). Figures 4.1 and 4.2 illustrate an example of this procedure. Figure 4.1 is the initial state of the point star unfolding, where the arrow $\Delta\lambda$ indicates the direction that the geodesic will cut into the surface from the source point. Figure 4.2 represents the change to the surface after extending the geodesic slightly. Pieces of the surface will rotate away as the moving endpoint becomes closer its two nearest vertices, forming segment images to either side.

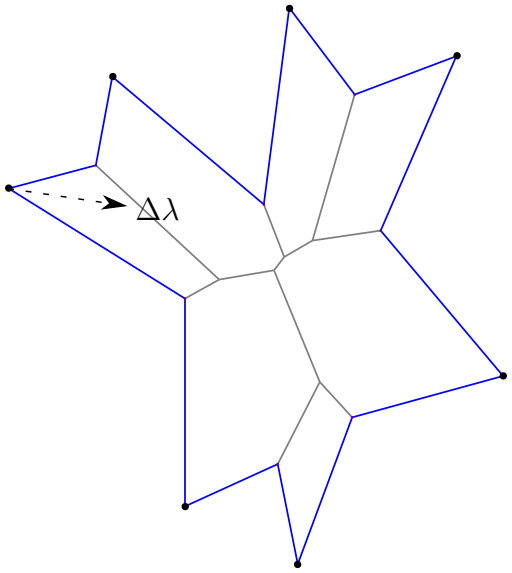

Figure 4.1: A 'zero length'-geodesic (i.e. point) star unfolding, with the direction of the geodesic cut indicated.

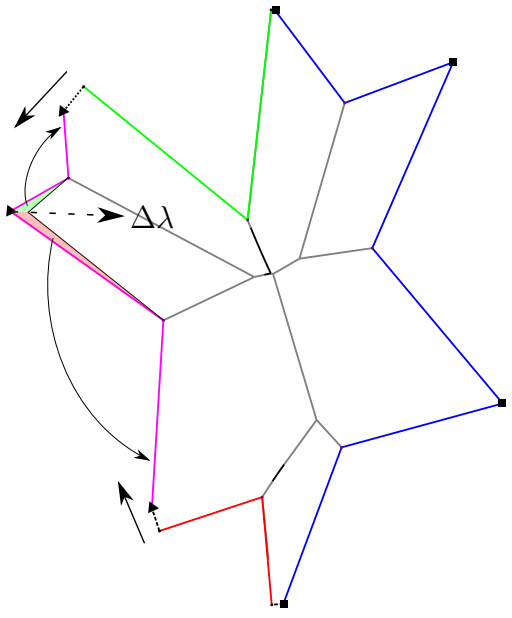

Figure 4.2: Extending the cut of Figure 4.1 slightly. No change to the rotational order of vertices around the unfolding has occurred.

One can show that, at least initially, the unfolding transforms continuously and without overlap as the curve extends. Unfortunately, this continuity ends when the rotational order of the vertex images around the unfolding changes. Modifications to the rotational order occur when the moving endpoint becomes closer to a vertex than any previous part of the geodesic curve (see Figures 4.3 and 4.4 for an example).

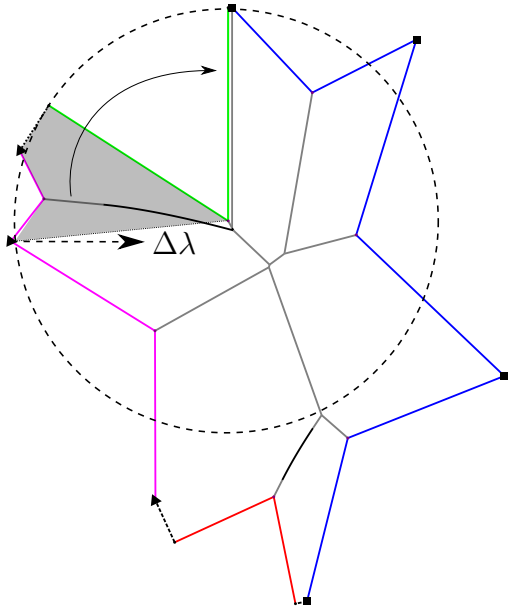

Figure 4.3: One endpoint of the geodesic segment is about to be closer to a vertex than its current reporting source segment image, meaning a rotation event is about to occur. The entire highlighted section will rotate in the direction indicated.

