Rheinische Friedrich-Wilhelms-Universität Bonn
Institut für Informatik I

Visibility Domains and Complexity

Dissertation

Zur Erlangung des Doktorgrades (Dr. rer. nat.)
der Mathematisch-Naturwissenschaftlichen Fakultät
der Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Alexander Gilbers

Moers

Bonn, 2013

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät
der Rheinischen Friedrich-Wilhelms-Universität Bonn

Gutachter: Prof. Dr. Rolf Klein, Universität Bonn
Prof. Dr. Heiko Röglin, Universität Bonn

Tag der Promotion: 13.01.2014

Erscheinungsjahr: 2014

Acknowledgements

I want to thank my advisor Prof. Rolf Klein. Without him this thesis would never have been written.
I also want to thank the other members of the committee, especially Prof. Heiko Röglin who agreed to be the second referee of this thesis.
Thanks to the many colleagues I had over the years for all the inspiring conversations. Without you, lunch would have been so much lonelier.
Thanks to Frank Hoffmann and Klaus Kriegel for many interesting discussions on topics that are not subject to this thesis but that we hopefully one day will explore further.
Many thanks to Magdalena Aretz and Ralf Valerien for proofreading this thesis. Your untiring efforts were really invaluable to me.
Thanks to my parents and my sisters for having been a source of support all my life and thanks to my niece and my nephews for giving me hope for the future.
Thanks to my friends for all the good times we had.
Finally, thanks to my beloved Olga for enriching my life in so many ways.
Words cannot express what you mean to me.

Contents

List of Figures

1 Introduction

In this thesis, we concern ourselves with *visibility domains* in computational geometry. A visibility domain can be thought of as the region of his environment an observer can see. In our model we assume that we are dealing with an observer that is equipped with all-round vision. Moreover, his sight is potentially unbounded, i.e. things that are arbitrarily far away can be seen by him, as long as no object blocks his view.

Now, two observers at two distinct locations will not see the same things: If they are at completely different places, like on different continents, their visibility domains will typically be disjoint. If the observers, by contrast, are in the same otherwise empty room, their visibility domains will hardly differ. The differences between visibility domains and the nature of intersections of visibility domains are the fundamental themes of this thesis. In this work we will encounter, in the context of two problems in discrete and computational geometry, questions about the combinatorial complexity of arrangements of visibility domains and about the hardness of path planning under cost measures defined by visibility domains.

The first problem is to estimate the VC-dimension of visibility domains and is addressed in Chapter 3. The VC-dimension is a fundamental parameter of every range space that is a standard tool in computational learning theory. The VC-dimension of visibility domains forms a measure for how complicated a set of visibility domains can intersect.

In computational geometry, it is typically used to derive upper bounds on the size of hitting sets via the Epsilon-Net Theorem. The VC-dimension contributes a constant factor to this upper bound, so better bounds on the VC-dimension directly translate into better bounds on the size of hitting sets. The special case of the hitting set problem where the ranges are visibility polygons is the famous *art gallery problem* that is a central problem in computational geometry since the very beginnings of the field until today.

Estimating the VC-dimension of other geometric range spaces is often very easy. The VC-dimension of visibility domains stands in sharp contrast to this. In the last fifteen years the question attracted many researchers, among them the most distinguished discrete geometers. Nevertheless, the VC-dimension of visibility domains could only be determined in some special cases while the question for the important case of visibility polygons in simple polygons is still open.

In this thesis we develop new tools to tackle this problem. We will exploit the symmetry of visibility domains and combine encircling arguments with a new kind of decomposition techniques. The main ingredient of our novel approach is the idea of relativization. This method makes it possible to replace in the analysis of intersections the complicated visibility domains by simpler geometric ranges such as wedges, half planes or intervals. This gives us the possibility to use upper bounds for the VC-dimensions of the dual range spaces of the simpler objects for our setting.

In Section 3.5 we will deal with the case of visibility polygons in simple polygons. For this case, the best known lower bound is 6. We derive an upper bound of 14 that improves significantly upon the previously known best upper bound of 23 and halves the range of possible values for the VC-dimension.

In [30], King and Kirkpatrick ask for an upper bound for the VC-dimension of perimeter visibility domains. In Section 3.4 we use the machinery that we developed to show an upper bound of 7 that leaves only a very small gap to the best known lower bound of 5.

The second problem is that of the barrier resilience of visibility domains that Chapter 4 is devoted to. In barrier resilience problems, we are given a set of barriers and two points s and t in $\mathbb{R}^d$. The task is to find the minimum number of barriers one has to remove such that there is a way between s and t that does not cross a barrier. In the field of sensor networks, the barriers are interpreted as sensor ranges and the barrier resilience of a network is a measure for its vulnerability. Barrier resilience problems have attracted growing interest also in the computational geometry community in recent years. The barrier resilience of arrangements of geometric objects such as circles and line segments have been considered.

We initiate the investigation of the very natural special case where the barriers are visibility domains. We will therefore consider the problem of finding a path connecting points s and t that crosses a minimum number of given visibility domains, a so-called minimum witness path.

In Section 4.2 we will consider the barrier resilience of arrangements of visibility domains inside polygons. For visibility domains in simple polygons we will take advantage of the topological structure of simple polygons that enables us to find an optimal path very efficiently. In polygons with holes the situation is more complicated. We will show for this case an approximation hardness result that is stronger than previous hardness results in geometric settings. In Section 4.4 we will consider two different three-dimensional settings and demonstrate their relations to the Minimum Neighborhood Path problem and the Minimum Color Path problem in graphs. We will be able to give a 2-approximation for one of the three-dimensional problems. For the general problem of finding minimum witness paths among polyhedral obstacles we will show that it is not approximable in a strong sense.

All figures in this thesis have been created with the help of the extensible drawing editor Ipe by Otfried Cheong.

2 Preliminaries

2.1 Geometric Preliminaries

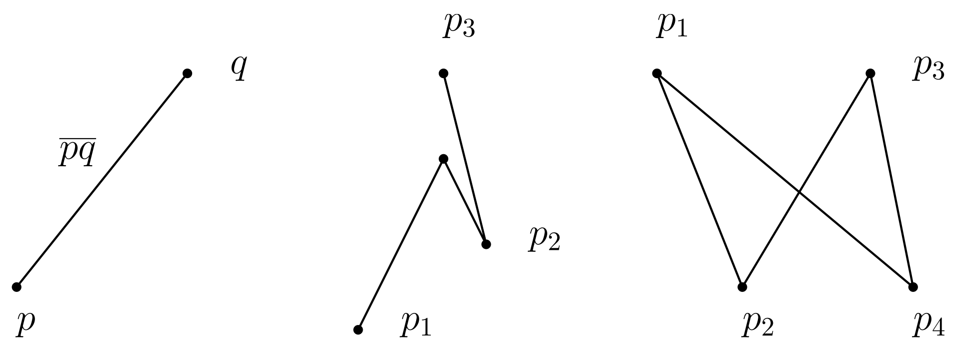

Figure 2.1: A line segment, a polygonal chain and a closed polygonal chain.

Let us now discuss some preliminaries that we will make use of in the later chapters and start with the remark that often in computational geometry, it is assumed that all finite point sets we encounter are in *general position*. This means (usually), that no three points lie on a common line and, in three dimensional settings, no four points lie on a common plane.

For two points $p, q \in \mathbb{R}^2$, we denote by $\overline{pq}$ the *line segment* between p and q, $\overline{pq} = \{\lambda p + (1-\lambda)q | \lambda \in [0,1]\}$. A *polygonal chain* is a subset of $\mathbb{R}^2$ of the form $B = \bigcup_{i=1}^{n-1} \overline{p_i p_{i+1}}$. The points $p_1, \ldots, p_n$ are called the *vertices* of B, the segments $\overline{p_i p_{i+1}}$ are called the *edges* of B. If $p_1 = p_n$ we call B a *closed polygonal chain*, see Figure 2.1. If all vertices $p_1, \ldots, p_{n-1}$ are pairwise distinct and if no two edges of B intersect except at their endpoints, then B divides $\mathbb{R}^2$ into two regions, one unbounded, the other one bounded. Then the topological closure P of the bounded region is called a *simple polygon*. Its boundary is exactly B.

Let now P_0 be a simple polygon with boundary B_0 and $P_1, P_2, \ldots, P_k \subset P_0$ be pairwise disjoint simple polygons with respective boundaries $B_1, B_2, \ldots, B_k$ that are all contained in P_0. For every subset X of $\mathbb{R}^d$, we let X° denote the topological interior of X. Then the set $P = P_0 \setminus \bigcup_{i=1}^{k} P_i^\circ$ is called a *polygon with holes* and its boundary is $B = \bigcup_{i=0}^{k} B_i$, see Figure 2.2.

A *convex set* $C \in \mathbb{R}^d$ is a set that contains for every two points $a, b \in C$ also the whole line segment $\overline{ab}$. The *convex hull* $ch(A)$ of a set $A\mathbb{R}^d$ is the intersection of all convex subsets of $\mathbb{R}^d$ that contain A. A *star-shaped set* $S \in \mathbb{R}^d$ is a set

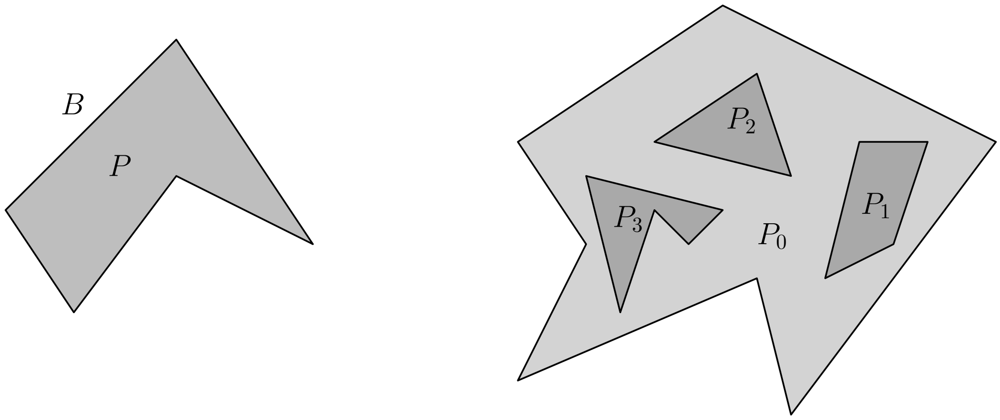

Figure 2.2: *Left:* A simple polygon P with boundary B. *Right:* A polygon with holes which is the set difference between P_0 and the polygons P_1, P_2, P_3.

that contains a point $s \in S$ such that for all points $p \in S$ the whole line segment $\overline{sp}$ is contained in S.

In our settings the visibility domains will always be of the following form: There is a point p in our ground space $X = \mathbb{R}^d$, where typically $d = 2$ holds. Moreover, we have a collection of objects $\{H_i\}_{i \in I}$ that block vision. The visibility domain of p, $\mathrm{vis}(p)$ is defined to be the set of all points q in X, that can be seen by p because the visibility segment $\overline{pq}$ does not intersect the interior of an obstacle. Note that usually it is allowed that part of a visibility segment follows the boundary of an obstacle, see Figure 2.3.

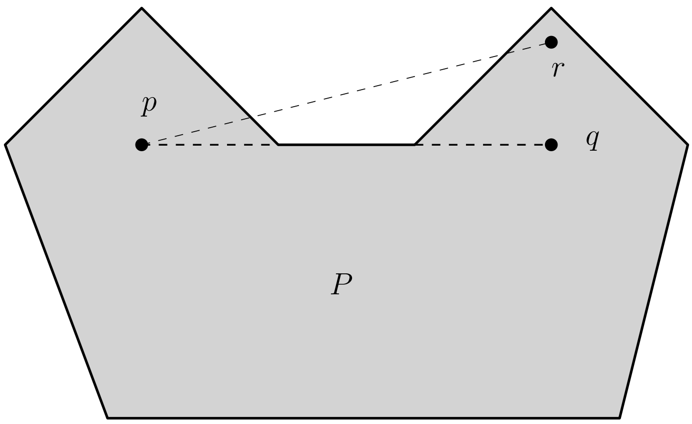

Figure 2.3: The segment $\overline{pq}$ partly follows the boundary but does not cross it. Consequently, p sees q. The segment $\overline{pr}$, on the other hand, crosses the boundary and intersects the complement of P. Therefore, p does not see r.

In chapter 3 we will have the case where p is a point inside the simple polygon $P \subset \mathbb{R}^2$. The (only) obstacle in this case is the complement of P. The visibility domain of p then is itself a simple polygon and is in this case also called the *visibility polygon* of p. By definition, a point $p \in P$ does not see any point of the complement of P.

The vertices of a simple polygon can be divided into *convex vertices* and *reflex vertices*. Convex vertices have an interior angle α with $\alpha \leq \pi$. Reflex vertices have an interior angle $\alpha > \pi$. If all vertices of a simple polygon are convex, it is itself a convex polygon and every point $p \in P$ sees all points of P, i.e. $\text{vis}(p) = P$. Consequently, in order that a point $p \in P$ can not see all points $q \in P$, there has to be some reflex vertex r on the boundary of P, see Figure 2.4.

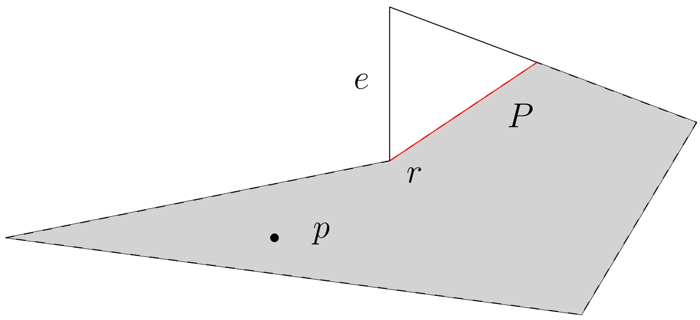

Figure 2.4: A simple polygon with four convex vertices and a reflex vertex r. In red, we see a visibility cut. No point of edge e is seen by p. The visibility polygon of p (gray) has the same number of edges as P.

It is a result of D.T. Lee [35] that the visibility polygon of a point p in a simple polygon P with n vertices can be computed in time $O(n)$. The edges of the visibility polygon are edges or parts of edges of the original polygon P or they are *visibility cuts*.

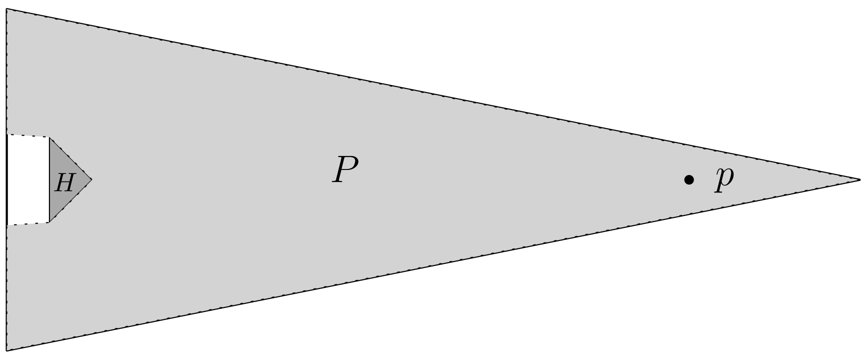

Figure 2.5: The visibility polygon of p (lightgray) has got more edges than the polygon with hole P.

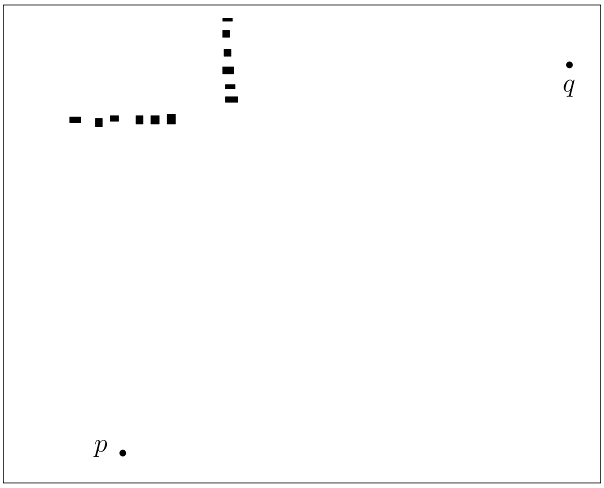

Figure 2.6: A configuration in which the intersection of two visibility polygons has got high boundary complexity.

What is a visibility cut? Let p be some point inside P and let r be a reflex vertex that p can see. Now take the ray with origin p and going through r. We follow this ray from p until we arrive at r. If we go on in the same direction there are two possibilities: It can be that the ray enters the complement of P immediately after r. Then no visibility cut is generated. Otherwise, we follow the ray from r to the last point f of it that still is contained in P, see Figure 2.4. (There is such a point, as P is a closed and bounded set.) The segment $\overline{rf}$ is now called a visibility cut of p as it cuts off the part of P lying behind $\overline{rf}$ from the visibility polygon of p.

Let us compare the number of edges of a visibility polygon of p with the number of edges in the original simple polygon P. Every new edge of vis(p), i.e. every visibility cut starting at reflex vertex r, cuts off a part of P and with this it cuts off completely one of the two old edges (edges of P) that are adjacent to r. This old edge is therefore not an edge of vis(p). Therefore the visibility polygon itself has got at most n edges.

The visibility domain of a point in a polygon with holes is also a simple polygon. Asano [3] showed that one can compute this visibility polygon in time $O(n \log h)$, where h is the number of holes and n is the total number of vertices of the surrounding polygon and the holes. Figure 2.5 shows an example for the fact that a visibility polygon in a polygon with holes can have more edges than the polygon and holes together. However, every vertex of the visibility polygon is a vertex of P (where we count the vertices of the holes as vertices of P) or the endpoint of a visibility cut. There are at most as many visibility cuts as there are vertices of P. Therefore, the visibility polygon can have at most twice as many vertices (and edges) as P. In this respect at least, a visibility polygon by itself is not exceedingly complex.

What about the intersection of two visibility polygons? Finke and Hinrichs showed in [15] that the intersection of two polygons with n edges can be computed in time $O(n+k)$ where k is the number of edge intersections. We know that the intersection of two arbitrary polygons with a total number of edges n can have boundary complexity $\Omega(n^2)$. It is not hard to construct a polygon with holes P with n vertices and containing two points p, q such that the intersection $\text{vis}(p) \cap \text{vis}(q)$ has $\Omega(n^2)$ vertices: In Figure 2.6, each of the k holes in the horizontal row generates two edges of $\text{vis}(p)$, each of the k holes in the vertical column generates two edges of $\text{vis}(q)$. The $2k$ edges belonging to p intersect all the $2k$ edges belonging to q, resulting in at least $4k^2$ vertices on the boundary of $\text{vis}(p) \cap \text{vis}(q)$. The total number of vertices of P in this construction is $8k+4$. The intersection of the visibility polygons of two points in the same polygon with holes can therefore be computed in time $O(n^2)$, and this time can be necessary, as the object to be computed can have this complexity.

On the other hand, the visibility polygons of two points in a simple polygon P cannot have that many intersections. Instead, every edge of $\text{vis}(p)$ can cross at most two edges of $\text{vis}(q)$ and these crossings can be thought of as entering and leaving $\text{vis}(q)$: The moment an edge e of $\text{vis}(p)$ leaves $\text{vis}(q)$, it enters the part P' of P that is cut off by the visibility cut of q it just crossed. This part is itself a simple polygon that is bounded by parts of the boundary of P and exactly one visibility cut of q. As e cannot cross this visibility cut for a second time, it will end in P' without having entered $\text{vis}(q)$ again, see Figure 2.7. Therefore the total number of edge intersections between the two visibility polygons is bounded by $2n$ and so the intersection of two visibility polygons in a simple polygon can be computed in time $O(n)$.

For a detailed description of many visibility algorithms in the plane see for example the book by Ghosh [18].

2.2 The Model of Computation

In this thesis we use the Real RAM as our model of computation. The Real RAM was introduced 1978 by Michael Ian Shamos in his doctoral dissertation [43] and has shortly afterwards become the standard model of computation in all of computational geometry. Shamos describes the Real RAM in his dissertation as follows:

"The model that we will adopt for most purposes is a random-access machine (RAM) similar to that described in [Aho(74)] but in which each storage location is capable of holding a single real number. The following operations are available at unit cost:

1. The arithmetic operations $+, -, \times, /$.

2. Comparisons between two real numbers $(<, \leq, =, \neq, \geq, >)$.

3. Trigonometric functions, EXP, and LOG. (In general, analytic functions.) While we normally will not employ these functions, they are included in the model to strengthen our lower bound results.

4. Indirect addressing of memory (integer addresses only).

This model is an amalgam of useful features of the straight-line, computation tree, and Integer RAM models, and we shall refer to it as a *real RAM*."
(Shamos, p. 20, [43])

(The citation '[Aho(74)]' refers to the book *The Design and Analysis of Computer Algorithms* by Aho, Hopcroft and Ullman, [1]).

Many quantities that appear naturally in geometric computations are non-algebraic numbers. The effect of using the Real RAM is that we do not have to worry about them. Another feature of our model of computation is that the arithmetic operations have unit cost, no matter how large or small the absolute value of the numbers involved may be.

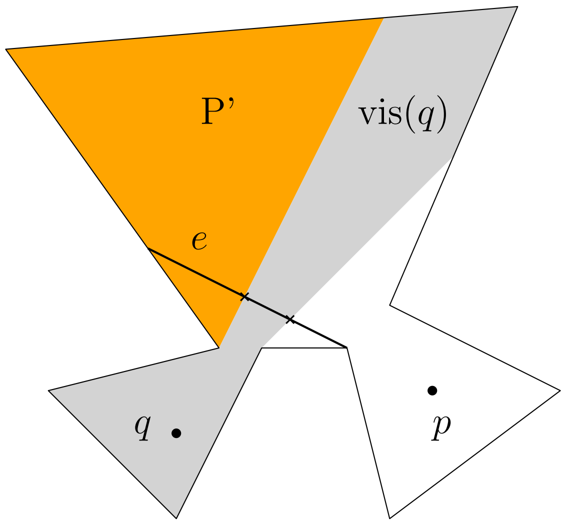

Figure 2.7: After having left vis(q) (grey), edge e will stay in P' (orange) and therefore never enter vis(q) again.

3 VC-Dimension of Visibility Domains

3.1 Introduction

3.1.1 Definition of the VC-Dimension

We first set up the abstract framework in which we can develop our questions.

Definition 1. A *range space* is a pair $(X, \mathcal{F})$ where X is a set and $\mathcal{F}$ is a subset of the powerset of X. The set X is called the *ground set* or, alternatively, the *universe* and the elements of $\mathcal{F}$ are called the *ranges* of the range space.

As mentioned, this definition is very general. In geometric settings, X will often be the $\mathbb{R}^d$ or some subset of $\mathbb{R}^d$. The ranges will usually be geometrically defined subsets of X such as half spaces, balls, boxes, etc. The VC-dimension is a fundamental parameter of a range space that, intuitively speaking, measures how differently the ranges intersect with subsets of the ground set. It is named after Vladimir Vapnik and Alexey Chervonenkis who introduced it in [47]. The following definition of VC-dimension is adopted from [37].

Definition 2. Let $(X, \mathcal{F})$ be a range space. A subset $S \subseteq X$ is said to be shattered by $\mathcal{F}$ if each of the subsets of S can be obtained as the intersection of some $F \in \mathcal{F}$ with S. We define the VC-dimension of $\mathcal{F}$, denoted by $\dim(\mathcal{F})$, as the supremum of the sizes of all finite shattered subsets of X. If arbitrarily large subsets can be shattered, the VC-dimension is ∞.

For many important geometric range spaces such as rectangles, disks or half spaces, the VC-dimension is not hard to estimate. Take for example the VC-dimension of disks in the plane. It is easy to see that there are point sets of size three that can be shattered by disks, see Figure 3.1.

To see that the VC-dimension of this range space is indeed 3, assume there were four points $a, b, c, d \in \mathbb{R}^2$ that were shattered by disks. It is clear that no three of them can lie on a common line. So the three points $a, b, c \in \mathbb{R}^2$ are in general position and determine a unique disk D on whose boundary they lie. If the fourth point $d \in \mathbb{R}^2$ lies in D, there is no disk D' that contains a, b, c but not d. By repeating this argument for all triples, we get that a, b, c, d are in convex position, w.l.o.g. in this order on the boundary of their convex hull.

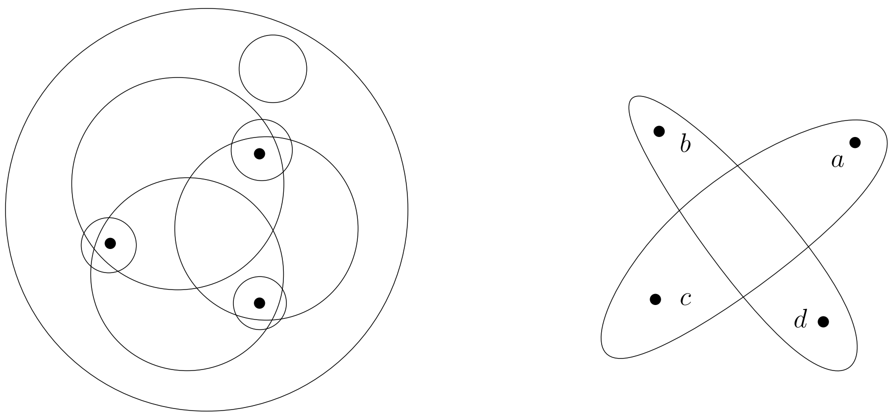

Figure 3.1: *Left:* An example of three points that are shattered by disks. *Right:* The bounding circles of two disks that would contain $\{a, c\}$ and $\{b, d\}$, respectively, would have to intersect at four points.

Assume there were disks $D_{a,c}$, $D_{b,d}$ containing exactly a, c and b, d, respectively. Then the bounding circles of these disks would have to intersect at four points, see Figure 3.1. A contradiction.

3.1.2 VC-Dimension in Computational Learning Theory

VC-dimension is essential in computational learning theory. There, we want to learn a target function from a set of training examples. In the setting of *Probably Approximately Correct Learning* (PAC-learning) we would like to be able to estimate the expected size of a training set of examples needed to learn our approximately correct hypothesis. When it comes to sample complexity bounds in the case of infinite hypothesis spaces, the use of VC-dimensions allows for deriving upper and lower bounds on the number of training examples necessary. For an excellent introduction to computational learning theory and the role of VC-dimensions therein see the classic textbook of Kearns and Vazirani, [28].

3.1.3 VC-Dimension of Visibility Polygons

Besides its importance in machine learning, VC-dimensions became significant to computational geometry, chiefly through their role in the Epsilon-Net Theorem by Haussler and Welzl [25].

In this chapter we are interested in the VC-dimension of visibility polygons. Let us be a bit more precise. What we will derive are not bounds on the VC-dimension of a single range space $(X, \mathcal{F})$ but rather bounds for the supremum of VC-dimensions over all possible range spaces of the form $(P, \mathcal{V})$ where P is a polygon (or, in another case, the boundary of a polygon) and $\mathcal{V}$ is the set of visibility domains defined by points of P. One aspect that makes this problem challenging is that the elements of $\mathcal{V}$ depend heavily on the shape of P. Let us think of the polygons as ground plans of art galleries. The size of

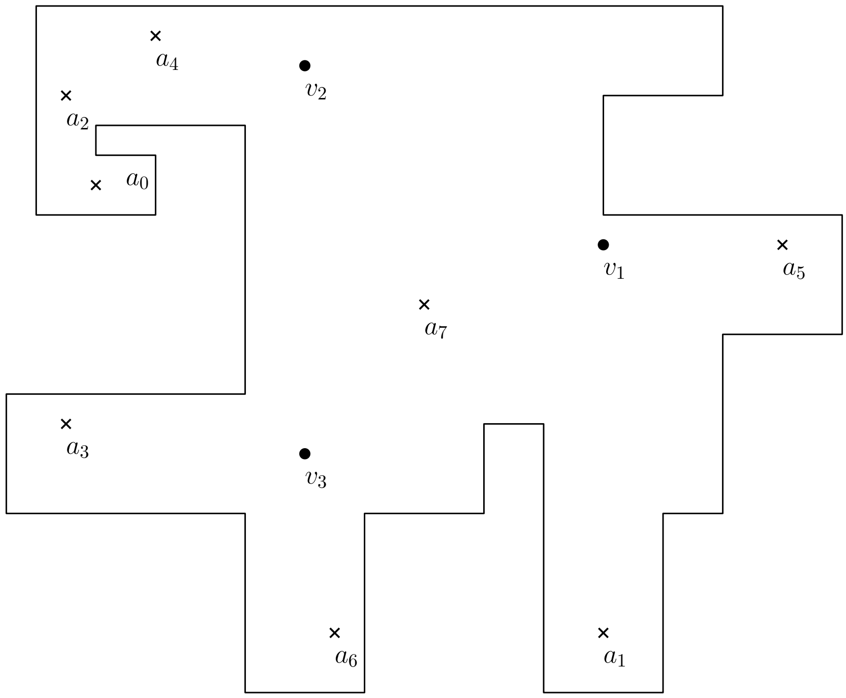

Figure 3.2: Every piece of art $a_0, a_1, \ldots, a_7$ in the art gallery is seen by a different subset of the set of contemplators $\{v_1, v_2, v_3\}$.

the VC-dimension then intuitively means the following. If the VC-dimension of Visibility polygons is d or greater, then there is an art gallery with 2^d sculptures in which we can position contemplators at d specific locations such that every one of the sculptures is seen by a different group of contemplators (see Figure 3.2).

3.1.4 Art Gallery Theory and ε-Nets

The interpretation of a polygon as the ground plan of an art gallery leads to the name *Art Gallery Problem* for the task of determining how many star-shaped regions are needed to cover a given polygon P. Around this fundamental problem, a rich theory has evolved in the last decades of the 20th century. The original motivation to consider the VC-dimension of visibility domains came from this theory of Art Galleries. Formally, Art Gallery Theory is concerned with the following setting:

Definition 3. Given a polygon P, a point set $G \subset P$ is said to *guard* the polygon P if $P = \bigcup_{g \in G} \mathrm{vis}(g)$ (every point in P is under surveillance by at least one point of G). The polygon P in this context is called the *art gallery* and the points in G are usually called *guards* for P. The Art Gallery problem is to find a guard set G of minimum size.

The Art Gallery problem was posed by Victor Klee. Variants of this problem restrict the positions of guards to the boundary of P or even to its vertices. In the 1980s, all three variants have been shown to be NP-hard for polygons with holes by O'Rourke and Supowit [42] and later for polygons without holes by Lee and Lin [36]. At the end of the last millennium, Eidenbenz [13, 14] showed that these problems are even APX-hard and so there is no polynomial time approximation scheme, unless $P = NP$.

On the other hand, there are tight upper and lower bounds on the number of guards needed to guard the polygon P in dependence on the number n of vertices of P. An early result by [10, 16] is that a simple polygon on n vertices can be guarded by at most $\lfloor \frac{n}{3} \rfloor$ guards positioned on vertices of P. Avis and Toussaint [4] were the first to provide an algorithm with $O(n \log n)$ running time to compute a guard set of this size. There are many classical related problems and algorithms in this area concerning variants such as orthogonal polygons, exterior visibility, mobile guards and many more. A wealth of results can be found, e. g., in O'Rourke [41] and Urrutia [45].

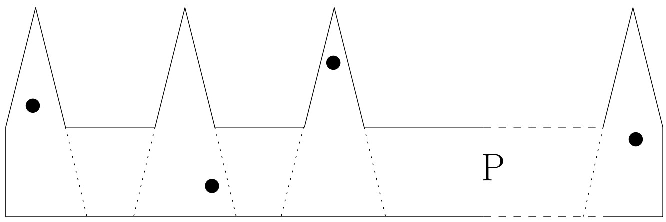

Figure 3.3: An art gallery over $3m + 2$ vertices that requires m guards.

The example depicted in Figure 3.3 shows that $\lfloor \frac{n}{3} \rfloor$ guards may also be necessary. On the other hand, there are many examples of polygons that have many vertices but can be guarded with very few guards. This holds for example for convex polygons that may have arbitrarily many vertices but can of course be guarded by a single guard, see Figure 3.4. It is therefore quite natural to ask for different parameters that we can relate the number of needed guards to. One idea inspired by the example of convex polygons could be to ask for the guards needed depending on the number r of *reflex* vertices of P. It can be shown ([8, 40]) that if there is a positive number r of reflex vertices of P then r guards are always sufficient. On the other hand, there are polygons with $n - 3$ reflex vertices that still can be guarded by one guard, see Figure 3.5. A different approach is based on the following observation. The art gallery in Figure 3.3 has the significant property that it contains many points from which only a small portion of the total area is visible. Now assume we have a gallery where every

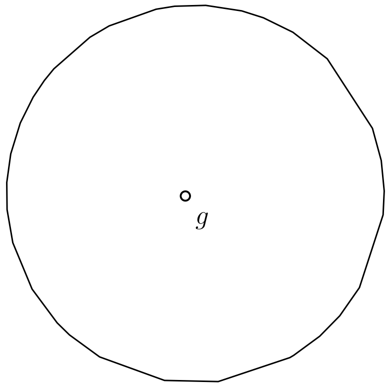

Figure 3.4: Convex polygons are examples of polygons that can have arbitrarily many vertices and still can be guarded by one guard.

