Portalgons gone curvy

Shortest paths in a subset of portalgons with a fitting portal

Bachelorthesis *Mathematics* (*WISB399*)
7.5 ECTS

Universiteit Utrecht

Bernadet Klein Wassink,
6677959,

under guidance of *Frank Staals* and *Marieke van der Wegen*

17/6/2022

Abstract

In Löffler et al.'s paper, they define portalgon as a set of fragments together with a set of portals that "glue" fragment together along the portals. A **portal** $e = (e^+, e^-)$ is a pair of directed, equal length edges of two, possibly the same, fragments of F. Let $e^+ = \overrightarrow{uv}$ and $e^- = \overrightarrow{wz}$, where u, v vertices of the fragment containing e^+ and w, z vertices of the fragment containing e^-. Let p^- be a point on e^- and p^+ on e^+. If there is a $\lambda \in [0, 1]$ for which $p^- = \lambda u + (1 - \lambda)v$, then $p^+ = \lambda w + (1 - \lambda)z$. Furthermore, no edge can be part of multiple portals. In their paper, they explore shortest paths in these portalgons and find out that there is a way to represent the portalgon in such a way that the complexity of shortest paths is bounded. For this, they introduce happiness h; a portalgon is h-happy if any shortest path crosses every portal at most h times. The complexity of shortest paths in an h-happy portalgon is $O(n + hm)$ where n is the number of vertices in the portalgon and m is the number of portals. Furthermore, they also discovered that every polyhedral surface has a representation as a portalgon where the happiness is constant. This means the complexity of shortest paths in the representation is $O(n)$.[9] Note that an object with input size n has asymptotic upper bound $O(f(n))$ if and only if $\exists N, c > 0, \forall n > N, g(n) \leq c \cdot f(n)$, where $g(n)$ is the amount of components of the object that need constant space [7].

In this thesis, we expand their research into the realm of portalgons where the portals do not have to be linear anymore. Instead we are going to research fitting portals specifically, which are those portals where two corresponding edges are described by the same function and orientation. In other words, exactly those non-linear portals that fit together. We introduce the concepts of portal line segment (line segment between the start and end vertex of the portal) and obstacles (the boundary of the fragment, not the portal itself, intersecting with the portal line segment). In this thesis, we show that any portalgon P with n vertices, m straight portals (portals where its form is described by the portal line segment) and one fitting portal without any intersections between the portal and its portal line in an obstacle has an equivalent portalgon with $O(n)$ vertices and straight portals. Then, it also follows that P has a representation as a portalgon where the happiness is constant. This means the complexity of shortest paths in the representation is $O(n)$.

Acknowledgements

I want to thank three people that have helped me immensely in this process. I first want to thank my supervisors for helping me with my thesis and especially for helping me limit the scope of the thesis. Furthermore, I want to thank my boyfriend for listening to me and bearing with me in all those moments where I was frustrated with overlooking minor details that were causing major problems.

Contents

1 Introduction

Portalgons are a concept from “Shortest Paths in Portalgons” [9]. In this thesis, we will be looking at a new kind of portalgon. Originally, the problem that was brought to the attention of the authors of [9] was that of Voronoi diagrams on clothing patterns, where the seam lines are the portals between the pieces of cloth. A Voronoi diagram is a partition of a plane into regions closest to each of a given set of objects/points. This problem was translated into subsets of the plane with an equivalence relationship defining the seam lines as portals. In the paper, the authors discuss the problem of the shortest paths, which are essential to Voronoi diagrams, in a space composed a set of polygons connected by portals.

A set of these simple polygons together with the portals connecting them is called a **portalgon**. A portal e is a pair of two straight edges e^+ and e^- of equal length, where if the interior of the polygon is on the left of one edge, it is one the right of the other. Let $e^+ = \overrightarrow{uv}$ and $e^- = \overrightarrow{wz}$, where u, v are vertices of the fragment containing e^+ and w, z are vertices of the fragment containing e^-. Let p^- be a point on e^-. If there is a $\lambda \in [0, 1]$ for which $p^- = \lambda u + (1-\lambda)v$, then $p^+ = \lambda w + (1-\lambda)z$ and p^+ on e^+. Two portalgons are equivalent when they are isometric, so when the quotient spaces created by “gluing” the portals are isometric. Refer to Figure 1 for some examples of the portalgons Löffler et al. researched.

A h-happy portalgon is a portalgon where any shortest path between two point in the portalgon cross any portal at most h times. Loffler et al. found out that any portalgon with n vertices (the points describing the polygons) and m portals has a representation where the complexity of a shortest path is $O(n + hm)$. The notation $O(n + hm)$ means that the number of component of the shortest path is bounded by a multiple of $n + hm$ for n and m big enough. They also discovered that a polyhedral surface, a surface composed of planar polygons [1], has a representation as a portalgon where complexity of any shortest path between two points on it is $O(n)$.

I have worked the last couple of months on a subset of portalgons where the portal edges are not straight anymore. To formulate what we are exactly going to prove, we first need a few concepts. The portal line segment is the line segment between the start and end vertex of the portal. Straight portals are those portals following the portal line segment fully, so portals that are a line segment between their start and end vertex. Let fragment f containing portal edge e^+, then an obstacle to e^+ is a part the boundary of f intersecting with the portal line segment of e^+ that does not contain any part of e^+.

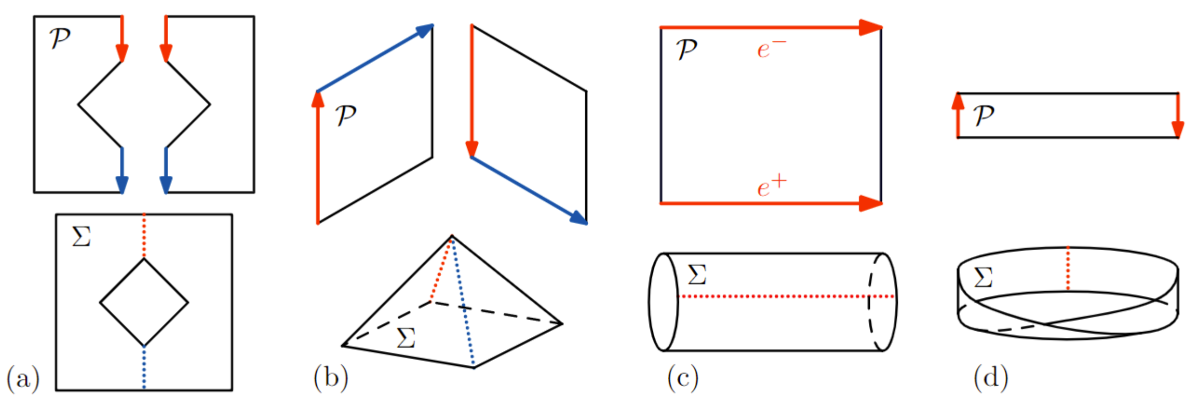

Figure 1: These are a few examples of normal portalgons. Here $\mathcal{P}$ is the portalgon represented as fragments connected by portals and Σ is quotient space created by “gluing” all corresponding portal edges. The portalgon (a) is a portalgon with a hole. The portalgon (b) is a pyramid without a floor. Portalgon (c) is a cylinder and portalgon (d) is a Möbius strip. [9]

Now we can define those portalgons essential in this thesis; the portalgons we will be looking at are those that have at most one portal where neither of the edges intersect with the portal line segment via an obstacle and where the edges fit together. The former intuitively means that there are no intersections of the portal line segment and the portal edge between the intersections of a obstacle with the portal line segment. The latter means the portal edges are described by the same function, so they "fit" like puzzle pieces.

Note that we still restrict the portals to those where the edges have the same length. Refer to Figure 2 for a simple example of one of the portalgons that will be considered, to Figure 3 for more complex examples of what portalgons will be considered and refer to Figure 4 for examples of what portalgons will not be considered. We are going to prove the following theorem in the following thesis.

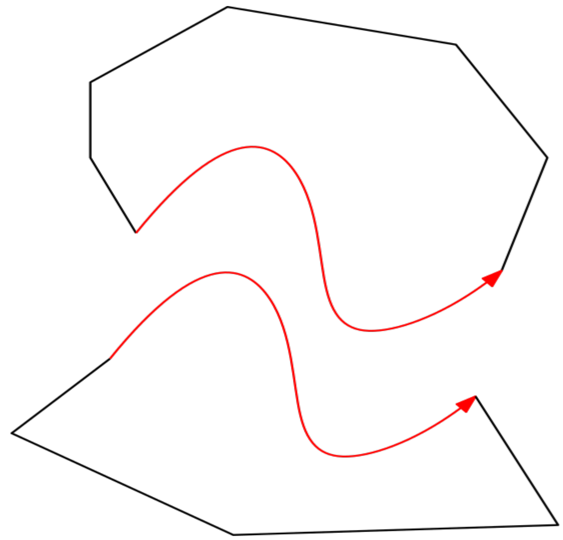

Figure 2: A simple example of a portalgon that has a fitting portal.

Theorem. Let $\mathcal{P}$ be a portalgon with n vertices, m straight portals and one non-straight fitting portal that does not intersect with the portal line segment via an obstacle. Then there is a portalgon equivalent to $\mathcal{P}$ that has no curvy portals, $O(n)$ vertices and $O(n)$ straight portals. Furthermore, $\mathcal{P}$ has an equivalent h-happy portal with $O(hn)$ the complexity of its shortest paths.

First, within this chapter we will discuss the related works. Then, in the chapter "Preliminaries", the theoretical foundations are laid down. Specifically, we will discuss the necessary theory on complexity; graph theory; simple polygons and the portalgon. Lastly, we explore a few of the open problems in this field and how they are related to this problem.

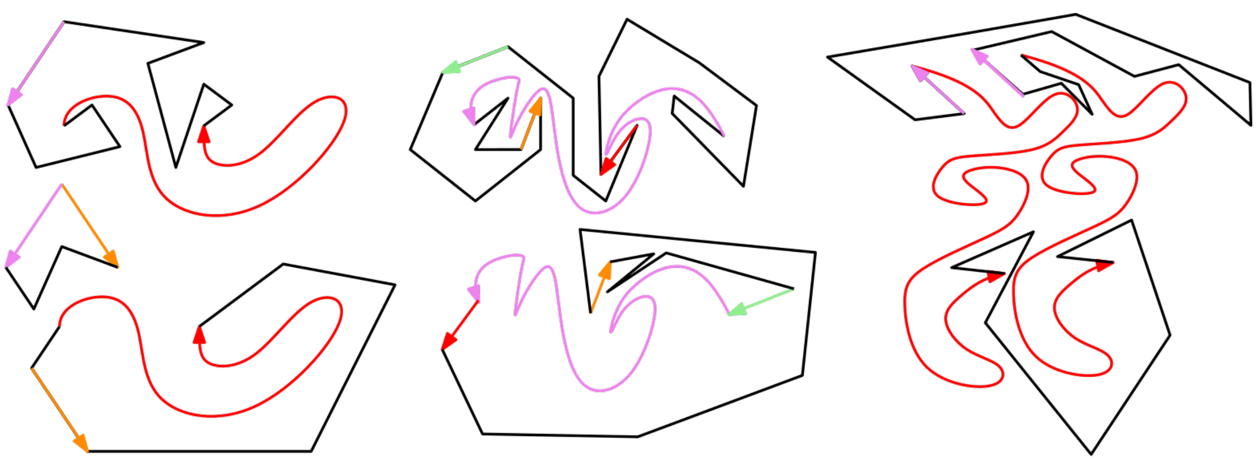

Figure 3: These for portalgons are prime examples of what we are going to explore in this thesis. The orientation of the portals (the direction from start to end vertex) is denoted with the arrow. Note that all portals of the same colours are the same length, have the same orientation and are the same form. We can have obstacles but note that there are no portal intersections within these obstacles.

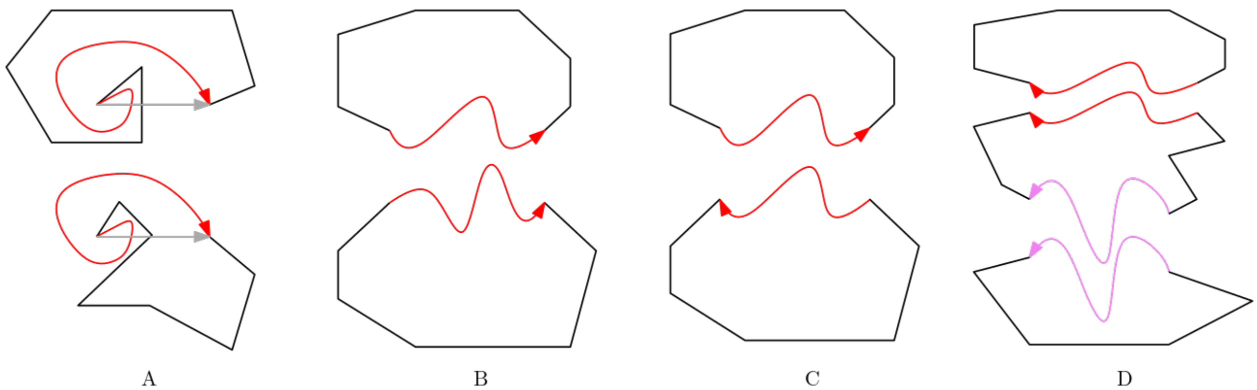

Figure 4: These for portalgons are prime examples of what we are not going to explore in this thesis. The orientation of the portals (the direction from start to end vertex) is denoted with the arrow. In portalgon A, there is a fitting portal that intersects with the portal line segment via an obstacle. The grey arrow in portalgon A is the portal line segment of portalgon A. In portalgon B, there is a non-fitting portal. In portalgon C, there is a portal where the edges do not have the same orientation (and so, in this case, they are not the same form). In portalgon D, there is a fragment with multiple non-straight portal edges on one fragment.

1.1 Related work

There has been done a lot of research into computing the shortest paths in all kind of different spaces, such as polygons or a normal 2D space but with polygonal obstacles.

In the paper "An Optimal Algorithm for Euclidean Shortest Paths" by Herschberger and Suri [6], they worked out an optimal structure, the shortest path map, for shortest path queries for a given source in the plane with polygonal obstacles. They compute the structure by maintaining a wavefront, a collection of waves (arcs) signifying at all points of the distance travelled in that time to the source, which is updated every time there is an event: wave-wave collision, wave-obstacle collision or disappearance of a wave. Later H. Wang [12] improved the space complexity of this to also be optimal.

The shortest path map can also be computed in other ways. L. Gewali and V. Roman looked into setting up a shortest path map, but in a way that a similar structure could also be used to figure out shortest paths without any sharp turns. [4]

F. Li and R. Klette have written the book "Euclidean Shortest paths". [8] The chapters "Euclidean Shortest Paths", "Partititioning a Polygon in the Plane" and "ESPs [Euclidean Shortest Paths] in Simple Polygons" are the most applicable to our problem of finding shortest paths around and in simple polygons. They explore the possibilities of using graphs to describe the problem, but also using an alternate form of Dijkstra, continuous Dijkstra.