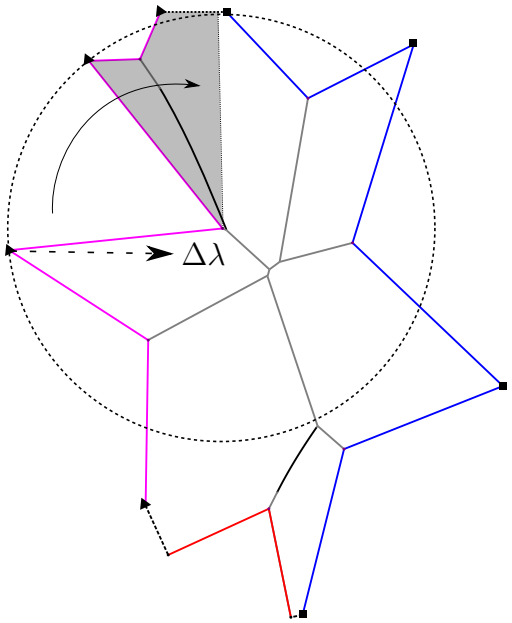

Figure 4.4: The state of the unfolding after the rotation event of Figure 4.3. The rotational order of the vertices around the unfolding has now changed.

How the unfolding changes at these *rotation event* is still well-defined; an entire section of the unfolding will rotate about the image of that vertex in the plane. Unfortunately, we were unable to show that non-overlap was preserved in the unfolding after a sequence of these events. In particular, the result could be a fully-extended C-shaped geodesic (for which we still have no proof of non-overlap). Perhaps if we could show that the ultimate result of any geodesic extension is non-overlapping, it might be simpler to argue that no intermediate state overlaps as well?

4.2.2 Generalization of Aronov and O'Rourke's technique

There may be some hope of applying a more general variant of Aronov and O'Rourke's proof of non-overlap of the star unfolding in [7] to the geodesic setting (this idea is hinted at in the proof of Lemma 3.14). The same vertex-merging technique of Alexandrov can be applied to a geodesic unfolding, with the additional restriction that the shortest path between the two vertices to be merged must not cross the geodesic cut, in addition to not crossing the ridge tree. The primary barrier to this approach would be establishing base cases for the induction. We have observed examples where the simplest possible polyhedron

that could be reduced to would be at best a tetrahedron (fully-extended C-shaped geodesics have this property), and other instances might require even more vertices. Furthermore, a stronger variant of Lemma 3.15 that bounds the planar size of non-geodesic loop sections of the unfolding would also be necessary.

Figures 4.6 and 4.6 illustrate this idea, and its potential interplay with the idea presented in Section 4.2.1. We conjecture that as long as the rotational order of the vertices remains the same, so do the relative sectors of emptiness.

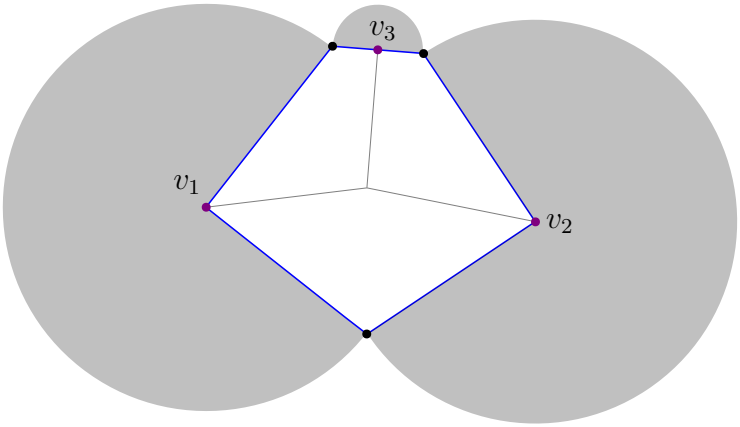

Figure 4.5: A point star unfolding with its sectors of emptiness highlighted.

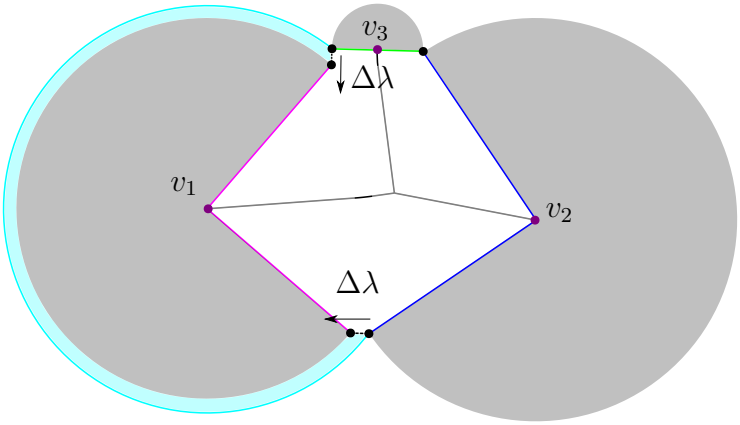

Figure 4.6: An unfolding similar to Figure 4.6 except the source point has been extended to a short geodesic. Note how the original sector of emptiness around v_1 remains empty (except for local intrusion by the shortest path cuts from v_1 itself) as the geodesic grows.

