



Department of Mathematics and Computing Science

On the Flip Distance of Pointed Pseudotriangulations

by
Remko René van der Vossen

Supervisor: Bettina Speckman

Eindhoven, August 2008

ABSTRACT

In the past 15 years a new planar subdivision has emerged, the pseudotriangulation. The pseudotriangulation is a planar subdivision reminiscent of the triangulation, but having various properties advantageous to certain applications. For instance pseudotriangulations are minimally rigid and can be constructed with a bound on both the vertex and the face degree.

One of the basic properties of both triangulations and pseudotriangulations is that one is able to transform one triangulation into another of the same point set by executing a number of edge flips. The minimum number of flips required to transform one triangulation into another is the flip distance of these two triangulations. The flip distance of two triangulations on the same point set was proved to be bounded by the number of intersections between the two triangulations. This thesis explores the possibility of a similar input-sensitive bound for the flip distance of pseudotriangulations.

The existence of a strictly decreasing distance measure as can be obtained for triangulations is examined. A number of different distance measures are proposed and evaluated. Furthermore we find that certain sub-cases have a flip distance of $\mathcal{O}(n)$ flips. Finally we introduce a new distance measure which measures the distance between two pseudotriangulations in a radically different way than was thought of up till now, making the measure more reminiscent of the distance measure used in the case of triangulations.

CONTENTS

1. Introduction	1
1.1 Pseudotriangulations	1
1.2 Edge Flips	4
1.3 Goal	6
1.4 Notation and Definitions	7
1.5 Bounding the Problem	8
1.6 Thesis Overview	11
2. Distance Measures	13
2.1 Introduction	13
2.2 Insertion and Overlay Measures	14
2.3 Removal Measure	16
2.4 Graph Measure	16
2.5 Extended Graph Measure	17
2.6 Selective Graph Measure	18
2.7 Weighted Graph Measure	19
2.8 Extended Graph Measure 2	19
2.9 Edge Measure	20
2.10 Half-Degree Measure	21
2.11 Conclusions	22
3. Software	23
3.1 Genptu Overview	23
3.2 Pseudotriangulation Generation	23

3.3	Implementation Details	24
3.4	Generated Problem Instances	25
4.	Pseudo-k-gons	27
4.1	Characteristic Graphs	28
4.2	Flipping Geodesics	32
4.3	Flip Distance	34
5.	Flips	35
5.1	Intersection Reducing Flips	35
5.2	Pairwise Perfect Flips	35
5.3	Sequence of Known Flips	36
6.	Conclusions	38
6.1	Thesis Objectives	38
6.2	Future Work	38
6.3	Summary	41
6.4	Acknowledgments	41
	Bibliography	43

1. INTRODUCTION

1.1 Pseudotriangulations

Pseudotriangulations, and pointed pseudotriangulations in particular, are of interest in various fields of computational geometry. Pseudotriangulations have been used in rigidity theory, collision detection, dynamic graph drawing and shape morphing, furthermore, pseudotriangulations have been used for proofs of the carpenter’s ruler problem and art gallery guarding with vertex π -guards.

Introduced in their present form in 1993, pseudotriangulations are a relatively young field of study and there are still many open problems related to pseudotriangulations. However, a closely related field of study, that of triangulations, has been a very well researched topic. The natural question arises whether there exist properties for pseudotriangulations similar to the many proved properties of triangulations.

A survey paper was recently written by Rote et al [Rote 08] detailing many of the results of past years and various applications of pseudotriangulations.

The Pseudotriangle




Fig. 1.1: Mutually tangent convex sets.

Pseudotriangles as used in this thesis were introduced by Pocchiola and Vegter in their contribution to the ninth Annual Symposium on Computational Geometry, “The visibility complex” [Pocchiola 93]. A pseudotriangle is the simply connected subset of the plane that lies between any three mutually tangent convex sets. See Figure 1.1. The boundary of a pseudotriangle consists of three mutually tangent convex curves. The endpoints of these curves form the three corner vertices of the pseudotriangle and have an exterior reflex angle.

Of particular interest are pseudotriangles which are simple polygons. In this case the convex curves

are convex chains of straight line edges between the corners. The vertices of these chains have interior reflex angles except at their endpoints. We call these interior reflex angles non-corners, in contrast to the three corners of the pseudotriangle which are interior non-reflex angles. Figure 1.2 shows some examples of various polygonal pseudotriangles. In the rest of this document we will consider only polygonal pseudotriangles.




Fig. 1.2: Example of various polygonal pseudotriangles.

Pseudotriangulations

A pseudotriangulation for a set S of n points in the plane is a planar partition of the convex hull of S into pseudotriangles whose vertex set is exactly S [Kettner 03]. A minimum pseudotriangulation has the least number of edges among all pseudotriangulations on the same point set, which is exactly $2n - 3$ edges as will be proved shortly. There exist several minimum pseudotriangulations on the same point set. Figure 1.3 shows an example of three different minimum pseudotriangulations on the same point set.




Fig. 1.3: Three different minimum pseudotriangulations on the same point set.

There always exists at least one (minimum) pseudotriangulation on every point set, we can give the proof by an incremental construction. First order the vertices of the pointset in the x direction and make a triangle of the leftmost three vertices. Then in order from left to right add vertices incrementally to the pseudotriangulation by adding the tangents from the new vertex to the existing pseudotriangulation. All faces such added are pseudotriangles, as two of the three sides are the two straight line tangents and the third side is a convex chain of the convex hull of the existing pseudotriangulation. The resulting minimum pseudotriangulation is called the incremental pseudotriangulation, an example of which is shown in Figure 1.3(c).

Pointedness

A vertex v is pointed if and only if there exists a line l through v such that all edges incident to v are strictly contained in a half-plane defined by l , see Figure 1.4. Equivalently, some consecutive pair of edges (in circular order around v) spans a reflex angle [Streinu 05]. A pointed pseudotriangulation is a pseudotriangulation whose vertices are all pointed.




Fig. 1.4: Example of a pointed and non-pointed vertex.

Theorem 1.1. *Let \mathcal{T} be a pseudotriangulation on point set S , then the following statements are equivalent;*

- \mathcal{T} is a minimum pseudotriangulation,
- \mathcal{T} is a pointed pseudotriangulation,
- \mathcal{T} has $2n - 3$ edges.

The proof is provided by Streinu in [Streinu 05]. Presented here is a slightly simplified version of this proof.

Proof. Let n, e and f denote the number of vertices, edges and faces of \mathcal{T} respectively. The total vertex degree $2e$ equals the number of corners plus the number of non-corners. Since \mathcal{T} is a pseudotriangulation all faces are pseudotriangles and thus there are exactly $3f$ corners. Since a non-corner spans a reflex angle, a non-corner is incident to a pointed vertex and each pointed vertex is incident to exactly one non-corner, thus the number of non-corners is the number of pointed vertices, or equivalently $n - b$, where b is the number of non-pointed vertices. This gives us the following equality:

$$2e = 3f + n - b$$

Applying Euler's formula $n - e + f = 1$ we obtain:

$$e = 2n - 3 + b$$

Clearly \mathcal{T} is minimum iff $e = 2n - 3$ iff $b = 0$, i.e. \mathcal{T} is a pointed pseudotriangulation.

□

Note that the number of edges in a pointed pseudotriangulation is invariant with respect to the number of vertices on the boundary of the convex hull, in contrast to triangulations for which the number of edges is $3n - k - 3$, where k is the number of vertices on the boundary of the convex hull.

1.2 Edge Flips

Two adjacent pseudotriangles T_1 and T_2 can be merged into a pseudoquadrilateral Q by removing an edge e that is shared by T_1 and T_2 . A unique edge $e' \neq e$ can be added to Q to form two different pseudotriangles T'_1 and T'_2 as illustrated in Figure 1.5. This process is called an edge flip. The diagonals e and e' are uniquely defined within Q by the geodesics $u \leftrightarrow w$ and $v \leftrightarrow x$ respectively.




Fig. 1.5: Example of an edge flip.

In contrast with triangulations, all inner edges of a pointed pseudotriangulation can be flipped to form a different pointed pseudotriangulation on the same point set. In a triangulation edges can only be flipped if their containing quadrilateral is convex, as illustrated in Figure 1.6.




Fig. 1.6: An edge flip in a triangulation and an example of an unflippable edge in a triangulation.

Note that there does not have to be a unique edge e between two incident pseudotriangles as illustrated in Figure 1.7. This means that two incident pseudotriangles can have more than one containing pseudoquadrilateral. There can however, be at most two edges shared by two pseudotriangles, in which case both can be flipped. Removing one of these two edges, \overline{ux} in our example, forms a degenerate pseudoquadrilateral in which two corners of the pseudoquadrilateral are actually the same vertex w . A dangling edge \overline{wx} which protrudes from w is the fourth “side” of the degenerate quadrilateral. One can consider this as a normal pseudoquadrilateral by splitting the dangling edge \overline{wx} , thereby splitting the vertex w shared by two of the corners. Naturally removing \overline{wx} instead of \overline{ux} is a symmetric case.