In the paper "On the Geodesic Voronoi Diagram of Point Sites in a Simple Polygon" by B. Aronov [2], the concept of Voronoi diagrams in simple polygons and shortest paths in a simple polygon are explored. A Voronoi diagrams is a partition of a plane into regions closest to each of a given set of objects/points.

Then, in my opinion, the most important paper, "Shortest Paths in Portalgons" by M. Löffler, T. Ophelders, F. Staals and R.I. Silveira. This paper looks at our problem, but then specifically for only straight portals. The most important result is that they have proven that the shortest path any portalgon (defined as a set of polygons with "portals" connecting them) with n vertices and m portals has a bounded complexity, namely $\Theta(n + hm)$. And that any polyhedral surface has a representation as a portalgon where the happiness is constant. In such portalgons, a shortest path has complexity $O(n)$.[9]

2 Preliminaries

This section is on the theory that is the foundation of our discoveries.

2.1 Complexity

We are going to explain complexity based on the book "Introduction to Algorithms" by Cormen et al [7]. Complexity is a measure of the resources (time/space) used in an algorithm or by an object. This is a way to effectively compare algorithms to each other.

Definition 2.1.1. A **basic operation** is an operation that takes a constant amount of time or space.

Definition 2.1.2. **Basic operations in time complexity** are addition, subtraction, multiplication and division.

The upper bound of an algorithm or object is a way to measure what it needs at most in time or space. It is a measure of how good it at least is.

Definition 2.1.3. An algorithm or object with input size n has **asymptotic upper bound** $O(f(n))$ if and only if there exists a $N, c > 0$, for all $n > N, g(n) \leq c \cdot f(n)$, where $g(n)$ is the number of basic operations of the algorithm or the number of components of the object that need constant amount of space.

Remark 2.1.4. A basic operation is $O(1)$.

There are infinitely many types of complexities, yet, three are the most common with input n: exponential $(O(a^{f(n)}), a \in \mathbb{R})$, polynomial $(O(\sum_0^r a_r n^r), r \in \mathbb{N}, a_r \in \mathbb{R})$ and logarithmic $(O(\log_b(f(n))), b \in \mathbb{R})$. When comparing two complexities, the "="-mark is an inequality, where the right complexity has to be asymptotically bigger than or equal to the left complexity. For example, $O(n) = O(n^2)$ as there exists a N and a c, where for all $n > N$, $n \leq c \cdot n^2$ and so, $O(n^2)$ is asymptotically bigger than $O(n)$. [11]

Example 2.1.5. We will show that $O(n^{30}) = O(2^n)$. So, we want to show that there exists a $N, c > 0$, for all $n > N, , n^{30} \leq c \cdot 2^n$. Take $c = 10^{28}$, then it is easy to verify that $n^{30} \leq 10^{28} \cdot 2^n, \forall n > 50$.

Example 2.1.6. We will show that $O(\log^3(n)) = O(n^{1/2})$. So, we want to show that there exists a $N, c > 0$, for all $n > N, \log^3(n) \leq c \cdot n^{1/2}$. Take $c = 10^{20}$, then it is easy to verify that $\log^3(n) \leq 10^{20} \cdot n^{1/2}, \forall n > 0$.

Between exponential terms, a bigger base is leading over a smaller base, so $O(2^n) = O(3^n)$. Exponential terms are always asymptotically bigger than polynomial ones, even if the polynomial has a term n^{100} and the exponential is 2^n. Between polynomial terms, the term with the highest exponent is the asymptotically biggest. Polynomial terms are always asymptotically bigger than logarithmic ones, even if the polynomial has a term $n^{1/2}$ and the logarithmic is $\log^3(n)$.

When comparing mixed terms you compare the parts that are different, so in the case of $O(n^2)$ and $O(n \log(n))$, we should compare $O(n)$ and $O(\log(n))$ to figure out which is asymptotically bigger. When we have multiple terms in the upper bound, we take the term that is asymptotically bigger, so $O(n^2+n) = O(n^2)$. But if we have multiple terms, where we do not know which is asymptotically bigger, we leave them both in, so $O(n + k) = O(n + k)$. Note also that $O(5n^2) = O(n^2)$ as we are not interested in the coefficient, but in the leading term.

The lower bound of an algorithm or object is a way to measure what it at least needs in time or space. It is a measure of how bad it at least is.

Definition 2.1.7. An algorithm or object with input of size n has **asymptotic lower bound** $\Omega(f(n))$ if and only if there exists a $N, c > 0$, for all $n > N, g(n) \geq c \cdot f(n)$, where $g(n)$ is the number of basic operations of the algorithm or the number of components of the object that need constant amount of space.

Then, we also have the asymptotic tight bound for when it is both $\Omega(f(n))$ and $O(f(n))$.

Definition 2.1.8. (Asymptotic tight bound $\Theta(f(n))$) An algorithm or object with input of size n has asymptotic optimal bound $\Theta(f(n))$ if and only if there exists a $N, c_1, c_2 > 0$, for all $n > N, c_2 \cdot f(n) \leq g(n) \leq c_1 \cdot f(n)$, where $g(n)$ is the number of basic operations of the algorithm or the number of components of the object that need constant amount of space.

Example 2.1.9. [11] See the algorithm below.

```
program sum_of_squares:
    s = 0;                     O(1) - assigning a value can be done in constant time
    Do i = 0 to n
        j = i * i;          ||  O(1) - this is a basic operation
        s = s + j;             O(1) - this is a basic operation
    return s
```

We loop n times over a constant body of code, so the time complexity is $O(n)$.

Example 2.1.10. Suppose you have n values and every value takes $O(1)$ space, then the total complexity is $O(n)$. Suppose we have n rows with each m values, where we do not know whether m is bigger or smaller than n, then we have complexity $O(nm)$.

2.2 Graph Theory

For the theory in this chapter, we follow the book "Graph Theory" by Keijo Ruohonen [10]. A graph is a set of vertices connected by edges. Edges are pairs of vertices, denoting a connection between those vertices. We denote a graph with $G = (V, E)$, where V is the set of vertices and E is the multiset of edges. Edges can occur more than once in E. A **directed graph** is a graph where the edges have directions, so any movements through the graph have to commit to these directions. In an **undirected graph**, the edges (v, u) and (u, v) are the same. A **weighted graph** is a graph where the edges have weights. Depending on the graph these weights can be defined differently, but as we will be working with distance, our graphs will be positively weighted. Then, some definitions on sequences of vertices and edges in a graph.

Definition 2.2.1. A **walk** is a finite sequence of alternating edges and vertices $v_0, e_1, v_1, ..., e_n, v_n$, where every edge e_i has end vertices v_{i-1} and v_i. The first vertex v_0 is the initial vertex and the last vertex v_n is the terminal vertex. The length of a walk is the sum of the weights of the edges and if it is not a weighted graph, the length is the number of edges.

Definition 2.2.2. A **path** is a walk where no edge or vertex is visited twice. A circuit is a path where the initial and terminal vertex is the same. A walk can contain circuits if there is a part of the walk that is a circuit.

Now, with these definitions we can define some types of graphs.

Definition 2.2.3. Two vertices in a graph are **connected** when there exists a path between them. A graph is **connected** when all pairs of vertices are connected.

Definition 2.2.4. A graph is a **tree** when the graph is connected and there are no circuits. A tree can have a designated root. When a root is deleted the remaining connected components are called branches. The size of a branch is determined by the number of vertices in that component.

2.3 Simple polygons

To establish anything new, we first need to establish what we already know. In this chapter, we discuss the current theory on the topic of simple polygons. First, we need the definition of a few curves.

Definition 2.3.1. [3] A **simple closed curve** is a subset of the plane that is homeomorphic to the unit circle. Alternatively, a simple closed curve is the image of a continuous injective function from the unit circle into the plane.

Definition 2.3.2. [3] A closed curve is **polygonal** if it is the union of a finite number of line segments. The place where line segments come together is called vertices.

So, a simple closed curve means it has to be a continuous loop and no self-intersections. To define the simple polygon, we also need the **Jordan Curve Theorem**:

Definition 2.3.3. [5, Definition 1] A **Jordan Curve** is a simple closed curve. Alternatively, a Jordan curve is the image of a continuous function $\phi : [0, 1] \mapsto \mathbb{R}^2$ such that $\phi(0) = \phi(1)$ and the restriction of ϕ to $[0, 1)$ is injective.

Theorem 2.3.4. *[5, Theorem 1] (Jordan Curve Theorem) Let C be a Jordan curve in the plane $\mathbb{R}^2$. Then its complement, $\mathbb{R}^2 \backslash C$, consists of exactly two connected components. One of these components is bounded (the interior) and the other is unbounded (the exterior), and the curve C is the boundary of each component.*

Finally, we will define the simple polygon.

Definition 2.3.5. Let $\phi : [0, 1] \mapsto \mathbb{R}^2$ be a simple polygonal closed curve. The **simple polygon** is the interior created by ϕ together with ϕ itself.

Example 2.3.6. Refer to Figure 5, there are a few examples of forms created by line segments. The right one is a simple polygon as it has no self-intersections and is closed. The middle one is not a simple polygon as it is not just a closed loop, it has a extremity. The left one is not a simple polygon as it self-intersects.

Definition 2.3.7. Two points are visible to each other within a polygon if there is a line segment between the two points and this edge is within the polygon. Two edges, e_i and e_j are visible to each other if there exists $p_i \in e_i, p_j \in e_j$ for which the points p_i and p_j are visible to each other.

Then, to be able to define what a shortest path is, we first need to know what a path is, which is more or less a sequence of lines between points.

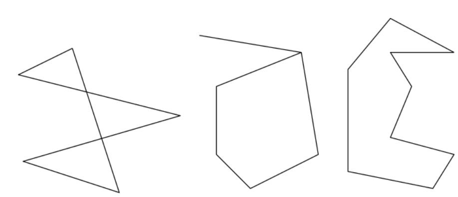

Figure 5: The different kinds of Figures made out of line segments.

Definition 2.3.8. Let R a subset of $\mathbb{R}^2$. A **path** from p to q in R is an alternating sequence of vertices and lines between vertices where every line and vertex lies within R. The length of every segment is its Euclidean length, the length of a path is the sum of the lengths of its components. We say the size of a path is the number of segments.

The shortest path between p and q is the path between p and q with the smallest length. The length of the shortest path is called the distance between p and q. Distance between p and q is denoted with $d(p, q)$. If there is no path between p and q, there is no shortest path and the distance between the two points is ∞.

2.4 Portalgons

The theory in this chapter is based on the theory in the paper "Shortest Paths in Portalgons" by Löffler et al [9]. First, we start with some basic definitions.

Definition 2.4.1. A **portal** $e = (e^+, e^-)$ is a pair of directed, equal length edges of two, possibly the same, fragments.

Let $e^+ = \overrightarrow{uv}$ and $e^- = \overrightarrow{wz}$, where u, v are vertices of the fragment containing e^+ and w, z are vertices of the fragment containing e^-. Let p^- be a point on e^-. If there is a $\lambda \in [0, 1]$ for which $p^- = \lambda u + (1 - \lambda)v$, then $p^+ = \lambda w + (1 - \lambda)z$ and p^+ on e^+.

Definition 2.4.2. A portalgon $\mathcal{P}$ is a pair (F, P), where F is a collection of fragments embedded in $\mathbb{R}^2$, which are simple polygons, and P a collection of portals.

Furthermore, no edge can be part of multiple portals. Thus, the upper limit on the number of portals is $m/2$, so $O(m) = O(n)$, where m is the amount of portals and n the amount of vertices of the portalgon. Then, if we "glue" the portals together, we get the surface Σ. We can write Σ as a quotient space $(\bigcup_{f \in F} f)/\sim$, where the equivalence relationship is exactly those of the portal edges.

We call two portalgons $\mathcal{P}$ and $\mathcal{Q}$ **equivalent** if the surfaces they represent are isometric. This is the case if there is a bijective isometry g between them, so, if the distance between $p, q \in \mathcal{P}$ is the same as the distance between $g(p), g(q) \in \mathcal{Q}$, for all p and q.

We call a portalgon happy if all shortest paths in the portalgon revisits any fragment a constant number of times. In other words, let $\mathcal{P} = (F, P)$ be a portalgon, and let $f \in F$. We define $c(X)$ as the number of connected components in a set $X \subset f$, where $f \in F$. The **happiness** $\mathcal{H}(f) = \max_{p,q \in \mathcal{P}} \max_{\pi \in \pi(p,q)} c(f \cap \pi)$ of fragment f is defined as the maximum number of times a minimum complexity shortest path π between any pair of points $p, q \in \mathcal{P}$ can go through the fragment. The happiness of $\mathcal{P}$ is then defined as the maximum happiness of its fragments $\mathcal{H}(\mathcal{P}) = \max_{f \in F} \mathcal{H}(f)$. A portalgon $\mathcal{P}$ is h-happy if $\mathcal{H}(\mathcal{P}) \leq h$.

Then, the following theorem is their most important result of [9].

Theorem 2.4.3. *Let $\mathcal{P}$ be a h-happy portalgon. Every shortest path in $\mathcal{P}$ has complexity $\Theta(n + hm)$, where n is the total amount of vertices of the fragments of $\mathcal{P}$ and m the number of portals.*

3 Portalgons with fitting portals

The last chapter explored the theory of portalgon with straight portals. In this chapter we will be looking at portalgons with non-straight portals, specifically fitting portals (non-straight portals where both portal edges still have the same form). We start by laying the groundwork with definitions on some general concepts of portalgon when portal edges are allow not to be straight. Next, we explore then we explore a subset of fitting portals, where we are going to prove that a portalgon $\mathcal{P}$ with only straight portals except for one fitting portal always have an equivalent portalgon with only straight edges with a bounded complexity. Then with the theory on these portalgons, we know that $\mathcal{P}$ has an equivalent portalgon where the complexity of the shortest path is bounded.

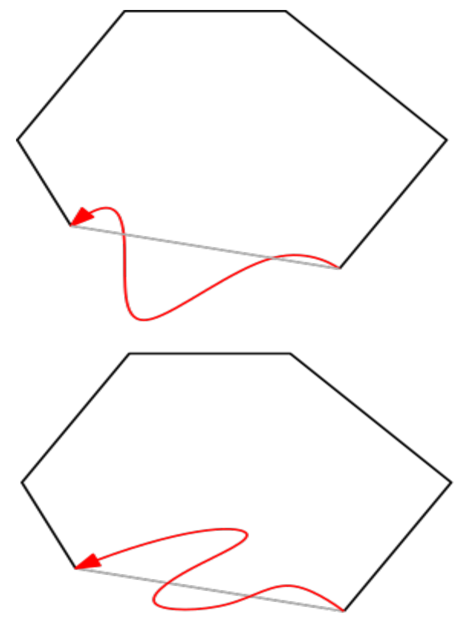

Figure 6: Here are two fragments with different portal edges. The gray lines are the portal line segments. The upper fragment has a monotone portal edge and the lower fragment has a non-monotone portal edge.