4.2.3 A More Generalized Star Unfolding

One interesting feature of Lemma 3.4 is that the nature of the proof does not seem to rely on the source being a point or a geodesic. Rather, it only requires some kind of ridge tree on the surface, and some notion of source angle on each edge. Therefore, it might be possible to use this technique to argue non-overlap of star unfoldings from an even wider class of 'source' objects, such as polylines (or closed polygons made up of piecewise geodesic segments). This would be a step closer towards the original question of non-overlap of a dual unfolding to the Sun Unfolding posed by [15].

Appendix A

Generating the Figures

Figures of all unfoldings and perspective views of polyhedra presented in this thesis were generated with a custom program written using CGAL [2], OpenGL [3], and Cairo [1]. The program implements an algorithm to compute shortest paths from any point to geodesic curve on a convex polyhedron. Using the information given by the shortest paths algorithm, the star unfolding from a geodesic can be computed. Since the ridge tree is simply the Voronoi diagram of the star unfolding layout, it can be simply be computed on top of the planar layout of the unfolding. The figures were drawn using the Cairo graphics library to achieve vector drawings. This program was also used to test several geodesic star unfoldings for potential overlap, but no counterexamples were found.

A.1 Computing shortest paths to a Geodesic

The shortest paths algorithm used was a variation of Chen and Han [11], modified to compute the geodesic offset. An example of such a modification can be found in [27] (though we did not implement any of the additional optimizations outlined in this paper). Chen and Han's algorithm builds what is called a *cone tree* from one (or more) source points, which are collections of geodesic (and therefore locally shortest) paths which all share a common edge sequence. Naive construction of this tree would result in an exponential size structure. Chen and Han give a simple method to prune nodes which cannot possibly define shortest paths, and make a structure that can be constructed in $O(n^2)$ time, and (theoretically) $O(n)$ space. To compute geodesic shortest paths, a new type of cone tree node is introduced, called a *Segment Source*. Segment source nodes function much the same as normal nodes; they live on a given face, and have a pair of left and right side

rays which delimit the cone inside which they can 'see' parts of the surface. The only difference is that the two rays no longer share the same source point, instead they are parallel rays originating from either end of a line segment. Shortest paths in each planar layout are then computed as the perpendicular distance to this segment, but otherwise the execution of the algorithm is unchanged. Each face that the geodesic curve crosses generates 2 segment source nodes, one for each side of the geodesic. The two endpoints of the geodesic curve simply act as two independent point source nodes, and shortest paths to them are computed using the standard Chen and Han algorithm. The theoretical running time of this extended algorithm is now a function of the number of faces crossed by the geodesic, call it k, with an upper bound of $O((n+k)n)$.

A.2 Computing the Unfolding

The only combinatorial information required to compute the unfolding from a geodesic segment is

- The set of shortest paths from each vertex to the geodesic curve (interior points along l/r, or endpoints a/b), and and their respective lengths.
- The rotational order of those shortest paths around the geodesic.

With this information, one can iteratively build the unfolding polygon by iterating through the rotational order of the shortest paths around the geodesic curve, and laying out each source and vertex image in order. We can then easily check if the resulting polygon is simple. Building the ridge tree is simply a matter of constructing the Voronoi diagram over all of the point and segment images once they are laid out in the plane. This is, however, assuming that the unfolding itself is non-overlapping. If there were self-overlap in the unfolding, one would have to divide the unfolding into non-overlapping sub-components, compute the Voronoi diagram over each component, and then merge along their boundaries. Computation of point-and-line-segment Voronoi diagrams was accomplished using the associated package in CGAL [22].

References

[1] Cairo 2d graphics library. http://cairographics.org/.

[2] CGAL, Computational Geometry Algorithms Library. http://www.cgal.org.

[3] OpenGL. http://www.opengl.org/.

[4] Pankaj K. Agarwal, Boris Aronov, and Catherine A. Schevon. Star unfolding of a polytope with applications. *SIAM J. Comput*, 26:1689–1713, 1997.

[5] A. D. Alexandrov. *Convex Polyhedra*. Springer, 2005.

[6] Greg Aloupis, Erik D. Demaine, Stefan Langerman, Pat Morin, Joseph O'Rourke, Ileana Streinu, and Godfried Toussaint. Edge-unfolding nested polyhedral bands. *Computational Geometry*, 39(1):30–42, 2008.

