Decomposing Polygons into r-stars or α-boundable Subpolygons

A Thesis Submitted to the
College of Graduate Studies and Research
in Partial Fulfillment of the Requirements
for the degree of Master of Science
in the Department of Computer Science
University of Saskatchewan
Saskatoon

By

Chris Worman

Permission to Use

Abstract

To make computations on large data sets more efficient, algorithms will frequently divide information into smaller, more manageable, pieces. This idea, for example, forms the basis of the common algorithmic approach known as Divide and Conquer. If we wish to use this principle in planar geometric computations, however, we may require specialized techniques for decomposing our data. This is due to the fact that the data sets are typically points, lines, regions, or polygons. This motivates algorithms that can break-up polygons into simpler pieces. Algorithms that perform such computations are said to compute *polygon decompositions*. There are many ways that we can decompose a polygon, and there are also many types of polygons that we could decompose. Both applications and theoretical interest demand algorithms for a wide variety of decomposition problems.

In this thesis we study two different polygon decomposition problems. The first problem that we study is a polygon decomposition problem that is equivalent to the Rectilinear Art Gallery problem. In this problem we seek a decomposition of a polygon into so-called r-stars. These r-stars model visibility in an orthogonal setting. We show that we can compute a certain type of decomposition, known as a Steiner-cover, of a simple orthogonal polygon into r-stars in polynomial time. In the second problem, we explore the complexity of decomposing polygons into components that have an upper bound on their size. In this problem, the size of a polygon refers to the size of its bounding-box. This problem is motivated by a polygon collision detection heuristic that approximates a polygon by its bounding-box to determine whether an exact collision detection computation should take place. We show that it

is NP-complete to decide whether a polygon that contains holes can be decomposed into a specified number of size-constrained components.

Acknowledgements

This thesis would not have been possible without the supervision of Dr. Mark Keil. He provided invaluable insight into each of the problems dealt with in this thesis. He also helped me *nearly* solve many problems that do not appear in thesis, but still plague me to this day.

There are many members of the faculty from the Department of Computer Science that indirectly helped me with this thesis. In particular, Dr. Grant Cheston and Dr. Jean-Paul Tremblay have taught me nearly everything I know about fundamental computer science concepts. I would also like to thank the graduate students, and especially Tzvetalin Vassilev and Mark Watson, for providing an interesting workplace.

I must also thank Dr. David Mould, Dr. Chris Soteros, and (again) Dr. Grant Cheston for being on my thesis committee. Each committee member made important comments that affected the final draft of this thesis.

I also need to thank all the members of my family for supporting me during my post-secondary education. Finally, I need to thank Geneise for putting up with me while I wrote this thesis. I love you.

Contents

List of Figures

Chapter 1

Introduction

1.1 Overview of Results

A *decomposition* of a polygon P is a set of polygons whose geometric union is exactly P. A polygon decomposition problem is typically of the form: Given a polygon P with property X, compute a decomposition of P where each member of the decomposition has property Y. Perhaps the most intensely studied decomposition problem is polygon triangulation. In this problem the polygon P is typically a simple polygon, and the decomposition is a set of non-overlapping triangles. In this thesis we study the computational aspects of two decomposition problems.

The first problem that we study is a polygon decomposition problem that is equivalent to the Rectilinear Art Gallery problem. In this problem we seek a minimum cardinality decomposition of a polygon into so-called r-stars. These r-stars model visibility in an orthogonal setting. We show that even if we allow overlap among the members of the decomposition, we can still compute such decompositions in polynomial time.

In the second problem we explore the complexity of decomposing polygons into components that have an upper-bound on their size. We consider such decompositions on polygons that contain holes. We show that several variants of this problem are NP-complete.

1.2 Motivation

Polygons can be used to model many different kinds of objects. In computer graphics, for example, polygons are used as the building blocks for objects in a virtual world. Thus it is not surprising that motivation for designing polygon decomposition algorithms often comes from from other areas of computer science. These include pattern recognition, computer graphics, VLSI layout and design, data compression, database systems, and image processing [26].

Although decomposition problems are often studied for practical reasons, theoretical interest also motivates researchers. Like many combinatorial problems, a group of related decomposition problems may fall into quite different complexity classes. Locating the boundary where a problem changes from being hard to being tractable has found interest among the algorithmic research community (see, for example, [15]). In an attempt to locate this boundary for decomposition problems, researchers have considered artificial restrictions to the decomposition. In [11, 8] the authors show that even if we only need to decompose the *boundary* of the polygon, certain decomposition problems remain NP-complete.

It is this theoretical interest in decomposition problems that motivates the study of the first problem in this thesis. Art gallery problems form the core of a large area of research in combinatorial and computational geometry. This research was initiated when Victor Klee posed the now famous Art Gallery Problem. He asked how one positions the minimum number of guards needed to watch the interior of a n-wall art gallery. This problem was subsequently tackled by Vasek Chvátal. He proved that $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient and sometimes necessary to guard an n-wall art gallery that is modeled by a simple polygon. Since this result was discovered, many variations of this problem have been consider from both a combinatorial and an algorithmic point of view. If the walls of the art gallery are all either horizontal or vertical, for example, then $\lfloor \frac{n}{4} \rfloor$ guards are sometimes necessary and always sufficient to guard the art gallery [22]. O'Rourke provides the definitive compendium of results in this area [38]. Of particular interest are the algorithmic results regarding art gallery problems. In [2], Aggarwal shows that deciding where to position the minimum number of guards is NP-complete. Thus unless **P=NP**, any hope for feasible algorithmic solutions lies in restricting the general art gallery problem. This motivates the first problem that we study herein. In this problem we wish to compute the placement of guards in an art gallery where each wall is either horizontal or vertical. Furthermore, the guards have a restricted form of visibility known as *r-visibility* (see Section 1.3). As with the general art gallery problem, this restricted version can be formulated as a polygon decomposition problem.

The motivation for the second problem that we study comes from the area of computer graphics. Consider a real-time 2-dimensional animation where the objects have been modeled by polygons. At some point we may need to compute whether or not two objects overlap for the sake of collision detection. It would be beneficial if we could avoid this operation for polygons that are far apart. One way to achieve this is to use a bounding box heuristic [41]. Using this heuristic, we only compute whether two objects intersect if their bounding boxes intersect. Since computing whether a pair of bounding boxes intersect is trivial, we can often avoid relatively expensive object intersection computations. Zhou and Suri have shown that this heuristic performs better if the boundings boxes of each polygon have similar size. Unfortunately, the polygons used to represent the objects may have arbitrary size. This implies that their respective bounding boxes may have large size discrepancy. Damian [14] suggests a solution to this problem: if we decompose the polygons into pieces that all have an upper bound on their size then we have a set of polygons whose bounding boxes are more likely to have similar size. Since this operation will likely increase the number of polygons we need to consider, minimizing the cardinality of the decomposition becomes important. We will consider the problem of finding a decomposition where the members of the decomposition all have an upper bound on their size.

1.3 Types of Polygons

We will consider geometric computations that take place in the plane. Since we are dealing with polygon decomposition problems, the most important geometric object to consider is the polygon. Before considering the various relevant types of polygons, we pause to consider some useful notation.

We will often need to consider different operations on polygons and other geometric objects. These operations will correspond to set theoretic operations such as union, intersection, and difference. For the sake of brevity, we will use the set theoretic operator symbols on polygons. For two polygons P_1 and P_2, we will write their union, intersection, and difference as $P_1 \bigcup P_2$, $P_1 \bigcap P_2$, and P_1/P_2 respectively. We will also make use of the notation for set inclusion. When a point p is contained in a polygon P we will write $p \in P$. If some other geometric object g that is not a point (i.e. it is a line segment, a collection of points, another polygon, etc.) is in a polygon P then we write $g \subseteq P$.

1.3.1 Simple and Non-simple Polygons

We will define a polygon to be a closed region in the plane defined by a sequence of *edges*. We define a polygon as a *closed* region (as opposed to an *open* region) so that any point on the boundary is considered to be within the interior of the polygon. This definition will be important when discussing decompositions defined in terms of visibility. Edges are defined by a sequence of two *vertices*. We will often refer to the edges of a polygon as $E(P)$ and the vertices as $V(P)$. Since we have defined polygons as closed regions, we have that for all $v \in V(P)$, $v \in P$. Similarly, for each $e \in E(P)$, $e \subset P$. A vertex is called *reflex* if its internal angle with P is greater than $180°$ (Figure 1.1). If each edge of a polygon is perpendicular to one of the coordinate axis then the polygon is called *orthogonal* or *rectilinear*. If no non-consecutive pair

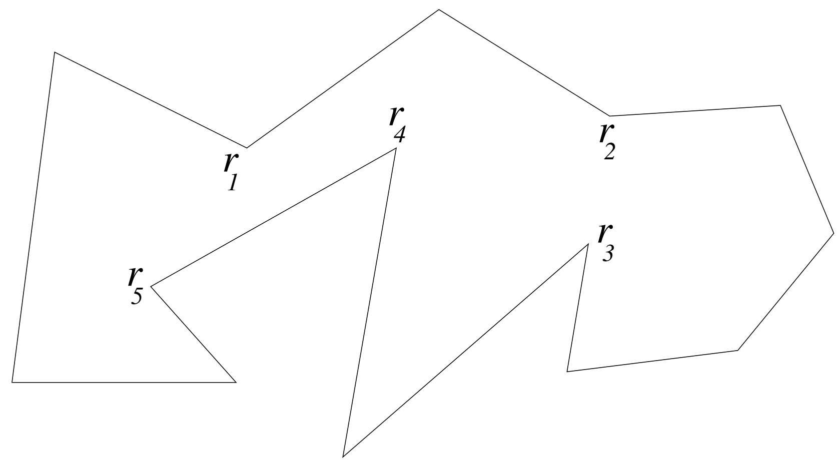

Figure 1.1: A simple polygon with 5 reflex vertices r_1, r_2, r_3, r_4, and r_5.

of edges overlap then the polygon is said to be *simple*. We will consider a polygon with holes to be non-simple. A *polygon with holes* is an otherwise simple polygon that has simple polygons "removed" from its interior (Figure 1.2).

1.3.2 Visibility Induced Polygons

We can further classify both simple and non-simple polygons based upon other geometric properties. Visibility is a property defined for two points with respect to some polygon. Typically we say that two points p and q are *visible* (or p *sees* q) in P if the line segment $\overline{pq}$ lies entirely within P. Since other relevant notions of visibility exist, we will call this notion of visibility *l-visibility* or *line* visibility. We also recognize two other forms of visibility. Two points are *r-visible* if the orthogonal[1] bounding rectangle for p and q lies within P. The last relevant definition of visibility, called *s-visibility*, uses the notion of a staircase path. A *staircase path* is an orthogonal chain of edges with bends that alternate between *exactly two* orthogonal orientations. We say p and q are *s-visible* if there is a staircase path starting at p and ending at q that is entirely within the interior of P. We can use these different notions of visibility to define various types of polygons. Of particular

[1]A rectangle is orthogonal if each of its edges are perpendicular to one of the coordinate axis.

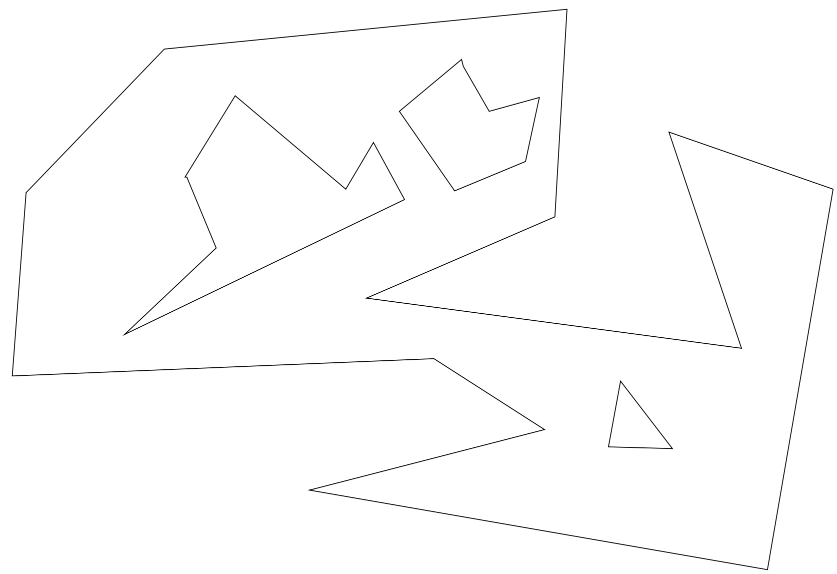

Figure 1.2: A polygon with 3 holes.

interest are so-called *star-shaped* or simply *star* polygons. A polygon P is an x-*star*, where $x \in \{l, s, r\}$, if there exists a point $k \in P$ such that for each point $q \in P$, q is x-*visible* from k. The set of all points k from which the polygon is x-visible is known as a *kernel* of the x-star (see Figure 1.3). Since l-star polygons are the most common, we will simply refer to them as star-shaped polygons. There is another commonly used type of polygon defined in terms of visibility. A polygon is x-*convex* if for each pair of points $p, q \in P$, p is x-visible from q, where $x \in \{l, s, r\}$. Since we are only concerned with l-convex polygons, we will use the term convex to mean l-convex.

It will often be useful to extend the notion of visibility to other geometric objects besides points. We say that two geometric objects g_1 and g_2 are x-visible for $x \in \{l, s, r\}$ iff for all points $p_1 \in g_1$ and $p_2 \in g_2$ we have that p_1 x-sees p_2.

We can also define subpolygons in terms of visibility. We define the x-*visibility polygon* of a point k in polygon P to be the subpolygon consisting of the set of points that are x-visible from k within P, where $x \in \{l, s, r\}$. We will denote the x-visibility polygon for a point k as $\nu_x(k)$. We note here that $\nu_x(k)$ will always be an x-star with k in the kernel.

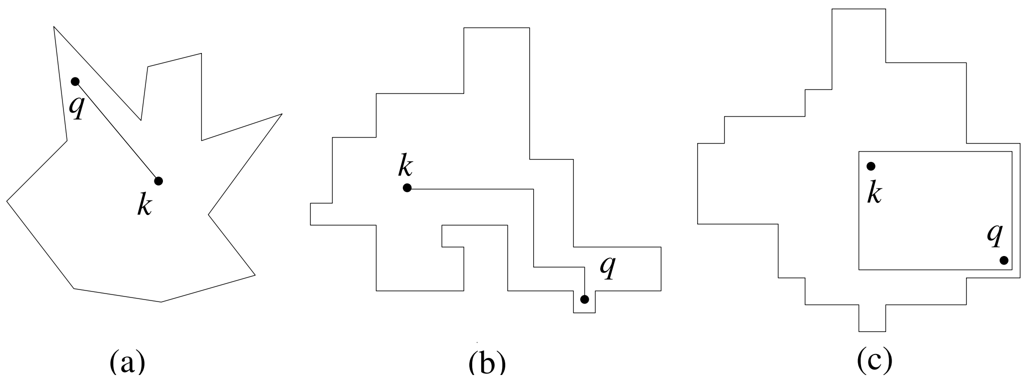

Figure 1.3: The different types of stars (a) An l-star or star-shaped polygon (b) An s-star (c) An r-star

1.3.3 Size-bounded Polygons

There are also ways of classifying polygons based upon some measure of size. A common example of this type of measure is the *area* of a polygon. Another example is the *diameter* of a polygon, which is traditionally defined to be the largest distance between any two vertices of the polygon. A different definition of diameter is defined and explored in [14]. Here the diameter of a polygon is the diameter of the smallest circle that can circumscribe the polygon. A very similar definition of diameter is the longest side-length of the smallest orthogonal rectangle that can bound the polygon (Figure 1.4). So as not to conflict with the traditional "farthest vertices" definition of diameter, we will refer to the "circle-bounded" or "rectangle bounded" definition simply as the *span* of a polygon. It will always be clear from the context which notion of span (i.e. either circle or rectangle bounded) is currently being used. If a polygon has a span no larger than α then we say the polygon is *α-boundable*. We will refer to subpolygons that are α-boundable as *α-subpolygons*.

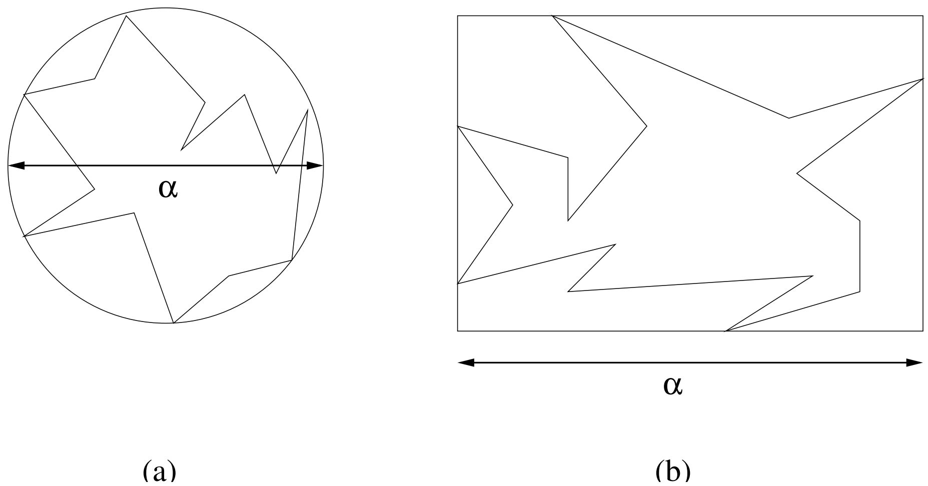

Figure 1.4: α-boundable polygons (a) Circle bounded (b) Rectangle bounded

1.3.4 Classes of Orthogonal Polygons

Culberson and Reckhow [10, 9] propose a method for classifying an orthogonal polygon based upon the types of *dents* it has. A *dent* is an edge whose endpoints are both reflex vertices of the polygon. Four types of dents are identified based upon their orientation (Figure 1.5). Imagine aligning the polygon with a compass such that north (N) extends upwards along the positive y-axis. We define a dent to be *north*, or simply N, if while traversing the polygon clockwise, we travel along the dent from west to east. The other dent orientations (i.e. E, S, and W) are defined analogously so that dent orientations have a correspondence with the compass directions (see Figure 1.5). A *class k* orthogonal polygon is defined to have k different dent orientations. The class 2 orthogonal polygons can be further subdivided. If the two dent orientations are parallel along the compass (i.e. {N,S} and {E,W}) we say the polygon is class 2a. All other class 2 polygons are deemed to be class 2b. These classes of orthogonal polygons have different algorithmic properties that will be explored in Section 1.5.

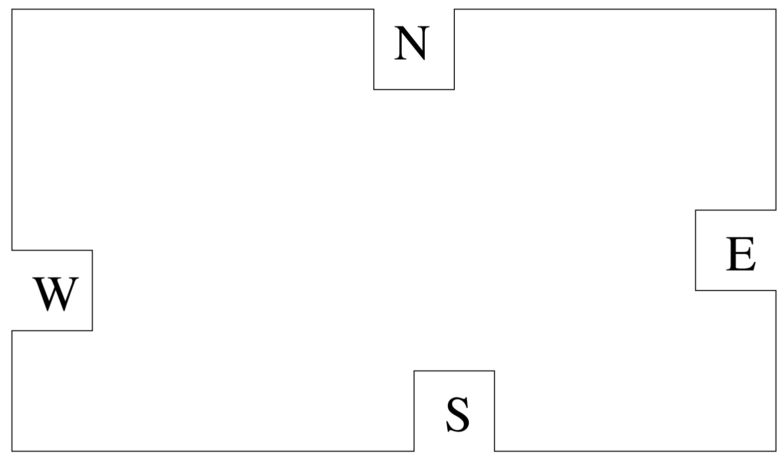

Figure 1.5: The four different dent orientations for orthogonal polygons

1.4 Types of Decompositions

A *decomposition* of a polygon P is a set of polygons whose geometric union is exactly P. The members of the decomposition are often referred to as *subpolygons* or *components*. Decompositions are typically categorized by restrictions placed on their components (Figure 1.6). If the components are allowed to overlap, then we call the decomposition a *cover*. If the components are mutually disjoint (except along boundaries) then we have a *partition*. There is also a distinction made among decompositions based upon how their components are formed. If the set of vertices used by all the components is a subset of $V(P)$, then the decomposition is said to be *Steiner-free*. *Steiner decompositions* are allowed to use *Steiner points*, which are points not in $V(P)$. We note here that a partition of a polygon is always a valid cover, while a cover may not be a valid partition. Similarly, a Steiner-free decomposition is always a valid Steiner-decomposition, while the converse does not necessarily hold.

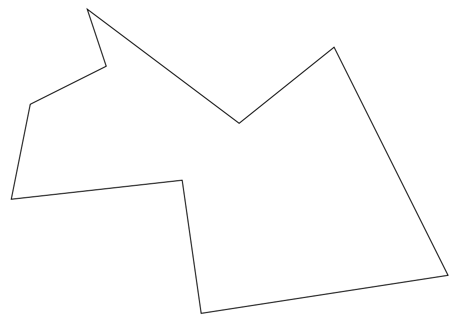

(a)

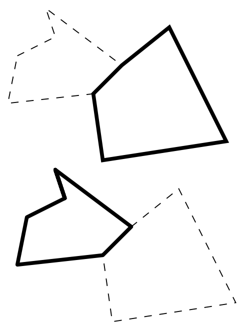

(b)

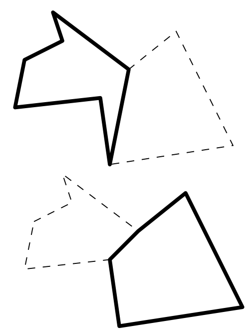

(c)

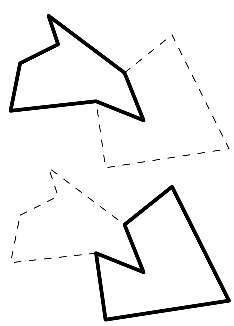

(d)

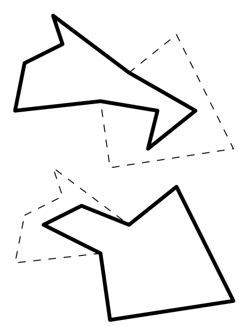

(e)

Figure 1.6: The different categories of decompositions. The members of the decompositions have been drawn with solid lines. (a) A simple polygon P (b) A Steiner-free partition of P (c) A Steiner-free cover of P (d) A Steiner-partition of P (e) A Steiner-cover of P

1.5 Algorithms and Complexity

Researchers study decomposition problems by placing restrictions on the subject polygon and the components of the decomposition. Since there are many types of polygons, researchers have been able to study a wide range of decomposition problems. In [26], Keil explores the results that have been discovered concerning polygon decomposition. Many of these results concern *minimizing* the cardinality of decompositions. This is not surprising since decomposition is an operation that usually increases the number of objects to consider. In this context, coverings are better than partitions since the minimum cardinality cover is at least as small as the minimum cardinality partition. This is because any partition is a valid cover, while a cover may or may not be a valid partition. This observation is reflected in the fact that covers are typically harder to compute than partitions. Finding a Steiner-free partition of a simple polygon into the minimum number of star-shaped polygons, for example, is known to be computable in polynomial time [24], while the equivalent covering problem is NP-complete [2]. This disparity between partitioning and covering simple polygons forms a trend in the literature; partitioning is often tractable, while covering tends to be NP-hard. There are, however, some interesting counterexamples to this observation.

Covering *orthogonal* polygons with certain types of stars has proven to be tractable, even if Steiner-points are allowed. These algorithms take advantage of properties that are exhibited by different classes of orthogonal polygons (see Section 1.3). When the components are restricted to s-stars, Culberson and Reckhow [10] have developed an $O(n^2)$ time algorithm for finding the minimum cardinality cover of class 2 orthogonal polygons. In [39, 36] a powerful tool known as the *visibility graph* (see Section 2.3) is used to tackle class $k > 2$ polygons. The authors develop an $O(n^3)$ time algorithm for class 3 polygons, while class 4 polygons take $O(n^8)$ time to decompose using their techniques. Some results have also been discovered for covering with r-stars. Keil [25] shows how to minimally cover the class 2a orthogonal polygons with r-stars in

$O(n^2)$ time. This result is subsequently improved in [16], where an $O(n)$ time algorithm is provided. These results suggest that there may be hope in creating efficient algorithms for covering other classes of orthogonal polygons with r-stars.

Many objects can naturally be modeled by polygons with holes. They are used, for example, to represent a region that contains obstacles in many motion planning [6] and art gallery problems [1, 19]. Despite their usefulness in modeling, polygons with holes have proven to be difficult to decompose. Partitioning or covering with Steiner-free convex or star-shaped polygons is NP-complete for polygons with holes [23, 29, 37]. The situation is similar even if we restrict the subject polygon to being orthogonal. Finding a Steiner-cover of an orthogonal polygon that may contain orthogonal holes with rectangles was shown to be NP-complete by Masek [33]. Perhaps the only positive results concern partitioning orthogonal polygons into rectangles. Even if the polygon contains orthogonal holes, polynomial time algorithms exist for finding the minimum cardinality Steiner-partition [31, 20, 30]. Thus there exists an interesting relationship between covering polygons with holes and partitioning polygons with holes.

Decomposition problems have traditionally focused on restricting the components by *shape*. Examples of this type of restriction include convex, star-shaped, spiral, or monotone subpolygons. Although these restrictions induce decompositions containing polygons that are *simpler*, they are not necessarily *smaller*. Damian and Pemmaraju have introduced a decomposition that fills this void [14, 12, 13]. This decomposition, which is known as an α-*decomposition* (i.e. α-*partition* or α-*cover*), contains components that all have a span less than or equal to $\alpha \in \mathbb{R}$. Damian and Pemmaraju's results all use the circle-bounded definition of span and most of their results are for partitioning polygons into α-boundable components. A 4.5-approximation algorithm is given for minimally partitioning a convex polygon with Steiner points. The time complexity of this algorithm is not given. For Steiner-free partitions, a 3-approximation algorithm is given for approximating a minimum partition of a simple polygon into α-subpolygons. The time taken by this algorithm is $O(mn)$, where n is the number of vertices and m is the number of edges in the vertex visibility graph of P[2]. Without increasing the running time, this algorithm can be extended to produce partitions whose members are all convex. An $O(mn^3)$ algorithm is given for finding the minimum cardinality α-partition of a simple polygon. This algorithm can be modified to find partitions with components that are both α-boundable and convex. An interesting approximation scheme is presented in [12]. Let OPT be the cardinality of a minimum α-partition of some convex polygon P. Damian and Pemmaraju provide an $O(n(1+\frac{\alpha}{\epsilon})^8)$ time algorithm that produces an $(\alpha+\epsilon)$-partition of a convex polygon P that is of size at most OPT. This research opens the door for studying the computational complexity of computing α-decompositions for different types of polygons.

[2]The number of edges in the visibility graph (i.e. m) is $O(n^2)$.

Chapter 2

Covering Simple Orthogonal Polygons with r-stars

2.1 Introduction

In this chapter we study a polygon decomposition problem that is a restatement of an Art Gallery problem. The Art Gallery problem that we investigate is a version of the rectilinear art gallery problem. As already mentioned in Section 1.2, the Art Gallery problem seeks the placement of the minimum number of guards needed to watch an n-wall art gallery. Specifically, we want to place the guards so that each point in the gallery is visible from at least one guard (Figure 2.1). There are many factors that affect the interpretation of this problem. We can, for example, place restrictions on how a guard views its surroundings. In the most general version of this problem, no restriction is placed on the vision of the guards except that they cannot see through walls. We can also place restrictions on the layout of the art gallery. For instance, we could consider the case where the art gallery is littered with visual obstacles. We will study an art gallery problem with restrictions placed on the vision of the guards and on the orientation of the gallery walls.

Definition 1. *We say that a set of points G, which are known as guards, can watch a polygon P if for each point $q \in P$, there exists some guard $g \in G$ such that g sees q. In the* **Rectilinear Art Gallery** *problem we wish to compute a minimum set of*

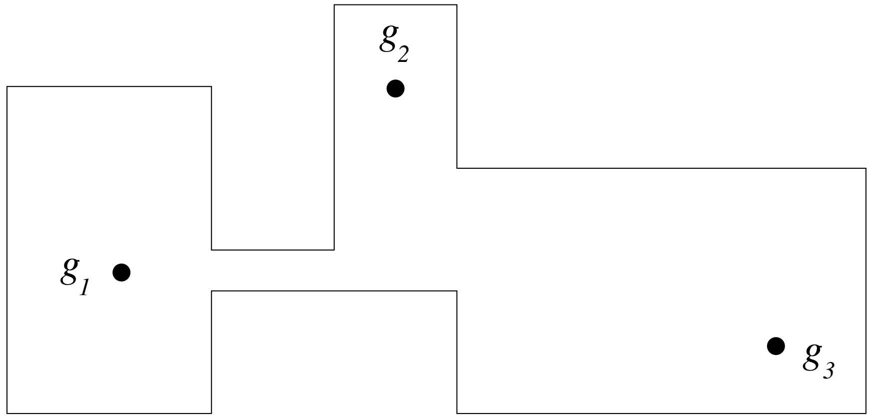

Figure 2.1: An example of an art gallery being watched by three guards with r-visibility.

guards that can watch a simple orthogonal polygon P with r-visibility.

This problem has already been studied in [10] and mentioned in [39], but no solution is known. In [25, 16], the authors solve a simpler version of this problem where P is restricted to be class 2a. This result relies heavily on geometric properties that are specific to class 2a polygons, and is therefore not applicable to guarding a general orthogonal polygon. We can, however, use the general approach used by the authors that have studied similar problems. We will formulate the **Rectilinear Art Gallery** problem as a polygon decomposition problem. Consider a set G of guards that can watch an orthogonal polygon P with r-visibility. For each guard $g \in G$ there is an associated visibility polygon $\nu_r(g)$. Recall that $\nu_r(g)$ will be an r-star. Since the guards in G can watch P, we know that the collection of visibility polygons for the guards must cover P. Since $\nu_r(g)$ is an r-star for any guard g, the cover associated with the guards is a Steiner r-star cover. Now consider a Steiner r-star cover of P. We can construct a set of guards that can watch P as follows. For each r-star in the cover, we associate a guard that is positioned in the kernel of the r-star. Thus the **Rectilinear Art Gallery** problem is equivalent to the problem of finding a minimum cardinality r-star cover of P. Now we can focus on the following polygon decomposition problem.