Just as before, a **portalgon** is a set of fragments; a fragment starts as a simple polygon, but parts of the boundary are substituted by portals. We still have that a portal e has two edges e^+ and e^- of the same length. Yet, we add the form of a portal edge e^+, which is described by its function, $g_{e^+} : [0, 1] \to e^+$ where $g_{e^+}(0)$ is the beginning vertex of the portal edge e^+ and $g_{e^+}(1)$ is the end vertex of the portal edge e^+. The function for e^- is set up analogously. We restrict this function such that it is injective and continuous. Furthermore, if the interior of the fragment containing e^+ is on the left of e^+ when going from the start vertex to the end vertex, then the interior of the fragment containing e^- is on the right of e^-.

Then, we still need a few definitions to be able to classify different portals. First of all, we want to be able to distinguish easily between portalgons with and without the newly allowed portals. A **straight portal edge** is a portal edge e^* where the image of the function $g_{e^*}(x)$ is the line segment between the start and end vertex of the portal edge.

A **straight portal** is a portal where both edges are straight edges. Then, a **straight portalgon** is a portalgon with only straight portals. A **curvy portalgon** is a portalgon with at least one portal that is not straight.

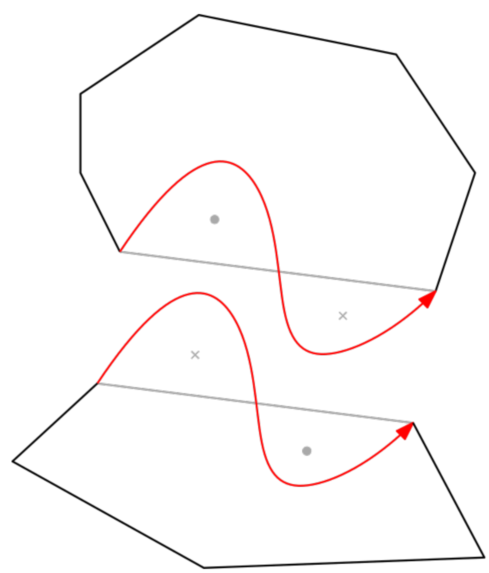

Figure 7: Here the interior territory is labelled with a dot and the exterior with a cross.

Next, with the newly allowed portals, we still classify further to a base case (monotone) and a more complex case (non monotone). The **portal line segment** l_{e^*} is the straight line segment between the start- and endpoint of the portal edge e^*. Let l'_{e^*} be a line in the plane of the fragment (belonging to this portal edge) perpendicular to the portal line segment. A **monotone portal edge** is a portal edge where every line l'_{e^*} through the portal line segment intersects the portal edge at most once. A **monotone portal** is a portal where both edges are monotone portal edges.

A monotone portalgon only has monotone portals. Refer to Figure 6 for an example of monotone and non-monotone portals.

Now, that we have the concept of the portal line segment, we can also define an obstacle.

Definition 3.0.1. Let e be a portal edge where the interior of its fragment lays to the left of it. **Obstacles** to e are those parts of the boundary of the fragment, which are not the portal edge itself, that cross the portal line segment. When obstacles cross the portal line segment from the left, they are interior obstacles. When obstacles cross the portal line segment from the right, they are exterior obstacles. When the interior lays to the right of e, we know without loss of generality that the definition is the other way around.

Refer to Figure 8 for an example of a portalgon with an interior obstacle. Refer to Figure 11 for an example of a portalgon with an exterior obstacle. In the next subsection we will discuss the specifics of dealing with obstacles.

Next, we have the definition of a fitting portal.

Definition 3.0.2. A **fitting portal** $e = (e^+, e^-)$ is a portal where the portal edge functions adhere to:

- let g_{e^+} and g_{e^-} be the portal edge functions of e. Then there is a rigid transformation A (transformation that does not affect the shape or size of the input) for which $\forall z \in [0,1], A(g_{e^+}(z)) = g_{e^-}(z)$.

In this chapter, we are going to proof that for a portalgon with n vertices, m straight portals and one fitting portal (with some restrictions) always has an equivalent straight portalgon with $O(n)$ vertices and $O(n)$ straight portals. We start proving this for a portalgon with a monotone fitting portal where the edges are on two different fragments, then continue to a portalgon with a monotone fitting portal where the edges are on the same fragment. Then, for a subset of non-monotone portals, we first show this for a portalgon with a non-monotone fitting portal out of this subset where the edges are on two different fragment, then continue to a portalgon with a non-monotone fitting portal out of this subset where the edges are on the same fragment. We will prove this by cutting of pieces of the fragments and “glueing” portals in such a way that we get an equivalent portalgon without the curvy portal.

3.1 Portalgons with a monotone fitting portal

For monotone fitting portals we need to define certain areas called territory. More specifically, we define the **interior** and **exterior territory** of a monotone portal edge. Let e^* be a monotone portal edge and let f be the fragment on which e^* is an edge. Let $S = \{s_1, s_2, ..., s_p\}$ be the intersections of the portal line segment l_{e^*} with the portal edge e^*, in order on the portal line segment with the first intersection being the starting vertex and the last intersection being the end vertex of the portal edge. These can be either a point or a line segment. For any line segment $l \in S$, substitute l in S for its end points. Let $g' = g_{e^*}|_{[s_i, s_{i+1}]}$ the restricted function for an $i \in [1, ..., p-1]$. The function g' either completely intersect with the portal line segment or it only intersects the portal line segment at the start and end points of g'. In the latter situation it is because in this case g' is on just one side of the portal edge; if it were not, it would lead to a contradiction: it would either have more intersections in $[s_i, s_{i+1}]$, or would intersect the rest of the boundary of the fragment, or the edge would not be monotone, or it would intersect completely with the portal line segment.

The former pieces are left out of consideration for the territory pieces, instead it is called a **straight territory segment** of the portal. The latter pieces are territory pieces. A territory piece is bounded by g' and $l_{e^*}|_{[s_i, s_{i+1}]}$. If the interior of the fragment lies on the left of e^*, then: an interior territory piece is a territory piece which lies left of the portal line segment. An exterior territory piece is a territory piece which lies right of the portal line segment. If the interior of the fragment lies on the right of e^*, then the definitions of interior and exterior are switched around. The collection of all interior territory pieces together is the interior territory of the portal edge. The collection of all exterior territory pieces together is the exterior territory of the portal edge. Refer to Figure 7 for an example of territories and to Figure 8 for an example of how territory pieces fit together. This lead us to the following lemma.

Lemma 3.1.1. *Let $e = (e^+, e^-)$ be a monotone fitting portal. Let $g_{e^+}(s), g_{e^+}(t)$ be two consecutive intersections of l_{e^+} with e^+ and let y be the corresponding territory piece. Assume without loss of generality that y is an exterior territory piece. Then, z is an interior territory piece bounded by $g_{e^-}([s,t])$ and l_{e^-}; and $z = A(y)$, where A the rigid transformation from Definition 3.0.2 for e.*

Proof. Both edges have the same edge function by definition. So, they have the same portal line segment and the same intersections. We have s^+, t^+ two consecutive intersections of l_{e^+} and e^+, thus s^-, t^- are two consecutive intersections of l_{e^-} and e^- where $g_{e^+}([s^+, t^+]) = g_{e^-}([s^-, t^-])$.

Without loss of generality, assume that the interior of the fragment containing e^+ is on the left of e^+, then we know the interior of the fragment containing e^- is on the right of e^-. First assume that y is an exterior territory piece of e^+, then it is on the right of l_{e^+}. It follows that the territory piece belonging to $l_{e^-}|_{[s^-,t^-]}$ is on the right of l_{e^-}. So, it is an interior territory edge of the fragment containing e^-.

As A maps the portal edges onto each other, the same transformation also maps the portal line segment to each other without changing its form. Therefore, the territory pieces can also be mapped with A to each other perfectly without changing it size or form. ☐

Now that we have gotten some basic concepts down that we will need, we are going to prove that any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone fitting portal, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals. We will do this by first proving this for the edges on two fragments without any obstacles, then we look at the situation with interior obstacles. We continue to the case with exterior obstacles too. Lastly, we are going to look at the situation where both edges are on the same fragment without any obstacle and then with obstacles. We will be doing this by cutting off pieces along the portal line segment and joining those to the other fragment in a certain order.

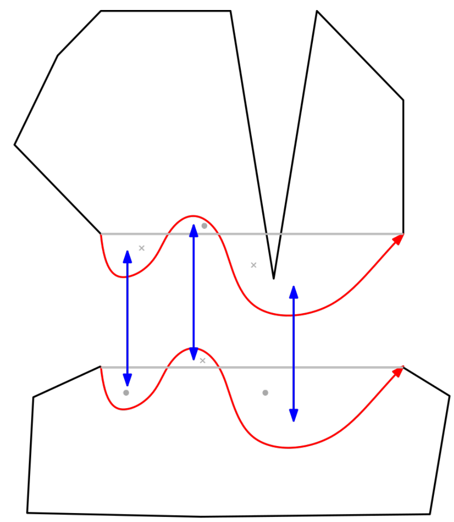

Figure 8: This is an example of a portalgon of two fragments with one monotone fitting portal and an interior obstacle. As visible, any exterior territory piece fits within the corresponding interior territory piece.

Lemma 3.1.2. *Let portalgon $\mathcal{P}$ have n vertices in total, m straight portals and one monotone non-straight fitting portal e, where the edges are on different fragments. Let the e be without any obstacles. Let y the first territory piece of e^+ and let x be the corresponding territory piece on e^- given by Lemma 3.1.1. Then the exterior territory piece between x and y can be cut off and "glued" to the interior piece in such a way that the resulting portalgon is equivalent to $\mathcal{P}$, but now has n vertices in total, m straight portals and one monotone fitting portal with less territory pieces than e.*

Proof. We are going to prove the lemma by first showing we can cut off the exterior territory piece and glue it to the interior territory piece. Then we are going to analyse the complexity after operation.

Assume without loss of generality that x is the exterior piece, hence, y is the corresponding interior piece. There are a couple of cases:

1. The first territory pieces start at the start vertex, but do not end at the end vertex.

2. The first territory pieces start at the start vertex and end at the end vertex.

3. The first territory pieces don't start at the start vertex, but end at the end vertex.

4. The first territory pieces don't start at the start vertex and don't end at the end vertex.

Note that for all cases, it follows that the cut-off piece x fits perfectly within y as we have no obstacles. Then, the cases are solved as follows:

1. We cut off x over the portal line segment, where the new portal follows the orientation of the start vertex to the next intersection (viewed from the start vertex) of the portal edge e^+ with the portal line segment. Note that e has now been cut into two pieces, one part on x and y and one part e' remaining on the fragments. Then, we merge x into y by Lemma 3.1.1. Now, we have the straight portal over the portal line segment of which its end vertex intersects with the starting vertex of e'. We merge these together into a new portal, which can be done as their orientation aligns. It is clear that this new portal is still a monotone fitting portal, as from the starting vertex to the first intersection between the portal line and the edge we have a straight part which is monotone fitting and the rest of the portal (e') still has the same form as before and therefore is also monotone fitting.
2. We cut off x over the portal line segment, where the new portal follows the orientation of the start vertex to the next intersection (viewed from the start vertex) of the e^+ with the portal line segment. Then, we merge x into y by Lemma 3.1.1. Now, we end up with a straight portal from the start vertex to the end. A straight portal is per definition monotone fitting.
3. We cut off x over the portal line segment, where the new portal follows the orientation of the first intersection s of the territory piece to the next intersection t (viewed from the start vertex) of the

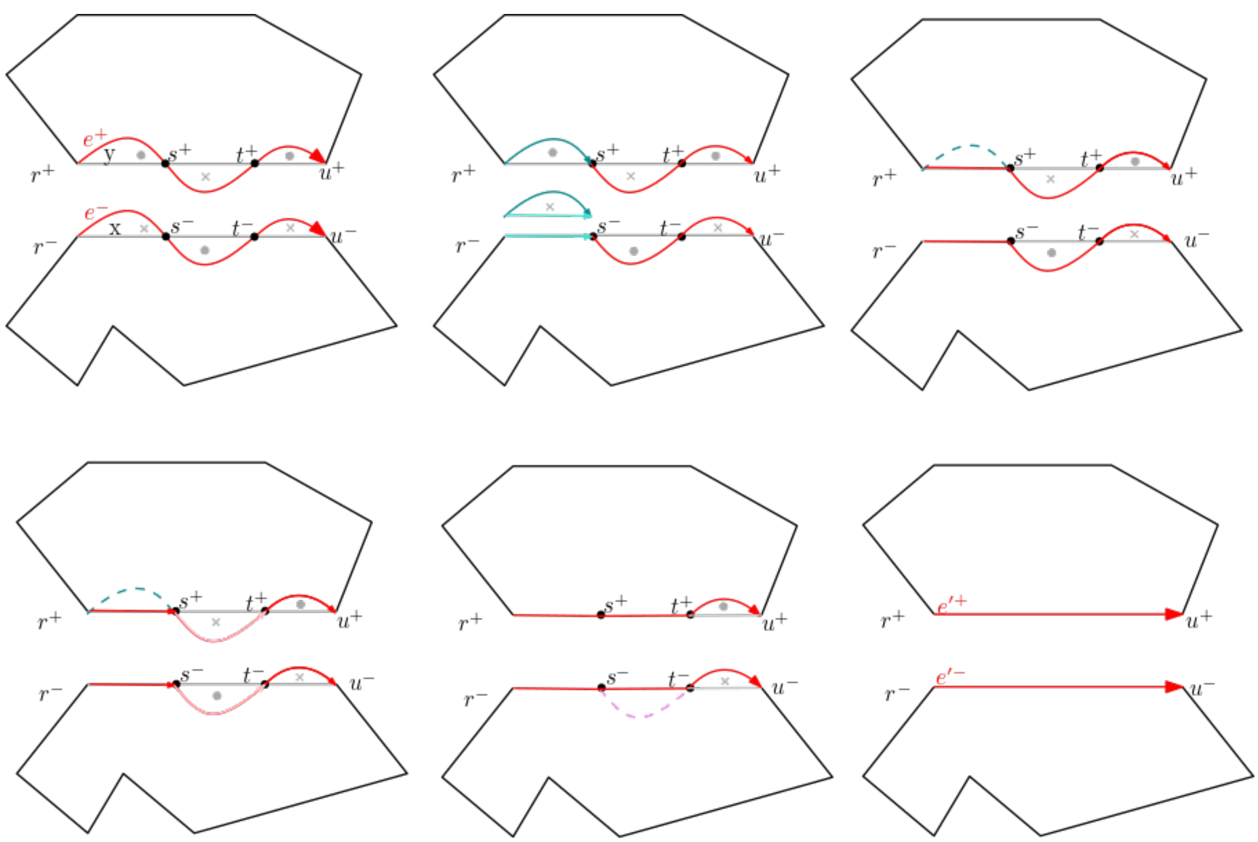

Figure 9: This figure shows the method of Lemma 3.1.2 and 3.1.4 in a simple example. The top left is a pair of fragments with a fitting portal. The portal line segment and territory pieces are marked with the established notation. We have the start vertex r and the end vertex u. Furthermore, we have two intersection between them, s and t, which split the portal into three pieces. In the following pair of fragments, we have cut off the exterior territory piece of e^+ and glued it to the other fragment, we then do the same for all other territory pieces in order from the r to u. In the fourth pair we end up with multiple portals that can be substituted for one portal as seen in the last pair.

e^+ with the portal line segment. Note that e has now been cut into three pieces, a straight piece e^* from the start vertex to s, one part on x and y and one part e' remaining on the fragment from t to the end vertex. Then, we merge x into y by Lemma 3.1.1. Now, we have the straight portal over the portal line segment of which its start vertex intersects with the end vertex of e^* and its end vertex intersects with the starting vertex of e'. We merge these together into a new portal, which can be done as their orientation aligns. It is clear that this new portal is still a monotone fitting portal, as from the starting vertex to the first intersection we have straight part which is monotone fitting and the rest of the portal (e' and e^*) still has the same form as before and therefore is also monotone fitting.