Fig. 1.7: Edge flips in a degenerate pseudoquadrilateral.

Flip Distance

Given two pointed pseudotriangulations \mathcal{T}_1 and \mathcal{T}_2 on the same point set, one can transform \mathcal{T}_1 into \mathcal{T}_2 and vice versa by a series of edge flips. Figure 1.8 presents two such pointed pseudotriangulations. The flip distance between two pointed pseudotriangulations \mathcal{T}_1 and \mathcal{T}_2 is defined as the minimum number of flips required to make this transformation between \mathcal{T}_1 and \mathcal{T}_2 . The flip distance problem asks for a bound on this flip distance.




Fig. 1.8: Example of two pointed pseudotriangulations on the same point set, transformed into one another through a sequence of edge flips.

The best currently known upper bound is $\mathcal{O}(n \log n)$ flips. The transformation is done by first transforming one pointed pseudotriangulation into a canonical pointed pseudotriangulation and subsequently transforming the canonical pointed pseudotriangulation into the other pointed pseudotriangulation. Both transformations take $\mathcal{O}(n \log n)$ flips as shown by Sergey Bereg [Bereg 04].

For triangulations the lower and upper bound is $\Theta(n^2)$ flips [Hurtado 96]. However, there also exists an input-sensitive upper bound. Hanke, Ottman, and Schuierer prove that the flip distance between

two triangulations \mathcal{T}_1 and \mathcal{T}_2 on the same point set is at most the number of intersections in the union of \mathcal{T}_1 and \mathcal{T}_2 [Hanke 96]. This bound is proved by showing that there always exists an edge flip which decreases the number of intersections in the union of the two triangulations. In other words, the number of intersections is a strictly decreasing distance measure between the two triangulations.

1.3 Goal

The goal of this thesis is to explore a similar input-sensitive bound for flip distances in pointed pseudotriangulations. The first thing that immediately strikes us however is that the union of two different pointed pseudotriangulations on the same point set does not have to include any intersections at all as illustrated in Figure 1.9(a). This means that for pseudotriangulations intersections alone cannot be enough to form the distance measure, there needs to be another concept from which we can construct the distance measure, possibly in tandem with the intersections. Although the union in Figure 1.9(a) does not have any intersections, there is another feature that is of interest, there is a non-pointed vertex in the union. In the general case we will see both intersections and non-pointed vertices as is the case with the union of the two pseudotriangulations of Figure 1.8 as presented in Figure 1.9(b). This seems to suggest that non-pointed vertices in the union should play role in defining the distance measure.




Fig. 1.9: Two examples of problem instances.

The problem considered in this thesis can be stated as such:

- does there exist an easily computable distance measure such that there always exists a flip path between any two pointed pseudotriangulations which strictly decreases this measure, or
- can the existence of such a measure be disproved

The next section shall introduce a special notation for the problem instances we encounter and define a

number of concepts related to the problem. This will make it easier to consider the problem instances and discuss the development of a bound.

1.4 Notation and Definitions

Throughout the rest of this document, whenever we refer to a pseudotriangulation, a pointed pseudotriangulation is implied unless explicitly noted otherwise. Also intersections between edges are meant as interior intersections, i.e. two edges meeting at a common vertex do not intersect. Furthermore intersections of geodesics are meant to be proper intersections, which means that there has to be a segment on the geodesic whose interior is entirely incident with the other geodesic and whose endpoints lie on opposite sides of the other geodesic. Incidence of the interior alone is not enough to constitute the intersection of two geodesics.

When considering a problem instance the two pseudotriangulations are usually presented overlaid in the same graph, an example of which, overlaying the two leftmost pseudotriangulations of Figure 1.3, is presented in Figure 1.10. The solid non-fat edges belong to the solid pseudotriangulation and the dashed non-fat edges belong to the dashed pseudotriangulation. The fat edges are shared by both pseudotriangulations. In other words, the fat and solid edges are one pseudotriangulation and the fat and dashed edges are the other pseudotriangulation.




Fig. 1.10: Example of two overlaid pseudotriangulations.

The reason for presenting the problem instances in this way is because it makes it easier to visually determine the distance between the two pseudotriangulations. The representation makes the non-pointed vertices in the union evident, while still allowing us to see the two different pseudotriangulations, which is important to be able to see the edge flips in the pseudotriangulations.

Perfect Flip

To simplify the reasoning about flip distances a special type of edge flip is introduced, the perfect flip. When transforming pseudotriangulation \mathcal{T}_1 into pseudotriangulation \mathcal{T}_2 , an edge flip $e \rightarrow e'$ in \mathcal{T}_1 is named a perfect flip if and only if $e' \in \mathcal{T}_2$. In other words, a perfect flip is a flip whose target edge is part of the target pseudotriangulation. We see some examples of perfect flips and non-perfect flips in Figure 1.11. Solid edge \overline{ae} flips to \overline{be} , which is a dashed edge, which means the flip is a perfect flip. On the other hand solid edge \overline{bd} flips to \overline{ac} , which is not a dashed edge, hence this flip is not a perfect flip.




Fig. 1.11: Example of a perfect flip and a non-perfect flip.

Perfect flips are a clear improvement over the previous situation as they unequivocally lessen the difference between two pseudotriangulations. Whenever a problem instance permits a flip path of all perfect flips the flip distance is linear in the number of differing edges. In many examples we have observed that many of the flips in a flip path for pseudotriangulations are perfect flips providing a close to linear flip distance.

Good and Bad Flips

We say a flip is a good flip if the flip reduces the distance measure under consideration and a bad flip otherwise. Obviously what exactly constitutes a good or a bad flip depends on the distance measure at hand.

This definition allows us to restate the first item of our problem statement as “does there exist a distance measure such that every problem instance permits at least one good flip.” Alternatively, if we can find a problem instance that does not permit any good flips for a certain distance measure, we can conclude the distance measure is not viable.

1.5 Bounding the Problem

Triangulations

As mentioned in Section 1.2, for triangulations there exists a known input-sensitive bound on the flip distance between two triangulations on a point set, which is the number of intersections in their union. An algorithm is presented by Hanke, Ottman and Schuierer [Hanke 96] to perform this sequence of flips. The algorithm repeatedly flips one of the edges in the first triangulation with the highest number of intersections with the other triangulation. Such a flip always brings the total number of intersections with the other triangulation down. When no more intersections remain the triangulations must be equal. The flips taken are the flip path from the first to the second triangulation.

This algorithm is a unidirectional iterative flip algorithm. In each step a single edge flip is made, only in the first triangulation of the pair. The number of intersections provide a strictly decreasing distance measure and thus a bound on the complexity of the algorithm. Also note that the algorithm will never

flip shared edges, as these edges never have any intersections with the other triangulation.

Intuition and problem symmetry suggest there is a similar approach to the flip distance of pseudotriangulations. Intersections alone however cannot be enough to define the distance measure, as shown earlier. It would seem reasonable, however, that the distance measure for pseudotriangulations would consist of the number of intersections in the union plus some other element. As mentioned earlier problem instances that do not show any intersections seem to indicate that non-pointedness in the union would be this element, as that is the only remarkable feature in these problem instances.

The question is what the contribution of these non-pointed vertices to the distance measure would be. There exist simple examples that show that the number of non-pointed vertices alone is simply not enough. One such example is shown in Figure 1.12, there are three intersections and three non-pointed vertices, giving a distance measure of six. There are however six differing edges, six edges in each pseudotriangulation not in the other, meaning that at least six perfect flips will have to be made in the flip path, since a perfect flip is the only way to create a shared edge. However since there are no perfect flips to be made from the initial state there has to be at least one non-perfect flip in the flip path, giving a flip path of at least seven flips in length. Thus a distance measure defined by the number of intersections and the number of non-pointed vertices in the union is not a valid distance measure.




Fig. 1.12: Example of two pseudotriangulations needing more flips than the number of intersections and non-pointed vertices.

Staged Algorithm

Initially one might be tempted to work with a staged algorithm, where first all intersections are resolved and then all pointedness violations or vice versa. These methods do not seem to work however. First removing all intersections does not work because there is an example of a problem instance where there exist only pointedness violations, but where any flip will introduce an intersection.

Figure 1.13 presents an instance where there exist no intersections. Here any flip will incur an intersection, which makes it also impossible to first remove all intersections and then all pointedness violations.




Fig. 1.13: Example of a problem instance where any flip introduces an intersection from a state without intersections.

Note that this example is symmetrical, which means that there is no progress from either direction. Hence, this reasoning does not only preclude a unidirectional staged algorithm, but also a bidirectional staged algorithm, in which the two pseudotriangulations would be 'flipped towards each other' instead of transforming one pseudotriangulation into the other.