Definition 2. *In the* **Minimum r-star Cover** *problem we seek a minimum cardinality Steiner-cover of an orthogonal polygon, where each of the subpolygons in the cover is an* r*-star.*

This problem is known to be in **P** if the subpolygons are s-stars rather than r-stars [39, 36, 10]. In [10], the authors present an $O(n^2)$ time algorithm for the **Minimum** s**-star Cover** problem when P is restricted to be a class 2 orthogonal polygon. Motwani et al. provide algorithms for higher order orthogonal polygons [39, 36]. For class 3 polygons, they develop an $O(n^3)$ time algorithm. Using similar techniques, they also show that class 4 orthogonal polygons can be optimally covered with s-stars in $O(n^8)$ time. Like many other algorithms for covering orthogonal polygons, the results in [39, 36] use a graph theoretic transformation. In this technique, a graph is constructed, where nodes correspond to important regions within the polygon. Two nodes are adjacent in the graph if their corresponding regions have some type of visibility relationship. If this graph is constructed carefully then it will contain all the relevant information needed to find a cover of the polygon. Specifically, the original *polygon covering* problem will be equivalent to a *graph clique covering* problem on the graph that was constructed. Without further insight, this approach is not useful: graph clique covering problems are very difficult to solve and are typically NP-complete. We can only use this idea if the graph has some properties that we can exploit in order to create an efficient graph clique covering algorithm.

Fortunately, the graph constructed in [39] has the necessary properties to be a useful tool. A minimum clique cover of this graph corresponds to a minimum cover of P with s-stars. The graph is also shown to belong to the class of *perfect* graphs. These graphs have properties that allow for efficient clique cover computations. In this chapter we develop a polynomial-time algorithm that solves the **Minimum r-star Cover** problem by using a similar graph theoretic transformation.

Visibility Symbol	Meaning
$p \equiv q$	p sees q
$p \wedge q$	p indirectly sees q
$p \overline{\wedge} q$	p strictly indirectly sees q
$p \not\wedge q$	p does not indirectly see region q

Figure 2.2: The various types of visibility.

2.2 Visibility

For the remainder of this chapter we will be using r-visibility exclusively. Because of this, *visibility* will always refer to *r-visibility* from now on. We will adopt some notation regarding visibility that is similar to that found in [39]. Let p and q be two geometric objects (i.e. points, lines, rectangles, etc.). If p is visible from q then we write $p \equiv q$ (read *p sees q*). If p and q both see at least one point $k \in P$ then we write $p \wedge q$ (read *p indirectly sees q*). We will refer to the set of all points that both p and q see as $\mathcal{H}(p, q)$. We will often refer to this set as p and q's *helper region*. If $p \wedge q$ and p does not see q then we write $p \overline{\wedge} q$ (read *p strictly indirectly sees q*). If there is no point that p and q both see then we write $p \not\wedge q$ (read *p does not indirectly see q*). The various notions of visibility are summarized in the table of Figure 2.2. If $p \wedge q$ then we refer to a rectangle that covers p and at least one point from $\mathcal{H}(p, q)$ as $R^p_{\mathcal{H}(p,q)}$, and a rectangle that covers q and any set of points from $\mathcal{H}(p, q)$ will be called $R^q_{\mathcal{H}(p,q)}$. We further classify indirect visibility depending on how $R^p_{\mathcal{H}(p,q)}$ and $R^q_{\mathcal{H}(p,q)}$ may overlap. If for all $R^p_{\mathcal{H}(p,q)}$ and $R^q_{\mathcal{H}(p,q)}$ such that $R^p_{\mathcal{H}(p,q)} \bigcap R^q_{\mathcal{H}(q,p)} \neq \emptyset$, $R^p_{\mathcal{H}(p,q)} / R^q_{\mathcal{H}(q,p)}$ is *not* a rectangle, then we say p has indirect *diagonal* visibility with q. In all other cases we say that p has indirect *orthogonal* visibility with q. This distinction is depicted in Figure 2.3.

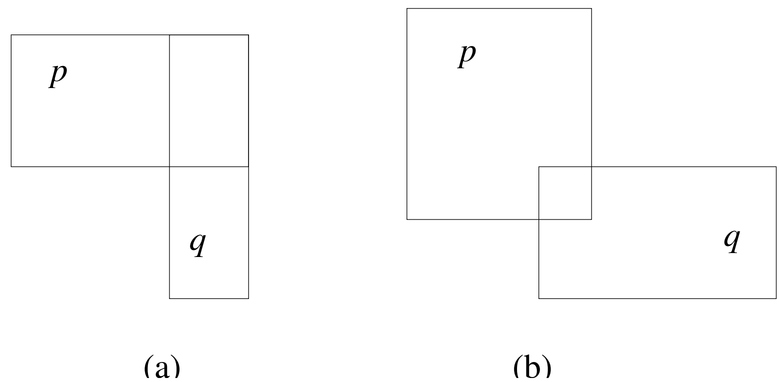

Figure 2.3: (a) Indirect orthogonal visibility. (b) Indirect diagonal visibility.

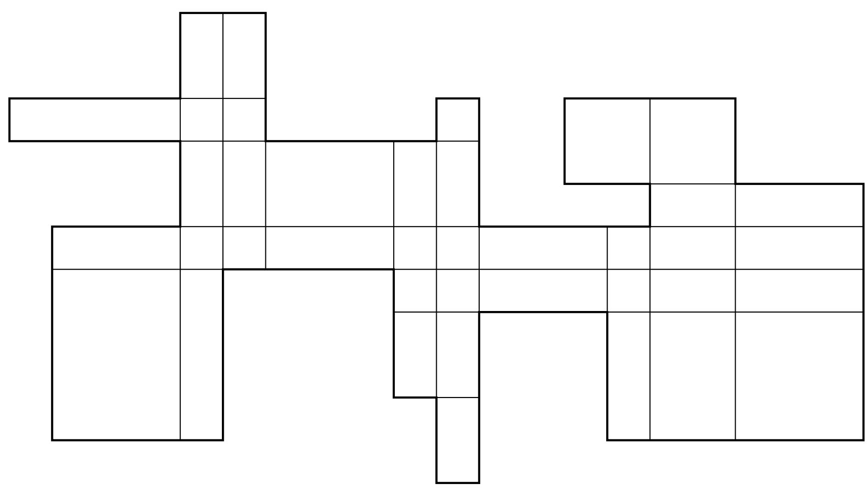

Figure 2.4: A simple orthogonal polygon partitioned into the basic regions that constitute the dent diagram.

2.3 The Region Visibility Graph

Now we define and discuss the graph that will be used in solving the **Minimum r-star Cover** problem. This graph will be built from P. Each node in the graph will correspond to a region within P. These regions must be atomic in the sense that they must represent groups of points that share the same visibility relationships with all the other points within P. We will use the notion of *dent diagrams* to generate these regions. *Dent diagrams* were developed by Culberson and Reckhow [10] as a tool for dealing with orthogonal polygon covering problems.

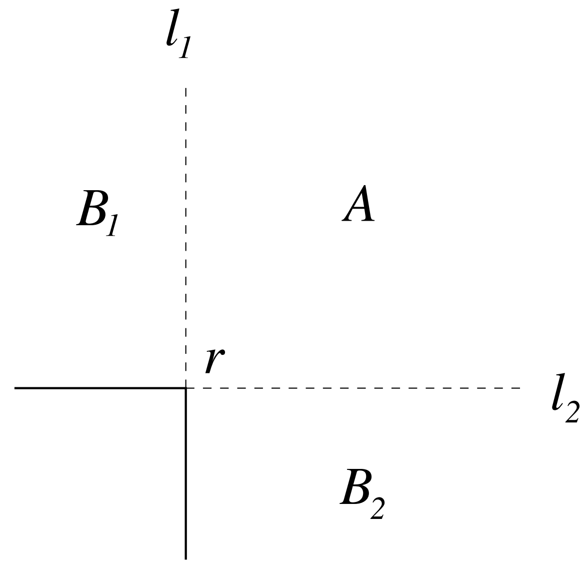

Figure 2.5: The regions surrounding the reflex vertex r.

Let P be a simple orthogonal polygon. To facilitate the identification of important regions within P, we will first partition P using a simple technique. For each reflex vertex of P we will extend two orthogonal lines. We will call these lines *dent lines*. Since these dent lines were extended from reflex vertices, they will eventually collide with the boundary of P. We will only consider the portion of these lines that are within P. The dent lines induce a subdivision within P that constitutes a partition of P into rectangles (Figure 2.4). We will refer to the rectangles in this partition as *basic regions*. The partition itself will be referred to as the *dent diagram* of P. Let us take a closer look at the regions directly surrounding a reflex vertex (Figure 2.5). The dent lines l_1 and l_2 extending from r will induce three regions that we will label A, B_1, and B_2. No point in region B_1 is directly visible to any point in B_2 since the rectangle implied by r-visibility would necessarily be partially outside of the interior of P. Unless prevented by some reflex vertex other than r, every point in region A sees every point in B_1 and B_2. Thus the set of points contained in a basic region will have the same visibility relationship to all other regions.

Now we turn our attention to explaining an important observation about helper regions for indirectly visible basic regions. The following Lemma will be useful in

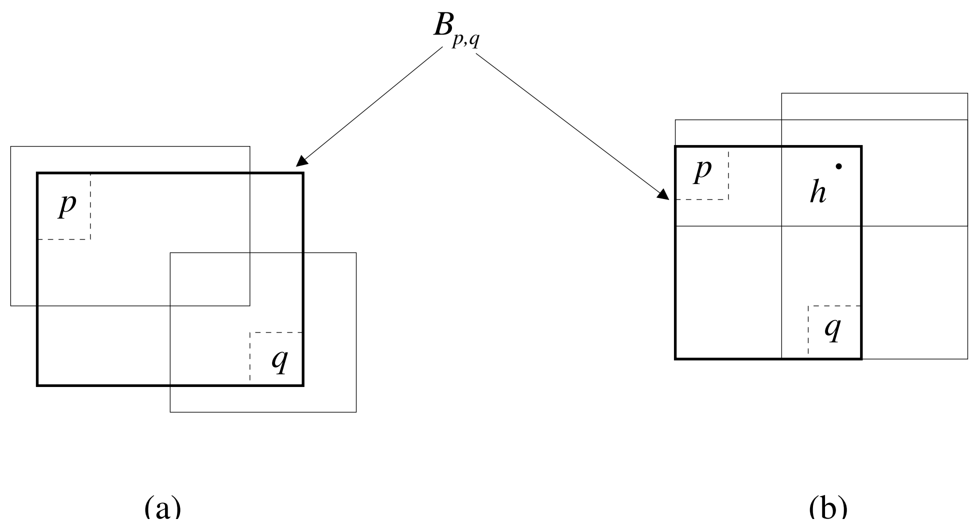

Figure 2.6: The regions p and q have been shown as dashed-line rectangles. (a) Regions p and q have indirect diagonal visibility. (b) Regions p and q have indirect orthogonal visibility.

subsequent proofs.

Lemma 1. *If p and q are two basic regions such that $p \wedge q$ then there is some portion of $\mathcal{H}(p,q)$ in the bounding-rectangle of p and q.*

Proof. The bounding-rectangle for p and q is the smallest orthogonal rectangle that contains both p and q. Notice that such a rectangle may not lie completely within P. We will refer to this bounding-rectangle as $B_{p,q}$. We must show that there is some portion of $\mathcal{H}(p,q)$ in $B_{p,q}$. If $p \equiv q$ then the Lemma obviously holds. Thus for the remainder of the proof we have that $p \overline{\wedge} q$. Since $p \overline{\wedge} q$ we can, without loss of generality, assume that p is above and to the left of q. We will consider two cases based upon the type of indirect visibility p and q have.

First suppose that region p and q have indirect diagonal visibility. Due to our assumption about the relative location of p and q, we have that p is to the left of and above $\mathcal{H}(p,q)$, while region q is to the right of and below $\mathcal{H}(p,q)$ (Figure 2.6(a)). This implies that $\mathcal{H}(p,q) \subseteq B_{p,q}$.

Now suppose that p and q have indirect orthogonal visibility. Due to our as-

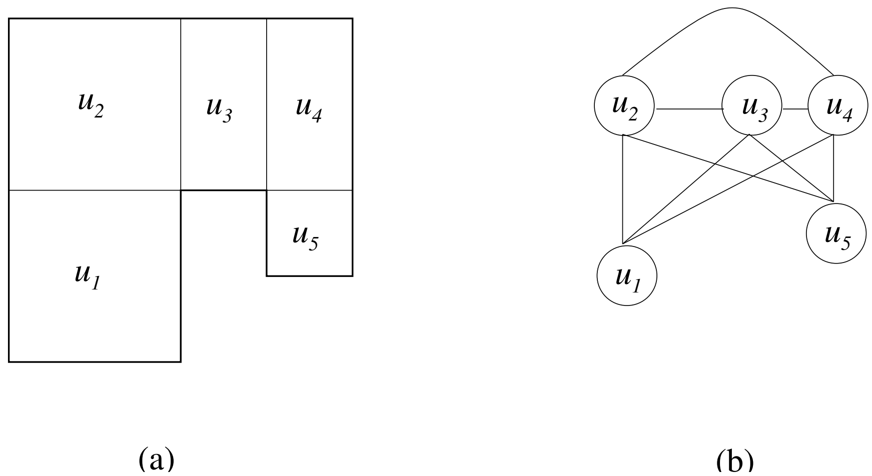

Figure 2.7: (a) An orthogonal polygon P partitioned into a dent diagram. (b) $RVG(P)$.

sumption about the relative location of p and q we have that there is some point $h \in \mathcal{H}(p, q)$ that is in the same column as q and the same row as p within the dent diagram of P, or in the same row as q and the same column as p in the dent diagram of P (Figure 2.6(b)). If no such point exists then p and q cannot have indirect orthogonal visibility. Due to the location of h, we have that $h \in B_{p,q}$, as required. $\square$

We will use the dent diagram to build the *region visibility graph*, denoted $RVG(P)$, for the polygon P (Figure 2.7). The nodes in this graph will be the basic regions in the dent diagram. Two nodes u_1 and u_2 will be adjacent in the region visibility graph iff $u_1 \wedge u_2$. In other words, u_1 and u_2 are adjacent iff there is a star that contains both u_1 and u_2.

Now consider the clique formed by nodes u_1, u_2, and u_3 from the graph in Figure 2.7 (b). This clique corresponds to a star formed by the union of regions u_1, u_2, and u_3 within P. Analogously, the star formed by regions u_1, u_2, and u_3 corresponds to a clique in the region visibility graph. In other words, cliques in the region visibility graph are equivalent to stars in P. A *clique cover* of a graph is a set C of cliques, where each node in the graph belongs to at least one of the cliques in C. Thus a minimum cardinality clique cover of $RVG(P)$ corresponds to a minimum star cover of P. Thus if we compute a minimum clique cover of $RVG(P)$, we can easily determine the minimum star cover of P. This relationship is essential if we are to use the region visibility graph to solve the **Minimum r-star Cover** problem. Although this property holds for the example given in Figure 2.7, it is not obvious that it will hold for an arbitrary orthogonal polygon P and its region visibility graph $RVG(P)$. Therefore we must show that the following Theorem holds if we are to use this graph theory approach to solving the **Minimum r-star Cover** problem.

Theorem 1. *Given a simple orthogonal polygon P and its region visibility graph RVG(P), there is a star containing a set of regions U iff the members of U form a clique in RVG(P).*

The remainder of this section is devoted to proving Theorem 1. The Theorem obviously holds in one direction: if there is a star covering a set of regions U, then there must be a clique in $RVG(P)$. This observation follows directly from the definition of $RVG(P)$.

The proof for the other direction requires further analysis. As per the results in [39], we will use a Theorem from topology [35] to assist in establishing Theorem 1. In the following, a cell denotes a simply connected subset of the plane. A subset C of the plane is simply connected if both C and $\mathbb{R}^2/C$ are connected.

Theorem 2. *[35] Let $\mathcal{C}$ be a set of cells in the plane. If $C \bigcap C'$ is a cell for every $C, C' \in \mathcal{C}$ and $C \bigcap C' \bigcap C'' \neq \emptyset$ for $C, C', C'' \in \mathcal{C}$, then $\bigcap \{C \in \mathcal{C}\} \neq \emptyset$.*

We will use Theorem 2 to show that the intersection of a set of pairwise overlapping visibility polygons is non-empty. Then we will show how we can use this result to prove Theorem 1.

In order to apply Theorem 2 to mutually overlapping visibility polygons, we need some basic results. First we need to show that if two visibility polygons overlap, then their intersection is a simply connected region. Thus we must show the following Lemma.

Lemma 2. *Let p and q be two basic regions from the dent diagram of P such that $\nu(p) \bigcap \nu(q) \neq \emptyset$. Then $\nu(p) \bigcap \nu(q)$ is a simply connected region.*

Proof. Assume for the sake of contradiction that $S = \nu(p) \bigcap \nu(q)$ is not a simply connected region. We will consider two cases.

Case(i): S contains at least one region with at least one hole. Let H be any of the regions in S that contains a hole. Let h be any of the holes that exist in H. By the definition of H, $H \subseteq \nu(p)$, which implies that h is also a hole in $\nu(p)$, which contradicts the simple connectedness of $\nu(p)$.

Case(ii): S is composed of two or more disjoint simply connected regions. Let u and v be two points from different components of S. Without loss of generality we will assume that u is above and to the left of v[1]. We will consider the locations of p and q relative to the subdivision induced by the orthogonal bounding box for points u and v. Let $B_{u,v}$ be the orthogonal bounding box of points u and v. We can subdivide P into 8 zones by extending the edges of $B_{u,v}$ (Figure 2.8). Region p (resp. q) is either in (or partially in) one of these zones, or it is in (or partially in) $B_{u,v}$. Before considering where regions p and q are located with respect to $B_{u,v}$, we make an important observation that will be critical to the rest of the proof. Let t, r, b, and l be the top edge, right edge, bottom edge, and left edge of $B_{u,v}$, respectively. If we show that both t and r (resp. b and l) are in both of $\nu(p)$ and $\nu(q)$ then we contradict our assumption that S is composed of disjoint regions. This is because the inclusion of t and r (resp. b and l) in both $\nu(p)$ and $\nu(q)$ implies that points u and v are connected in S. Thus our approach in what follows will be to show that either t and r or b and l are within $\nu(p)$ and $\nu(q)$ regardless of where regions p and q are located with respect to $B_{u,v}$. We also make the following observations. In the following, $s \in \{p, q\}$: (Recall that $u, v \in \nu(s)$, and $B_{s,u}$ and $B_{s,v}$ are contained within P.)

[1]This includes the case where u has the same x coordinate (resp. y coordinate) as v.

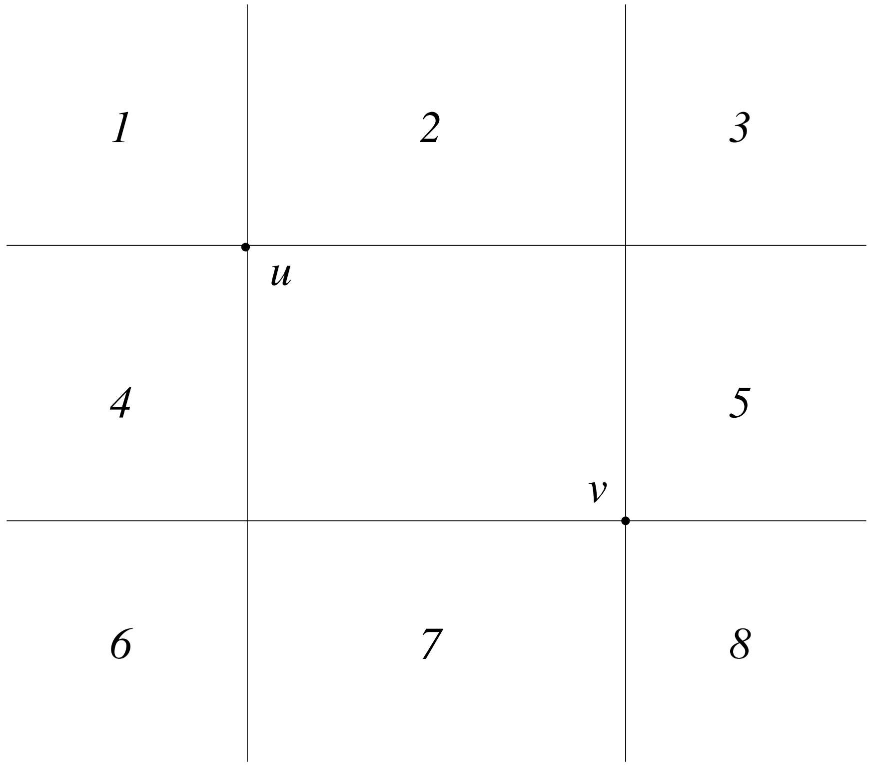

Figure 2.8: The 8 zones induced by the bounding box $B_{a,b}$.

- If s is in zone 1, then s sees all of t, r, b, and l.
- If s is in zone 2, then s sees t and r.
- If s is in zone 3, then s sees t and r.
- If s is in zone 4, then s sees l and b.
- If s is in zone 5, then s sees t and r.
- If s is in zone 6, then s sees l and b.
- If s is in zone 7, then s sees l and b.
- If s is in zone 8, then s sees all of t, r, b, and l.

These observations follow from the visibility relationships that exist among p, q, u, and v. The first observation, for example, follows from the fact that $B_{s,v}$ is contained within P and all of t, r, b, and l are contained within $B_{s,v}$ for $s \in \{p, q\}$. We will use these observations implicitly throughout the remainder of the proof. Now we will consider where regions p and q may be located with respect to $B_{u,v}$.

Subcase(i): Neither region p nor q is in $B_{u,v}$. Without loss of generality, we can assume that p is in zone 1, 2, or 3. We can immediately eliminate the possibility that region p is in zone 1 since no matter what zone q is in, u and v will be connected by either t and r or b and l. Thus p is in either zone 2 or zone 3. This implies that region q cannot be in zones 1, 2, 3, 5, or 8, since in each of these cases u and v are connected in S. Thus q must be in one of zone 4, 6, or 7, while region p is in zone 2 or 3. We can argue against these remaining locations for regions p and q using the following argument. Let $R_{q,b}$ be the smallest orthogonal sub-rectangle of P that contains both region q and the edge b. We are guaranteed that $R_{q,b}$ exists since $q \equiv v$ and $q \equiv u$. Since region p sees both t and r, and q sees l and b, we have that $B_{u,v}$ is free of edges from P. This is due to the fact that P does not contain holes. This implies that we can translate the top edge of $R_{q,b}$ upwards so that $R_{q,b}$ includes t and r. This implies that q sees t and r, and hence u and v are connected in S.

Subcase(ii): Exactly one of p or q is in $B_{u,v}$. Without loss of generality assume that region p is in $B_{u,v}$. Let $R_{p,u}$ (resp. $R_{p,v}$) be the smallest orthogonal sub-rectangle of P that includes region p and the point u (resp. v). No matter what zone region q is in, q will either see t and r, or it will see l and b. If q sees l and b, then since P does not contain holes, we can translate the bottom edge of $R_{p,u}$ downwards until it meets with b. This implies that region p sees l and b, which implies that u and u are connected in S. We can handle the case where q sees t and r in a similar manner. This implies that u and v are connected in S for this subcase.

Subcase(iii): Both p and q are in $B_{u,v}$. Since we are assuming that both p and q are in $B_{u,v}$, we will ignore the zones mentioned in the previous subcases and introduce a new set of zones. These zones are induced by orthogonal extensions of the edges of the rectangle that defines the basic region p (Figure 2.9). We only need to consider the portion of the zones that are within $B_{u,v}$. Due to obvious symmetries we can limit region q to being in zone 1, 2, or 3.

Suppose that region q is in zone 1 with respect to p. Due to the relative position

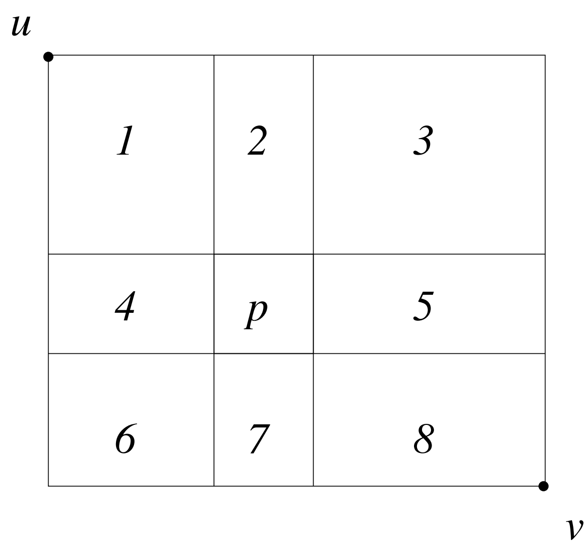

Figure 2.9: The 8 zones induced by the basic region p.

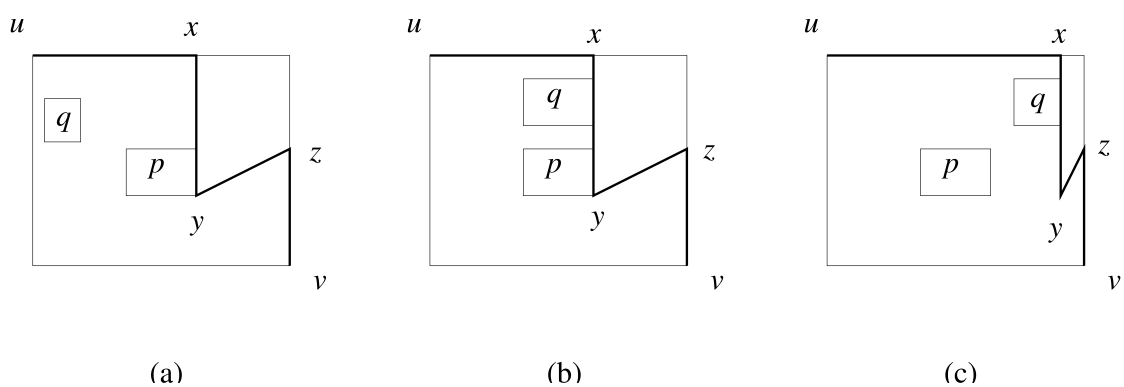

Figure 2.10: A sequence of lines segments connecting u and v in S. (a) q is in zone 1. (b) q is in zone 2. (c) q is in zone 3.

of p and q, we have that $B_{q,v}$ contains $B_{p,v}$. Also, $B_{q,v}$ will contain $B_{p,v}$ even if q is in zone 2 since p and q have the same width. If p and q did not have the same width then either p or q would not be a basic region in the dent diagram of P. Also, since $q \equiv v$, $B_{q,v}$ is contained within P, which implies that $B_{q,v} \subseteq \nu(p)$. Similarly, due to the relative position of p and q, we have that $B_{p,u}$ contains $B_{q,u}$, which implies $B_{q,u} \subseteq \nu(q)$. Let the top right corner of $B_{p,u}$, the bottom right corner of $B_{p,u}$ and the top right corner of $B_{p,v}$ be called x, y and z, respectively. Then the following set of line segments connects u and v in S: $(u, x), (x, y), (y, z), (z, v)$ (see Figures 2.10(a) and (b)). Thus region q cannot be in zone 1 or 2.

The only remaining possibility is that q is in zone 3 with respect to p. Since P does not contain holes, we can translate the right edge of $B_{p,u}$ to the right so that $B_{p,u}$ completely contains $B_{q,u}$. Call the rectangle that is created by this translation $B'_{p,u}$. Since $B'_{p,u}$ was constructed in such a way that it contains both p and q, we have that $B'_{p,u} \subseteq \nu(p)$ and $B'_{p,u} \subseteq \nu(q)$. We also have that the top right corner of $B_{p,v}$ is in $B_{q,v}$ since p is below and to the left of q. Let the top right corner of $B'_{p,u}$, the bottom right corner of $B'_{p,u}$, and the top right corner of $B_{p,v}$ be called x, y, and z, respectively. Then $(u, x), (x, y), (y, z), (z, v)$ connects u and v in S (Figure 2.10(c)). □

We need another result so that we can apply Theorem 2 to mutually overlapping visibility polygons.

Lemma 3. *If p, q, and r are basic regions from the dent diagram of P such that $\nu(p) \bigcap \nu(q) \neq \emptyset$, $\nu(q) \bigcap \nu(r) \neq \emptyset$, and $\nu(r) \bigcap \nu(p) \neq \emptyset$ then $\nu(p) \bigcap \nu(q) \bigcap \nu(r) \neq \emptyset$.*

Proof. We will consider the relative positions of regions p, q, and r. Let $x(p)$ (resp. $x(q)$ and $x(r)$) be the leftmost x coordinate of region p (resp. q and r). Without loss of generality we can assume that $x(p) \leq x(q) \leq x(r)$. Let $y(p)$ (resp. $y(q)$ and $y(r)$) be the highest y coordinate of region p (resp. q and r). Due to obvious symmetries and isomorphisms, we only need to consider two arrangements for region p, q, and r with respect to their highest y coordinates.

Case(i): $y(p) \geq y(q) \geq y(r)$. We will refer to the orthogonal bounding-box for two regions s_1 and s_2 as B_{s_1,s_2}. We will show that there exists a point $h \in \mathcal{H}(p, r)$ which is contained in each of $\nu(p)$, $\nu(q)$, and $\nu(r)$, which will imply the Lemma for this case. Due to the assumed arrangement of p, q, and r, $B_{p,q}$ and $B_{q,r}$ will be completely contained in $B_{p,r}$. This scenario is shown in Figure 2.11. Invoking Lemma 1, we know that there is a point $h \in \mathcal{H}(p, r)$ that is contained in $B_{p,r}$. Now we consider refinements to the position of h within $B_{p,r}$. If h is in $B_{p,q}$ then $B_{h,r}$ will necessarily contain region q, which implies that $q \equiv h$. If h is in $B_{q,r}$ then $B_{h,p}$ will

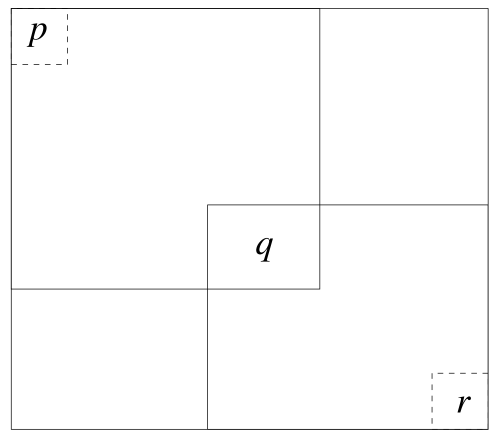

Figure 2.11: The regions p, q, and r arranged such that $y(p) \geq y(q) \geq y(r)$. The bounding boxes for each pair of regions has been shown.

contain region q, which implies that $q \equiv h$. Thus h is neither in $B_{p,q}$ nor $B_{q,r}$.