4. We cut off x over the portal line segment, where the new portal follows the orientation of the penultimate intersection to the end vertex. Note that e has now been cut into two pieces, a straight portal e' from the start vertex to the penultimate intersection and one part on x and y. Then, we merge x into y by Lemma 3.1.1. Now, we have the straight portal over the portal line segment of which its end vertex intersects with the starting vertex of e'. We merge these together into a new portal, which can be done as their orientation aligns. It is clear that this new portal is still a monotone fitting portal, as both pieces are straight portals which are by definition monotone fitting.

Note that in all cases we end up with less territory pieces as we merge one piece and the rest of the form of the portal stays the same.

Then for the complexity, it is clear that while we might have added vertices or portals as intermediate steps, we always end up with a monotone fitting portal in place of e without changing anything else on the portalgon. As the initial and final portalgon describe the same space just with portals on different spots, they are equivalent. This same logic holds for all future proofs. Thus, we end up with a portalgon that is equivalent to $\mathcal{P}$, and still has n vertices in total, m straight portals and one monotone fitting portal with less territory pieces than e. ☐

For an example on the above Lemma, refer to the first three columns of Figure 9. Note that after we merge the exterior territory piece x into the interior piece y, we merge the turquoise portal together with the remaining red portal to create our new portal with less territory pieces. This can be done as they intersect on the end vertex of the turquoise portal and the starting vertex of the shortened red portal. Also see that if the first territory piece does not start at the start vertex, we temporarily split the portal up into three pieces, solve for the territory piece and then merge the newly aligning portals.

Remark 3.1.3. Note that in Lemma 3.1.2 we specifically pick the first territory pieces from the start vertex of the portal edges. It is of course also possible to merge it in a different order without any impact to the result.

Next we will show that by repeated usage of Lemma 3.1.2, we can show that a portalgon with a monotone fitting portal without obstacles has an equivalent straight portalgon with bounded complexity.

Lemma 3.1.4. *Let $\mathcal{P}$ be a portalgon with n vertices in total, m straight portals and one monotone fitting portal e, where there are no obstacles in the interior or exterior territory of e and the edges are on different fragments, then there is an equivalent straight portalgon $\mathcal{P}'$ with n vertices in total and $m+1$ straight portals.*

Proof. We are going to use Lemma 3.1.2 to show that we can merge all territory pieces and that we then end up with a straight portal in the place of e. Then, we shortly analyse the complexity.

There are the following cases of e:

1. The portal e is straight. We immediately have an equivalent straight portalgon $\mathcal{P}'$ with n vertices in total and $m+1$ straight portals.

2. The portal e has one territory piece. By Lemma 3.1.2, there is an equivalent portalgon $\mathcal{Q}$ with n vertices in total, m straight portals and a monotone fitting portal with less territory pieces than e. Note that the only valid number of territory pieces smaller than 1 is 0. Therefore, the resulting portalgon is an equivalent straight portalgon $\mathcal{P}'$ with n vertices in total and $m+1$ straight portals.

3. The portal e has more than one territory piece. We will first show that a portal always has a finite number of territory pieces. Afterwards, we will prove by induction that we end up with a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with n vertices in total and $m+1$ straight portals.

   We will now argue that there are r is finite. Suppose the smallest portal part on a territory piece is from intersection s to t, then it takes up x distance of the portal. If d is the length of the portal, then d/x is the biggest number of territory pieces possible. This is finite as d is finite and $x \leq d$. Note that there always has to be a smallest portal part, as every portal on a territory piece takes up a certain length of the portal and you can simply that the minimum of the set of these lengths. Hence, r is finite.

   The base step is that we can merge the first territory pieces from the starting vertex, after which we end up with an equivalent portalgon $\mathcal{Q}$ with n vertices in total, m straight portals and a monotone fitting portal with less territory pieces than e. This is proven by Lemma 3.1.2.

   Then the induction step is that after merging the first k territory pieces of e where we now have the portalgon $\mathcal{R}$, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and a monotone fitting portal e' with less territory pieces than e, we can always merge the $k+1^{\text{th}}$ territory pieces p of e if they exist and end up with a portalgon $\mathcal{R}'$, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and a monotone fitting portal e^* with less territory pieces than e'. Let a and b be the intersections of p. As we have merged the first k piece we know that our intermediate portal e' has straight segment from the start vertex to a and so, in e' p are the first territory pieces from the start vertex of the portal. By Lemma 3.1.2, we can merge p and end up with the portalgon $\mathcal{R}'$ as described.

   As r is finite, we always have that at some point the induction step will mean merging the last territory pieces (so, if the $k+1^{\text{th}}$ territory pieces p of e do not exist). In this case, the k^{th} territory pieces were the last territory pieces. This means after merging the k^{th} territory pieces, we have ended up with a portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and one monotone fitting portal. But now the monotone fitting portal is a straight portal. Thus, we have a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with n vertices in total and $m+1$ straight portals.

So, in all cases, we have proven that there is a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with n vertices in total and $m+1$ straight portals.

□

Refer to Figure 12 for an example of Lemma 3.1.4 in action. Next, we are going to prove that any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone non-straight fitting portal, where there might be interior obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals.

Lemma 3.1.5. *Let $\mathcal{P}$ be a portalgon with n vertices in total, m straight portals and one monotone non-straight fitting portal e, where the edges are on different fragments and there might be interior obstacles, then there is an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices and straight portals.*

Proof. We start construct the portalgon $\mathcal{P}'$ the same way as the proof of Lemma 3.1.4. We will be cutting off the exterior territory pieces and merging them into their corresponding interior territory pieces in order of them appearing from the start vertex to the end vertex. We will show that we are still able to merge exterior territory pieces with obstacles into the corresponding interior territory pieces and then, that we after merging the territory pieces, we end up with a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with $O(n)$ vertices and straight portals.

We are still able to merge exterior territory pieces with interior obstacles into the corresponding interior territory pieces as interior obstacles only take area away and do not intersect with the portal edge. Then, for the total amount of new portals and vertices, we will prove by induction that after we have merged all territory pieces. Note beforehand that the amount of interior obstacles is bounded by n as any interior obstacle at least needs one vertex of the fragment.

The base case is that we can merge the first exterior territory piece x with y interior obstacles into its corresponding interior territory piece and that the resulting portalgon is equivalent to $\mathcal{P}$; has a maximum of $n - y$ interior obstacles; and has a maximum of $n + 4y$ vertices, a maximum of $m + 2y + y + 1$ straight portals and a maximum of one monotone non-straight fitting portal.

We now need $y+1$ portals to cut off the territory piece from the rest of the fragment, as the portal line segment is interrupted y times by the obstacles. The last portal however is merged with the remaining segment rs of e after the last intersection on the portal line segment of x, if there rs does not exist or is straight, we have no remaining non-straight portal and instead end up with an extra straight portal. For the $y+1$ portals, we need a maximum of $4y$ additional vertices, namely on the intersections of the portal line segment with the interior obstacles: each obstacle has two intersections, if they had more they would be multiple obstacle, and each portal edge have a duplicate of these intersections. We also need a maximum of $2y$ new portals, as we need a new portal for every portal we split by cutting over the intersections of the portal line segment with the interior obstacles. These new portals do not create additional vertices as those are the same as the intersections of the portal line segment with the interior obstacles, and we have already counted these. Let r be the last intersection of the portal line segment with the interior obstacles. After merging the exterior territory piece, we ended up with y straight portals between the start vertex of e and r and the segment from r to the last vertex of e with contains a maximum of $n-y$ interior obstacles. Refer to Figure 10 for an example.

Then the induction step is that when we have merged the first k territory pieces, which in total had y interior obstacles, we can now merge the next exterior territory piece x with z interior obstacles into its corresponding interior territory piece and that the resulting portalgon is equivalent to $\mathcal{P}$; has a maximum of $n-y-z$ interior obstacles; and has a maximum of $n+4y+4z$ vertices, a maximum of $m+2y+y+2z+z+1$ straight portals and a maximum of one monotone non-straight fitting portal.

We now need $z + 1$ portals to cut off the territory piece from the rest of the fragment, as the portal line segment is interrupted z times by the obstacles. The last portal however is merged with the remaining segment rs of e after the last intersection on the portal line segment of x, if there rs does not exist or is straight, we have no remaining non-straight portal and instead end up with an extra straight portal. For the $z + 1$ portals, we need a maximum of $4z$ additional vertices, namely on the intersections of the portal line segment with the interior obstacles: each obstacle has two intersections, if they had more they would be multiple obstacle, and each portal edge have a duplicate of these intersections. We also need a maximum of $2z$ new portals, as we need a new portal for every portal we split by cutting over the intersections of the portal line segment with the interior obstacles. These new portals do not create additional vertices as those

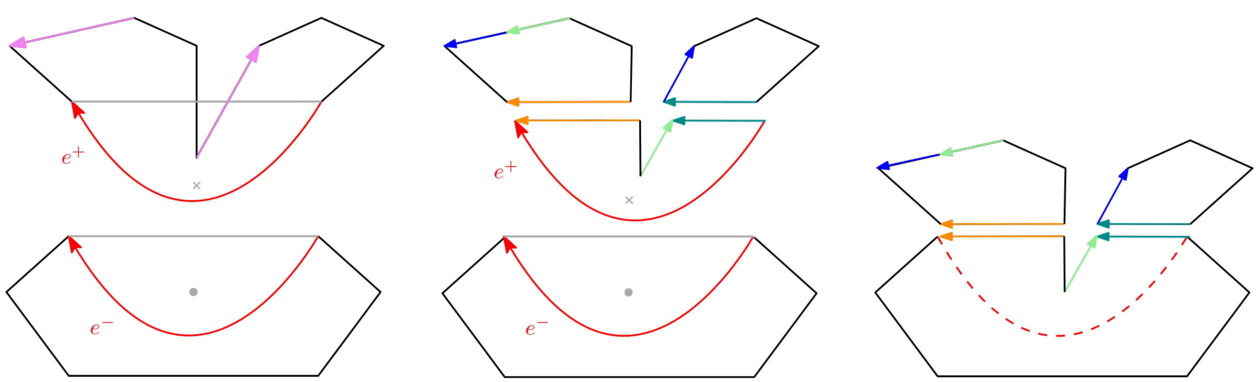

Figure 10: This Figure shows the theory of Theorem 3.1.6 in a simple example. The left is a pair of fragments with a fitting portal and a obstacle in the exterior territory piece. The portal line segment also intersects with a portal. We solve this situation by method of Theorem 3.1.6 as seen in the middle pair and the right pair of fragments.

are the same as the intersections of the portal line segment with the interior obstacles, and we have already counted these. Let r be the last intersection of the portal line segment with the interior obstacles. After merging the exterior territory piece, we ended up with a maximum of $n - y - z$ interior obstacles in the segment from r to the last vertex of e.

We show that we have finite amount of territory pieces. Suppose the smallest portal part on a territory piece is from intersection s to t, then it takes up x distance of the portal. If d is the length of the portal, then d/x is the biggest number of territory pieces possible. This is finite as d is finite and $x \leq d$. Note that there always has to be a smallest portal part, as every portal on a territory piece takes up a certain length of the portal and you can simply that the minimum of the set of these lengths.

Now, as we have a finite amount of exterior pieces, we know that at some point the final pieces are the k^{th} territory pieces. To we also have a final case to consider. In the final case, we have merged all r territory pieces and have ended up with a portalgon equivalent to $\mathcal{P}$; has no interior obstacles; and has a maximum of $n + 4n = 5n$ vertices, a maximum of $m + 3n + 1$ straight portals. This is exactly the straight portalgon $\mathcal{P}'$ we wanted.

□

Next, we are going to prove that any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone non-straight fitting portal, where there might be obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals. We are going to prove this by first cutting off the exterior obstacles and then we end up with a situation that can be solved with Lemma 3.1.5. Thus, we need to know how to cut off exterior obstacles.