In the case of the reverse staging, i.e. first making all non-pointed vertices pointed and subsequently removing all intersections we would need to show the absence of a similar case as in Figure 1.13 in which any flip would introduce a non-pointed vertex. We have however not been able to prove that such an example can not be constructed. The problem is further compounded by the fact that by making the non-pointed vertices in the problem instance pointed we are likely to introduce many new intersections. This means that we need a number of flips related to the degree of the non-pointed vertices to make those vertices non-pointed and subsequently need to resolve a number of intersections possibly far greater than the number of intersections in the original problem instance. The obvious question here is how this distance should be measured on the original problem instance.

Another example against a staged algorithm is given by Figure 1.12, in which there simply is no flip which reduces the number of intersections. We can at this point not yet give an analogous example for pointedness violations as we have not yet defined how non-pointedness contributes to the distance measure.

These examples do not eliminate the possibility of an iterative staged algorithm in which one might alternately remove as many intersections as possible and remove as many pointedness violations as possible until no intersections and pointedness violations remain.

To prove the existence of such an algorithm, one would need to prove the absence of the possibility of a loop in which the algorithm would, from some state s , return to the same state s after a certain number of iterations. To prove such without a large number of case distinctions, it would be necessary to show that some distance measure is decreased by each flip made. In this case it is more convenient to do without the staged algorithm and search for a distance measure for which there exists a flip decreasing the distance measure in all possible problem instances. Chapter 2 presents an overview of a number of different distance measures which were considered.

Bidirectional Algorithm

The algorithm for triangulations is a unidirectional algorithm. It takes one of the two triangulations as a starting point and then progressively makes the triangulation closer to the second triangulation, the target triangulation, by making edge flips which decrease the distance measure. In pseudotriangulations this quickly presents a problem as presented in Figure 1.14.




Fig. 1.14: A problem instance that seems to indicate the need for a bidirectional algorithm.

All possible edge flips in the solid pseudotriangulation in Figure 1.14(a) seem to make matters worse. Each of the flips increases the number of intersections by at least two and removes at most one edge from a non-pointed vertex in the union, which stays non-pointed after the flip. The dashed pseudotriangulation presents a completely different picture however, here there are many perfect flips to perform, which always seems to be a good idea as it reduces the number of differing edges between the two pseudotriangulations.

This problem instance seems to indicate that the problem cannot be solved with a unidirectional algorithm or at least shows that a bidirectional approach is more sensible as one has more freedom in choosing edges to flip. A bidirectional algorithm will flip the two pseudotriangulations towards each other, i.e. will make an edge flip in either of the two pseudotriangulations which decreases the distance measure between them. When the two pseudotriangulations \mathcal{T}_1 and \mathcal{T}_2 are both flipped to some third pseudotriangulation \mathcal{T}_3 , possibly the same pseudotriangulation as either \mathcal{T}_1 or \mathcal{T}_2 , the flip path from \mathcal{T}_1 to \mathcal{T}_2 is the flip path from \mathcal{T}_1 to \mathcal{T}_3 concatenated by the reverse of the flip path from \mathcal{T}_2 to \mathcal{T}_3 .

1.6 Thesis Overview

This chapter already covered a lot of ground on the problem at hand, we have seen a few strategies that will not work and have seen why we need to look at a bidirectional algorithm working with a strictly decreasing distance measure. Chapter 2 presents a number of these measures as they were considered during the course of the the thesis work and cites a number of observations on each of these measures.

Chapter 3 describes the research software that was written to aid in the research into this problem. The results that came forth from the findings obtained with the help of the research software are described in Chapters 4 and 5. Chapter 4 deals with a specific sub-case of the flip distance problem, namely that of the pseudotriangulated pseudo- k -gon, while Chapter 5 looks at the problem by identifying various types of flips. Finally Chapter 6 provides an overview of all the findings of the thesis work and how the work might be expanded upon.

2. DISTANCE MEASURES

2.1 Introduction

As we have seen in Chapter 1, the initial observations on the problem and various problem instances leads us to believe that a suitable distance measure would need to depend on both the intersections and the non-pointedness in the union of the two pseudotriangulations composing the problem instance. The contribution of the intersections is quite straightforward, the question is however, what the correct contribution of non-pointedness should be to a suitable distance measure.

During the course of the research we investigated a number of different distance measures that could fit the problem, which are presented in this chapter. All of these distance measures always count all intersections and have various different ways to count the contribution that non-pointedness represents in the flip distance. The reason for always counting the number of intersections is because there are problem instances with only intersections and no pointedness violations, a trivial example of which would be the case of a problem instance on a point set in convex position.

Each section describes a specific distance measure, why this distance measure was conceived and various observations on the distance measure such as possible strengths and weaknesses of the distance measure and in a number of cases a counterexample which shows the distance measure not to be viable.

Definitions

First we provide a number of definitions to aid in the discussion of the various distance measures.

Let \mathcal{T}_1 and \mathcal{T}_2 be two pseudotriangulations on a point set S .

Let \mathcal{T}_{12} be $\mathcal{T}_1 \cup \mathcal{T}_2$.

Let $\widehat{\mathcal{T}}_{12}$ be $\mathcal{T}_{12} \setminus (\mathcal{T}_1 \cap \mathcal{T}_2)$.

Let $\mathcal{D}_m(\mathcal{T}_1, \mathcal{T}_2)$ be the flip distance between \mathcal{T}_1 and \mathcal{T}_2 given distance measure m .

Let $\mathcal{I}(G_1, G_2)$ be the set of intersections between graphs G_1 and G_2 .

Let $\mathcal{NP}(G_1, G_2)$ be the set of non-pointed vertices incident to graph G_1 in graph $G_1 \cup G_2$.

Let $\mathcal{NP}(e, G_2)$ be the set of non-pointed vertices incident to graph G_2 with edge e added to it.

Let $\text{deg}(S, G)$ be the total degree of all vertices in S in graph G .

Let $\text{edges}(S, G)$ be the set of edges incident to vertex set S in graph G .

2.2 Insertion and Overlay Measures

Two of the most “natural” distance measures are the Insertion and Overlay Measures, which count the number of intersections and pointedness violations introduced when inserting edges into the target pseudotriangulation or overlaying the origin and target pseudotriangulations.

Distance Measure 2.1 (Insertion Measure).

$$\mathcal{D}_{insert}(\mathcal{T}_1, \mathcal{T}_2) = \sum_{e \in \mathcal{T}_1} (|\mathcal{I}(e, \mathcal{T}_2)| + |\mathcal{NP}(e, \mathcal{T}_2)|)$$

The distance from \mathcal{T}_1 to \mathcal{T}_2 is the sum over all edges $e \in \mathcal{T}_1$, of the number of intersections and non-pointed vertices in $e \cup \mathcal{T}_2$.

Distance Measure 2.2 (Overlay Measure).

$$\mathcal{D}_{overlay}(\mathcal{T}_1, \mathcal{T}_2) = \sum_{e \in \mathcal{T}_1} (|\mathcal{I}(e, \mathcal{T}_2)| + |\mathcal{NP}(e, \mathcal{T}_2)|)$$

The distance from \mathcal{T}_1 to \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus for each non-pointed vertex in \mathcal{T}_{12} the number of incident edges from \mathcal{T}_1 .

Observations




Fig. 2.1: Example of two pseudotriangulations with asymmetric measures.

These measures were one of the first to be considered and were done so before the need of a bidirectional algorithm became apparent, hence one of the first observations on these measures was that they were asymmetric, i.e. $\mathcal{D}_{insert}(\mathcal{T}_1, \mathcal{T}_2) \neq \mathcal{D}_{insert}(\mathcal{T}_2, \mathcal{T}_1)$. Note that the contribution of the intersections to the distance measure is always symmetric, this is however not always the case with the pointedness violations.

The Insertion Measure is asymmetric as can be seen in Figure 2.1. No edges of the solid pseudotriangulation incur pointedness violations when inserted into the dashed pseudotriangulation, while six edges of the dashed pseudotriangulation incur pointedness violations when inserted into the solid pseudotriangulation, three of which at both ends.

As the Insertion Measure, the Overlay Measure is also asymmetric, the same example of Figure 2.1 holds. This time six edges of the solid pseudotriangulation are incident to one pointedness violation for a total of six. While three edges of the dashed pseudotriangulation are incident to one pointedness violation and an additional three edges are incident to two pointedness violations for a total of nine.

As these measures are not symmetric the question is how these measures should be applied to the problem. In Section 1.5 we already saw that a unidirectional algorithm would be very problematic and indeed when we apply these measures to the same problem instance, repeated in Figure 2.2, we see there is no distance decreasing flip for either measure in the solid pseudotriangulation in Figure 2.2(a).




Fig. 2.2: A problem instance that seems to indicate the need for a bidirectional algorithm.

Note that only solid edges ag , bh and ci contribute to the Insertion Measure by way of creating non-pointed vertices upon inserting them into the dashed pseudotriangulation and do so on only one end, this means that each edge contributes at most 1 to the Insertion Measure considering non-pointedness. However every edge flip in the solid pseudotriangulation will increase the number of intersections by at least 2, and thus will always increase the Insertion Measure. There is no good flip to be made.