This implies that h is either in the top-right quadrant or the bottom-left quadrant of $B_{p,r}$ (refer to Figure 2.11). Without loss of generality, assume that h is in the top right quadrant of $B_{p,r}$. Consider the orthogonal bounding-box $B_{q,h}$ for q and h. If a portion of $B_{q,h}$'s top edge is not in P then p does not directly see h, which is a contradiction. If a portion of $B_{q,h}$'s right edge is not in P then r does not directly see h, which is a contradiction. If a portion of $B_{q,h}$'s bottom edge is not in P then $q \not\forall r$, which is a contradiction. If a portion of $B_{q,h}$'s left edge is not in P then $p \not\forall q$, which is a contradiction. Thus all the edges of $B_{q,h}$ are in P and hence $B_{q,h}$ is in P. This implies that $q \equiv h$. Since all three of p, q, and r see a common region, we have $\nu(p) \bigcap \nu(q) \bigcap \nu(r) \neq \emptyset$.

Case(ii): $y(p) \geq y(r) \geq y(q)$**.** Let $h \in \mathcal{H}(p, q)$ be any point from the helper region of p and q that is within the bounding-box of p and q. Lemma 1 guarantees that h exists. Define $B_{p,h}$ and $B_{q,h}$ to be the orthogonal bounding-box of p and h, and q and h, respectively. By definition $B_{p,h}$ and $B_{q,h}$ are sub-rectangles of P. Now we describe a transformation of $B_{p,h}$ and $B_{q,h}$. Imagine translating the bottom edge of $B_{p,h}$ downwards until an edge e_p from $E(P)$ is encountered. Let us call the rectangle created by translating this edge $B'_{p,h}$. Similarly, we translate the top edge

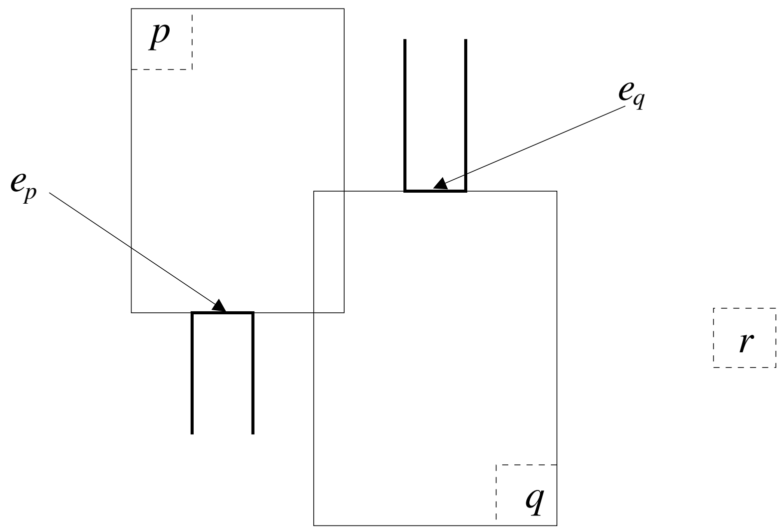

Figure 2.12: The rectangles $B'_{p,h}$ and $B'_{q,h}$.

of $B_{q,h}$ upwards until an edge e_q from $E(P)$ is encountered to create a new rectangle $B'_{q,h}$ (Figure 2.12). We will refer to $B'_{p,h} \bigcap B'_{q,h}$ as h^*. By definition $h^* \subseteq \mathcal{H}(p, q)$. By construction, all points from h^* are no higher than e_q and no lower than e_p.

Let $R_{r,p}$ be an orthogonal sub-rectangle of P that contains r in the lower right corner, and at least one point from $\mathcal{H}(r, p)$. Also, let $R_{r,p}$ be *maximally extended to the left*. Due to the configuration of p, q, and r we have that p sees nothing above and to the right of e_q and nothing below and to the right of e_p. This implies that the top edge of $R_{r,p}$, which we will call t, must be below e_q and above e_p since p sees at least the left endpoint of the top edge of $R_{r,p}$.

We now show that t (i.e. the top edge of $R_{r,p}$) intersects with h^*. This will imply that r directly sees a member of $\mathcal{H}(p, q)$, which will imply the Lemma. We will show that t intersects with h^* indirectly: assume for the sake of contradiction that t is disjoint from h^*.

First suppose that t intersects with $B'_{q,h}$. Since $y(r) \geq y(q)$ and r is in the lower right corner of $R_{r,p}$, we have that the left edge of $R_{r,p}$ is in $B'_{q,h}$. Since $B'_{q,h}$ is completely contained in P, $R_{r,p}$ is not maximally extended to the left, which is a contradiction.

The only remaining possibility is that t is disjoint with $B'_{q,h}$. In this case t must be to the right of region q (recall that $x(q) \leq x(r)$). Then we can translate the left edge of $R_{r,p}$ to the left until this edge lies within $B'_{q,h}$ since $q \wedge r$ and $p \wedge r$. This contradicts the fact that $R_{r,p}$ is maximally extended to the left. $\square$

Lemma 2 and 3 allow us to apply Theorem 2 to the visibility polygons of basic regions in the following way. Lemma 2 tells us that visibility polygons overlap simply. In terms of Theorem 2, this means that $C \bigcap C'$ is a cell (i.e. a simply connected region) for all visibility polygons C and C'. Lemma 3 tells us that three mutually overlapping visibility polygons will not be disjoint. In terms of Theorem 2, we have $C \bigcap C' \bigcap C'' \neq \emptyset$. Thus we can apply Theorem 2 to cells that are the visibility polygons of basic regions. This yields that $\bigcap\{C \in \mathcal{C}\} \neq \emptyset$. In other words, mutually overlapping visibility polygons of basic regions will have a common intersection region.

We can use this result to finish the proof of Theorem 1. A clique in the region visibility graph corresponds to basic regions with mutually overlapping visibility polygons. Thus we can apply Theorem 2 to the visibility polygons for basic regions in cliques: for each clique in the region visibility graph there is a non-empty region within P that is contained in each visibility polygon for the basic regions in the clique. This non-empty region therefore sees all basic regions in the clique, and is therefore the kernel of an r-star that covers all the regions in the clique. This completes the proof for Theorem 1.

2.4 Perfect Graph Theory

Since Theorem 1 holds, finding an efficient algorithm for computing the minimum clique cover of $RVG(P)$ becomes the next step in solving the **Minimum r-star Cover** problem. Unfortunately, the minimum clique cover problem has been shown to be NP-complete for general graphs [15]. One can immediately recognize that the set of region visibility graphs taken over all simple orthogonal polygons, which we will denote as $\mathcal{RVG}$, has special properties that distinguishes itself from general graphs. At the core of the results shown in this chapter is the proof that $\mathcal{RVG}$ is a subset of a class of graphs known as the *perfect graphs*. Perfect graphs have many properties that have been studied from both a combinatorial and an algorithmic point of view [17, 5]. We will only explore those properties of perfect graphs that are relevant to the current discussion.

Perfect graphs were introduced by Claude Berge [4]. The original description of perfect graphs involved an equality between the chromatic number and the clique number, as well as an equality between the stability number and the clique cover number [17]. This definition was later refined by Lovász [32] in the *Perfect Graph Theorem*. We will present a definition of perfect graphs that is a direct implication of this theorem.

A *coloring* of a graph is an assignment of colors to nodes such that adjacent nodes are assigned different colors. The *chromatic number* of a graph G, which is denoted $\chi(G)$, is the minimum number of colors needed to color G. The *clique number* of a graph G, which is denoted $\omega(G)$, is the size of the largest clique in G. We will denote the subgraph induced by the set of nodes A as G_A. We can define perfect graphs as follows:

Definition 3. *[32] A graph G is perfect iff $\omega(G_A) = \chi(G_A)$ for all $A \subseteq V(G)$.*

Given some class $\mathcal{G}$ of graphs, it may be difficult to use Definition 3 alone to show that $\mathcal{G}$ is a subset of the perfect graphs. This task can be made easier if we know some structural properties of perfect graphs. There are two fundamental structural results that we will use. The first result is the recently proven *Strong Perfect Graph Theorem* [34]. This theorem describes perfect graphs in terms of holes and antiholes. A *hole* is a chordless cycle of length 4 or greater, while an *antihole* is the complement graph of a hole. We say a hole or antihole is *odd* if it contains an odd number of vertices.

Theorem 3. *[34] A graph G is perfect iff it contains no odd hole and no odd antihole.*

This theorem was conjectured to be true by Berge in 1960. Since then there has been a flourish of research all aimed at proving this conjecture. The conjecture was finally confirmed by Chudnovsky et. al. [34]. We will use this characterization of perfect graphs when showing that $\mathcal{RVG}$ is a subset of the perfect graphs.

We will also use a structural result involving star cutsets. A *cutset* is a set of nodes whose removal renders the graph disconnected. A *star cutsut* is a cutset S that contains some node s, which is known as the *universal node*, that is adjacent to every other node in S. We will combine the two following theorems to get the result we need. The first result is a restatement of the strong perfect graph theorem [5].

Theorem 4. *[5, 34] The only minimally imperfect graphs are the odd holes and odd antiholes.*

A graph is *minimally imperfect* if it is not perfect, but all its proper induced subgraphs are perfect. A very important result concerning minimally imperfect graphs comes from Chvátal [7].

Theorem 5. *[7] No minimally imperfect graph has a star cutset.*

If we combine Theorems 4 and 5 then we get the following.

Theorem 6. *No odd hole or odd antihole has a star cutset.*

Theorem 6 will be used extensively in order to show that $\mathcal{RVG}$ is a subclass of the perfect graphs.

Knowing that $\mathcal{RVG}$ is a subclass of the perfect graphs is only useful if we can solve the minimum clique cover problem on perfect graphs. The following result, which is due to Grötschel et. al., is essential to our current discussion.

Theorem 7. *[18] A minimum clique cover can be computed on a perfect graph in polynomial time.*

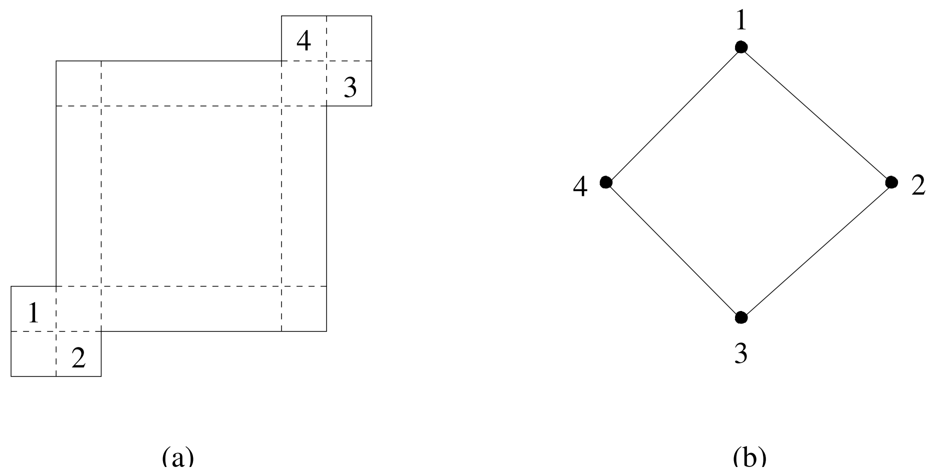

Figure 2.13: A polygon P where $RVG(P)$ is not chordal. (a) An orthogonal polygon P and its associated dent diagram. (b) An induced 4-hole in $RVG(P)$.

There are many subclasses of perfect graphs that have been studied [17, 5]. Of particular interest are the *chordal* and *weakly chordal* graphs. A graph is *chordal* if it contains no holes. A graph is *weakly chordal* if it does not contain any holes or antiholes of size 5 or larger. Problems can often be solved more efficiently on these subclasses than on general perfect graphs. It is therefore useful to know if $\mathcal{RVG}$ is in one of these subclasses. Figures 2.13 and 2.14 show that $\mathcal{RVG}$ is neither chordal nor is it weakly chordal.

2.5 Constriction Regions

Now we explore a geometric tool that we will use in the proof that all graphs in $\mathcal{RVG}$ are perfect. This tool, called the *constriction region*, was introduced in [39]. The constriction region is defined in terms of two basic regions p and q in P that are not adjacent in $RVG(P)$. The fact that p and q are not adjacent in $RVG(P)$ implies that the visibility polygons $\nu(p)$ and $\nu(q)$ are disjoint. If they were not disjoint then their overlap would constitute a kernel for an r-star covering both p and q. Let $C_1, C_2, ..., C_k$ be the maximal connected components of $\mathcal{C} = P/\{\nu(p) \bigcup \nu(q)\}$. The constriction region will be defined with the aid of the following Lemma.

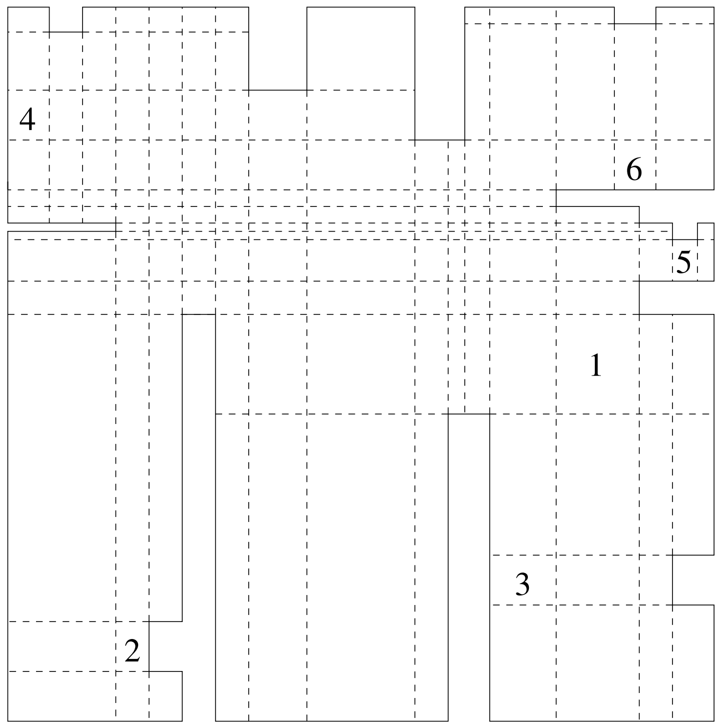

(a)

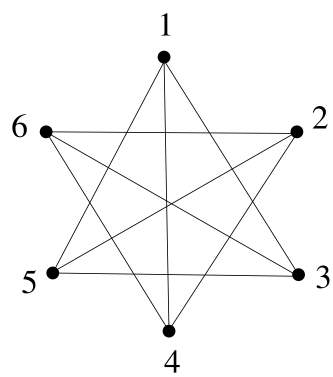

(b)

Figure 2.14: A polygon P where $RVG(P)$ is not weakly chordal. (a) An orthogonal polygon P and its associated dent diagram. (b) An induced 6-antihole in $RVG(P)$.

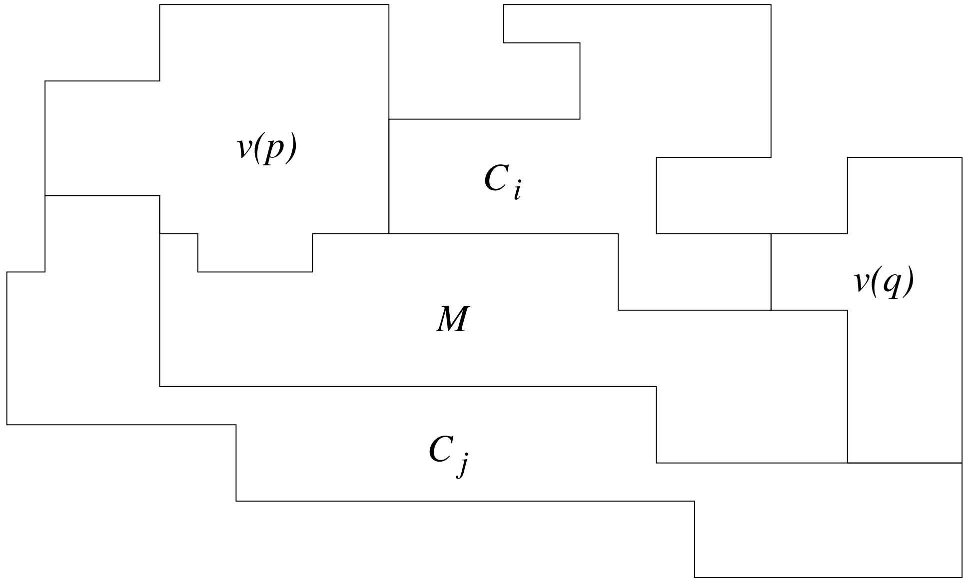

Figure 2.15: The union of $\nu(p)$, $\nu(q)$, C_i, and C_j.

Lemma 4. *There is a unique component $C_i \in \mathcal{C}$ such that $\nu(p) \bigcup \nu(q) \bigcup C_i$ is a connected region.*

Proof. We will define P_i to be $C_i \bigcup \nu(p) \bigcup \nu(q)$ for each $i = 1, ..., k$. Clearly all components from $\mathcal{C}$ are mutually disjoint, even along their boundaries. Since P is simply connected, there is at least one component of $\mathcal{C}$, say C_i, such that P_i is simply connected. Now suppose that there are two components, say C_i and C_j, such that both P_i and P_j are simply connected. We will define S to be the union of $\nu(p)$, $\nu(q)$, C_i, and C_j. Then S contains some region M that prevents C_i and C_j from sharing a boundary (Figure 2.15). M shares a boundary with C_i and C_j and thus $M \notin \mathcal{C}$. This implies that M is a hole (in P) or it contains a hole, which implies that P contains a hole, which contradicts the simple connectedness of P. □

Lemma 4 motivates the definition of a constriction region.

Definition 4. *Let p and q be the regions corresponding to two non-adjacent vertices from $RVG(P)$. The constriction between p and q, denoted $Cons(p, q)$, is the unique region such that $\nu(p) \bigcup \nu(q) \bigcup Cons(p, q)$ is a connected region.*

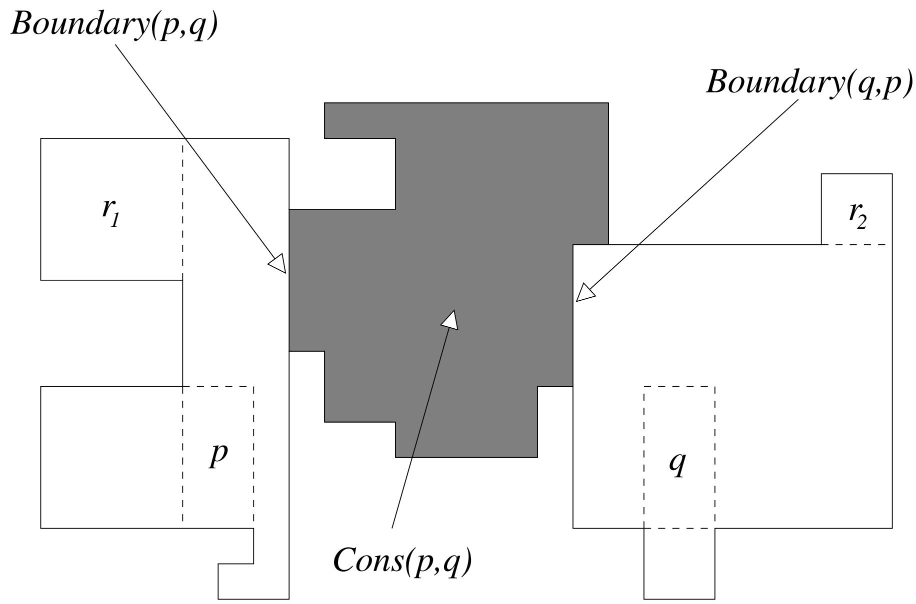

Figure 2.16: The constriction region $Cons(p, q)$. The region r_1 is on the same side as p (i.e. $r_1 \subset P(p, q)$), while r_2 is on the same side as q (i.e. $r_2 \subset P(q, p)$) with respect to $Cons(p, q)$.

This definition implies that p (resp. q) does not see any point in $Cons(p, q)$ (Figure 2.16). We can use constriction regions to partition P with respect to two non-adjacent regions from $RVG(P)$.

Definition 5. *Let p and q be two non-adjacent regions from $RVG(P)$. We define $P(p, q)$ to be the region from $P/Cons(p, q)$ that contains p.*

When a region r is in $P(p, q)$ we say that r is on the same side as region p with respect to $Cons(p, q)$.

We will call the boundary $\nu(p)$ has with $Cons(p, q)$ the *Boundary*(p, q) (Figure 2.16). It turns out that these boundaries are well-behaved. Let $O =< o_1, o_2, ..., o_{t+1} >$ be an ordered set of points from $\mathbb{R}^2$. Then O defines an *orthogonal t-chain* if the line segment (o_i, o_{i+1}) is orthogonal for all $1 \leq i \leq t$, this set of line segments is not self-intersecting, and if the line segments (o_i, o_{i+1}) and (o_{i+1}, o_{i+2}) are perpendicular.

Lemma 5. *Let p and q be two basic regions such that $p \not\sim q$. Then Boundary(p, q) is* *either an orthogonal 1-chain (i.e. a single line segment) or an orthogonal 2-chain.*

Proof. By definition, $Boundary(p, q)$ is a portion of the boundary of $\nu(p)$ and hence is an orthogonal chain of some length. Visibility polygons are defined to be maximal in the sense that every edge of a visibility polygon is touching at least a portion of an edge of $E(P)$. Thus every edge in the chain that defines $Boundary(p, q)$ is touching at least a portion of an edge from $E(P)$. We will use this observation to derive a contradiction in the following argument.

By the definition of $Boundary(p, q)$, we have the following: except at its endpoints, $Boundary(p, q)$ is properly contained in P. In other words, except at its endpoints, $Boundary(p, q)$ will be disjoint from all edges in $E(P)$. Therefore if $Boundary(p, q)$ is an orthogonal *t-chain* for $t \geq 3$ then one of the edges in the chain cannot be touching an edge from $E(P)$. This contradicts our observation from the previous paragraph that every edge in the chain that defines $Boundary(p, q)$ is touching at least a portion of an edge from $E(P)$. Thus $Boundary(p, q)$ cannot be a t-chain for $t \geq 3$.

An example of the cases where the boundaries between visibility polygons and constriction regions are t-chains for $t \leq 2$ are shown in Figure 2.16. In Figure 2.16, $Boundary(p, q)$ is an orthogonal 1-chain, while $Boundary(q, p)$ is an orthogonal 2-chain. This implies that $Boundary(p, q)$ is either an orthogonal 1-chain or an orthogonal 2-chain. □

Lemma 5 tells us that $Boundary(p, q)$ may be an orthogonal 2-chain for two regions $p \mathrel{\mathrm{A}\!\!-} q$. Now we introduce some terminology for dealing with boundaries that are orthogonal 2-chains. Let p and q be two basic regions such that $p \mathrel{\mathrm{A}\!\!-} q$ and $Boundary(p, q)$ is an orthogonal 2-chain. We will refer to the three vertices that define $Boundary(p, q)$ as o_1^p, o_2^p, and o_3^p. (Figure 2.17(a)).

We will use line segments extending from o_1^p and o_3^p to partition $P(p, q)$. We extend an orthogonal line segment from o_1^p (resp. o_3^p) towards p until an edge from

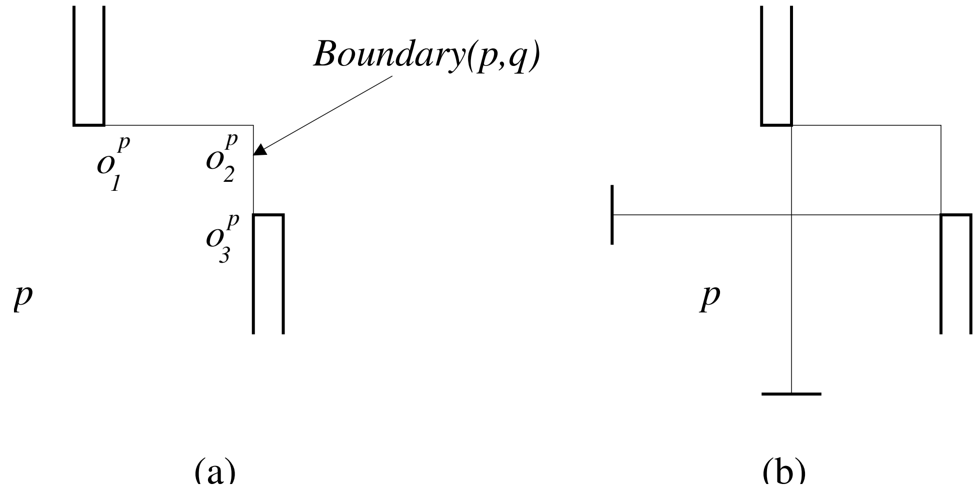

Figure 2.17: The different zones with respect to a 2-chain boundary. (a) The end-points of the 2-chain. (b) Partitioning p's side.

$E(P)$ is encountered (Figure 2.17(b)). Each of the regions in this partition have unique properties that we will exploit in Section 2.7. We will give these regions names so that we can easily refer to them. The region that p is located in will simply be referred to as p's *zone* (Figure 2.18). The region that is bounded on two sides by $Boundary(p, q)$ will be called p's *bulge*. The remaining two zones will be labeled based upon their orientation with $Boundary(q, p)$. To do this we will first consider an extension of $Boundary(q, p)$. If this boundary is an orthogonal line segment then we will consider the *extension* of $Boundary(q, p)$ to be the line that is collinear with $Boundary(q, p)$. If $Boundary(q, p)$ is an orthogonal 2-chain then we will consider the extension of $Boundary(q, p)$ to be the pair of orthogonal rays that terminate at o_2^q and travel along the vectors $< o_2^q, o_1^q >$ and $< o_2^q, o_3^q >$. If there exists an orthogonal line that passes through one of the remaining zones and the bulge that intersects the extension of $Boundary(q, p)$ at a point, then that zone will be called p's *perpendicular zone*. The remaining zone will be called p's *parallel zone*. If no such ray exists for either remaining zone then the zones are labeled arbitrarily either parallel or perpendicular. We now point out two observations about these zones:

Observation 1. *Any region in p's zone cannot see any region in $Cons(p, q)$ or $P(q, p)$.*

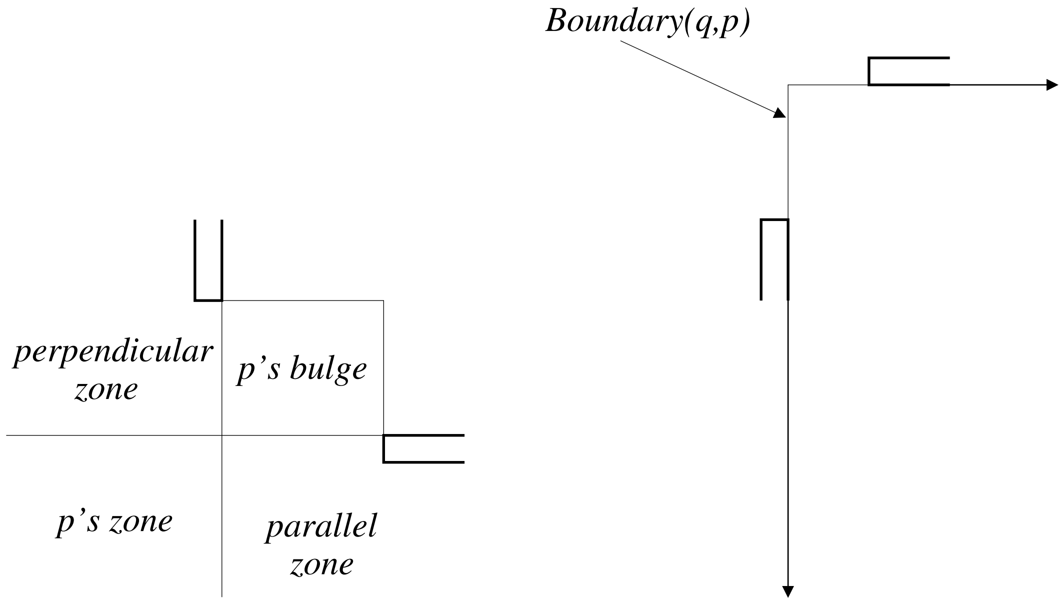

Figure 2.18: The various zones with respect to p.

Observation 2. *If $Boundary(q, p)$ is a line segment then any region in p's parallel zone cannot see any portion of $Boundary(q, p)$.*

Depending on the visibility relationship a basic region r has with p and q, we can often restrict where r may be positioned with respect to $Cons(p, q)$. Specifically, we will often be able to say that r must be on the same side as p or q with respect to $Cons(p, q)$. Restricting where r may be positioned will be an important step in showing that subgraphs of $RVG(P)$ are not odd holes or odd antiholes. We now describe a fundamental observation about constriction regions and regions that are indirectly visible.