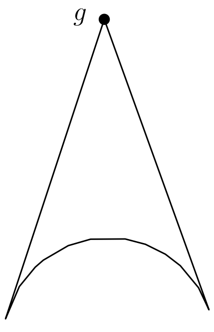

Figure 3.5: There are polygons with arbitrarily many reflex vertices that can be guarded by one guard.

point sees at least a fixed fraction $\frac{1}{r}$ of the area of P. Can we bound the number of guards needed by a function of r? This question leads to the area of *ε-nets*.

Definition 4. Given a range space $(X, \mathcal{F})$ and a measure μ on X such that $\mu(X)$ is finite and every $F \in \mathcal{F}$ is μ-measurable and given a real number ε with $0 \leq \varepsilon \leq 1$. A (strong) ε-net for $(X, \mathcal{F})$ with respect to μ is a subset $N \subset X$ such that for each range $F \in \mathcal{F}$ with $\mu(F \cap X) \geq \varepsilon\mu(X)$ the intersection $N \cap F$ is nonempty.

How does this relate to our question? Let us set $\varepsilon = \frac{1}{r}$. By our assumption, every point in P sees at least $\varepsilon \cdot \text{Area}(P)$. Thus, if we set $\mathcal{V} = \{\text{vis}(p) | p \in P\}$ and take $(P, \mathcal{V})$ as range space and μ as the two-dimensional Lebesgue-measure in $\mathbb{R}^2$, an ε-net N for $(P, \mathcal{V})$ will intersect *all* the visibility polygons in $\mathcal{V}$ and therefore every point in P is seen by some point in N. So the points of N guard P. But how big can an ε-net be?

3.1.5 ε-Nets of Geometric Range Spaces

There are some range spaces where one can derive the existence of small ε-nets directly. Take as an easy example the closed interval $[a, b] \subset \mathbb{R}$ and the set $\mathcal{S}$ of closed subintervals $S \subseteq I$ measured by the Lebesgue-measure. It is clear that the set of points $(a_i)_{i=0}^{\lfloor \frac{1}{\varepsilon} \rfloor}$ with $a_0 = a$ and $a_{i+1} = a_i + \varepsilon(b - a)$ for $i < \frac{1}{\varepsilon}$ will hit every interval of measure at least $\varepsilon(b - a)$.

In most cases, however, it is not easy at all to construct a small ε-net. In this situation, the VC-dimension of the range space comes into play. In [25] Haussler and Welzl showed, that every range space with finite VC-dimension d has an ε-net of size $(\frac{8d}{\varepsilon} \log \frac{8d}{\varepsilon})$ where d denotes the VC-dimension. In [33], Komlós et al. improve upon this result:

Theorem 1 (Komlós et al., [33]). *Every range space of VC-dimension $d < \infty$ has got an ε-net of size*

$$C \cdot d \cdot \frac{1}{\varepsilon} \log \frac{1}{\varepsilon}.$$

The constant C found by Komlós et al. is very small. It follows that, if one can prove that the VC-dimension d of visibility polygons in a simple polygon is finite, then one has shown an upper bound of $C \cdot d \cdot r \log r$ for the number of guards needed to guard any polygonal art gallery where every point sees at least an r-th part of the polygon.

3.1.6 Previous Work on the VC-Dimension of Visibility Domains

The VC-dimension of range spaces of visibility regions was first considered by Kalai and Matoušek [27]. They showed that the VC-dimension of visibility regions of a simply connected gallery (i.e. a compact set in the plane) is finite. In their proof they start with assuming that a large set (of size about 10^{12}) of points A inside a gallery is shattered by the visibility regions of the points of a set B. They then derive a configuration as in Figure 3.6. Here, points a and b should not see each other but the segment $\overline{ab}$ is encircled by visibility segments, a contradiction. This kind of encircling argument also plays an important role in our proof of the new bound in the general case. They also gave an example

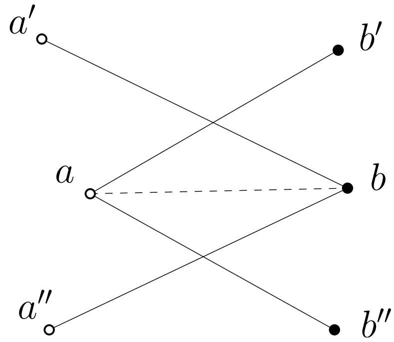

Figure 3.6: Segment $\overline{ab}$ is encircled by visibility segments.

of a gallery with VC-dimension 5. Furthermore, they showed that there is no constant that bounds the VC-dimension if the gallery has got holes.

For simple polygons, Valtr [46] proved an upper bound of 23. In the same paper he gave an example of a gallery with VC-dimension 6 and showed an upper bound for the VC-dimension of a gallery with holes of $O(\log h)$ where h is the number of holes. In this thesis we show that the VC-dimension of Visibility polygons of a Simple Polygon is at most 14. This result has been published in [23].

Isler et al. [26] examined the case of exterior visibility. In this setting the points of S lie on the boundary of a polygon P and the ranges are sets of the form vis(v) where v is a point outside the convex hull of P. They showed that the VC-dimension is 5. They also considered a more restricted version of exterior visibility where the view points v all must lie on a circle around P, with VC-dimension 2. For a 3-dimensional version of exterior visibility with S on the boundary of a polyhedron Q they found that the VC-dimension is in $O(\log n)$ where n is the number of vertices of Q. King [29] examined the VC-dimension of visibility regions on terrains. For 1.5-dimensional terrains, i.e. x-monotone polygonal chains, he proved that the VC-dimension equals 4 and on 2.5-dimensional terrains, i.e. polyhedral surfaces that intersect every vertical line only once, there is no constant bound.

Kirkpatrick ([31]) showed that it is possible to guard a polygon with $O(r \log \log r)$ many perimeter guards (i.e. guards on the boundary of the polygon) if every point on the boundary sees at least an r-th part of the boundary. In [30]

King and Kirkpatrick extended this work and obtained an $O(\log \log \mathrm{OPT})$-approximation algorithm for finding the minimum number of guards on the perimeter that guard the polygon. As an open question they left whether it is easier to find the VC-dimension in the case of perimeter guards than for general visibility polygons. They show that the corresponding VC-dimension is at least 5. In this thesis we show that, using the technique from Gilbers and Klein [21], one can obtain an upper bound of 7 for this VC-dimension.

3.1.7 Guarding a Polygon vs. Guarding its Boundary

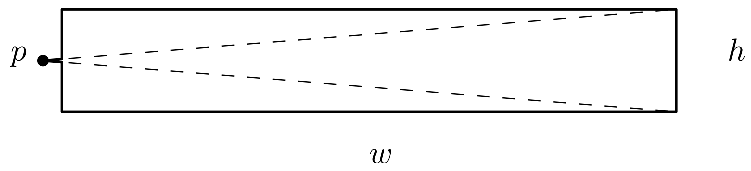

Figure 3.7: As the ratio $\frac{w}{h}$ goes to infinity, the fraction of the boundary seen by p goes to 0. On the other hand, every point in P sees at least half of the polygon. This latter fraction can even be made to be arbitrarily close to 1 if we choose a polygon whose boundary follows the shape of $\mathrm{vis}(p)$ more closely instead of a rectangle.

It is a basic fact, that a point $p \in P$ sees the whole polygon if and only if it sees all of its boundary. It is not the case, however, that a set of points always guard the whole polygon if they guard its boundary. For every $r > 1$ there are examples of polygons in which every point sees an r-th part of the polygon but some points see only an arbitrarily small fraction of its boundary (see Figure 3.7). For $r \geq 3$ there are also examples of polygons, where every point sees at least an r-th part of the boundary but only an arbitrarily small fraction of the polygon (see Figure 3.8). Therefore, bounds for the number of guards needed in one setting do not automatically carry over to the other one.

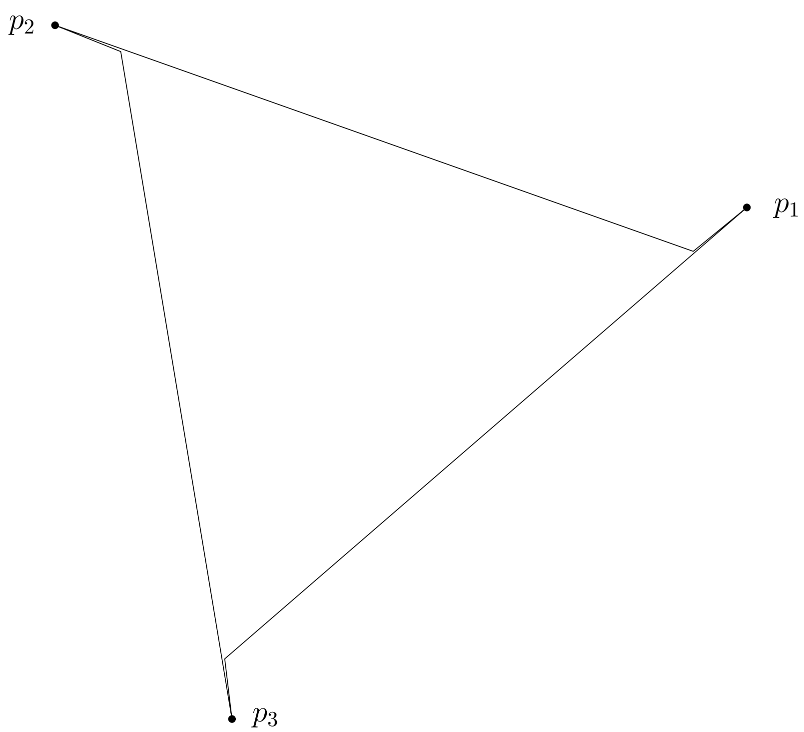

Figure 3.8: Every point in P sees at least a third of the boundary of P. The points p_1, p_2 and p_3 see only a tiny fraction of the area of P, each.

3.2 Properties of the VC-Dimension

As we discussed in the introduction, the VC-dimension is a combinatorial property of a range space $\mathcal{F}$ on a ground set X (also called a *set system*) that asks for the maximum number of points in X such that all subsets of this point set can be obtained as intersections with sets $F \in \mathcal{F}$. In the introduction, we also gave a formal definition of VC-dimension.

A refinement of the VC-dimension is given by the *Shatter Function*.

Definition 5. Let for every given finite subset $Y \subseteq X$ of the ground set X $\Pi_{\mathcal{F}}(Y) = \{Y \cap F | F \in \mathcal{F}\}$ denote the set of subsets of Y that can be represented as intersection of Y with a range $F \in \mathcal{F}$. The *shatter function* $\pi_{(X,\mathcal{F})} \colon \mathbb{N} \longrightarrow \mathbb{N}$ of a range space is then defined by

$$\pi_{(X,\mathcal{F})}(m) = \max_{Y \subseteq X, |Y|=m} |\Pi_{\mathcal{F}}(Y)|.$$

If X is clear from context we will also simply write $\pi_{\mathcal{F}}$ instead of $\pi_{(X,\mathcal{F})}$.

The shatter function measures how many distinct intersections a subset of X of size m can have with ranges of $\mathcal{F}$. For range spaces of finite VC-dimension there is a bound on the size of $\pi_{(X,\mathcal{F})}$ that is fundamental for the theory of ε-nets, the so-called *Shatter Function Lemma.*

Lemma 1. *Let $(X, \mathcal{F})$ be a range space with finite VC-dimension d. Then for all $m \in \mathbb{N}$ we have*

$$\pi_{(X,\mathcal{F})}(m) \leq \sum_{i=0}^{d} \binom{m}{i}.$$

A proof of this lemma can be found for example in [37].

Of course, there can be subsets $Y \subseteq X$ of size m with $\Pi_{\mathcal{F}}(Y) \ll \pi_{(X,\mathcal{F})}(m)$. If a finite subset $Y \subseteq X$ with $|Y| = n$ is shattered by $\mathcal{F}$, then the set $\Pi_{\mathcal{F}}(Y)$ has 2^n elements. Otherwise it has got $p < 2^n$ elements, leaving a deficit of $2^n - p$ subsets of Y that are not intersections of Y with ranges of $\mathcal{F}$. Because we will need this quantity in our proofs, for every finite $Y \subseteq X$ we define the *coarseness* of $\mathcal{F}$ with respect to Y to be $c_{\mathcal{F}}(Y) = 2^n - |\Pi_{\mathcal{F}}(Y)|$. Obviously, Y is shattered by $\mathcal{F}$ if and only if $c_{\mathcal{F}}(Y) = 0$. For a set of ranges that is the union of two subsets, the following relation holds.

Lemma 2. *Let Y be shattered by $\mathcal{F}$ and $\mathcal{F} = \mathcal{F}' \cup \mathcal{F}''$. Then $|\Pi_{\mathcal{F}''}(Y)| \geq c_{\mathcal{F}'}(Y)$.*

Proof. Let Y be shattered by $\mathcal{F}$ and $\mathcal{F}$ be the union of $\mathcal{F}'$ and $\mathcal{F}''$. If $c_{\mathcal{F}'}(Y) = k > 0$ there are k subsets of Y that are not shattered by $\mathcal{F}'$. It is clear that for each of these subsets Z there must be some set $F'' \in \mathcal{F}''$ such that $Z = F'' \cap Y$. Therefore $\Pi_{\mathcal{F}''}(Y) \geq k$. $\square$

To every range space one can assign its *dual range space.*

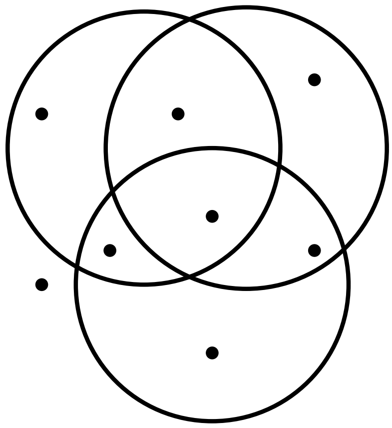

Figure 3.9: The dual range space: All subsets of the set of three disks are stabbed by points.

Definition 6. Let $\mathcal{R} = (X, \mathcal{F})$ be a range space. Its *dual range space* $\mathcal{R}^*$ is the range space $(\mathcal{F}, \{\mathcal{F}_x\}_{x \in X})$ where $\{\mathcal{F}_x\}_{x \in X} = \{F \in \mathcal{F} | x \in F\}$. The original range space is then called the *primal range space*. We sometimes call the VC-dimension of the range space that is dual to $(X, \mathcal{F})$ the *dual VC-dimension of* $(X, \mathcal{F})$.

This means that the elements of the ground set of $\mathcal{R}^*$ are the ranges of $\mathcal{R}$. The ranges of $\mathcal{R}^*$ are subsets of the set of ranges of $\mathcal{R}$: Every range $\mathcal{F}_x$ contains as elements exactly the ranges $F \in \mathcal{F}$ that contain the element $x \in X$ (of course there can be $x \neq y$ with $\mathcal{F}_x = \mathcal{F}_y$). A finite subset $\mathcal{F}' = \{F_1, F_2, \ldots, F_m\}$ of the ground set of $\mathcal{R}^*$ is shattered if and only if for every subset $T \subseteq \{1, 2, \ldots, m\}$ there is an element x_T of X that is contained in all ranges F_i with $i \in T$ but in no range F_j with $j \notin T$. In this case we say that the subsets of $\mathcal{F}'$ are *stabbed* by elements of X.

The range spaces of visibility domains that we are interested in have a remarkable property. If the range space has got the form $(X, \mathcal{V}_X)$, where $\mathcal{V}_X = \{\text{vis}(x) | x \in X\}$, i.e. if the ranges are exactly the visibility domains of points in X (as will be the case in the considered range spaces in sections 3.4 and 3.5) then there is a canonical surjection $\phi\colon P \longrightarrow \mathcal{V}(P)$ of the ground set onto the set of ranges. It is simply given by $p \mapsto \text{vis}(p)$. From the symmetry of the visibility relation it follows now that $p \in \phi(q)$ if and only if $q \in \phi(p)$.

If we want to know if a finite point set $S \subseteq P$ is shattered by visibility domains, we normally would have to look for one point v_T per subset T of S such that $\text{vis}(v_T) \cap S = T$. By the symmetry mentioned above, we can take a different path. For a given subset $T \subseteq S$ we can simply build the intersection of all visibility domains of points in T, $\bigcap_{t \in T} \text{vis}(t)$ and all of the complements of all points in $S \setminus T$, $\bigcap_{s \in S \setminus T} (\text{vis}(s))^c$. If this intersection is nonempty, then every element in this intersection is a point v_T as needed.

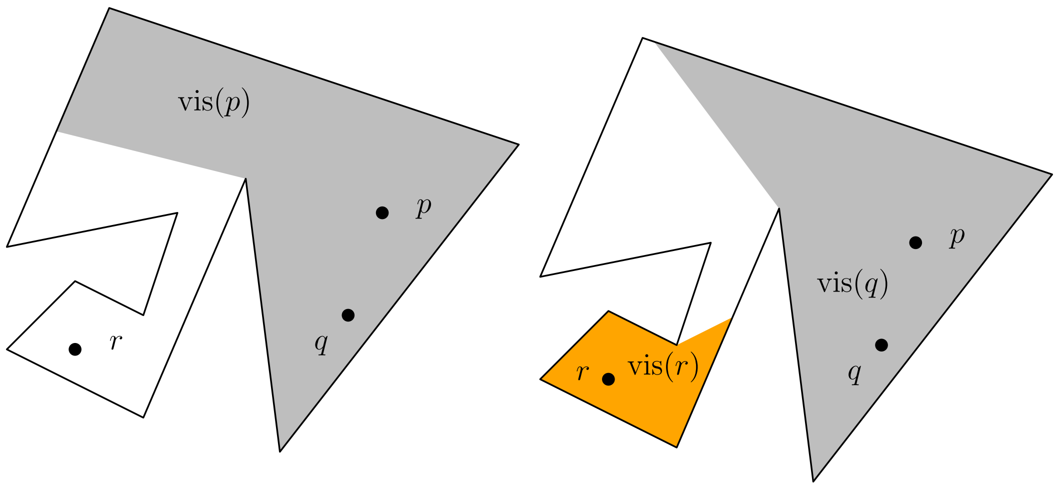

Figure 3.10: The symmetry of visibility: $q \in \text{vis}(p)$ iff $p \in \text{vis}(q)$. The visibility polygons $\text{vis}(q)$ and $\text{vis}(r)$ are disjoint.

Note that, as the intersections of visibility polygons and their complements can be computed efficiently, the above method yields a way of effectively testing, wether a given candidate set S is shattered or not.

The second thing that is notable about this *symmetry* is the following. Suppose you have a set $S = \{s_1, s_2, \ldots, s_m\}$ of m points in P. This set is shattered if and only if the set $S' = \{\text{vis}(s_1), \text{vis}(s_2), \ldots, \text{vis}(s_m)\}$ of m visibility polygons in the dual range space is also shattered. To see this, suppose that S is shattered. That means that for each of the subsets $T \subset S$ there is a point v_T whose visibility polygon contains all points of T but no point of $S \setminus T$. Every subset of S' is of the form $T' = \{\text{vis}(t) | t \in T\}$ for some $T \subseteq S$. The set $\mathcal{V}_{v_T} = \{\text{vis}(p) | p \in P \text{ and } v_T \in \text{vis}(p)\}$ is a range of the dual space. As $\{\text{vis}(p) | p \in P \text{ and } v_T \in \text{vis}(p)\} = \{\text{vis}(p) | p \in P \text{ and } p \in \text{vis}(v_T)\}$ and $s \in \text{vis}(T)$ if and only if $s \in T$ for all $s \in S$ it turns out that $T' = S' \cap \mathcal{V}_{v_T}$.

Let in the other direction S' be a set of m distinct visibility domains of the form $\{\text{vis}(s) | s \in S\}$ (for some $S \subseteq P$) that is shattered. As before, there is for every $T' \subseteq S'$ of the form $T' = \{\text{vis}(t) | t \in T\}$ (for some $T \subseteq S$) a point $v_{T'}$ that lies exactly in the visibility polygons in T' but not in the visibility polygons in $S' \setminus T'$. By symmetry, $\text{vis}(v_{T'})$ contains exactly the points of T but not the points of $S \setminus T$.

It follows that for these range spaces, every shattered point set in the primal range space corresponds uniquely to a shattered set of the same size in the dual range space. In general, if we have a polygon P and two subsets $X, Y \subseteq P$ and a range space $(X, \mathcal{V}_Y)$ where the ground set is X and the ranges are the visibility polygons of points in Y, an analogous proof yields the following lemma.

Lemma 3. *Let P be a simple polygon and $X, Y \subseteq P$ subsets of P, let $\mathcal{V}_X$ and $\mathcal{V}_Y$ denote the sets of visibility polygons of points in X and of points in Y,* *respectively. A finite subset $S \subseteq X$ of X is shattered by visibility domains of points in Y if and only if the set $\{\mathrm{vis}(s) | s \in S\}$ of visibility polygons of points in S is stabbed by points in Y. Therefore the VC-dimension of $(X, \mathcal{V}_Y)$ equals the VC-dimension of $(Y, \mathcal{V}_X)^*$, where $(Y, \mathcal{V}_X)^*$ is the dual of $(Y, \mathcal{V}_X)$.*

For the special case $X = Y$ we get a corollary.

Corollary 1. *Let P be a simple polygon and $X \subseteq P$. Then the VC-dimension of $(X, \mathcal{V}_X)$ is the same as the VC-dimension of its dual range space $(X, \mathcal{V}_X)^*$.*

This Corollary is in sharp contrast to the general case. In general, the VC-dimension of the range space can be exponentially large in the VC-dimension of the primal space.

Exploiting the symmetry between primal and dual range spaces of visibility domains by using Lemma 3 will become one key step in all our proofs. This is because it turns out that we can bound the VC-dimension of the dual range space easier than that of the primal space.

Another fundamental step that we will encounter more than once in our proofs is to find a suitable *decomposition* of a given range space into subspaces and deriving bounds for the entire space from bounds for the subspaces. Note, that Lemma 2 is a first tool that one can use for this task.

We next formulate a property of shattered sets that will be a cornerstone in our proofs:

Lemma 4. *Let Y', Y'' be disjoint subsets of the finite set Y and $\mathcal{F}'$ and $\mathcal{F}''$ subsets of $\mathcal{F}$ whose (not necessarily disjoint) union equals $\mathcal{F}$, i.e. $\mathcal{F}' \cup \mathcal{F}'' = \mathcal{F}$. Then $c_{\mathcal{F}}(Y) \geq c_{\mathcal{F}'}(Y') \cdot c_{\mathcal{F}''}(Y'')$.*

Proof. Let $Z' \subset Y', Z'' \subset Y''$ such that for no $F' \in \mathcal{F}'$: $Z' = Y' \cap F'$ and for no $F'' \in \mathcal{F}''$: $Z'' = Y'' \cap F''$. Then there can be no $F \in \mathcal{F}$ such that $Z' \cup Z'' = Y \cap F$. That means that for every such combination of subsets of Y', Y'' there is a distinct subset of Y that does not have a representation as an intersection of Y with some $F \in \mathcal{F}$. There are $c_{\mathcal{F}'}(Y') \cdot c_{\mathcal{F}''}(Y'')$ such combinations. The inequality follows. $\square$

As a special case of this, we get that for shattered Y, $c_{\mathcal{F}}(Y) = 0$. Therefore in this case $c_{\mathcal{F}'}(Y')$ or $c_{\mathcal{F}''}(Y'')$ must equal zero.

Corollary 2. *Let Y', Y'' be disjoint subsets of the finite set Y and $\mathcal{F}' \cup \mathcal{F}'' = \mathcal{F}$. If $\mathcal{F}$ shatters Y then $\mathcal{F}'$ shatters Y', or $\mathcal{F}''$ shatters Y''.*

Another property that we will make use of is the following.

Lemma 5. *Let X be a set that is shattered by $\mathcal{F}$, $Y \subsetneq X$ and $\mathcal{F}_Y = \{F \in \mathcal{F} | Y \subset F\}$. Then $X \setminus Y$ is shattered by $\mathcal{F}_Y$.*

Proof. Suppose not. Then there is some $X' \subset X \setminus Y$ such that $F \cap X \setminus Y = X'$ for no $F \in \mathcal{F}_Y$. But then there can be no $F \in \mathcal{F}_Y$ with $X \cap F = X' \cup Y$. As there can also be no such set in $\mathcal{F} \setminus \mathcal{F}_Y$, X cannot be shattered by $\mathcal{F}$. $\square$

The next conceptual building block of our proofs besides exploiting the symmetry and finding a good decomposition is something we shall name *relativization.*

Definition 7. Let $(X, \mathcal{F})$, $(X, \mathcal{R})$ be two range spaces on the same ground set and let $Y \subseteq X$ and $F \in \mathcal{F}$. If there is some $R \in \mathcal{R}$ such that $Y \cap F = Y \cap R$ then we say that F is $\mathcal{R}$-like relative to Y.

The general method of relativization we will use to show upper bounds on the VC-dimension of visibility domains can be described as follows.

1. We assume that the VC-dimension of our range space $(X, \mathcal{V}_Y)$ is some integer d.

2. We derive the existence of a point set S of size d that is shattered by visibility domains.

3. We use Lemma 3 to obtain that the visibility domains of points in S are stabbed.

4. As a consequence of Lemma 5, for each $S' \subsetneq S$ the visibility domains of S' are stabbed by points in $\bigcap_{s \in S \setminus S'} \operatorname{vis}(s)$.

5. We find some subset $S' \subseteq S$ such that the visibility domains of points in S' are $\mathcal{R}$-like relative to $\bigcap_{s \in S \setminus S'} \operatorname{vis}(s)$, where $\mathcal{R}$ is a set of ranges such that the VC-dimension of $(Y, \mathcal{R})^*$ is less than d.

6. We conclude, that the visibility domains of points of S' cannot be stabbed. A contradiction.

While this general structure is common to all of our proofs for bounding VC-dimensions, there is no standard way for finding an appropriate S'. Especially in Section 3.5 proving that such a subset always exists is the most challenging part of the proof.

From the assumption that S is shattered by visibility domains it follows that for every subset $T \subseteq S$ there is a *view point* v_T that sees all points $t \in T$ but no point of $S \setminus T$. We will often choose one such set V of $2^{|S|}$ view points. Then for every subset T of S we will denote by V_T the subset of V containing exactly the points that see at least T, $V_T = \{v \in V | T \subset \operatorname{vis}(v)\}$, see Figure 3.2.

After we have found S', we still have to show that the visibility domains of S' are $\mathcal{R}$-like. For this, we make use of encircling type arguments. We mentioned before, that Kalai and Matoušek used an encircling argument in their proof in [27]. Encircling arguments are used to show that two points p, w see each other. To show this, we will provide a polygonal chain consisting of visibility segments (i.e. line segments connecting points that see each other) that encircles the segment $\overline{pw}$. It follows that p and w see each other. Otherwise there had to be some point $x \in \mathbb{R}^2 \setminus P$ on the segment. As P is simple, there must be a

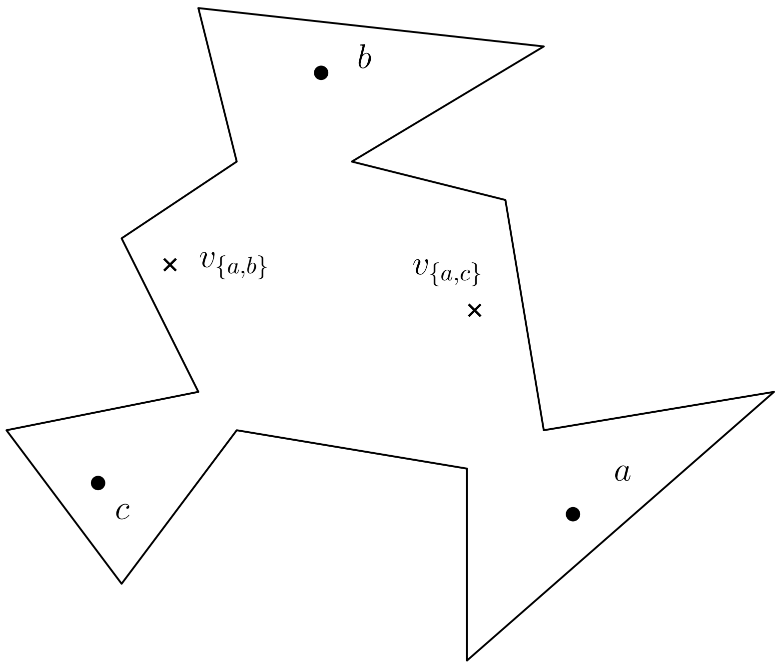

Figure 3.11: The view point that sees a and b but not c is called $v_{\{a,b\}}$. View points $v_{\{a,b\}}$ and $v_{\{a,c\}}$ are both in $V_{\{a\}}$.

path from x to some point y outside the convex hull of P that is completely contained in $\mathbb{R}^2 \setminus P$. But this path must cross at least one of the encircling visibility segments. A contradiction.

Note, that it does not suffice to have a cyclic polygonal chain starting and ending in p and going through w, but one has to ensure that the polygonal chain really encircles the segment. We will show later on how that can be achieved.

The kind of ranges by which we will replace visibility domains most often is wedges. Therefore, the following Theorem by Isler et al. [26] will become important. If we let $\mathcal{W}$ be the set of closed wedges in $\mathbb{R}^2$, it states that the VC-dimension of the range space $(\mathbb{R}^2, \mathcal{W})^*$ (i.e., the range space dual to $(\mathbb{R}^2, \mathcal{W})$) is at most 5.

Theorem 2. *(Isler et al.) For any arrangement of six or more wedges, there is a subset T of wedges for which no cell exists that is contained in exactly the wedges of T.*

For convenience, we include a short proof based on the ideas in [26].

Proof. By Euler's formula, an arrangement of n wedges has $n + k + 1$ many cells, where k denotes the number of half-line intersections. Since two wedges intersect in at most 4 points – in which case they are said to *cross* each other – we have $k \leq 4\binom{n}{2}$. Thus, an arrangement of 6 wedges has at most 67 cells. We are going to provide an accounting argument which shows that for each

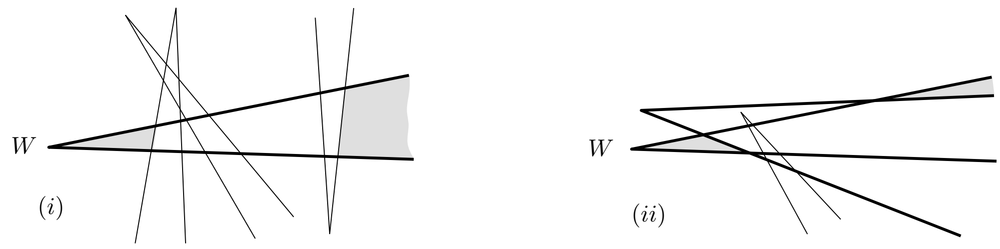

Figure 3.12: In (i) and (ii), respectively, the shaded cells are contained in wedge W only.

wedge one cell is missing from a maximum size arrangement (due to a shortage of intersections), or one of the existing cells is redundant (because it stabs the same subset of wedges as some other cell does). This will imply that at most $67 - 6 = 61$ many of all $2^6 = 64$ different subsets can be stabbed by a cell, thus proving the theorem.

If at least one of the two half-lines L_i, R_i bounding the wedge W_i intersects all the other half lines $L_j, R_j, j \neq i$, then the (bounded) cell in W_i incident to its apex and an infinite cell of W_i are stabbing the same subset of wedges, namely the singleton set $\{W_i\}$. So, if there is at most one cell stabbing $\{W_i\}$, then both half-lines L_i, R_i do not cross all the half-lines of all other wedges. Therefore there are at least two crossings missing compared to a maximally crossing arrangement. Let d denote the number of wedges W_i that contain at least two cells stabbing $\{W_i\}$ and let m denote the number of pairs of half-lines from different wedges that do not cross.

Then $m \geq 6 - d$, ($6 - d$ is the number of wedges that do not contain two singleton-cells. Every of the bounding half-lines of those wedges lacks at least one crossing and if we sum up the missing crossings of all the half-lines every crossing is counted at most twice). Plugging this into the formula above, we get that the arrangement has got $n + k + 1 = 67 - m \leq 61 + d$ many cells.

As the d additional cells are redundant, we have at most 61 combinatorially distinct cells. $\square$

3.3 Points on the Boundary

In this section we will start by considering the question how many points on the boundary of a simple polygon R can be shattered by visibility polygons of points in the whole of R. In this context we are able to present the decomposition technique used in Section 3.4 without having to care about intricate subtleties of boundaries of ranges and without having to invoke coarseness. On the other hand, we will see relativization to wedges as in the general case that is described in Section 3.5, but in a much less complicated setting. We will also explain the encircling arguments in great detail and try to make very clear how the encircling arguments enable the relativization to simpler ranges. So this section may be seen as preparatory for the later sections on VC-dimensions. On the other hand, the result proven is interesting in its own right and does not follow from the results in sections 3.4 and 3.5. Most of the material in this section has been published in [21].

We will prove the following Theorem.

Theorem 3. *No set of 14 points on the boundary of of a simple polygon can be shattered by interior visibility regions.*

We first explain the main idea of the proof of Theorem 3. We shall illustrate how Corollary 2 will be put to work in this proof.

Kalai and Matoušek [27] obtained a finite upper bound on the VC-dimension of interior visibility domains by proving that, for shattered point sets beyond a certain size, the impossible configuration depicted in Figure 3.13 (i) would occur.