Note that each solid edge is incident to exactly one non-pointed vertex and that every edge flip in the solid pseudotriangulation does not make any non-pointed vertices pointed. Since every edge flip in the solid pseudotriangulation increases the number of intersections by at least 2 and removes at most 1 edge from a non-pointed vertex, any flip will always increase the Overlay Measure. Again, there is no good flip to be made.

A possible solution to these problems would be to define a new measure by simply adding the distances from both directions to each other. This measure would however be counting intersections twice and it would be easier to define the measure on the union graph instead of defining the measure as a comparison between the two pseudotriangulations. This is done in the later distance measures such as the Extended Graph Measure.

2.3 Removal Measure

As a complement to the Insertion Measure we have the Removal Measure in which we give an edge more weight if removal of that edge would reduce the number of non-pointed vertices.

Distance Measure 2.3 (Removal Measure).

$$\mathcal{D}_{removal}(\mathcal{T}_1, \mathcal{T}_2) = \sum_{e \in \mathcal{T}_1} (|\mathcal{I}(e, \mathcal{T}_2)| + |\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2)| - |\mathcal{NP}(\mathcal{T}_1 \setminus e, \mathcal{T}_2)|)$$

The distance from \mathcal{T}_1 to \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus for each edge $e \in \mathcal{T}_1$ the number of non-pointed vertices in \mathcal{T}_{12} which become pointed in \mathcal{T}_{12} after removal of edge e .

Observations

In Figure 2.2(c) we can see that there are instances where there are simply no flips which decrease the distance measure. Neither in the solid nor in the dashed pseudotriangulation are there edges which will make a vertex non-pointed by removal of that edge and since there are no intersections to remove we are left to conclude there simply is no flip which can decrease the Removal Measure, hence this distance measure is not viable.

2.4 Graph Measure

Another possible distance measures is the Graph Measure. In contrast to the previous measures, this measure does not attach a measure to every edge, but instead gives a distance measure between two pseudotriangulations based on their union graph. The measure came from the original idea that non-pointed vertices in the problem instance should somehow contribute to the distance as explained in Section 1.5. The simplest way to make this contribution would be to count the number of non-pointed vertices in the union, which is what this distance measure does. We have already demonstrated in Section 1.5 that this measure does not provide a measure high enough to account for all necessary flips, hence this measure can not be used as a bound in the flip distance problem. It is only included here for completeness and to introduce the distance measure name which it originally received.

Distance Measure 2.4 (Graph Measure).

$$\mathcal{D}_{graph}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\mathcal{T}_1, \mathcal{T}_2)| + |\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2)|$$

The distance between \mathcal{T}_1 and \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus the number of non-pointed vertices in \mathcal{T}_{12} .

Observations

Figure 2.3 presents the same case as used in Section 1.5, in which there are more flips in the shortest flip path than the distance measure indicates. Each triangulation contains six edges not in the other




Fig. 2.3: Example of two triangulations needing more flips than the measure indicates and which does not have an intersection reducing flip.

triangulation, hence at least six flips are needed. However since there are no perfect flips to be made, irrespective of which flip we execute first, there will still be 6 differing edges between the two pseudotriangulations after the flips. Hence at least one more flip is needed, giving a minimum of seven flips. The measure however gives a distance of only six flips, three intersections and three non-pointed vertices, leaving us to conclude that this measure is not viable.

2.5 Extended Graph Measure

This is an extension of the Graph Measure that arose from the observation that simply counting the number of non-pointed vertices does not provide a distance measure high enough to accommodate the shortest flip path in all problem instances. In the Extended Graph Measure each non-pointed vertex contributes its degree to the measure, giving us a much higher distance measure than the Graph Measure.

Distance Measure 2.5 (Extended Graph Measure).

$$\mathcal{D}_{\text{extended graph}}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\mathcal{T}_1, \mathcal{T}_2)| + \text{deg}(\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2), \mathcal{T}_{12})$$

The distance between \mathcal{T}_1 and \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus the total degree of all non-pointed vertices in \mathcal{T}_{12} .

Observations

Whenever there are non-pointed vertices in the problem instance, the distance measure immediately shoots up, for non-pointed vertices each have degree at least three. Furthermore several interconnected non-pointed vertices exacerbate the problem, as the edges connecting the non-pointed vertices will add to the distance measure twice, once at each end.

While no counter examples have been found which show this measure not to be viable, we have also not been able to prove that this measure permits at least one good flip for every problem instance.

2.6 Selective Graph Measure

One of the methods examined to tighten the Extended Graph Measure was to only count non-intersected edges adjacent to the non-pointed vertices. The reasoning being that the intersected edges were already contributing towards the distance measure by virtue of their intersections. Another change from the Extended Graph Measure is that edges which are shared by both pseudotriangulations are not counted towards the degree of non-pointed vertices.

Distance Measure 2.6 (Selective Graph Measure).

Let G be \widehat{T}_{12} minus all intersected edges.

$$\mathcal{D}_{\text{selective graph}}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\mathcal{T}_1, \mathcal{T}_2)| + \text{deg}(\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2), G)$$

The distance between \mathcal{T}_1 and \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus the non-intersected degree of all non-pointed vertices in \widehat{T}_{12} .

Observations

This measure presents a problem, as removal of an edge from \mathcal{T}_{12} can actually increase the measure. Let $e \in \mathcal{T}_1, e \notin \mathcal{T}_2$ be an edge with both endpoints non-pointed in \mathcal{T}_{12} . If e is intersected by a single edge f , removal of edge f will reduce the measure by one by virtue of removing the intersection, but it will increase the measure by two due to the non-pointed endpoints of the now non-intersected edge e . If the flip $f \rightarrow f'$ does not induce any other decreases of the measure, the measure over the problem instance will have been increased by the flip.




Fig. 2.4: Example of a perfect flip which increases the Selective Graph Measure.

Among other things this allows a perfect flip to leave the measure invariant or even increase the measure. Figure 2.4 presents an example of a perfect flip increasing the measure. Flipping edge \overline{ae} to

\overline{bi} removes three intersections, but introduces two non-intersected edges, each having two non-pointed endpoints giving a net increase of one to the measure.

Although this does not completely discount the measure as viable, it will be much more difficult to prove a flip bound with it. It would also seem rather unlikely that a suitable distance measure would be increased by performing a perfect flip.

2.7 Weighted Graph Measure

This measure is a refinement of the Selective Graph Measure. By counting only one for each edge with pointedness violations on both ends, removal of an edge from one of the pseudotriangulations can not raise the distance measure anymore.

Distance Measure 2.7 (Weighted Graph Measure).

Let G be $\widehat{T_{12}}$ minus all intersected edges.

$$\mathcal{D}_{weighted\ graph}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\mathcal{T}_1, \mathcal{T}_2)| + |\text{edges}(\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2), G)|$$

The distance between \mathcal{T}_1 and \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus the number of non-shared, non-intersected edges incident to a non-pointed vertex in \mathcal{T}_{12} .

Observations

Going back to the example in Figure 2.3, we see that there is no flip which reduces the measure from either direction. First observe that no flip changes the pointedness of any of the vertices, this means that for a flip to be a good flip, the flip target of the flip needs to contribute less to the distance measure than the original edge. Edge \overline{ag} contributes 1 to the distance measure as it is a non-intersected edge incident to a non-pointed vertex. The flip target of \overline{ag} , edge \overline{df} , also contributes 1 to the distance measure by virtue of an intersection hence the flip is a bad flip. Edge \overline{dg} , contributing 1 to the distance measure, flips to \overline{ae} , contributing 3; also a bad flip. Edge \overline{ah} , contributing 1 to the distance measure, flips to \overline{df} , also contributing 1 to the distance measure; another bad flip. Finally edge \overline{gh} contributes 1 to the distance measure as it is a non-intersected edge incident to a non-pointed vertex. Its flip target, edge \overline{ai} , also contributes 1 to the distance measure by virtue of an intersection, a bad flip as well. The other edges are merely symmetric cases of the above. Hence this measure can not be used to prove an upper bound on the number of flips required, the distance measure is not viable.

2.8 Extended Graph Measure 2

Since both the Selective Graph Measure and the Weighted Graph Measure did not yield very positive results, this measure is only a slight refinement over the Extended Graph Measure, in that it does not count edges incident to a non-pointed vertex, which are contained in both pseudotriangulations.

Distance Measure 2.8 (Extended Graph Measure 2).

$$\mathcal{D}_{extended\ graph\ 2}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\mathcal{T}_1, \mathcal{T}_2)| + deg(\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2), \widehat{\mathcal{T}}_{12})$$

The distance between \mathcal{T}_1 and \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus the degree of each non-pointed vertex in \mathcal{T}_{12} only counting non-shared edges.