Lemma 6. *Let p, q, and r be three basic regions such that $p \not\wedge q$ and $p \wedge r$. If $r \subseteq Cons(p, q)$ or if $r \subset P(q, p)$ then r sees some portion of $Boundary(p, q)$.*

Proof. It cannot be the case that $p \equiv r$ since $r \subseteq Cons(p, q)$ or $r \subset P(q, p)$. Therefore it must be the case that $p \overline{\wedge} r$. If some portion of $\mathcal{H}(p, r)$ lies in $P(p, q)$ then the visibility rectangle for r and $\mathcal{H}(p, r)$ must cross $Boundary(p, q)$. This causes r to see a portion of $Boundary(p, q)$. Suppose that $\mathcal{H}(p, r)$ is in $Cons(p, q)$ or $P(q, p)$. But

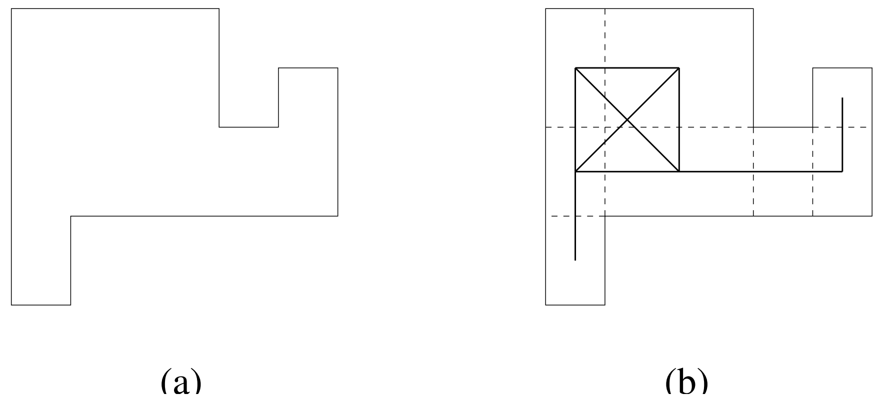

Figure 2.19: (a) A simple orthogonal polygon. (b) The line embedding of $RVG(P)$.

this implies that p sees a region (i.e. $\mathcal{H}(p, r)$) in $Cons(p, q)$ or in $P(q, p)$, which is a contradiction. □

2.6 Properties of Region Visibility Graphs

In this chapter we explore the various relevant properties of $\mathcal{RVG}$. The ultimate goal is to show that all graphs in $\mathcal{RVG}$ are perfect. Since such graphs have a relationship to our polygon P, we will often want to embedd $RVG(P)$ in P using the following technique.

Definition 6. *The line embedding of $RVG(P)$ is a set of line segments Γ created in the following way. If an edge $e = (u_1, u_2)$ from $RVG(P)$ joins two nodes that have direct visibility, then we add a line segment to Γ whose endpoints are the centers of the regions that correspond to u_1 and u_2. If edge e joins two nodes that are indirectly visible then we add two line segments to Γ to represent e: one line segment from the center of u_1 to a point $k \in \mathcal{H}(u_1, u_2)$, and one line segment from k to the center of u_2.*

Figure 2.19 depicts an example of a line embedding of a region visibility graph. Although a line embedding Γ is simply a set of line segments, due to the construction of Γ, certain line segments will share endpoints. Thus we can treat the set of line segments $l_1, l_2, ..., l_k$ from Γ such that l_i shares an endpoint with l_{i+1} for $i = 1, 2, ..., k-1$ as a path in P.

Since there is an important relationship between star-cutsets and perfect graphs, we will explore how star-cutsets can be formed within Region Visibility Graphs.

Lemma 7. *Let p, q, and r be three basic regions from a connected vertex subgraph $G \subseteq RVG(P)$ such that $(p,q) \notin E(G)$, $(p,r) \in E(G)$, and $(q,r) \in E(G)$. If $r \subseteq Cons(p,q)$ then G contains a star-cutset with r as a universal node.*

Proof. First we will point out some simple observations. Since $(p,q) \notin E(G)$ we know that $p \not\wedge q$. Likewise, since (p,r) and (q,r) are edges in G we have $p \wedge r$ and $q \wedge r$. Moreover, we know that $p \overline{\wedge} r$ and $q \overline{\wedge} r$ since $r \subseteq Cons(p,q)$.

Now we make some observations about $\nu(r)$. The visibility polygon $\nu(r)$ is maximal in the sense that it is the largest star with r in the kernel. We also know that r strictly indirectly sees both region p and region q. These two facts imply that $P/\nu(r)$ induces at least two polygons. One of these polygons contains region p, while some other polygon from $P/\nu(r)$ contains region q. We will refer to the polygon from $P/\nu(r)$ that contains region p as P_p. Likewise, we will refer to the polygon from $P/\nu(r)$ that contains region q as P_q.

Now we will consider paths in G between p and q that do not use r. If no such path exists then r is an articulation point and $\{r\}$ is a trivial star-cutset. Let $L =< p = l_1, l_2, ..., l_k = q >$ be a path in G from p to q that does not contain r. We will show that there must be a node that is neither p nor q along path L that indirectly sees r. Then we will show how to construct a star-cutset with r as a universal node from this information.

We will consider the line embedding Γ of L in P (see Definition 6). By the definition of P_p and P_q, we know that Γ starts in P_p, passes through $\nu(r)$, then continues into P_q. There are two possibilities for how Γ passes through $\nu(r)$: either an endpoint from a line segment from Γ is in $\nu(r)$, or a line segment from Γ is

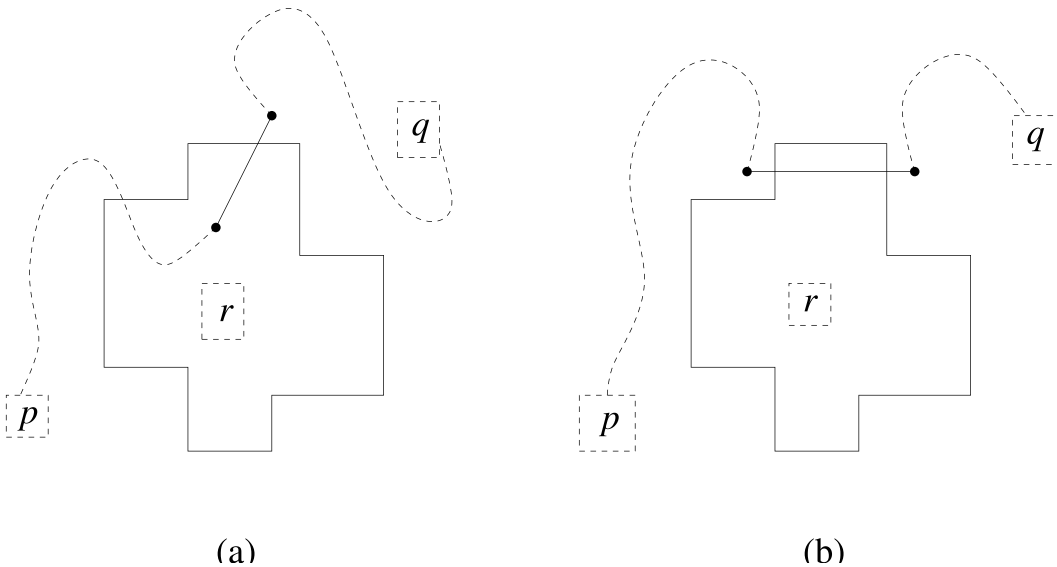

Figure 2.20: The two ways that a line segment from Γ may pass through $\nu(r)$. (a) Γ has an endpoint in $\nu(r)$. (b) A portion of a line segment from Γ is in $\nu(p)$.

partially contained in $\nu(r)$ (Figure 2.20).

First let us consider the case where an endpoint from a line segment from Γ is in $\nu(r)$ (Figure 2.20(a)). We will refer to the first endpoint from Γ that is in $\nu(r)$ as l_i. By the definition of Γ, l_i corresponds to either a basic region or a helper region for a pair of regions along L. If l_i corresponds to a basic region then $r \equiv l_i$. Since $r \equiv l_i$, $r \overline{\wedge} p$, and $r \overline{\wedge} q$, we know that l_i corresponds to neither p nor q. If l_i corresponds to a helper region then r indirectly sees both of the basic regions associated with the helper region that l_i represents. Since $p \not\wedge q$, l_i cannot correspond to a helper region for p and q, and hence r sees a basic region from L that is neither p nor q.

Now consider the case where a line segment from Γ is partially contained in $\nu(r)$ (Figure 2.20(b)). We will refer to the endpoints of this line segment as l_i and l_{i+1}. By definition of Γ, at least one of these endpoints corresponds to a basic region of P. First suppose that l_i corresponds to a basic region. Then $r \wedge l_i$ since the edge (l_i, l_{i+1}) crosses $\nu(r)$. The basic region associated with l_i cannot be region q since l_i is in P_p. Now suppose l_i corresponds to the basic region p. Since the line segment (l_i, l_{i+1}) crosses $\nu(r)$, we have that p sees a basic region (i.e. the one associated with l_{i+1})

that is in P_q, which is a contradiction. Thus l_i cannot be associated with p. Now suppose that l_{i+1} is associated with a basic region. We can use the same argument as was given for the case where l_i was assumed to be associated with a basic region to show that $r \wedge l_{i+1}$ and l_{i+1} corresponds to neither p nor q.

We have shown that r will indirectly see some region from L that is neither p nor q. Now we can construct a star-cutset as follows. Let $\{L_1, L_2, ..., L_t\}$ be the set of paths in G connecting p and q that do not use r. Let the region that r indirectly sees that is neither p nor q from the i^{th} path be called $s(L_i)$. Now consider the set $S = \{s(L_i) | i = 1, ..., t\} \bigcup \{r\}$. By the definition of S, r is universal in S. Furthermore, the removal of S from $V(G)$ will force p and q into different components since all paths from p to q have been severed by the removal of $s(L_i)$ and r. Thus S is a star-cutset with r as a universal node, as required. $\square$

Now we prove a lemma that will be especially useful when showing that no graphs from $\mathcal{RVG}$ contain odd antiholes of size 5 or larger.

Lemma 8. *If $p, q, r,$ and s are four basic regions from a vertex induced subgraph $G \subseteq RVG(P)$ such that $p \not\wedge q$, $r \wedge p$, $r \not\wedge q$, $s \not\wedge p$, $s \wedge q$, $r \wedge s$, (i.e. $p, q, r,$ and s form a chordless 4-path in $RVG(P)$) and $r, s \subseteq Cons(p, q)$ then G contains a star-cutset with either r or s as a universal node.*

Proof. We will divide this proof into two major cases based upon the visibility relationship that exists between region r and region s. We will assume that p is to the left of q[2].

Case(i): $r \overline{\wedge} s$. For this case we intend to show that $r \overline{\wedge} s$ implies that either $r \subseteq Cons(s, p)$ or $s \subseteq Cons(r, q)$. We can then employ Lemma 7 to get our result for Case(i).

[2] We can make this assumption since p is either left, right, below, or above with respect to region q, which are all symmetric.

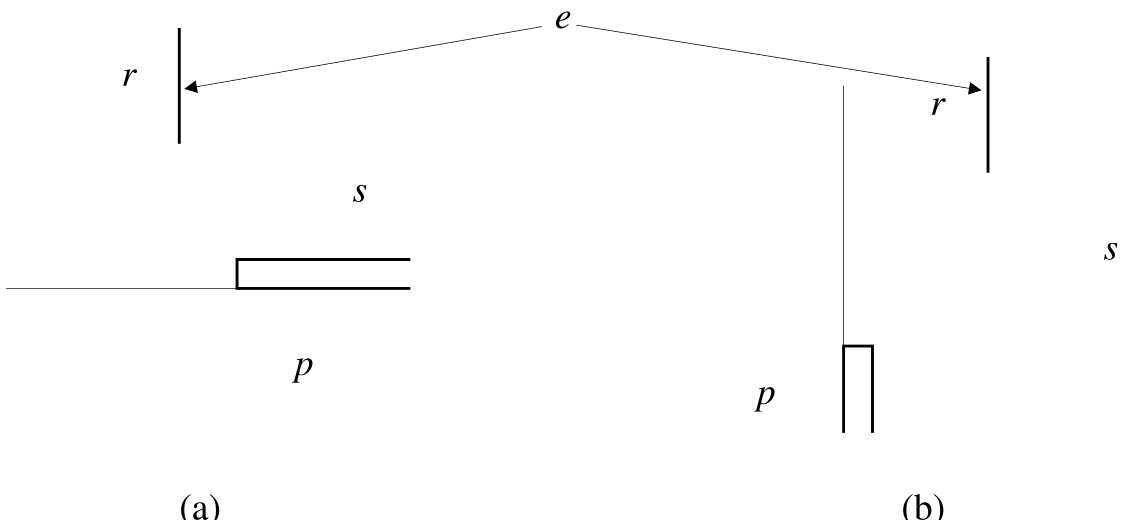

Figure 2.21: Region r is to the left of region s within $Cons(p,q)$. (a) Region r is to the left of region p.

Either r is to the left of s, or r is to the right of s since $r \barwedge s$.

First suppose r is to the left of s. Since $r \barwedge s$, there is some edge $e \in E(P)$ in $Cons(p,q)$ that interferes with r and s's direct visibility. We now consider where e is located with respect to region r by considering where region r is located with respect to region p. Suppose region r is to the left of region p. It is easy to see that this can only occur if $Boundary(p,q)$ is a line segment that is either above or below r (Figure 2.21(a)) [3]. In this case e must be located to the right of region r. This forces r to be in $Cons(s,p)$, and by Lemma 7 we have that G has a star-cutset with r as a universal node. Now suppose that p is to the left of r (Figure 2.21(b)). Again e must be to the right of region r and again we have that r must be in $Cons(s,p)$. Invoking Lemma 7 once more we get that r is a universal node in a star-cutset of G.

Now suppose r is to the right of s. In this case there must be edges from $E(P)$ that are part of $Cons(p,q)$ that force $s \not\wedge p$ and some other chain of edges that are part of $Cons(p,q)$ that cause $r \not\wedge q$. There must also be an edge $e \in E(P)$ that is

[3] We pause here to note that in Figure 2.21, the outlines of the basic regions have not been shown. In order to simply subsequent figures, we will discontinue showing the boundary of basic regions in figures for the remainder of this Chapter.

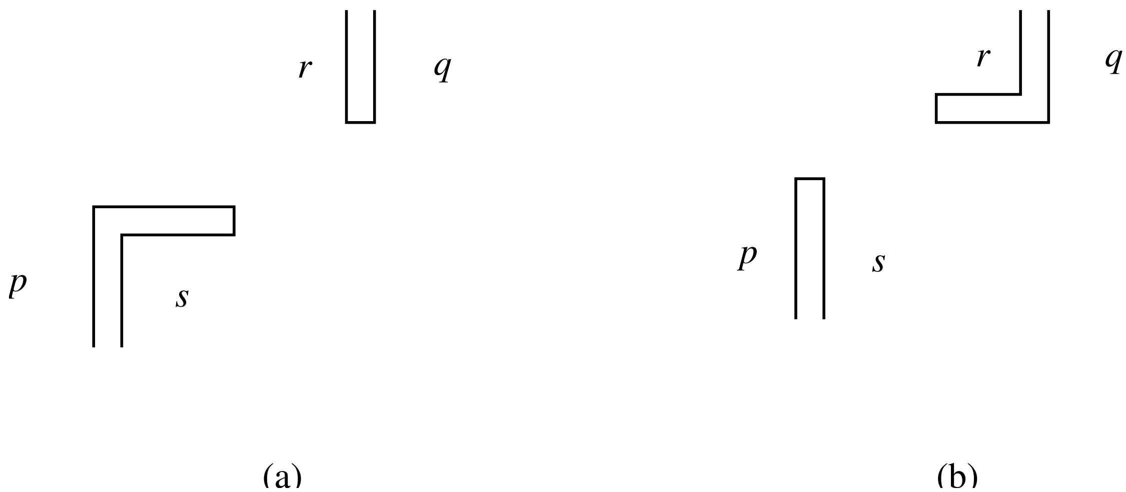

Figure 2.22: Region r is to the right of region s within $Cons(p,q)$. (a) The edge that forces $r\overline{\wedge}s$ is part of the chain that causes $s \not\wedge p$. (b) The edge that forces $r\overline{\wedge}s$ is part of the chain that causes $r \not\wedge q$.

part of *one* of these two chains that causes $r\overline{\wedge}s$. Suppose that e is part of the chain of edges that causes $s \not\wedge p$ (Figure 2.22(a)). In this case r must be within $Cons(s,p)$. If e is part of the chain that causes $r \not\wedge q$ then region s is within $Cons(r,q)$. In either case, Lemma 7 implies that G has a star-cutset.

Case(ii): $r \equiv s$. Let $B_{r,s}$ be any maximal orthogonal sub-rectangle of P that contains both r and s. Recall that since $r \subseteq Cons(p,q)$, $r\overline{\wedge}p$. Also, since $s \subseteq Cons(p,q)$, $s\overline{\wedge}q$. These observations imply that $P/B_{r,s}$ disconnects P into 2 or more polygons. Moreover, p and q will be in different polygons from $P/B_{r,s}$. Now we consider the line embeddings of paths in G from region p to region q that do not use region r. If no such path exists, then r is an articulation point and hence $\{r\}$ is a star-cutset, as required.

We can analyze all line embeddings of paths $L =< p = l_1, l_2, ..., l_k = q >$ as they cross $B_{r,s}$. Using the same analysis from Lemma 7, we can easily show that region r indirectly sees some region along this path that is neither p nor q. This implies that there exists a star-cutset with r as a universal node. $\square$

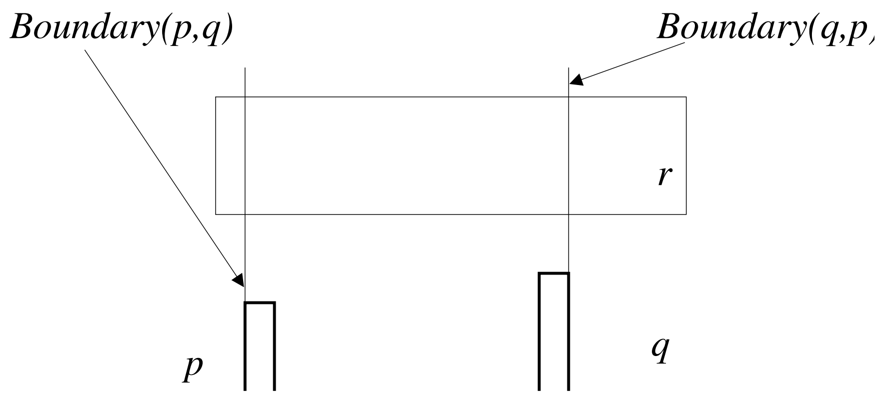

Figure 2.23: A depiction of Lemma 9.

In our proof that all graphs from $\mathcal{RVG}$ are perfect, we will use another result. As in Lemma 7, this result is regarding 3 basic regions from the dent diagram. But unlike the previous Lemma, one of the three regions considered does not indirectly (or directly) see the other two regions.

Lemma 9. *Let p, q, and r be three basic regions from a vertex induced subgraph $G \subseteq RVG(P)$ such that $(p, q) \notin E(G)$, $(p, r) \in E(G)$, and $(q, r) \notin E(G)$. Then $r \subseteq Cons(p, q)$ or $r \subset P(p, q)$.*

Proof. Assume $r \subset P(q, p)$. We know that r must see a portion of $Boundary(p, q)$ since $r \wedge p$ and $r \not\subset P(p, q)$ (Figure 2.23). Thus there is a rectangle covering r that crosses $Boundary(p, q)$. Since it was assumed that $r \subset P(q, p)$, this rectangle must also cross over $Boundary(q, p)$. But this implies that $r \wedge q$, a contradiction. $\square$

2.7 Graphs in $\mathcal{RVG}$ are Perfect

In this section we will show that $\mathcal{RVG}$ is a subclass of the perfect graphs. We will use the hole-antihole characterization of perfect graphs for this proof. Throughout the remainder of this Section, we use a *cyclic* labeling of holes and antiholes. In a *cyclic* labeling of a hole, we label the vertices of the hole with integers from 1 to n so that the vertex with label i is adjacent to the vertices with labels $i - 1$ and $i + 1$

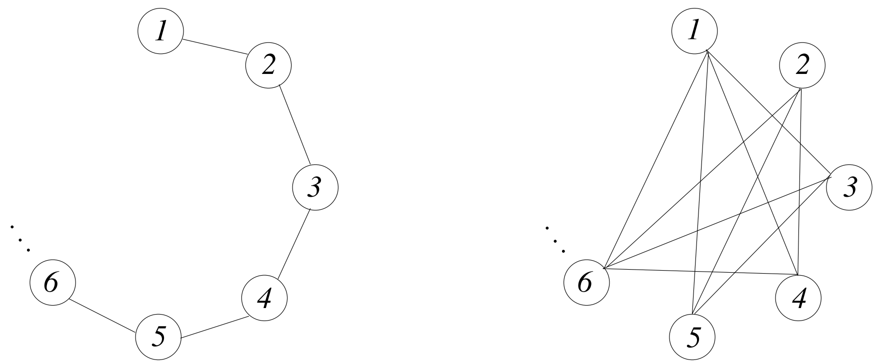

Figure 2.24: A cyclic labeling of hole and an antihole.

(Figure 2.24)[4]. We label the vertices of an antihole H so that they are cyclic in the complement graph $\overline{H}$ (Figure 2.24).

Now we present a lemma that restricts the position of certain regions that are part of an antihole.

Lemma 10. *If $p, q, r, s, t,$ and u are 6 consecutive nodes from a vertex induced antihole $H \subseteq RVG(P)$ of size 7 or larger, then $q \subseteq Cons(r, s)$ or $t \subseteq Cons(r, s)$.*

Proof. Lemma 9 tells us that $q \subseteq Cons(r, s)$ or $q \subset P(s, r)$. Similarly, $t \subseteq Cons(r, s)$ or $t \subset P(r, s)$. Thus we can prove the Lemma indirectly by assuming that $t \subset P(r, s)$ and $q \subset P(s, r)$. In this case, q and t cannot have direct visibility since this would imply that, for example, $q \wedge r$. Thus $q \overline{\wedge} t$ and hence q and t are not is the same column in the dent diagram of P. Thus without loss of generality we can assume that region t is to the left of region q. Also, we can limit where $\mathcal{H}(q, t)$ is located: if there exists a point $h \in \mathcal{H}(q, t)$ such that $h \in P(r, s)$ then $q \wedge r$, and if $h \in P(s, r)$ then $t \wedge s$. Thus $\mathcal{H}(q, t) \subseteq Cons(r, s)$. Let $R^t_{\mathcal{H}(t,q)}$ and $R^q_{\mathcal{H}(t,q)}$ be maximally horizontally extended orthogonal sub-rectangles of P that contain t and $h \in \mathcal{H}(t, q)$, and q and h, respectively. This implies that there is some edge from $E(P)$ touching the right edge of $R^t_{\mathcal{H}(t,q)}$. Let us call this edge e_t. There must also be some other edge of $E(P)$ touching the left edge of $R^q_{\mathcal{H}(t,q)}$, which we will call e_q (Figure 2.25). Now we will

[4]In the special case where $i = 1$, $i - 1$ is interpreted as n. Also, if $i = n$ then node $i + 1$ is interpreted as 1.

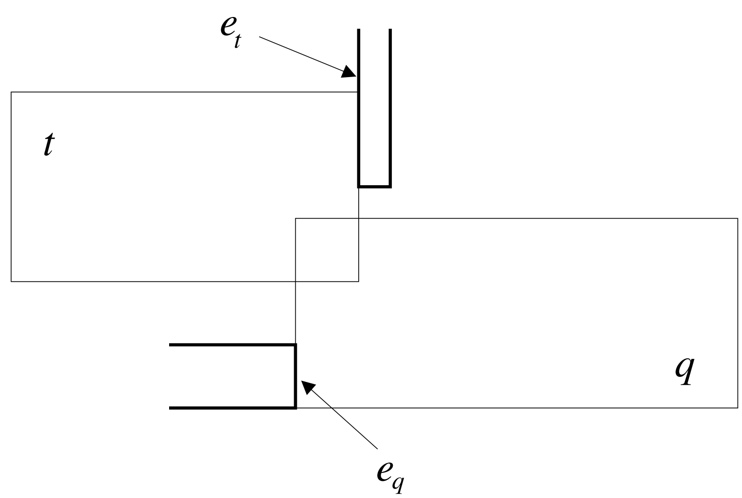

Figure 2.25: An edge from $E(P)$ touches the edges of $R^q_{\mathcal{H}(t,q)}$ and $R^t_{\mathcal{H}(t,q)}$.

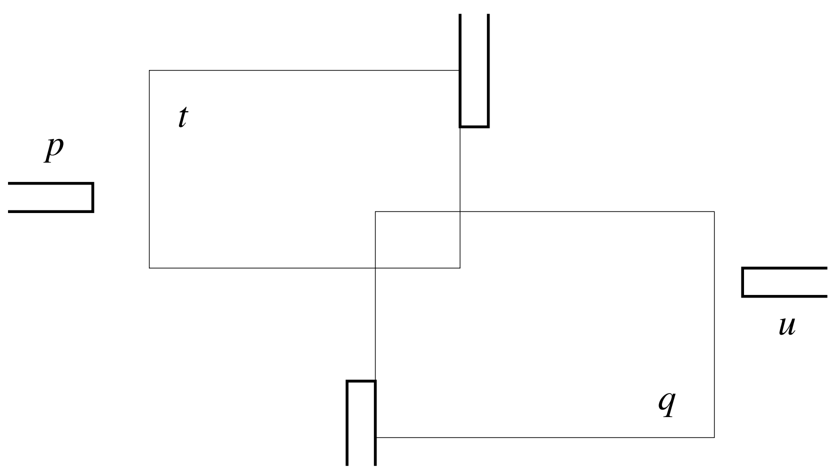

Figure 2.26: Regions p and u are on opposite sides of $Cons(r, s)$.

consider where these edges are located.

Case(i): Both e_t and e_q are in $Cons(r, s)$. We can argue against this case by considering where regions p and u are located. We know $u, p \not\subseteq Cons(r, s)$ since otherwise Lemma 7 would imply that H contains a star-cutset, which is a contradiction by Theorem 7. Suppose $u \subset P(s, r)$ and $p \subset P(r, s)$. Since $p \not\mathrel{A} q$, there must be some edge from $E(P)$ that prevents $R^p_{\mathcal{H}(p,s)}$ from overlapping with $R^q_{\mathcal{H}(t,q)}$. Likewise, since $u \not\mathrel{A} t$, there must be some edge from $E(P)$ that blocks $R^u_{\mathcal{H}(u,r)}$ from overlapping with $R^t_{\mathcal{H}(t,q)}$ (Figure 2.26). These edges will necessarily force $u \not\mathrel{A} p$, which is a con-

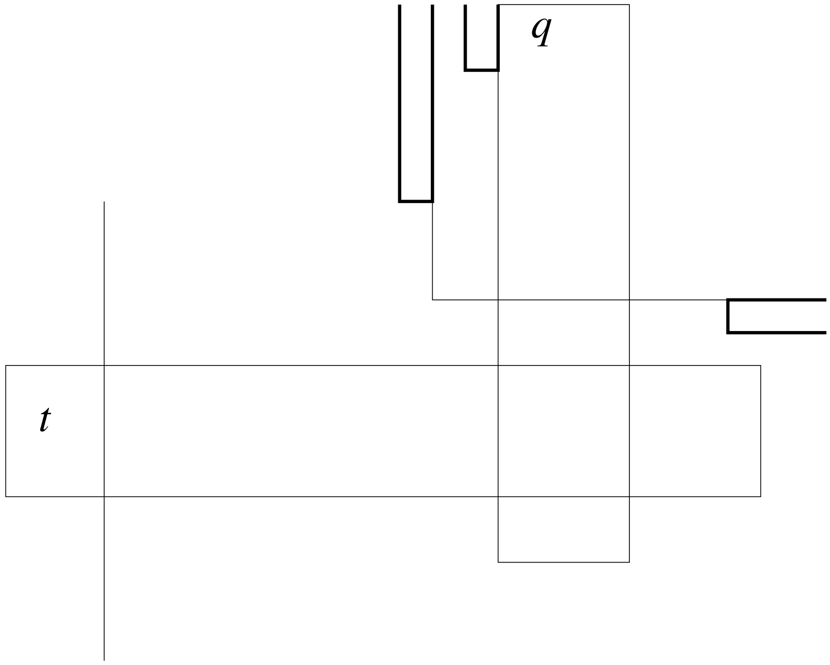

Figure 2.27: The edge e_q is not in $Cons(r, s)$.

tradiction. A similar argument can be used against the case where $u \subset P(r, s)$ and $p \subset P(s, r)$.

Next consider the case where p and u are on the same side of $Cons(r, s)$. First suppose p and u are both in $P(s, r)$. There must be an edge either above or below u that restricts $R^u_{\mathcal{H}(u,r)}$ from overlapping with $R^t_{\mathcal{H}(t,q)}$. Likewise, there is an edge that restricts $R^p_{\mathcal{H}(p,r)}$ from overlapping with $R^q_{\mathcal{H}(t,q)}$. These edges will force, $p \not\Join u$, which is a contradiction. A similar argument can be used to show that $p \not\Join u$ if p and u are both in $P(r, s)$. In this case, the edges that force $R^u_{\mathcal{H}(u,s)}$ from overlapping with $R^t_{\mathcal{H}(t,q)}$ and $R^p_{\mathcal{H}(p,s)}$ from overlapping with $R^q_{\mathcal{H}(t,q)}$ will block the indirect visibility between region p and u.

Case(ii): At least one of e_t and e_q is not in $Cons(r, s)$. Without loss of generality, let us assume that e_q is not in $Cons(r, s)$. Figure 2.27 depicts such a situation. Since $e_q \subset P(s, r)$, $R^q_{\mathcal{H}(t,q)}$ must extend downwards (resp. upwards) so that it can overlap with $R^t_{\mathcal{H}(t,q)}$. This implies that region q is in s's parallel zone. This situation implies that $R^q_{\mathcal{H}(t,q)}$ vertically spans s's bulge. Now consider where region p is located. We know $p \not\subseteq Cons(r, s)$, otherwise Lemma 7 implies that H contains a star-cutset. If $p \subset P(r, s)$, then for p to be able to indirectly see region s, it must also indirectly see region q. If $p \subset P(s, r)$, then for p to be able to indirectly see region r, it must also indirectly sees region q. In either case we arrive at a contradiction. □

Now we are ready to prove our main result of this Section. We will show that $\mathcal{RVG}$ only contains perfect graphs. As already stated, we will use the hole-antihole characterization of perfect graphs that was given in Theorem 3. This Theorem states that a graph is perfect iff it does not contain any odd holes or odd antiholes of size 5 or larger. Based upon this characterization, we will subdivide the proof into two parts. First we consider odd holes in graphs from $\mathcal{RVG}$. We note here that a proof sketch for the following Lemma has been found by Culberson and Keil [27].

Lemma 11. *No graph in $\mathcal{RVG}$ has an odd hole of size 5 or larger.*

Proof. We will prove Lemma 11 indirectly: let $G \in \mathcal{RVG}$ be any region visibility graph that contains a vertex induced odd hole H of size 5 or larger. For the remainder of the proof we will employ a cyclic labeling of the vertices in the hole H.

We will divide the proof into 4 main cases. Each of these cases will deal with different types of boundaries. Lemma 5 tells us that boundaries with constriction regions are either orthogonal line segments, or they are orthogonal 2-chains.