We shall use similar configurations, but in a different way. Let us consider a situation like in Figure 3.13 (ii), where the view points w', w, w'' all can see the points l, r and a radial sweep around any of w', w, w'' would hit the points l, p, r in the same order. If w' and w'' can also see point p, then w must be able to see p, too, as this figure illustrates. Consequently, those view points of $V_{\{l,r\}}$ that are able to see p form a contiguous subsequence, in radial order around p. Our proof of Theorem 3 draws on this property, that will be formally defined in Definition 9 below. But how can we guarantee this property? To ensure that segment $\overline{pw}$ is indeed encircled by visibility segments, as shown in Figure 3.13 (ii), it suffices to show that the following condition is fulfilled. Points l and w' lie on one side of the line $\mathcal{L}(p, w)$ through p and w, and r and w'' lie on the other side of this line. Then we obtain polygonal chains $p - w' - l - w$ and $p - w'' - r - w$ that consist of visibility segments and that together encircle the segment $\overline{pw}$, therefore guaranteeing that p must see w. If the points l, p, r lie on the boundary of the polygon, this condition is ensured by providing a line L that separates $\{l, p, r\}$ from the view points.

Now we formally define the property.

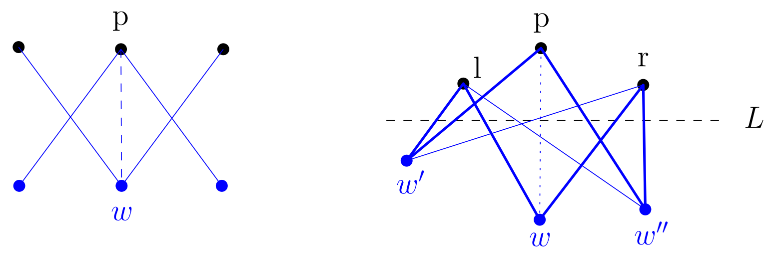

Figure 3.13: In (i), segment $\overline{pw}$ cannot be accessed by the boundary of R which, by assumption, encircles these six points. Hence, this configuration is impossible. For the same reason, the dotted segment in (ii) connecting p to w must be a solid visibility segment.

Definition 8. A (closed) *half plane* is a set $H \subset \mathbb{R}^2$ of the form $H = \{p \in \mathbb{R}^2 | \langle p, q \rangle \geq r\}$ for some $q \in \mathbb{R}^2$ and $r \in \mathbb{R}$. A (closed) *wedge* $W \subset \mathbb{R}^2$ is a set of the form $H_1 \cap H_2$ for two closed half planes H_1, H_2.

We will usually define a wedge W rather by its apex and two segments lying on the rays that bound W or just by two nonparallel segments on its boundary starting in its apex than by providing two half planes.

Definition 9. Let R' be a subset of R. We call a subset R'' of R' *wedge-like relative to* R', if it is the intersection of R' with some wedge $W \subset \mathbb{R}^2$.

Now we will give a showcase example for a situation in which the visibility polygon of a point is wedge-like relative to some subset of P.

Lemma 6. *Let l, p, r, L, and Q_p be situated as shown in Figure 3.14 (i) or (ii). Let V be a finite point set that lies completely in wedge Q_p that is defined by lines through p and l and through p and r, respectively. Then* $\text{vis}(p)$ *is wedge-like relative to* $\text{vis}(l) \cap \text{vis}(r) \cap V$.

Proof. Let W_p be the smallest subwedge of Q_p that contains all points of $V \cap \text{vis}(p)$. Suppose that w' and w'' lie on the two rays that bound W_p and that $\{w', w, w''\} \subseteq \text{vis}(l) \cap \text{vis}(r) \cap V \subset Q_p$ appear in counterclockwise order around p, as shown in Figure 3.15. Then $w', w'' \in \text{vis}(p)$ holds, that is, w' and w'' can see l, p, r. If $w \in V$ then segment $\overline{pw}$ is contained in an interior domain of the cycle

$$p - w' - l - w - r - w'' - p.$$

We observe that this fact is independent of the position of w' with respect to the line through l, w. Similarly, it does not matter if w'' lies to the left or to the right of the line through w and r. Hence, w sees p, that is, $w \in V \cap \text{vis}(p)$. This proves that $(\text{vis}(l) \cap \text{vis}(r) \cap V) \cap W_p \subseteq (\text{vis}(l) \cap \text{vis}(r) \cap V) \cap \text{vis}(p)$. On the other hand, there can be no point $w \in \text{vis}(l) \cap \text{vis}(r) \cap V$ that is in $V \cap \text{vis}(p)$ that is not in W_p as by definition of W_p, W_p contains all these points. □

Notice, that for our proof it does not matter, if $\text{vis}(l) \cap \text{vis}(r) \cap V$ is empty. Notice further, that in Figure 3.14 the points of V are separated from $\{l, p, r\}$ by a line L, but that this is not necessary for the proof. It is, however, essential for our proof that all points of $\text{vis}(l) \cap \text{vis}(r) \cap V$ lie in Q_p. Intuitively, this ensures that all points of $\text{vis}(l) \cap \text{vis}(r) \cap V$ see l, p, r in the same radial order. Figure 3.16 shows an example where this condition is not met and $\text{vis}(p)$ is not wedge-like relative to $\text{vis}(l) \cap \text{vis}(r) \cap V$.

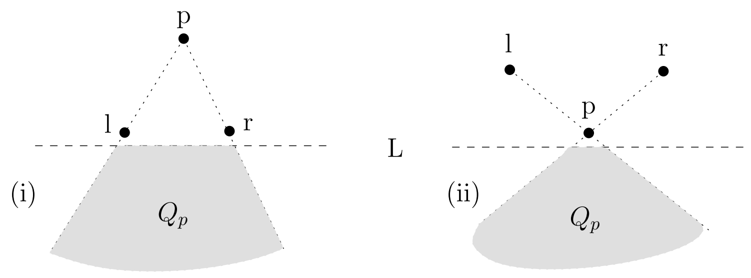

Figure 3.14: Two configurations addressed in Lemma 6.

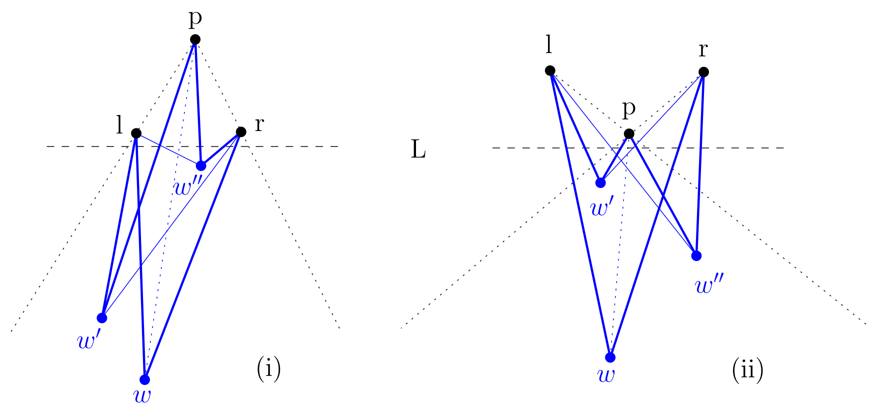

Figure 3.15: In either case, segment $\overline{pw}$ is encircled by visibility edges. Thus, w sees p.

Now we are ready to give the proof of Theorem 3. We will assume that there are 14 points on the boundary of a simple polygon that can be shattered by visibility domains of points in R and derive a contradiction.

Proof. Let R be a simple polygon, let S be a set of 14 points on the boundary D of R.

We choose two points $l, r \in S$ such that the two boundary segments, τ and β, of R between l and r contain 6 points of S each. Let $T := \tau \cap S$ and $B := \beta \cap S$. We may assume that the line L through l and r is horizontal; see Figure 3.17 for a sketch.

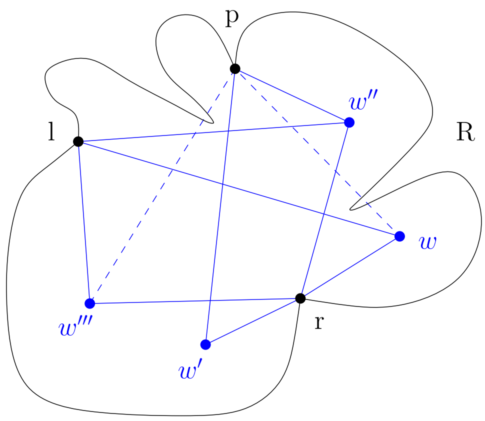

Figure 3.16: Here the view points seeing l, p, r do not appear consecutively around p.

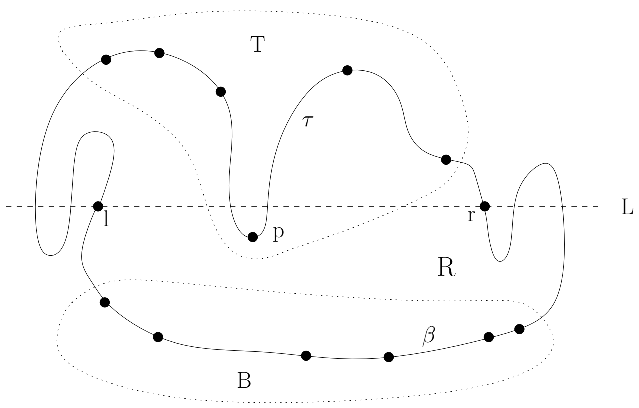

Figure 3.17: Some notations used.

Now let us assume that S can be shattered by the set $\mathcal{V} = \{\text{vis}(w)|\ w \in R\}$ of visibility regions of points in R. By Lemma 5 this implies the following. The set $T \cup B$ is shattered by the ranges of $\mathcal{V}_{\{l,r\}}$, the set of visibility polygons of points in R that see l and r. In other words, for each $U \subseteq T \cup B$ there is a view point $w_U \in R$ such that $\text{vis}(w_U) \cap S = \{l, r\} \cup U$.

Now we define two subsets of $\mathcal{V}_{\{l,r\}}$, namely P' and P'', where $P' = \{\text{vis}(w)|w$ lies below $L\}$ and $P'' = \{\text{vis}(w)|w$ lies above $L\}$ (both sets including visibility polygons of points lying *on* L). By Corollary 2, T is shattered by P' or B is shattered by P''. Let us w.l.o.g. assume that T is shattered by P'. That means that for each $U' \subseteq T$ there exists a view point $w \in R$ below line L that sees l and r such that $U' = T \cap \text{vis}(w)$.

Now we fix a set of view points from P', i.e. we choose for every $U' \subset T$ one point $v_{U'} \in R$ that

(i) lies *below* line L

(ii) sees l and r, and

(iii) sees of T exactly the subset U'.

We then define

$$V := \{v_{U'}|U' \subseteq T\} \quad \text{and, for all } p \in T, \quad V_p := \{v_{U'} \in V|p \in U'\}.$$

Thus, from each view point in V_p at least l, r and p are visible. By Lemma 3 it follows that the set of visibility polygons of points in T can be stabbed by points in V. Now, our main task is in proving the following fact.

Lemma 7. *For each $p \in T$, $\text{vis}(p)$ is wedge-like relative to V, i.e. there exists a wedge W_p (with apex p) such that $V_p = W_p \cap V$ holds.*

Proof. We distinguish two cases.
Case 1: Point p lies above line L. We define the wedge W_p by p and the two half-lines from p through those view points $v_l, v_r \in V_p$ that maximize $\measuredangle(v_l, p, v_r)$. Clearly, direction "$\subseteq$" of our lemma holds by definitions of W_p and V_p.

Since p and v_l, v_r are separated by line L, both visibility segments $\overline{pv_l}$ and $\overline{pv_r}$ must cross L. We claim that both crossing points, c_l, c_r must be situated on L between l and r. In fact, all other orderings would lead to a contradiction. Namely, if r was lying in between c_l and c_r, then r could not be contained in the boundary of R, which cannot intersect visibility segments. In Figure 3.18 the combinatorially different configurations are depicted. If both c_l and c_r were to the right of r, then v_l and v_r could not be contained in R; see Figure 3.19. This settles the claim: c_l, c_r must indeed be between l and r, see Figure 3.20. We are in the situation depicted in Figure 3.14 (i), and conclude from Lemma 6 that $\text{vis}(p)$ is wedge-like relative to V, and the proof of Case 1 is complete.
Case 2: p lies below line L. Let Q_p denote the wedge with apex p defined by the half lines from l and r through p; see Figure 3.21 (i). We claim that $V_p \subset Q_p$

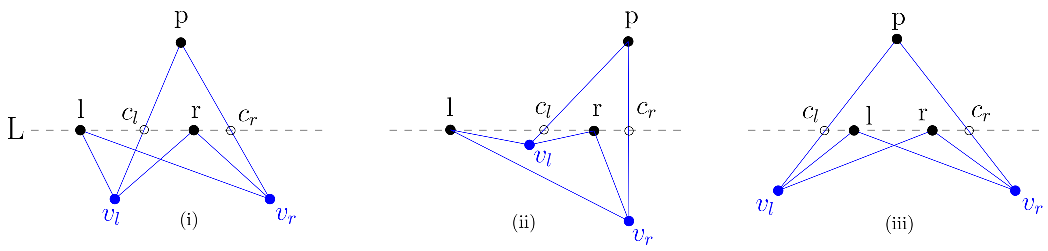

Figure 3.18: If r was between c_l and c_r it could not be reached by the boundary of R, which passes through p.

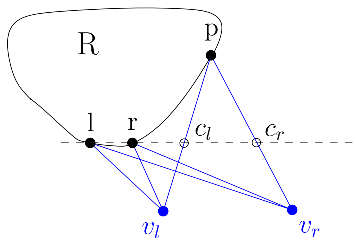

Figure 3.19: If c_l, c_r were to the right of r then the boundary of R, which visits l, r, p in counterclockwise order, could not encircle v_l, v_r.

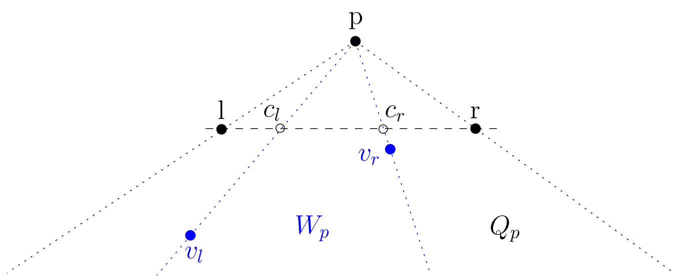

Figure 3.20: View point set V is contained in wedge $W_p \subset Q_p$.

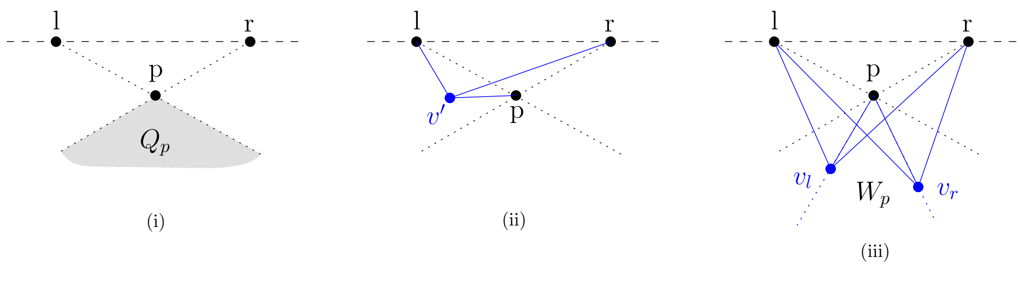

Figure 3.21: (i) Wedge Q_p is defined by p and the half lines from l, r through p. (ii) As the boundary of R visits l, r, p in counterclockwise order, it could not encircle v'. (iii) Wedge W_p is defined by the view points v_l, v_r of V_p that have a maximum angle at p.

holds. Indeed, if some point $v' \in V_p$ were situated outside Q_p then R could not contain v', because the boundary of R visits l, r, p in counterclockwise order; see Figure 3.21 (ii).

Now let v_l, v_r be those view points in V_p that maximize $\measuredangle(v_l, p, v_r)$. As in Case 1, let W_p be the wedge with apex p defined by the half lines from p trough v_l, v_r, respectively; see Figure 3.21 (iii). We have $W_p \subseteq Q_p$ and are in the situation depicted in Figure 3.14 (ii). As in Case 1, Lemma 6 implies that vis(p) is wedge-like relative to V. This completes the proof of Lemma 7. $\square$

Now let $U' \subseteq T$ and $p \in T$. By Lemma 7 we obtain the following equivalence for the view point $v_{U'} \in V$ of U'.

$$v_{U'} \in W_p \iff v_{U'} \in V_p \iff p \in U'.$$

That is, for each subset U' of T there exists a point (namely, $v_{U'}$) that is contained in exactly those wedges W_p where $p \in U'$. But this is impossible because of $|T| = 6$. A contradiction to Theorem 2. This concludes the proof of Theorem 3. $\square$

We can interpret this result as a result about the VC-dimension of an appropriately chosen range space. For the boundary D of our polygon R we can consider the set of ranges $\mathcal{V}_D = \{\text{vis}(p) \cap D\}_{p \in R}$ and the resulting range space $(D, \mathcal{V}_D)$. By Theorem 3 it follows that the VC-dimension of this range space is at most 13. As mentioned in the Introduction, every result about the VC-dimension of a range space leads to an upper bound on the size of ε-nets for this space and this in turn can be converted into a theorem about art gallery guarding. Let C denote the constant from the Epsilon-Net Theorem.

Theorem 4. *If each point in a simple polygon R sees at least an rth part of the the polygon's boundary D, then R can be covered by $13 \cdot C \cdot r \log r$ many guards on the boundary.*

Proof. With the notation from above and Theorem 3 we infer that the VC-dimension of $(D, \mathcal{V}_D)$ is at most 13. By the Epsilon-Net Theorem, there exists an $1/r$-net $N \subset D$ of size at most $13 \cdot C \cdot r \log r$. By assumption, each set in $\mathcal{V}_D$ contains a point p of N. That is, the visibility polygon of each $w \in R$ contains some point $p \in D$ and therefore each $w \in R$ is visible from some point $p \in N$ in the set of guards N. $\square$

3.4 Points and Guards on the Boundary

3.4.1 Perimeter Visibility Domains

Let P be a simple polygon with boundary B. Remember that for a point $p \in P$ its *visibility polygon* $\mathrm{vis}(p)$ is the set of points v such that the whole segment $\overline{pv}$ is contained in P. Now we restrict our attention to the portions of visibility polygons on the boundary, $\mathrm{vis}(p) \cap B$. For every $p \in B$ we will call this boundary portion its *perimeter visibility domain* and denote it by $V(p)$. In this section we are only concerned with this kind of visibility domains, so we will refer to them simply as *visibility domains*, below. We are now interested in the VC-dimension of the range space $(B, \mathcal{V})$ with $\mathcal{V} = \{V(b)\}_{b \in B}$.

As in the section before, we will first assume that a point set S of a certain size can be shattered. Then we will show that the visibility domains of most of these points look like much simpler ranges if they are relativized to the right subset of B. This time the simpler ranges are not wedges but certain subsets of B that we call *intervals*.

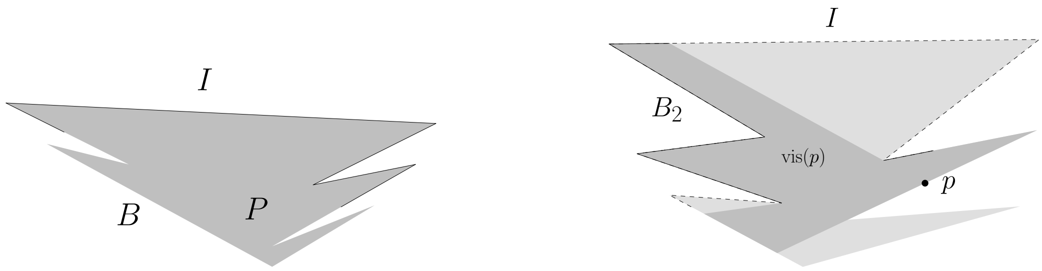

Figure 3.22: *Left:* I is an interval in B. *Right:* B_2 is interval-like relative to $\mathrm{vis}(p)$ as it is its intersection with I.

Definition 10. A subset I of B is called an *interval in* B if there is a continuous injective function $\pi \colon [0,1] \longrightarrow B$ with image I or if $I = \emptyset$.

Definition 11. Let B_1 be a subset of B. We call a subset B_2 of B *interval-like relative to* B_1, if its intersection with B_1 equals the intersection of B_1 with some interval in B.

The key geometric insight that we will need is formulated in the following lemma:

Lemma 8. *Let $a_1, a_2 \in B$ be two distinct boundary points such that $B \setminus \{a_1, a_2\}$ splits into two connected components C and D. Then for every point $c \in C$, $V(c)$ is interval-like relative to $D \cap V(a_1) \cap V(a_2)$.*

Proof. We have to find an interval I in B such that $I \cap D \cap V(a_1) \cap V(a_2)$ equals $V(c) \cap D \cap V(a_1) \cap V(a_2)$.

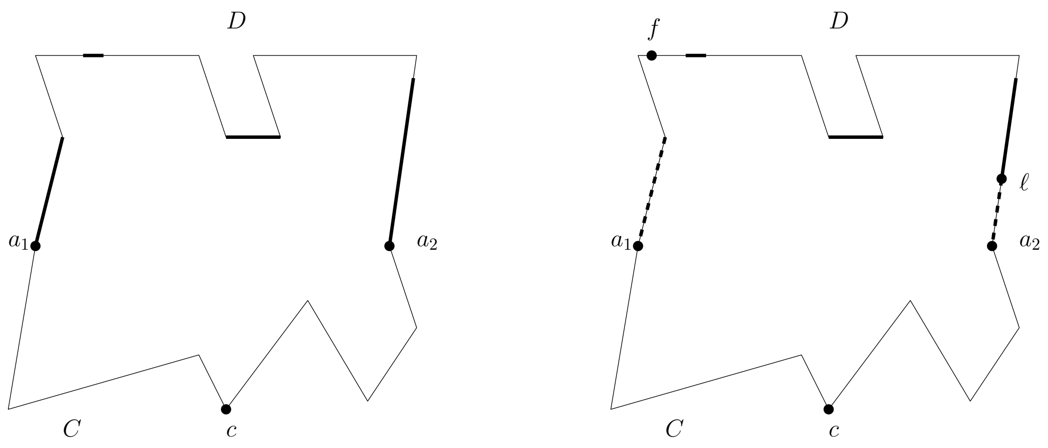

Figure 3.23: *Left:* The intersection of the visibility domains of a_1 and a_2 and D. *Right:* The parts of $V(a_1) \cap V(a_2) \cap D$ that are also in $V(c)$ (not dashed) are exactly the ones between the two points f and ℓ (that define an interval).

Let to this end $\pi\colon [0,1] \longrightarrow D \cup \{a_1, a_2\}$ be a continuous bijective map with $\pi(0) = a_1$ and $\pi(1) = a_2$.

If the intersection $D \cap V(c)$ of D with the visibility domain of c is empty, we set $I = \emptyset$. Otherwise there are values $f = \inf\{x \in (0,1) : \pi(x) \in V(c)\}$ and $\ell = \sup\{x \in (0,1) : \pi(x) \in V(c)\}$. We set I as the image of $[f, \ell]$ under π, $I = \pi[[f, \ell]]$. It remains to show that the two intersections $I \cap D \cap V(a_1) \cap V(a_2)$ and $V(c) \cap D \cap V(a_1) \cap V(a_2)$ are indeed equal.

By the definition of I it is clear that $V(c) \cap D \cap V(a_1) \cap V(a_2) \subset I \cap D \cap V(a_1) \cap V(a_2)$, as $V(c) \cap D \subset I \cap D$.

So let $v \in I \cap D \cap V(a_1) \cap V(a_2)$. Assume to a contradiction that $v \notin V(c) \cap D \cap V(a_1) \cap V(a_2)$ (i.e. c does not see v). The two visibility segments $\overline{a_1 v}, \overline{a_2 v}$ together with the boundary portion C bound a polygon P' that has no proper intersection with D (see Figure 3.4.1). Now take a look at the shortest path in P from c to v. By our assumption, the path bends at some point of B, because otherwise c would see v. As the visibility segments $\overline{a_1 v}, \overline{a_2 v}$ are straight line segments with endpoint v, the shortest path from c to v cannot properly cross one of these segments.

As c lies in P' it follows that the whole shortest path from c to v stays inside P'. Therefore the path does not bend at a point belonging to D. Consequently, it bends at some point of $C \cup \{a_1, a_2\}$.

Let $x \in C \cup \{a_1, a_2\}$ be the first such point on the shortest path from c to v. If we prolong the segment $\overline{cx}$ until it hits B at point y this prolongation cuts off a part of P and thereby cuts off a connected interval B_{xy} of B (the one between x and y) that contains v. If $y \in C \cup \{a_1, a_2\}$ it would follow that $D \cap V(c)$ was empty. That is because, by assumption, v is a point in D. If it lies between x and y, all points of D must lie between x and y as D is a connected part of

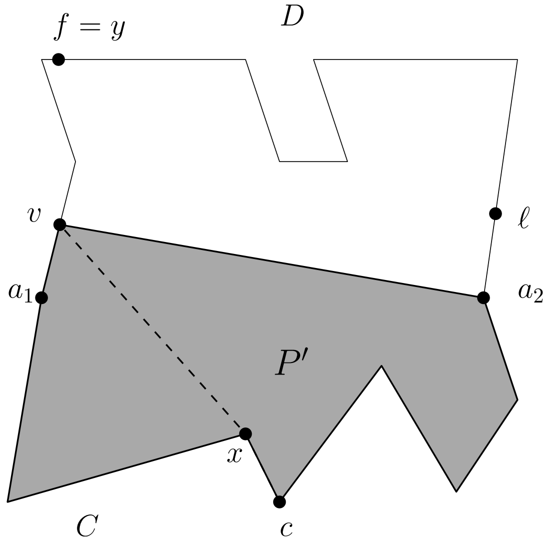

Figure 3.24: The segments to v from a_1 and a_2, respectively, together with the boundary portion C bound a polygon P' that contains the shortest path from c to v. So the shortest path does not bend at a point of D.

the boundary. So D would lie completely in B_{xy}. But that cannot be, because then $V(c) \cap D$ would be empty in contradiction to the fact that we are just considering the case where this intersection is nonempty.

So $y \in D$. As $x \in C \cup \{a_1, a_2\}$, the removal of x and y splits D into two parts and so one of a_1, a_2 lies in B_{xy}. Let w.l.o.g. a_2 lie in B_{xy}. There is some $t \in (0,1)$ with $\pi(t) = y$. The points of $\pi[(t,1]]$ all lie in B_{xy}. It follows that $y = \ell$ and v does not lie in I, a contradiction to our assumption.

This yields $V(c) \cap D \cap V(a_1) \cap V(a_2) \supset I \cap D \cap V(a_1) \cap V(a_2)$ and therefore $V(c) \cap D \cap V(a_1) \cap V(a_2) = I \cap D \cap V(a_1) \cap V(a_2)$. □

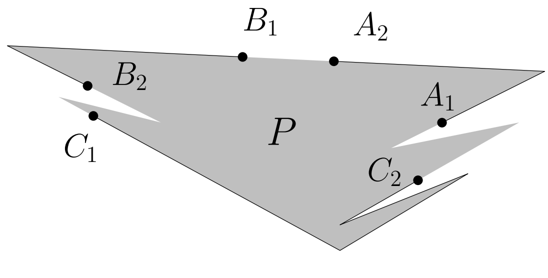

Figure 3.25: The Intervals $[A_1, A_2], [B_1, B_2], [C_1, C_2]$ divide the boundary of P into six cells. This number does not increase if the intervals intersect.

Together with the following lemma, this will yield our new upper bound.

Lemma 9. *Let $V(p_1), V(p_2), V(p_3)$ for three points $p_1, p_2, p_3 \in B$ be interval-like relative to some subset B' of B. Let $\mathcal{V}' = \{V(b')|b' \in B'\}$ be the set of visibility domains of points in B'. Then $\{p_1, p_2, p_3\}$ is not shattered by $\mathcal{V}'$. In particular, we can lower bound the coarseness by $c_{\mathcal{V}'}(\{p_1, p_2, p_3\}) \geq 2$.*

Proof. Let I_1, I_2, I_3 be intervals such that $I_j \cap B' = V(p_j) \cap B'$ for $j = 1, 2, 3$. The arrangement of the three intervals divides B in at most 6 cells. Every two points of B' that are in the same cell see the same subsets of $\{p_1, p_2, p_3\}$. As there are 8 subsets of $\{p_1, p_2, p_3\}$, at least two sets are missing and so $\{p_1, p_2, p_3\}$ is not shattered. □

3.4.2 The Upper Bound for the VC-Dimension

Theorem 5. *The VC-dimension of visibility domains on the boundary is at most* 7.

Proof. We will show that there can be no set of eight boundary points that is shattered by sets of $\mathcal{V}$.

Let $S = \{s_1, \ldots, s_8\} \subset B$ (where the points are in counterclockwise order along the boundary) is a set on the boundary of P. Assume to a contradiction that S is shattered by the set of visibility domains $\mathcal{V}$ of points in B.

Let us define the special points $l = s_1$ and $r = s_5$. Let C the part of the boundary between l and r traversed in counterclockwise direction, D be the part of B between r and l in counterclockwise direction. Let $C' = C \cap V(l) \cap V(r)$ and $D' = D \cap V(l) \cap V(r)$, i.e. the parts of C and D, respectively, that see both l and r. We split S into the sets $S_C = \{s_2, s_3, s_4\}$, $S_D = \{s_6, s_7, s_8\}$ and the two points l, r. We also define $\mathcal{V}_{C'} = \{V(b)\}_{b \in C'}$ and $\mathcal{V}_{D'} = \{V(b)\}_{b \in D'}$.

By the assumption that S is shattered by $\mathcal{V}$ and Lemma 5 we get that $(S_C \cup S_D)$ is shattered by $\mathcal{V}_{\{l,r\}} = \{V \in \mathcal{V}|l, r \in V\}$. If l sees r then $\mathcal{V}_{\{l,r\}} = \mathcal{V}_{C'} \cup \mathcal{V}_{D'} \cup \{\text{vis}(l), \text{vis}(r)\}$. Otherwise $\mathcal{V}_{\{l,r\}} = \mathcal{V}_{C'} \cup \mathcal{V}_{D'}$. Lemma 10 below shows that $\mathcal{V}_{C'} \cup \mathcal{V}_{D'}$ does not shatter $(S_C \cup S_D)$, so $\mathcal{V}_{\{l,r\}} = \mathcal{V}_{C'} \cup \mathcal{V}_{D'} \cup \{\text{vis}(l), \text{vis}(r)\}$. By Lemma 2 we get that $|\Pi_{\{\text{vis}(l),\text{vis}(r)\}}(S_C \cup S_D)| \geq c_{\mathcal{V}_{C'} \cup \mathcal{V}_{D'}}(S_C \cup S_D)$. Lemma 10 says that $c_{\mathcal{V}_{C'} \cup \mathcal{V}_{D'}}(S_C \cup S_D) \geq 4$ and so $|\Pi_{\{\text{vis}(l),\text{vis}(r)\}}(S_C \cup S_D)| \geq 4$. But $|\Pi_{\{\text{vis}(l),\text{vis}(r)\}}(S_C \cup S_D)| \leq 2$, as there are only two sets in $\{\text{vis}(l), \text{vis}(r)\}$. A contradiction.

Therefore S is not shattered by visibility domains of points in B. □

Only the following lemma remains to be shown.

Lemma 10. $c_{\mathcal{V}_{C'} \cup \mathcal{V}_{D'}}(S_C \cup S_D) \geq 4$.

Proof. By Lemma 8, $V(c)$ is interval-like relative to D' for every $c \in S_C$ and $V(d)$ is interval-like relative to C' for every $d \in S_D$. By Lemma 9, it follows that $c_{\mathcal{V}_{C'}}(S_D) \geq 2$ and $c_{\mathcal{V}_{D'}}(S_C) \geq 2$. By Lemma 4, $c_{\mathcal{V}_{C'} \cup \mathcal{V}_{D'}}(S_C \cup S_D) \geq 4$. □

This completes the proof. We have shown that the VC-dimension of perimeter visibility domains is at most 7. With the best known lower bound it follows that this VC-dimension is between 5 and 7.

3.5 The General Case

Most of the material in this section has been published in [23]. An extended abstract appeared in [22].

In this section we return to the original problem considered by Kalai and Matoušek. It is the most challenging case, as the points of the set that we want to shatter as well as the view points can lie anywhere in the simple polygon P. In particular, there is no obvious way to guarantee that view points are seen in the same cyclic order by all points of S.

In the former sections we chose a subset $S' = \{l, r\}$ of S and relativized the visibility domains of the other points to subsets of $\text{vis}(l) \cap \text{vis}(r)$. It was not hard to find the points l, r, as every pair of points of S that splits S into two equally large subsets would do. In contrast, in this section there is no canonical way anymore to find such a subset of S. We will show, that one can nevertheless find such a set and use our relativization technique to prove an upper bound that improves considerably on the best upper bound known before.

We first state the theorem. In this section, we are considering range spaces $(P, \mathcal{V})$ where our ground set is a simple polygon P and the set of ranges is $\mathcal{V} = \{\text{vis}(p) | p \in P\}$ the set of all visibility polygons of points in P. The *VC-dimension of visibility polygons in simple polygons* d is the supremum of the VC-dimensions of all range spaces $(P, \mathcal{V})$ taken over all possible simple polygons P.