Observations

As with the Extended Graph Measure, we have not found any counter examples showing the Extended Graph Measure 2 not to be viable, but neither were we able to prove that there is always at least one good flip in every problem instance.

2.9 Edge Measure

The idea here is to take the Extended Graph Measure 2, but count only one for the edges with pointedness violations on both ends.

Distance Measure 2.9 (Edge Measure).

$$\mathcal{D}_{edge}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\mathcal{T}_1, \mathcal{T}_2)| + |\text{edges}(\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2), \widehat{\mathcal{T}}_{12})|$$

The distance between \mathcal{T}_1 and \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus the number of non-shared edges incident to a non-pointed vertex in \mathcal{T}_{12} .

Observations




Fig. 2.5: Example of two pseudotriangulations for which no edge flip reduces the Edge Measure.

This measure also presents a lack of good flips in certain situations. Figure 2.5 show such a situation, there exists no flip which reduces the Edge Measure. Edge \overline{dg} flips to \overline{eg} making vertex g pointed.

Edges \overline{gh} and \overline{gi} are still incident to another non-pointed vertex however and since the flip target \overline{eg} intersects \overline{hi} the distance measure remains invariant, making this a bad flip. Edge \overline{gh} flips to \overline{di} , which does not change the non-pointedness of any of the interior vertices, and since the flip target is still incident with one of these interior vertices the distance measure does not decrease by executing this flip; another bad flip. As the other possible flips are only symmetric cases of the above two, this problem instance does not permit any good flips to be made under the Edge Measure. This leaves us to conclude that the Edge Measure is not a viable measure.

2.10 Half-Degree Measure

This measure is another attempt to reduce the contribution of non-pointed vertices to the measure to a more realistic level. Given that every non-pointed vertex v can be made pointed by flipping away all edges contained in any given half-plane whose boundary contains v , it seems reasonable to count the least number of such edges towards the distance measure for every non-pointed vertex.

Distance Measure 2.10 (Half-Degree Measure).

Let the half-degree of a vertex v , in some graph $G = (V, E)$, be the least number of edges in E incident to v contained in any half-plane whose boundary contains v .

Let $halfdeg(S, G)$ be the total half-degree of all vertices in S in graph G .

$$\mathcal{D}_{halfdegree}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\mathcal{T}_1, \mathcal{T}_2)| + halfdeg(\mathcal{NP}(\mathcal{T}_1, \mathcal{T}_2), \mathcal{T}_{12})$$

The distance between \mathcal{T}_1 and \mathcal{T}_2 is the number of intersections in \mathcal{T}_{12} plus the total half-degree of all non-pointed vertices in \mathcal{T}_{12} .

Observations

There are a few problems with this measure, it does not take into account where the shared edges are located. Since we only wish to flip non-shared edges, any situation such as in Figure 2.6 presents a problem. The half-degree of the center vertex is very low, one in fact, but we will need to flip many more edges away from the vertex to make it pointed than the half-degree gives us.

This raises the question whether it is a wise decision not to flip shared edges, or whether the half-degree should be defined differently. Perhaps the half-degree could be defined not by the least, but by the most number of edges incident to the vertex in any half-plane whose boundary contains the vertex. Alternatively only the half-planes could be considered which do not contain any of the shared edges incident to the vertex.




Fig. 2.6: Example of a situation where the Half-Degree Measure can give a misleading distance measure.

2.11 Conclusions

We have seen quite a few different distance measures, some of which were proved not to be viable, some of which present some obvious difficulties and finally some which remain possibilities for a viable distance measure.

It has become clear in trying to prove or disprove the viability of some of these distance measures that we do not yet have a clear understanding of what exactly is going on in the interaction between two pseudotriangulations when flipping edges in them. Intuitively we feel that the Extended Graph Measure and Extended Graph Measure 2 are viable distance measures, however in trying to construct a counter example for these measure we have not been able to put a handle on the thing preventing us from constructing such an example. A better understanding of the larger picture might lead us to be able to either construct a counter example for these measures, or develop an insight on how to prove that the measure is viable. Once we have a viable input-sensitive distance measure we can then see if it is possible to define a tighter measure.

3. SOFTWARE

Many of the existing examples were relatively small and given the difficulty of constructing counterexamples that show a particular distance measure not to be viable, yet also the difficulty of proving the viability of a distance measure it might be possible that the existing examples are boundary cases of the flip distance problem. In order to gain a better overview of the common cases encountered in larger problem instances and the possibility of examining different examples of problem instances given a certain set of domain restrictions a program called `genptu` generating pseudorandom problem instances was written.

3.1 Genptu Overview

`Genptu` constructs a pseudo random problem instance, allowing restrictions to be laid on the problem instance, such as minimizing the number of intersections between the two pseudotriangulations. The created problem instance is saved as an IPE XML picture file, allowing easy editing and printing of the problem instance with IPE [Cheong 06].

3.2 Pseudotriangulation Generation

First a pseudorandom point set S is generated by the program, making sure that no two points are too close to each other. This is accomplished by first generating a pseudorandom point set and subsequently checking whether the point set complies with the restrictions we desire, if not a new point set is generated and checked again. The restrictions imposed are simply that the distance between any two points in the point set are to be at least a given minimum distance apart. The reason for this restriction is to improve the readability of the generated problem instances. Of course this restriction can limit the possible topologically different point sets, however since the chosen minimum distance between points will normally be small compared to the area in which the point set is generated, this restriction will be very weak. If desired, the minimum distance can also be chosen to be 0, such that, effectively, there is no restriction on the generated point set. Another possible restriction to enhance readability would be to limit near colinearity in the point set, but this restriction would be much stronger than the minimum distance restriction.

The program then proceeds to generate a random minimal pseudotriangulation \mathcal{T}_1 on S by a trivial $\mathcal{O}(n^2)$ algorithm. The algorithm creates a random permutation of all edges in the complete graph on S and proceeds to insert these edges into the new pseudotriangulation as long as these edges do not violate the minimal pseudotriangulation property. In other words edges which do not intersect

the existing graph and for which both incident vertices remain pointed after inserting said edge are inserted.

A second pseudotriangulation \mathcal{T}_2 is then created, subject to a set number of constraints. The basic algorithm to do this is the same as for the first pseudotriangulation, with the exception of additional constraints on inserting the edges. Since these constraints may prohibit the generation of a complete minimal pseudotriangulation, the rejected edges are collected and reconsidered with a different set of constraints when appropriate. A chain of such constrained insert passes may be necessary to guarantee the construction of a complete minimal pseudotriangulation. The implemented constraint sets are described below, the software is set up such that additional constraint sets can be easily added.

Fully Random

The fully random generation does not impose any additional constraints for the generation of \mathcal{T}_2 . It generates two independently random minimal pseudotriangulations on the point set S .

Minimal Intersections

The minimal intersections constraint set tries to minimize the number of intersections between the two pseudotriangulations by greedily inserting edges into \mathcal{T}_2 that do not intersect \mathcal{T}_1 . Thus the first pass has the additional constraint of edges not contained in and not intersecting \mathcal{T}_1 , the second pass only the additional constraint of edges not intersecting \mathcal{T}_1 and a final pass with no additional constraints.

3.3 Implementation Details

The basis for the genptu program is an object oriented (C++) based model to represent graphs with simple vertex sets and edge lists. Each vertex has references to all incident edges and each edge has references to both incident vertices. A Mersenne Twister [Matsumoto 98] is used to generate the pseudorandom data used in generating the pseudorandom point set and pseudotriangulations.

The vertex class is based on a custom 2d vector class which allows comparison of vectors, vector arithmetic, etc. The vertex class also provides for a way to test whether the vertex is pointed or not.

The graph class contains the vertex set and edge list and provides functionality to check whether a graph contains an edge, whether it intersects an edge, to add the convex hull to the graph using the quickhull algorithm and to write the graph to an output stream in the IPE XML format [Cheong 06]. Custom 2d line and 2d line segment classes are available to do the pointedness, intersection and hull calculations.

The following two chapters deal with the research results obtained with the use of this software and further study of the problem.

3.4 Generated Problem Instances

Figures 3.1 and 3.2 show some of the generated problem instances in which the number of intersections was minimized. These examples show many perfect flips. We also see many shared pseudo-k-gons which gave rise to a new insight that allowed us to define a result for certain sub-cases. Pseudo-k-gons are introduced in Chapter 4 along with a discussion of this new result and how it could be applied to the problem as a whole. Finally we also see perfect flip pairs, which are pairs of edges from the two pseudotriangulations which flip to the same target edge, more on this in Chapter 5.





Fig. 3.1: The first problem instance generated with genptu.

pseed: 904826568 aseed: 3819619198 bseed: 3039372339



pseed: 3860965500 aseed: 2340043557 bseed: 1220010351




Fig. 3.2: Another two examples of problem instances generated with genptu.