Case(i): $Boundary(1, 3)$ **and** $Boundary(3, 1)$ **are both orthogonal line segments, and these boundaries are parallel to each other.** Now consider the location of region 2. By Lemma 7, we know that either $2 \subset P(1, 3)$ or $2 \subset P(3, 1)$. These two situations are isomorphic, so we can assume that $2 \subset P(1, 3)$ without loss of generality. This situation is depicted in Figure 2.28. Now consider where region 4 may be located. According to Lemma 9, either $4 \subseteq Cons(1, 3)$, or $4 \subset P(3, 1)$.

Subcase(i): $4 \subseteq Cons(1, 3)$**.** So that region 4 does not indirectly see region

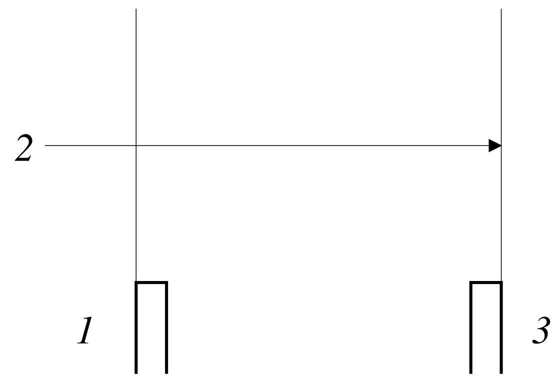

Figure 2.28: Region 2 is in $P(1,3)$. The sight line from region 2 to $Boundary(3,1)$ has been shown.

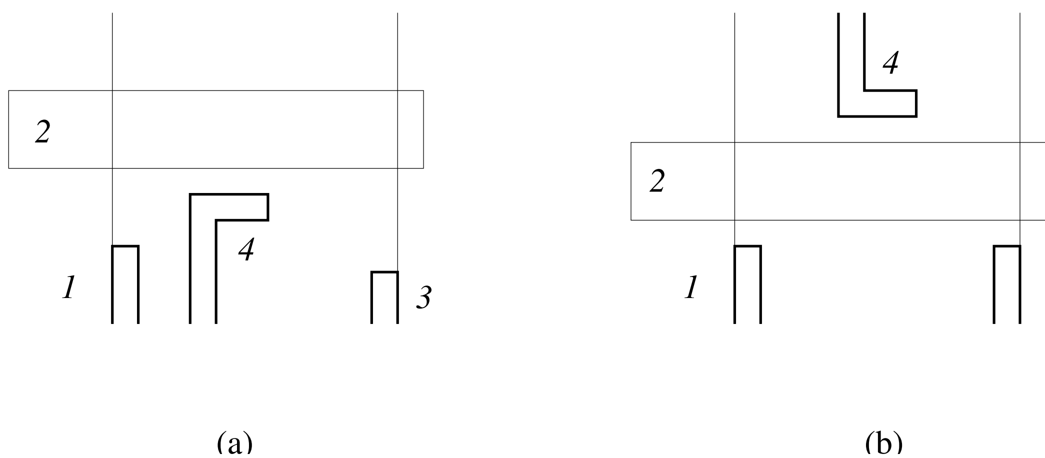

Figure 2.29: Region 4 is in $Cons(1,3)$. (a) Region 4 is above $R^2_{\mathcal{H}(2,3)}$. (b) Region 4 is below $R^2_{\mathcal{H}(2,3)}$.

2, there must be a chain of edges from $E(P)$ within $Cons(1,2)$ that blocks region 4 from seeing $R^2_{\mathcal{H}(2,3)}$[5] since $4 \not\!A 2$. These edges must not interfere with region 4's visibility with region 3. Since these edges horizontally separate region 4 and $R^2_{\mathcal{H}(2,3)}$, we know that region 4 is either below $R^2_{\mathcal{H}(2,3)}$ or region 4 is above $R^2_{\mathcal{H}(2,3)}$. These two situations are depicted in Figure 2.29. Now we can consider where region 5 may be located with respect to these two configurations. We can immediately notice that $5 \not\subset P(1,3)$ since this would imply that $4 \not\!A 5$. This lack of indirect visibility between

[5] Recall that $R^2_{\mathcal{H}(2,3)}$ is any orthogonal sub-rectangle of P that contains region 2 and at least one point from $\mathcal{H}(2,3)$.

regions 4 and 5 would be caused by the chain of edges that block region 4 from seeing $R^2_{\mathcal{H}(2,3)}$ (refer to Figure 2.29).

First consider the case where $n = 5$ (i.e. H is a 5-hole). This implies that region 5 indirectly sees region 1. This forbids region 5 from being in $P(3,1)$ due to Lemma 9. This implies $5 \subseteq Cons(1,3)$ for the configurations depicted in Figure 2.29. Since $5 \wedge 1$, and since region 5 is now assumed to be in $Cons(1,3)$, we know that region 5 sees a portion of $Boundary(1,3)$. It must be located in such a way that it cannot indirectly see region 2. This implies that region 5 cannot see $R^2_{\mathcal{H}(2,3)}$. This implies that region 5 is either strictly above or strictly below $R^2_{\mathcal{H}(2,3)}$. If region 5 is above $R^2_{\mathcal{H}(2,3)}$, then there is an edge below region 5 that blocks region 5 from seeing $R^2_{\mathcal{H}(2,3)}$. In this situation, region 4 may also be above $R^2_{\mathcal{H}(2,3)}$. In this scenario, the fact that $5 \wedge 1$ will imply that $4 \not\wedge 3$, which is a contradiction (Figure 2.30(d)). If region 4 is below $R^2_{\mathcal{H}(2,3)}$, while region 5 is above, then $5 \not\wedge 4$ since both 4 and 5 cannot indirectly see $R^2_{\mathcal{H}(2,3)}$ (Figure 2.30(a)). These same arguments forbid the possibility that region 5 is below $R^2_{\mathcal{H}(2,3)}$ (Figure 2.30(b) and (c)). This implies that region 5 cannot be positioned anywhere if $4 \subseteq Cons(1,3)$ and $n = 5$.

Now suppose that $n \geq 7$ (recall that we are dealing with *odd* holes). Since $n \geq 7$, we have that $5 \subset P(3,1)$ or $5 \subseteq Cons(1,3)$. First suppose $5 \subseteq Cons(1,3)$. Figures 2.30 (a) and (c) depict cases that do not rely on any assumption about the value of n, so we can rule out the possibility that region 4 and region 5 are on opposite sides of $R^2_{\mathcal{H}(2,3)}$ even if $n \geq 7$. Figures 2.31 (a) and (b) show the two configurations where region 4 and 5 are on the same side as $R^2_{\mathcal{H}(2,3)}$. Now consider the line embedding Γ of the path $L = < 5, 6, ..., n, 1 >$. Due to our current configuration of regions 4 and 5, a line segment from Γ will cross over $R^4_{\mathcal{H}(3,4)}$ (Figure 2.31 (c) and (d)). We can easily argue that line segment from Γ that crosses $R^4_{\mathcal{H}(3,4)}$ will have an endpoint that corresponds to a basic region from the set $\{6, 7, ..., n\}$ since region 5 has an edge below it which is lower (resp. higher) that region 4. This causes a chord in H, which is a contradiction. Thus $5 \not\subseteq Cons(1,3)$. The only remaining possibility for this case is that $5 \subset P(3,1)$. Now consider the line embedding Γ of the path $L = < 5, 6, ..., n, 1 >$.

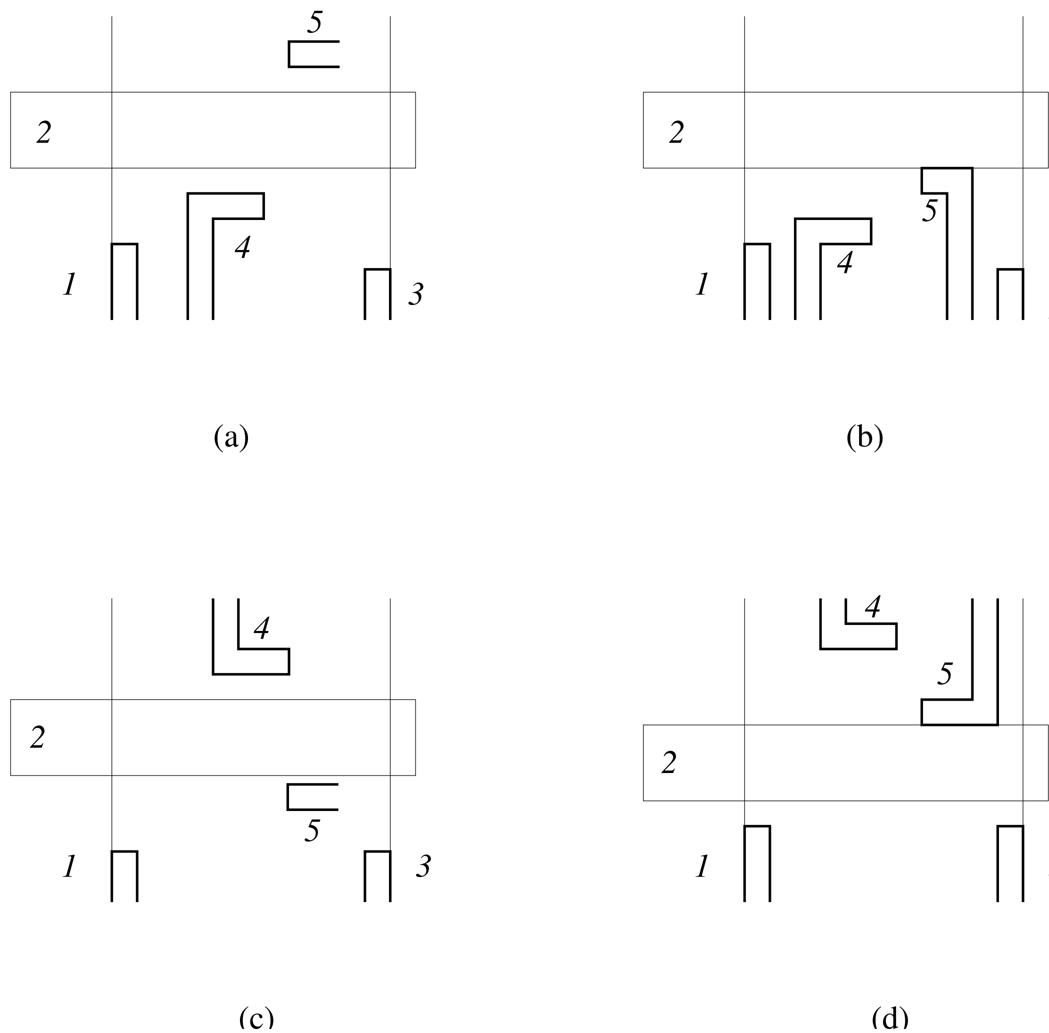

Figure 2.30: Region 5 is in $Cons(1,3)$. (a) Region 5 is above $R^2_{\mathcal{H}(2,3)}$, while region 4 is below $R^2_{\mathcal{H}(2,3)}$. (b) Region 5 and region 4 are below $R^2_{\mathcal{H}(2,3)}$. (c) Region 5 is below $R^2_{\mathcal{H}(2,3)}$, while region 4 is above $R^2_{\mathcal{H}(2,3)}$. (d) Region 5 and region 4 are above $R^2_{\mathcal{H}(2,3)}$.

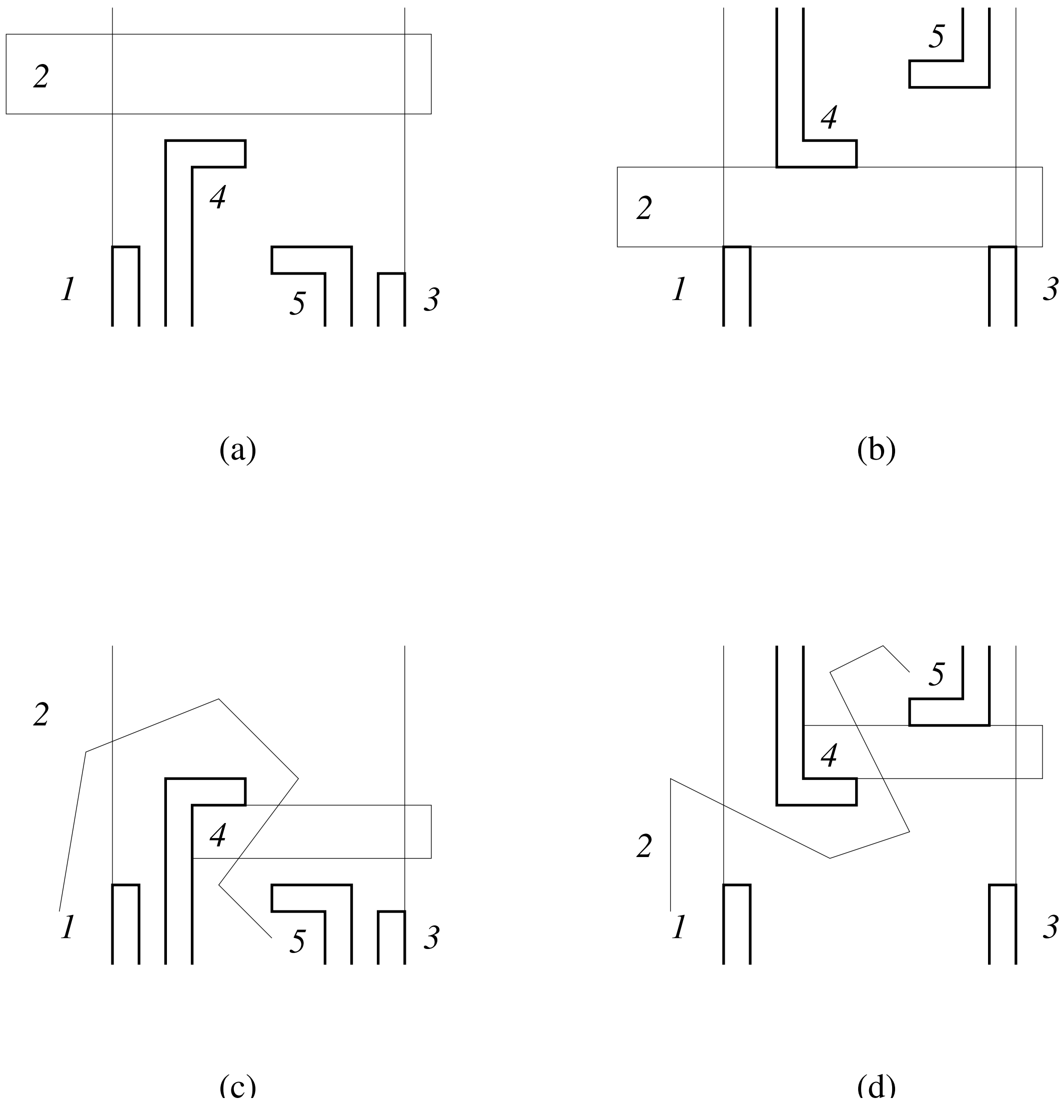

Figure 2.31: Region 4 and 5 are in $Cons(1,3)$ for $n \geq 7$. (a) Region 4 and region 5 are below $R^2_{\mathcal{H}(2,3)}$. (b) Region 4 and region 5 are above $R^2_{\mathcal{H}(2,3)}$. (c) and (d) The path from region 5 to region 1 that uses region 6 crosses over $R^4_{\mathcal{H}(3,4)}$.

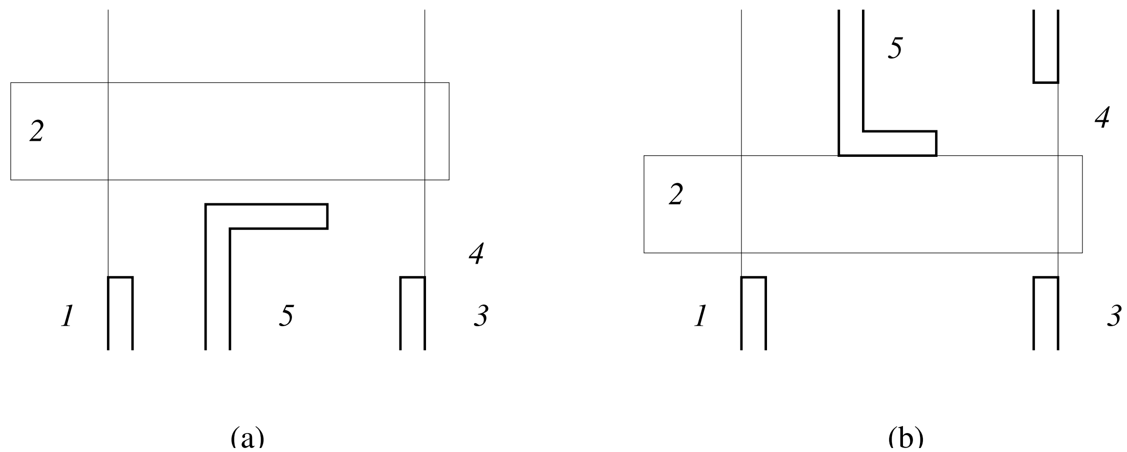

Figure 2.32: Region 5 is in $Cons(1,3)$. (a) Region 5 is below $R^2_{\mathcal{H}(2,3)}$. (b) Region 5 is above $R^2_{\mathcal{H}(2,3)}$.

Since $5 \subset P(3,1)$, a line segment from Γ must cross $Boundary(3,1)$. This implies that some region along L indirectly sees region 3, which is a contradiction. Since this is the final location for region 5 for $n \geq 7$, we have that $4 \not\subseteq Cons(1,3)$. This ends this subcase.

Subcase(ii): $4 \subset P(3,1)$**.** Now let us consider the possible locations for region 5. First suppose $5 \subset P(1,3)$. If there exists a point $h \in \mathcal{H}(4,5)$ within $P(1,3)$ then $4 \wedge 1$, which is a contradiction. If there exists a point $h \in \mathcal{H}(4,5)$ within $P(3,1)$ then $5 \wedge 3$, which is a contradiction. The only remaining possibility is that $\mathcal{H}(4,5) \subseteq Cons(1,3)$. This implies that region 4 and region 5 have indirect diagonal visibility, for if they did not then $4 \wedge 1$ or $5 \wedge 3$. Since regions 4 and 5 have indirect diagonal visibility, there are edges from $E(P)$ within $Cons(1,3)$ that force $R^2_{\mathcal{H}(2,3)}$ to overlap with either $R^4_{\mathcal{H}(4,5)}$ or $R^5_{\mathcal{H}(4,5)}$ if it is to also overlap with $Boundary(3,1)$. In either case we have a contradiction, so we have $5 \not\subset P(1,3)$.

Now suppose that $5 \subseteq Cons(1,3)$. In this case, there must be a chain of edges from $E(P)$ that is in $Cons(1,3)$ that blocks region 5 from seeing $R^2_{\mathcal{H}(2,3)}$. This implies that region 5 is either strictly above or strictly below $R^2_{\mathcal{H}(2,3)}$ (Figure 2.32). From this Figure we can see that we can rule out the possibility that $n = 5$. This is due to the fact that there is now a edge that blocks region 5 from indirectly seeing region 1. Thus

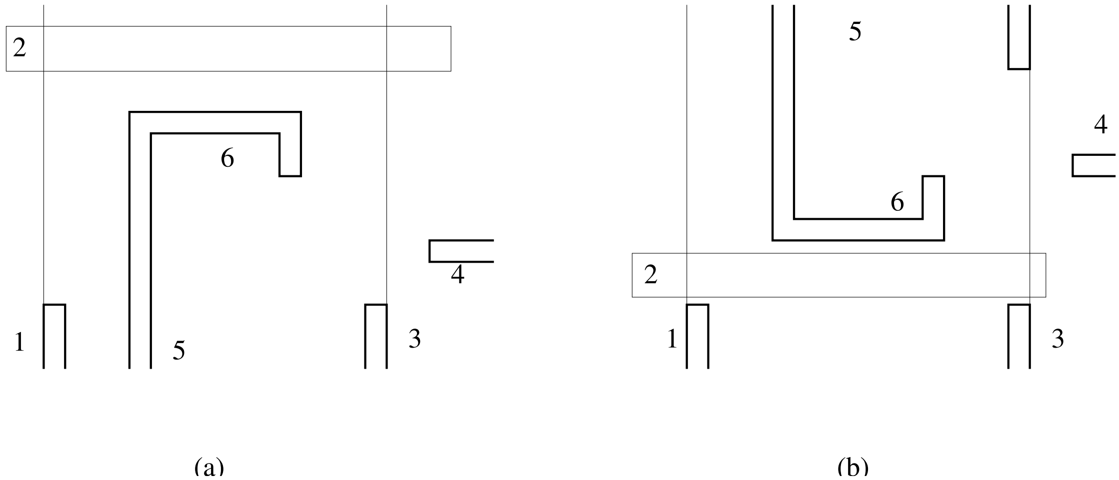

Figure 2.33: (a) Region 5 and 6 are below $R^2_{\mathcal{H}(2,3)}$. (b) Region 5 and 6 are above $R^2_{\mathcal{H}(2,3)}$.

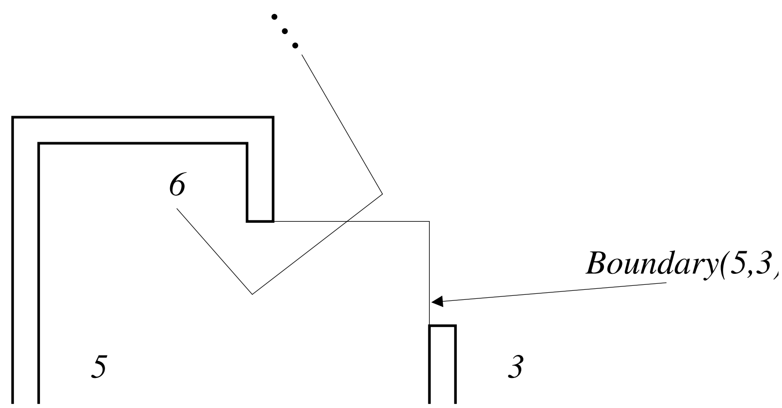

Figure 2.34: A closeup of Γ as it crosses $Boundary(5,3)$.

$n \geq 7$. Now let us consider where region 6 may be located. We know $6 \not\subset P(1,3)$, otherwise $5 \not\wedge 6$. If $6 \subset P(3,1)$ then region 5 or 6 will be forced to see a portion of $Boundary(3,1)$ since $6 \wedge 5$. Since this is a contradiction, $6 \subseteq Cons(3,1)$. If region 5 is below $R^2_{\mathcal{H}(2,3)}$ (Figure 2.32(a)) then region 6 must also be below $R^2_{\mathcal{H}(2,3)}$, otherwise $2 \wedge 6$. For the same reason, if region 5 is above $R^2_{\mathcal{H}(2,3)}$, then region 6 is also above $R^2_{\mathcal{H}(2,3)}$. The case where both region 5 and region 6 are below $R^2_{\mathcal{H}(2,3)}$ is depicted in Figure 2.33(a). Now consider the line embedding Γ of the path $L = < 6, 7, ..., n, 1 >$. Figure 2.34 shows a closeup of the portion of Γ at $Boundary(5,3)$. We can see that a line segment from Γ will cross $Boundary(5,3)$. This implies that region 5 will see

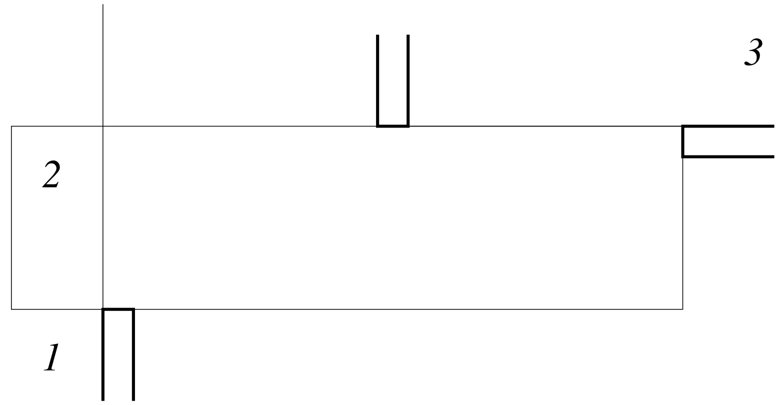

Figure 2.35: Region 2 sees all of $Boundary(3, 1)$.

some region along this path. We also know that this region is not 6 since region 6 cannot directly see $Boundary(5, 3)$. This is a contradiction, so $5 \not\subseteq Cons(1, 3)$. The only remaining possibility $5 \subset P(3, 1)$. Consider the line embedding Γ of the path $L =< 5, 6, ..., n, 1 >$. Since $5 \subset P(3, 1)$, a line segment from Γ will cross $Boundary(3, 1)$. This implies that some region along this path indirectly sees region 3, which is a contradiction. Thus region 5 cannot be located anywhere if $4 \subset P(3, 1)$. This implies that this case cannot exist.

Case(ii): $Boundary(1, 3)$ **and** $Boundary(3, 1)$ **are both orthogonal line segments, and these boundaries are perpendicular do each other.** Lemma 7 implies that region 2 may be in one of two different positions: $2 \subset P(1, 3)$ or $2 \subset P(3, 1)$. These two cases are isomorphic, thus without loss of generality we can assume that $2 \subset P(1, 3)$. Since $2 \subset P(1, 3)$, Lemma 6 tells us that region 2 sees a portion of $Boundary(3, 1)$. Moreover, since $Boundary(1, 3)$ and $Boundary(3, 1)$ are perpendicular, region 2 sees *all* of $Boundary(3, 1)$ (Figure 2.35). We will make extensive use of this fact in the following paragraph.

Now consider the location of region 4. Lemma 9 tells us that either $4 \subseteq Cons(1, 3)$ or $4 \subset P(3, 1)$. If $4 \subseteq Cons(1, 3)$ then Lemma 6 implies region 4 sees a portion of $Boundary(3, 1)$. Since region 2 sees all of $Boundary(3, 1)$, $2 \wedge 4$, which is a contradiction. Thus $4 \subset P(3, 1)$. Now that we have fixed the location of region 4,

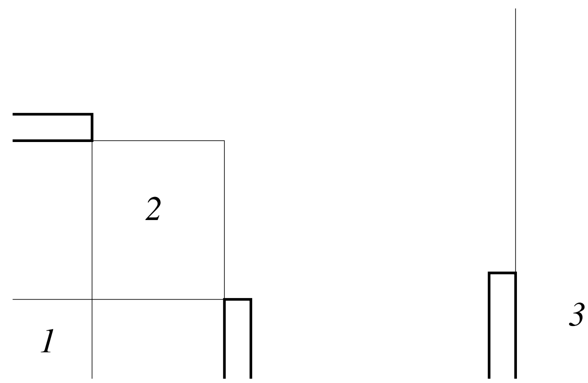

Figure 2.36: Region 2 is in 1's bulge.

we will consider the location of region 5. If $5 \subset P(1,3)$ or $5 \subseteq Cons(1,3)$ then either region 4 or region 5 sees a portion of $Boundary(3,1)$ since $4 \subset P(3,1)$. This implies $5 \wedge 2$ or $4 \wedge 2$, which are both contradictions. This implies that $5 \subset P(3,1)$. If we consider the line embedding Γ of the path $L =< 5, 6, .., n, 1 >$, we see that Γ must contain a line segment that crosses $Boundary(3,1)$. This implies that a basic region from the set $\{5, 6, ..., n, 1\}$ indirectly sees region 3, which is a contradiction. This ends Case(ii).

Case(iii): One of $Boundary(1,3)$ or $Boundary(3,1)$ is an orthogonal line segment, and the other is an orthogonal 2-chain. Without loss of generality, we will assume that $Boundary(1,3)$ is an orthogonal 2-chain, and $Boundary(3,1)$ is an orthogonal line segment. According to Lemma 7, $2 \subset P(1,3)$ or $2 \subset P(3,1)$. First we will assume $2 \subset P(1,3)$. Since $2 \wedge 3$ and $2 \subset P(1,3)$, Lemma 6 tells us that region 2 sees a portion of $Boundary(3,1)$. Using Observations 1 and 2, this implies that region 2 cannot be in 1's zone or in 1's parallel zone. Thus region 2 must be in 1's bulge or in 1's perpendicular zone. First assume that region 2 is in 1's bulge. This situation is shown is Figure 2.36. Now consider the location of region 4. According to Lemma 9, $4 \subseteq Cons(1,3)$ or $4 \subset P(3,1)$. We will consider these two cases simultaneously by considering the line embedding Γ of the path $L =< 4, 5, ..., n, 1 >$. Due to our current configuration, there is some line segment from Γ that crosses $Boundary(1,3)$. This line segment will have an endpoint that corresponds to a region

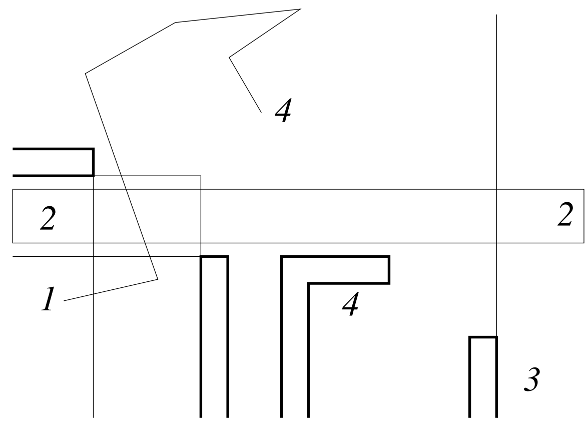

Figure 2.37: Region 4 is in $Cons(1,3)$: region 4 is either above $R^2_{\mathcal{H}(2,3)}$ or it is below.

from the set $\{4, 5, ..., n\}$. Notice that region 1 is excluded from this set since region 1 cannot see any points from $Cons(1,3)$ or $P(3,1)$. This implies that the region corresponding to this endpoint will see some portion of $Boundary(1,3)$. Recall that region 2 is in 1's bulge, and hence it sees all of $Boundary(1,3)$. Thus we have $2 \wedge r$ for $r \in \{4, 5, ..., n\}$, which is a contradiction. This implies that region 2 cannot be in region 1's bulge.