Theorem 6. *For the VC-dimension d of visibility polygons in simple polygons, $d \leq 14$ holds.*

3.5.1 Proof Technique

Theorem 6 will be proven by contradiction. Throughout Sections 3.5.1 and 3.5.2, we shall assume that there exists a simple polygon P containing a set S of 15 points that is shattered by visibility polygons. That is, for each $T \subseteq S$ there is a view point v_T in P such that

$$T = \text{vis}(v_T) \cap S \tag{3.1}$$

holds, where, as usual, $\text{vis}(v) = \{x \in P;\ \overline{xv} \subset P\}$ denotes the visibility domain of a point v in the (closed) set P.

We may assume that the points in S and the view points v_T are in general position, by the following argument. If $p \notin T$, then segment $\overline{v_T p}$ is properly crossed by the boundary of P, that is, the segment and the complement of P have an interior point in common. On the other hand, a visibility segment $\overline{v_U q}$, where $q \in U$, can be touched by the boundary of P, because this does not block visibility. By finitely many, arbitrarily small local enlargements of P we can remove these touching boundary parts from the visibility segments without losing any proper crossing of a non-visibility segment. Afterwards, all points and view points can be perturbed in small disks.

By the symmetry of visibility, property 3.1 can be rewritten as

$$T = \{p \in S |\ v_T \in \text{vis}(p)\} \tag{3.2}$$

This means, if we form the arrangement Z of all visibility regions vis(p), where $p \in S$, then for each $T \subseteq S$ there is a cell (containing the view point v_T) which is contained in exactly the visibility regions of the points in T. To obtain a contradiction, one would like to argue that the number of cells in arrangement Z is smaller than 2^{15}, the number of subsets of S. But as we do not have an upper bound on the number of vertices of P, the complexity of Z cannot be bounded from above.

For this reason we shall again replace visibility domains with *wedges*; for we can use the upper bound of 5 for the dual VC-dimension of the range space of wedges, see Theorem 2. To illustrate the technique that we will be using this time, let S be a shattered set containing at least 15 points and let a be a point of S. We assume that there are

1. points b_1, b_2 of S,
2. a view point v that sees b_1 and b_2, but not a, such that
3. a is contained in the triangle defined by $\{v, b_1, b_2\}$;

see Figure 3.26 (i). We denote by U the wedge containing v formed by the lines through a and b_1 and b_2, respectively. Any view point w that sees b_1 and b_2 must be contained in wedge U. Otherwise, the chain of visibility segments $v - b_1 - w - b_2 - v$ would encircle the line segment $\overline{va}$ connecting v and a, preventing the boundary of P from blocking the view from v to a; see Figure 3.26 (ii).

Let w_1, w_2 denote the outermost view points in U that see a, b_1, b_2 and include a maximum angle (by assumption, such view points exist; by the previous reasoning, they lie in U). Then w_1, w_2 define a sub-wedge W of U, as shown in Figure 3.26 (iii). We claim that in this situation

$$V_{\{b_1,b_2\}} \cap \text{vis}(a) = V_{\{b_1,b_2\}} \cap W \tag{3.3}$$

holds, where $V_{\{b_1,b_2\}}$ denotes the set of all view points that see at least b_1 and b_2. Indeed, each view point that sees b_1, b_2 lies in U. If it sees a, too, it must lie in W, by definition of W. Conversely, let v' be a view point in W that sees b_1, b_2. Then line segment $\overline{v'a}$ is encircled by the visibility segments $v' - b_1 - w_1 - a - w_2 - b_2 - v'$, as depicted in Figure 3.26 (iv). Thus, $v' \in \text{vis}(a)$.

Fact 3.3 says that vis(a) is wedge-like relative to $V_{\{b_1,b_2\}}$. So if we were able to find two points like b_1 and b_2, and restrict ourselves to studying only those 2^{13} view points $V_{\{b_1,b_2\}}$ that see both b_1, b_2, as a benefit, the visibility region vis(a) behaves like a wedge and we can apply Theorem 2.

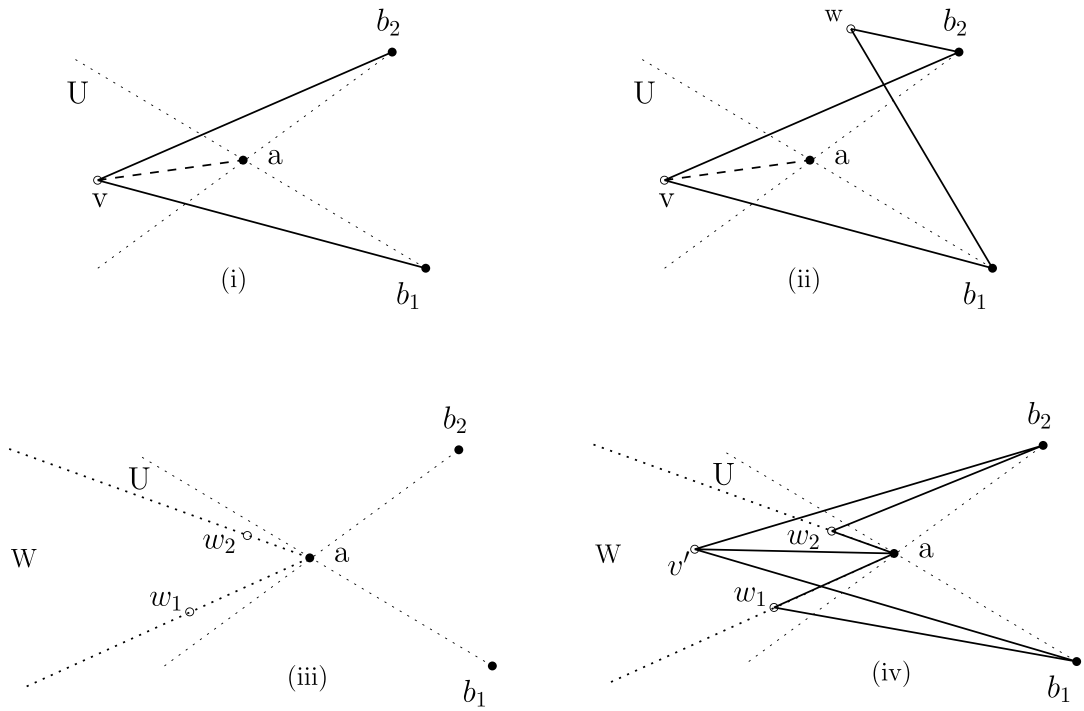

Figure 3.26: Solid lines connect points that are mutually visible; such “visibility segments” are contained in polygon P. Dashed style indicates that the line of vision is blocked; these segments are crossed by the boundary of P.

Arguments as above will be applied as follows. In Section 3.5.2 we prove, as a direct consequence, that at most 5 points can be situated in the *interior* of the convex hull of S. Then, in Section 3.5.3, we show that at most 9 points can be located on the *boundary* of the convex hull. Together, these claims imply Theorem 6.

3.5.2 Interior Points

The goal of this section is in proving the following fact.

Lemma 11. *At most five points of S can lie inside the convex hull of a shattered set S.*

Proof. Suppose there are at least six interior points a_i, $1 \leq i \leq 6$, in the convex hull. Each of the remaining points of S is a vertex of the convex hull of S. Let $b_0, \dots b_{m-1}$ be an enumeration of these points in cyclic order.

As S is shattered, for every subset $T \subseteq S$ there is a view point that sees exactly this subset. We fix a set V that contains for each subset of S one such view point.

Let v_B (where $B = \{b_0, \dots, b_{m-1}\}$) be the view point that sees only these vertices but no interior point; see Figure 3.27. Each interior point a_i is contained

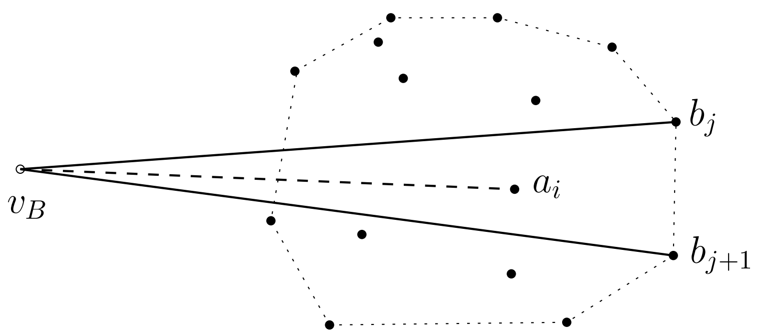

Figure 3.27: Each interior point a_i is contained in some triangle defined by $\{v_B, b_j, b_{j+1}\}$.

in a triangle defined by $\{v_B, b_j, b_{j+1}\}$, for some j (where the indices are taken modulo m). Since properties 1.–3. mentioned in Section 3.5.1 are fulfilled, Fact 3.3 implies that there exists a wedge W_i such that $V_{\{b_j, b_{j+1}\}} \cap \text{vis}(a_i) = V_{\{b_j, b_{j+1}\}} \cap W_i$ holds. If V_B denotes the set of view points that see at least the points of B, we obtain

$$V_B \cap \text{vis}(a_i) = V_B \cap W_i \text{ for } i = 1, \dots, 6,$$

which implies the following. For each subset T of $A = \{a_1, \dots, a_6\}$ the view point $v_{T \cup B}$ lies in exactly those wedges W_i where $a_i \in T$. So every subset of this set of six wedges is stabbed by a point in P. This is a contradiction to the fact that the dual VC-dimension of wedges is not larger than 5, see Theorem 2. Thus, the convex hull of S cannot contain six interior points. $\square$

Therefore, at least 10 points of S must be vertices of the convex hull of S.

3.5.3 Points on the Boundary of the Convex Hull

Ignoring interior points, we prove, in this section, the following fact.

Lemma 12. *Let S be a set of 10 points in convex position inside a simple polygon, P. Then S is not shattered by visibility polygons inside P.*

Proof. Again, the proof is by contradiction. So let S be a set of 10 points in convex position inside a simple polygon P. Assume that S is shattered. Let V be a fixed set of view points that contains for every subset T of S one point that sees exactly T.

First, we enumerate the points around the convex hull. (Notice, that the edges of the convex hull of S may intersect the boundary of P.) Let E denote the set of even indexed points. Let v_E be the view point that sees exactly the even indexed points. If v_E lies outside the convex hull, $\mathrm{ch}(S)$, of S, we draw the two tangents from v_E at $\mathrm{ch}(S)$. The points between the two tangent points facing v_E are called *front points*, all other points are named *back points* of S; see Figure 3.28. In other words, the back points are exactly the points of S that lie on the boundary of the convex hull of $S \cup \{v_E\}$. If $v_E \in \mathrm{ch}(S)$ then all points of S are called back points.

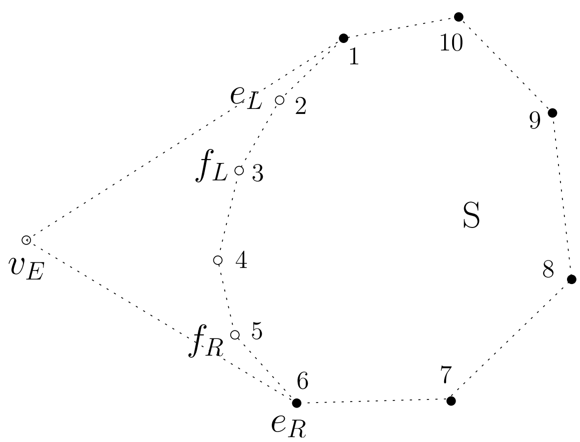

Figure 3.28: Front points appear in white, back points in black. View point v_E sees exactly the points of even index.

We are going to discuss the case depicted in Figure 3.28 first, namely:

Case 1: There exists an odd-indexed front point.

It follows from the definition of front points that in this case v_E lies outside the convex hull of S. Let f_L and f_R be the outermost left and right front points with odd index, from the position of v_E; and let e_L and e_R denote their outer neighbors, as shown in Figure 3.28. While $f_L = f_R$ is possible, we always have $e_L \neq e_R$. Observe that e_L and e_R may be front or back points; this will require some case analysis later on.

Notation. For two points a, b, let $H^+(a,b)$ denote the open half-plane to the left of the ray $L(a,b)$ from a through b, and $H^-(a,b)$ the open half-plane to its right.

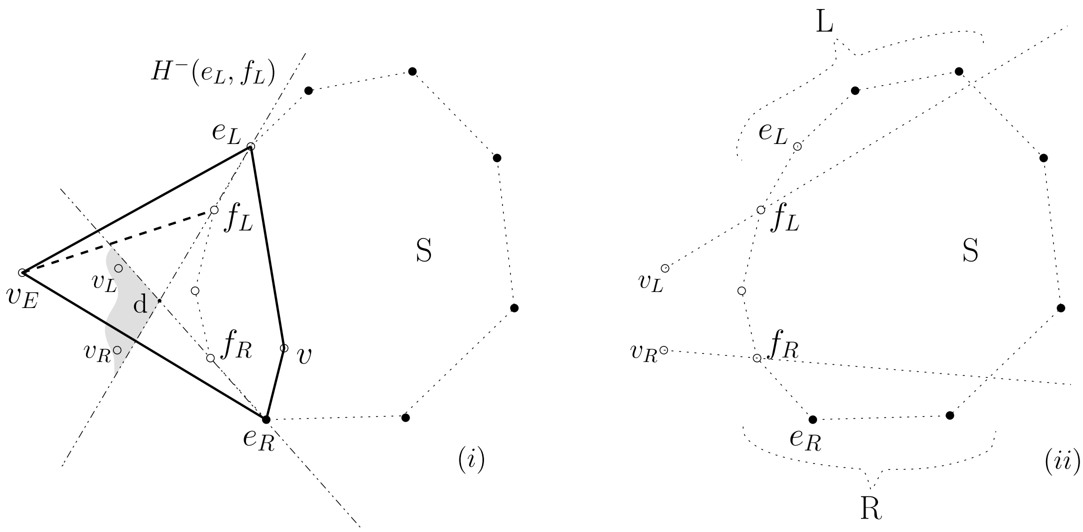

Figure 3.29: (i) As segment $\overline{v_E f_L}$ must be intersected by the boundary of P, it cannot be encircled by visibility segments. (ii) Defining subsets L and R of S.

Claim 1. *Each view point v that sees e_L and e_R lies in $H^-(e_L, f_L) \cap H^-(f_R, e_R)$.*

Proof. If v were contained in $H^+(e_L, f_L)$ then the chain of visibility segments $e_L - v - e_R - v_E - e_L$ would encircle the segment $\overline{v_E f_L}$—a contradiction, because v_E does not see the odd indexed point f_L; see Figure 3.29 (i). So $v \in H^-(e_L, f_L)$. A similar argument implies $v \in H^-(f_R, e_R)$. □

We now define two subsets L and R of S that will be crucial in our proof.

Definition 12. (i) Let $v_L := v_{S\setminus\{f_L\}}$ and $v_R := v_{S\setminus\{f_R\}}$ denote the view points that see all of S except f_L or f_R, respectively.
(ii) Let $L := S \cap H^+(v_L, f_L)$ and $R := S \cap H^-(v_R, f_R)$.

By Claim 1, the points of S contained in the triangle (e_R, e_L, v_E) are front points with respect to v_R, v_L, too; see Figure 3.29.

Claim 2. *None of the sets $L, R, S \setminus (L \cup R)$ are empty. The sets L and R are disjoint.*

Proof. By construction, we have $e_L \in L$, $e_R \in R$, and $f_L, f_R \notin L \cup R$. If $v_L = v_R$ (because f_L could be f_R) then $L \cap R = \emptyset$, obviously. Otherwise, there is at least one even indexed point, e, between f_L and f_R on $\text{ch}(S)$. Assume that there exists a point q of S in the intersection of L and R. Then segment $\overline{v_R f_R}$ would be encircled by the visibility chain $q - v_R - e - v_L - q$, contradicting the fact that v_R sees every point *but* f_R; see Figure 3.30. □

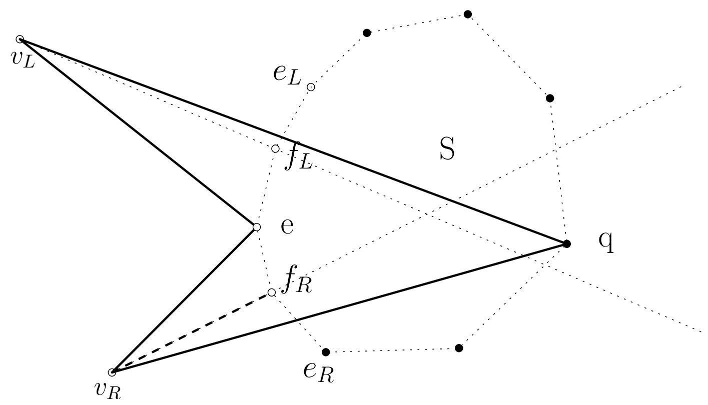

Figure 3.30: L and R are disjoint.

The purpose of the sets L and R will now become clear: They contain points like b_1, b_2 in Section 3.5.1, that help us reduce visibility regions to wedges. The precise property will be stated for R in Lemma 13 below; a symmetric property holds for L. The proof of Lemma 13 will be postponed. First, we shall derive a conclusion in Lemma 15, and use it in completing the proof of Lemma 12 in Case 1.

In the next lemmas, D^c denotes the complement of a set D.

Lemma 13. *There exist points r_1, r_2 in R such that the following holds either for $Q = vis(r_1) \cap vis(r_2)$ or for $Q = vis(r_1)^c \cap vis(r_2)$. For each $p \in S$ different from r_1, r_2, each view point that (i) sees p, (ii) lies in Q, and (iii) sees at least one point of L, is contained in the half-plane $H^-(p, r_2)$.*

And analogously:

Lemma 14. *There exist points l_1, l_2 in L such that the following holds either for $Q' = vis(l_1) \cap vis(l_2)$ or for $Q' = vis(l_1)^c \cap vis(l_2)$. For each $p \in S$ different from l_1, l_2, each view point that (i) sees p, (ii) lies in Q, and (iii) sees at least one point of R, is contained in the half-plane $H^-(l_2, p)$.*

From Lemma 13 and Lemma 14 we obtain the following conclusion.

Lemma 15. *Let $p \in S \setminus \{l_1, l_2, r_1, r_2\}$. Then each view point in $Q \cap Q'$ that sees p lies in the wedge $U_p = H^-(p, r_2) \cap H^-(l_2, p)$.*

Now we can proceed as in Section 3.5.1; see Figure 3.26 (iii) and (iv). Within wedge U_p we find a sub-wedge W_p satisfying

$$Q \cap Q' \cap \mathrm{vis}(p) = Q \cap Q' \cap W_p, \tag{3.4}$$

with the same arguments that led to Fact 3.3, replacing (a, b_1, b_2) with (p, r_2, l_2). Since membership in Q, Q' only prescribes the visibility of $\{l_1, l_2, r_1, r_2\}$, Fact 3.4 implies the following. For each subset $T \subseteq S \setminus \{l_1, l_2, r_1, r_2\}$ there exists a cell in the arrangement of the remaining six wedges W_p, where $p \in S \setminus \{l_1, l_2, r_1, r_2\}$, that is contained in precisely the wedges related to T. As in Section 3.5.2, this contradicts Theorem 2 and proves Lemma 12 in Case 1. □

It remains to show how to find r_1, r_2 and Q in Lemma 13.

Proof. (of Lemma 13) Before starting a case analysis depending on properties of R and e_R we list some helpful facts.

Claim 3. *If a view point v sees a point $r \in R$ and a point $s \notin R \cup \{f_R\}$ then $v \in H^-(s, r)$. A symmetric claim holds for L.*

Proof. Otherwise, $\overline{v_R f_R}$ would be encircled by $r - v - s - v_R - r$, since f_R lies in the triangle defined by (v_R, r, s); see Figure 3.31 (i). □

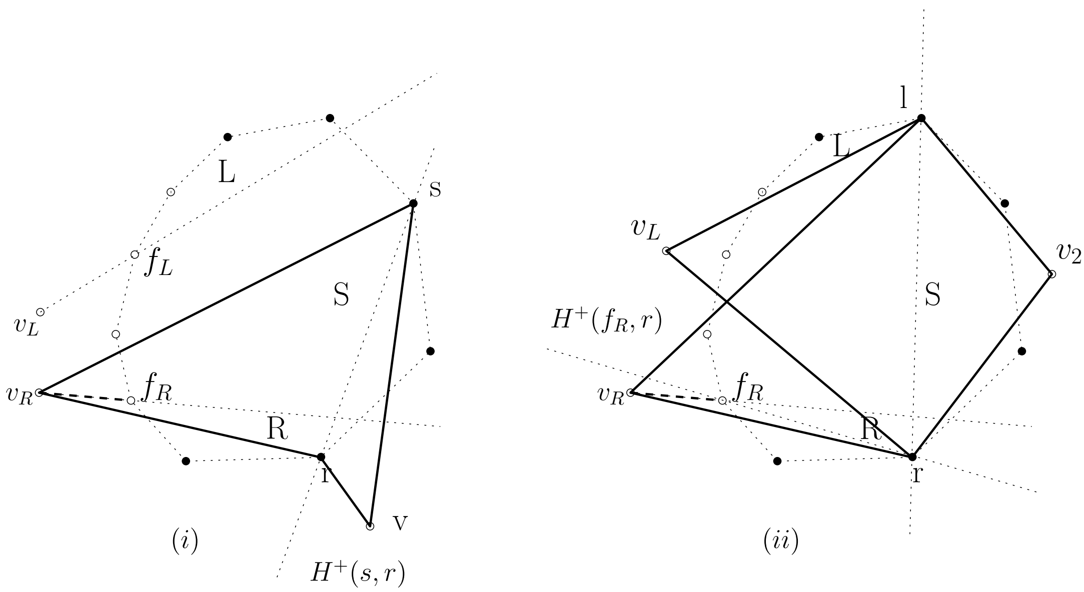

Figure 3.31: Illustration to Claims 3 and 4.

The next fact narrows the locus from which two points, one from L and R each, are visible.

Claim 4. *If a view point v sees points $r \in R$ and $l \in L$ then v lies in the wedge $H^-(f_R, r) \cap H^-(l, f_L)$, and on the same side of $L(r, l)$ as v_R and v_L do.*

Proof. If $v \in H^+(f_R, r)$, or if v were situated on the opposite side of $L(r, l)$, then $\overline{v_R f_R}$ would be encircled by $r - v - l - v_R - r$; see points $v = v_1$ and $v = v_2$ in Figure 3.31 (ii). □

Now we start on the case analysis. In each case, we need to define $r_1, r_2 \in R$ and a set $Q = \text{vis}(r_1) \cap \text{vis}(r_2)$ or $Q = \text{vis}(r_1)^c \cap \text{vis}(r_2)$. Then we must prove that the following assertion of Lemma 13 holds.

Assertion

If $p \in S$ is different from r_1, r_2, and if $v \in Q$ is a view point that sees p and some point $l \in L$, then $v \in H^-(p, r_2)$.

Case 1a: Point set R contains at most two points.
We define $\{r_1, r_2\} := R$ and let $Q := \text{vis}(r_1) \cap \text{vis}(r_2)$.
Let p and v be as in the Assertion. If $p \neq f_R$ then Claim 3 implies $v \in H^-(p, r_2)$.
If $p = f_R$ we obtain $v \in H^-(p, r_2)$ by the first statement in Claim 4.

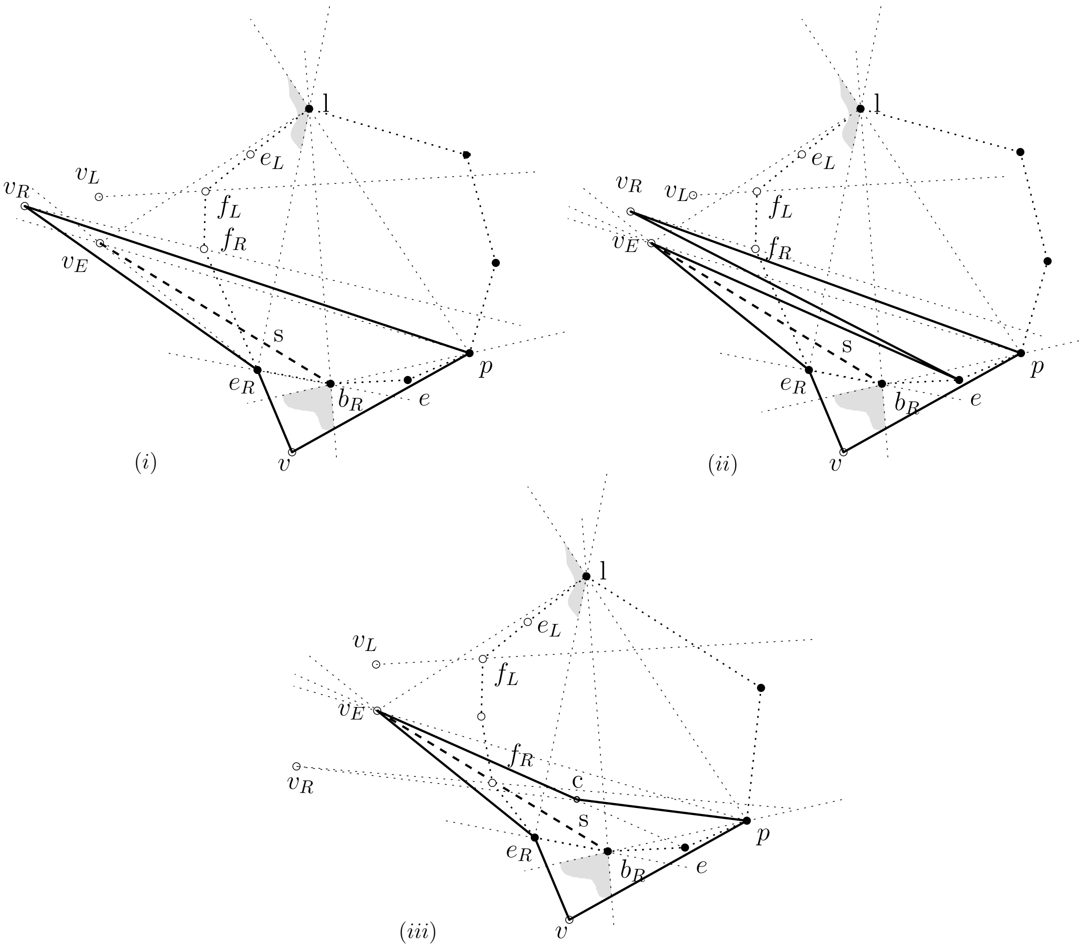

Figure 3.32: Illustrations of Case 1b.

Case 1b: Point set R contains more than two points, and e_R is tangent point of $\text{ch}(S)$ as seen from v_E; compare Figure 3.28.
We set $r_1 := e_R$ and let r_2 be the odd indexed back point b_R counterclockwise next to e_R. Moreover, $Q := \text{vis}(r_1) \cap \text{vis}(r_2)$.
For each $p \notin R$ the proof of Case 1a applies. Let $p \in R$ be different from r_1, r_2.

Assume, by way of contradiction, that $v \in H^+(p, r_2)$ holds. Since the second statement of Claim 4 implies $v \in H^-(l, e_R) \cap H^-(l, p) \subset H^-(l, b_R)$, we obtain $v \in H^-(l, b_R) \cap H^+(p, b_R)$; see Figure 3.32. Now we discuss the location of view point v_R. If it lies in the wedge $H^+(e_R, v_E) \cap H^+(v_E, p)$ then segment $s := \overline{v_E b_R}$ is encircled by $e_R - v_R - p - v - e_R$; see Figure 3.32 (i). If v_R does not lie in this wedge, let e be the counterclockwise neighbor of b_R in R. If v_R lies on the same side of $L(e, v_E)$ as p, then $e_R - v_E - e - v_R - p - v - e_R$ protects segment s; see (ii). If it lies on the opposite side, then $\overline{v_E e}$ intersects $\overline{v_R p}$ at some point c, and $e_R - v_E - c - p - v - e_R$ encircles segment s; see (iii). In either situation, we obtain a contradiction.

Before continuing the case analysis we prove a simple fact.

Lemma 16. *Let a, b, c denote the vertices of a triangle, in counterclockwise order. Suppose there exists a view point w in $H^+(b, a) \cap H^-(c, b)$ that sees a and c. Then, each view point $v \in H^+(b, a)$ that sees a and c but not b lies in $H^-(c, b)$.*

Proof. Otherwise, segment $\overline{vb}$ would be encircled by $c - v - a - w - c$; see Figure 3.33. □

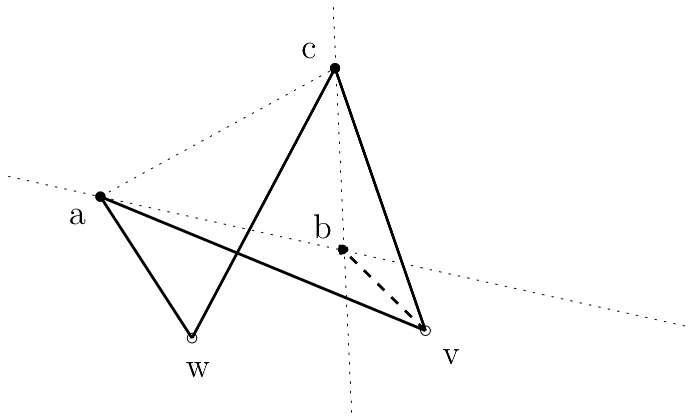

Figure 3.33: Proof of Lemma 16

Case 1c: Point set R contains more than two points, and the counterclockwise neighbor, b_R, of e_R, is tangent point as seen from v_E. Let e denote the counterclockwise neighbor of b_R, and let $w_R := v_{S \setminus \{e_R\}}$ denote the view point that sees all of S except e_R. We consider three subcases, depending on the location of w_R.

(1ci) If $w_R \in H^-(b_R, e_R)$, we set

$$(r_1, r_2, Q) := (e_R, b_R, \text{vis}(e_R)^c \cap \text{vis}(b_R)).$$

To prove the Assertion, let $p \neq e_R, b_R$, and let v be a view point that sees p, b_R, l but not e_R, for some $l \in L$.

As both w_R and v see $l \in L$ and $b_R \in R$, Claim 4 implies $w_R, v \in H^-(f_R, b_R) \cap H^-(l, f_L) \subset H^+(e_R, l)$. The latter inclusion allows us to apply Lemma 16 to $(a, b, c, w) = (l, e_R, b_R, w_R)$, which yields $v \in H^-(b_R, e_R)$. Now $v \in H^-(p, b_R)$ follows; see Figure 3.34 (i).

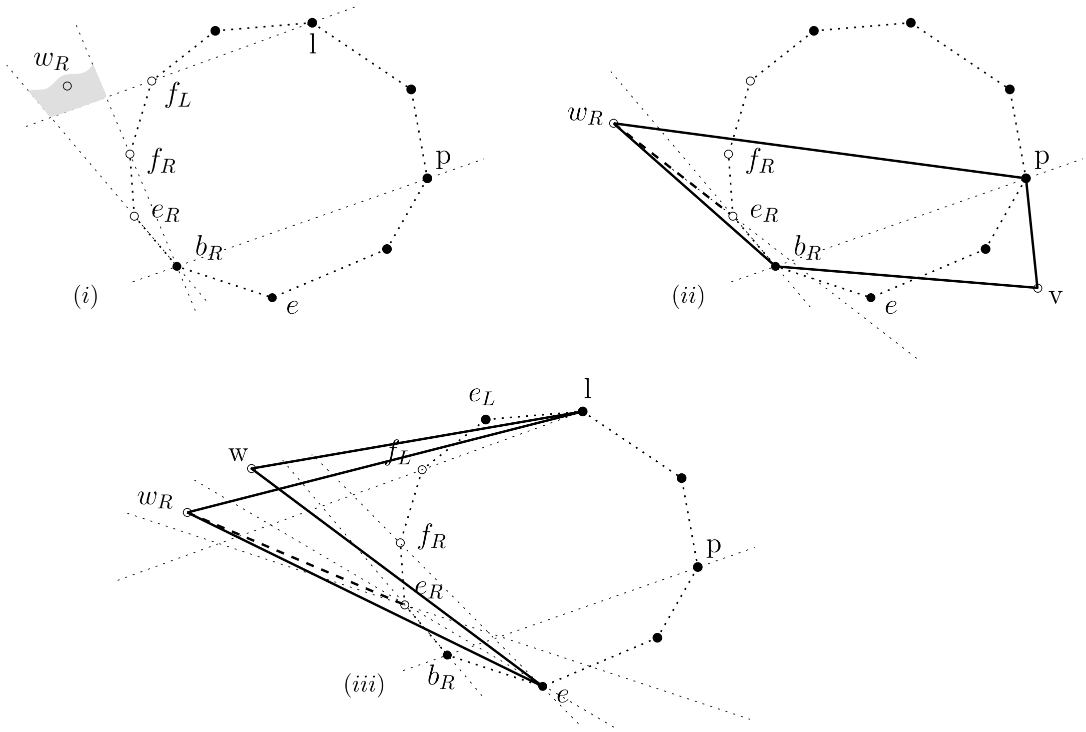

Figure 3.34: Illustrations of Case 1c.