When we cut off an exterior obstacle, we first add a new straight portal at every the intersection of the fragment and the line l parallel to the portal line segment with an offset away from the obstacle. Note the offset is chosen to be bigger than zero and smaller than the smallest distance between the vertices, which are not the vertices of the portal) and the portal line segment. If the obstacle is immediately attached to one of the end vertices of the portal, we instead cut from the end vertex to the furthest intersection between the portal edge and the offsetted line. If we do not do this, we risk not cutting off the obstacle as it would

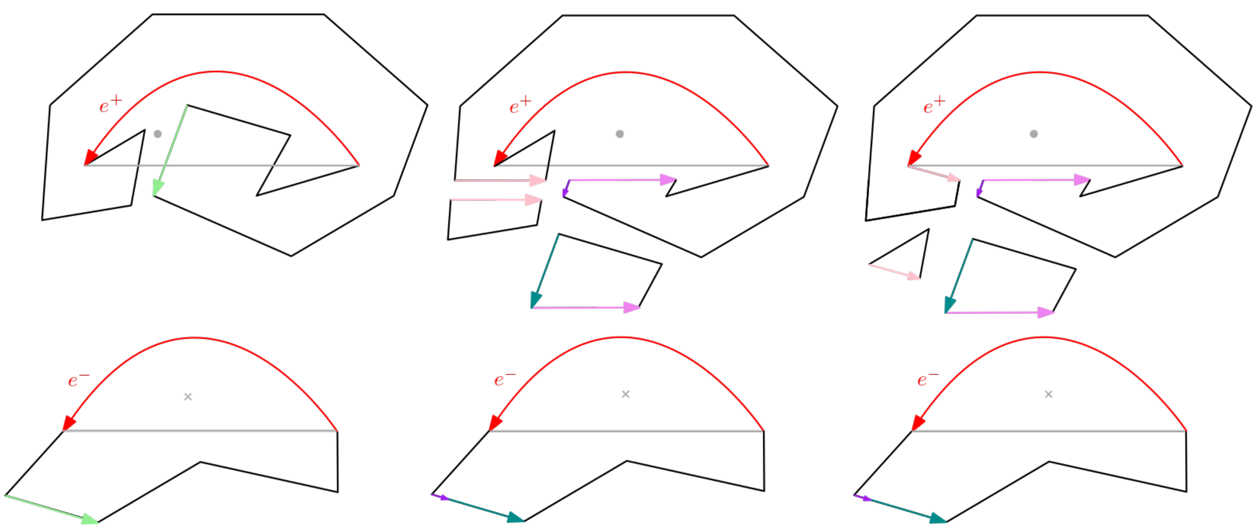

Figure 11: In the first column, there are two fragments that are connected by the red monotone fitting portal. The gray lines are the portal line segments. There are two obstacles, namely the left one connected to one of the vertices of the portal and the obstacle that is not directly connected to either vertex. In the second column, the obstacles are both cut off in the same way, it is clear that the left obstacle has not been removed by this operation. In the third column, we cut off the left obstacle in the augmented way as described.

still be attached to the rest of the fragment by a small strip. Refer to Figure 11 for an example.

We can also already say something about the complexity of this operation. As the maximum amount of exterior obstacles is bounded by the amount of vertices n belonging to the fragment, the amount of new portals created by cutting all obstacles off is also bounded by n. Note that any portals that we cut over will be split into two new portals. The number of new portals created by splitting original straight portals m on the obstacle is bounded by the amount of original straight portals on the fragment (as it is impossible to split more portals then there are on the fragment). So, after cutting off all exterior obstacles, we have at most $n + m$ new portals.

Note first that this does also work in the same way for non-monotone portals. Next, we will prove the theorem.

Theorem 3.1.6. *Let $\mathcal{P}$ be a portalgon with n vertices in total, m straight portals and one monotone non-straight fitting portal e, where the edges are on different fragments, there might be obstacles. There is an equivalent straight portalgon $\mathcal{P}'$ have $O(n)$ vertices in total and $O(n)$ straight portals.*

Proof. We start by constructing a portalgon $\mathcal{Q}$, equivalent to $\mathcal{P}$, that has no exterior obstacles. Then, we can use Lemma 3.1.5 to prove this Theorem.

We are going to construct $\mathcal{Q}$ by removing the exterior obstacles as described before and shown in Figure 11. The portal edges are on two different fragments f and g, denote their number of vertices by n_f and n_g. We will remove the obstacles of f, the results are the same for g. Then we will combine the results.

We will describe what happens when we cut off one obstacle and then what happens when all are cut off. Cutting one these obstacles off by the process of Figure 11 we add one new portal between the obstacle and the rest of the fragment. This portal has a start s and end vertex t on the obstacle and the remaining fragment. These vertices can already be an original vertex, but this is not guaranteed. Thus, we have for every obstacle we cut off a maximum of four new vertices. When we cut off an obstacle, there might be a maximum two portals we split; namely, those on the boundary where we enter the obstacle and on the boundary where we leave the obstacle. So, there are an additional two portals that are created by cutting off an obstacle. Notice that here no new vertices are introduced, since the new start and end vertices are the same as s and t. So, cutting off one obstacle results in less then five new vertices and less than four

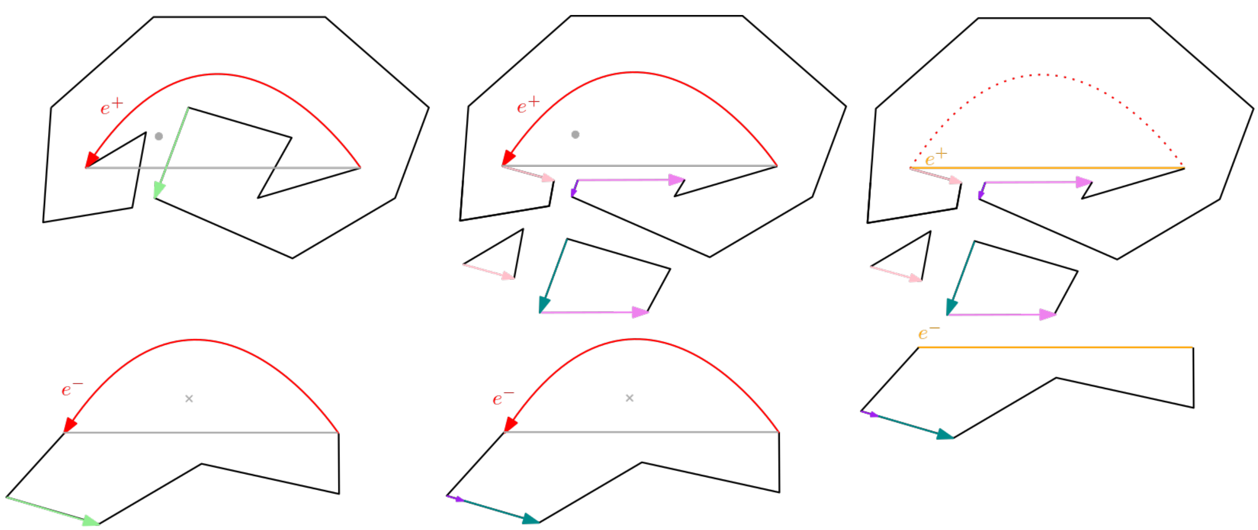

Figure 12: This figure shows the method of Theorem 3.1.6 in a simple example. The left is a pair of fragments with a fitting portal. The portal line segment and territory pieces are marked with the established notation. There is an obstacle that we cut off and then we merge the territory pieces following Lemma 3.1.4.

new portals. There are a maximum of n_f obstacle in f as any obstacle needs at least one vertex. Thus, after cutting off all obstacles, we have instead of fragment f an equivalent collection of fragments f' with less than or equal to $5n_f$ vertices and less than or equal to $4n_f$ straight portals. Note that as we have cut off all exterior obstacles on fragment f, the portal edge on f' is now free of any exterior obstacles.

We do the same for g. Therefore, we end up with a portalgon $\mathcal{Q}$, which is equivalent to $\mathcal{P}$. However, fragment f has been interchanged with f' with $O(n_f)$ vertices and straight portals, where there are no exterior obstacles in the portal edge of e contained within f'; and fragment g has been interchanged with g' with $O(n_g)$ vertices and straight portals, where there are no exterior obstacles in the portal edge of e contained within g'. So, a portalgon $\mathcal{Q}$ has no exterior obstacles in the portal edges of e and has $O(n_f + n_g) = O(n)$ vertices and portals.

Notice that we meet the requirements of Lemma 3.1.5. So, by Lemma 3.1.5, the portalgon $\mathcal{Q}$ has an equivalent straight portalgon $\mathcal{P}'$ where the vertices and the number of straight portals are $O(n)$. □

Refer Figure 12 for an example of the above theorem. Lastly, we will in this subsection look at the situation where both monotone edges are on the same fragment. We are going to do this in a few steps, first we are going define a new concept, then we will prove the situation without any obstacles. Penultimately, we prove it with interior obstacles. Ultimately, we prove it with any obstacles.

We start with defining portal interference, refer to Figure 13 for an example of portal interfering.

Definition 3.1.7. Let e^+ and e^- be portal edges of the same portal e on fragment f. We say that an interior territory piece x of e^+ is **portal interfering** with exterior territory piece y of e^- if x crosses the portal line segment of e^- from the interior of the fragment and there is a point p on the part of the portal e^+ on x and a point q on the part of the portal e^- on y that are visible to each other.

The latter requirement is simply because the portal edges can be solved as if they are on different fragments if they are not visible to each other, as that would mean there is an obstacle keeping them apart and so, they are not interfering.

Remark 3.1.8. Notice that if two portal edges are not interfering. Then any operation involving only the portal line segment and the part of the fragment on the exterior side (compared to the portal line segment) can be done as if they are on different fragments even if they are on the same fragment.

It is most interesting to talk about when the territory pieces belonging together are interfering with each other. In other words, that the same part of the portal edges is portal interfering.

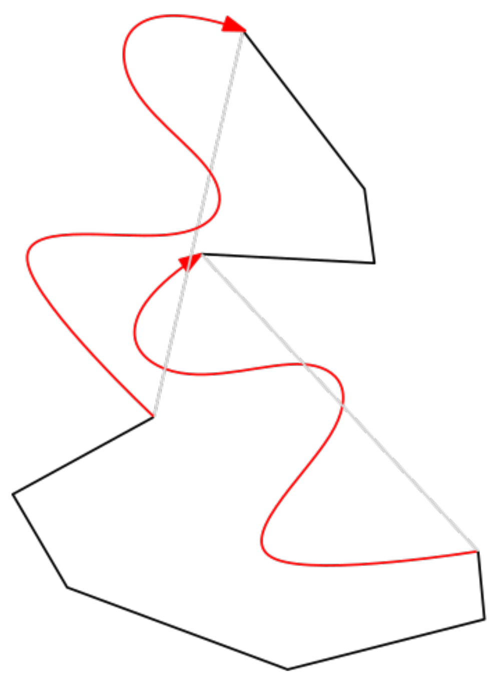

Figure 13: This is an example of where the portal edges are interfering as an interior territory piece of the lower edge interferes with an exterior territory piece of the other.

We then continue to the case of two edges on the same fragment without any obstacles. For some examples on how the Lemma works, refer to Figure 14.

Lemma 3.1.9. *Any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone non-straight fitting portal e, where the edges are on the same fragment and there are no obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with n vertices in total and $m+1$ straight portals.*

Proof. When the portal edges are not portal interfering, they can be solved as a normal monotone non-straight fitting portal. No cut over the portal line segment of e^- with have any impact on the territories of

e^+. So, after solving it with Lemma 3.1.4 we get a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with n vertices in total and $m+1$ straight portals.

When the portal edges are portal interfering, we are going to do the following. We are going to cut up the fragment over a line and then show that the newly created fragments can be merged in such a way that we end up with a portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with n vertices in total and $m+1$ straight portals. We are going to do this by induction.

Let l be the line segment connecting the start and end vertex of the portal in the space after "gluing" together portal. Portal interfering pieces are interior territory pieces, so if we enter there, we end up in exterior territory pieces, the line can only intersect a boundary earlier if there is an obstacle that is in the exterior territory piece or on the portal line segment. Yet, that is a contradiction with the fact that there are no obstacles. Furthermore, notice also that the visibility line the shortest length between two points, so l has the same length as the portal line segment. When we say the components of l, they are the components of l in $\mathcal{P}$. Though, we only consider those components of l that are not intersecting fully with a straight part of the portal, so we filter those out and are left with a set of components, of which the sum of its length is less than the length of the portal line segment.

Note that the amount of components r of the line l has is finite, as every component has a length. Suppose the smallest component is from intersection s to t of the portal line segment of one portal edge with the other portal edge, then it takes up x distance of the portal line segment. If d is the length of the portal line segment, then d/x is the biggest number of components possible. This is finite as d is finite and $x \leq d$. Note that there always has to be a smallest component, as every component takes up a certain length of the portal and you can simply that the minimum of the set of these lengths. Hence, r is finite.

Then, we also need to determine that there are not components of l perpendicular to the portal line segments. If this was the case, the portal line segments would be perpendicular to each other. If we have portal-interfering edges with perpendicular portal line segment, then because of the perpendicular portal line segments the only place they could be portal interfering are the pieces containing the start or end vertices of the portal edges. Any piece that is portal interfering at one of the ends of the portal in a monotone portal would partially be made up out of a part of the boundary of the fragment. So, it also introduces an interior obstacle, and we have determined that these do not exist in this situation.

Now we will cut over the line l selectively to get the desired portalgon. A section is a connected subset of the fragment bounded by e^+ and a component of l. The base case is that we can cut off the sections bounded by one component of l and part of the portal edge and merge them into the corresponding parts of the interior territory of the other portal edge and end up with a portalgon, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and one monotone non-straight fitting portal e.

We will show that the base case works for every section x bounded by one component of l and a part of the portal edge. A section x is the furthest away from the portal and are without interference (as if there was, it would not be just one component of l). As these sections are parts of the exterior territory pieces of the portal edges and that there are no obstacles, we know that we cut x off over the line component in the same orientation as the portal edges and merge them into the corresponding portal part of the other edge. Then we merge the new straight portals with the two monotone portals, which is possible as they have the same orientation and extension of each other (end/begin vertices are intersecting).

It is clear that the portal edges are still fitting as we have changed the form the same way on both edges. It is also still monotone as there is no component of l perpendicular to the portal line segments, which means every component is monotone to the portal line segments and the curved part on the section that has been substituted for a straight part is still monotone, as well as that the rest has not changed. So, the augmented e is still monotone. It is also still not straight as if it would be there would not have been any interference. Note there are no new components of l, in fact if we were to setup l again, then it would intersect the portal over this line segment. So, this part would be filtered out. We end up with a portalgon, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and one monotone non-straight fitting portal e.

The induction step is that after we have cut off the sections bounded by the components furthest from the portal for k times, where we ended up with a portalgon, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and one monotone non-straight fitting portal e, then we can cut off the sections bounded by the component(s) of l furthest away from the portal line segments and parts of the portal edge and merge them into the corresponding parts of the interior territory of the other portal edge and end up with a portalgon, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and one monotone non-straight fitting portal e.

The step is either the same as the base case if the section is bounded by just one component of l and one part of e. In the situation that the section x is bounded by more components of l and/or more components of e, we will do the following. We cut the piece off along the the components of l. We know we can merge the non-straight portal parts on x to its corresponding part in the corresponding interior territory component. Then we merge the new straight portals with the two monotone portals, which is possible as they have the same orientation and extension of each other (end/begin vertices are intersecting). This reduces the amount of portals we end up with significantly as we do not count all these straight parts of the final portal separately.

It is clear that the portal edges are still fitting as we have changed the form the same way on both edges. It is also still monotone as there is no component of l perpendicular to the portal line segments, which means every component is monotone to the portal line segments and the curved part on the section that has been substituted for a straight part is still monotone, as well as that the rest has not changed. So, the augmented e is still monotone. It is does not have to not straight as there could still be components of l, though it could also be straight if we have merged the last pieces (for what to do after merging the last piece, refer to the next paragraph). Note there are no new component of l, in fact if we were to setup l again, then it would intersect the portal over this line segment. So, this part would be filtered out. We end up with a portalgon, equivalent to $\mathcal{P}$, with n vertices in total, m straight portals and one monotone fitting portal e with straight parts.