Hence region 2 must be in region 1's perpendicular zone if $2 \subset P(1,3)$. We will consider this case at the same time as we consider the case where $2 \subset P(3,1)$. We will be able to do this since in either case either $R^2_{\mathcal{H}(2,3)}$ or $R^2_{\mathcal{H}(1,2)}$ will span $Cons(1,3)$ from $Boundary(1,3)$ to $Boundary(3,1)$. Now consider possible positions for region 4. Lemma 9 tells us that either $4 \subseteq Cons(1,3)$ or $4 \subset P(3,1)$. First consider the case where $4 \subseteq Cons(1,3)$. The two ways that this situation may occur are shown in Figure 2.37: either region 4 is above $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$) or it is below. If it is above $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$), then the line embedding Γ of the path $L =< 4, 5, ..., n, 1 >$ will contain a line segment that crosses over $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$). This implies that some basic region from the set $\{4, 5, ..., n\}$. will indirectly sees region 2, which is a contradiction. Notice that region 1 is excluded from the set $\{4, 5, ..., n\}$ since region 1 cannot see any point from $Cons(1,3)$.

Thus region 4 must be below $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$). In this situation, there

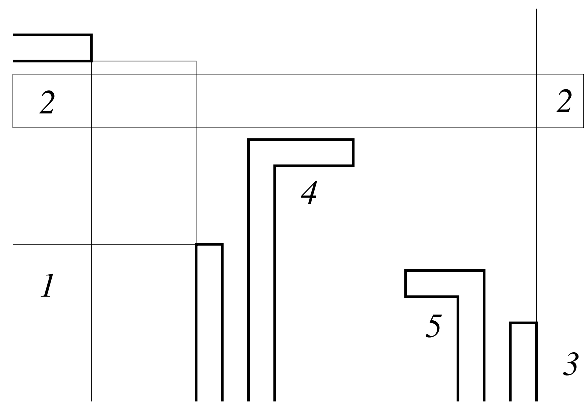

Figure 2.38: Region 5 is in $Cons(1,3)$.

must be an edge from $E(P)$ that is above region 4. This edge will block region 4 from seeing $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$) (Figure 2.37). Now consider where region 5 may be located. If $5 \subset P(1,3)$, then $5 \not\wedge 4$. If $5 \subset P(3,1)$ then the line embedding Γ of the path $L =< 5, 6, .., n, 1 >$ will contain a line segment that crosses over $Boundary(3,1)$. This implies that some region along L indirectly sees region 3, which is a contradiction. Thus $5 \subseteq Cons(1,3)$. If region 5 is above $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$) then $5 \wedge 2$ since $5 \wedge 4$, which is a contradiction. This implies that region 5 is below $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$). Region 5 must be positioned in such a way that is can indirectly see region 4, but cannot indirectly see region 3. Also, the edge that blocks region 5's indirect visibility to 3 must not block 4's indirect visibility to region 3. There must also be an edge above region 5 that blocks it's visibility with $R^2_{\mathcal{H}(2,3)}$ (resp. $R^2_{\mathcal{H}(1,2)}$). Figure 2.38 depicts this scenario. Now consider the line embedding Γ of the path $L =< 5, 6, ..., n, 1 >$. Due to our current configuration of regions 4 and 5, a line segment from Γ must cross over $R^4_{\mathcal{H}(3,4)}$ (Figure 2.39). This implies that some region from L indirectly sees region 4, which is a contradiction. Thus region 5 cannot be located anywhere if $4 \subseteq Cons(1,3)$.

This implies that $4 \subset P(3,1)$. Now consider where region 5 may be located. Suppose that $5 \subset P(1,3)$. If region 5 is in 1's bulge then $5 \wedge 2$ since region 2 sees a portion of $Boundary(1,3)$ and region 5 sees all of $Boundary(1,3)$. If region 5 is in 1's zone, then $5 \not\wedge 4$. Suppose region 5 is in 1's perpendicular zone. If there

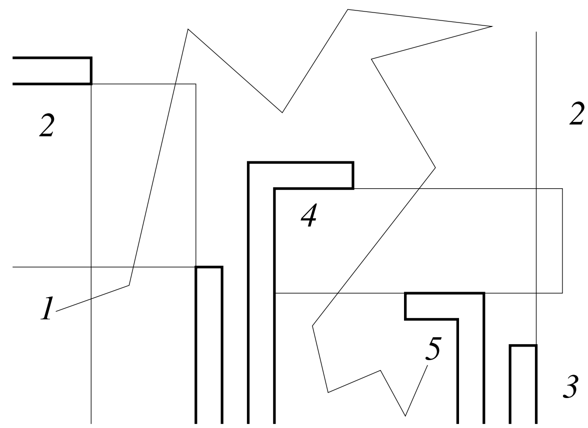

Figure 2.39: A line segment from Γ must cross over $R^4_{\mathcal{H}(3,4)}$.

exists $h \in \mathcal{H}(4,5)$ that is in $P(1,3)$ then $4 \wedge 1$, which is a contradiction. If there exists $h \in \mathcal{H}(4,5)$ that is in $P(3,1)$ then $5 \wedge 3$, which is a contradiction. Thus $\mathcal{H}(4,5) \subseteq Cons(1,3)$, and hence 5 has indirect *diagonal* visibility with region 4. The two different chains of edges that force this visibility to be diagonal will also block $2 \wedge 3$ or $2 \wedge 1$ depending on whether $2 \subset P(1,3)$ or $2 \subset P(3,1)$ (recall we are considering these two cases simultaneously). Thus if $5 \subset P(1,3)$ then it must be in 1's parallel zone. This situation implies that region 5 indirectly sees region 2, which is a contradiction. Thus $5 \not\subset P(1,3)$ if we assume that $4 \subset P(3,1)$. If $5 \subset P(3,1)$ then the line embedding Γ of the path $L =< 5, 6, ..., n, 1 >$ will contain a line segment that crosses $Boundary(3,1)$. This implies that some region from L will indirectly sees region 3, which is a contradiction. Thus $5 \subseteq Cons(1,3)$ for our current configuration of region 4 and region 2. We know that region 4 and 5 cannot directly see each other since this would imply that $5 \wedge 3$. So region 4 and 5 have indirect visibility. The region $\mathcal{H}(4,5)$ must be in $Cons(1,3)$, for if it were in $P(3,1)$ then region 5 would indirectly see region 3. Our current situation is shown is Figure 2.40. One of the key aspects of this situation is that region 5 must have an edge from $E(P)$ above it so that it does not see $R^2_{\mathcal{H}(2,3)}$. We can see that this implies that $5 \not\wedge 1$, which implies that $n \geq 7$. So we can consider the possible locations for region 6. If $6 \subset P(1,3)$ then $6 \not\wedge 5$. If $6 \subset P(3,1)$ then the line embedding Γ of the path $L =< 6, 7, ..., n, 1 >$ will contain a line segment that crosses $Boundary(3,1)$.

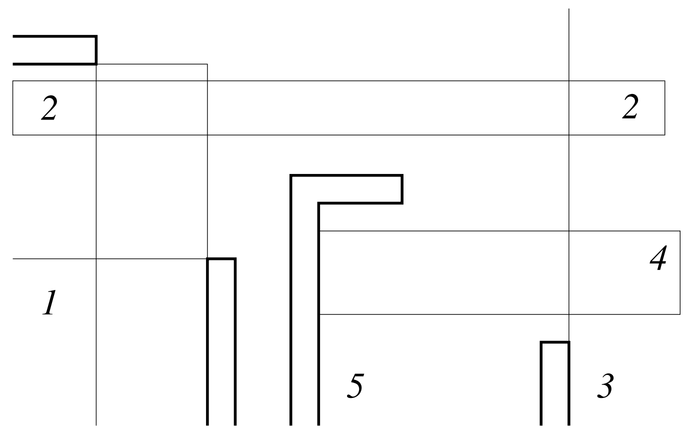

Figure 2.40: Region 5 is in $Cons(1,3)$, while $4 \subset P(3,1)$.

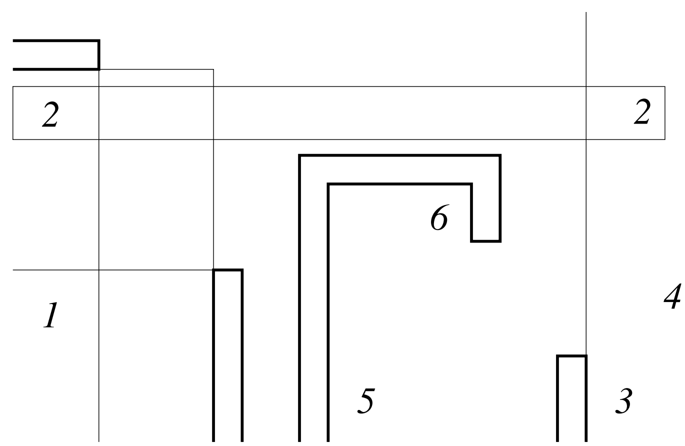

Figure 2.41: Region 6 is in $Cons(1,3)$.

This implies that region 3 indirectly sees some region along this path, which is a contradiction. Thus $6 \subseteq Cons(1,3)$. For this to occur, region 6 must be below $R^2_{\mathcal{H}(2,3)}$, otherwise region 6 will see $R^2_{\mathcal{H}(2,3)}$ since it indirectly sees region 5. Region 6 must be positioned so that region 5 can still indirectly see region 4. This situation is shown in Figure 2.41. Figure 2.42 shows a closeup of $Boundary(5,3)$ with respect to Figure 2.41. Due to our current configuration of regions 5 and 6, we have that the line embedding Γ of the path $L =< 6, 7, ..., n, 1 >$ will contain a line segment that crosses $Boundary(5,3)$. This corresponds to a region from L that is not region 6 indirectly seeing region 5. This is a contradiction, which implies that region 5 cannot be in $Cons(1,3)$. Thus region 4 cannot be positioned anywhere. This ends this case.

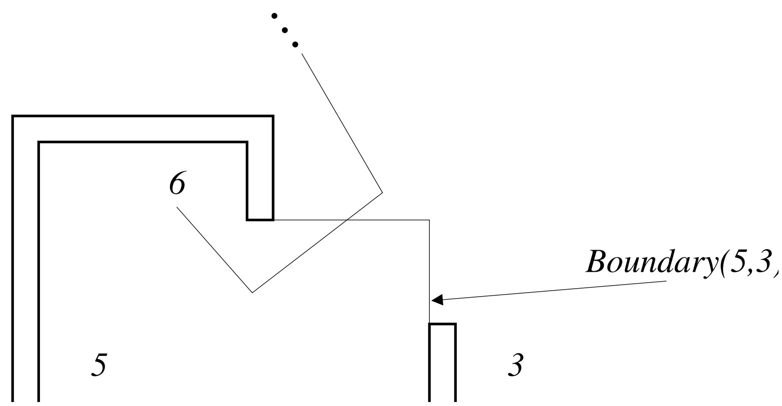

Figure 2.42: A closeup of Γ as it crosses $Boundary(5,3)$.

Case(iv): $Boundary(1,3)$ **and** $Boundary(3,1)$ **are orthogonal 2-chains.** First we will consider the position of region 2. Lemma 7 tells us that either $2 \subset P(1,3)$ or $2 \subset P(3,1)$. These two possibilities are isomorphic, so we can, without loss of generality, assume $2 \subset P(1,3)$. Since $2 \wedge 3$, we can limit where it may be positioned with respect to region 1. Region 2 cannot be in 1's zone or in 1's parallel zone, otherwise region 2 could not indirectly see region 3. Thus region 2 is either in 1's bulge or in 1's perpendicular zone.

By Lemma 9, we know either $4 \subseteq Cons(1,3)$ or $4 \subset P(3,1)$. Now consider the line embedding Γ the path $L =< 4, 5, ..., n, 1 >$. Due to the location of region 4, Γ must contain a line segment that crosses $Boundary(1,3)$. This implies that some region from L sees a portion of $Boundary(1,3)$. If we suppose that region 2 is in 1's bulge, then region 2 indirectly sees some region along this path, which is a contradiction. Thus region 2 must be in 1's perpendicular zone. Now let us consider the two possible locations for region 4 for this position of region 2.

Subcase(i): $4 \subseteq Cons(1,3)$**.** Let us consider the relative positions of region 4 and $R^2_{\mathcal{H}(2,3)}$. Suppose region 4 is above $R^2_{\mathcal{H}(2,3)}$. Then the line embedding Γ of the path $L =< 4, 5, ..., n, 1 >$ must contain a line segment that crosses over $R^2_{\mathcal{H}(2,3)}$. This implies that region 2 indirectly sees some region from L, which is a contradiction. Thus region 4 must be below $R^2_{\mathcal{H}(2,3)}$. This situation is shown in Figure 2.43. Now consider where region 5 is located. If $5 \subset P(1,3)$ then region 5 will not be able

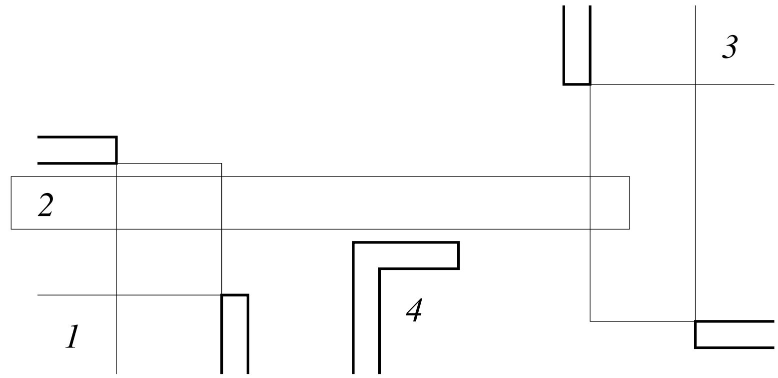

Figure 2.43: Region 4 is in $Cons(1,3)$ and is situated below $R^2_{\mathcal{H}(2,3)}$.

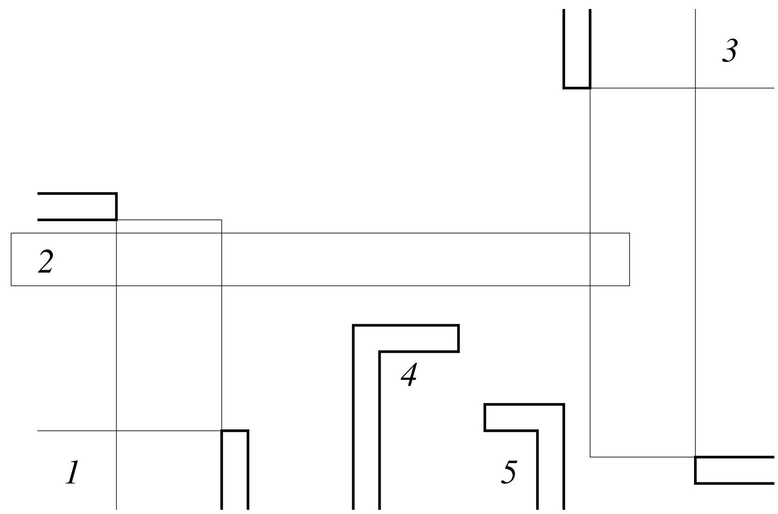

Figure 2.44: Region 5 is in $Cons(1,3)$ and is situated below $R^2_{\mathcal{H}(2,3)}$.

to indirectly see region 4, which is a contradiction. If $5 \subset P(3,1)$ then the line embedding Γ of the path $L =< 5, 6, ..., n, 1 >$ will contain a line segment that crosses $Boundary(3,1)$. This implies that region 3 indirectly sees a region from L, which is a contradiction. Thus $5 \subseteq Cons(1,3)$. Region 5 cannot be above $R^2_{\mathcal{H}(2,3)}$ since it indirectly sees region 4, but does not indirectly see region 2 (refer to Figure 2.43). Thus region 5 must be below $R^2_{\mathcal{H}(2,3)}$. As in the previous cases, we can see that the line embedding Γ of the path $L =< 5, 6, ..., n, 1 >$ contains a line segment that crosses $R^4_{\mathcal{H}(3,4)}$. This line segment will cross in such a way that a region that is not region 5 will indirectly see region 4, which is a contradiction. This implies that $4 \not\subseteq Cons(1,3)$.

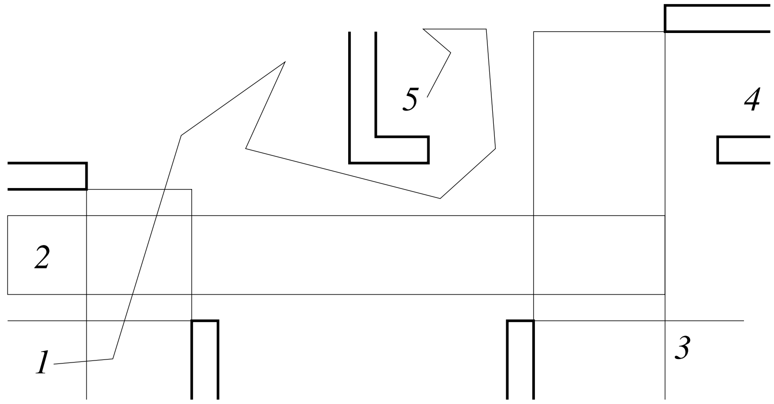

Figure 2.45: $4 \subset P(3,1)$, $5 \subseteq Cons(1,3)$, and 5 is situated above $R^2_{\mathcal{H}(2,3)}$.

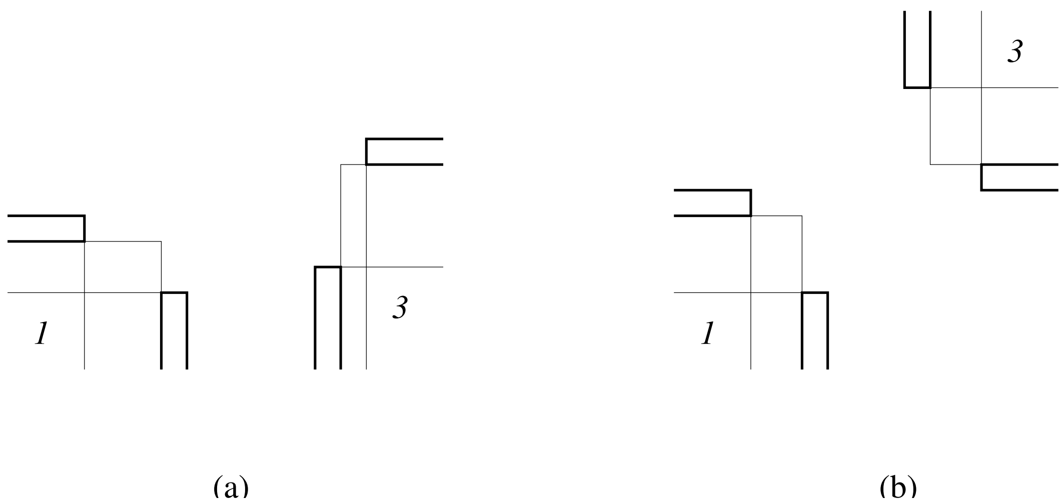

Figure 2.46: The respective positions of $Boundary(1,3)$ and $Boundary(3,1)$. (a) The boundaries are aligned. (b) The boundaries are flipped.

Subcase(ii): $4 \subset P(3,1)$**.** As in the previous subcase, we can restrict region 5 to be located in $Cons(1,3)$. Now consider the relative positions of region 5 and $R^2_{\mathcal{H}(2,3)}$. Suppose that region 5 is above $R^2_{\mathcal{H}(2,3)}$ (Figure 2.45). This implies that the line embedding Γ of the path $L =< 5, 6, ..., n, 1 >$ will contain a line segment that crosses over $R^2_{\mathcal{H}(2,3)}$. This implies that region 2 sees some region from L, which is a contradiction. Thus region 5 must be below $R^2_{\mathcal{H}(2,3)}$.

At this point we need to consider how $Boundary(1,3)$ and $Boundary(3,1)$ are positioned with respect to each other. These boundaries may be "aligned" (Figure

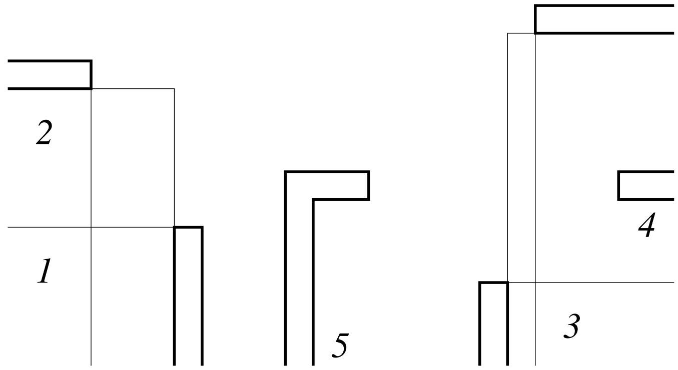

Figure 2.47: Region 5 is in $Cons(1,3)$ and is situated below $R^2_{\mathcal{H}(2,3)}$, while $4 \subset P(3,1)$.

2.46(a)) or they may be "flipped" (Figure 2.46(b)). We will consider these two cases separately.

First let us consider the case where $Boundary(1,3)$ and $Boundary(3,1)$ are "aligned". Recall that we know that $5 \subseteq Cons(1,3)$ and that 5 is below $R^2_{\mathcal{H}(2,3)}$. We can also say something about the location of $\mathcal{H}(4,5)$. If there exists $h \in \mathcal{H}(4,5)$ that is in $P(3,1)$ then region 5 sees some portion of $Boundary(3,1)$. This implies $5 \wedge 3$, which is a contradiction. Thus $\mathcal{H}(4,5)$ must be in $Cons(1,3)$. This observation, along with the fact that region 2 sees through 3's bulge, implies that region 4 must be in 3's perpendicular zone. This situation is depicted in Figure 2.47. We can see that this situation implies that $5 \not\wedge 1$, which implies that $n \geq 7$. Now we can consider the possible locations for region 6. If $6 \subset P(1,3)$ then $6 \not\wedge 5$. If $6 \subset P(3,1)$ then the line embedding Γ of the path $L =< 6, 7, ..., n, 1 >$ will contain a line segment that crosses $Boundary(3,1)$. This implies that region 3 indirectly sees some region from L, which is a contradiction. Thus $6 \subseteq Cons(1,3)$. For this to occur, region 6 must be below $R^2_{\mathcal{H}(2,3)}$, otherwise region 6 will see $R^2_{\mathcal{H}(2,3)}$ since it indirectly sees region 5. As in previous cases (see Figures 2.41 and 2.42 for example), we can see that this situation implies that the line embedding Γ of the path $L =< 6, 7, ..., n, 1 >$ will contain a line segment that crosses $Boundary(5,3)$. This line segment will cross this boundary in such a way that one of the regions from L that is not region 6 will indirectly see region 5. This causes a contradiction.

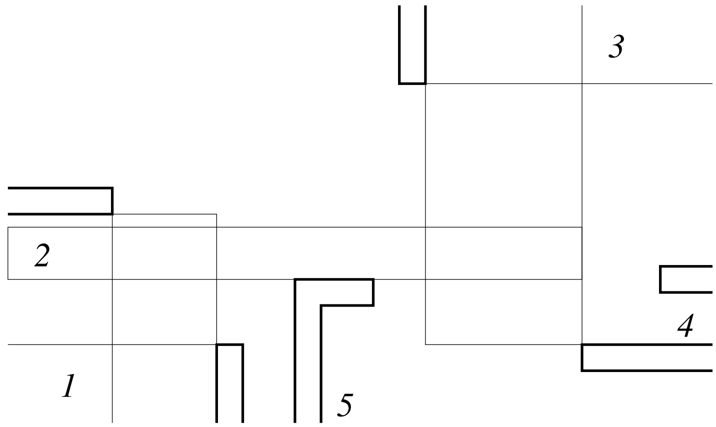

Figure 2.48: $Boundary(1,3)$ and $Boundary(3,1)$ are "flipped". Region 5 is in $Cons(1,3)$.

The only remaining situation to consider for this case is the situation where $Boundary(1,3)$ and $Boundary(3,1)$ are "flipped". Recall that we know that $5 \subseteq Cons(1,3)$ and that region 5 is below $R^2_{\mathcal{H}(2,3)}$. We also know that $\mathcal{H}(4,5)$ is in $Cons(1,3)$. If we combine this fact with the fact that region 2 sees through 3's bulge, we have that region 4 is in 3's perpendicular zone. An example of this configuration is shown in Figure 2.48. Now consider the line embedding Γ of the path $L =< 2, 1, n, n-1, ..., 6, 5 >$. Our current configuration implies that Γ contains a line segment that will cross $R^4_{\mathcal{H}(4,5)}$. This implies that some region from L indirectly sees region 4. This implies a contradiction. $\square$

Now we move onto the next step in showing that all graphs in $\mathcal{RVG}$ are perfect. Since we've shown that graphs from $\mathcal{RVG}$ do not contain odd holes, we must now show that these graphs do not contain odd antiholes.

Lemma 12. *No graph in $\mathcal{RVG}$ has an odd antihole of size 5 or larger.*

Proof. Since 5 holes and 5 antiholes are isomorphic, Lemma 11 implies that all graphs in $\mathcal{RVG}$ are free of 5 antiholes.

Now we only need to show that no graph from $\mathcal{RVG}$ has odd antiholes of size $n \geq 7$. We will prove the Lemma indirectly: let G be any graph from $\mathcal{RVG}$ that

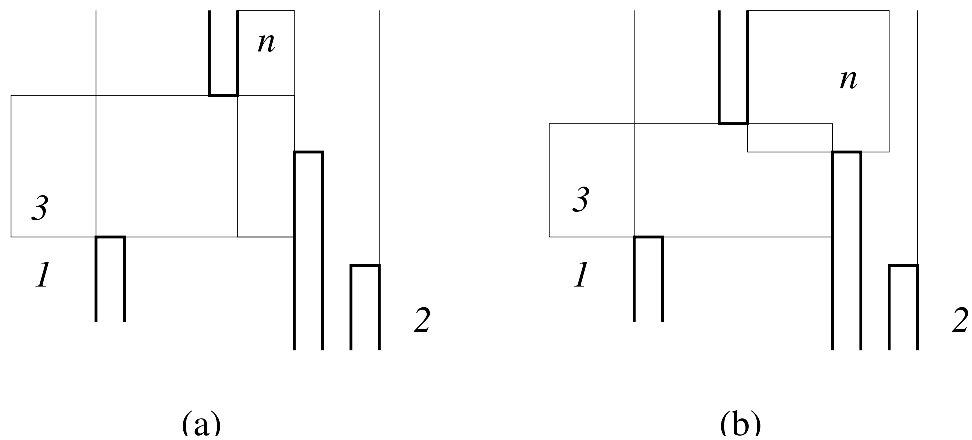

Figure 2.49: Region n is in $Cons(1,2)$. (a) Regions 3 and n have indirect orthogonal visibility. (b) Regions 3 and n have indirect diagonal visibility.

contains a vertex induced odd antihole H of size $n \geq 7$. We will use a cyclic labeling of the vertices in the hole H.

First we make an observation about the location of regions 3 and n. Consider where region 3 and region n can be located with respect to $Cons(1,2)$. Lemma 8 tells us that region 3 and n cannot both be in $Cons(1,2)$ since they are part of an antihole. Also, Lemma 10 tells us that either $n \subseteq Cons(1,2)$ or $3 \subseteq Cons(1,2)$. In other words, *exactly one of region n or region 3 is in* $Cons(1,2)$. We will use this fact throughout the remainder of the proof.

We will break the proof up into 4 major cases based upon the type of boundary $\nu(1)$ and $\nu(2)$ have with $Cons(1,2)$. Lemma 5 tells us that the boundary between a visibility polygon and a constriction region is either a single line segment or it is a orthogonal 2-chain. Thus we have 4 cases to consider:

Case(i): Both $Boundary(1,2)$ and $Boundary(2,1)$ are orthogonal line segments, and these boundaries are parallel. Recall that exactly one of 3 or n is in $Cons(1,2)$. Without loss of generality let us assume that $n \subseteq Cons(1,2)$. This implies that $3 \subset P(1,2)$. We can immediately exclude the possibility that $n \equiv 3$, since this would cause $n \wedge 1$. Thus we only need to consider cases where $n \overline{\wedge} 3$. Region

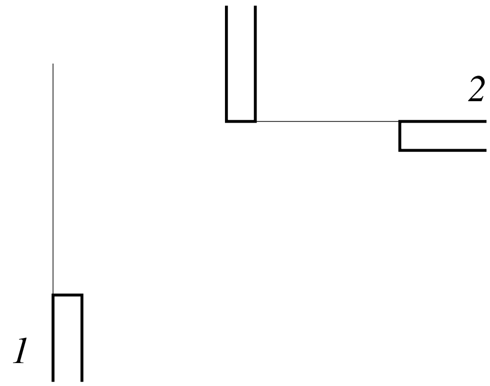

Figure 2.50: $Boundary(1,2)$ and $Boundary(2,1)$ are perpendicular line segments.

3 and n may have indirect orthogonal or diagonal visibility (Figure 2.49).

Now let us consider where region n is with respect to $Cons(3,2)$. Firstly, region 3 sees into $Cons(1,2)$, since $\mathcal{H}(3,n)$ is in $Cons(1,2)$. Thus $Boundary(3,2)$ is at least partly within $Cons(1,2)$. Since $3 \not\wedge 2$, and since $3 \subset P(1,2)$, $Boundary(2,3)$ will be the same as $Boundary(2,1)$. This implies that region n is in $Cons(3,2)$. Applying Lemma 7 to regions 2, 3, and n, we have that H contains a star-cutset, which is a contradiction. This implies that $Boundary(1,2)$ and $Boundary(2,1)$ cannot be parallel line segments.

Case(ii): Both $Boundary(1,2)$ and $Boundary(2,1)$ are orthogonal line segments, and these line segments are perpendicular. This situation is shown in Figure 2.50. First let us consider the case where 4 does not have diagonal visibility to neither 1 nor 2. Lemma 7 implies that region 4 may be in one of two different positions; either $4 \subset P(1,2)$ or $4 \subset P(2,1)$. Due to Lemma 6, if $4 \subset P(1,2)$ then 4 must directly see a portion of $Boundary(2,1)$. Since we've assumed that this visibility is not diagonal, this visibility is not possible since $Boundary(1,2)$ and $Boundary(2,1)$ are perpendicular line segments. Thus $4 \not\subset P(1,2)$ if region 4 does not have diagonal visibility with region 2. If $4 \subset P(1,2)$ then region 4 cannot indirectly see region 1 due to the arrangement of $Boundary(1,2)$ and $Boundary(2,1)$.