4. PSEUDO-K-GONS

We can generalize the concept of a pseudotriangle and a pseudoquadrilateral to a pseudo-k-gon. A pseudo-k-gon is a simple polygon on n vertices, with k corners and $n - k$ non-corners. In other words, a simple polygon with k vertices with exterior reflex angle and $n - k$ vertices with interior reflex angles. The k chains of edges between consecutive corners on the polygon boundary can be considered k pseudosides to the pseudo-k-gon. Of course, as with pseudoquadrilaterals, a side of the pseudo-k-gon can be in a degenerate form as a dangling edge interior to the pseudo-k-gon. In this case, the vertex from which the edge dangles represents either two corners, or a corner and a non-corner at the same time. The degenerate case can be viewed as non-degenerate by splitting open the dangling edge at this vertex. We call a pseudo-k-gon with k corners a rank k pseudo-k-gon, e.g. a pseudotriangle is a rank 3 pseudo-k-gon and a pseudoquadrilateral is a rank 4 pseudo-k-gon.

In the larger examples generated by genptu we often see shared pseudo-k-gons appearing as subproblems of the larger problem instance, where each pseudotriangulation provides a different pseudotriangulation of the interior of the pseudo-k-gon. Figure 4.1 shows an example problem instance with several of these shared pseudo-k-gons.




Fig. 4.1: Example of two pseudotriangulations with shared pseudo-k-gons.

Pseudo-k-gon $\langle c, d, e, l, f, k, m \rangle$ is the non-degenerate case of a shared pseudo-k-gon with no dangling

edges and no interior vertices. Pseudo-k-gon $\langle a, b, n, o, n, c, m, k \rangle$ is a degenerate case with dangling edge \overline{no} . We can treat this as the non-degenerate case if we split open the dangling edge by splitting vertex n and moving both sides away from each other by an infinitesimal amount. This is allowed because we assume general position for the vertex sets we are working on. The last case is pseudo-k-gon $\langle a, k, f, g, h, i \rangle$ which has interior objects within the pseudo-k-gon, in this particular case, vertex j .

4.1 Characteristic Graphs

We can relate the pseudotriangulation of a non-degenerate pseudo-k-gon to a graph on a point set with k points in convex position, where k is the number of corners in the pseudo-k-gon. The idea of this graph is to represent the geodesics which exist in the pseudotriangulation as edges in the graph, we shall call such graphs characteristic graphs. Before we introduce the Characteristic Graphs themselves we shall first get more familiar with the geodesics in the pseudotriangulated pseudo-k-gon by exploring the relationship between geodesic triangulations and pseudotriangulations.

Geodesic Triangulations

A geodesic triangulation as introduced by Chazelle et al. [Chazelle 91] is a decomposition of the interior of a simple polygon into geodesic triangles. A geodesic triangle is formed by connecting three vertices of the polygon with the three geodesics between them as illustrated in Figure 4.2(a). Some of the vertices in Figure 4.2(a) are slightly blown up to show the different geodesics, normally the geodesic will run tightly around the vertices. Note the close relation of a geodesic triangle to a pseudotriangle.




Fig. 4.2: A geodesic triangle and a geodesic triangulation

Creating a geodesic triangulation of a simple polygon amounts to adding such geodesic triangles to the polygon as long as the geodesics defining these geodesic triangles do not intersect any of the other geodesics. An example of a geodesic triangulation is given in Figure 4.2(b). Note that a geodesic

triangulation of a simple polygon is also a (not necessarily pointed) pseudotriangulation of that polygon, they are however not the same. As can be seen in Figure 4.3, different geodesic triangulations can result in the same decomposition of the interior of the polygon. This is not the case with pseudotriangulations as pseudotriangulations are defined by their edges, not the geodesics that run from vertex to vertex.




Fig. 4.3: An example of two geodesic triangulations that yield the same decomposition of the interior of the polygon.

We can characterize the geodesic triangulation of a simple polygon by a triangulation of a convex polygon with the same number of vertices corresponding to the vertices of the polygon. Two vertices of the convex polygon are connected by an edge iff they are connected by a geodesic in the geodesic triangulation of the corresponding polygon.




Fig. 4.4: An example of a geodesic triangulation with its characteristic graph.

Figure 4.4 presents an example of a geodesic triangulation with its characteristic graph. Note that a

characteristic graph uniquely defines the geodesic triangulation of the corresponding polygon

A geodesic triangulation is not necessarily pointed, because the geodesics are not only connecting corners of the polygon, but also the non-corners of the polygon. As we are dealing with pointed pseudotriangulations we need to slightly modify the notion of the geodesic triangulation and introduce the pointed geodesic triangulation. In a pointed geodesic triangulation not just any two vertices of the polygon can be connected by their geodesics, only corners of the polygon will be connected with geodesics. Note that this will still result in valid geodesic triangles as the geodesic between two adjacent corners of the polygon will lie along the boundary of the polygon. As with geodesic triangulations we can make a characteristic graph of a pointed geodesic triangulation. In this case however we do not need all vertices of the polygon on the characteristic graph, only the corners of the polygon are needed since only these vertices can be connected by geodesics.

Pseudotriangulations

We now turn our attention back to pseudotriangulations. We have already noted the similarity between geodesic triangulations and pseudotriangulations, but also that they are not exactly the same, for every pseudotriangulation there are a number of different geodesic triangulations. This means that we cannot define a unique characteristic graph for pseudotriangulations as we did for geodesic triangulations, we have a slightly different definition.

A characteristic graph of a pseudotriangulation of a pseudo-k-gon with k corners is a convex polygon with k vertices corresponding to the k corners of the pseudo-k-gon, with interior edges such that;

- a. for every interior edge \overline{uv} of the characteristic graph, the corresponding geodesic $u \leftrightarrow v$ is in the pseudotriangulation; and
- b. for every interior edge e of the pseudotriangulation, e is incident with the geodesic corresponding to at least one of the interior edges of the characteristic graph.

We shall call the interior edges of the characteristic graph geodesics, as they directly represent geodesics in the related pseudo-k-gon.

As a geodesic between two corner vertices of the pseudo-k-gon can be incident with more than one of the interior edges of the pseudotriangulated pseudo-k-gon and more than one geodesic can be incident with an interior edge of the pseudo-k-gon, there can be more or fewer geodesics in the characteristic graph than there are interior edges in the pseudotriangulated pseudo-k-gon.

There is both a unique maximum and unique minimum characteristic graph of each pseudotriangulation of a pseudo-k-gon. The maximum characteristic graph can simply be found by inserting all geodesics which are in the pseudotriangulation of the pseudo-k-gon. Note that the a maximum characteristic graph is a union of the characteristic graphs of all pointed geodesic triangulations that impose the same decomposition of the interior of the pseudo-k-gon as the pseudotriangulation. Conversely, the minimum is the intersection of the characteristic graphs of all pointed geodesic triangulations that impose the same decomposition as the pseudotriangulation. An example of a maximum and minimum characteristic graph is given in Figure 4.5.




Fig. 4.5: An example of a pseudotriangulated pseudo-k-gon with its maximum and minimum characteristic graphs.

Note that while a pseudotriangulation permits several different characteristic graphs, a characteristic graph uniquely defines the pseudotriangulation of a given pseudo-k-gon, this because of the requirement that each interior edge of the pseudotriangulated pseudo-k-gon is covered by at least one of the geodesics in the characteristic graph.




Fig. 4.6: Relation between two pseudotriangulations of a non-degenerate pseudo-k-gon and their maximum characteristic graphs.

Figure 4.6 shows the two unique characteristic graphs of the two pseudotriangulations of pseudo-k-gon $\langle c, d, e, l, f, k, m \rangle$ in Figure 4.1. These characteristic graphs are unique, since the maximum characteristic graph for both pseudotriangulations is also the minimum characteristic graph.

In order to make characteristic graphs of the pseudotriangulations of pseudo-k-gon $\langle a, b, n, o, n, c, m, k \rangle$ in Figure 4.1, vertex n first needs to be split in two as illustrated in Figure 4.7(a). The resulting maximum characteristic graphs of both pseudotriangulations are shown in Figure 4.7(b) and 4.7(c). The maximum characteristic graphs are non-planar in this case.

Note that a characteristic graph can not be made for every pseudotriangulated pseudo-k-gon with interior vertices as the edges incident to any such interior vertices will not necessarily be on a geodesic




Fig. 4.7: Relation between two pseudotriangulations of a degenerate pseudo-k-gon and their maximum characteristic graphs.

between any two corner vertices of the pseudo-k-gon. For example the pseudotriangulations of pseudo-k-gon $\langle a, k, f, g, h, i \rangle$ in Figure 4.1 do not have characteristic graphs.

4.2 Flipping Geodesics

Naturally since we are dealing with flip operations in pseudotriangulations we wish to relate the edge flips in the pseudotriangulations with the geodesics in the characteristic graph. To do this we need a special type of characteristic graph, namely the triangular characteristic graph. A triangular characteristic graph is a maximally planar characteristic graph and hence a triangulation. This means that a triangular characteristic graph is really a characteristic graph of one of the geodesic triangulations that gives the same decomposition as the pseudotriangulation.