As there is a finite amount of components of l, there is also a limited amount of times we can do the steps of the induction step. The moment the new furthest components of l are on the portal line segment, it is the last time we can do the induction step, because afterwards there are no components of l anymore. Note that now, e intersects fully with the portal line segment, since if it did not, we would still have components of l. So, e has been substituted by a straight portal along the portal line segment with the orientation as the portal line segment.

Thus, we get a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with n vertices in total and $m+1$ straight portals. Thus, we have proven the theorem, any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone non-straight fitting portal e, where the edges are on the same fragment and there are no obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals. □

Remark 3.1.10. If there are two points on the normal non-portal boundary of the fragment that are visible to each other, where the line in between them split up the fragment such that both new fragments contain one portal edge, then adding a portal over that line immediately reduces it to a problem of two portal edges on their own fragment.

Lemma 3.1.11. *Any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone non-straight fitting portal, where the edges are on the same fragment and there might be interior obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total, $O(n)$ straight portals.*

Proof. We are first going to distinguish different cases, then we are going to solve them on by one. In every of these cases we also analyse the complexity. Then, the cases are:

1. The portal edges e^+ and e^- are not portal interfering at any part of the portal edges.

2. The portal edges e^+ and e^- are portal interfering.

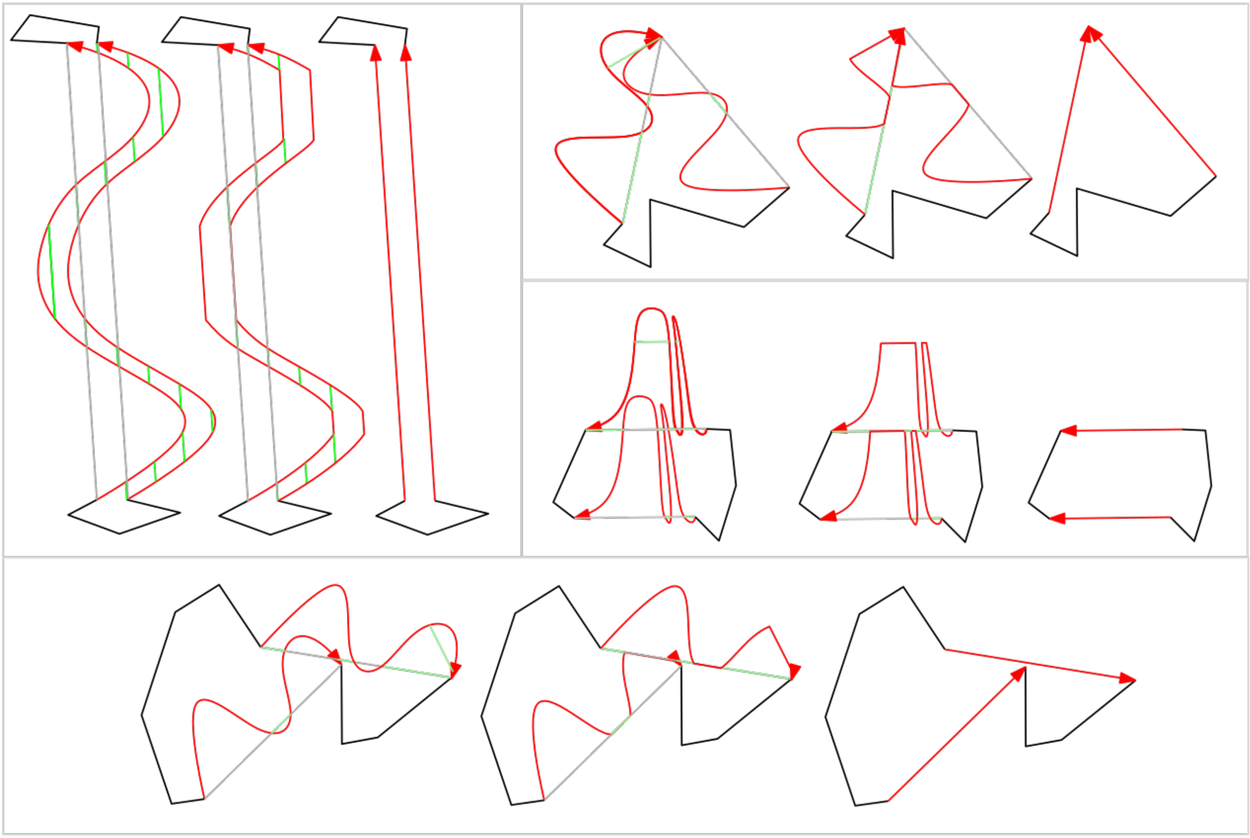

Figure 14: These are all example of how Lemma 3.1.9 is applied in various cases. Here the green line is the line l, the gray line is the portal line segment and the red line is the non-monotone fitting portal.

Now we are going to solve the cases one by one.

1. When the portal edges are not portal interfering, they can be solve as a normal monotone non-straight fitting portal. No cut over the portal line segment of e^- with have any impact on the territories of e^+. So, after solving it like normal we get a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with $O(n)$ vertices in total and $O(n)$ straight portals.

2. We are going to distinguish a few more cases in this case. Namely:

   (a) The portal edges which are parallel to each other. For an example, refer to the upper left example of Figure 14. Note that the middle right example in the figure is not parallel.

   (b) The portal edges which are attach at a one point. For an example, refer to the upper right example of Figure 14.

   (c) Any other portal edges.

   We construct the line l constructed in a same way to l in Lemma 3.1.9, but now every time we encounter an interior obstacle, we start up the next component of l along the portal line segment at the next intersection of the interior obstacle and the portal line segment. It now has different types of components as it is also obstructed by interior obstacles now. There are therefore also a few different types of series of components:

Type 1) A series of components from either the start vertex or the end vertex to an obstacle.

Type 2) A series of components between obstacles.

Type 3) A series of components from the start vertex to the end vertex.

Now we will solve each of these cases:

(a) We will first remove the rest of the fragment, so that these cannot intersect when we start "gluing" portals, then we can use Lemma 3.1.9.

When the portal edges which are parallel to each other, we will first add portals along the line between the start vertices and add portals along the line between the end vertices. As obstacles and portals on obstacles are bounded by $3n$, the amount of new portals is too. The amount of new vertices for this is bounded by $4n$.

Now there are interior obstacles coming from the line between the start vertices, we call those obstacles start obstacles, and similarly, we call the interior obstacles from the line between the end vertices end obstacles. We are going to filter out additional pieces of l, namely if we keep the right components of l, we can first solve as much as we can following the method of Lemma 3.1.9 and then as a last step we can cut the fragment into a maximum of three fragments where the remaining curved parts of e are on different fragments that can be merged immediately.

In the case of a series of components of Type 3, we default to the proof of Lemma 3.1.9 as there are no interior obstacles in this setting. When there are no obstacles from just one side we want the series of components of Type 1. When there are obstacles from both sides, we want a series of components of Type 2. Note that any of these series do not need the rest of the component as the obstacles themselves split up the rest of corresponding parts of the portal onto different fragments.

Now we can simply use the same proof as Lemma 3.1.9, though, had we not first removed the rest of the fragment, we could now have intersecting parts, as the parts with an obstacle edge are not necessarily bounded in the exterior territory. Yet, there is no rest of the fragment to interfere with, so we are safe.

(b) When the portal edges are attach at a one point, assume that this is their start vertex, we follow the proof of case 1, but with the following differences. We cut off the rest of the fragment not over the line between the end vertices, we cut it off over the line segments between the end vertices and the intersection of the lines perpendicular to the portal line segment thought the end vertices. As obstacles and portals on obstacles are bounded by $3n$, the amount of new portals is too. The amount of new vertices for this is bounded by $4n + 1$.

In the case of a series of components of Type 3, we default to the proof of Lemma 3.1.9 as there are no interior obstacles in this setting. Else, when there are no obstacles from just one side we want the series of components of Type 1. There is no Type 2 in this case. The rest of the proof is the same.

(c) We we have any other portal edges, we have to cut off the rest of the portal differently than in the parallel case, namely, after we have already used Lemma 3.1.9 and have multiple fragments. Now for the two outer fragments, we cut off all extra territory created by the obstacles over the portal line segments of that fragment. As obstacles and portals on obstacles are bounded by $3n$, the amount of new portals is too. The amount of new vertices for this is bounded by $4n$.

Thus, we have proven the theorem, any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone non-straight fitting portal, where the edges are on the same fragment and there might be interior obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total, $O(n)$ straight portals. □

Theorem 3.1.12. *Any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one monotone non-straight fitting portal, where the edges are on the same fragment and there might be obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals.*

Proof. We will first cut off all exterior obstacles and then we can use Lemma 3.1.11 to solve the remaining situation.

As determined before, in the proof of Theorem 3.1.6 we know that cutting off all obstacles give at most an additional $5n$ vertices and at most an additional $4n$ straight portals. So, after this operation, we are left with the portalgon portalgon $\mathcal{Q}$, equivalent to $\mathcal{P}$, with at most $6n$ vertices in total and $4n + m$ straight portals and one monotone non-straight fitting portal, where the edges are on the same fragment and there might only be interior obstacles.

Now $\mathcal{Q}$ meets the requirements of Lemma 3.1.11. Thus, we know we have a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with $O(n)$ vertices in total and $O(n)$ straight portals. □

3.2 Portalgons with a non-monotone fitting portal

In this section, we will prove the same situations as in the section above for a subset of the non-monotone portals. We start again with a few definitions. First, for some of the proofs we need a dual graph $G(f)$ for portal edge f. We set-up the dual graph in the following way. The portal line segment of e splits up the fragment in pieces. For every piece, we add a vertex to $G(f)$ and for every newly added portal (after splitting up the fragment into pieces over the portal line segment), we add an edge between the vertices belonging to the fragments containing the portal edges. When the start or end vertex ends up in a fragment on the exterior side of the portal line segment we add an extra node to the dual graph representing that node, this is then also added to $R_{G(f)}$. Refer to Figure 15 for an example of a simple dual graph.

Lemma 3.2.1. *Let f be a portal edge. Dual graph $G(f)$ is a tree.*

Proof. We will prove this Lemma by contradiction. Suppose $G(f)$ is not a tree. Then $G(f)$ either has a circuit or is not connected. The former case would imply that f has a hole, this contradicts with the definition that fragments should be simple polygons and that the portal edges should be continuous. The latter case would imply that the f before splitting up is not one fragment, that is also a clear contradiction. So, $G(f)$ is a tree. □

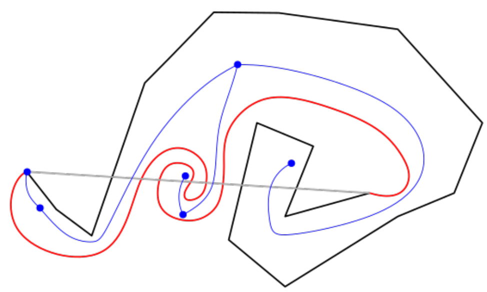

Figure 15: This is an example of a simple dual graph of this fragment. The dual graph is denoted in blue.

Then there is also the set $R_{G(f)}$ which is the path between the fragments containing the start and end vertex of the portal.

Remark 3.2.2. In a portalgon with a monotone portal this would result in a dual graph where the $R_{G(f)}$ has only branches of size one.

As non-monotone portals can have more forms than monotone portals, we need a way to differentiate then based on a characteristic. We have chosen interference, in layman's terms, this is the way the portal interacts with the portal line segment. The simplest case is that it interacts in the same way as in the monotone case and then we have non-interference. Then we also have a kind of interference that starts out a territory pieces of a monotone portal and then is transformed such that it intersect with the portal line segment multiple times. The first is self interference, in this case it happens without any piece being encapsulated by an obstacle.

Definition 3.2.3. There is **self interference** on portal edge f when

- $R_{G(f)}$ has a branch of size two or more, where the first intersection x of the branch and the portal line segment is from the interior to exterior, the branch after x is bounded by the portal line segment and the portal and (a part of) the branch does only crosses the portal line segment from the interior side after crossing it from the exterior side (situation B in Figure 16). For ease, we call this type 1 self interference.; or

- $R_{G(f)}$ crosses the portal line segment it does not necessarily have to mean self-interference. There are two cases where it does mean self-interference. The first is when the $R_{G(f)}$ has a branch of size one or bigger while in the exterior compared to the portal line segment, where the branch crosses the portal line segment from the exterior side, we have a case A or C of Figure 16. The second is when we have case D of Figure 16. Now $R_{G(f)}$ crosses the portal line segment multiple times and goes through the pieces actually interfering.

  Then, there is also the possibility of a combination of these cases, but in both scenarios the whole piece of the tree is self-interfering (in Figure 16 in situation A and D from the first intersection (denoted as a blue cross) to the last in within that situation and in situation C, from the start or end vertex to the first intersection). For ease, we call this type-2-self-interference.

When there is self interference, the order of intersections of the portal edge with the portal line segment are not the same if you compare their order on the portal edge to the order on the portal line segment. Refer to Figure 18 for examples.

Remark 3.2.4. It is clear that if one edge is self-interfering, it necessarily creates an empty space on the exterior side of the portal line segment (as it needs to cross the portal line segment again and for this it first needs to get of the portal line segment). This empty space is also necessarily connected to an empty space in the between the portal edge and the portal line segment (in terms of the monotone portal, an interior territory piece). So, this creates a self-interfering branch on the other fragment in this space. Refer to Figure 19 for an example. So, if a self-interfering branches are solve in respect to on of the edges, it is clear that that point the portal as a whole does not have any self-interfering parts anymore.

Then there is the last of the interferences where a branch is encapsulated by an obstacle.