This means that 4 must have indirect diagonal visibility with at least one of 1 or

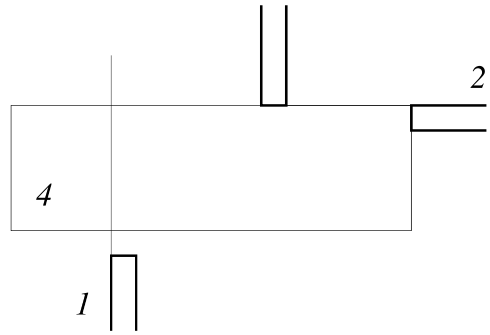

Figure 2.51: Region 4 has diagonal visibility with region 2.

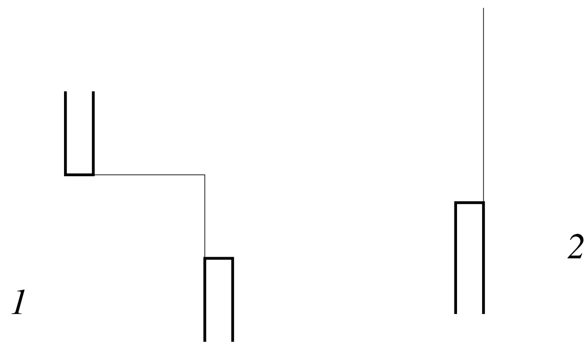

Figure 2.52: $Boundary(1,2)$ is an orthogonal 2-chain, while $Boundary(2,1)$ is a line segment.

2. Without loss of generality we will assume that 4 indirectly sees 2 diagonally. This case is depicted in Figure 2.51. This case only occurs when $\mathcal{H}(2,4)$ is a degenerate rectangle (i.e. a line segment). In fact, $\mathcal{H}(2,4)$ is the same as $Boundary(2,1)$. Now consider where region 5 may be positioned with respect to $Cons(1,2)$ (refer to Figure 2.51). If $5 \subset P(2,1)$ then directly sees a portion of $Boundary(2,1)$ since $5 \wedge 1$. This implies that $5 \wedge 4$, which is a contradiction. If $5 \subset P(1,2)$ then region 5 must see a portion of $Boundary(2,1)$ since $5 \wedge 2$. This implies that $5 \wedge 4$, which is a contradiction. Lemma 7 excludes the possibility that $5 \subseteq Cons(1,2)$ since this would imply that H contains a star-cutset. Thus region 5 cannot be positioned anywhere with respect to $Cons(1,2)$. This excludes the possibility that $Boundary(1,2)$ and $Boundary(2,1)$ are perpendicular line segments.

Case(iii): One of $Boundary(1,2)$ or $Boundary(2,1)$ is an orthogonal line

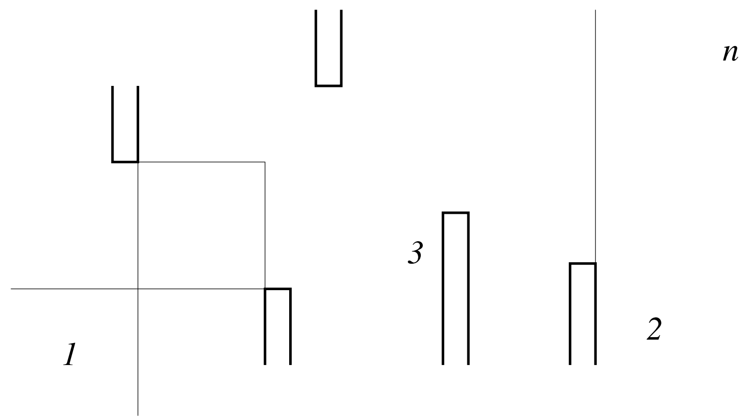

Figure 2.53: Region 3 is in $Cons(1, 2)$ and $n \subset P(2, 1)$.

segment, and one of them is an orthogonal 2-chain. Without loss of generality let us assume that $Boundary(1, 2)$ is an orthogonal 2-chain, while $Boundary(2, 1)$ is a line segment. This situation is depicted in Figure 2.52. Recall that exactly one of region 3 or region n is in $Cons(1, 2)$. Since $Boundary(1, 2)$ and $Boundary(2, 1)$ are different types of boundaries, we have two subcases to consider based upon which one of region 3 and region n is in $Cons(1, 2)$:

Subcase(i): $3 \subseteq Cons(1, 2)$ and $n \subset P(2, 1)$. This situation is depicted in Figure 2.53. We will argue against this subcase by showing that $3 \subseteq Cons(n, 1)$. Once we have shown that $3 \subseteq Cons(n, 1)$, we can use Lemma 7 to show that H contains a star-cutset, which is a contradiction.

We begin with some simple observations for this subcase. If $3 \equiv n$ then region 3 must see a portion of $Boundary(2, 1)$ since $n \subset P(2, 1)$. This implies that region 3 indirectly sees region 2, which is a contradiction. Thus $3 \overline{\wedge} n$. We can also restrict where $\mathcal{H}(3, n)$ may be. The region $\mathcal{H}(3, n)$ must be in $Cons(1, 2)$ otherwise $3 \wedge 2$ or $n \wedge 1$. Now consider where region 3 is with respect to $Cons(1, n)$. Since $n \subset P(2, 1)$, and since $n \not\wedge 1$, we know that $Boundary(1, 2)$ and $Boundary(1, n)$ are the same. We also know that $Boundary(n, 1)$ is partially within $Cons(1, 2)$ since region n sees into $Cons(1, 2)$ (i.e. $n \equiv \mathcal{H}(n, 3)$). Since we've established that region 3 does not directly see region n, region 3 is not contained in the portion of $Cons(1, 2)$ that region n

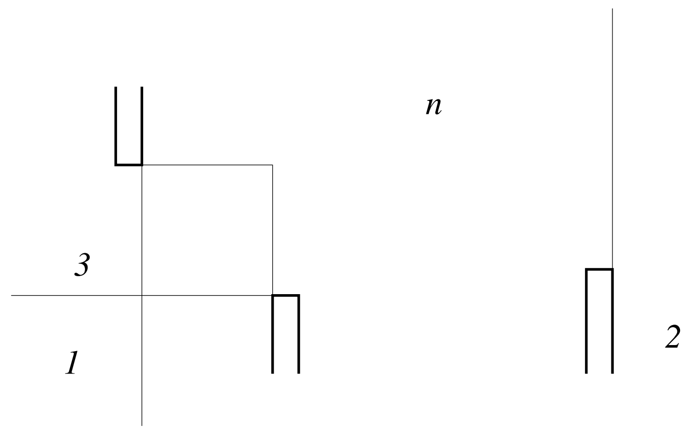

Figure 2.54: Region n is in $Cons(1,2)$, while $3 \subset P(1,2)$.

can see directly, which places region 3 within $Cons(n,1)$. Applying Lemma 7, we have that H contains a star-cutset, which contradicts the assumption that H is an antihole.

Subcase(ii): $n \subseteq Cons(1,2)$ **and** $3 \subset P(1,2)$**.** This situation is depicted in Figure 2.54. We will argue against this subcase in a similar manner to that of the previous subcase.

First consider where region 3 may be positioned with respect to region 1. Region 3 cannot be in 1's zone, otherwise it will not be able to indirectly see region n. Hence region 3 is in one of the following zones: 1's parallel zone, 1's bulge, or 1's perpendicular zone (Figure 2.55). As in the previous subcase, we can limit where $\mathcal{H}(3,n)$ may be positioned: it must be in $Cons(1,2)$, otherwise either $3 \wedge 2$ or $n \wedge 1$. This implies that $Boundary(3,2)$ is partially within $Cons(1,2)$. Since we've established that region 3 does not directly see region n, region n is not contained in the portion of $Cons(1,2)$ that region 3 can see directly, which places region n within $Cons(3,2)$ (Figure 2.55). Applying Lemma 7 to region 2, 3, and n, we have that H contains a star-cutset, which is a contradiction. This eliminates the possibility that $Cons(1,2)$ has one orthogonal line segment boundary and one orthogonal 2-chain boundary.

Case(iv): $Boundary(1,2)$ **and** $Boundary(2,1)$ **are both orthogonal 2-chains.**

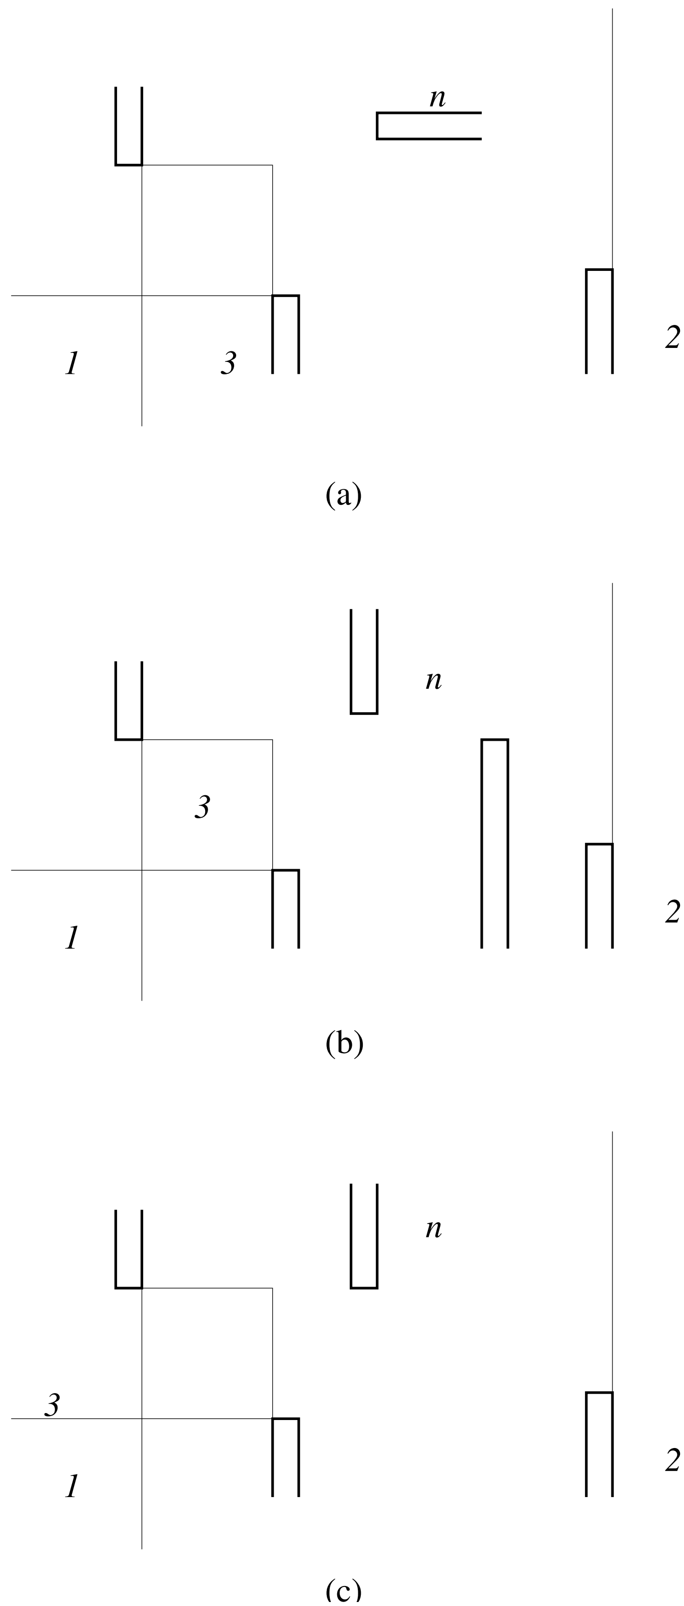

Figure 2.55: The possible positions for region 3 with respect to region 1 (a) Region 3 is in 1's parallel zone. (b) Region 3 is in 1's bulge. (c) Region 3 is in 1's perpendicular zone.

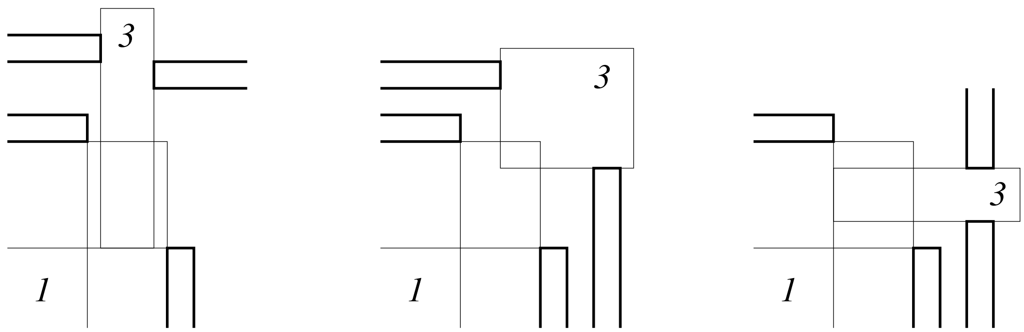

(a) (b) (c)

Figure 2.56: The ways that region 3 can see $Boundary(1, 2)$.

Without loss of generality, let us assume that $3 \subseteq Cons(1, 2)$ and $n \subset P(2, 1)$. There are 3 possible ways that region 3 can indirectly see region 1: it may only see a portion of the "top" edge of $Boundary(1, 2)$, it may only see a portion of the "left" edge of $Boundary(1, 2)$, or it may see a portion of both edges of $Boundary(1, 2)$ (Figure 2.56).

Now let us consider different locations for region 4 with respect to the three configurations shown in Figure 2.56. First let us assume that $4 \subset P(1, 2)$. Since region 4 sees region 2, it cannot be in 1's zone or in 1's parallel zone. If region 4 is in 1's bulge then $3 \wedge 4$, which is a contradiction. So region 4 must be in 1's perpendicular zone. This immediately excludes the case shown in Figure 2.56(a) since this would cause $3 \wedge 4$. For the remaining to cases (i.e. Figures 2.56(b) and (c)), the arrangements force 4 to indirectly see region 3 if it is to indirectly see region 2. This is a contradiction since 4 does indeed indirectly see region 2, but not region 3.

Lemma 7 implies that $4 \not\subseteq Cons(1, 2)$ since if $4 \subseteq Cons(1, 2)$ then H would contain a star-cutset.

So we have that $4 \subset P(2, 1)$. Region 4 cannot be in 2's zone or in 2's parallel zone since $4 \wedge 1$. Now consider the visibility rectangle $R^4_{\mathcal{H}(1,4)}$ with respect to the

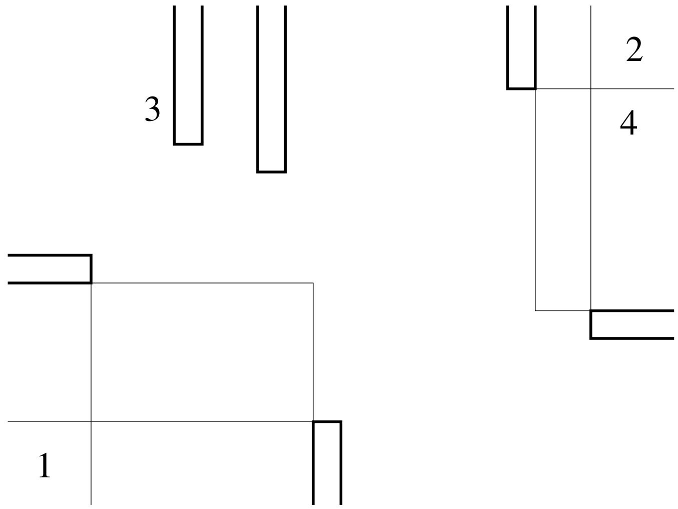

Figure 2.57: One possible way to place region 4 on the same side as region 2 if $3 \subseteq Cons(1,2)$.

three cases shown in Figure 2.56. For case (b), $R^4_{\mathcal{H}(1,4)}$ must overlap with $R^3_{\mathcal{H}(1,3)}$ if it is to overlap with $Boundary(1,2)$. This excludes case (b) from being a possibility. Figure 2.57 shows how we can place region 4 in $P(2,1)$ for the case depicted in Figure 2.56(a). Now consider possible positions for region 5 with respect to $R^4_{\mathcal{H}(1,4)}$. If region 5 is below $R^4_{\mathcal{H}(1,4)}$ then there must be some edge from $E(P)$ above region 5 that blocks region 5 from seeing $R^4_{\mathcal{H}(1,4)}$ since $5 \not\leftrightarrow 4$. This edge will also force $5 \not\leftrightarrow 3$ since region 3 is above $R^4_{\mathcal{H}(1,4)}$. If region 5 is above $R^4_{\mathcal{H}(1,4)}$ then there must be an edge from $E(P)$ that blocks region 5 from seeing $R^4_{\mathcal{H}(1,4)}$. This edge will force $5 \not\leftrightarrow 1$ since region 1 is below $R^4_{\mathcal{H}(1,4)}$. Also, 5 cannot be to the left of $R^4_{\mathcal{H}(1,4)}$ as it sees region 2. In all cases we have a contradiction.

Now we can go on to considering how region 4 may be positioned with respect to Figure 2.56(c). This situation is shown in Figure 2.58. Again we will consider where region 5 is located with respect to $R^4_{\mathcal{H}(1,4)}$. If region 5 is below $R^4_{\mathcal{H}(1,4)}$ then there must be an edge from $E(P)$ that blocks region 5 from seeing $R^4_{\mathcal{H}(1,4)}$. This edge will force $5 \not\leftrightarrow 2$ since region 2 is above $R^4_{\mathcal{H}(1,4)}$. If region 5 is above $R^4_{\mathcal{H}(1,4)}$ then there is an edge from $E(P)$ below region 5 that blocks region 5 from seeing $R^4_{\mathcal{H}(1,4)}$. This edge will force $5 \not\leftrightarrow 3$ since region 3 is below $R^4_{\mathcal{H}(1,4)}$. This ends this subcase. $\square$

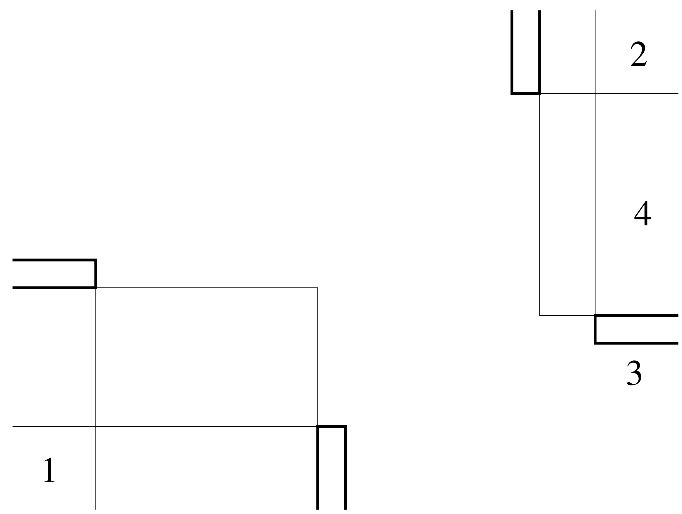

Figure 2.58: The final way to place region 4 on the same as region 2 if $3 \subseteq Cons(1, 2)$.

Algorithm cover(a polygon P) returns R, a minimum r-star cover

1. Construct the region visibility graph of P: $G \leftarrow RVG(P)$.
2. Compute a minimum clique cover of G. Store this clique cover in C.
3. Compute R, the subpolygons associated with the cliques in C.
4. Return R

Figure 2.59: The algorithm for solving **Minimum r-star Cover**.

If we combine Lemma 11 and Lemma 12 with Theorem 3, then we get the following:

Theorem 8. *$\mathcal{RVG}$ is a subclass of the perfect graphs.*

2.8 Solving Minimum r-star Cover in Polynomial Time

In this section we describe our main result of this chapter. Specifically, we will use the results from the previous section to show that we can compute minimum cardinality r-star covers of simple orthogonal polygons in polynomial time.

Theorem 9. *There is a polynomial time algorithm for solving* **Minimum r-star Cover**.

Proof. Now we describe our algorithm for computing minimum r-star covers of simple orthogonal polygons in polynomial time. This algorithm is outlined in Figure 2.59. First we show that this algorithm correctly computes a minimum cardinality r-star cover of P. In step 1 of this algorithm we compute the region visibility graph of P. In step 2 we compute a minimum clique cover of G. Recall that in Theorem 1 we showed that minimum clique covers on graphs in $\mathcal{RVG}$ correspond to minimum r-star covers on polygons. Thus the clique cover computed in step 2 will have a direct correspondence to a minimum r-star cover of P. Step 3 makes this association among the cliques and the r-stars. The resulting r-stars are returned in step 4 of the algorithm.

Now we analyze the time complexity of this algorithm. We will discuss the timing as a function of n, which is the number of vertices on the polygon P. Step 1 of the algorithm computes the region visibility graph of P. We can perform such an operation in polynomial time by using the techniques presented in [39]. We note here that the number of nodes and edges in $RVG(P)$ is polynomial in n. Now consider step 2 of the algorithm given in Figure 2.59. Theorem 8 tells us that all graphs in $\mathcal{RVG}$ are perfect. Theorem 7 tells us that the minimum clique cover problem can be computed in polynomial time on perfect graphs, and hence step 2 can be completed in polynomial time. Step 3 must compute the union of the set of rectangles for each clique in the minimum clique cover. We can easily accomplish this task in polynomial time. $\square$

Chapter 3

Decomposing Polygons into Span Bounded Components

3.1 Introduction

In Section 1.5 we introduced the concept of an α-decomposition. Recall that an α-decomposition of a polygon P is a decomposition of P into components that all have a span that is less than or equal to $\alpha \in \mathbb{R}$. The components in such a decomposition are called *α-subpolygons*. Damian and Pemmaraju have already made progress on various decomposition problems involving α-subpolygons [14, 12, 13]. These results are primarily concerned with decomposing simple polygons. Since these results are positive, it is interesting to see if they hold for polygons with holes. This is the focus of this chapter. We note here that the results from this Chapter appear in [40].

Definition 7. *In the* **Decide α-cover** *problem, we seek an answer to the question "does there exist a Steiner-free α-cover D of a polygon P such that $|D| = k$."*

Definition 8. *In the* **Decide α-partition** *problem, we seek an answer to the question "does there exist a Steiner-free α-partition D of a polygon P such that $|D| = k$."*

We will use the rectangle bounded definition of span for each of these problems. If we restrict the input polygon to be an orthogonal polygon with holes, then these problems have a similarity to the problem of decomposing orthogonal polygons into rectangles. This rectangle decomposition problem is tractable if we want a partition [31, 20, 30], while the equivalent covering problem is NP-complete [33]. Thus one might expect that the same duality holds for α-decompositions. To answer this question we need to know the complexity of **Decide α-cover** and **Decide α-partition**.

3.2 The Complexity

As with many decomposition problems on polygons with holes, **Decide α-cover** and **Decide α-partition** are both NP-complete. We begin with the covering version of the problem.

Theorem 10. Decide α-cover is NP-complete for non-simple polygons.

Proof. **Decide α-cover** is in NP since we can "guess" a set of k α-boundable subpolygons and verify that they cover P in polynomial time.

We now proceed to reduce **Planar 3,4SAT** [21] to **Decide α-cover**. **Planar 3,4SAT** is very similar to **Planar 3SAT**. In the **Planar 3SAT** problem we want to decide if a given boolean function ϕ is satisfiable. The problem stipulates that ϕ is in conjunctive normal form with exactly 3 literals per clause. The set of literals found in ϕ is referred to as U and the set of clauses is C. Another restriction is placed on ϕ in this problem: the graph $G(\phi) = (V, E)$ is planar, where $V = U \cup C$ and $E = \{(u, c) | u \in U, c \in C, u \textit{ or } \bar{u} \textit{ is a literal in } c\}$[1]. In the **Planar 3,4SAT** problem there is one more restriction that is imposed: all variables appear at most 4 times negated or unnegated within ϕ. The **Planar 3,4SAT** problem was shown to be NP-complete in [21].

In our reduction, various polygon components will be constructed in a manner that is similar to that presented in [3, 8]. These polygons will have a direct cor-

[1]We can safely assume that $G(\phi)$ is connected, for if it is not then we can consider each of its connected components independently.

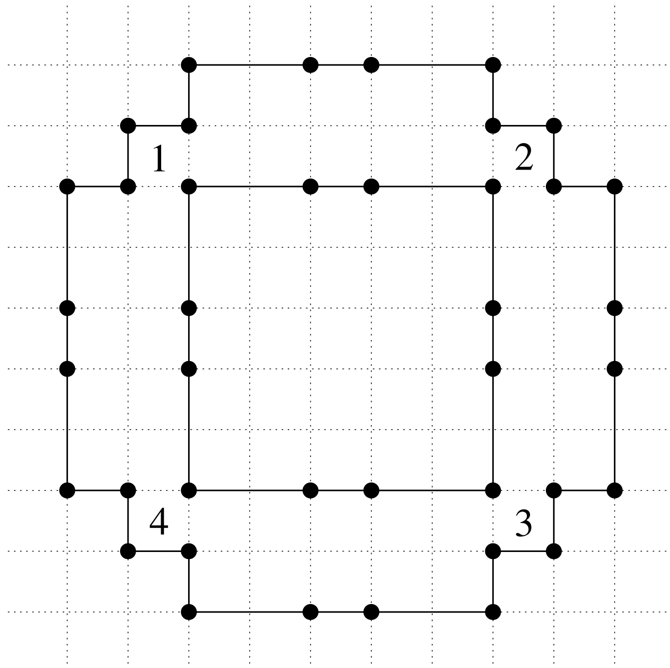

Figure 3.1: The variable polygon. The vertices of the polygon of been emphasized.

respondence to the components of the graph $G(\phi)$ that is described above. The polygons will be joined together to form one large polygon P that will have a minimum α-covering of size k iff ϕ is satisfiable. The value of k will be determined during the construction of P. The graph $G(\phi)$ expresses three main concepts: variables, clauses, and the inclusion of a variable in a particular clause. Each of these concepts will be represented by a different polygon component. For the remainder of this discussion we will fix α at 3 for reasons that will become apparent.

3.2.1 Variable Polygons

The polygon used to represent a variable, called the *variable polygon*, is given in Figure 3.1. Recall that we have fixed α at 3. This polygon can be minimally covered by 8 Steiner-free polygons in exactly 2 ways (Figure 3.2). One of the coverings will represent the variable being set to true and the other will represent false. Wires will attach to variable polygons at 1 or more of the 4 labeled terminals.

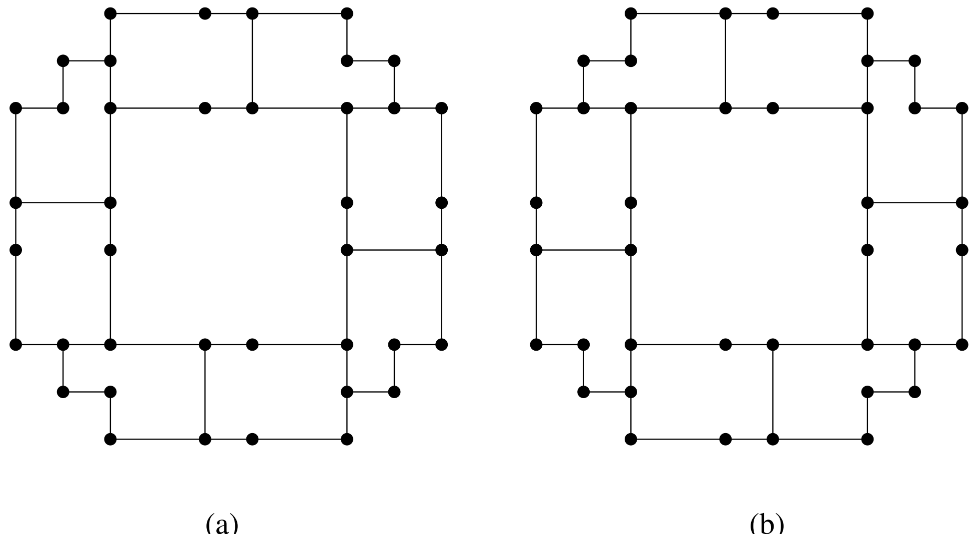

Figure 3.2: Representing truth assignments (a) True (b) False

3.2.2 Wire Polygons

The truth value of a variable will be "transmitted" from a variable polygon to a clause polygon using a sequence of one or more *wire polygons*, or simply *wires* (Figure 3.3). Wire polygons can be attached to variable polygons, other wires, and clause polygons. When attaching a wire to a variable polygon, we butt it up against a terminal. When wires are attached to other wires they do so at an overlap. Similarily, wires will overlap at the terminal of the clause polygon. A sequence of wires connecting a variable polygon to a clause polygon will represent an edge from $G(\phi)$. Wires need to be slimmed down near the terminals of variable polygons so they can attach properly (Figure 3.4). The orientation of the attachment will determine whether the variable is to be negated in the connecting clause. Figure 3.4(a) shows a variable polygon that has been set to true. The covering subpolygons within the variable polygon that overlap with terminals have been shaded in. Notice that these shaded subpolygons can extend over the tip of the outgoing wires. This is because the outgoing wires are all in the unnegated orientation. This is not the case in Figure 3.4(b), where the outgoing wires are in the negated position. Here the shaded polygons cannot extend over the tip of the wires since we have set α

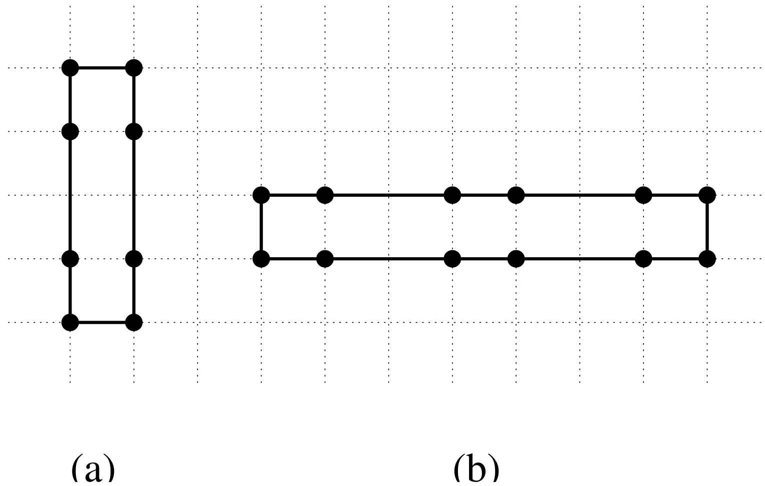

Figure 3.3: (a) A wire polygon (b) 2 wire polygons that are connected

to 3. In Figure 3.4(c), the variable polygon has been set to false, and the outgoing wires are all in the unnegated orientation. Notice that in this case the α-subpolygons that cover the terminals cannot be extended to cover the tips of the outgoing wires. Figure 3.4(d) demonstrates how this situation is reversed when the outgoing wires are orientated in the negated position. This illustrates how truth values travel along wires. Wires that carry true will have the first portion covered by a subpolygon from a variable polygon. When a wire is connected to another wire, the truth value will propagate to the next wire. This is due to the fact that wires are connected at an overlap. Thus wires that carry true will have the "advantage" of being coverable by one less polygon. Clause polygons will be constructed in a way that exploits this. For a wire to connect a variable polygon to a clause polygon it may have to be bent or shifted. (Figure 3.5). Wires will need to be bent or shifted so that they can be guided from variable polygons, onto wires and then towards clauses. A wire may also need to be offset if it is not the right distance away from a clause. Section 3.2.4 completely describes this process of bending, shifting, and offsetting wires.