(1cii) If $w_R \in H^+(b_R, e_R)$, and if b_R and e are situated on opposite sides of $L(w_R, e_R)$, we set $(r_1, r_2, Q) := (e_R, b_R, \text{vis}(e_R) \cap \text{vis}(b_R))$.
All points of $S' := S \setminus \{e_R, b_R\}$ lie on the same side of $L(w_R, e_R)$ as e. A view point v that sees some point $p \in S'$ and b_R must be in $H^-(p, b_R)$. Otherwise, $\overline{w_R e_R}$ would be encircled by $b_R - w_R - p - v - b_R$; see Figure 3.34 (ii).

(1ciii) If $w_R \in H^+(b_R, e_R)$, and if b_R and e are situated on the same side of $L(w_R, e_R)$, we set $(r_1, r_2, Q) := (b_R, e, \text{vis}(b_R)^c \cap \text{vis}(e))$.
Clearly, $w_R \in H^+(e, e_R)$. Each view point w that sees e and some $l \in L$ —in particular point v of the Assertion— must lie in $H^-(e_R, e)$, or $\overline{w_R e_R}$ would be enclosed by $w_R - l - w - e - w_R$; see Figure 3.34 (iii) (observe that w must be contained in $H^-(f_R, e) \cap H^-(l, f_L) \subset H^+(e, l)$, by Claim 4, as depicted in the figure).

Let x denote the view point that sees exactly e_R, e, e_L, l. By Claim 4, $x \in H^+(e_R, e_L) \cap H^+(e, e_L) \subset H^+(b_R, e_L)$. We file for later use that $x \in H^+(b_R, l)$ holds, for the same reason. Since b_R is tangent point from v_E, we have $v_E \in H^-(e, b_R)$. Thus, we can apply Lemma 16 to $(a, b, c, w) = (e_L, b_R, e, v_E)$ and obtain $x \in H^-(e, b_R)$.

We have just shown that $x \in H^+(b_R, l) \cap H^-(e, b_R)$ holds. Moreover, Claim 4 implies $v \in H^-(f_R, e) \cap H^-(l, f_L) \subset H^+(b_R, l)$ since v sees l and e. Since v does not see b_R we can apply Lemma 16 to $(a, b, c, w) = (l, b_R, e, x)$ and obtain $v \in H^-(e, b_R)$. Together with the first finding in (1ciii), this implies $v \in H^-(e_R, e) \cap H^-(e, b_R) \subset H^-(p, e)$ for all $p \neq b_R, e$.

This completes the proof of Lemma 12 in Case 1.

□

Now we discuss the second case of Lemma 12, thereby completing its proof. This also completes the proof of our main result, Theorem 6.

Case 2: There is no odd front point.

In this situation, view point v_E either lies inside $ch(S)$, so that no front point exists, or v_E lies outside $ch(S)$, and at most one front point is visible from v_E between the two tangent points on $ch(S)$; if so, its index is even.

Independently of the position of v_E, we introduce some notation. Let v_S denote the point that sees all points in S. The line G through v_E and v_S divides S into two subsets, L and R (not to be confused with L and R in case 1), one of which may possibly be empty. We cut G at v_E, and rotate the half-line passing through v_S over L; see Figure 3.35. The first and the last odd indexed points of L encountered during this rotational sweep are named l_1 and l_2, respectively. Similarly, r_1 and r_2 are defined in R.

We observe that, e. g., l_1 and l_2 need not exist, or that $l_1 = l_2$ may hold; these cases will be taken care of in the subsequent analysis. Also, the half-line rotating about v_E may cut through S in its start position, depending on the position of v_S. This is of no concern for our proof, which is literally the same for either situation.

Lemma 17. *If there are odd-indexed points in both L and R, then exactly one point lies between l_1 and r_1 on the boundary of the convex hull of S. This point has even index and will be called e_1. Similarly, there is exactly one even-indexed point between l_2 and r_2, called e_2.*

Proof. Since l_1, r_1 are both odd indexed points, there is at least one point between them on $\partial ch(S)$. If there was more than one point on $\partial ch(S)$ between l_1 and r_1, one of them would have an odd index. Let w.l.o.g o lie on the same side of $L(v_E, v_S)$ as l_1. Being odd, l_1 and o must be back points, since no odd front points exist in Case 2. As the order in which back points of L are encountered by the rotating rays coincides with their order on the boundary of the convex hull, o would have to be hit by the rotating ray before l_1, contradicting the definition of l_1. The same argument applies to l_2 and r_2. □

We will deal with a somewhat special subcase first.

Case 2A: Both of the following properties hold.

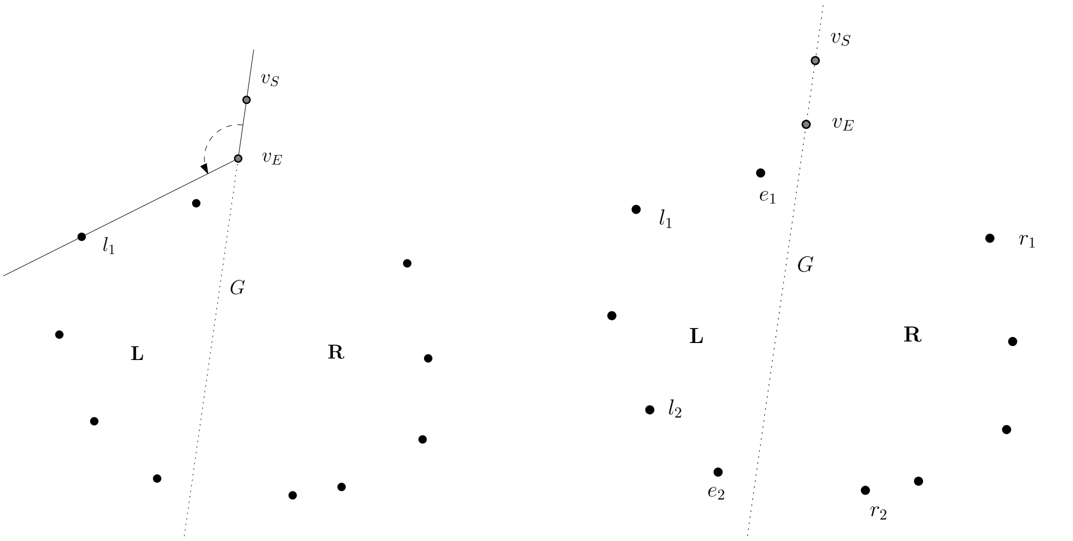

Figure 3.35: The half-line is rotated about v_E over L. The first odd point encountered is named l_1, the last one l_2.

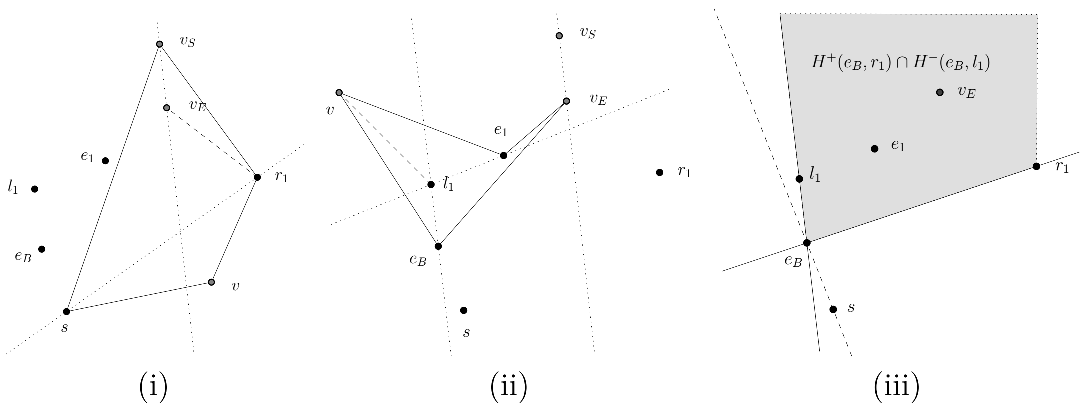

Figure 3.36: The proof of Case 2A.

1. One of L, R contains exactly one point of S; its index is odd.
2. Point e_1 is a front point with respect to v_E.

In this case, v_E must lie outside $ch(S)$, and e_1 is the only front point with respect to v_E, so that l_1 and r_1 are tangent points, as seen from v_E. Moreover, v_E lies in the triangle given by l_1, r_1 and v_S because otherwise e_1 would be a back point.

W.l.o.g., let r_1 be the only point in R, and let L be situated to the right of the directed line from v_S through v_E; see Figure 3.36. Then every view point v that sees r_1 and some point s of L lies on the same side of $L(r_1, v_E)$ as v_S does, or $\overline{r_1 v_E}$ would be encircled by $r_1 - v - s - v_S - r_1$. For the same reason every view point v that sees r_1 and some point s of L lies on the same side of $L(s, r_1)$ as v_S does, see Figure 3.36 (i).

Also, r_1 does not see any point $s \in S$, apart from itself, otherwise $\overline{r_1 v_E}$ would be encircled by $r_1 - s - v_S - r_1$.

Now let e_B be the even-indexed neighbor of l_1 that is a back point, and let us set $Q = \text{vis}(e_B) \cap \text{vis}(e_1) \cap \text{vis}(r_1) \cap \text{vis}(l_1)^c$.

Next, we want to show that every view point $v \in Q$ lies in $H^-(e_B, l_1)$; see Figure 3.36 (ii). If this were wrong, $v \in H^-(r_1, e_B) \cap H^-(r_1, e_1) \subset H^-(r_1, l_1)$ would imply $v \in H^+(l_1, e_1)$. Since v_E obviously lies in $H^+(l_1, e_1) \cap H^-(e_B, l_1)$, we could apply Lemma 16 to $(a, b, c, w) = (e_1, l_1, e_B, v_E)$, and obtain $v \in H^-(e_B, l_1)$—a contradiction.

Now let us assume that, in addition to being in Q, view point v sees a point $s \in S \setminus \{l_1, r_1, e_1, e_B\}$. As v lies in $H^-(r_1, e_B)$, it follows that $v \in H^-(s, e_B) \supset H^-(r_1, e_B) \cap H^-(e_B, l_1)$, see (iii).

On the other hand, v also lies in $H^-(r_1, s)$ as already shown.

Summarizing, we have obtained a result analogous to Lemma 15.

Lemma 18. *Let $s \in S \setminus \{l_1, r_1, e_1, e_B\}$. Then each view point in Q that sees s lies in the wedge $U_s = H^-(s, e_B) \cap H^-(r_1, s)$.*

Now the proof of Case 2A is completed by exactly the same arguments used subsequently to Lemma 15 in Section 3.5.3.

If one of the properties of Case 2A is violated, we obtain the following, by logical negation.

Case 2B: At least one of the following properties holds.

1. None of L, R is a singleton set containing an odd indexed point.
2. Point e_1 is a back point, as seen from v_E.

Other than in the previous cases, we will now reduce visibility regions to *half-planes*, rather than to wedges. We will show the existence of three points, p_1, p_2, p_3 in S, and of a half plane H_i for each, such that the following holds. Let Q denote the set of view points that see at least $S \setminus \{p_1, p_2, p_3\}$. Then,

$$\text{for each } v \in Q: \text{ for each } i = 1, 2, 3: \quad v \text{ sees } p_i \iff v \in H_i. \tag{3.5}$$

Property 3.5 leads to a contradiction, due to the following analogon of Theorem 2.

Lemma 19. *For any arrangement of three (or more) half-planes, there is a subset T of half-planes for which no cell is contained in exactly the half-planes of T.*

Proof. With three half-planes, we have eight subsets, but at most seven cells. □

While this fact is easier to prove, and somewhat more efficient, as we need only three points to derive a contradiction, it is harder to find points fulfilling Property 3.5. This will be our next task. Again, we consider points in L and points in R separately. Let us discuss the situation for L.

We start by defining two points, l_1' and e'. Suppose there is a point with odd index in R. We set $l_1' = l_1$. As we are in Case 2B, point e_1—situated between l_1 and r_1—is a back point, or there is some point e with even index in R. In the first case we set $e' := e_1$, in the latter case we set $e' := e$; see Figure 3.37 (i) and (ii).

If there is no point with odd index in R then there are five points with odd index in L. We then set l_1' to be the second point with odd index that was hit during the rotation of the half-line from v_E through v_S. Then l_1 and l_1' are distinct back points with respect to v_E (since there is no odd front point). Between l_1 and l_1' on the boundary of the convex hull there lies exactly one point e that has even index. We set $e' = e$. In this case, there are three points with even index on the convex hull between l_1' and l_2. Notice that e' is a back point with respect to v_E; see Figure 3.37 (iii).

In either case the points l_1' and l_2 have odd indices, and the point e' has even index and is either a back point with respect to v_E, or it lies in R.

We will now prove the following.

Lemma 20. *For all back points p with even index that lie in the wedge given by the rays from v_E through l_1' and l_2 the following holds. There is a half-plane H_p such that every view point v that sees l_1' and l_2 sees p if and only if $v \in H_p$. The analogue holds if we replace l_1' and l_2 by r_1' and r_2.*

Before we prove Lemma 20, we first use it to derive the following consequence. As explained before, it provides us with a contradiction, thus proving Case 2B of Lemma 12 and completing all proofs.

Lemma 21. *There are three points $p_1, p_2, p_3 \in S$ and half-planes H_1, H_2, H_3 that satisfy Property 3.5.*

Proof. If there are odd points in both L and R, then there is exactly one even point between l_1' and r_1' and one even point between l_2 and r_2 and all other even points lie between the rays from v_E through l_1' and l_2 and through r_1' and

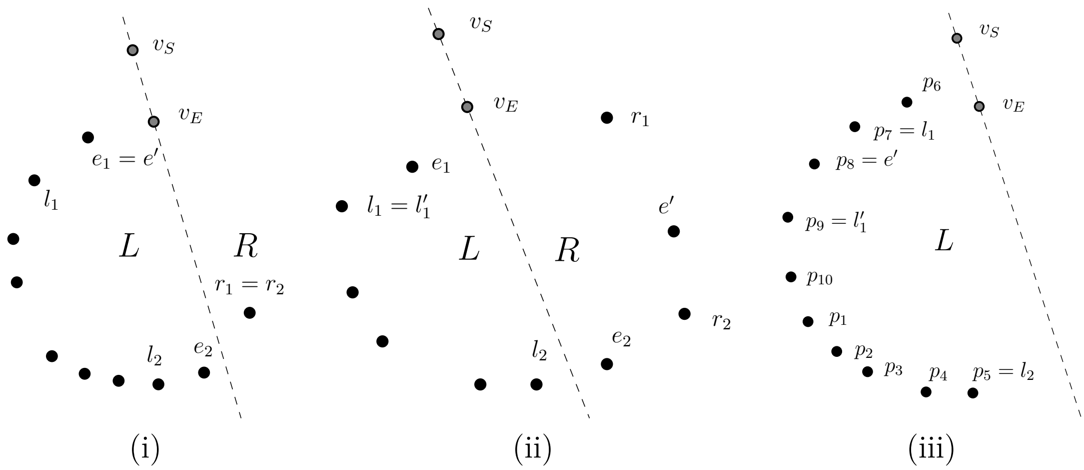

Figure 3.37: (i) If e_1 is a back point with respect to v_E we set $e' = e_1$. (ii) Otherwise there is an even indexed point in R we will call e'. (iii) If there is no (odd) point in R then l'_1 is the second odd indexed point and e' is the (back) point between l_1 and l_2.

r_2, respectively. By Lemma 20, we get that the remaining three even-indexed points have the desired property.

If there is no odd point in R or in L, then there are four even-indexed points between l_1 and l_2 or between r_1 and r_2 and therefore there are three points with the desired property between l'_1 and l_2 or between r'_1 and r_2. $\square$

Proof. To prove Lemma 20 let $e \in S$ be a point with even index that lies between l'_1 and l_2. Points e and v_S lie on opposite sides of $L(l'_1, v_E)$, by the definition of l'_1.

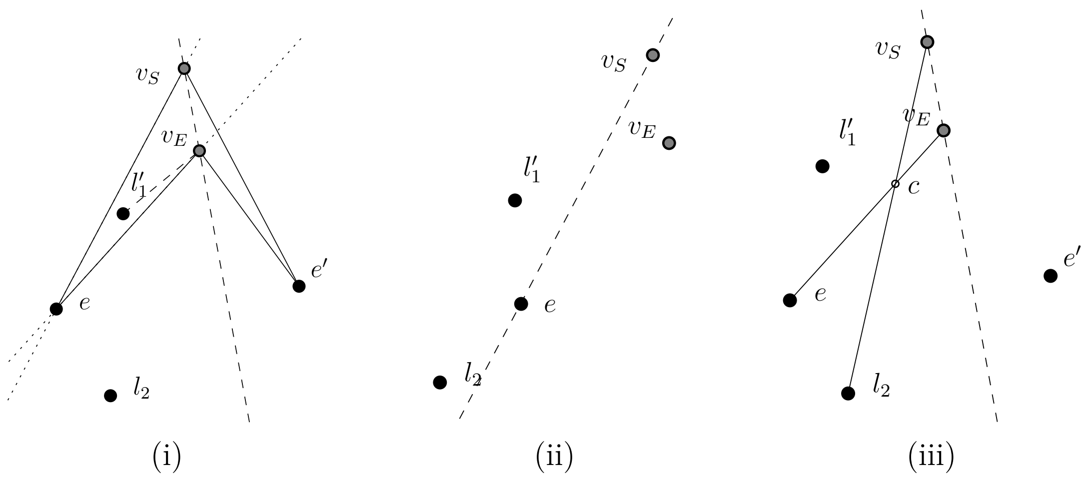

Figure 3.38: (i) l'_1 cannot lie between the rays $L(e, v_S)$ and $L(e, v_E)$. (ii) l'_1 and l_2 can not lie on the same side of $L(e, v_S)$. (iii) So there must be an intersection between $\overline{ev_E}$ and $\overline{l_2 v_S}$.

Claim 5. *In this situation the segments $\overline{ev_E}$ and $\overline{l_2 v_S}$ intersect in a point c.*

Proof. The segment $\overline{l_2 v_S}$ intersects the line $L(e, v_E)$ by definition of l_2. It remains to show that l_2 does neither lie on the side of $L(v_S, v_E)$ opposite to e nor on the side of $L(v_S, e)$ opposite to v_E. As l_2 and e both belong to L, the first assertion follows. For the second one, notice that l'_1 cannot lie on the same side of $L(e, v_S)$ as v_E does because otherwise $\overline{l'_1 v_E}$ would be encircled by $e - v_E - e' - v_S - e$, see Figure 3.38 (i). But l'_1 and l_2 cannot both lie on the side of $L(e, v_S)$ opposite to v_E: Because l'_1, l_2, e are backpoints, v_E, l'_1, e and l_2 are the corners of a convex quadrilateral. If l'_1 and l_2 lie on the same side of a line through e, this line must be a tangent to this quadrilateral and therefore l'_1, l_2 and v_E would have to lie on the same side of this line, see (ii). So the segment $\overline{ev_E}$ crosses $\overline{l_2 v_S}$ in a point c. ☐

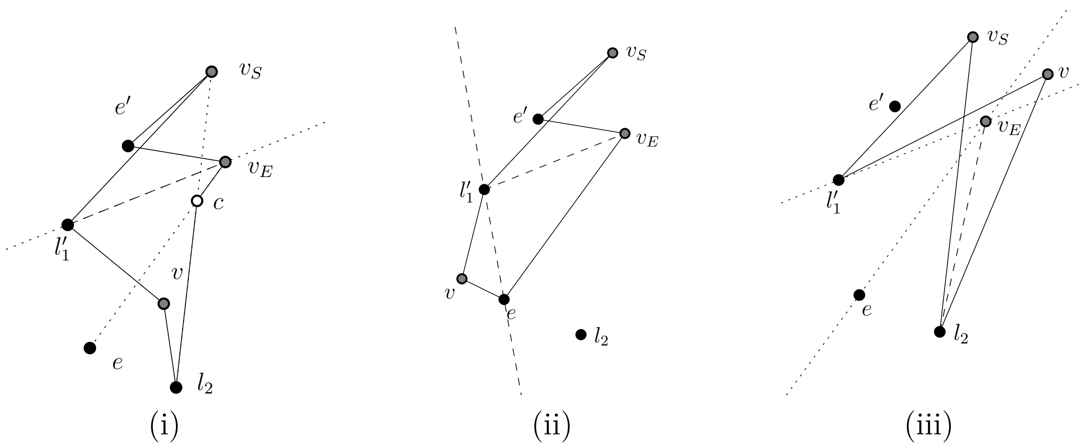

Figure 3.39: (i) v and v_S must lie on the same side of $L(l'_1, v_E)$. (ii) v and v_S must lie on the same side of $L(e, l'_1)$. (iii) v and v_S must lie on the same side of $L(e, v_E)$.

Now it follows that every view point v that sees l'_1 and l_2 lies on the same side of $L(l'_1, v_E)$ as v_S does, because otherwise the segment $\overline{l'_1 v_E}$ would be encircled by $v_E - c - l_2 - v - l'_1 - v_S - e' - v_E$, see Figure 3.39 (i).

It also follows that every view point that sees l'_1, l_2 and e has to lie on the same side of $L(e, l'_1)$ as v_S does, because otherwise $\overline{l'_1 v_E}$ would be encircled by $v_E - e' - v_S - l'_1 - v - e - v_E$, see Figure 3.39 (ii).

Claim 6. *Every view point v that sees l'_1 and l_2 lies on the same side of $L(e, v_E)$ as v_S does.*

Proof. Assume v and v_S lay on opposite sides of $L(e, v_E)$. We already showed that v must lie on the same side of $L(l'_1, v_E)$ as v_S does. So $\overline{l_2 v_E}$ would be encircled by $l_2 - v - l'_1 - v_S - l_2$, see Figure 3.39 (iii). ☐

Lemma 22. *All view points v that see $\{l_1', l_2, e\}$ lie in the wedge W given by the two rays originating in e and going through l_1' and l_2, respectively*

Proof. We just showed that all such view points v lie in the wedge W_e given by the two rays originating in e and going through v_E and l_1', respectively. As W_e is a subset of W, the lemma follows. □

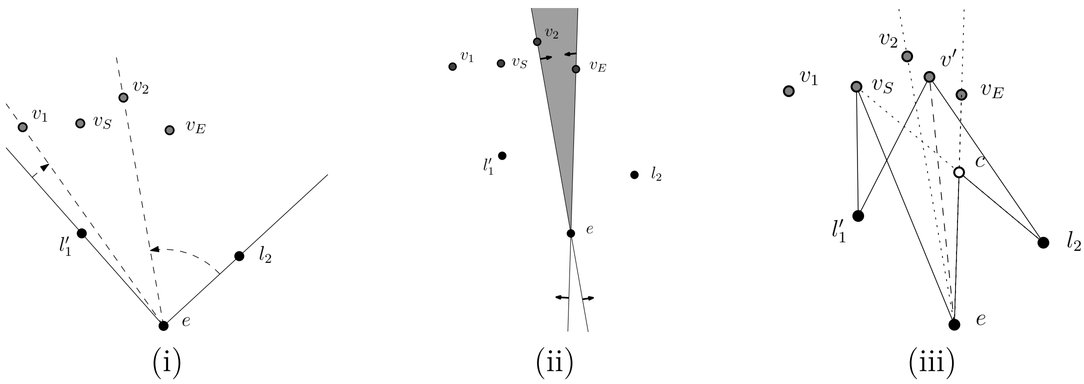

Figure 3.40: (i) We rotate the rays through l_1' and l_2 until they encounter v_1 and v_2. (ii) The area between $L(e, v_E)$ and $L(e, v_2)$. (iii) No point v' that lies in this area sees l_1' and l_2 but not e.

Let us now rotate the ray with origin e through l_2 over the wedge W, towards l_1'. (Notice that, by Claim 5, the rotating ray passes over v_E before it reaches v_S.) Let us denote the first view point we encounter that sees l_1', l_2 and e by v_2. Let us then rotate the ray with origin e through l_1' over the wedge W, towards l_2. Let us denote the first view point we encounter this time that sees l_1', l_2 and e by v_1, see Figure 3.40 (i).

We now obtain the two following facts.

Claim 7. *All view points that see l_1', l_2 and e lie in the wedge originating in e and going through v_1 and v_2.*

Proof. By Lemma 22 we know that all such points lie between l_1' and l_2. By construction of v_1 and v_2 there is no such point between l_1 and v_1 or between l_2 and v_2. □

Lemma 23. *There is no view point on the side of $L(e, v_2)$ opposite to v_S that sees l_1' and l_2 but not e.*

Proof. By Claim 6, all view points that see l_1' and l_2 lie on the same side of $L(e, v_E)$ as v_S does. As v_2 also lies on this side of the line and moreover inside the wedge between the rays from e through v_E and v_S, respectively, it follows, that a point that sees l_1' and l_2 and that lies on the side of $L(e, v_2)$ opposite to v_S must lie in the wedge given by the rays from e through v_E and v_2, which in turn is contained in the wedge between the rays from e through v_E and v_S, see Figure 3.40 (ii).

Now assume there was a view point v' in this wedge, that saw l_1' and l_2 but not e. If we take c to be the intersection of $\overline{ev_E}$ and $\overline{l_2 v_S}$, then the segment $\overline{ev'}$ would be encircled by $e - c - l_2 - v' - l_1' - v_S - e$, see Figure 3.40 (iii).

$\square$

Now we are able to complete the proof of Lemma 20.

We define H_e to be the closed half plane to the side of the line through e and v_1 in which v_2 lies. By Claim 7 all view points that see l_1', l_2 and e lie in H_e. Assume now there was a view point v in H_e that sees l_1' and l_2 but not e. By Lemma 23 and the assumption that v lies in H_e, it follows that then v must lie in the wedge with origin e and rays through v_1 and v_2.

This again leads to a contradiction because the segment $\overline{ev}$ then would be encircled by $e - v_1 - l_1' - v - l_2 - v_2 - e$. So a view point v that sees l_1' and l_2 sees e if and only if $v \in H_e$. $\square$

Now all proofs are complete.

4 Barrier Resilience of Visibility Domains

4.1 Introduction

4.1.1 Motivation

Put yourself in a smuggler's shoes. You want to deliver some goods to a fixed destination but you do not want to be seen by many witnesses. Unfortunately, there is no way to your destination that is completely unobserved, nor can you conceal your goods. Perhaps you just want to minimize the number of witnesses or perhaps there is some number k of witnesses that still is acceptable.

It is not important to you, how often or how long the witnesses see you on your way. You only care for their number.

You are given a map of your city, in which your starting and your target point are marked as well as the position of all the law-abiding people, see Figure 4.1. Can you compute the path that is seen by the minimum number of witnesses?

This turns out to be a special case of the Barrier Resilience problem. Given a start point s and a target point t as well as the positions and ranges of n sensors that are designed to detect intruders, we want to find a path from s to t that minimizes the number of its witnesses (i.e. the sensors that detect the agent traveling on this path, see Figure 4.3 for an instance of the Barrier Resilience problem for disk sensors). We call an optimal path in this respect a *minimum witness path*.

This problem can be seen from two sides: On the one hand, it is a path planning problem. On the other hand, the minimum possible number of sensors that detect a path of the agent is an important parameter of the sensor network. It is called the *barrier resilience* of the network. sensor networks with a low barrier resilience are more error-prone than those with high barrier resilience. In the analysis of a sensor network that is designed to detect an intruder, the minimum witness path points to the network's weak spot. Therefore, to optimize sensor networks it would be very helpful to have an efficient method at hand to compute the barrier resilience of the network or, even better, a minimum witness path.

There are many different types of sensor networks conceivable. We here restrict our attention to the very natural case where the sensor regions are visibility domains.

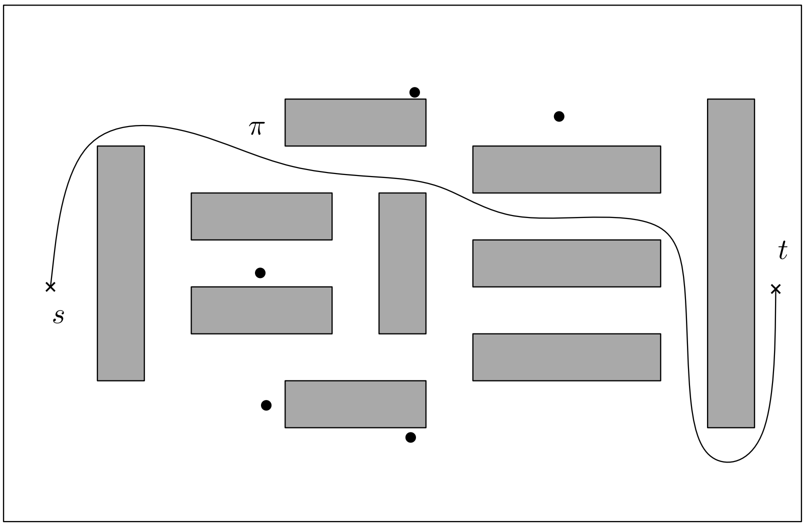

Figure 4.1: Path π from s to t is seen by two witnesses (black points).

The connection to the guarding problems mentioned earlier is evident. But while the guards in the art gallery problems had to guard every spot of the polygon, here it suffices that the witnesses observe some subset of the polygon that separates the points s and t. Also, in art gallery theory we try to position a minimum number of guards that can accomplish the task while here we are given a fixed guard set (or sensor network) and try to find its weak spot. Notice, that it would not be an interesting problem to find the minimum number of witnesses that can be placed somewhere in the polygon such that every path from s to t would be seen by at least k witnesses: Placing k witnesses at s (or at points that see s) would do and would obviously be optimal. On the other hand, if the visibility domains of a given set of guards in a simple art gallery P really cover P is not hard to compute either. Even if we do not make use of the special structure of the visibility polygons but simply compute their union as a union of simple polygons and then compare it with P, this can easily be done in polynomial time.

We can also view the Barrier Resilience problem in a different way: Let us interpret our witnesses as light bulbs that initially are turned on. The question now is, how many lights do we have to switch off, such that there is a path from s to t that lies in complete darkness. Notice, that the question again would be very easy, if the path π from s to t was prescribed. We would have to switch off exactly the lights whose visibility polygons are crossed by π.

Consider now the corresponding illumination problem. Again our witnesses are light sources, but this time, the lights are initially all turned off. Can we compute the minimum number of lights that have to be switched on such that there is a path from s to t that is completely lit-up? This question can be solved very easily. We formulate the result as a lemma. Notice, that the result is not

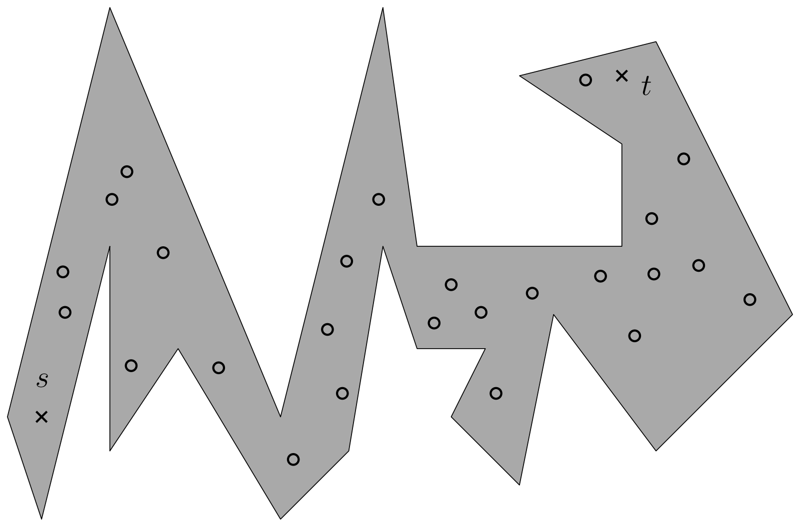

Figure 4.2: An illumination problem: How many lights (circles) do we have to turn on to illuminate a path from s to t?

restricted to simple polygons. The only properties of visibility polygons that the proof uses are that they are path connected, that in polynomial time we can compute visibility polygons, test whether a point (s or t) lies in a visibility polygon and test if two visibility polygons intersect.

Lemma 24. *Given a polygon P and the locations of a set of light sources L as well as a start point and a target point, a subset of L of minimum size that illuminates a path from s to t can be found in time polynomial in $|P|$ and $|L|$.*

The proof is simple. It seems, however, that this problem has not been considered before.

Proof. We can, after computing the visibility polygons of all lights, test every pair of them for intersection. We then build the graph with vertex set the set of lights L and an edge between l_1 and l_2 if and only if $\text{vis}(l_1)$ and $\text{vis}(l_2)$ have nonempty intersection. We then add one vertex for s and connect it to all vertices of witnesses that can see s and we add one vertex for t and connect it to the vertices of witnesses that can see t. Now we compute via breadth first search an edge minimal path from the vertex of s to the vertex of t. The light sources corresponding to the vertices on this path are a minimum solution for the problem of illuminating a path between s and t.

To see this, first observe that the vertices of our edge minimal path in the graph G indeed correspond to a set of light sources illuminating a path between s and t. Let to this end $v_1, v_2, \ldots, v_k$ be the vertices between s and t on our edge minimal path in this order and let $l_1, l_2, \ldots, l_k$ be the corresponding light sources. First, visibility polygons are path connected. Therefore, if a light is turned on, an agent can travel from any point of the visibility polygon of a light source to any other point on a completely lit-up subpath. Second, every two consecutive vertices v_i, v_{i+1} on our edge minimal path correspond to light sources whose visibility polygons intersect in an intersection point c_i. By construction, our agent can travel from s to the light source l_1 on a path completely lit-up by l_1. By our first observation, for every i our agent can travel from l_i to l_{i+1} via c_i on a path lit-up jointly by l_i and l_{i+1}. It follows by induction that the lights $l_1, \ldots, l_k$ illuminate the path from s to t.