Definition 3.2.5. Let f be a portal edge. Let a be an interior obstacle and s and t its intersections with

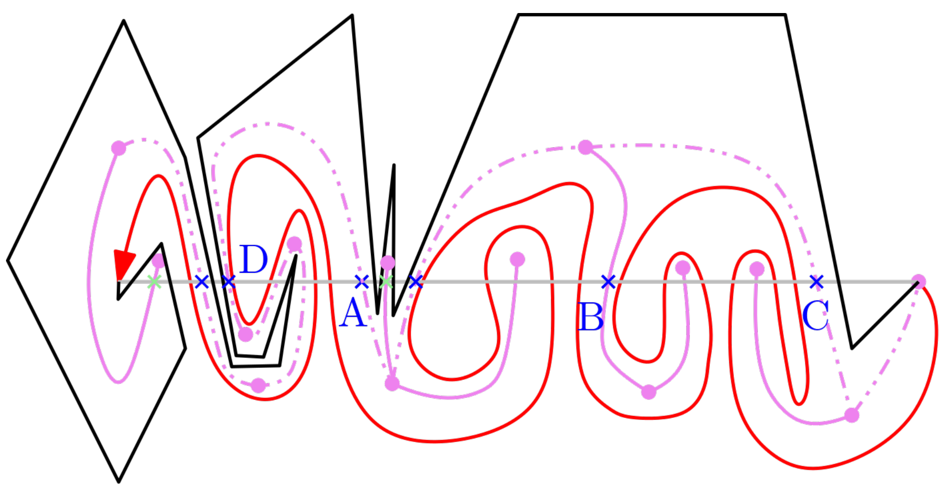

Figure 16: This is a dual tree of a fragment containing a self-interfering portal edge. The portal line segment as normal is denoted by a gray line and the dashed line in the dual graph is $R_{G(f)}$. There are a few things to consider in this new situation. Situation A is the situation of a self-interfering branch two interior obstacles. As the path now goes though the exterior, the self-interfering branch only has to have size one or more. Note that $R_{G(f)}$ now also goes though the exterior (as seen from the portal line segment), but not through the actual self-interfering piece of the branch. In situation B, there is a self interfering branch without any obstacles. In situation C, we have an interior obstacle that immediately join with one of the end vertices of the portal edge. Note that as we have described we have added a new vertex here. In situation D, the obstacle crosses the portal line segment multiple times as it follows the form of the self interfering portal. Note that $R_{G(f)}$ now also goes though the exterior (as seen from the portal line segment) and the self interfering pieces. We will need to keep the green crosses in mind when we are developing the algorithm.

the portal line segment. Let r be the intersection between portal edge and the portal line segment that lay in between s and t that is the closest to s and q the intersections between portal edge and the portal line segment that lay in between s and t that is the closest to t. **Outside interference** is when q and r exists and there is (part of) the interior of the fragment is between r and q and not between q and t (when ignoring other obstacles). In the dual graph of the edge, this is a branch the same as a self-interfering branch, but instead of intersection the portal line segment from the interior side, it intersect the portal line segment from the exterior side.

The exact opposite of outside interference. We have inside interference.

Definition 3.2.6. Let f be a portal edge. Let a be an exterior obstacle and s and t its intersections with the portal line segment. Let r be the intersection between portal edge and the portal line segment that lay in between s and t that is the closest to s and q the intersections between portal edge and the portal line segment that lay in between s and t that is the closest to t. **Inside interference** is when q and r exists and there is the interior of the fragment is between s and r and between q and t (when ignoring other obstacles). In the dual graph of the edge, this is resembled by $R_{G(f)}$ going through a.

When there is inside or outside interference, it is possible that the order of **all** intersection of the portal edge with the portal line segment are not in the same, but this is not a given. Though, when there is inside or outside interference without self-interference, the intersections of the portal edge with the portal line segment **within the obstacle** are in the same order on both the portal edge and the portal line segment. Refer to Figure 18 for examples.

There is an important connection between inside and outside interfere as they are exact opposites. Suppose we have a portal e where there is a branch that is interfering with the portal from the outside. We know that the interior of the fragment containing e^- is on the opposite side of the portal edge compared to the interior of the fragment containing e^+. We also know that for both portal edge have the same form and that an interfering part of the portal needs to be encapsulated as if it is not, then there are parts of the fragment with the interior on both sides of the boundary. As such it is clear that inside interference of the portal on one portal edge means the outside interference on the other portal edge on the part of the portal. This is summarised in the following corollary and there is an example in Figure 17.

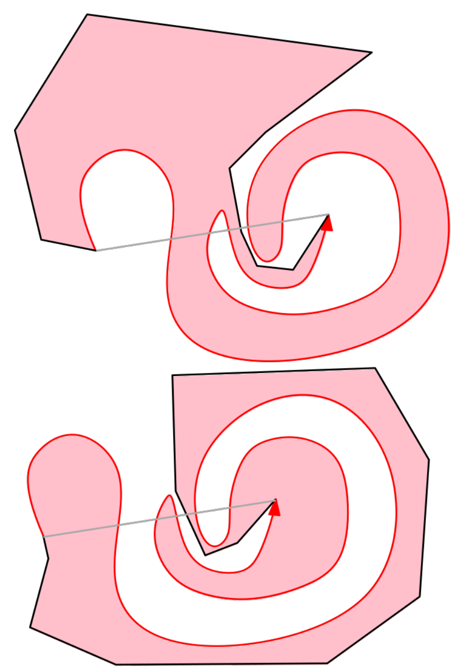

Figure 17: This is an example where the first fragment has outside interference and it is clear that the other fragment has inside interference on exactly the same place.

Corollary 3.2.7. Let e be a portal that is not interfering with itself. The fragment f containing the portal edge e^+ has inside interference if and only if the fragment g containing the portal edge e^+ has outside interference.

Furthermore, mixed interference is when a portal has multiple kinds of interference. Such as fragment G in Figure 18. Lastly, a non-monotone non-interfering portal is a non-monotone portal without any kind of interference.

Remark 3.2.8. Note that when we say that there is in inside interference or outside interference, we have to specify on which edge. So, when we say there is only inside interference, we mean there is only inside interference on either of the edges and only outside interference on the other edge. If an edge has both inside and outside interference, it is mixed interference.

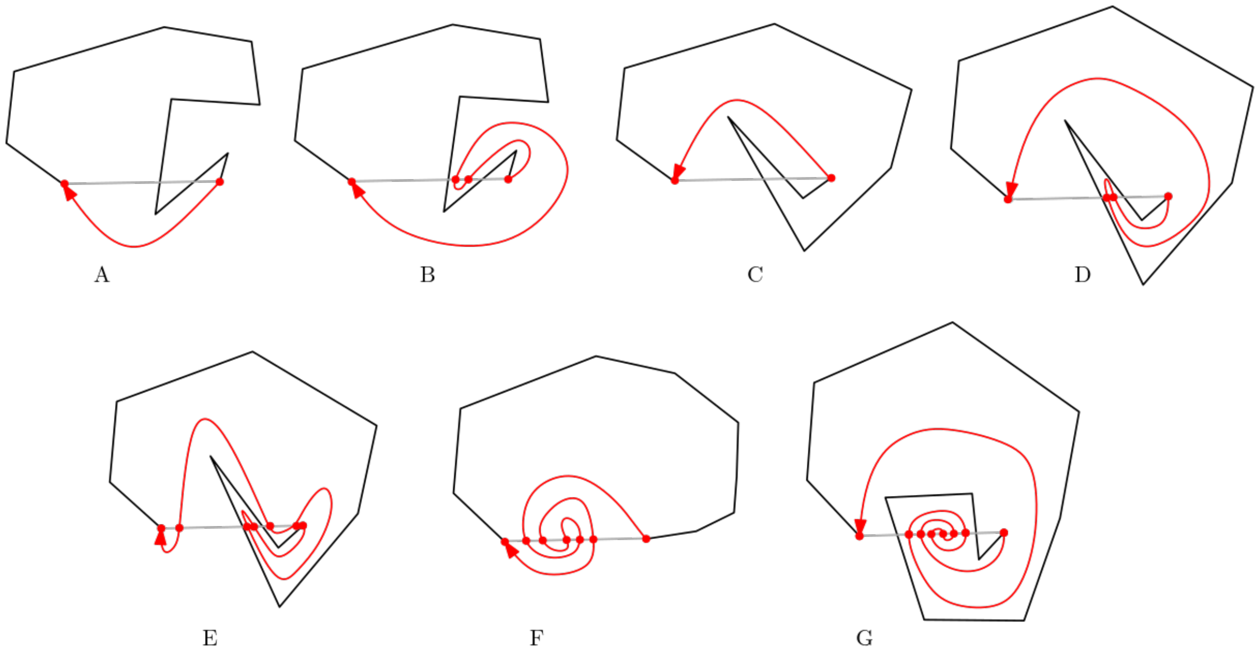

Figure 18: In all fragments, the portal edge is drawn in red, the portal line segment in gray and the intersections of the portal edge and the portal line segment in red dots. In fragment A, there is a fitting portal with an interior obstacle. The order of intersections of the portal line segment and the portal edge are the same and there is no interference of the portal edge with itself. In fragment B, there is a fitting portal with an interior obstacle. The order of intersections of the portal line segment and the portal edge are the same and there is outside interference of the portal edge with itself. Though, the order of the intersections on the portal edge within the obstacle appear in the same order as they appear on the portal line segment. In fragment C, there is a fitting portal with an exterior obstacle. The order of intersections of the portal line segment and the portal edge are the same and there is no interference of the portal edge with itself. In fragment D, there is a fitting portal with an exterior obstacle. The order of intersections of the portal line segment and the portal edge are the same and there is inside interference of the portal edge with itself. Though, the order of the intersections on the portal edge within the obstacle appear in the same order as they appear on the portal line segment. In fragment E, there is a fitting portal with an exterior obstacle. The order of intersections of the portal line segment and the portal edge are not the same and there is inside interference of the portal edge with itself. Though, the order of the intersections on the portal edge within the obstacle appear in the same order as they appear on the portal line segment. In fragment F, there is a fitting portal with self interference. The order of intersections of the portal line segment and the portal edge are not the same and there is no outside or inside interference. In fragment G, there is a fitting portal with self interference in the inside inference. The order of intersections of the portal line segment and the portal edge are not the same. There are many other options, such as self interference in outside interference.

In the rest of this section, we are going to prove that any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one non-monotone fitting portal, that is either self-interfering or not-interfering, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals. We will first deal with the case that it is a non-interfering portal and then with the case that it is self-interfering.

In the case of a non-interfering portal where the edges are on different fragments, we do not have any new situations compared to the monotone case, since the edge does not interact differently with its portal line segment. As the definition of territory pieces is still valid in exactly the same way as for monotone fitting portals, since the portal edges (thus also the resulting territory pieces) do not interfere with each other. So, the results can also be applied to this situation. This is summarise in the following corollary.

Corollary 3.2.9. Let portalgon $\mathcal{P}$ have n vertices in total, m straight portals and one non-monotone not-interfering fitting portal, where there are no obstacles in the fragments of the portalgon which contain the portal edges and the edges are on different fragments, then there is an equivalent straight portalgon $\mathcal{P}'$ with n vertices in total and $m + 1$ straight portals.

Corollary 3.2.10. Let portalgon $\mathcal{P}$ have n vertices in total, m straight portals and one non-monotone not-interfering fitting portal, where there might be obstacles in the fragments of the portalgon which contain the portal edges and the edges are on different fragments, then there is an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals.

But when the edges are on the same fragment, it can interact differently with the other's portal line segment. It can either have a self-interfering form on the other's portal line segment or can cross the portal line segment via an obstacle (so it an inside/outside interfering from on the other portal line segment). We are going to show that for these cases we can still transform the portalgon into an equivalent straight portalgon with $O(n)$ vertices in total and $O(n)$ straight portals.

We have two new situations to consider compared to the monotone portals. Namely, if one portal edge is "self-interfering" w.r.t. the portal line of the other portal edge and if one portal edge is "inside/outside-interfering" w.r.t. the portal line of the other portal edge. In the former case, all the pieces of the self interfering part are cut-off in the same loop as their components of the line l are all the same distance from the portal line. The only problem that can come up here is when the base of a branch gets merged before a leave of a branch, such that there is a hole in our fragment for a moment. Though, we can either argue that after merging the leaves too, the hole is gone, so the step into the forbidden realm is no problem, or we can enforce the order of first merging leaves and progressively merging towards the base of a branch.

The latter only happens when there are obstacles w.r.t. the second portal edge. When we add obstacles into the mix, interior obstacle can become outside interference from one portal edge to the other and an exterior obstacle can become inside interference. As any interior obstacle already forms a natural empty boundary area between the interfering portal part and the other portal, it is clear that the two exterior territory pieces that are interfering will be solve independently and are not any additional problem, compared to the monotone case.

Any interior obstacle can become a problem as they could block the order of operations, as we cannot cut off an obstacle with a piece of the portal inside it. So, we first fill up these parts with an offset to the interior side of the portal line it is intersecting with, we also fill up any holes this offset creates. In these parts there is a maximum of n interior obstacles, so we have at most an additional $4n$ vertices and $3n$ portals from this precomputation operation. After this, the rest of the computation is the same as before. So we get the following corollary.

Corollary 3.2.11. Any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one non-monotone no-interfering fitting portal, where the edges are on the same fragment and there might be obstacles, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals.

Now, we continue to the case of a portalgon with one self-interfering portal, where the edges are on different fragments and there are no obstacles.

Lemma 3.2.12. *Let portalgon $\mathcal{P}$ have n vertices in total, m straight portals and one non-monotone self-interfering fitting portal e, where there are no obstacles in the fragments of the portalgon which contain the portal edges and the edges are on different fragments, then there is an equivalent straight portalgon $\mathcal{P}'$ with n vertices in total and $m + 1$ straight portals.*

Proof. We are going to prove this Lemma by solving the Type 1 self interfering (as we do not have Type 2 without any interior obstacles) parts of the fragment, then we show we have ended up with a non-monotone non-interfering or monotone portal instead of e. Lastly, we will use Corollary 3.2.10. Note that by Remark 3.2.4 we only need to solve the self-interference on one of the edges.

Let f^+ and f^- be the fragments containing the portal edges e^+ and e^- respectively. Let $G(e^+)$ be a dual graph as described earlier in this chapter. We will start by showing we have finite amount of self-interfering branches by showing that we have a finite amount of branches. Suppose the smallest branch takes up x distance of the portal. If d is the length of the portal, then d/x is the biggest amount of branches possible. This is finite as d is finite and $x \leq d$. Note that there always has to be a smallest branch, as every branch takes up a certain amount of the portal and you can simply that the minimum of the set of these amounts.

Suppose we have a branch b on portal edge e^+ in the dual graph that we cut off from $R_{G(e^+)}$. We do this by adding a portal x between $R_{G(e^+)}$ and the branch along the part of the portal line segment that connects b to the rest of the fragment. We make x with the same orientation of the portal line segment. Let s and t be the start and end vertex of this new portal.

The branch b is now made up out of two portal edges, a straight portal edge between s and t with the same orientation of e between s and t and a non-straight portal edge s and t that follows the edge function of e between s and t with the same orientation as e. There cannot be any obstacles in the space bordered by the portal edge between s and t and the portal line segment on the fragment containing e^-, since we have no obstacles and any obstacle further down the branch would also be an obstacle earlier (see Figure 19 for an example). So, we can merge our branch into the other fragment without any problems.