Fig. 4.8: Every pseudotriangulation of a pseudo-k-gon has exactly $k - 3$ interior edges.

Note that a pseudotriangulated pseudo-k-gon has exactly $k - 3$ interior edges, a simple inductive argument suffices. A pseudotriangulated pseudo-k-gon of rank 3, i.e. a pseudotriangle, has no interior edges. Every rank the pseudotriangulated pseudo-k-gon increases an edge from the boundary of the pseudo-k-gon becomes an edge of the pseudotriangulation as illustrated in Figure 4.8.

A triangulation has exactly $3n - k - 3$ edges, where n is the number of vertices in the triangulation

and k is the number of vertices on the convex hull of the triangulation. All vertices of a characteristic graph are on the convex hull as the vertices are in convex position, hence $n = k$ for any characteristic graph. Given $n = k$ and the fact that there are k edges on the boundary of the characteristic graph, a triangular characteristic graph has exactly $k - 3$ geodesics.




Fig. 4.9: How faces and edges are matched between a triangular characteristic graph and its corresponding pseudotriangulation of a pseudo-k-gon.

We can now establish a matching between the interior edges of the pseudotriangulated pseudo-k-gon and the geodesic in the triangular characteristic graph. Note that we can view the boundary edges of the characteristic graphs as geodesics as well, since these geodesics will simply follow the boundary of the pseudo-k-gon. Observe that each pseudotriangle in the pseudotriangulation is contained in an area defined by the geodesics represented by the three sides of one of the triangles in the triangular characteristic graph as is defined in the geodesic triangulation, in other words, each triangle in the characteristic graph matches to one of the pseudotriangles in the pseudotriangulation, which we can view as a geodesic triangle. Now it easily follows, as illustrated in Figure 4.9 that a geodesic $u \leftrightarrow x$ in the characteristic graph can be matched with the interior edge \overline{rx} of the pseudotriangulation between the two geodesic triangles corresponding with pseudotriangles T' and S' , which match with the two triangles T and S adjacent to geodesic $u \leftrightarrow x$ in the geodesic graph. It is slightly less obvious that the geodesic $u \leftrightarrow w$ matches with edge \overline{ur} . Observe however that the geodesic triangle formed by the three geodesics $u \leftrightarrow x$, $w \leftrightarrow x$ and $u \leftrightarrow w$ of S' extends past r to u , hence \overline{ur} lies between the geodesic triangles corresponding to S' and R' .

Observe that all geodesics of a triangular characteristic graph can be flipped, since the vertices of the triangular characteristic graph are in convex position. Flipping one of these geodesics will either change the matching and leave the pseudotriangulation unaffected, or flip an interior edge of the pseudotriangulation in the related pseudo-k-gon. When we flip geodesic $u \leftrightarrow x$ in the geodesic graph of Figure 4.9 to $w \leftrightarrow y$, the pseudotriangulation will be changed such that the edge \overline{rx} is flipped to edge \overline{sw} .

In contrast, if we flip geodesic $u \leftrightarrow w$ in the characteristic graph to $v \leftrightarrow x$, the pseudotriangulation can not be changed as the matching edge \overline{ru} in the pseudotriangulation is also incident with geodesic $u \leftrightarrow x$ which is still in the characteristic graph. Instead the matching changes, such that geodesic $u \leftrightarrow x$ matches with edge \overline{ru} instead of \overline{rx} .

4.3 Flip Distance

Since a characteristic graph uniquely defines the pseudotriangulation of a pseudo-k-gon, two pseudotriangulations must be the same pseudotriangulation if they permit the same characteristic graph. We can use this fact and the technique of flipping geodesics as just described to bound the flip distance between two pseudotriangulations of the same pseudo-k-gon. We can take a triangular characteristic graph of each pseudotriangulation and then flip the characteristic graphs towards each other, implicitly flipping the pseudotriangulations towards each other instead of directly.

Theorem 4.1. *A pseudotriangulated pseudo-k-gon of rank k can be flipped to any other pseudotriangulation of the same pseudo-k-gon in time linear in k .*

Proof. By overlaying the triangular characteristic graphs of two pseudotriangulations of the same pseudo-k-gon we can obtain a distance measure between the two pseudotriangulations. Sleator et al. proved that any triangulation of a point set of n points in convex position can be flipped to any other triangulation of the same point set in at most $2n - 10$ diagonal flips [Sleator 86]. Hence we can flip the characteristic graphs into one another by at most $2n - 10$ diagonal flips. As we know that each geodesic flip will either alter the matching or flip the edge in the pseudotriangulation matched to the geodesic, the flip distance of the two pseudotriangulations of a pseudo-k-gon of rank k is at most the flip distance of the characteristic graphs, which is at most $2k - 10$. \square

Note however that this does not give us a bound for the flip distance problem of pseudotriangulations in general. As we have already observed, it is usually not possible to create a characteristic graph of a pseudo-k-gon with interior vertices. Unfortunately problem instances in general do exhibit such interior vertices and hence we can not apply this technique to the problem as a whole. It is however possible to transform the original problem instance by performing edge flips, such that any interior vertices are connected to the boundary of the pseudotriangulation, creating a large degenerate pseudo-k-gon of the whole problem. The resulting subproblem in the degenerate pseudo-k-gon can then subsequently be solved in $\mathcal{O}(n)$ flips. The question is, if such a transformation can be bounded by some input-sensitive bound.

5. FLIPS

In an earlier chapter good and bad flips were introduced which were defined in terms of their effect on the distance measure. Given the difficulty of devising a measure for which the existence of an ever decreasing flip path can be proved, a new viewpoint on the problem was needed. Since a large part of the problem can often be solved by performing a series of perfect flips it stood to reason to examine what different types of flips exist. These flips are defined outside of the scope of the distance measure, for some distance measures these types of flips will reduce the distance measure, but this is certainly not always the case.

One of the types of flips, the perfect flip, has already been introduced earlier, it is a flip in one of the two pseudotriangulations, whose target is an edge in the other pseudotriangulation. A perfect flip creates a new shared edge between the two pseudotriangulations.

5.1 Intersection Reducing Flips

An intersection reducing flip is a flip which, as its name implies, reduces the number of intersections between the two different pseudotriangulations. This is the basic flip which we encounter in the triangulation flip distance problem. These flips can clearly introduce new pointedness violations and thus will not always be good flips for every distance measure.

5.2 Pairwise Perfect Flips

From the problem instances generated with genptu it became apparent that in many cases, where there were no perfect flips, it was possible to flip one edge from the one pseudotriangulation and a second edge from the second pseudotriangulation towards the same target edge, a so-called pairwise perfect flip as illustrated in Figure 5.1. This figure is the same figure as the one used to illustrate the perfect flip in Chapter 1, but this time we focus on the non-perfect flips. Note that the flips of edges \overline{bd} and \overline{be} are both non-perfect flips, but that the flip target is \overline{ac} in both cases. This means that the edges \overline{bd} and \overline{be} constitute a pairwise perfect flip.

Many of the problems could be almost completely solved by only executing perfect flips and pairwise perfect flips, thus the natural question arose whether it would be possible to solve all problem instances with only perfect flips, pairwise perfect flips and intersection reducing flips.




Fig. 5.1: Example of a pairwise perfect flip.

5.3 Sequence of Known Flips

An algorithm which at each point would use an intersection reducing flip, a perfect flip or a pairwise perfect flip, would obviously take $\mathcal{O}(|\text{intersections}| + |\text{differing edges}|)$ flips and thus would define a suitable input dependent distance measure. The question was whether it would be possible to prove the existence of one of these three flips for any given problem instance.

Günter Rote presented an example where none of these types of flips could be made, disproving the possibility of an algorithm such as sketched above. This example is given in Figure 5.2.

Observe that there is no perfect flip in Figure 5.2. Furthermore there does not exist a pairwise perfect flip; edge \overline{kl} flips to \overline{eh} , however edge \overline{dg} , being the only possibly edge of the other triangulation to flip to the same target flips to \overline{hi} . Note that due to the symmetric nature of the example no other pair of edges can flip to the same target edge. Finally observe that both unique possibly edge flips disregarding symmetric edges do not decrease the number of intersections. Edge flip \overline{kl} to \overline{eh} keeps the number of intersections constant by removing one and introducing another, while flip \overline{dg} to \overline{hi} increases the number of intersections by two.

The existence of this problem instance gives rise to a completely new way to view the problem and define the Geodesic Measure. Instead of looking at the edges themselves and the non-pointedness of vertices we consider the geodesics defining each edge within the pseudoquadrilateral that contains that edge. This is further discussed in Chapter 6.




Fig. 5.2: Example showing the need for a geodesic distance measure.