To see that our solution indeed is optimal, consider a smaller set of light sources $L_{OPT} = l'_1, l'_2, \ldots, l'_m$ with corresponding vertices $v'_1, v'_2, \ldots, v'_m$ that also illuminates a path π_{OPT} from s to t and that has optimum size. We can assume that for every $l \in L_{OPT}$ there is some point $p \in \pi_{OPT}$ such that p lies in the visibility polygon of l but in the visibility polygon of no other light source of L_{OPT} (as otherwise we could eliminate such an l, obtaining an even smaller illuminating set). The set of corresponding vertices in our graph is connected (and connects s to t). It therefore contains a path in the graph from s to t that contains every of the vertices $v'_1, v'_2, \ldots, v'_m$ at most once and is therefore shorter than our solution, contradicting the fact that our solution is edge minimal. □

In the following sections, we will show that we can find minimum witness paths in polynomial time in simple polygons and in polygons with one hole. On the other hand we prove that the Barrier Resilience problem for visibility polygons in polygons with holes is APX-hard. In particular, we get a stronger inapproximability factor than the hardness results known for line segments.

4.1.2 Related Work

Finding minimum witness paths is related to several other tasks. Algorithms that are concerned with the search for shortest paths in polygons (see for example [24]) or minimum cost paths in graphs, where weights are assigned to the edges of the graph [12] are among the best-researched topics in the field of algorithms.

While this section is about getting somewhere without being seen by too many people, there are many works concerning itself with deploying guards or cameras so that everything of interest is seen at least once or at least a certain number of times. In this category fall the many variations of the art gallery problem, see for example [41].

Also problems that combine path planning questions with guarding problems have been examined. In the Watchman Route problem, introduced by Chin and Ntafos [9] the task is to find a shortest closed path π from a given starting point through a polygon P such that every point of P can be seen from some point of π (or equivalently, such that every point in P sees the path π). Chin and Ntafos prove that the problem is $NP-hard$ for polygons with holes and there is an $O(n \log \log n)$ time algorithm for simple rectilinear polygons. Since

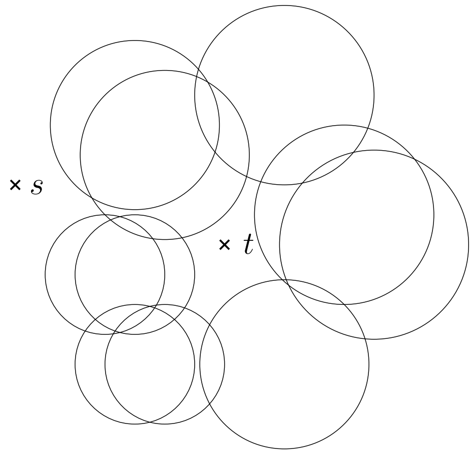

Figure 4.3: An instance of the Barrier Resilience problem for disks

then, various versions of the Watchman Route problem have been defined. The one most strongly related to our problem is the Robber problem that was defined by Ntafos in 1990 [39]. Given a set of edges S and a set of threats T a robber in a simple polygon P wants to find a shortest cycle from which he can see all of S, while not being seen by any of the threats. In this setting, the problem has got a solution only if there exists such a path outside the visibility polygons of the threats. Ntafos gives an algorithm that solves the problem in time $O(n^4 \log\log n)$.

In [17] Gewali et al. define a special case of the Weighted Regions problem [38] and apply it to the following problem. Given a polygon with holes, a starting point s, a target point t and a set of k threats. Find the *least risk path* from s to t. The authors give an algorithm that computes a least risk path in time $O(k^4 n^4)$. The risk is measured by the total length of the subpaths that are inside the visibility polygon of some threat. Here lies the main difference to our model in which the cost of using a witnesses visibility region is fixed, no matter how often or how long the path traverses this region.

For a different kind of related problems, notice that we do not require the witnesses to see the whole path. That the witnesses need only to see a single point of the path, establishes a connection to *Hitting Set* problems. Our question could also be formulated as follows: We are implicitly given a space containing a lot (magnitudes more than one could check one-by-one) of (geometric) objects and a collection of subsets (visibility regions of witnesses) of the ground set. Find the object that is hit by the fewest subsets.

In 2005, in the environment of sensor networks Kumar et al. [34] introduce the notion of a *k-barrier coverage*. In their setting, somebody wants to cross a belt region over which a sensor network is deployed. The belt region is called *k-barrier covered* if every path that crosses the belt is detected by at least k sensors.

Bereg and Kirkpatrick [5] introduce the notion of *barrier resilience*: Given a collection of geometric objects that model the ranges of sensors and two points s, t in the plane, find the minimum number of objects one has to remove such that s and t are in the same component of the complement of the remaining objects. I.e. the barrier resilience is the maximum k such that the region is k-covered. They give an approximation algorithm for this problem when the sensor ranges are unit disks. Until today it is unknown if this original problem is NP-hard. In [2] Alt et al. show that the Barrier Resilience problem for line segments is APX-hard and they also define related problems. In [44] Tseng and Kirkpatrick strengthen the result to *unit* line segments. Gibson et al. [19] give an approximation algorithm for a path that visits multiple points and tries to avoid as many unit disks as possible. Chan and Kirkpatrick [7] give a 2-approximation algorithm for the case of Non-identical Disk Sensors.

One can also view the barrier resilience problem in a very abstract graph-theoretic setting where an agent wants to travel from some start vertex of a graph G to some target vertex. In this setting the barriers are arbitrary subsets of the edge set of G. The barriers can also be interpreted as colors that are assigned to the edges. This problem is then called the Minimum Color Path problem. Carr et al. [6] show that unless P = NP, the optimal solution cannot be approximated to within a factor $O(2^{\log^{1-\delta(|C|)}|C|})$, where $|C|$ is the number of colors and $\delta(|C|) = \frac{1}{\log\log^a |C|}$, for any constant $a < 1/2$. In [48], Yuan et al. use the Minimum Color Path model to analyze reliability in mesh networks.

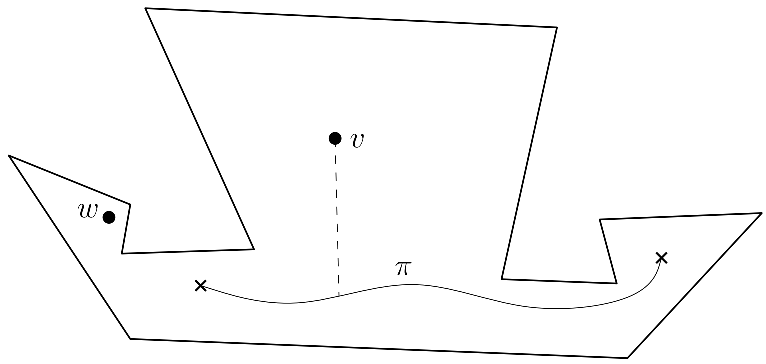

Figure 4.4: Path π is seen by witness v but not by witness w.

4.2 Minimum Witness Paths in Polygons

Most of the results in the next two sections will appear in the extended abstract [20].

4.2.1 Simple Polygons

In our first setting the starting point s and the target point t lie inside a simple polygon P, and we are given a finite set of witness points $W \subset P$. We want to find a path from s to t that is seen by as few as possible witnesses. Let us restate this formally.

Definition 13. Let a polygon P, two points $s, t \in P$ and a set of so-called *witness points* $W = \{w_1, \ldots, w_n\} \subset P$ be given. The *barrier resilience* of W is the minimum cardinality of a subset V of W such that there is an $s-t$-path in P that does not touch any visibility polygon of a point in $W \setminus V$. A path that attains this minimum is called a *minimum witness path*.

We use the usual notion of visibility inside simple polygons that is also illustrated in Figure 4.4.

Definition 14. Let P be a simple polygon. We say that $p_1 \in P$ *sees* $p_2 \in P$ iff the line segment $\overline{p_1 p_2}$ is a subset of P. We say that a witness point $w \in P$ *sees the path* π iff there is a point p on π that is seen by w.

It turns out that in this setting one can find an optimal path very efficiently. The key insight is the following structural lemma.

Lemma 25. *Let P a simple polygon, points $s, t \in P$ and a witness point $w \in P$. If there is a path π in P from s to t that is not seen by w, then the shortest path from s to t in P is not seen by w.*

Before we prove the lemma, we draw the following conclusions that settle the problem for simple polygons.

Theorem 7. *Given a simple polygon P with n edges, two points $s, t \in P$ and a set of witness points $W \subset P$, the shortest path between s and t is an optimal solution to the minimum witness path problem.*

Proof. Let π' denote the shortest path from s to t. By Lemma 25, for every path π between s and t the set $W' = \{w \in W \mid w \text{ sees } \pi'\}$ is a subset of $W(\pi) = \{w \in W \mid w \text{ sees } \pi\}$ and consequently $|W'| \leq |W(\pi)|$. □

Corollary 3. *Given a simple polygon P with n edges, two points $s, t \in P$ and a set of witness points $W \subset P$, we can determine a minimum-witness path in time $O(n)$.*

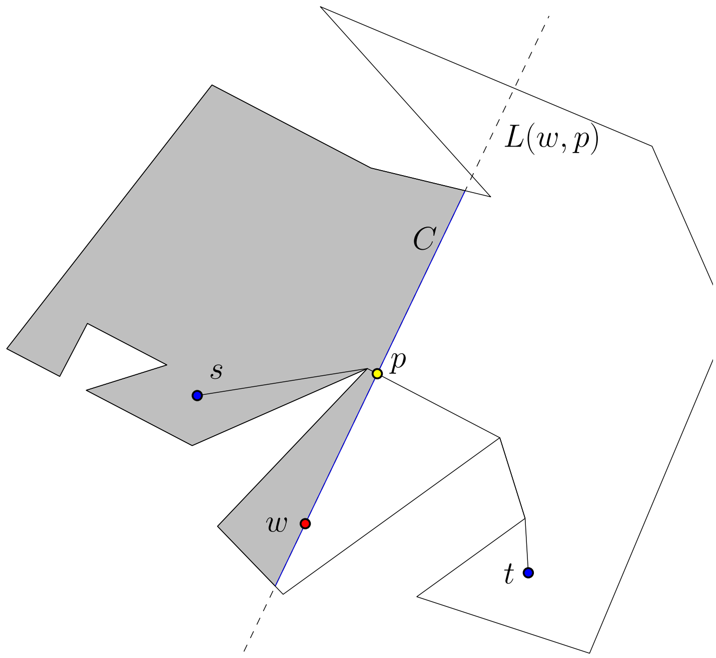

Figure 4.5: The connected component C of $L(w,p) \cap P$ that contains w and p splits P into two connected components, one containing s, the other containing t.

Proof. The shortest path between two points inside a simple polygon with n edges can be computed in time $O(n)$ [24]. □

The proof of the lemma uses the simple topological structure of the polygon.

Proof of Lemma 25. Let π' be the shortest path between s and t and $w \in P$ a point that sees the point p on π'. If w sees s or t it obviously sees every path from s to t. Otherwise consider the line $L(w,p)$ through w and p.

The points w and p lie in the same connected component C of $L(w,p) \cap P$. Now $P \setminus C$ splits into at least two connected components. As π' is the shortest path, s and t lie in different components (otherwise π' could be shortened to a path that is completely contained in the common component of s and t).

It follows that every path from s to t must pass C and is therefore seen by w. $\square$

4.2.2 Polygons with Holes

The next step is looking at polygons with holes. So now we have a simple polygon P' and a collection of simple polygons $H_1, \ldots, H_m$, called the holes, where every hole lies in the interior of P' and $H_i \cap H_j = \emptyset$ for all $1 \leq i < j \leq m$. The polygon with holes P then is defined to be $P = P' \setminus \bigcup_{i=1}^{m} \mathring{H}_i$, where $\mathring{H}_i$ denotes the topological interior of H_i. Let $|P|$ denote the total number of edges of P. Two points $p_1, p_2 \in P$ *see* each other if and only if the line segment $\overline{p_1 p_2}$ is completely contained in P.

Again we are given two points $s, t \in P$ and witnesses $w_1, \ldots, w_n \in P$ in general position, and we want to find a path π inside P from s to t minimizing the number of witnesses who can see π.

First we show that the problem is APX-hard by a reduction from Vertex Cover that provides a stronger factor than other hardness proofs in the context of barrier resilience.

Theorem 8. *Estimating the barrier resilience of a set of visibility polygons inside polygons with holes is APX-hard. In particular, unless $P = NP$, the barrier resilience of visibility polygons with holes cannot be approximated within a factor of* 1.3606*. If the Unique Games Conjecture is true, then the barrier resilience cannot be approximated within any constant factor better than* 2*.*

Proof. We show this by an approximation factor preserving reduction from Minimum Vertex Cover.

Let $G = (V, E)$ be an instance of vertex cover. Let $e_1, e_2, \ldots, e_m$ an enumeration of the edges, $v_1, v_2, \ldots v_n$ an enumeration of the vertices.

We now construct a polygon with holes P in the plane that contains a start point s, a target point t and n witness points $w_1, \ldots w_n$ such that every path from s to t in P corresponds to a vertex cover of G.

To this end we build a big surrounding rectangle $P' = [-2(m+n+1), m+2] \times [-m-n-1, m+n+1]$. We place the start point at the origin, $s = (0, 0)$ and the target point at $t = (m+1, 0)$.
For every edge e_j in E, we add a thin rectangular hole $R_j = [j, j+0.5] \times [-j, j]$. Then we place the witness points at $w_i = (-2(m+n), i - \lceil \frac{n}{2} \rceil)$. If v_k and v_l (with $k \leq l$) are the vertices incident to edge e_j we define $L(j) = w_k$, $H(j) = w_l$ to be the witnesses corresponding to the vertices with lower and with higher index, respectively. We also define $f\colon \{w_1, \ldots, w_n\} \longrightarrow \{v_1, \ldots, v_n\}$ to be the bijection that maps every w_i to v_i.

To construct the holes that model the vertex-edge incidences we proceed as follows:

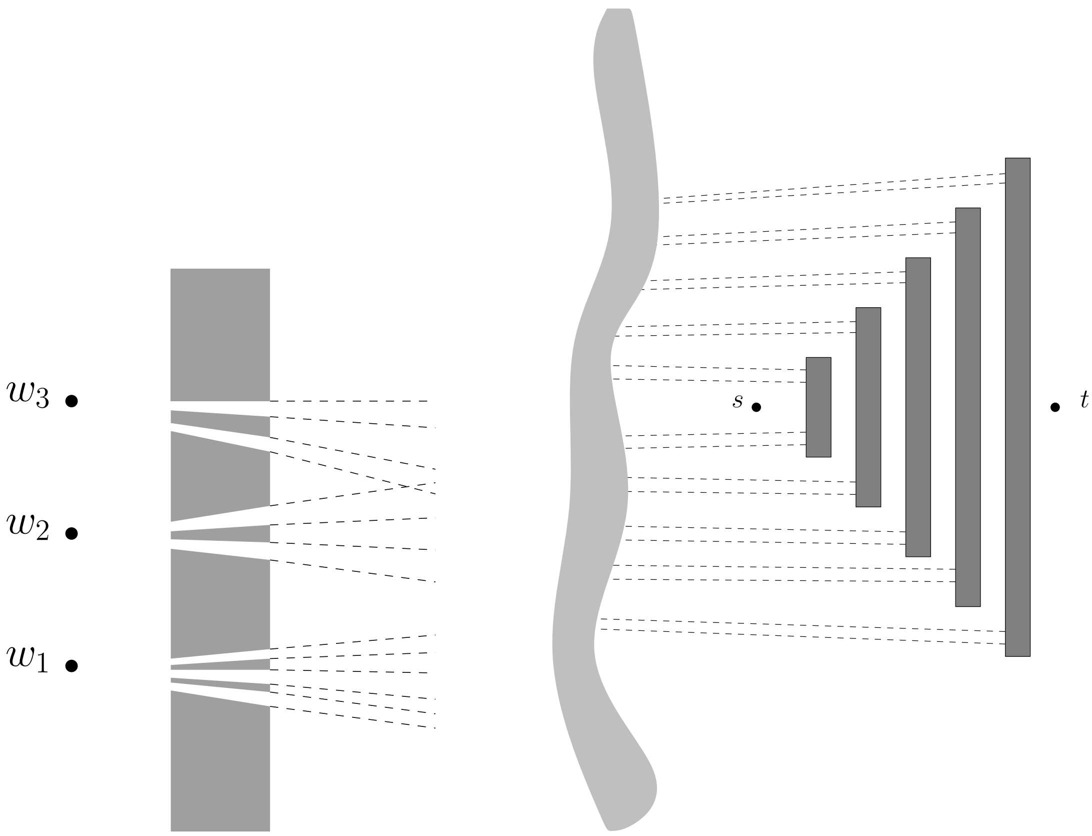

Figure 4.6: On the left side of the polygon there are only narrow slits between the holes through which the witnesses (which correspond to the vertices) can peek. Far away on the right side portions of visibility regions hit the rectangles corresponding to edges.

We start with one rectangle
$Z = [-2(m+n)+0.5, -2(m+n)+1] \times [-m-n, m+n]$ and split it into $2m+1$ pieces.
For every edge e_j we define the two triangles

$$TH_j = \Delta(H(j), (j, j-0.25), (j, j-0.5))$$

and

$$TL_j = \Delta(L(j), (j, 0.25-j), (j, 0.5-j)).$$

Now we construct the $2m+1$ holes by simultaneously cutting the interiors of all these triangles out of Z. We set

$$Z' = Z \setminus \bigcup_{j=1}^{m} (T\mathring{H}_j \cup L\mathring{H}_j)$$

We add the connected components of Z' as holes to our scene.
By this construction every witness w_i sees a rectangle R_j iff the vertex v_i is incident to e_j.

We first notice that this reduction is clearly polynomial-time. The total number of edges of P is $12m+8$ and the number of points (witnesses and start/target) is n+2, each of which can easily be computed in polynomial time.

To see that every path from s to t that is seen by k witnesses corresponds to a vertex cover of G, observe the following: For every edge e_j the quadrilateral with corners $(j, 0.5 - j), (j, j - 0.5), H(j), L(j)$ contains s and does not contain t. Thus every path from s to t must cross one of its four sides. One of the sides is the edge of a hole that cannot be crossed. The other three sides are visibility segments of $L(j)$ and $H(j)$, respectively, and thus crossing them means to be seen by $L(j)$ or $H(j)$. Therefore, if π is a path from s to t that is seen by the set of witnesses $W(\pi)$ then the image of $W(\pi)$ under f is a vertex cover of G. As f is a bijection, the set of witnesses has the same cardinality as the resulting vertex cover.

On the other hand, if $C \subset V$ is a vertex cover of G we can construct a path from s to t with at most the same number of witnesses. From s we first go to the point $(1, 0)$. Now we are on the boundary of R_1 that corresponds to edge e_1. By definition, $f(H(1))$ or $f(L(1))$ are in C. If $f(H(1))$ is in C, our path proceeds to $(1, 1)$, crossing the visibility region of $H(1)$ (but no other visibility region), and then to $(1.5, 1)$. Otherwise, the path proceeds to $(1, -1)$ (crossing the visibility region of $L(1)$) and then to $(1.5, -1)$. In both cases, the next way point is $(2, 0)$.
We continue in this manner, getting, for every j, from $(j, 0)$ to $(j + 1, 0)$ by crossing the visibility region of $H(j)$ if $f(H(j)) \in C$ and crossing the visibility region of $L(j)$ otherwise, until we reach t. The resulting set $W(\pi)$ of witnesses has at most as many elements as C.
It follows that an α-approximation for the Barrier Resilience problem yields an α-approximation for Minimum Vertex Cover. □

Next we show that in the case of one convex hole either one can ignore the hole (Lemma 26) or one can compute two paths, one of which is a minimum witness path (Theorem 9).

Lemma 26. *Let P be a polygon with one convex hole H, (i.e. $P = P' \setminus \mathring{H}$ for some simple polygon P' and a convex polygon $H \subset P$). Assume that for every point $h \in H$ and for every two line segments $S_1, S_2 \subset P \cup H$ that both have as one endpoint h and the other endpoint on $\partial(P \cup H)$, s and t lie in the same connected component of $(P \cup H) \setminus (S_1 \cup S_2)$. Then there is a unique shortest path from s to t in P and it is a minimum witness path.*

Proof. Take the shortest path π between s and t in $P' = P \cup H$. As this is a simple polygon, π is unique. By assumption, there is no point $h \in H$ and line segments $S_1, S_2 \subset P'$ that connect h to the boundary of P' such that s, t lie in different components of $P' \setminus (S_1 \cup S_2)$, see Figure 4.7. Then π does not intersect H.

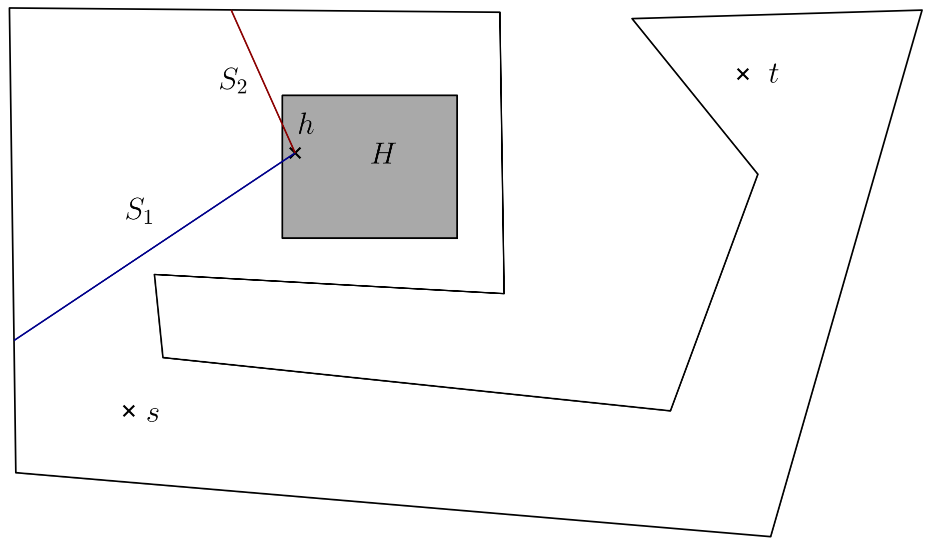

Figure 4.7: The removal of the segments S_1 and S_2 splits P into two connected components. s and t lie in the same connected component.

Otherwise we could take a point $h \in \pi \cap H$ and draw a line segment $S \subset P'$ that crosses π in h and ends in the two points b_1, b_2 on the boundary of P'. Then setting $S_1 = \overline{b_1 h}$ and $S_2 = \overline{b_2 h}$ yields a contradiction to the assumption as s and t lie in different connected components of $P' \setminus (S_1 \cup S_2)$ (because the shortest path crosses the segment S exactly once).

Therefore, π is completely contained in P and is the unique shortest path between s and t. Now suppose, there was a path that was seen by less witnesses π'. Then there was in particular one witness w that sees π but not π'. Let p be a point on the path π that is seen by w. Let further S be the connected component of the intersection $L(w,p) \cap P'$ of the line through w and p with P' that contains p and S_1 be the connected component of $L(w,p) \cap P$ that contains p. If both endpoints of S_1 lay on the boundary of P' then s and t were in distinct components of $P \setminus S_1$. Then every path from s to t would have to cross S_1 and therefore be seen by w, a contradiction.

Thus, one of the endpoints must lie on the boundary of H, let us call this endpoint h. If we now set S_2 to be the topological closure of $S \setminus S_1$, then h, S_1, S_2 are as above and s, t are in different connected components of $P' \setminus (S_1 \cup S_2)$, a contradiction.

It follows, that there can be no path π from s to t and witness w such that w sees π but not π'. Thus the shortest path π is optimal. $\square$

Theorem 9. *Let $P = P' \setminus \mathring{H}$ a polygon with one convex hole, s, t be start and target point, respectively. Let there be line segments S_1, S_2, each of them connecting a point on an edge (not a vertex) of H to a point on an edge (not*

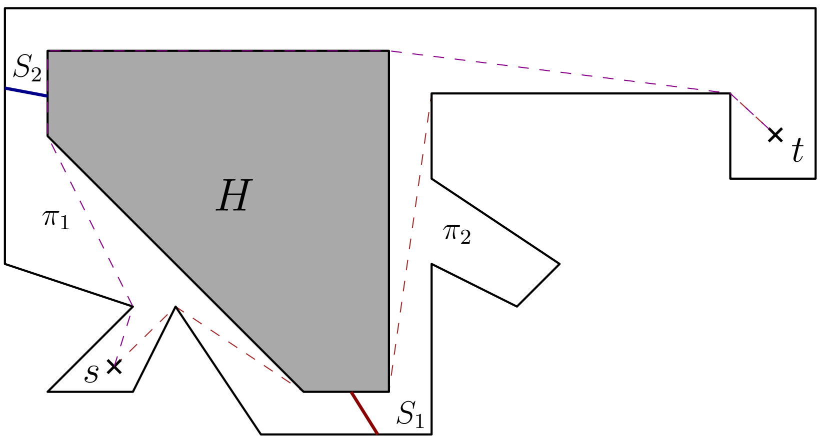

Figure 4.8: The removal of either S_1 or S_2 leaves polygons with unique shortest paths π_1, π_2 one of which is a minimum witness path

a vertex) of the boundary of $P \cup H$, so that s and t lie in different connected components of $P \setminus (S_1 \cup S_2)$. Either the shortest path π_1 from s to t in $P \setminus S_1$ or the shortest path π_2 from s to t in $P \setminus S_2$ is a minimum witness path in P.

Proof. Suppose none of them were optimal. Then there exist witnesses w_1, w_2 (possibly $w_1 = w_2$) and a path π', such that w_1, w_2 do not see π', but w_1 sees point p_1 on π_1 and w_2 sees p_2 on π_2. Let T_1 and T_2 denote the line segments from boundary to boundary of $P \cup H$ through w_1 and p_1 and through w_2 and p_2, respectively. The segments S_1 and T_1 together with H as well as S_2, T_2, H separate the points s and t. By the existence of π', T_1, T_2 and H together do not separate s and t. The connected component of s in $P \setminus (T_1 \cup T_2)$ is simply connected and contains t. As π' does not cross T_1, T_2 it crosses both S_1 and S_2. s and t lie in different components of $P \setminus (S_1 \cup S_2)$, so $(S_1 \cup S_2)$ is crossed an odd number of times. Now we can repeatedly replace subpaths between two crossings of the same segment S_i by the direct paths along the segment (this does not add witnesses) until only one crossing is left, contradicting the fact, that π' crosses S_1 and S_2. $\square$

It follows that in this case the barrier resilience can be computed in polynomial time by computing S_1 and S_2 and then the respective shortest paths.
One can show that this also holds if P contains many convex holes that are strictly separated in a sense made precise below.

Theorem 10. *Let $P = P' \setminus \bigcup_{i=1}^{m} \mathring{H}_i$ a polygon with convex holes, $s, t \in P$, $W = \{w_1, \dots, w_n\} \subset P$ a set of witness points. Let for every $i \neq j$ there be a line segment $S_{ij} \subset P$ s.t. H_i and H_j lie in distinct connected components of $P' \setminus S_{ij}$ and S_{ij} is not seen by any witness $w \in W$. Then one can find a minimum witness path from s to t in polynomial time.*

Proof. Let C_{ij} denote the connected component of $P' \setminus S_{ij}$ that contains H_i. Then for every $1 \leq i \leq m$, $C_i = \bigcap_{j \neq i} C_{ij}$ is a simple polygon, that contains H_i but no H_j for any other index $j \neq i$. For $j \neq i$ $C_i \cap C_j = \emptyset$. We now compute in $O(|P'|)$ time the shortest path π' from s to t in P'. The parts of π' outside $\bigcup_{i=1}^{m} C_i$ are already optimal by Theorem 9. To get the witness-minimal path π through P, we replace the parts of π' inside the C_i by witness-optimal paths, according to Lemma 26 or Theorem 9. The different parts do not affect each other. To this end let us call the point where π' enters C_i s_i and the point where it leaves C_i t_i. We then draw a segment B_1 from an arbitrary point on the boundary of H_i to the boundary of C_i. Then we compute in time $O(|P'|+m) = O(|P|)$ the shortest path π_i^1 from s_i to t_i in $C_i \setminus (\mathring{H}_i \cup B_1)$. We then choose a second segment B_2 from H_i to the boundary of C_i, that intersects π_i^1 (if such a segment exists; otherwise π_i^1 is optimal in C_i by Lemma 26) and compute in time $O(|P'|+m) = O(|P|)$ the shortest path π_i^2 from s_i to t_i in $C_i \setminus (\mathring{H}_i \cup B_2)$. We choose the path less seen by witnesses in P to replace the part of π' inside C_i. (Testing all possible path π_i^1 or π_i^2 with all possible witnesses can be done in total time $O(n|P|^2)$ after the construction of the witnesses' visibility polygons.)

By sewing together the thus computed parts we get a witness-minimal path π from s to t in P. The running time is dominated by the visibility tests that can be carried out in $O(n|P|^2)$ The computation of the polygons C_i can be computed in time $O(m^2|P'|)$. (Shoot a ray from H_i to find the boundary of C_i in $O(|P'| + m)$. Then follow the boundary, turning at every intersection. Testing for the intersections of the m many S_{ij} is in total time $O(m^2)$, following the boundary of $|P'|$ and testing if it meets a segment is in $O(m|P'|)$.) $\square$

We note that as usual for fixed k the question if the barrier resilience is at most k is polynomially solvable by checking all k-element subsets of the set of visibility polygons of witnesses.

4.3 Minimum Neighborhood Paths in Graphs

Let us introduce a graph-theoretic version of the problem: Let $G = (V, E)$ a simple connected graph. For every vertex $v \in V$ we say that v *sees* $w \in V$ iff $\{v, w\} \in E$ or $v = w$. In other words, the set of vertices that are seen by v is exactly its closed neighborhood $N(v)$.

Now let there be two fixed vertices $s, t \in V$ where s denotes the start and t denotes the target. The task now is to find the path π from s to t that is seen by the least number of vertices.

For every path π in G we define the *Neighborhood of* π to be

$$W(\pi) = \{v \in V \mid \text{There is a vertex } p \text{ on } \pi \text{ such that } v \text{ lies in } N(p)\}$$

and we will call the number of neighbors in $W(\pi)$ by $w(\pi)$.

Note that for every path π from s to t all vertices of π and especially s and t always lie in $W(\pi)$.

In this notation, the Minimum Neighborhood Path problem gets the following formulation:

Definition 15. Given an undirected, simple, connected graph $G = (V, E)$ and two vertices $s, t \in V$, the problem Minimum Neighborhood Path is to find a path π from s to t in G that minimizes $w(\pi)$.

As the order in which the vertices are traversed on the path does not matter for the number of neighbors, we identify the path π with its set of vertices and call this set π, too.

4.3.1 Examples

In this section we will have a look at some simple examples and settings in which the problem is easily solved and will on the other hand try to develop an intuition for why the problem may be hard in general. In the following section we will show that in the case of planar graphs there is a simple 3-approximation algorithm. In a later section we will have a look at the complexity of the general problem.

For some particularly simple graphs the problem is solved soon: For $\pi_1 \subset \pi_2$ always holds $w(\pi_1) \leq w(\pi_2)$. It follows that one optimal path between s and t is the shortest path, if G is itself a path or if G is a tree.

Also, if G is a cycle, the shorter one of the two inclusion-minimal paths is the better.

Until now it seems as if the problem could in general be solved using one of the well-known algorithms for shortest paths in graphs. But there are very simple examples in which a shortest path is not the right solution. Consider Figure 4.11.

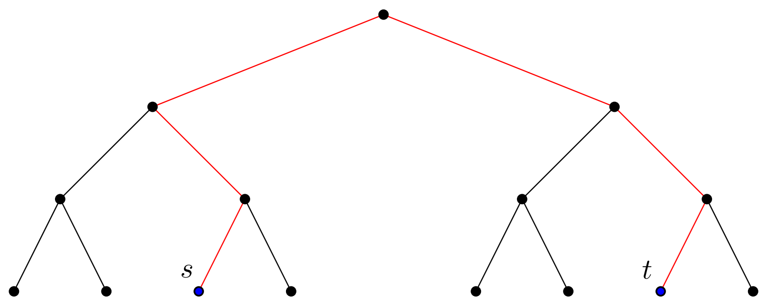

Figure 4.9: Every path from s to t must contain the red edges and will therefore result in at least the same number of neighbors.

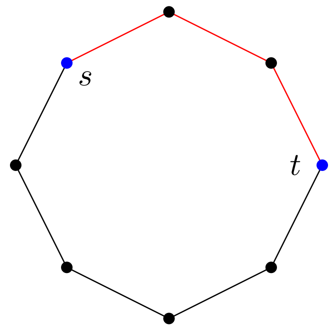

Figure 4.10: The red path has got 6 neighbors, the path on the other side 8 neighbors.