If we use one of the components from Figure 3.5 then k must be updated accordingly. If we use an offset component, for example, then we must increase k by 4. The remaining segment of a wire carrying false will be covered by a clause polygon.

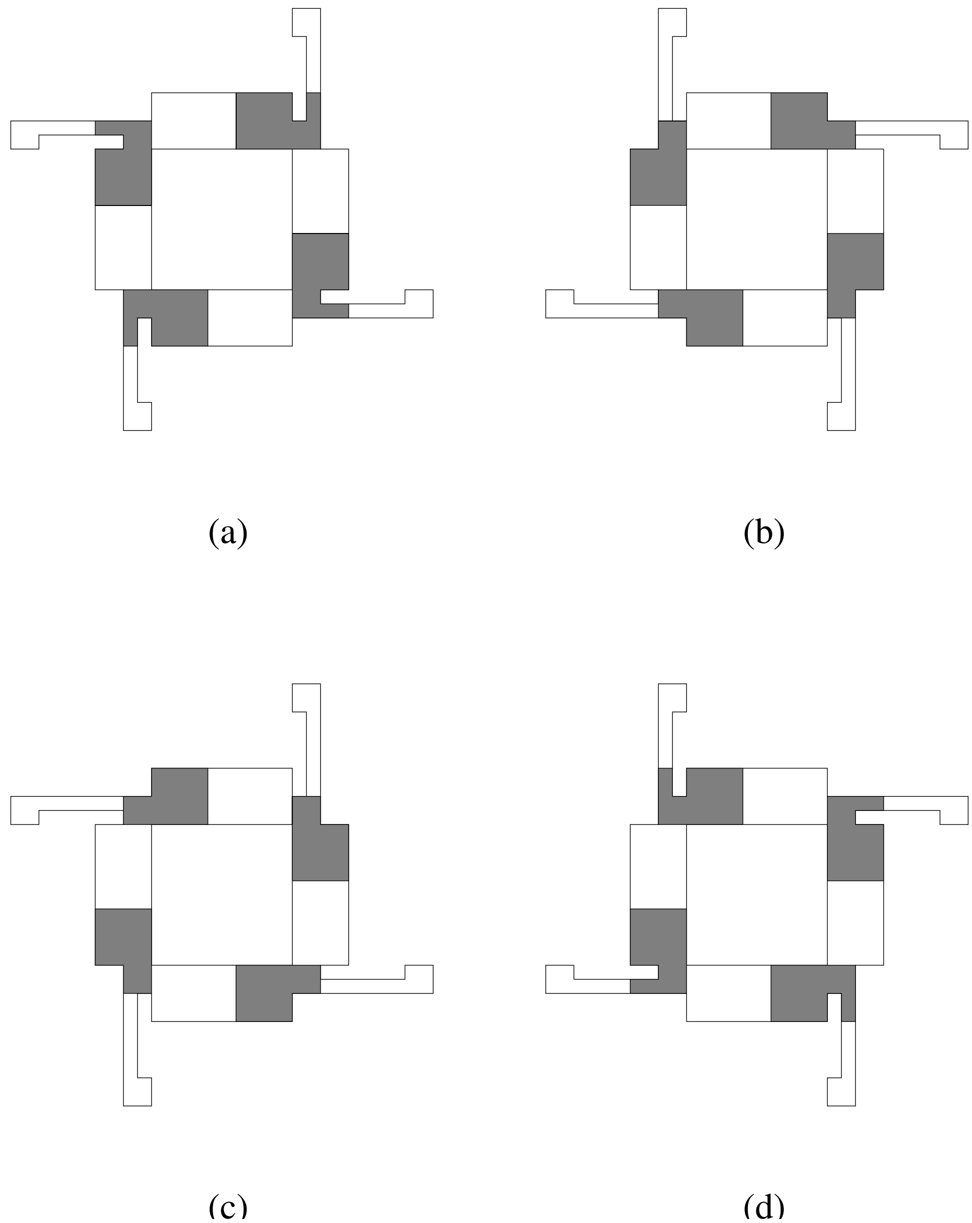

Figure 3.4: (a) A variable polygon set to true with outgoing wires orientated in the unnegated position. (b) A variable polygon set to true with outgoing wires orientated in the negated position. (c) A variable set to false with outgoing wires orientated in the unnegated position. (d) A variable set to false with outgoing wires orientated in the negated position.

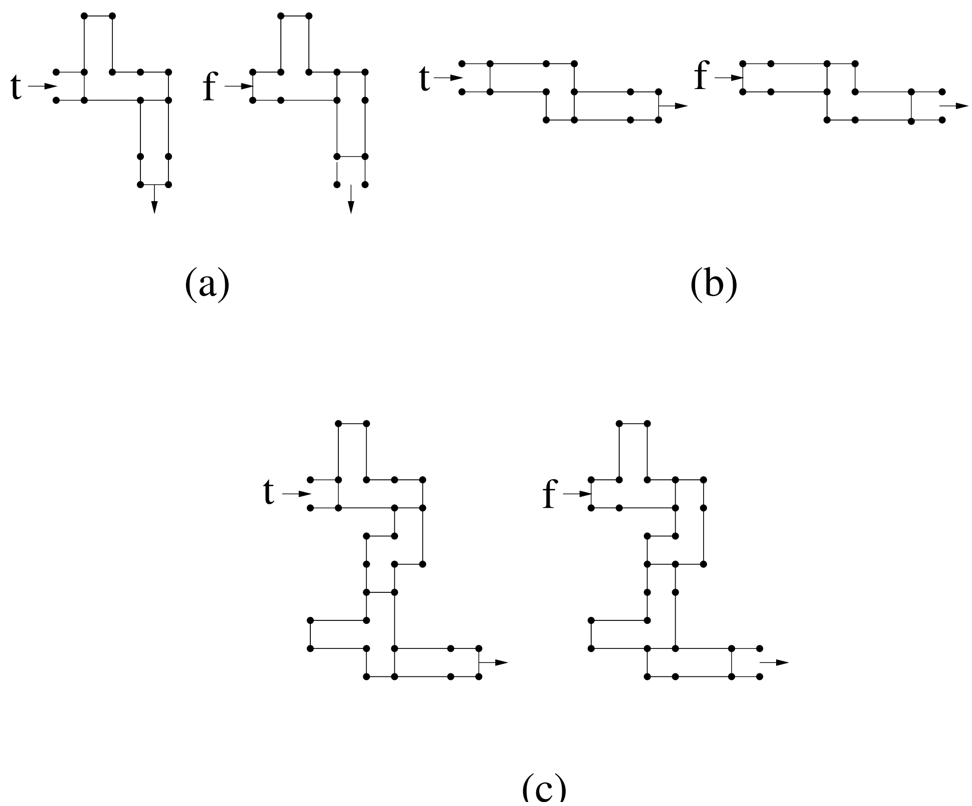

Figure 3.5: (a) Bending a wire. (b) Shifting a wire. (b) Offsetting a wire.

3.2.3 Clause Polygons

The *clause polygon* is shown in Figure 3.6(a). A clause polygon has 3 terminals where incoming wires will be attached. Since each clause has 3 literals, each clause polygon will have 3 incoming wires. The size of the minimum covering for some clause polygon will depend on the wires that are attached to it (Figure 3.6). If 1 or more of the incoming wires is carrying true then the clause polygon will require 3 polygons to cover it. If all incoming wires are carrying false then the polygon will require 4 polygons to be covered. When calculating k, each clause polygon will contribute 3 polygons.

3.2.4 Attaching the Components

We have described the polygonal components that will represent the various aspects of $G(\phi)$. Now we briefly describe the algorithm for building P from these components. We will construct P by approximating a *planar orthogonal grid drawing* of $G(\phi)$. A *planar orthogonal grid drawing* of a planar graph G is a planar embedding of G with the following properties:

- Each edge of G is drawn as a chain of orthogonal lines.
- Each bend in the drawing of an edge lies at an integer valued coordinate.
- Each node of G is drawn as a point that is located at an integer valued coordinate.
- All edges are disjoint except when their endpoints meet at a node.

Figure 3.7 shows an example of a planar orthogonal grid drawing. Notice that for a graph to have a planar orthogonal grid drawing, each node in the graph must have

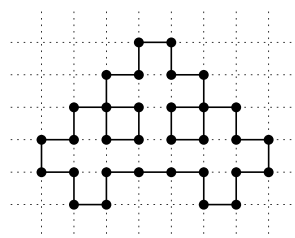

(a)

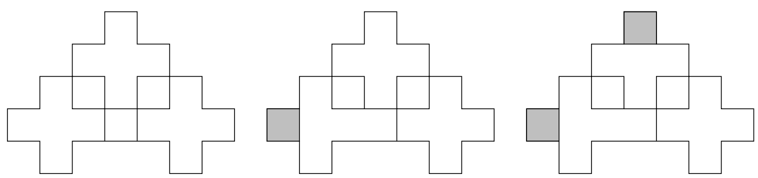

(b) (c) (d)

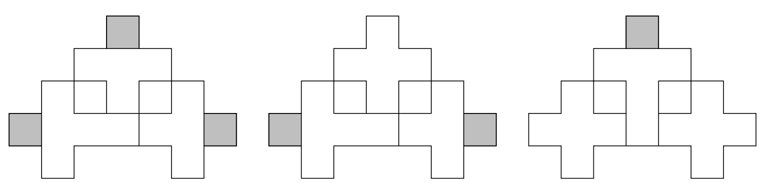

(e) (f) (g)

Figure 3.6: (a) The clause polygon (The symmetric cases are not shown). (b) false, false, false (c) true, false, false (d) true, true, false (e) true, true, true (f) true, false, true (g) false, true, false

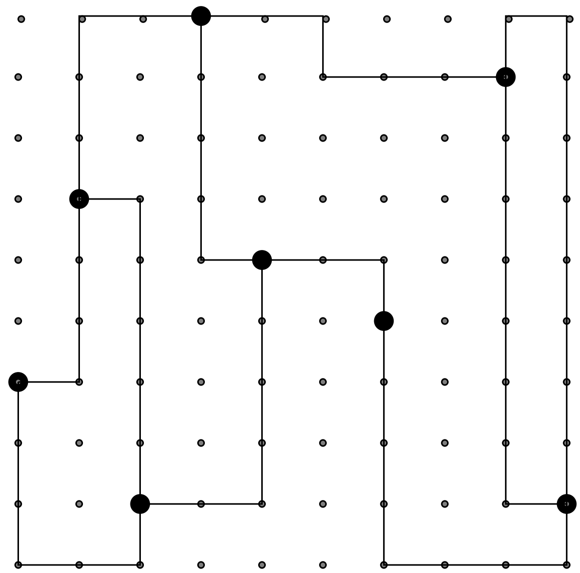

Figure 3.7: An example of a planar orthogonal grid drawing.

degree less than or equal to 4. Such an embedding can be computed in polynomial time for those graphs that admit such an embedding [28].

Let $\mathcal{G}(\phi)$ be a planar orthogonal grid drawing of $G(\phi)$. We will use $\mathcal{G}(\phi)$ as an outline to create P. Figure 3.8 shows how we use $\mathcal{G}(\phi)$ to position the variable and clause polygons. In Figure 3.8 the clause and variable polygons are positioned so that they do not overlap. This may not be achievable for a given embedding $\mathcal{G}(\phi)$. We can get around this problem if we scale $\mathcal{G}(\phi)$ appropriately. The scaling factor can easily be computed by analyzing the width and height of variable and clause polygons.

Once we have our clauses and variables positioned on $\mathcal{G}(\phi)$, we need to attach variables to their appropriate clauses. We will have to bend and shift the wires as the exit variables so that they can match up with the edges in the drawing. Figure 3.9 shows how we can achieve this for negated and unnegated wires. We may need to scale $\mathcal{G}(\phi)$ to accommodate such constructions, but the value of the scaling factor can easily be computed since such constructions can be standardized.

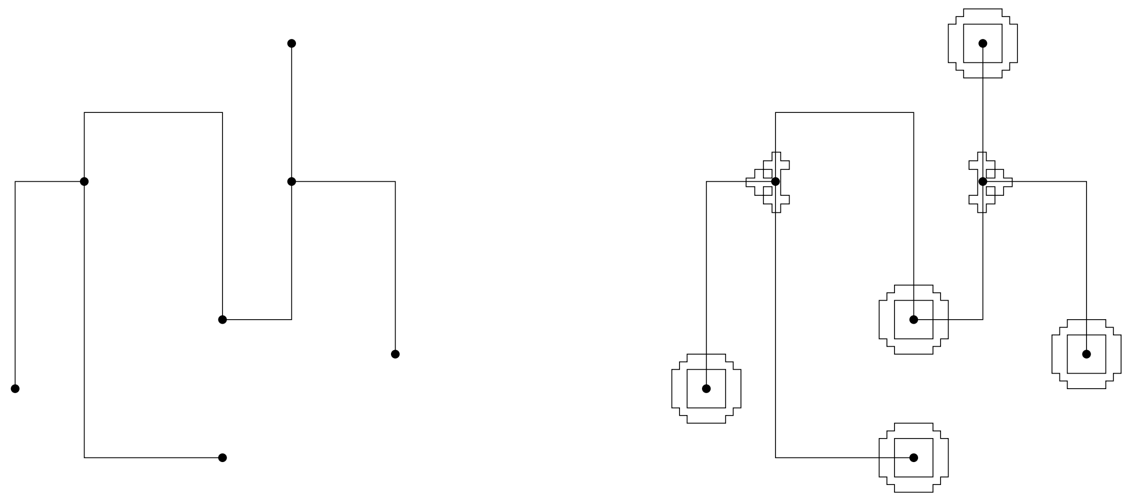

(a) (b)

Figure 3.8: (a) A planar orthogonal grid drawing of a graph. (b) The variable and clause polygons positioned on the drawing.

A few more problems may arise as wires make their way to clauses. A sequence of wire polygons may not be able to strictly follow an edge in the drawing around a bend. We can easily get around this problem by using a sequence of bent and shifted wires. Also, as a sequence of wires approaches its destination clause polygon, it may be the case that they are 1 or 2 units too far away to match up with a terminal of the clause polygon. We can remedy this by using one or two offset wires and an appropriate amount of shifted wires. We can initially scale $\mathcal{G}(\phi)$ by an appropriate amount to accommodate such offsets.

Figure 3.10 shows the complete polygon P for the boolean expression $\phi = (x_1 \vee x_2 \vee \overline{x_3})$. To calculate k, each component must be accounted for. The 3 variable polygons contribute a total of $8 \times 3 = 24$ polygons, the 6 wires contribute 6, and the clause polygon adds 3. Thus we have $k = 24 + 6 + 3 = 33$. The minimum covering has cardinality 33 and this covering corresponds to a satisfying assignment for ϕ. We still must show that for an arbitrary instance of **Planar 3,4SAT** , P will have a minimum covering of size k if and only if ϕ is satisfiable.

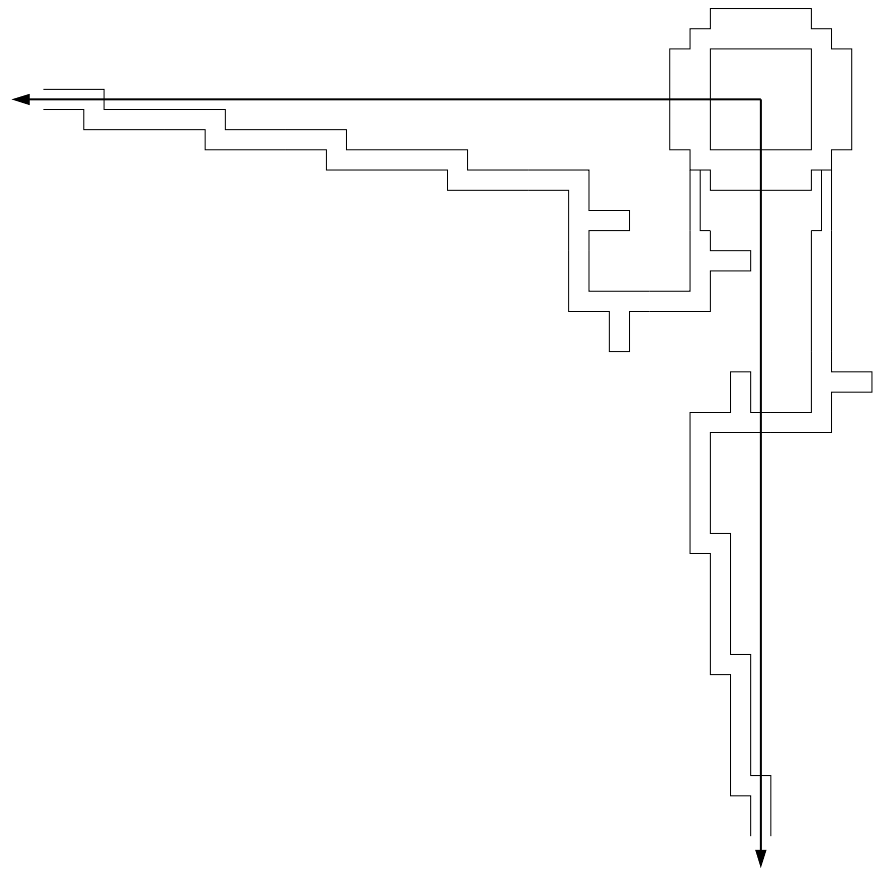

Figure 3.9: Coaxing wires onto the appropriate edges of $\mathcal{G}(\phi)$.

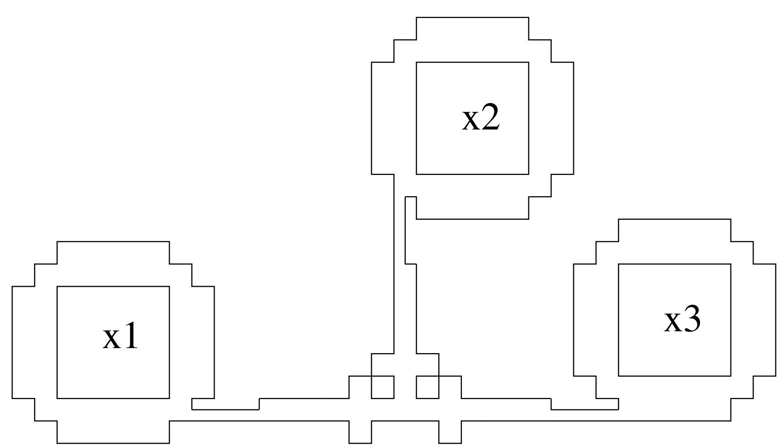

Figure 3.10: The polygon for the boolean expression $\phi = (x_1 \vee x_2 \vee \overline{x_3})$

Suppose that P has a minimum covering M of size k, where k is the value that was calculated during the construction of P. Consider the truth assignment associated with the minimum covering M. Variable polygons can only be minimally covered 2 ways, which ensures that each variable has a truth value. The value of k can be expressed as $8v+1w+3c$, where v is the number of variable polygons, w is the number of wire polygons, and c is the number of clause polygons. Since variable polygons can only be minimally covered 2 ways using 8 polygons we know that $8v$ of the polygons in M are used for covering variable polygons. A sequence of l wires will need l subpolygons to cover them regardless of the truth value that is being transmitted. If they are transmitting true, then they will need l subpolygons to cover the remaining portion of the wires that were not covered by the variable polygon attached to the wire. If the wires are transmitting false then l subpolygons will be necessary to cover all but the last unit of the last wire in the sequence. Thus we know that at least w of the polygons in M must be covering wire polygons. The remaining $3c$ polygons in M are used for covering clause polygons. Recall that a clause polygon needs at least 3 polygons to cover it. Since there are c clause polygons and only $3c$ polygons left in M, we know that each of the clause polygons was covered using 3 polygons. This corresponds to each clause being satisfied, and hence ϕ is satisfiable.

Assume that ϕ is satisfiable under some truth assignment T. We will show how to cover P with k polygons using T. First we cover the variable polygons with 8 polygons each according to T. Use w polygons to cover the wires, which may leave the tips of some wires uncovered. These tips will be within clause polygons. Since we know that T satisfies ϕ, every clause will have at least 1 incoming wire that does not have an uncovered tip. Since this is the case, each clause polygon can be covered with 3 polygons. Thus our covering has size $8c + w + 3c = k$. □

Let us consider the covering discussed in the proof of Theorem 10. We can see that the minimum cover has the same cardinality as the minimum partition: we can replace any subpolygons that overlap, for example in clause polygons, with appropriately shaped ones that do not overlap without affecting the cardinality of the decomposition. This implies the following:

Theorem 11. Decide α-partition *is NP-complete for non-simple polygons.*

Thus the covering and partitioning versions of this problem are both NP-complete. We can extend Theorems 10 and 11 in two ways. Firstly, the polygons constructed in the proof of 10 are all orthogonal. Thus we have the following:

Corollary 1. Decide α-cover *is NP-complete for non-simple orthogonal polygons.*

Corollary 2. Decide α-partition *is NP-complete for non-simple orthogonal polygons.*

Secondly, we can extend our results to hold for the circle-bounded definition of span. The proofs of Theorems 10 and 11 can be restated for this problem if we fix α to be the diameter of the smallest bounding circle for a square whose sides are 3 units long. Although the radius of such a circle will be $\sqrt{18}$, and hence irrational, we can use an approximate value without loss of correctness.

Chapter 4

Conclusion and Open Problems

In this thesis we have studied two polygon decomposition problems. In Chapter 2 we showed that the **Minimum r-star Cover** problem can be solved in polynomial time. Recall that this also implies that **Rectilinear Art Gallery** can also be solved in polynomial time. Although this result is positive, it comes with a caveat. The algorithm we describe for solving these problems makes use of the minimum clique cover algorithm for perfect graphs that is described in [18]. This algorithm computes the minimum clique cover by numerical methods and is impractical as a "real-world" algorithm. This leads to our first open problem.

Open Problem 1. Is there a polynomial time combinatorial algorithm for solving **Minimum r-star Cover**?

Open Problem 1 could be solved in two different ways. One could solve this problem by discovering a combinatorial algorithm that uses the geometric properties of the problem directly (i.e. does not use a graph theoretic transformation). This problem could also be solved if there exists an efficient combinatorial algorithm for computing minimum clique covers on perfect graphs. There is hope for such an algorithm now that the Strong Perfect Graph Theorem has been proven.

It may be the case that the building being guarded in the **Rectilinear Art Gallery** contains visual obstacles that are not walls. For example, a building may contain pillars or furniture that would interfere with visibility. In this situation, we can model, or at least approximate, such obstacles as orthogonal holes in the input polygon P.

Open Problem 2. What is the complexity of **Rectilinear Art Gallery** when orthogonal visual obstacles are present?

This problem can be restated as **Minimum r-star Cover** for orthogonal polygons that contain orthogonal holes. It is likely that this problem is NP-complete.

In Chapter 3 we investigated the complexity of decomposing non-simple polygons into span bounded polygons. We showed that both the covering and partitioning versions of this problem are NP-complete. We also showed that this result holds even if we restrict input polygons to be non-simple orthogonal polygons. The partitioning version of this problem is known to be in **P** if we restrict the input to be a simple polygon, but the time complexity of the covering version of this problem is unknown.

Open Problem 3. What is the time complexity of computing α-covers of simple polygons?

Of particular interest is the optimization version of this problem. This problem is important since decomposing a polygon increases the number of objects we must consider. Thus it is important to minimize the number of subpolygons in a decomposition.

In [14], Damian describes a 4.5-approximation algorithm for minimally partitioning a convex polygon with the use of Steiner points. This is the only known result for Steiner α-decompositions. Steiner decompositions are interesting since every polygon has a Steiner α-decomposition, while not every polygon has a Steiner-free α-decomposition.

Open Problem 4. What is the time complexity of computing Steiner α-decompositions?

Bibliography

[1] J. Abello, V. Estivill-Castro, T. Shermer, and J. Urrutia. Illumination of orthogonal polygons with orthogonal floodlights. *International Journal of Computational Geometry and Applications*, 8:25–38, 1998.

[2] A. Aggarwal. *The Art Gallery Theorem: Its Variations, Applications and Algorithmic Aspects*. PhD thesis, Deptartment of Electrical Engineering and Computer Science, John Hopkins University, 1984.

[3] Larry J. Aupperle. Covering regions by squares. Master's thesis, University of Saskatchewan, July 1987.

[4] Claude Berge. Farbung von graphen deren samtliche bzw. ungerade kreise starrsind (zumsammenfassung). *Wiss. Z. Martin-Luther-Univ. Halle-Winterberg, Math.-Natur. Reihe*, 114, 1961.

[5] Andreas Brandstädt, Van Bang Lee, and Jeremy P. Spinrad. *Graph Classes: A Survey*. SIAM Monographs on Discrete Mathematics and Applications. 1999.

[6] W. Chin and S. Ntafos. Optimum watchman routes. In *Proceedings of the second annual symposium on Computational geometry*, pages 24–33. ACM Press, 1986.

[7] V. Chvátal. Star-cutsets and perfect graphs. *J. Combin. Theory Ser. B*, 39:189–199, 1985.

[8] H. E. Conn and J. O'Rourke. Some restricted rectangle covering problems. In *Allerton Conference*, pages 898–907, 1987.

[9] Joseph Culberson and Robert Reckhow. Covering a simple orthogonal polygon with a minimum number of orthogonally convex polygons. In *3rd Annual ACM Symposium on Computational Geometry*, pages 268–277, 1987.

[10] Joseph Culberson and Robert Reckhow. Dent diagrams: A unified approach to polygon covering problems. Technical Report TR 87-14, Department of Computer Science, University of Alberta, July 1987.

[11] Joseph Culberson and Robert Reckhow. Covering polygons is hard. *Journal of Algorithms*, 14:2–44, 1994.

[12] Mirela Damian and Sriram Pemmaraju. Computing optimal diameter-bounded polygon partitions. *Algorithmica*. To Appear.

[13] Mirela Damian and Sriram Pemmaraju. Computing optimal α-fat and α-small decompositions. In *Proceedings of the 12th Annual ACM Symposium on Discrete Algorithms*, pages 338–340, 2001.

[14] Mirela Damian-Iordache. *Shape Constrained Polygon Decomposition and Graph Domination Problems*. PhD thesis, University of Iowa, July 2000.

[15] Micheal R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[16] L. Gewali, J. M. Keil, and S. Ntafos. On covering orthogonal polygons with star shaped polygons. 7:45–63, 1992.

[17] Martin Charles Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Computer Science and Applied Mathematics. Academic Press, New York, 1980.

[18] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences on combinatorial optimization. *Combinatorica 1*, pages 169–197, 1981.

[19] Frank Hoffmann. On the rectilinear art gallery problem. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, pages 717–728. Springer Verlag, 1990.

[20] H. Imai and Ta. Asano. Efficient algorithms for geometric graph search problems. *SIAM Journal of Computing*, 15:478–494, 1986.

[21] Klaus Jansen and Haiko Müller. The minimum broadcast time problem for several processor networks. *Theoretical Computer Science*, 147(1–2):69–85, 1995.

[22] J. Kahn, M. Klawe, and D. Kleitman. Traditional galleries require fewer watchmen. *SIAM Journal of Algebraic Discrete Methods*, 4(2):194–206, 1983.

[23] J. M. Keil. *Decomposing a polygon into simpler components*. PhD thesis, University of Toronto, Toronto, Canada, 1983.

[24] J. Mark Keil. Decomposing a polygon into simpler componenets. *SIAM Journal of Computing*, 14:799–817, 1985.

[25] J. Mark Keil. Minimally covering a horizontally convex orthogonal polygon. In *2nd Annual ACM Symposium on Computational Geometry*, pages 43–51, 1986.

[26] J. Mark Keil. Polygon decomposition. In *Handbook of Computational Geometry*, chapter 11. Elsevier Science B. V, 2000.

[27] Mark Keil and Joe Culberson. Personal Communication.

[28] T. Lengauer. *Combinatorial Algorithms for Integrated Circuit Layout*. Teubner/Wiley & Sons, 1990.

[29] A Lingas. The power of non-rectilinear holes. In *Proceedings of the 9th International Colloquium on Automata, Languages, and Programming*, volume 140 of *Lecture Notes in Computer Science*, pages 369–383. Springer Verlag, 1982.

[30] W. Lipski. An O($nlogn$) manhattan path algorithm. *Information Processing Letters*, 19:99–102, 1984.

[31] W. Lipski, Jr. E. Lodi, and F. Luccio. On two-dimensional data organization ii. *Fund. Inform.*, 2:245–260, 1979.

[32] L. Lovász. A characterization of perfect graphs. *Journal of Combinatorial Theory*, 2:253–267, 1972. MR46 #1624.

[33] W. J. Masek. Some NP-complete set covering problems. Manuscript, 1979.

[34] M.Chudnovsky, N. Robertson, P.D.Seymour, and R.Thomas. The strong perfect graph theorem. Submitted for publication.

[35] J. Molnár. Uber den zweidimensionalen topologischen satz von helly. *Mat. Lapok*, 8:108–114, 1957.

[36] R. Motwani, A. Raghunathan, and H. Saran. Covering orthogonal polygons with star polygons: The perfect graph approach. *Comput. Syst. Sci.*, 40:19–48, 1990.

[37] J. O'Rourke and K. J. Supowit. Some NP-hard polygon decomposition problems. *IEEE trans. Information theory*, 29(2):181–190, 1983.

[38] Joseph O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, August 1987.

[39] Arvind Raghunathan. *Polygon Decomposition and Perfect Graphs*. PhD thesis, University of California, Berkeley, 1988.

[40] Chris Worman. Decomposing polygons into diameter bounded components. In *Proceedings of the 15th Canadian Conference on Computational Geometry*, pages 103–106, 2003.

[41] Yunhong Zhou and Subhash Suri. Analysis of a bounding box heuristic for object intersection. *Journal of the ACM (JACM)*, 46(6):833–857, 1999.