6. CONCLUSIONS

6.1 Thesis Objectives

In Chapter 1 we defined our problem statement as:

- does there exist a distance measure such that there always exists a flip path between any two pointed pseudotriangulations which strictly decreases this measure, or
- can the existence of such a measure be disproved

Unfortunately we have not been able to show the viability of a suitable distance measure, nor have we been able to disprove the existence of such a measure. Despite various efforts, no way of finding a counter example for the Extended Graph Measure or proving the viability of the distance measure yielded any results. This shows that we do not yet have a sufficient understanding of the interactions between two pseudotriangulations while flipping edges in these pseudotriangulations.

When this became apparent a new objective was formulated, namely to achieve a better understanding of the interactions between two pseudotriangulations while flipping edges in these triangulations. To this end supporting research software was written to provide new problem instances, the reason being that we had only been looking at relatively small problem instances and could possibly be unintentionally restricting ourselves when drawing up these instances. This software has given us new insights into how the problem could be approached and with further study helped us to better understand what happens at the geodesic level in the pseudotriangulations in which we are flipping edges.

Although we were unsuccessful in proving the viability of a suitable distance measure, we have opened up a number of different new possibilities for finding such a distance measure and have strengthened the understanding of the problem at hand. Some possible directions further research into this problem could take are presented in the next section.

6.2 Future Work

Unfortunately some of the ideas and insights conceived during the course of the thesis work could not be further researched due to lack of time. This section describes these ideas and insights such that they may be further researched at a future point in time. Furthermore some of the researched matter deserves further attention when a better insight into the problem has been developed.

Transforming Into Pseudo-k-gons

In Chapter 4 we saw that it is possible to flip a pseudotriangular embedding of a pseudo-k-gon into another pseudotriangular embedding of the same pseudo-k-gon in $\mathcal{O}(n)$ flips. This implies that, if we are able to construct an algorithm which connects all interior vertices of a problem instance to its boundary, such that the algorithm takes a number of flips bounded by some k , which is input dependent, we could derive an algorithm that flips two pseudotriangulations into one another in $\mathcal{O}(n + k)$ flips. It would therefore be worthwhile to research the possibility of such an algorithm.

Geodesic Measure

The experiences with the pseudo-k-gons and their geodesic graphs combined with the problem instance that Günter Rote provided as shown in Chapter 5 lead to a new viewpoint on the original problem. Instead of looking at non-pointedness directly, we will be looking at the geodesics defining the edges incident to the non-pointed vertices. These geodesics will be intersecting one another at the non-pointed vertices.

The geodesic measure does not count non-pointedness by the number of non-pointed vertices, their degree or some comparable method. Instead non-pointedness is 'translated' into intersections between the geodesics defining the edges of the pseudotriangulations. This way the distance between the pseudotriangulation can be expressed solely by the number of intersections between these geodesics and the problem reverts to a form more reminiscent of the original triangulation flip distance problem.




Fig. 6.1: The geodesic that defines an edge.

As mentioned in Chapter 1, each interior edge is a diagonal in a containing pseudoquadrilateral. A diagonal, and hence every interior edge, lies on the geodesic connecting two opposing corners of this pseudoquadrilateral, as illustrated in Figure 6.1. In our example edge \overline{yz} is contained by the pseudoquadrilateral $\langle u, z, x, w, y, v \rangle$ and lies on the geodesic $v \leftrightarrow x$. We say this geodesic defines the edge.

We can now view a problem instance not as two sets of edges in the plane, but instead as two sets of geodesics in the plane. Intersections between the edges remain intersections between the geodesics

as the edges are always part of the geodesic. Non-pointedness is represented by geodesics wrapping around vertices in two different directions and thus intersecting as illustrated in Figure 6.2. Figure 6.2 deviates slightly from our normal notational convention for problem instances; vertex e is blown up slightly so that the geodesics can be shown to wrap around this vertex. Shown in the figure are the geodesics defining edges \overline{ae} and \overline{be} , instead of those edges themselves. Note how these two geodesics intersect at vertex e , such intersections of geodesics around a vertex only occur at non-pointed vertices.




Fig. 6.2: An example showing how the two geodesics of edges incident to a non-pointed vertex intersect.

Distance Measure 6.1 (Geodesic Measure).

Let $\text{geodesics}(\mathcal{T})$ be the set of geodesics defining the interior edges of pseudotriangulation \mathcal{T} .

$$\mathcal{D}_{\text{geodesic}}(\mathcal{T}_1, \mathcal{T}_2) = |\mathcal{I}(\text{geodesics}(\mathcal{T}_1), \text{geodesics}(\mathcal{T}_2))|$$

The distance from \mathcal{T}_1 to \mathcal{T}_2 is the number of intersections between $\text{geodesics}(\mathcal{T}_1)$ and $\text{geodesics}(\mathcal{T}_2)$.

Since this measure takes a completely different viewpoint to the problem at hand a new look at how this measure behaves under various different flips is required to establish whether or not this measure provides a viable input-sensitive bound on the flip distance of pointed pseudotriangulations.

Extended Graph Measures

The Extended Graph Measures introduced in Chapter 2, have not yet been disproved as viable and our intuition tells us that these measures are very likely candidates for a suitable distance measure. The developments with the pseudo-k-gons and the Geodesic Measure have provided us with a better understanding of the problem however. This new understanding strengthened by further research in these matters could open up fresh possibilities on either proving or disproving the viability of the Extended Graph Measures.

6.3 Summary

Although a suitable distance measure was not found, much has been learned during the course of the thesis work. The evaluation of the various distance measures has shown us what does not work and what might work in addition to providing us with a better understanding of the problem at hand. The explorations into the pseudo-k-gons and geodesic graphs have opened up another way in which a suitable flip distance bound could be established. Furthermore the pseudo-k-gons and their relations with the geodesic graphs have provided us with valuable information on how the edge flips in a pseudotriangular embedding behave with respect to the geodesics in the pseudotriangulation.

6.4 Acknowledgments

First and foremost I would like to thank my parents for all their support during my studies. Furthermore I would like to thank the following people for their help and support in the research work and the development of this thesis.

- My supervisor, Bettina Speckman, for all her patience and support,
- The attendees of the Fourth European Pseudotriangulation working week for all their input on my research work and insights into the problem, and
- Brian Pauw for his invaluable feedback on this thesis.

BIBLIOGRAPHY

- [Bereg 04] Sergey Bereg. *Transforming Pseudo-Triangulations*. Information Processing Letters, vol. 90, no. 3, pages 141–145, 2004.
- [Chazelle 91] Bernard Chazelle, Herbert Edelsbrunner, Michelangelo Grigni, Leonidas Guibas, John Hershberger, Micha Sharir & Jack Snoeyink. *Ray Shooting in Polygons Using Geodesic Triangulations*. In Proceedings of the 18th International Colloquium on Automata, Languages and Programming, pages 661–673, New York, NY, USA, 1991. Springer-Verlag New York, Inc.
- [Cheong 06] Otfried Cheong. *The IPE manual*, 2006.
- [Hanke 96] Sabine Hanke, Thomas Ottman & Sven Schuierer. *The Edge-flipping Distance of Triangulations*. Journal of Universal Comp. Science, vol. 2, no. 8, pages 570–579, 1996.
- [Hurtado 96] F. Hurtado, M. Noy & J. Urrutia. *Flipping Edges in Triangulations*. In SCG '96: Proceedings of the Twelfth Annual Symposium on Computational Geometry, pages 214–223, New York, NY, USA, 1996. ACM.
- [Kettner 03] Lutz Kettner, David Kirkpatrick, Andrea Mantler, Jack Snoeyink, Bettina Speckmann & Fumihiko Takeuchi. *Tight Degree Bounds for Pseudo-Triangulations of Points*. Computational Geometry: Theory and Applications, vol. 25, no. 1&2, pages 1–12, 2003.
- [Matsumoto 98] Makoto Matsumoto & Takuji Nishimura. *Mersenne Twister: a 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator*. ACM Transactions on Modeling and Computer Simulation, vol. 8, no. 1, pages 3–30, 1998.
- [Pocchiola 93] Michel Pocchiola & Gert Vegter. *The Visibility Complex*. In SCG '93: Proceedings of the Ninth Annual Symposium on Computational Geometry, pages 328–337, New York, NY, USA, 1993. ACM.
- [Rote 08] Günter Rote, Francisco Santos & Ileana Streinu. *Pseudo-Triangulations - A Survey*. vol. 453, pages 343–410, 2008.
- [Sleator 86] Daniel Sleator, Robert Tarjan & William Thurston. *Rotation Distance, Triangulations, and Hyperbolic Geometry*. In STOC '86: Proceedings of the Eighteenth Annual

ACM Symposium on Theory of Computing, pages 122–135, New York, NY, USA, 1986. ACM.

- [Streinu 05] Ileana Streinu. *Pseudo-Triangulations, Rigidity and Motion Planning*. Discrete and Computational Geometry, vol. 34, no. 4, pages 587–635, 2005.