Then we end up with the new portals on both fragments, where the portal from start point to s, from s to t and from t to the end point are all fitting portals. The portal made up out of these pieces is also a fitting portal. Note that we can join these pieces as they do not intersect and their orientations align. There will no be new branches added after merging a branch, as the part of the portal where the branch was is now a straight edge and the rest of the form of the portal remains untouched. In short, after cutting off and merging a branch, we still have just one fitting portal without adding new portals or vertices and we will have less branches.

After merging the last branch, we have no self-interfering branches left on e^+, as we also have no obstacles or inside/outside interfering, we have only branches of size one. Thus, we are left with a non-monotone non-interfering portal; a straight portal or monotone portal. As we have added no extra portals or vertices, we end up either with portalgon $\mathcal{Q}$ with n vertices in total, m straight portals and one non-monotone non-interfering portal or one monotone portal or immediately with the desired $\mathcal{P}'$.

Now $\mathcal{Q}$ meets the requirements of Corollary 3.2.9 if it has a non-monotone non-interfering portal and Theorem 3.1.6 if it is a monotone portal. Thus, we know we have a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with $O(n)$ vertices in total and $O(n)$ straight portals. ☐

Next, we continue by proving the claim for a portalgon with one non-monotone self-interfering fitting portal e where there might be interior obstacles.

Lemma 3.2.13. *Let portalgon $\mathcal{P}$ have n vertices in total, m straight portals and one non-monotone self-interfering fitting portal e, where there might be interior obstacles in the fragments of the portalgon which contain the portal edges and the edges are on different fragments, then there is an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals.*

Proof. We are going to prove this Lemma by first solving the type-2-self-interfering parts of the fragments. We will end up with multiple portals, that might be portal interfering. We solve the portal interfering parts and then solve the portals independently as they are restricted to forms we have seen earlier. Then we bring the results together to give the desired result.

Let there be z interior obstacles in the type-2-self-interfering parts. This means that the type-2-self-interfering parts is bounded by z. We will show how to solve one of these parts and show that there are no new type-2-self-interfering parts are created. By induction, this means that we can solve all type-2-self-interfering parts of the portal. Suppose we have type-2-self-interfering part v with $c \leq z$ obstacles. We add portals along the portal line segment from the first intersection v_1 to the last intersection v_2 of v with the portal line segment. These new portal have the same orientation as the portal line segment. The cut-off piece v' is now bounded by the portal line segment and the portal edge from v_1 to v_2 as the interior obstacle only take area away and does not add any area. The piece v' also has the whole of the portal edge between v_1 to v_2 as obstacle cannot intersect with portals and such cannot take area away there. There cannot be any exterior obstacles in the space bordered by the portal edge between v_1 and v_2 and the portal line segment on the fragment containing the other portal edge, since we have no exterior obstacles and any exterior obstacle further down the branch would also be an exterior obstacle earlier (see Figure 19 for an example). So, we can merge our branch into the other fragment without any problems. After merging we merge the first straight portal (containing v_1) with the portal part from the start vertex of e to v_1 and the last straight portal (containing v_2) with the portal part from v_2 to the end vertex of e

Now after solving for v we have a maximum of $3c + 1$ new straight portals (maximum $c + 1$ new portals for cutting over the base of v but being interrupted c times and $2c$ for every portal that can be on the either edge of the obstacles) and $4c$ new vertices (the start and end vertices of the new portals). After solving v, we have only straight pieces in the place of the curved portal parts that was self interfering, the rest of the portal has remained the same and so, after solving v the amount of type-2-self-interfering parts decrease by one and the amount of interior obstacles in type-2-self-interfering parts decreases by c. Now, as we have a maximum of z interior obstacles in all type-2-self-interfering parts together, the maximum amount of new vertices and portals after solving every type-2-self-interfering part is $4z$ for both.

Now, some of the fragments contain portal edges of different parts of the original portal. These can be portal interfering, if they are we first need to fill up the parts that are portal interfering and then we know by Remark 3.1.8 that we can still solve the situation as if these portal edges are on different fragments, so we can solve them independently.

Now the different portals might be portal interfering, where they were first type-1-self-interfering. Let portal k be interfering with l. We solve this by setting up an augmented dual tree of k, in which we also add an edge for every cut made by the portal edge of l, we now cut off every leave of the branches of two or bigger of the dual tree (exactly these branches where type-1-self-interfering in the original portal) with an arbitrarily small offset to the exterior side of the portal line segment such that after merging it we will not still have an intersection between the portal k and the portal line segment of l. This offset is smaller than the smallest distance between the portal edge and the point with a tangent parallel to the portal line segment to circumvent the same problem as we had with cutting off obstacles. This new portal y has the same orientation as the portal it is cutting it off from. These leaves do not have interior obstacles, if they had one, they would have been type-2-interfering, and those have solved. We merge them into the corresponding parts of the other portal edge and we merge the curved portal pieces with the portal edges of y. Note that the former is possible as they have the same form and we have no exterior obstacles (so, also not in the places of these leaves as those obstacle would have already needed to pass the portal line segment and so they would classify as exterior obstacles). And the latter is possible as they have the same orientation. This does not impact the rest of the portal forms, such that it cannot created any new instances of portal interferences. As such after every operation the amount of instances reduces by one.

As every portal interfering instance in this case is between two portals and as we have a finite amount of these portals, we have a finite amount of cases in which there is portal interference. When we have solved every case, which is possible as every operation reduces the amount of cases by one, we have not-portal-interfering portals P that are parts of the original fragment. By Remark 3.1.8 we can solve every resulting portal independently. Any of these pairs are either:

- Straight;
- Monotone;
- Non-monotone non-interfering;
- Non-monotone self-interfering, but only type 1.

The other kinds of interference are not possible as we have just solved all type 2 self-interference and as the portal edges still are subsets of the original portal they are necessarily neither inside nor outside interfering.

Any straight portal, we leave as is, we do not need to solve it any further. In the case of the non-monotone self-interfering portal, we meet the requirements of Lemma 3.2.12. In this cases we end up with one straight portal and no other extra vertices or portals.

In the case of the monotone portal, we meet the requirements of Theorem 3.1.6. In the case of the non-monotone non-interfering portal, we meet the requirements of Corollary 3.2.10. Let x be the portal. Let x have r interior obstacles. Then in these cases after transforming x, we have $r+1$ straight portals instead of x (the original straight portal that would have substituted x now has been interrupted r times), a maximum of $2r$ extra portals (by cutting up those on the boundary of the obstacles) and a maximum of $4r$ new vertices (the start and end vertices of the new portals).

We know that all interior obstacles in these resulting portals P together with z is bounded by n as any vertex can be just one obstacle to just one portal edge. Hence, the total amount of new vertices created by P are bounded by $4(n-z)$ and the amount of new portals is bounded by $4(n-z)$. Together with the new vertices and portals created when removing type-2-self-interfering branches, which are bounded by $4z$ and $3z$ respectively, we get that the final portalgon has $O(n)$ portals, $O(n)$ vertices and has no non-straight portals. It is clear that it is equivalent to $\mathcal{P}$. So, we have proven the Lemma as desired.

□

Then, for a portalgon with one non-monotone self-interfering fitting portal and obstacles.

Theorem 3.2.14. *Let portalgon $\mathcal{P}$ have n vertices in total, m straight portals and one non-monotone self-interfering fitting portal e, where there might be obstacles in the fragments of the portalgon which contain the portal edges and the edges are on different fragments, then there is an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals.*

Proof. We will first cut off all exterior obstacles and then we can use Lemma 3.2.13 to solve the remaining situation.

As determined before, in the proof of Theorem 3.1.6 we know that cutting off all obstacles give at most an additional $5n$ vertices and at most an additional $4n$ straight portals. So, after this operation, we are left with the portalgon portalgon $\mathcal{Q}$, equivalent to $\mathcal{P}$, with at most $6n$ vertices in total and $4n+m$ straight portals and one monotone non-straight fitting portal, where the edges are on the same fragment and there might only be interior obstacles. In the case of a non-monotone portal, the portal could be in one of the obstacle we are trying to cut off, but since that is inside/outside interference and we are only in a self-interfering case, that is not possible.

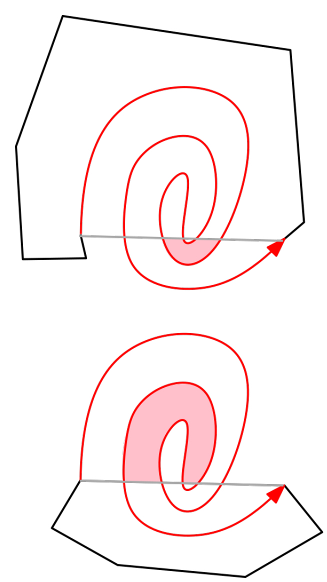

Figure 19: In the shaded areas it is impossible for an obstacle to be without it also passing over the portal line segment twice.

Now $\mathcal{Q}$ meets the requirements of Lemma 3.2.13. Thus, we know we have a straight portalgon $\mathcal{P}'$, equivalent to $\mathcal{P}$, with $O(n)$ vertices in total and $O(n)$ straight portals. □

We still need to prove that we get the right equivalent portalgon when both edges are on the same fragment. The only new situation now is when the portal edges are self interfering with respect to their own portal line. As we have already dealt with self-interference with the other portal edge, we can now use the same techniques for self-interference with its own portal edge and as such, we have the following result.

Corollary 3.2.15. Any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one non-monotone self-interfering fitting portal, where the edges are on the same fragment, can be transformed into an equivalent straight portalgon $\mathcal{P}'$ with $O(n)$ vertices in total and $O(n)$ straight portals.

3.3 Conclusion

We now know that any portalgon $\mathcal{P}$ with n vertices in total, m straight portals and one fitting portal, that is either monotone, not-interfering non-monotone or self-interfering non-monotone, can be transformed into an equivalent portalgon $\mathcal{P}'$ with $O(n)$ vertices in total, $O(m+n) = O(n)$ straight portals and no fitting portals. So, now we can apply Löffler et al.'s theory [9]; if h is the happiness of $\mathcal{P}'$, we know that the complexity of the shortest path is $O(n + hn) = O(hn)$ in $\mathcal{P}'$.

4 Open problems

There remain many open problems. The original problem brought to the attention of the authors of “Shortest Paths in Portalgons” was that of Voronoi diagrams on clothing patterns, where the seam lines are the portals between the pieces of cloth. This was translated into subsets of the plane with an equivalence relationship defining the seam lines as portals. In “Shortest Paths in Portalgons”, the authors explored the problem of the shortest paths in this space where the subsets were simple polygons. The authors mention that the shortest non-simple polygons could be solved in the same way as you merely could add portals to divide the non-simple polygon up into multiple simple polygons and then the original solution is used to solve this problem. [9]

In this thesis, we have explored the problem of the shortest paths in the case of a subset of fitting portals by finding out that we any portalgon with one fitting portals has an equivalent straight portalgon. It is now still an open question whether or not all portalgons with fitting edges can be transformed into straight portalgon, so, in the cases of a portalgon with a non-monotone inside/outside-interfering portal, a portalgon with a non-monotone mixed-interfering portal or a portalgon with multiple fitting portals. For the former two cases, the main problem is recursive interference. It quickly becomes unclear what the inside is of the polygon and what the root is of its dual graph. One of my ideas, which I have not been able to develop in this thesis, includes setting up a dual tree, like the one in the case of self-interference, and pruning the branches or leaves in a specific order. My supervisor also mentioned maybe having different kinds of leaves. Another idea was splitting the whole portalgon up over the extended portal line, the whole line through the start and end vertex instead of the portal line between the start and end vertex. It might also be needed to subdivide this kind of portal into more kinds.

Another open problem is the computation time for transforming the discussed portalgons to straight portalgons and what is the optimal computation time for computing the shortest path in the original portalgon. Along the same lines, in what situation would transforming the portalgons first and then performing the shortest path query be faster than immediately performing the shortest path query?

Then some open problems are further removed from this thesis, yet are still relevant to the subject. In the clothing industry, not only seam lines can be curvy, but boundary edges can also be curvy. It would be interesting to see how that would complicate the problem and to see whether there are shortcuts that are available to simplify the problem. Some of my ideas are temporarily interchanging the curve with a polygonal curve or interchanging it with a straight boundary edge and then computing the extra distance that now has been removed.

Then there is the obvious open problem of what about the non-fitting portals. Is it always possible to transform portalgons with non-fitting portals, where the portal edges still have the same length, to a portalgon with a bounded complexity for shortest paths? Is there a way of transforming them into straight portalgons? The latter might be possible by describing the curve as an infinite amount of straight edges. It might also be beneficial to approximate a curve to a simpler form.

Penultimately, there is also the question whether it is useful to transform a straight portal into a curved portal in certain situations. It might be possible that creating a curved portal in certain cases can be a helpful intermediate step, just like imaginary numbers sometimes can be. Lastly, an interesting question which N. Liefrink is now exploring is what happens when you allow two straight portal edges to have different lengths. He is looking at two portals that are mapped uniformly to each other, it would also be interesting to find out what happens if you allow different mappings. This, of course, brings up the new open problem of what happens when you allow two non-straight edges to have different lengths.

References

[1] G. Amoruso. *Handbook of research on visual computing and emerging geometrical design tools*. IGI Global, 2016.

[2] B. Aronov. "On the geodesic Voronoi diagram of point sites in a simple polygon". In: *Algorithmica* 4.1 (1989), pp. 109–140.

[3] Jeff Erickson. *Simple polygons*. 2013.

[4] L. Gewali and V. Roman. "Generalization of shortest path map". In: *2010 Seventh International Conference on Information Technology: New Generations*. IEEE. 2010, pp. 296–300.

[5] T.H. Hales. "Jordan's proof of the Jordan curve theorem". In: *Studies in logic, grammar and rhetoric*. Vol. 10. 23. 2007, pp. 45–60.

[6] John Hershberger and Subhash Suri. "An optimal algorithm for Euclidean shortest paths in the plane". In: vol. 28. 6. SIAM, 1997, pp. 2215–2256. DOI: 10.1137/S0097539795289604.

[7] C.E. Leiserson, R.L. Rivest, T.H. Cormen, and C. Stein. *Introduction to algorithms*. Vol. 3. MIT press, 2009, pp. 16–64. ISBN: 9780262033848.

[8] F. Li and R. Klette. *Exact and approximate algorithms for the calculation of shortest paths*. Springer, 2011. ISBN: 9781447122555. DOI: 10.1007/978-1-4471-2256-2.

[9] M Löffler, Rodrigo Ignacio Silveira, Frank Staals, and Tim Ophelders. "Shortest paths in portalgons". unpublished. 2022.

[10] K. Ruohonen. *Graph Theory*. Tampere University Of Technology, 2008, pp. 1–12.

[11] G. Tel. *Datastructuren - Analyse van Algorithmen*. https://www.cs.uu.nl/docs/vakken/ds/. (Visited on 05/17/2022).

[12] H. Wang. "A new algorithm for Euclidean shortest paths in the plane". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 975–988.