But also in this case the path that is least seen is easily discovered: The vertex on the red path is particularly public, so to speak. Perhaps it is a marketplace. If we assigned every vertex as a weight the number of its neighbors then, in this example, the vertex on the red path would be very expensive so we would never have set foot on it.

Is this the simple answer that addresses the general problem? Assign to each vertex its degree as a weight and find a minimum cost path in the new model? What has the path thus found to do with an optimal path in the original model? Figure 4.12 shows that this approach does not work. It does not even give us a good approximation: Let us adjust the example in Figure 4.12 such that the red path consists of n vertices and the green path and the row above the green path consist of $\lceil\sqrt{2n}+1\rceil$ vertices each. Then in the modified model the red path has cost $2n+4$ and the green path has cost at least $(\sqrt{2n}+1)^2+4 > 2n+4$. So the algorithm would choose the red path. But in the original model the red path has cost $n+4$ while the green path has cost $\leq 2\sqrt{2n}+6$. The result is that for this scheme of examples we get a factor of $> \frac{n+4}{3\sqrt{n}+6} > \frac{1}{6}\sqrt{n}$. Note that in the construction of the examples the number of all vertices is linear in n.

The crucial point that separates our model from others and that prevents

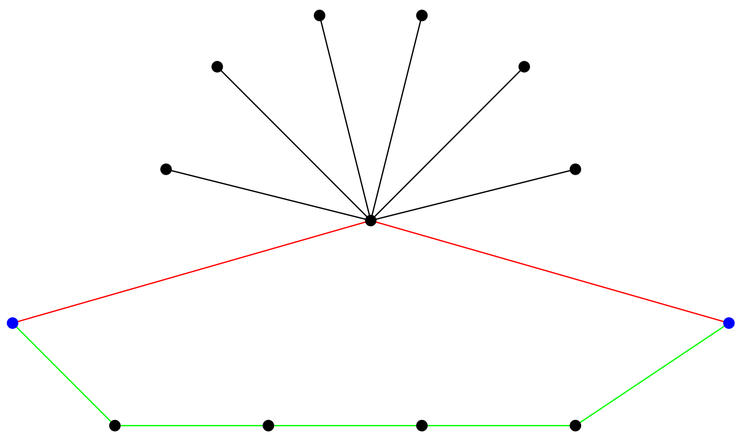

Figure 4.11: While the red path is shorter than the green one, it has got 11 neighbors. The longer green path has got only 7 neighbors.

shortest or minimum cost path algorithms from succeeding is that every neighbor is only counted once.

However, there is another class of graphs containing a start vertex s and a target vertex t, that admit an efficient solution.

Definition 16. Let $G = (V, E)$ simple connected graph and $s, t \in V$. A vertex $v \in V$ is called *essential for* (G, s, t) iff there is no path in $G \setminus \{v\}$ from s to t. We let $\mathrm{Ess}(G, s, t)$ denote the set of vertices that are essential for (G, s, t).

Remark 1. *Obviously, s and t are essential.*

With the help of this definition we can define a class of graphs applied to which the procedure from above yields the desired result. As all solvable classes mentioned so far are special cases of this class, the following theorem summarizes the results of this section:

Theorem 11. *Let $G = (V, E)$ be a simple connected graph and $s, t \in V$ the start and target vertex, respectively. If $G \setminus Ess(G, s, t)$ is a forest, then there is a polynomial-time algorithm that computes a path π from s to t with a minimum number of neighbors $w(\pi)$.*

Proof. First, we notice that one can easily test in polynomial time, if a vertex is essential: Just remove the vertex and check if s and t are still in the same connected component. Also searching for a cycle in $G \setminus \mathrm{Ess}(G, s, t)$ is polynomial.

We now mention some facts that our algorithm will rely on. Every path from s to t must use all vertices from $\mathrm{Ess}(G, s, t)$ by definition. The next thing to notice is that all simple paths from s to t visit all vertices of $\mathrm{Ess}(G, s, t)$ in the same order. (Otherwise there would be a pair of vertices $a, b \in \mathrm{Ess}(G, s, t)$ such that a precedes b in one path but b precedes a in another path. Then there is

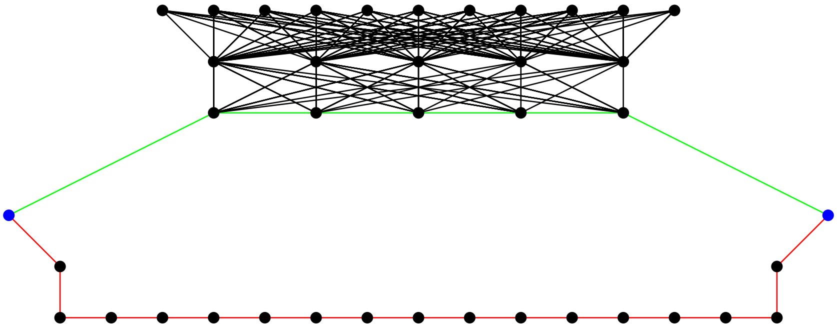

Figure 4.12: The algorithm that finds a minimum-cost path in the weighted graph prefers the red path. Here, every vertex has weight 2, summing up to a total cost of 38. On the green path every vertex except s and t has weight 7, so the green path has total weight 39. In our original model the red path has cost 21 and the green path has cost 14.

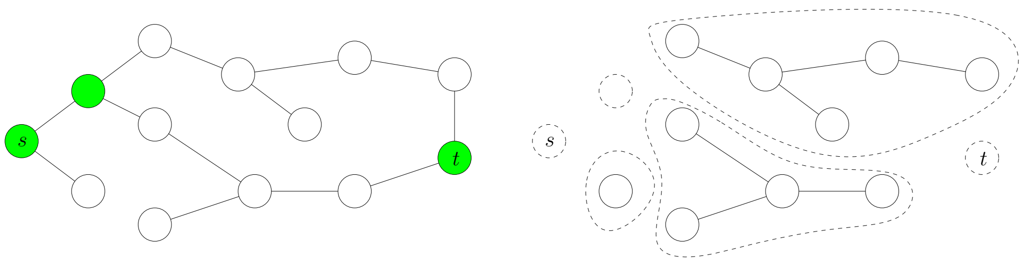

Figure 4.13: *Left:* The green vertices are essential. *Right:* Here, removing the essential vertices yields a forest.

a path from s to a that does not use b in the first path and a path from a to t that does not use b in the second one. Joining the two partial paths at a gives a path from s to t that does not use b. A contradiction to the fact that b is essential.)

Now consider the connected components of $G \setminus \text{Ess}(G, s, t)$. Each connected component is adjacent to at most two essential vertices. If there were three essential vertices a, b, c (in their order on every simple path from s to t) adjacent to one connected component C, then we could shortcut every path from s to t through C to obtain a path from s to t not using b, a contradiction because b is essential.

By assumption, all connected components of $G \setminus \text{Ess}(G, s, t)$ are trees. By the facts we just showed, any vertex of a connected component C of $G \setminus \text{Ess}(G, s, t)$ can be adjacent to at most two vertices (that are essential). If there are less than two essential vertices adjacent to C then the component is useless for building a path from s to t. Otherwise let v_1, v_2 denote the two essential vertices adjacent to

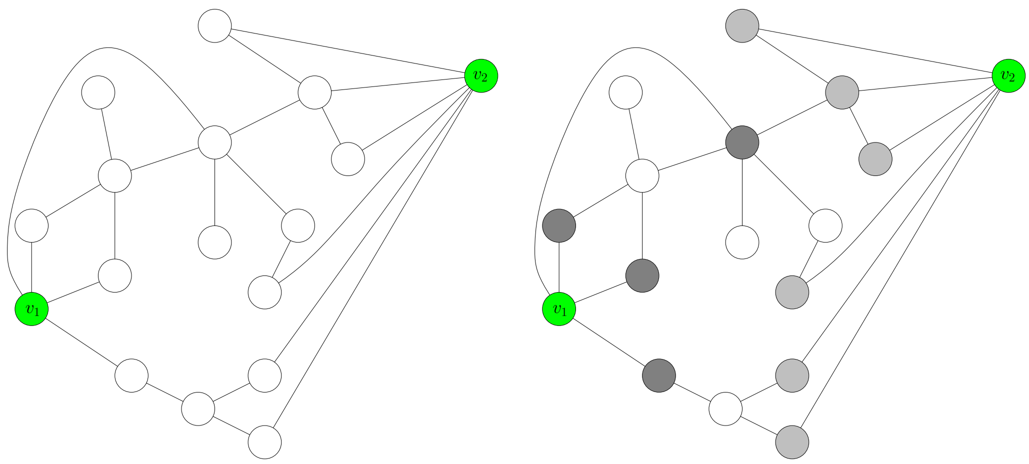

Figure 4.14: *Left:* Trees between two essential vertices. *Right:* The essential vertices' neighborhood are neighbors for every path from s to t.

C. We assign to the connected component the number of neighbors in G, $w(C)$, of the path $\pi_C(v_1, v_2)$ from v_1 to v_2 through C that is least observed by vertices not in $N(v_1) \cup N(v_2)$ (which see every path from v_1 to v_2 anyway). To this end we simply compute the shortest path in C between every $x \in N(v_1) \cap C$ and every $y \in N(v_2) \cap C$. By counting for every such path the number of neighbors that are not in $N(v_1) \cup N(v_2)$ we find a best path from v_1 to v_2 through C. Notice, that the number $w(C)$ can be 0.

Now we build a new graph $G' = (V', E')$ (that contains multiple edges). The set of vertices contains exactly the essential vertices of G. First, for every two vertices $v, w \in V'$ we add an edge e_{vw} if there is an edge between v and w in G. We assign the weight 0 to e_{vw}. Then we add an edge e_C between $v_1, v_2 \in V'$ for every connected component of $G \setminus \mathrm{Ess}(G, s, t)$ to which v_1 and v_2 are adjacent. We assign the weight $w(C)$ to the edge e_C.

We now search for the optimal path π' from s to t in this new graph G'. This search boils down to choosing one lightest edge between each pair of consecutive essential vertices. To obtain the corresponding path π in G we just replace each edge of the form e_C by the path $\pi_C(v_1, v_2)$ from above. If the set of essential vertices is given by $\{x_1, x_2, \ldots, x_n\}$ in their order on every path from s to t and e_i is the edge between x_i and x_{i+1} chosen by our algorithm, then the number of neighbors of π is given by $\#(\bigcup_{i=1}^{n} N(x_i)) + \sum_{j=1}^{n-1} w(e_j)$.

To see that our algorithm really finds a minimum-neighborhood path, observe the following facts.

Every path from s to t visits all $x \in \mathrm{Ess}(G, s, t)$ in a fixed order $x_1, x_2, \ldots, x_n$. A minimum-neighborhood path is given by the concatenation of paths from x_i to x_{i+1}.

Every inclusion-minimal path from x_i to x_{i+1} visits only x_i and x_{i+1} and vertices of at most one connected component of $G \setminus \mathrm{Ess}(G, s, t)$.

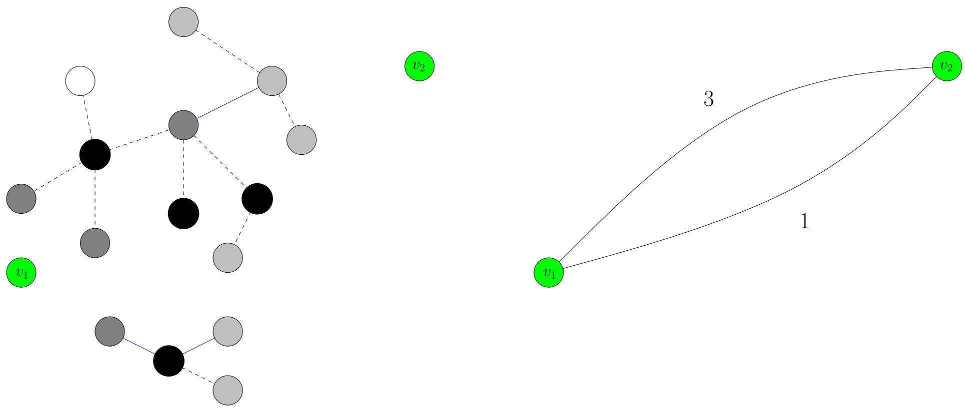

Figure 4.15: *Left:* The chosen paths inside the components and their neighbors. *Right:* In G' each component of $G \setminus \mathrm{Ess}(G,s,t)$ is replaced by an edge that is weighted with the number of neighbors it adds.

Every inclusion-minimal path between x_i and x_{i+1} inside a connected component C of $G \setminus \mathrm{Ess}(G,s,t)$ is seen only by vertices in C or in $N(x_i)$ or $N(x_{i+1})$. Therefore each subpath from x_i to x_{i+1} has no effect on the other parts of the path. Therefore, a concatenation of optimal subpaths π_i from x_i to x_{i+1} for each $1 \leq i \leq n-1$ yields an optimal path from s to t.

For every $1 \leq i \leq n-1$ our algorithm finds the path π_i from x_i to x_{i+1} that minimizes the number of vertices in $W(\pi_i) \setminus \bigcup_{x \in \mathrm{Ess}} N(x)$. In total we get that our algorithm computes the minimum-neighborhood path from s to t. $\square$

Remark 2. *The running time of the algorithm above could easily be improved by traversing the trees in a more clever way. We wanted to give a simple proof of the main point, i.e. that this special case is polynomially solvable.*

4.3.2 Planar Graphs

We now turn to the class of planar graphs. Consider again the following idea for an approximation algorithm. Assign to every vertex v in the graph the cost $c(v) = \deg(v) - 1$, where $\deg(v)$ denotes the degree of v in G. Then find a path $\pi = p_1, p_2, p_3, \ldots, p_l$ from s to t in G that minimizes the cost of π, $c(\pi) = \sum_{i=1}^{l} c(p_i)$, see Figure 4.16.

Computing $c(v)$ for all vertices of G can be done in time $O(|V|+|E|)$. We then build a directed graph $G' = (V, E')$ by replacing every undirected edge $\{v, w\}$ of E by the two directed edges (v, w) and (w, v). All those new directed edges together form the edge set E' of G'. We now assign to every edge $(v, w) \in E'$ the weight $c(w)$. These steps can be done in time $O(|E|)$. We compute a minimum weight path from s to t in G' by using Dijkstra's algorithm [12]. This algorithm computes a path π from s to t that minimizes $\sum_{v \in \pi} c(v)$ as the only vertex whose weights were not counted is the start vertex that is common to all paths

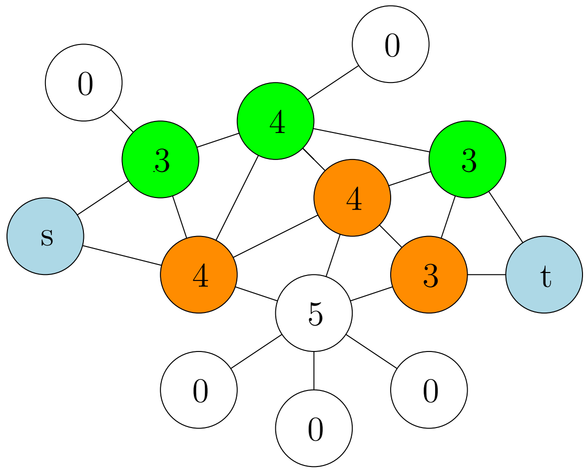

Figure 4.16: Vertices are labeled with $c(v)$. Our approximation algorithm finds the green path that minimizes the number of adjacent edges. The orange path is optimal.

from s to t. As is well-known, this algorithm can be implemented to run in time $O(|V|\log|V| + |E|)$, see for example the book of Cormen et al. [11]. As we are dealing with simple planar graphs, we have $|E| \leq 3|V| - 6$ (see e.g. [32]) and so the overall running time sums up to $O(|V|\log|V|)$.

For simple planar graphs this simple algorithm yields a factor-3-approximation of the minimum neighborhood path.
To see this, observe the following facts. Let π be a path from s to t, $G_\pi = (V_\pi, E_\pi)$ be the subgraph of G that is induced by $V_\pi = N(\pi)$, $w(\pi)$ the number of neighbors of π.

1. $w(\pi) - 2 \leq c(\pi)$.
2. $c(\pi) \leq |E_\pi|$.
3. $|E_\pi| \leq 3|V_\pi| - 6$ because G_π is simple and planar (assuming $|V| \geq 3$).
4. $|V_\pi| = w(\pi)$.

That means that for a path π that minimizes c and any other path π' the following holds.

$$w(\pi) \leq c(\pi) + 2 \leq c(\pi') + 2 \leq 3w(\pi') - 6 + 2 \leq 3w(\pi')$$

Therefore the path found has got at most 3 times as many neighbors as the optimal path.

We summarize:

Theorem 12. *Given a simple planar graph G on n vertices one can find a factor-3-approximate solution to the minimum neighborhood path problem in time $O(n\log n)$.*

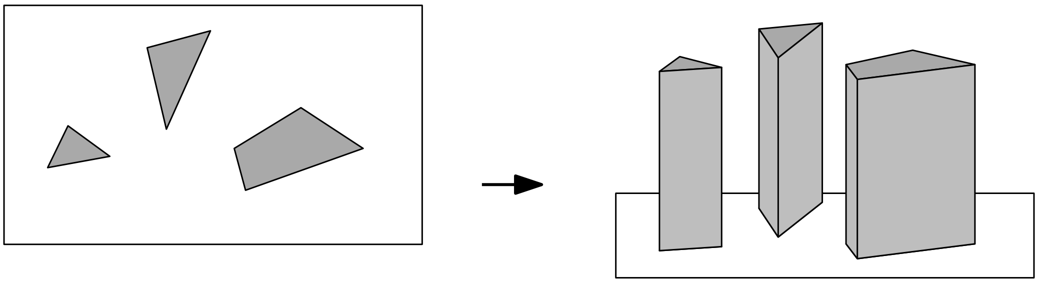

Figure 4.17: The holes of the polygon are transformed into cylindrical parts of the polyhedron.

4.4 Barrier Resilience of Visibility Domains in 3D

4.4.1 Visibility Domains of Polyhedra

In the two-dimensional case we saw that it was easy to find minimum witness paths in simple polygons. What would be the natural analogon for a simple polygon in three dimensions? It seems that a path between two points inside a polyhedron would correspond to our two-dimensional problem in a polygon most appropriately.

Definition 17. A general polyhedron is a three-dimensional solid that is bounded by a finite number of flat polygonal faces.

It is, however, not hard to transfer the hardness result from the last section to polyhedra and to show that finding a minimum witness path inside a nonconvex polyhedron is at least as hard as finding a minimum witness path in a polygon with holes.

Lemma 27. *Estimating the barrier resilience of a set of visibility domains inside nonconvex polyhedra is APX-hard.*

Proof. We transform a given polygon with holes P into a polyhedron R without holes. Let $P = Q \setminus \bigcup_{i=1}^{k} H_i^\circ$ be the polygons with holes, where Q and $H_1, \ldots, H_k$ are simple polygons and the H_i, $i = 1, \ldots, k$, are pairwise disjoint subsets of Q. Let (x_s, y_s) is the start point, (x_t, y_t) the target point and for every witness w in the set W of witnesses, (x_w, y_w) be the location of w of our problem instance of the Barrier Resilience problem in the original polygon with hole. We will construct the boundary of R and thereby define the polyhedron and then place start point, target point and witnesses inside this polyhedron.

First P is embedded canonically into the xy-plane in $\mathbb{R}^3$: We define the set $B = \{(x, y, 0) \in \mathbb{R}^3 | (x, y) \in P\}$ to be the *base* of R. The set B forms a part of the boundary of R. As P as every polygon has a triangulation, we can also see B as a union of triangles.

We continue our construction of the boundary of R by replacing every hole of P by a generalized cylinder. To this end, for every hole H_i of P let $H_i' = \{(x, y, 1)|(x, y) \in H_i\}$. Denote by V_i the union of the segments $\{d + (0, 0, \lambda)|\lambda \in [0, 1]\}$ for $d \in \partial H_i$ (so $V_i = \{d + (0, 0, \lambda)|\lambda \in [0, 1], d \in \partial H_i\}$), see Figure 4.17.

We add the sets H_i' and V_i to the boundary of R. Notice, that for every edge e of H_i the set $\{d + (0, 0, \lambda)|\lambda \in [0, 1], d \in e\}$ is a rectangle.

We form the outer cylinder of Q as the union of $Q' = \{(x, y, 2)|(x, y) \in Q\}$ and $V_Q = \{d + (0, 0, \lambda)|\lambda \in [0, 2], d \in \partial Q\}$. Again, for every edge e of Q the set $\{d + (0, 0, \lambda)|\lambda \in [0, 2], d \in e\}$ is a rectangle.

We define $B \cup Q' \cup V_Q \cup \bigcup_{1 \leq i \leq k}(H_i' \cup V_i)$ to be the boundary of our polyhedron.

To complete the instance of our three-dimensional Barrier Resilience problem, set start and target point at $(x_s, y_s, 0)$ and $(x_t, y_t, 0)$, respectively, and put for every witness w at position (x_w, y_w) in the original instance a witness at $(x_w, y_w, 0)$. Additionally, if there is at least one hole H_1, we choose a point $h = (x_h, y_h) \in H_1^\circ$ in the interior of H_1 and place $2|W|$ witnesses at $h' = (x_h, y_h, 1)$.

The whole reduction is polynomial. It is not hard to see, that for every path π in R one can always compute efficiently a path from s to t that is contained completely in the xy-plane and that is at least as good as π: If a path in R avoids points that lie above holes then it is seen by all witnesses that see its projection to the xy-plane. Every other path is seen by the $2|W|$ witnesses at h' and every path in the xy-plane is better than π. Every polynomial-time algorithm that computes a path in R that is seen by C witnesses can therefore be transformed into an algorithm that computes a path in P that is seen by at most C witnesses. Moreover, an optimal path in R that is contained in the xy-plane canonically corresponds to an optimal path for the corresponding problem in the polygon with holes. Therefore, the hardness result carries over to general polyhedra. □

Let us now have a look at paths on the surface of a convex polyhedron P instead, where the witnesses are placed on the surface of P and visibility now is *exterior visibility*, i.e. the polyhedron itself is breaking the visibility.

Definition 18. A *convex polyhedron* is a bounded subset of $\mathbb{R}^3$ that is the intersection of finitely many closed half spaces $H_1, \ldots, H_k$.

Definition 19. A proper *face* of a convex polyhedron P is every nonempty intersection of P with a plane E such that P is fully contained in one of the closed half spaces determined by E. The 2-dimensional faces are called the *facets*, the 1-dimensional faces are called the *edges* and the 0-dimensional faces are called the *vertices* of P.

We call two facets *adjacent*, if their respective boundaries share a common edge.

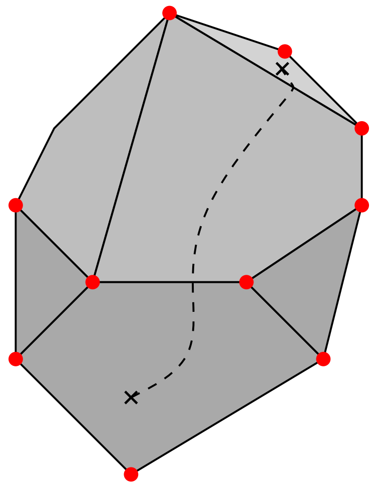

Figure 4.18: A path on a convex polyhedron P (dashed). Facets are grey, edges are black vertices are red.

Definition 20. The *exterior visibility domain*, evis(p), of a point p on the boundary of a convex polyhedron is the set of points v in $\mathbb{R}^3$ such that the segment between p and v does not intersect the interior of P,

$$\text{evis}(p) = \{v \in \mathbb{R}^3 | \overline{pv} \cap P^\circ = \emptyset\}.$$

The exterior visibility domain generally is the union of half spaces. If p is a point in the inner of a facet F, it is the closed half space whose boundary contains F and whose interior does not intersect P. We will call this half space the *outer half space of* F. If it lies on the interior of an edge that bounds the facets F and G, evis(p) is the union of the two outer half spaces defined by F and G. And if p lies on a vertex that is on the boundary of facets $F_1, \dots, F_k$, it is the union of the outer half spaces defined by those facets. As the vertices can overlook all these facets, it would be very reasonable to place witnesses there.

In two dimensions, the analogon to the problem that we will now consider would be a convex polygon P and two designated points s, t on its boundary together with witnesses on the boundary. The task would be to find the minimum witness path from s to t on the boundary. In two dimensions there are only two candidates for the minimum witness path. As $\partial P \setminus \{s, t\}$, the boundary of P without start and end point splits into two connected components C_1, C_2, one of these components is the optimal path, as every path from s to t contains a subset of C_1 or a subset of C_2. The witnesses that see C_i are exactly the witnesses on the edges that contain s and t plus the witnesses that lie on the parts of C_i without these edges. Notice the similarity to the Minimum Neighborhood Path problem on cycles. In this case, the barrier resilience can therefore be computed in linear time.

In three dimensions, however, there are many different paths from s to t, so the problem is more challenging. We will restrict our attention to the case in which there is a witness at every vertex of the polyhedron.

Our problem setting this time formally is as follows:

Let P be a convex polyhedron with boundary B and two distinguished points $s, t \in B$. Let V be the set of vertices of P and let W be a set of witness locations with $V \subseteq W$. For every path π on B, the cost $c(\pi)$ of π is the number of witnesses that see π,

$$c(\pi) = |\{w \in W | \exists p \in \pi : p \in \text{evis}(w)\}|.$$

We want to find a path on B from s to t that minimizes c.

While a shortest path between s and t might visit some vertices of P, in our case the intruder will typically want to avoid vertices, as at a vertex v he is visible to all vertices of all facets adjacent to v. On the other hand, at every point of the same facet, he is exposed to the views of the same set of witnesses. It is therefore reasonable to model the problem by defining a (planar) multigraph $G = (N, E)$, where to every facet we assign a node of N and for every two adjacent facets, we add an edge between their corresponding nodes. We can then label every node with the set of witnesses overseeing the corresponding facet. The edges do not need to be labeled additionally, because their corresponding edges are watched exactly by the union of the sets of witnesses of the two common facets it bounds.

We will assume that no two facets lie in a common plane and that the boundaries of every two adjacent facets have exactly one edge in common (it follows that they have at least two witnesses in common). We observe the following facts.

Lemma 28. *There is always a path π from s to t that minimizes c and that has the following properties:*

1. *The path π does not cross or touch itself.*

2. *If at some point in time π passes from facet F_1 to facet F_2 it will, after leaving F_2 never enter a facet adjacent to F_1.*

Proof. The first property follows from the fact, that we can delete from π the loop between the first and second visit of the crossing point and get a path from s to t that is seen by at most as many witnesses as π. The second property holds because if π enters such a facet F_3 at some point after leaving F_2, it could have entered it right after visiting F_1 while being seen by at most as many witnesses. We can therefore modify the path π by replacing the subpath between some point p_1 in F_1 and some point p_3 in F_3 by a path that goes from p_1 to p_3 without visiting any facet in addition to F_1 and F_3 and obtain a path π' with $c(\pi') \leq c(\pi)$. ☐

From now on, we will assume that our path π has these two properties.

If we think of the path π as the sequence of facets he visits (what perfectly grasps the combinatorial core of the problem), we can describe it by that sequence $(T_1, T_2, \ldots, T_k)$ where T_1 contains the point s and T_k contains t. We will call this sequence the *facet sequence of* π and will in slight abuse of notation call this sequence also π and for every facet T_i of the facet sequence, we will say $T_i \in \pi$. Lemma 28 above then means that the only way in which an intruder who follows such an optimum path can be seen by the same witness in two non-consecutive steps is that there are two facets T_i and T_j with $j \geq i+2$ such that T_i and T_j have exactly one vertex in common (they belong to the set of facets incident to this vertex but do not share a common edge).

Let us call this contact at exactly one vertex between T_i and T_j a *nudge* and define for all indices $1 \leq i < j \leq k$ of the k-step path π

$$N(i,j) = \begin{cases} 1 & \text{if there is a nudge between } T_i \text{ and } T_j \\ 0 & \text{otherwise} \end{cases} \tag{4.1}$$

By definition, $N(i, i+1) = 0$. By our assumptions, consecutive facets on the path have at least two vertices in common. Non-consecutive facets T_i, T_j have exactly $N(i,j)$ vertices in common. It is of course possible that for one facet T_j there are several preceding facets T_i such that for all of them $N(i,j) = 1$. For two adjacent facets T_i, T_{i+1} of the path, we define E_i to be the number of witnesses on the interior of the edge common to T_i and T_{i+1} (i.e. we do not count the two witnesses that always lie on the end points of this edge),

$$E_i = |\{w \in W | w \in T_i \cap T_{i+1}\}| - 2. \tag{4.2}$$

Let us now try to design an approximation algorithm similar to the one for the Minimum Neigborhood Path Problem in planar graphs in Section 4.3.2. To this end, assign to every facet T of P as weight w_T the number of vertices on its boundary minus two. For facets T_i in the facet sequence of a path $\pi = (T_1, T_2, \ldots, T_k)$ we will write w_i instead of w_{T_i}.

Let $\pi = (T_1, T_2, \ldots, T_k)$ be a path from s to t on P that adheres the rule to not visit two facets adjacent to T after having visited T. It is seen by all the vertices on the boundaries of the T_i. We define the weight of π to be

$$w(\pi) = 2 + \sum_{i=1}^{k} w_i. \tag{4.3}$$

While we sum up the weights of these facets, we of course count some vertices twice. More precisely, for every pair of consecutive facets T_i, T_{i+1}, we counted exactly the witnesses on their common edge twice. Additionally, for all nudges caused by the same vertex in one step, we added one. For the true cost $c(\pi)$ of our path we have

$$w(\pi) \geq c(\pi) \geq w(\pi) - \sum_{1 \leq m \leq k-1} E_m - \sum_{1 \leq i < j \leq k} N(i,j). \tag{4.4}$$

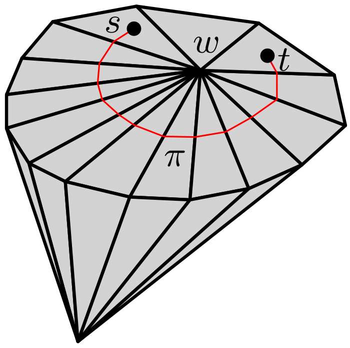

Figure 4.19: Path π that visits many facets incident with witness w. Every pair of non-adjacent facets that π visits creates a nudge.

But this estimate is too weak: Every vertex on an edge between two adjacent facets is counted at most twice by w, thus $\sum_{i=1}^{k-1} E_i < c(\pi)$. But how large can $\sum_{1 \leq i < j \leq k} N(i,j)$ be? Obviously, there is an upper bound of $\binom{k}{2}$. But it can indeed be as large as $\Omega(k^2)$ as a simple example shows, see Figure 4.19. That can of course be much larger than the optimal value of c.

We make use of an accounting argument to bound the excess in our cost measure. Let us count for every witness w the number of its appearances along our path π:

$$A(w) = |\{T_i \in \pi | w \in T_i\}| \tag{4.5}$$

Now we will try to split the appearances of w into groups that we can find upper bounds for. For every witness w that sees π, there is a first appearance of w along the path π. We count the number of times a witness sees π for the first time by $F(w)$. Obviously,

$$F(w) = \begin{cases} 1 & \text{if } w \text{ sees } \pi \\ 0 & \text{otherwise} \end{cases} \tag{4.6}$$

It follows that $c(\pi) = \sum_{w \in W} F(w)$.

The other appearances are *repeated* appearances. To be able to give a good upper bound on the number of repeated appearances, we will split it into two parts.

Definition 21. Let $\pi = (T_1, T_2, \ldots, T_k)$. For every $1 < i \leq k$ we define R_i to be the set of witnesses on T_i that lie also on some facet T_j of π for $j < i$. For every $2 \leq i < k$, our path π leaves T_i via some point p on the boundary of T_i. Let r_a, r_b be the two witnesses in R_i that one encounters first when following the boundary of T_i starting from p in the two possible directions. Let us call these two vertices the *repeated extreme witnesses* of T_i.

Notice, that there are always two such witnesses, as the endpoints of the edge between T_{i-1} and T_i are in R_i. Notice further, that r_a and r_b are always vertices of P.

Let us now define for every witness $w \in W$

$$RE(w) = |\{i \in \{1, \dots, k\}| w \text{ is repeated extreme witness of } T_i\}| \tag{4.7}$$

and

$$R(w) = |\{i \in \{1, \dots, k\}| w \in R_i\}| - RE(w) \tag{4.8}$$

So $R(w)$ is the number of times that w appears as a repeated witness that is not repeated extreme.

Now we can count the appearances of the witnesses: For every witness w_f that lies in the interior of a facet, we have:

$$A(w_f) = F(w_f)$$

For witnesses w_e that lie in the inner part of an edge, we get:

$$A(w_e) \leq 2F(w_e)$$

To bound the number of appearances of witnesses at vertices, we need the following lemma:

Lemma 29. *Let $w \in R_i$ be a repeated witness of T_i that is not a repeated extreme witness of T_i. Then for all $i < j \leq k$, $w \notin T_j$. It follows that $R(w') \leq 1$ for all $w' \in W$.*

Proof. Let T_a, T_b be the last facets from which π has seen the repeated extreme witnesses of T_i, r_a and r_b, respectively. (W.l.o.g., π visited T_a, T_b in this order.) The witnesses r_a, r_b lie on vertices of P. The facets that π visited, starting with T_a until T_b together with the boundary of T_i cut off the access to all faces that have any point of $R_i \setminus \{r_a, r_b\}$ on their boundaries, see Figure 4.20. As π does not have any self-crossings and will never return to T_i, no vertex of $R_i \setminus \{r_a, r_b\}$ will ever be seen by π again. □

Let us now count the number of times the vertex witnesses are seen by π. As always, we have $A(v) = F(v) + R(v) + RE(v)$ for every witness $v \in W$. As for every witness, v is seen at most once for the first time. It is also seen at most once as a non-extreme repeated witness, by Lemma 29. It may be seen many times as extreme repeated witness.