[7] Boris Aronov and Joseph O'Rourke. Nonoverlap of the star unfolding. *Discrete Comput. Geom.*, 8(3):219–250, 1992.

[8] Marshall Bern, Erik D. Demaine, David Eppstein, Eric Kuo, Andrea Mantler, and Jack Snoeyink. Ununfoldable polyhedra with convex faces. *Computational Geometry*, 24(2):51–62, 2003.

[9] Therese C. Biedl, Erik D. Demaine, Martin L. Demaine, Anna Lubiw, Mark H. Overmars, Joseph O'Rourke, Steve Robbins, and Sue Whitesides. Unfolding some classes of orthogonal polyhedra. In *Proceedings of the 10th Canadian Conference on Computational Geometry*, page 7071, August 1998.

[10] David Bommes and Leif Kobbelt. Accurate computation of geodesic distance fields for polygonal curves on triangle meshes. In *Proceedings of the Vision, Modeling, and Visualization Conference, Saarbrücken, Germany, November 7-9, 2007*, pages 151–160, 2007.

[11] Jindong Chen and Yijie Han. Shortest paths on a polyhedron, part 1: Computing shortest paths. *International Journal of Computational Geometry & Applications*, 06(02):127–144, 1996.

[12] Peter R. Cromwell. *Polyhedra*. Cambridge University Press, first edition, 1997.

[13] Erik D. Demaine, Martin L. Demaine, Anna Lubiw, Arlo Shallit, and Jonah Shallit. Zipper unfoldings of polyhedral complexes. In *Proceedings of the 22nd Annual Canadian Conference on Computational Geometry*, pages 219–222, August 2010.

[14] Erik D. Demaine, Martin L. Demaine, and Ryuhei Uehara. Zipper unfolding of domes and prismoids. In *Proceedings of the 25th Canadian Conference on Computational Geometry*, pages 43–48, August 2013.

[15] Erik D. Demaine and Anna Lubiw. A generalization of the source unfolding of convex polyhedra. In *Revised Papers from the 14th Spanish Meeting on Computational Geometry (EGC 2011)*, volume 7579 of *Lecture Notes in Computer Science*, pages 185–199, Alcalá de Henares, Spain, June 27–30 2012.

[16] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, New York, NY, USA, 2007.

[17] Réné Descartes. Progymnasmata de solidorum elementis, 1596-1650.

[18] Julie DiBiase. *Polytope Unfolding*. 1990. Undergraduate Thesis, Smith College, Northampton, Massachusetts, 1990.

[19] Albrecht Dürer. *The Painter's Manual: A Manual of Measurement of Lines, Areas, and Solids by Means of Compass and Ruler Assembled by Albrecht Dürer for the Use of All Lovers of Art with Appropriate Illustrations Arranged to be Printed in the Year MDXXV*. The literary remains of Albrecht Dürer. Abaris Books, 1977. English translation of "Unterweysung der Messung mit dem Zirkel un Richtscheyt in Linien Ebnen uhnd Gantzen Corporen", 1525, by Walter L. Strauss.

[20] Kouki Ieiri, Jin-ichi Itoh, and Costin Vîlcu. Quasigeodesics and farthest points on convex surfaces. *Advances in Geometry*, 11(4):571–584, 2011.

[21] Jin-ichi Itoh, Joseph O'Rourke, and Costin Vîlcu. Star unfolding convex polyhedra via quasigeodesic loops. *Discrete & Computational Geometry*, 44(1):35–54, 2010.

[22] Menelaos Karavelas. 2D segment Delaunay graphs. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.5 edition, 2014.

[23] Joseph O'Rourke. Unfolding face-neighborhood convex patches: Counterexamples and positive results. In *Proceedings of the 25th Canadian Conference on Computational Geometry*, pages 79–84, August 2013.

[24] Joseph O'Rourke and Costin Vîlcu. Development of curves on polyhedra via conical existence. *Computational Geometry*, 47(2, Part A):149–163, 2014.

[25] Micha Sharir and Amir Schorr. On shortest paths in polyhedral spaces. *SIAM J. Comput.*, 15(1):193–215, 1986.

[26] G. C. Shephard. Convex polytopes with convex nets. *Mathematical Proceedings of the Cambridge Philosophical Society*, 78:389–403, November 1975.

[27] Shi-Qing Xin, Xiang Ying, and Ying He. Efficiently computing geodesic offsets on triangle meshes by the extended Xin-Wang algorithm. *Comput. Aided Des.*, 43(11):1468–1476, November 2011.