But in every step T_i for $i \geq 2$, only two vertices can be repeated extreme witnesses, by definition. Furthermore, we know that in every step T_i, $i \geq 2$ there are at least two repeated vertices, so $\sum_{w \in W} RE(w) = 2(k-1)$. We therefore

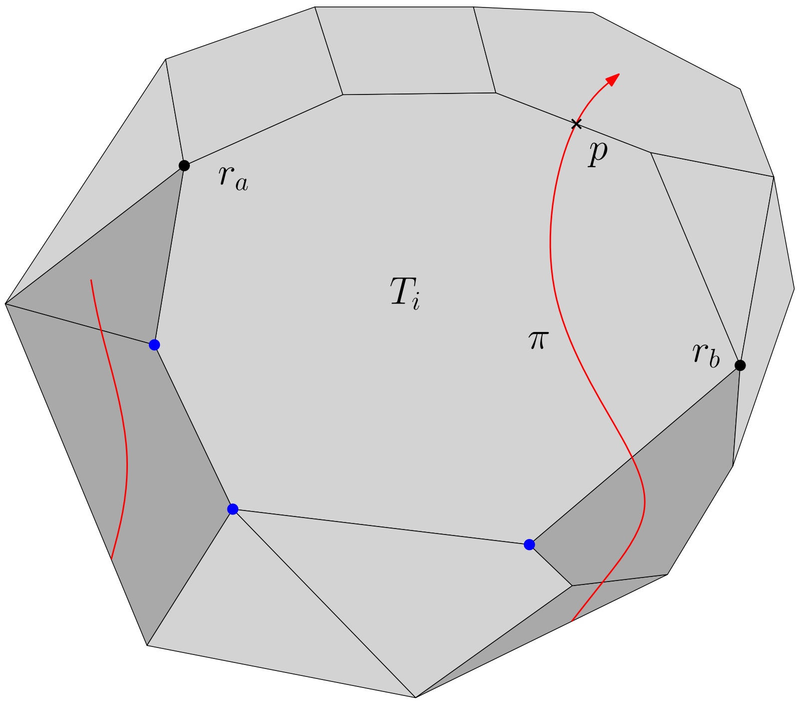

Figure 4.20: The repeated extreme witnesses r_a and r_b are closest to p along the boundary of T_i. Facet T_i together with the facets that π visited before encircle all other repeated vertices (blue).

see that for the computed costs of a path $c(\pi)$ the following inequalities hold.

$$\begin{aligned} w(\pi) &= 2 + \sum_{1\le i\le k} w_i \\ &= \sum_{w\in W} A(w) - 2(k-1) \\ &= \sum_{w\in W} (F(w) + R(w)) + \sum_{w\in W} RE(w) - 2(k-1) \\ &\le 2\sum_{w\in W} F(w) + 0 \\ &= 2c(\pi) \end{aligned}$$

On the other hand, $w(\pi) \ge c(\pi)$. Especially for the path that minimizes w, π^w_{OPT} and the optimal path with respect to the real cost function c, π_{OPT}, the following inequality holds:

$$c(\pi^w_{OPT}) \le w(\pi^w_{OPT}) \le w(\pi_{OPT}) \le 2c(\pi_{OPT})$$

Given the polyhedron and the witness locations, computing the weights can easily be done in polynomial time. The optimization itself can again be done using Dijkstra's algorithm as in Section 4.3.2. We have therefore proven:

Theorem 13. *One can compute a 2-approximate solution to the* Barrier Resilience *problem on a convex polyhedron in polynomial time if on every vertex there is exactly one witness.*

4.4.2 Visibility Domains between 3-dimensional Obstacles

We now turn our attention to minimum witness paths in three dimensional space with polyhedral obstacles. We will see that this problem is at least as hard to approximate as the Minimum Color Path problem. This problem is inapproximable in a strong sense as shown by Carr et al. [6]. We will also give a reduction from the Minimum Color Path problem to a variant of the Minimum Neighborhood Path problem that serves as a prototype for our reduction to the Barrier Resilience problem for visibility domains in three dimensions among polyhedral obstacles.

Definition 22. In the Minimum Color Path problem we are given a graph $G = (V, E)$, a start node $s \in V$, a target node $t \in V$, a finite set C of "colors" and a function $c\colon E \longrightarrow C$ assigning a color to each edge. The task is to find a path π from s to t in G that minimizes the number of colors assigned to the edges traversed by π.

Carr et al. ([6]) showed that unless P = NP, Minimum Color Path cannot be approximated to within a factor $O(2^{\log^{1-\delta(|C|)}|C|})$, where $|C|$ is the number of colors and $\delta(|C|) = \frac{1}{\log\log^a |C|}$, for any constant $a < 1/2$. It is easy to see, that

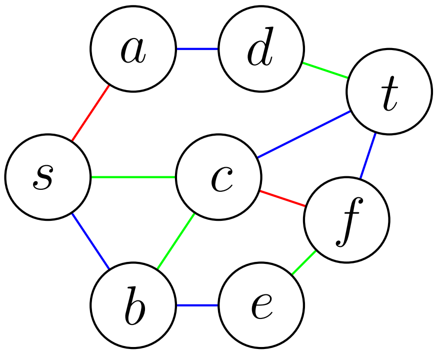

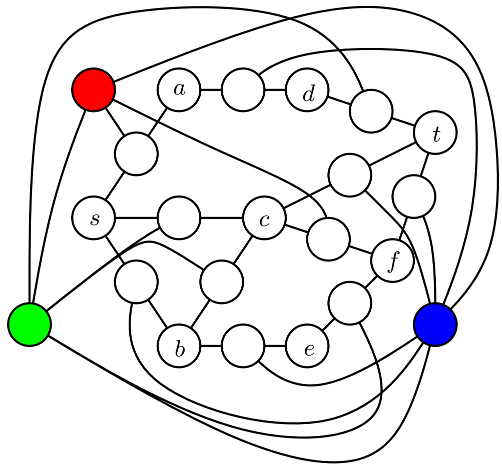

Figure 4.21: Left: An instance of Minimum Color Path where every edge has got a color and a path with the least number of colors is sought. Right: The corresponding instance of 0-1-Weighted Minimum Neighborhood Path. The color vertices have weight 1, all other vertices have weight 0. The color vertices are connected to vertices in G' that correspond to edges in G.

the factor of approximation carries over to a weighted variant of the Minimum Neighborhood Problem by a straightforward reduction:

The 0-1-Weighted Minimum Neighborhood Path problem is a variant of the Minimum Neighborhood Path problem, where all vertices of the graph $G = (V, E)$ are divided into two classes A and Z ($A \dot{\cup} Z = V$). We no longer want to minimize the size of the neighborhood of a path π from s to t. Instead we assign every vertex $z \in Z$ the weight $w(z) = 0$ and every vertex $a \in A$ the weight $w(a) = 1$ and then want to find a path π of minimum weight. In other words, we want to minimize the number of vertices in the intersection $W(\pi) \cap A$, where $W(\pi)$ denotes the neighborhood of the path π and we do not care about how many elements of Z are in the neighborhood of π.

We can reduce Minimum Color Path to 0-1-Weighted Minimum Neighborhood Path by building for an instance graph $G = (V, E)$ with color function $c\colon E \longrightarrow C$ the graph $G' = (V', E')$ with weight function $w\colon V \longrightarrow \{0, 1\}$.

We first describe how the vertex set V' of G' is built. All vertices of V are also vertices of V'. For every $e \in E$ there is a new vertex $v_e \in V'$, we define $V_E = \{v_e \in V' | e \in E\}$ to be the set of all the vertices added for an edge in the original graph. Also for every color $k \in C$, one vertex v_k is added to V', and we define $V_C = \{v_k \in V' | k \in C\}$ to be the set of vertices corresponding to colors in the original setting.

The weight function assigns positive weights only to the vertices corresponding to the colors in the original setting.

$$w(v) = \begin{cases} 0 & \text{if } v \in V \cup V_E \\ 1 & \text{if } v \in V_C \end{cases} \tag{4.9}$$

The set E' of edges of G' does not contain any of the edges of E. Instead, for every $e \in E$ that connects vertices $v, w \in V$ and is assigned the color $c(e)$, the three edges $\{v, v_e\})$, $\{w, v_e\}$ and $\{v_e, v_{c(e)}\}$ are added. Additionally, for every pair of colors $a, b \in C$ we put the edge $\{v_a, v_b\}$ into E', see Figure 4.21.

Now every path from s to t in G with cost r corresponds uniquely to a path in G' that does *not* visit any vertices of V_C and has total weight r. (Every path in G' that does visit a vertex of V_C automatically has total weight $|C| = |V_C|$.) Therefore the approximation hardness factor for Minimum Color Path carries over to to 0-1-Weighted Minimum Neighborhood Path. The reduction is clearly polynomial-time.

It is easy to see that we could also use this reduction to show approximation hardness for the original Minimum Neighborhood Path problem. To this end we could carry out the construction just described but assigning to every vertex the weight 1, obtaining a graph G as above. Then, by replacing every vertex v of V_C and its incident edges by multiple copies of it we get a graph that is equivalent to a weighted version of G in which the weights of all vertices in V_C are amplified while the weights of all other vertices are kept at one. For a large enough amplification factor, the weights of the color vertices dominate the weights of the vertices of V and V_E.

Now we show that the approximation hardness also carries over to the barrier resilience of visibility domains in 3D with polyhedral obstacles. To this end we are going to design for every instance of the Minimum Color Path problem a so-called *parcours*, i.e. a collection of disjoint polyhedra in $\mathbb{R}^3$ together with a set of locations of witnesses in $\mathbb{R}^3$ and a start position and target position. The reduction will mimic the reduction to 0-1-Weighted Minimum Neighborhood Path. We obtain the theorem:

Theorem 14. *Unless* $\mathrm{P} = \mathrm{NP}$*, the* Barrier Resilience *problem for three dimensional visibility domains among polyhedral obstacles cannot be approximated to within a factor* $O(2^{\log^{1-\delta(|W|)}|W|})$*, where* $|W|$ *is the number of witnesses and* $\delta(|W|) = \frac{1}{\log\log^a |W|}$*, for any constant* $a < 1/2$*.*

Proof. Let $G = (V, E)$ be the graph, $C = \{c_1, c_2, \dots, c_k\}$ the set of colors, $c\colon E \longrightarrow C$ the color function of our Minimum Color Path instance. As there are at most $\binom{n}{2}$ many edges that we can colors assign to, we may assume that $|C| \leq \binom{n}{2}$. Let $\{v_1, v_2, \dots, v_n\}$ be an enumeration of the vertices of V such that $v_1 = s$ and $v_n = t$.

We will now specify locations in space where we put witnesses corresponding to the colors of our Minimum Color Path instance. We will also specify

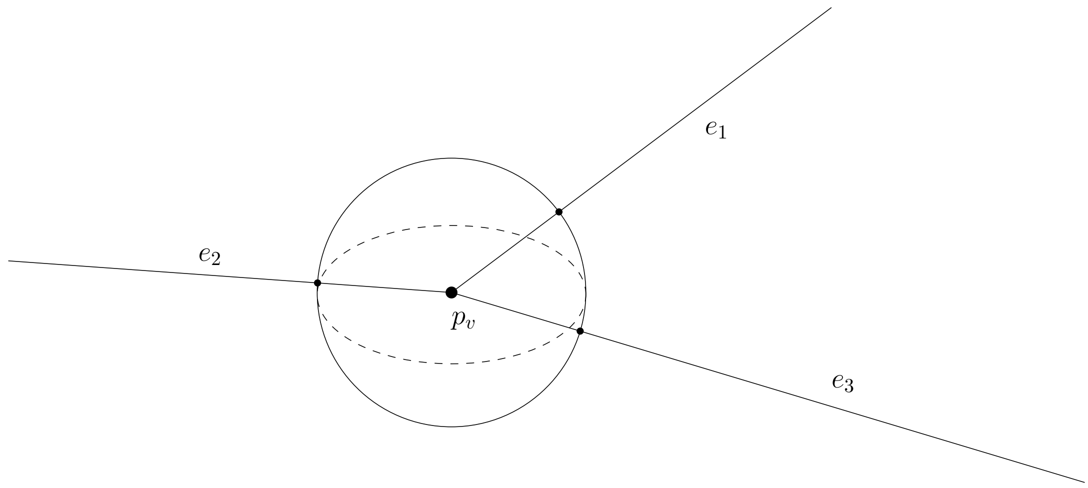

Figure 4.22: A sphere of radius 1 around p_v and entrance points of the edges e_1, e_2, e_3 at the intersections of their corresponding line segments with the sphere.

locations that we will associate with the vertices and edges of our graph G. We will make sure that the location associated with an edge e will be seen by the witness associated with the color of e, $c(e)$, and only by this witness. We will finally ensure, that every path that does not stick to the regions associated with vertices and edges of G will be seen by a huge amount of (additional) witnesses, thereby enforcing that all paths from s to t that may count as candidates for being optimal have to follow the routes prescribed by our construction.

We start by finding is to find for every vertex $v \in V$ a location p_v in $\mathbb{R}^3$ such that no three of these locations are on a common line. We will sometimes say p_i instead of p_{v_i}. We make sure, that no two line segments between these locations intersect, except at their endpoints. Also, we prevent the slopes of the line segments between the points p_i to become to high. For concreteness, we create a *graph zone* that is the box $[-n, n] \times [-n, n] \times [0, 1]$ that will contain all locations corresponding to vertices and edges of G. The vertex locations are now chosen such that their projections to the xy-plane are equally spaced points on the circle with radius n around the origin. So, the location for vertex v_i is placed at point $p_i = (n \cdot \cos(i\frac{2\pi}{n}), n \cdot \sin(i\frac{2\pi}{n}), h_i)$ where the $h_i \in [0, 1]$ are chosen such that no two line segments $\overline{p_i p_j}$, $\overline{p_k p_l}$ for pairwise distinct locations p_i, p_j, p_k, p_l intersect. Notice, that for $n \geq 6$ the distance between two vertex locations is at least 6 so the steepness of every edge is at most $\frac{1}{6}$. We call the set of all these positions $P_V = \{p_v \in \mathbb{R}^3 | v \in V\}$. The locations of the start vertex s and the target vertex t for the Minimum Color Path problem mark the starting and target points p_s, p_t for our Barrier Resilience problem for visibility domains among three dimensional obstacles.

In our reduction to 0-1-Weighted Minimum Neighborhood Path we built for every edge $e_{ij} = \{v_i, v_j\}$ in E a vertex $v_{e_{ij}}$. In this reduction here, the location that corresponds to $v_{e_{ij}}$ is the midpoint of the line segment $\overline{p_i p_j}$, $\frac{1}{2}(p_i + p_j)$. For every edge $e \in E$ we will call this midpoint of the corresponding line segment m_e.

We now generate for every color c of our Minimum Color Path instance one witness w_c. We position w_c in the *witness zone* that is the set $[-n, n] \times [-\frac{1}{n^2}, \frac{1}{n^2}] \times \{2n + 1\}$.

We will now go into the details of the placement of witnesses and polyhedral obstacles. Around every vertex location p_v we will build an umbrella polyhedron that shields p_v and the region around it from the views of all witness points. Around every line segment corresponding to edge $e \in E$, we will build a tunnel pair that shields the line segment from all witnesses but the one corresponding to color $c(e)$. The following specifications of our parcours are a bit technical, but they make sure, that no two obstacles intersect and that every line segment corresponding to an edge e is seen by exactly the witness corresponding to color $c(e)$. The value ρ that we are going to define will serve as the diameter of the circumcircles of the tunnel pairs. The value ε will serve as the width of the gaps between the two tunnel polyhedra of a tunnel pair.

The witnesses are placed in the witness zone in a way such that

1. for every two different colors $a, b \in C$ the witnesses w_a, w_b have distance at least $1/n$ to each other.

2. for every edge $e \in E$ with $c(e) = a$ the line segement $\overline{w_a m_e}$ does not intersect any other line segment $\overline{p_i p_j}$ between points of P_V.

The first property can easily be achieved as there are at most $\binom{n}{2}$ colors. The second property can also easily be achieved as for every midpoint m_e the projection of the set $\bigcup_{p,q \in P_V} \overline{pq}$ from m_e onto the plane $\{z = 2n + 1\}$ is an arrangement of line segments. It follows that arbitrarily close to every point in the witness zone, there is a point in the witness zone that fulfills the second property.

As the line segments are closed sets and as such have open complements, we can now choose a $\rho_1 > 0$ that is so small that for every color a and every edge e with $c(e) = a$ the set $B_{\rho_1}(\overline{w_a m_e})$ does not intersect $B_{\rho_1}(\overline{p_i p_j})$ for any pair of points $p_i, p_j \in P_V$, where for a set $X \subset \mathbb{R}^3$ the set $B_{\rho_1}(X) = \{y \in \mathbb{R}^3 | \exists x \in X : \|x - y\| \leq \rho_1\}$ is the set of points within distance ρ_1 to X. Choosing the diameter of the tunnels smaller than ρ_1 makes sure that the vision of no witness onto one of its edges is blocked by a tunnel around some other edge.

Let $\rho_2 = \min\{\operatorname{dist}(\overline{p_i p_j}, \overline{p_k p_l}) | p_i, p_j, p_k, p_l$ pairwise distinct elements of $P_V\}$ be the smallest distance between line segments that do not have a common end point (where the distance between two nonempty sets A, B is defined to be $\operatorname{dist}(A, B) = \inf_{a \in A, b \in B} \|a - b\|_2$). Choosing the diameter of tunnel pairs smaller than ρ_2 will guarantee that no two tunnels that have no endpoint in common intersect. Let $\rho_3 = \sin(\frac{\pi}{n^2})$. We set $\rho := \min\{\rho_1, \rho_2, \rho_3\}$. Choosing the diameter smaller than ρ will guarantee that no two tunnels that have one endpoint in common intersect. It also rules out the possibility that an umbrella polyhedron and a tunnel intersect. We also set $\varepsilon := \rho/2n^3$. This will make sure, that only the right witness can see through the gap of width ε between the tunnels.

Now that we fixed these values, we can specify the locations of the umbrella and tunnel polyhedra. The umbrella polyhedron of p_i shields it from all witnesses of W. The only apertures in the umbrella polyhedron of p_i will be in the direction of the vertex locations corresponding to vertices v' that are adjacent to v_i in G. To find the location and shape of these apertures, we do the following:

We draw spheres C_i^1 of radius 1 around every p_i. The intersection of this sphere with the segment S_e corresponding to an edge e incident to p_i is called an *entrance point* of this edge, see Figure 4.22. Let m_e be the midpoint of this segment. For every edge $e_{ij} = \{v_i, v_j\}$ we now build small equilateral triangles Δ_{ij}^1, Δ_{ji}^1, that are translates of each other and that are normal to $S_{e_{ij}}$, such that the diameter of their circumcircles is ρ and the centers of their circumcircles are the entrance points of S_e on C_i^1 and C_j^1, respectively. We additionally build small triangles δ_{ij}^1, and δ_{ji}^1 that are shrinked versions of Δ_{ij}^1 and Δ_{ji}^1 whose diameter is only ρ/n but that have the same center and lie in the same plane as their big counterparts . We also define for every $0 < r < 1$ the triangles $\Delta_{ij}^r = \{d \in \mathbb{R}^3 | \exists d_1 \in \Delta_{ij}^1 : d = rd_1 + (1-r)p_i\}$ and
$\Delta_{ji}^r = \{d \in \mathbb{R}^3 | \exists d_1 \in \Delta_{ji}^1 : d = rd_1 + (1-r)p_j\}$.

Now we define the umbrella polyhedron U_i of p_i. To ensure that p_i lies inside U_i, we first take a set $A \subset \mathbb{R}^3$ of four points with distance ρ from p_i such that p_i lies in the interior of $ch(A)$. Define for every $1 \le i \le n$ the set

$$E_i := \Big\{ j \in \{1, \ldots, n\} | \{v_i, v_j\} \in E \Big\}.$$

We define

$$Z_{ij} := ch(\Delta_{ij}^1 \cup \Delta_{ij}^\rho),$$
$$T_{ij} := ch(\{p_i\} \cup \delta_{ij}^1)$$

and

$$U_i := \Big(\bigcup_{j \in E_i} Z_{ij} \cup ch(A \cup \bigcup_{j \in E_i} \Delta_{ij}^\rho) \Big) \setminus \Big(ch(\{p_i\} \cup \bigcup_{j \in E_i} \Delta_{ij}^{\frac{\rho}{2}}) \cup \bigcup_{j \in E_i} T_{ij} \Big)$$

Then, we add for every pair of vertices $v_i, v_j \in V$ that is connected by edge $e = \{v_i, v_j\}$ in E a pair of *tunnels*, see Figure 4.23. Between these two tunnels, there is only a small slit and the ratios between the thickness of the tunnel walls on the one hand and the diameter of the tunnel tubes and the width of the slit on the other hand prevents every path that uses these tunnels from being seen by any witness except $w_{c(e)}$ that lies in a plane going right through this slit.

The tunnel pair is given as a set

$$\Big(ch(\Delta_{ij}^1 \cup \Delta_{ji}^1) \setminus ch(\delta_{ij}^1 \cup \delta_{ji}^1) \Big) \setminus \{q \in \mathbb{R}^3 | \exists p \in P_{ec} : \|q - p\| < \varepsilon\}$$

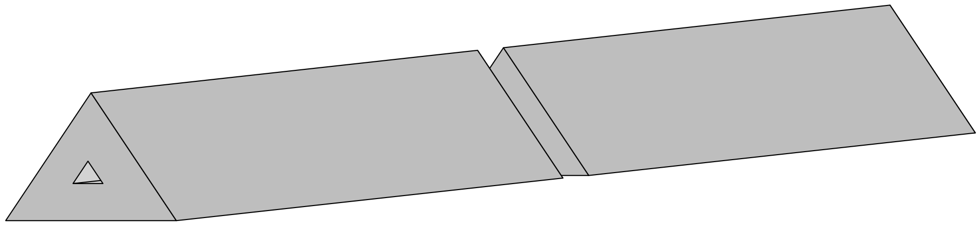

Figure 4.23: A tunnel pair consists of two tunnels but can also be seen as one tunnel that is trenched by a small slit. In general, the slit is not parallel to the entrance triangles.

where P_{ec} is the plane through $w_{c(e)}$, m_e and $w_{c(e)} + a$, where a is the vector of $(1, 1, 0)$ and $(1, -1, 0)$ that maximizes the smaller angle to the projection of $(p_j - p_i)$ onto the xy-plane.

Finally, we add at every outside face of every tunnel pair k additional witnesses that are positioned so that they can not see into the interior tubes of any tunnel.

Following from all these constructions a path starting in p_s and proceeding through the tunnel pairs along the locations $p_s = p_{v_{\sigma(1)}}, p_{v_{\sigma(2)}} \ldots, p_{v_{\sigma(m)}} = p_t$ is seen at most by the witnesses corresponding to the colors of $\{v_{\sigma(i)}, v_{\sigma(i+1)}\}$ for $i = 1, \ldots m - 1$. It is also clear that every path from s to t that leaves the shelter of the umbrellas and tunnel pairs is immediately seen by $2k$ witnesses. The reduction is polynomial. Therefore the factor of approximation carries over from Minimum Color Path to the Barrier Resilience problem for visibility domains with polyhedral obstacles.

□

In particular, for the general Barrier Resilience problem of visibility domains in three dimensions there can be no constant factor approximation unless P = NP.

5 Conclusion

In this thesis we considered visibility domains and problems related to their complexity.

In the field of VC-dimensions, we set about finding new lower bounds for the VC-dimensions of visibility domains. To this end, we developed a novel approach that combines encircling arguments with decomposition techniques and as a main ingredient the new idea of relativization. Relativization makes it possible to replace visibility domains in the analysis of intersections by less complex geometric ranges such as wedges, half planes or intervals.

In the case of visibility polygons, a sophisticated analysis together with a custom-made decomposition enabled the application of relativization. This resulted in Theorem 6 which establishes an upper bound of 14 that significantly improves upon the previously known best upper bound of 23. In the case of perimeter visibility domains, we were able to exploit the simpler structure of these visibility domains to establish an even better upper bound of 7 that is already very close to the lower bound of 5. While the decomposition in this case was less hard to find, to obtain the mentioned result we had to examine the parts closer. In particular, we looked at the coarseness of subsets of the set of visibility domains, that measures how far away a set system is from shattering a point set.

In this thesis we considerably narrowed the gap between lower and upper bounds for the VC-dimensions of visibility domains, but the question arises naturally how much further the methods presented in chapter 3 can carry us in this direction.

In the case of visibility polygons, the first step of our proof was the decomposition of the point set into points in the interior of the convex hull and points on its boundary. To apply our techniques to the whole set rather than to the two subsets seems to be a promising approach for strengthening our result. To treat all points at once, would make extensive case distinctions necessary. On the other hand, it is very well conceivable that this effort would be rewarded with an upper bound that is close to 9, the upper bound we obtained for the number of shattered points on the boundary of the convex hull.

In the case of perimeter visibility domains it is much harder to imagine how further progress could be reached with the methods we employed: Our relativization to intervals is already very strong, and we also decomposed the point set in very small sets.

In Chapter 4 we initiated the research on a very natural special case of the Barrier Resilience problem. In the case of simple polygons we could show that the shortest path is always seen by a minimum number of witnesses. The reason for this lies in the topological structure of simple polygons. We showed that in the case of polygons with one convex hole we can also easily compute an optimal path. In the other direction we showed that the Barrier Resilience problem for visibility polygons in polygons with holes is APX-hard. The approximation factor we were able to show is greater than factors shown in the research on barrier resilience of other geometric objects such as line segments. We also considered three dimensional settings for the barrier resilience of visibility domains and showed their connections to the Minimum Neighborhood Path problem and the Minimum Color Path problem. In particular, we gave simple approximation algorithms for finding minimum neighborhood paths in planar graphs and minimum witness paths in a setting on a convex polyhedron and we could show that finding a minimum witness path in three dimensions with polygonal obstacles is hard to approximate in a strong sense.

The study of barrier resilience problems is a very interesting emerging field in which very few tight results have been shown so far. It would be a challenging task to design approximation algorithms for the barrier resilience of classes of visibility domains as well as for classes of other geometric objects.

Bibliography

[1] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.

[2] Helmut Alt, Sergio Cabello, Panos Giannopoulos, and Christian Knauer. Minimum cell connection and separation in line segment arrangements. *CoRR*, abs/1104.4618, 2011.

[3] Tetsuo Asano. An efficient algorithm for finding the visibility polygon for a polygonal region with holes. *IEICE TRANSACTIONS*, 68(9):557–559, 1985.

[4] David Avis and Godfried T. Toussaint. An efficient algorithm for decomposing a polygon into star-shaped polygons. *Pattern Recognition*, 13(6):395–398, 1981.

[5] Sergey Bereg and David G. Kirkpatrick. Approximating barrier resilience in wireless sensor networks. In *ALGOSENSORS*, pages 29–40, 2009.

[6] Robert D. Carr, Srinivas Doddi, Goran Konjevod, and Madhav Marathe. On the red-blue set cover problem. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, SODA '00, pages 345–353, Philadelphia, PA, USA, 2000. Society for Industrial and Applied Mathematics.

[7] David Yu Cheng Chan and David G. Kirkpatrick. Approximating barrier resilience for arrangements of non-identical disk sensors. In *ALGOSENSORS*, pages 42–53, 2012.

[8] Bernard Marie Chazelle. *Computational geometry and convexity*. PhD thesis, 1980.

[9] Wei-Pang Chin and Simeon Ntafos. Optimum watchman routes. *Information Processing Letters*, 28:39–44, 1988.

[10] Vašek Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39–41, 1975.

[11] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.

[12] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.

[13] Stephan Eidenbenz. Inapproximability results for guarding polygons without holes. In *ISAAC*, pages 427–436, 1998.

[14] Stephan Johannes Eidenbenz. *(In-) Approximability of Visibility Problems on Polygons and Terrains*. PhD thesis, Swiss Federal Institute of Technology Zurich, 2000.

[15] Ulrich Finke and Klaus Hinrichs. Overlaying simply connected planar subdivisions in linear time. In *Symposium on Computational Geometry*, pages 119–126, 1995.

[16] Steve Fisk. A short proof of chvátal's watchman theorem. *J. Comb. Theory, Ser. B*, 24(3):374, 1978.

[17] Laxmi Gewali, Alex Meng, Joseph S. B. Mitchell, and Simeon Ntafos. Path planning in $0/1/\infty$ weighted regions with applications. In *Symposium on Computational Geometry*, pages 266–278, 1988.

[18] Subir K. Ghosh. *Visibility Algorithms in the Plane*. Cambridge University Press, New York, 2007.

[19] Matt Gibson, Gaurav Kanade, and Kasturi R. Varadarajan. On isolating points using disks. In *ESA*, pages 61–69, 2011.

[20] Alexander Gilbers. Barrier resilience of visibility polygons. In *Proceedings of the Middle-European Conference on Applied Theoretical Computer Science*, 2013. To appear.

[21] Alexander Gilbers and Rolf Klein. New results on visibility in simple polygons. In *Proceedings of the 11th International Symposium on Algorithms and Data Structures*, WADS '09, 2009.

[22] Alexander Gilbers and Rolf Klein. A new upper bound for the VC-dimension of visibility regions. In *Proceedings of the 27th Symposium on Computational Geometry*, SoCG '11, pages 380–386, 2011.

[23] Alexander Gilbers and Rolf Klein. A new upper bound for the VC-dimension of visibility regions. *Computational Geometry: Theory and Applications*, 47(1):61–74, 2014.

[24] Leonidas Guibas, John Hershberger, Daniel Leven, Micha Sharir, and Robert Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2:209–233, 1987.

[25] David Haussler and Emo Welzl. Epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.

[26] Volkan Isler, Sampath Kannan, Kostas Daniilidis, and Pavel Valtr. VC-dimension of exterior visibility. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(5):667–671, 2004.

[27] Gil Kalai and Jiří Matoušek. Guarding galleries where every point sees a large area. *Israel Journal of Mathematics*, 101:125–140, 1997.

[28] Michael Kearns and Umesh Virkumar Vazirani. *An introduction to computational learning theory.* The MIT Press, 1994.

[29] James King. VC-dimension of visibility on terrains. In *Proceedings of the 20th Canadian Conference on Computational Geometry*, CCCG'08, pages 27–30, 2008.

[30] James King and David Kirkpatrick. Improved approximation for guarding simple galleries from the perimeter. *Discrete & Computational Geometry*, 46(2):252–269, 2011.

[31] David Kirkpatrick. Guarding galleries with no nooks. In *Proceedings of the 12th Canadian Conference on Computational Geometry*, CCCG'00, pages 43–46, 2000.

[32] R. Klein. *Algorithmische Geometrie: Grundlagen, Methoden, Anwendungen.* EXamen.press Series. Springer, 2005.

[33] János Komlós, János Pach, and Gerhard J. Woeginger. Almost tight bounds for epsilon-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.

[34] Santosh Kumar, Ten H Lai, and Anish Arora. Barrier coverage with wireless sensors. In *Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 284–298. ACM, 2005.

[35] D. T. Lee. Visibility of a simple polygon. *Computer Vision, Graphics, and Image Processing*, 22(2):207–221, 1983.

[36] D. T. Lee and Arthur K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[37] Jiří Matoušek. *Lectures on Discrete Geometry.* Springer-Verlag, New York, 2002.

[38] Joseph S.B. Mitchell and Christos H. Papadimitriou. The weighted region problem: finding shortest paths through a weighted planar subdivision. *Journal of the ACM*, 38(1):18–73, 1991.

[39] Simeon Ntafos. The robber problem. *Information Processing Letters*, 34:59–63, 1990.

[40] Joseph O'Rourke. The complexity of computing minimum convex covers for polygons. In *Proc. 20th Allerton Conf. Commun. Control Comput.*, pages 75–84, 1982.

[41] Joseph O'Rourke. *Art gallery theorems and algorithms.* Oxford University Press, New York, 1987.

[42] Joseph O'Rourke and Kenneth J. Supowit. Some np-hard polygon decomposition problems. *IEEE Transactions on Information Theory*, 29(2):181–189, 1983.

[43] Michael Ian Shamos. *Computational geometry.* PhD thesis, New Haven, CT, USA, 1978. AAI7819047.

[44] Kuan-Chieh Robert Tseng and David G. Kirkpatrick. On barrier resilience of sensor networks. In *ALGOSENSORS*, pages 130–144, 2011.

[45] Jorge Urrutia. Art gallery and illumination problems. In Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational Geometry*, pages 973–1027. North-Holland, Amsterdam, 2000.

[46] Pavel Valtr. Guarding galleries where no point sees a small area. *Israel Journal of Mathematics*, 104:1–16, 1998.

[47] Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.

[48] Shengli Yuan, S. Varma, and J.P. Jue. Minimum-color path problems for reliability in mesh networks. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, volume 4, pages 2658–2669 vol. 4, 2005.